University of Southampton

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

University of Southampton

Faculty of Engineering and Physical Sciences

Department of Aeronautical and Astronautical Engineering

Machine Learning for Tailored Multifunctional Properties of Lightweight Metamaterials

DOI: 10.5258/SOTON/PG/T166

by

Peter Anthony Hawes

ORCID ID 0009-0006-0909-2097

Thesis for the degree of Doctor of Philosophy

April 2026

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
Department of Aeronautical and Astronautical Engineering
Thesis for the degree of Doctor of Philosophy

Machine Learning for Tailored Multifunctional Properties of Lightweight Metamaterials

by

Peter Anthony Hawes

Metamaterials are man-made materials with periodic cellular architectures, offering properties that are defined by geometry rather than chemical composition. This allows significant potential for lightweighting and tunability of mechanical responses that can be tailored to specific practical loading scenarios. Furthermore, additional levels of multifunctionality can be activated and tuned according to requirements such as sound and energy absorption. The principal aim of this thesis is to explore the effectiveness of machine learning for designing metamaterial architectures that maximise mechanical performance while performing multi-objective optimisations to balance trade-offs between competing design objectives, enabling enhanced multifunctionality. A comprehensive literature review has been carried out to explore the current state of the art in architected metamaterials. Machine learning offers a vast array of potential applications across engineering disciplines. A selection of five studies was presented in which workflows were developed to allow tailoring of material properties by artificial intelligence. An inverse convolutional neural network was used to fine-tune sound absorption responses of micro-perforated panel sound absorber metamaterials for maximum absorption amplitude and range. The process was then applied to tune mechanical responses of truss-based lattice structures for enhanced structural efficiency. Networks were then trained to perform inverse tailoring of combined sound absorption and structural performance of a quarter-wavelength resonator labyrinth unit cell. An alternative approach was explored that uses neural networks to predict performance objectives in combination with evolutionary optimisation algorithms. This method was used to synthesise material architectures for enhanced stiffness, strength and high velocity impact resistance. Both inverse and predictor approaches were used to tailor layup sequences for optimal stress responses in composites. An early proof of concept was also obtained for use of artificial intelligence as a faster alternative to finite element analysis in parametric stress analysis across engineering design practices. The developed network architectures and methodologies may be interchangeable with other problems of similar complexity and can theoretically be applied to any multi-variable design problem for which a dataset can be produced to train a machine learning model.

Table of Contents

Table of Contents **3**
Table of Tables **8**
Table of Figures **11**
List of Accompanying Materials **29**
Research Thesis: Declaration of Authorship **30**
Acknowledgements **31**
Definitions and Abbreviations **32**
Chapter 1 Introduction **36**
1.1 Background **36**
1.1.1 Material Property Space 37
1.1.2 Architectured Cellular Materials 38
1.2 Research Aim and Objectives **41**
1.3 Outline of the Thesis **44**
Chapter 2 Literature Review **47**
2.1 Design Optimisation of Metamaterials for Mechanical Performance **47**
2.1.1 Bending Vs Stretch Dominated Structures 47
2.1.2 Manually Generated Lattice Unit Cell Configurations 48
2.1.3 Numerically Generated Lattice Structures 60
2.2 Machine Learning **70**
2.2.1 Artificial Neural Networks 71
2.2.2 Deep Learning 71
2.2.3 Convolutional Neural Networks 77
2.2.4 Potential for Applications of Deep Learning in Metamaterial Design 78
2.2.5 Generative Design 82
2.3 Materials and Manufacturing Methods **86**
2.3.1 Mechanical Assembly 86
2.3.2 Additive Manufacture 92
2.4 Multi-scale Hierarchical Design **98**

2.5 Multifunctionality 106

2.5.1 Energy Absorption and Impact Resistance 106

2.5.2 Sound Absorption 114

2.6 Conclusion 123

Chapter 3 Design of Subwavelength Wide Bandwidth Sound Absorbers by Inverse Convolutional Neural Networks 124

3.1 Dataset Generation 125

3.2 Convolutional Neural Network Architecture 128

3.3 Training Performance Analysis 130

3.4 Optimisation and Experimental Validation 135

3.5 Compensation for Materials and Manufacturing 140

3.6 Design of Subwavelength Wideband Absorber for Industry Application 145

3.7 Design of Experiments Study of Dataset Size Requirements 150

3.8 Conclusion 153

Chapter 4 Artificial Intelligence Based Optimisation of Metamaterials for Tailored Stress-Strain Response and Enhanced Structural Efficiency 154

4.1 Base Material Properties Characterisation 155

4.1.1 Experimental Setup 155

4.1.2 FDM Material Comparison 160

4.1.3 AM Manufacturability Tests 165

4.2 FEA Dataset Generation 167

4.2.1 Non-linear FEA Model 167

4.2.2 Experimental Validation 171

4.3 Convolutional Neural Network Training 184

4.4 Tailored Mechanical Responses 189

4.5 Conclusion 197

Chapter 5 Machine Learning Based Multi-objective Structural Optimisation of Lightweight Sound Absorber Metamaterials 199

5.1 Inverse Tailoring of Combined Sound Absorption and Structural Performance 201

5.1.1 Metamaterial Unit Cell Configuration 201

5.1.2 Dataset Generation 218

5.1.3 Inverse Convolutional Neural Network Architecture 226

5.1.4 Network Training Performance 228

5.1.5 Inverse Optimisation and Experimental Validation 231

5.2 Multi-objective Optimisation by Predictor Neural Network 241

5.2.1 Predictor Neural Network Model 242

5.2.2 Training and Test Performance 244

5.2.3 Multi-objective Genetic Algorithm 247

5.2.4 Experimental Validation 250

5.3 Reducing Overall Material Thickness 255

5.3.1 Dataset Generation 255

5.3.2 Inverse Convolutional Neural Network 259

5.3.3 Inverse Optimisation and Experimental Validation at Reduced Thickness 262

5.4 Conclusion 268

Chapter 6 Multi-objective Optimisation of Lightweight Metamaterials for Impact Resistance by Machine Learning 270

6.1 Parametric Unit Cell Configuration 271

6.2 Dataset Generation 274

6.2.1 Static FEA Model 274

6.2.2 High Velocity Dynamic FEA Model 276

6.3 Convolutional Neural Network Training 280

6.4 Multi-objective Genetic Algorithm Optimisation 284

6.4.1 Structural Properties Optimisation 284

6.4.2 Specific Energy Absorption Optimisation 294

6.5 Conclusion 299

Chapter 7 Layup Optimisation of Carbon Fibre Reinforced Polymer Composites by Machine Learning 300

7.1 Inverse Design Approach for Tailored Stress Distribution 301

7.1.1 Dataset Generation 301

7.1.2 Convolutional Neural Network Training and Performance Evaluation 303

7.1.3 Tailoring Composite Layup for Desired Stress Responses 308

7.2 Stress Optimisation Via Predictor Neural Networks 314

7.2.1 Shallow Artificial Neural Network Predictor Model 315

7.2.2 Deep Learning Network Predictor Models 321

7.2.3 Maximum Stress Optimisation 334

7.3 Artificial Intelligence Based Parametric Stress Analysis 342

7.4 Composite Layup Optimisation for Specific Geometry and Boundary Conditions 348

7.5 Conclusion 354

Chapter 8 Conclusion and Future Works 355

8.1 Conclusion 355

8.2 Future Works 357

Appendix A MPP CNN Model Design of Experiments Test Plans 359

Appendix B MPP CNN Model Design of Experiments Study Validation Error, Loss and Regression Performance Plots 362

Appendix C ASTM D638 Tensile Coupon Measurements 367

Appendix D ASTM D695 Compressive Coupon Measurements 368

Appendix E ASTM D638 Tensile Test Results 370

Appendix F ASTM D695 Compressive Test Results 380

Appendix G Engineering and True Stress-Strain Compressive Material Data Used in the ANSYS Non-linear Material Model 396

Appendix H Solid Body FEA Model Element Size Comparison 400

Appendix I Non-linear Eigenvalue Buckling Analysis Mode Shapes for 2mm Strut Radius Lattice Compression Block 401

Appendix J Solution Unit Cell Geometries of Inverse Optimised Metamaterials 404

Appendix K Material Properties of Genetic Algorithm Optimised Pareto Front Solutions 405

Appendix L Solution Geometries of Low Thickness Inverse Optimised Metamaterials 411

Appendix M Pareto Front Solution Parameters Optimised by GA to Maximise Relative Specific Stiffness and Strength 412

Appendix N Pareto Front Solution Parameters Optimised by GA to Maximise Relative Specific Stiffness and Strength While Minimising Mass ... 414

Appendix O Pareto Front Solution Parameters Optimised by GA to Maximise SEA, Relative Specific Stiffness and Strength While Minimising Mass 416

Appendix P Composite Stress Predictor Model Deep Learning Network Architecture Iterations 418

Appendix Q Comparison Between FEA and CNN Predicted Composite Stress Contours 420

Appendix R Composite Maximum Stress Predictor Network Performance Evaluations 425

Appendix S Composite Maximum Stress Optimisation Test Solution Layup Stacking Angles 432

Appendix T Composite Parametric Stress Predictor Model Error Evaluation Test Dataset 433

Appendix U Comparison of Maximum Von Mises Stress Values for Optimised Layups Against Values Obtained Via Exhaustive Search 434

Appendix V Composite Layup Stress Optimisation Consistency 436

List of References 437

Table of Tables

Table 1 Coefficients for relative density scaling laws of common lattice topologies [8]. 58

Table 2 Activations and learnable parameters at each stage of the convolutional neural network 129

Table 3 Summary of best performing DNN and CNN architectures 132

Table 4 Measurements of manufacturing accuracy 142

Table 5 Discrepancy and deviation values for design of experiments test plans of 1300 samples 150

Table 6 FDM slicer settings 158

Table 7 FDM tensile and compressive material properties comparison 163

Table 8 Comparison of tensile material properties with manufacturer datasheet values 163

Table 9 Manufacturer testing methods 163

Table 10 Deviations in stiffness and strength with build direction of Polymaker CoPA compression samples 165

Table 11 Non-linear FEA contact model solution times for 2mm strut radius lattice example 177

Table 12 CNN training results for lattice mechanical response inverse design model 185

Table 13 Summary of solutions for tailored mechanical responses of lattice compression block using CNN trained with beam element model 189

Table 14 Experimental energy absorption values measured for lattice compression block solutions targeting 500N design load 192

Table 15 Summary of solution iterations to satisfy design criteria at 5kN design load target using CNN trained with solid element model 193

Table 16 Summary of solution iterations to satisfy design criteria at 2kN design load target using CNN trained with solid element model 193

Table 17 Impedance tube geometry parameters 211

Table 18 Absorption peak locations for spiral quarter-wavelength resonator samples with varied aspect ratio 215

Table 19 Measured mechanical properties of ColorFabb XT-CF20 FDM base material 219

Table 20 Geometry parameters for mesh convergence study test cases 222

Table 21 Parameter limits for chosen design space 224

Table 22 Inverse convolutional neural network layers 227

Table 23 Summary of training performance and test regression for inverse CNNs trained for multi-objective and single-objective optimisation 230

Table 24 First harmonic sound absorption performance of example optimised metamaterials 234

Table 25 Mechanical properties of example optimised metamaterials 234

Table 26 Geometry parameters of example optimised metamaterials 234

Table 27 Training progress and test performance of predictor neural networks 244

Table 28 Sound absorption and structural properties of GA optimised Pareto front solutions 251

Table 29 Structural properties of solution unit cells determined via exhaustive search 254

Table 30 Revised geometry parameter limits for reduced thickness model dataset 256

Table 31 Summary of training performance and test regression for inverse CNN trained to predict the geometry parameters of the metamaterial unit cell designed for reduced thickness 260

Table 32 First harmonic sound absorption performance of example metamaterials optimised for overall material thicknesses below 6mm 264

Table 33 Mechanical properties of example metamaterials optimised for overall material thicknesses below 6mm 264

Table 34 Geometry parameters of example metamaterials optimised for overall material thicknesses below 6mm 264

Table 35 Geometry parameters for 4 microphone test setup to measure transmission loss ... 267

Table 36 Node location matrix 271

Table 37 Node connectivity matrix of strut radius parameters 272

Table 38 Summary of training performance and test regression for predictor CNNs 283

Table 39 FEA model orthotropic material data 302

Table 40 Convolutional neural network training iterations 304

Table 41 Summary of requested stress responses, layup sequences returned by the CNN and resulting actual stress responses measured via FEA 309

Table 42 ANN training progress and performance summary 318

Table 43 Predictor network DNN architecture 322

Table 44 Predictor network CNN architecture 324

Table 45 Summary of training performance for chosen stress predictor networks 325

Table 46 Composite stress predictor network models mean error comparison summary 331

Table 47 Maximum stress predictor network training performance 334

Table 48 Comparison summary of absolute error between predicted and known maximum stress values for 500 test layups 336

Table 49 Optimisation algorithm settings 337

Table 50 Best performing solution layup stacking sequences 339

Table 51 Parametric stress predictor model parameters 343

Table of Figures

Figure 1 Modulus-density material properties space. Carbon fibre reinforced polymer (CFRP) and aluminium lattice structures are predicted to extend from the location of the legacy material to fill holes in the design space by defining material properties through geometry rather than chemical composition. Adapted from [8] 38

Figure 2 Cross section of human femur showing trabecular and cortical bone (top) [11] and cross section of a bird wing (bottom) [12] 39

Figure 3 Examples of periodic lattice unit cell topologies commonly used as cores of composite sandwich panels. Adapted from [19] 48

Figure 4 Weight per unit area against dimensionless loading parameter for comparison between optimised truss cores and honeycombs, all with $\sigma_y/E = 0.007$ and Poisson ratio = 1/3. Both truss core plates are fully optimised. Two cases of honeycombs are considered, one with core thicknesses, H_c fixed at $H_c/l = 0.1$ and one with core thickness taken equal to that of the sandwich plate with truss core at the same value of $V/(EM)^{1/2}$. From [20] 49

Figure 5 (a) Diagram of hat-stiffened plate; (b) Normalised weight per unit area as a function of the dimensionless load parameter for optimally designed simply-supported wide panels under compression ($\sigma_y/E = 0.007$ and Poisson ratio = 1/3). The upper and lower curves are for the sandwich plates with a truss core, constrained and unconstrained respectively. The middle curve is an optimised hat-stiffened single layer plate, which is regarded as one of the most efficient structures for this application. From [20] 51

Figure 6 Non dimensional force/strain curves simulated in compression for tetragonal and Kagomé cores made from (a) Cu alloy and (b) Al alloy, where R_c is the truss radius [28]..... 54

Figure 7 Structure of the octet-truss lattice unit cell. Darkened struts represent the octahedral core while the lighter struts connected at nodes p1-4 form one of eight tetrahedral cells [31]. 55

Figure 8 Comparison between stiffness and strength of the FCC octet-truss lattice, metallic foams and the Hashin-Shtrikman upper bounds [31]. 56

Figure 9 Comparison between analytical, numerical and experimental stress-strain responses for octet-truss lattice material cast from LM25 aluminium alloy [31]. 57

Figure 10 The damage tolerance of a lattice material and the parent solid [8]. 59

Table of Figures

Figure 11 Uniform and conformal lattices. 60

Figure 12 Triply periodic minimal surface unit cell configurations [41]. 61

Figure 13 Scaling relationships between Young's modulus and relative density [41]. 62

Figure 14 Random models of open cell solids. (a) Open-cell Voronoi tessellation with ρ/ρ_s = 0.15 (b) Node-bond model with $Z = 5.5$ (c) Digital node-bond model with $Z = 12$ (d) Open-cell GRF with $\rho/\rho_s \approx 0.1$. From [46] 65

Figure 15 Artificial neural network architecture 71

Figure 16 Forward propagation 72

Figure 17 (a) Sigmoid activation function; (b) hyperbolic tangent activation function 73

Figure 18 (a) ReLU activation function; (b) leaky ReLU activation function 74

Figure 19 Stochastic gradient descent with momentum (blue) and without momentum (red) 74

Figure 20 Example of convolution process 78

Figure 21 Unit cell configurations, pink and black colours referring to stiff and soft building blocks respectively (top) and microstructure consisting of an assemblage of unit cells that is then encoded (blue=U1, orange=U2, yellow=U3) and converted to a data matrix (bottom). From [69] 81

Figure 22 General Motors generative design seat belt bracket iterations [75]. 84

Figure 23 BAC Mono generative design wheel [76]. 84

Figure 24 Pyramidal truss cored lattices manufactured by snap-fitting CFRP laminates. From [81] 87

Figure 25 Illustration of the fabrication approach for CFRP cylindrical tube-based octet lattice structures. (a) A micro-water jet cuts sub-components from CFRP plate, (b) CNC milling of the cylindrical tips of joint connectors, (c) a single CFRP tube-based octet lattice structure, and (d) a single unit cell of the octet lattice structure. From [82] 88

Figure 26 (a) "cuboct" lattice assembly; (b) carbon fibre–reinforced polymer composite cuboct lattice; (c) reversible joints; (d) micrograph showing fibre orientations; (e) slicing of components from wound fibre stock. From [84] 89

Figure 27 Glass fibre octet-truss lattice manufactured by continuous spatial weaving. Adapted from [34] 91

Figure 28 (a) Projection microstereolithography process; (b-e) fabricated octet-truss microlattices; (b) solid polymer hexanediol diacrylate (HDDA); (c) hollow-tube metallic nickel-phosphorus (Ni-P); (d) hollow-tube ceramic (alumina); (e) solid ceramic (alumina); (f-i) magnified views of the struts of the microlattices in (b) to (e), respectively. From [13] 93

Figure 29 Load-elongation curves for PLA pulley housings (left) and test fixtures (right) [97]. 95

Figure 30 Continuous Composites continuous fibre additive manufacturing. From [101] 96

Figure 31 CCFRP demonstration component – Team Bath Racing 2020 Formula Student car rear suspension upright..... 97

Figure 32 Scanning electron micrographs of structural hierarchy present in the unit cell of hollow nickel alloy metamaterial. From [103]..... 99

Figure 33 Hierarchical honeycombs; (a) Unit cell of the hierarchical honeycombs with regular structure and with 1st and 2nd order hierarchy; (b) Images of honeycombs with a = 2 cm fabricated using three-dimensional printing. From [105] 100

Figure 34 Second order hierarchical pyramidal truss structure; (a) without and (b) with face sheets. From [109]..... 102

Figure 35 Hierarchical truss lattices [110]. (a) Octet of octet (OT –OT); (b) octahedron of octet (ON –OT) and (c) tetrakaidecahedron of octet (TN –OT). Figs. (i) are CAD models of the integral structures; (ii) are macroscopic configurations showing nodal connectivity Z; (iii) are mesoscopic lattice configurations; (iv) are hierarchical lattice materials fabricated by SLS 3D printing and (v) are the corresponding graphs showing microstructures at nodes and struts. 103

Figure 36 Hierarchical triply periodic minimal surface. From [111]..... 104

Figure 37 Hierarchical metamaterial configurations; (a) orders of hierarchy; (b) cubic family (simple cube, face diagonal cube, face centred cube, body centred cube); (c) octahedron family; (d) hybrid family. From [112] 105

Figure 38 CFRP tube reinforced hierarchical foam metamaterial before (left) and after crushing (right). From [114]..... 107

Table of Figures

Figure 39 Hierarchical CFRP tube structures. From [115] 108

Figure 40 Tubular auxetic metamaterial design; (a) 2D planar unit cell design; (b) 3D RVE of four unit cells arranged along the circumference: (c) tubular lattice formed from 6 RVEs; (d) 3D-printed sample; (e) and (h) Effect of wave amplitude A/l on stress and Poisson's ratio, respectively. Here, t/l=h/l=0.1; (f) and (i) Effect of varying beam thickness t/l with A/l=0.3 and h/l=0.2; (g) and (j) Effect of varying beam depth h/l with A/l=0.3 and t/l=0.2. From [117] 109

Figure 41 Unit cell optimised for maximum energy absorption and volume fraction of 12.5% (left), 6x6 lattice sample (centre) and stress-strain response with energy absorption comparison (right). From [121] 112

Figure 42 Uniform and graded density lattices. From [122] 113

Figure 43 (a) Schematic of Helmholtz resonator with extended neck; (b) checkerboard arrangement of alternating resonators A and B with different length necks. From [127] 116

Figure 44 Hybrid HR and fractal structure metamaterial design. From [128] 116

Figure 45 3D view of the individual absorber and its theoretical model. From [131] 118

Figure 46 Absorption response of parallel coiled coplanar absorbers. From [131] 118

Figure 47 Optimisation of quadruple parallel arranged PPAs with different back cavity depths. From [135] 120

Figure 48 (a) Model of MPP absorber; (b) electro-acoustic equivalent circuit 125

Figure 49 Convolutional neural network architecture. Block colours correspond to specific layer types as follows: blue = sequence input, red = regression output, yellow = convolution and fully connected, orange = ReLU activation, green = normalisation, purple = pooling ... 129

Figure 50 Effect of logarithmically increasing dataset size on validation RMSE and loss ... 130

Figure 51 Effect of logarithmically increasing dataset size on regression coefficient 131

Figure 52 Training history of final CNN architecture over 100 epochs. Black markers indicate validation points 134

Figure 53 Regression plots of predicted vs actual MPP geometry parameters of 48,000 examples tested on unseen absorption responses 134

Figure 54 Impedance test rig (top) and MPP sample (bottom) 135

Table of Figures

Figure 55 Requested absorption responses vs actual responses returned from the CNN model and experimental validations of MPPs optimised for (a) 200 to 800Hz, where d=0.35mm, σ=1.0594%, D=57.7mm and (b) 700 to 1000Hz, where d=0.9mm, σ=0.42171%, D=9mm. In both cases panel thickness is kept constant at t=1mm 137

Figure 56 Requested spike inputs vs actual returned analytical solutions and experimental validations for MPP designs targeted to give maximum absorption at 800Hz. (a) narrow bandwidth requested, d=1.25mm, σ=0.60355%, D=13mm, t=1mm (b) medium bandwidth, d=0.75mm, σ=0.80806%, D=20mm, t=1mm (c) broad bandwidth, d=0.4mm, σ=1.6089%, D=40mm, t=1mm 138

Figure 57 Experimental absorption responses of MPPs optimised by the CNN to target specific frequencies between 400-800Hz 139

Figure 58 Absorption responses of MPPs optimised by the CNN to maximise absorption between 0-1000Hz using (a) spike inputs, d=0.4mm, σ=1.6089%, D=40mm, t=1mm and (b) step inputs, d=0.25mm, σ=2.3325%, D=56mm, t=1mm 140

Figure 59 Optical microscope scans of perforations manufactured by CNC (left), pillar drill (middle) and laser cutting (right). All scans are at the same scale of 150x magnification. Nominal perforation diameter in all cases is 750μm 141

Figure 60 Manufacturing method comparison for MPP designed for 800Hz; d=0.75mm, σ=0.80806%, D=20mm, t=1mm, no. perforations=37 142

Figure 61 Original absorption response vs response compensated using CNN model for (a) laser cut and (b) pillar drill experimental samples 143

Figure 62 Original vs compensated absorption responses for FDM (a) nylon and (b) PETG MPP samples designed to maximise absorption at 800Hz 144

Figure 63 (a) Model of triple parallel MPP (b) electro-acoustic equivalent circuit 145

Figure 64 Single MPP solution; d=0.45mm, σ=0.58803, D=20.6mm, t=1mm 146

Figure 65 CNN regression values for (a) perforation diameter and (b) perforation ratio 146

Figure 66 Experimental vs analytical absorption responses for parallel MPP configurations optimised via CNN model (a) double parallel MPP; d1=0.65mm, d2=0.6mm, σ1=0.5723, σ2=1.0655, D1=D2=22mm, t1=t2=1mm and (b) triple parallel MPP; d1=0.75mm, d2=0.65mm, d3=0.6mm, σ1=0.3571, σ2=0.5727, σ3=1.0655, D1=D2=D3=22mm, t1=t2=t3=1mm 147

Figure 67 Triple parallel MPP test samples for 50.8mm diameter circular impedance tube (top) and 185mm square tube (bottom) 148

Figure 68 Experimental validation of 185mm triple MPP samples for a variety of back cavity materials 149

Figure 69 Mean regression vs dataset size for design of experiments test plans 151

Figure 70 Tensile test coupon geometry. All dimensions in mm 156

Figure 71 Tensile test coupon layer orientation 156

Figure 72 Compression coupon nested build plates (a) 90° cylindrical (b) 0° cuboid (c) 45° cuboid (d) 90° cuboid 157

Figure 73 FDM equipment used for test specimen manufacture 158

Figure 74 Instron 3369 uniaxial tensile (left) and compressive (right) experimental setup ... 159

Figure 75 Material comparison - Tension ASTM D638 - Extensometer 161

Figure 76 Material comparison - Tension ASTM D638 - Load cell 161

Figure 77 Material comparison - Compression ASTM D695 162

Figure 78 Build direction comparison of Polymaker CoPA compression samples 164

Figure 79 Minimum strut radius test prints 166

Figure 80 Example of a failed build at strut radius below 0.5mm 166

Figure 81 Lattice design optimisation parameters 167

Figure 82 Solid body FEA model example boundary conditions (left) and mesh (right) 168

Figure 83 Line body FEA model example geometry (left) and mesh (right) 168

Figure 84 Frictional contact FEA model example boundary conditions (left) and mesh (right) 170

Figure 85 Non-linear buckling model project schematic implemented in ANSYS 170

Figure 86 Example of perturbed starting geometry in non-linear buckling analysis 171

Figure 87 Lattice compression testing experimental setup 171

Figure 88 Mean experimental responses of lattice blocks loaded in uniaxial compression ... 173

Table of Figures

Figure 89 Mass vs strut radius of 20x20x60mm octet-framed truss compression block 173

Figure 90 Experimental response vs solid body FEA model of lattice compression block - 1mm strut radius 174

Figure 91 Experimental response vs solid body FEA model of lattice compression block - 2mm strut radius 174

Figure 92 Experimental vs solid body FEA deformation mode shape for lattice compression block - 2mm strut radius 175

Figure 93 Experimental response vs solid body FEA model of lattice compression block - 3mm strut radius 175

Figure 94 Experimental response vs line body FEA model of lattice compression block - 1mm strut radius 176

Figure 95 Experimental response vs line body FEA model of lattice compression block - 2mm strut radius 176

Figure 96 Experimental response vs line body FEA model of lattice compression block - 3mm strut radius 177

Figure 97 Experimental vs solid body FEA response of 2mm strut radius lattice with and without contacts 178

Figure 98 Experimental vs FEA responses for beam element and solid element model with buckling for 2mm strut radius lattice 179

Figure 99 Non-linear stabilisation of lattice compression block with 1mm strut radius 180

Figure 100 Non-linear stabilisation of lattice compression block with 2mm strut radius 181

Figure 101 Non-linear stabilisation of lattice compression block with 3mm strut radius 181

Figure 102 Non-linear stabilisation on example of a previously unconverged design point at r1=0.986, r2=0.516 and r3=0.508mm 182

Figure 103 Pairwise plot of design space of final dataset of 4297 samples generated using the beam element FEA model with non-linear stabilisation 183

Figure 104 Training history of CNN for predicting lattice mechanical responses based on beam element FEA model 186

Figure 105 Regression plots of predicted vs actual lattice geometry parameters for network based on beam element FEA 186

Figure 106 Training history of CNN for predicting lattice mechanical responses based on solid element FEA model 187

Figure 107 Regression plots of predicted vs actual lattice geometry parameters for network based on solid element FEA 187

Figure 108 Training history of example network based on solid body FEA model with overfitting due to extended training time 188

Figure 109 Regression plots for example network based on solid body FEA model with overfitting due to extended training time 188

Figure 110 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 1 targeting 5kN design load 190

Figure 111 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 2 targeting 2kN design load 190

Figure 112 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 3 targeting 500N design load 191

Figure 113 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 4 targeting 500N design load 191

Figure 114 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 1 194

Figure 115 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 2 194

Figure 116 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 3 194

Figure 117 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 1 195

Figure 118 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 2 195

Figure 119 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 3 195

Figure 120 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 4 196

Figure 121 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 5 196

Figure 122 (a) Helmholtz resonator based honeycomb metamaterial configuration top, bottom, isometric and 3D views; (b) Section view of an individual Helmholtz resonator cavity with extended neck. 202

Figure 123 (a) Diagram of straight QWR with transverse inlet. Thickness of the front wall, l_{inlet} is considered separately to the length of the resonator cavity, l in analytical solution; (b) Equivalent QWR coiled into an Archimedes spiral; (c) Concentric circular labyrinth design; (d) Concentric hexagonal labyrinth design. 203

Figure 124 Triple parallel Helmholtz resonator specimen additively manufactured in two halves to allow for different neck and cavity lengths. Complete sample bonded with epoxy adhesive (left), back wall terminations are introduced at 3 different heights to allow for different cavity lengths. Specimen was assembled by joining honeycomb back section (middle) and front plate with inlet necks (right). 203

Figure 125 Parallel quarter-wavelength resonator configurations; spiral (left), concentric circular (middle) and concentric hexagonal (right). 204

Figure 126 Spiral QWR open specimen (left) and same specimen sealed over during the FDM process (right) 209

Figure 127 Impedance tube test rig (left) and parallel QWR specimen mounted in the tube (right) 210

Figure 128 Diagram of impedance tube with microphone placement configuration 210

Figure 129 Comparison between analytical and experimental absorption coefficient responses for single Helmholtz resonator-based honeycomb samples manufactured via FDM. In all cases r_n=0.7mm, A_0=67.065mm^2, l_c=5mm. For 400, 500, 600, 700 and 800Hz specimens l_n=18, 13.7, 11, 9 and 7.5mm respectively. 211

Figure 130 Comparison between specimens manufactured via FDM and SLA for HR targeted at 800Hz 212

Table of Figures

Figure 131 Experimental vs analytical absorption responses for parallel HR specimens; (a) 6 resonators, $A0$=25mm^2, rn=1mm, lc=5mm, $ln1$=36mm, $ln2$=31mm, $ln3$=26mm, $ln4$=24mm, $ln5$=20mm, $ln6$=17mm; (b) 3 resonators, $A0$=67.065mm^2, rn=0.7mm, lc=5mm, $ln1$=12.5mm, $ln2$=10.5mm, $ln3$=8.5mm. 213

Figure 132 Parallel configuration of 6 Helmholtz resonators arranged transverse to the sound wave direction. 213

Figure 133 Comparison of analytical and experimental absorption coefficient responses for spiral QWR specimens manufactured via FDM. In all cases w=h=10mm, t=1mm, $linlet = 1mm$ and overall specimen thickness = 12mm. 214

Figure 134 Comparison of analytical and experimental absorption coefficient responses for labyrinth QWR specimens; (a) hexagonal, w=5mm, h=20mm, l=270.13mm, $linlet$=1mm; (b) circular, w=5mm, h=20mm, l=297.89mm, $linlet$=1mm. 216

Figure 135 Analytical vs experimental absorption responses for labyrinth QWRs arranged in parallel; (a) 4 resonators, w=6mm, h=17mm, $l1$=86mm, $l2$=80mm, $l3$=75mm, $l4$=70mm, $linlet$=1mm; (b) 5 resonators, w=6mm, h=17mm, $l1$=77mm, $l2$=72.5mm, $l3$=67mm, $l4$=63mm, $l5$=58.5mm, $linlet$=1mm. 217

Figure 136 ColorFabb XT-CF20 base material uni-axial compression test data 219

Figure 137 (a) Cutaway example unit cell, t=1mm, w=4mm, h=10mm; (b) FEA model mesh and boundary conditions 221

Figure 138 Mesh convergence study for (a) lowest volume, (b) greatest volume, (c) weakest, (d) strongest unit cell designs possible within the chosen design space limits 222

Figure 139 FEA mesh element size vs simulation solve time 223

Figure 140 Pairwise plot of Latin hypercube design of experiments test plan of 4500 design points 224

Figure 141 Range of possible FEA responses for chosen design space limits. 27 design points were selected from a full factorial design of experiments test plan with 3 levels based on the limits of the 3 geometry parameters. 224

Figure 142 Example dataset design point observation signal used as input for the convolutional neural network; (a) analytical sound absorption-frequency response; (b) FEA force-displacement structural response for equivalent geometry; (c) final sequence vector

generated from combination of raw sound absorption and normalised structural characteristics 225

Figure 143 Inverse convolutional neural network architecture schematic 227

Figure 144 Training history of convolutional neural network 229

Figure 145 Regression between inverse network predictions and known unit cell geometry parameters 229

Figure 146 Manufactured compression test specimens 232

Figure 147 Test rig setup for experimental mechanical characterisation of optimised test samples 232

Figure 148 Sound absorption responses and mechanical characteristics of example optimised metamaterials 236

Figure 149 Non-linear buckling analysis FEA model schematic 237

Figure 150 Example geometry perturbed according to the first buckling mode shape using a scale factor of 1 (left), 2 (middle) and 4 (right) 237

Figure 151 (a) Comparison of absorption responses for example A recorded on an impedance tube with 185mm square cross section vs 50.8mm diameter circular tube; (b) cutaway view of 185mm square test specimen 239

Figure 152 Analytical vs experimental absorption responses for parallel QWR specimens based on optimised example A, t=0.98mm, w=6.19mm, h=18.56mm, l_{inlet}=1mm ; (a) 4 resonators, l_1=92mm, l_2=85mm, l_3=78mm, l_4=71mm; (b) 5 resonators, l_1=81mm, l_2=76mm, l_3=70mm, l_4=63mm, l_5=57mm 240

Figure 153 Geometry layout for parallel QWR specimens; (a) 4 parallel resonators; (b) 5 parallel resonators 240

Figure 154 Schematic of predictor ANN 242

Figure 155 Test mean squared error and regression vs number of hidden layer neurons for networks trained to predict (a) relative specific stiffness and (b) relative specific strength .. 243

Figure 156 Training state history of neural networks for predicting; (a) relative specific stiffness; (b) relative specific strength 245

Table of Figures

Figure 157 Training performance history of neural networks for predicting; (a) relative specific stiffness; (b) relative specific strength 245

Figure 158 Test regression between known target values and output predictions of (a) relative specific stiffness; (b) relative specific strength 246

Figure 159 Histograms of error between known target values and output predictions of (a) relative specific stiffness; (b) relative specific strength 246

Figure 160 Multi-objective genetic algorithm optimiser history 249

Figure 161 Pareto front of multi-objective GA solutions..... 249

Figure 162 Performance objective histograms of Pareto front solutions 250

Figure 163 Pareto solution 4 material properties; (a) sound absorption response; (b) mechanical response..... 252

Figure 164 Pareto solution 2 material properties; (a) sound absorption response; (b) mechanical response..... 252

Figure 165 Pareto solution 18 material properties; (a) sound absorption response; (b) mechanical response..... 252

Figure 166 Pareto solution 5 material properties; (a) sound absorption response; (b) mechanical response..... 253

Figure 167 Comparison between experimental structural properties of genetic algorithm optimised Pareto front solutions and equivalent maximum values found by exhaustive search of the dataset of 4248 design points for (a) relative specific stiffness; (b) relative specific strength 254

Figure 168 Revised metamaterial unit cell geometry for reduced overall thickness; (a) cutaway view (b) 3D view..... 255

Figure 169 Comparison of analytical and experimental sound absorption responses for unit cell revised to reduce overall thickness at greatest possible channel length, where t=0.9mm and w=3mm; (a) h=2mm, overall thickness=3mm; (b) h=3mm, overall thickness=4mm; (c) h=4mm, overall thickness=5mm 257

Figure 170 FEA model loading and boundary conditions for unit cell revised to reduce overall thickness 258

Table of Figures

Figure 171 Range of possible FEA responses for reduced thickness model design space. Design points were selected from a full factorial test plan with 3 levels based on the chosen geometry parameter limits. 259

Figure 172 Training history of inverse convolutional neural network for reduced thickness model 260

Figure 173 Regression between inverse network predictions and known unit cell geometry parameters for metamaterials with reduced overall thicknesses limited to below 6mm 261

Figure 174 Example manufactured specimens for acoustic impedance testing (top) and mechanical compression testing (bottom) 263

Figure 175 Sound absorption responses and mechanical characteristics of example metamaterials optimised for overall material thicknesses below 6mm 266

Figure 176 Impedance tube layout for measurement of transmission loss 267

Figure 177 Experimental transmission loss measurements for reduced thickness model test specimens; (a) example O, t=1.23mm, w=3.46mm, h=4.66mm, overall thickness=5.66mm; (b) example I, t=0.93mm, w=3.71mm, h=2.64mm, overall thickness=3.64mm 267

Figure 178 Parametric unit cell layout; (a) outer frame nodes with fixed positions; (b) strut connections between nodes at arbitrary default positions 271

Figure 179 Strut radius geometry parameters assigned to 2x2 symmetric supercell 273

Figure 180 Polymaker CoPA nylon non-linear material data used in static ANSYS model . 274

Figure 181 Static FEA model mesh and boundary conditions. Example shown has arbritary default positions and all strut radii set to 1mm 275

Figure 182 Tangent modulus extraction by least squares regression 276

Figure 183 High velocity explicit dynamic FEA model mesh and loading. Example shown has arbritary default positions and all strut radii set to 1mm. 277

Figure 184 High velocity explicit dynamic FEA model mesh and fixed support. Example shown has arbritary default positions and all strut radii set to 1mm. 277

Figure 185 Comparison of high velocity dynamic impact simulation without (top) and with (bottom) steel impactor modelled. Examples use the same geometry parameters with arbitrary default positions and all strut radii set at 0.5mm 278

Table of Figures

Figure 186 Training histories for CNNs to predict (a) relative specific stiffness, (b) relative specific strength, (c) specific energy absorption, (d) Mass 281

Figure 187 Regression between convolutional neural network predictions and known values for test dataset of 484 observations 282

Figure 188 Optimiser history of multi-objective genetic algorithm set to maximise relative specific stiffness and strength; (a) mean distance between individuals; (b) mean Pareto spread 285

Figure 189 Pareto front of solutions from GA set to optimise relative specific stiffness and strength simultaneously 286

Figure 190 Octet-truss supercell geometry 287

Figure 191 Comparison of Pareto front solutions against equivalent octet-truss at same relative density for normalised stiffness (left) and normalised strength (right) 287

Figure 192 Percentage increase in performance from AI generated unit cell compared to equivalent octet-truss for normalised stiffness (left) and normalised strength (right) 287

Figure 193 Optimiser history of multi-objective genetic algorithm set to minimise mass in addition to maximising relative specific stiffness and strength; (a) mean distance between individuals; (b) mean Pareto spread 288

Figure 194 Pareto front of solutions from GA set to minimise mass while simultaneously maximising relative specific stiffness and strength 289

Figure 195 Performance objective histograms for GA set to minimise mass while simultaneously maximising relative specific stiffness and strength 290

Figure 196 Comparison of revised Pareto front solutions against equivalent octet-truss at same relative density for normalised stiffness (left) and normalised strength (right) 290

Figure 197 Percentage increase in performance from revised AI generated unit cell compared to equivalent octet-truss for normalised stiffness (left) and normalised strength (right) 291

Figure 198 Polynomial regression models for prediction of relationships between performance objectives and relative density; (a) normalised stiffness; (b) normalised strength 292

Figure 199 Polynomial regression models for prediction of relationships between supercell mass and relative density 293

Table of Figures

Figure 200 Estimated percentage mass reduction between revised AI generated unit cell and equivalent octet-truss for a given normalised stiffness (left) and normalised strength (right) 293

Figure 201 Optimiser history of multi-objective genetic algorithm set to maximise specific energy absorption in addition to relative specific stiffness and strength while minimising mass; (a) mean distance between individuals; (b) mean Pareto spread 294

Figure 202 Pareto front of solutions from GA set to maximise SEA, relative specific stiffness and strength while minimising mass. Black markers indicate outliers. 294

Figure 203 Comparison of specific energy absorption of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in SEA from AI generated lattice (right) 295

Figure 204 Comparison of normalised stiffness of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in normalised stiffness from AI generated lattice (right) 296

Figure 205 Comparison of normalised strength of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in normalised strength from AI generated lattice (right) 296

Figure 206 Example Pareto solution 31 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 40.3%, 120.4% and 66.4% respectively at relative density of 53.8%. 297

Figure 207 Example Pareto solution 56 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 37.8%, 106.7% and 55.3% respectively at relative density of 53.5%. 297

Figure 208 Example Pareto solution 25 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 35.9%, 120.5% and 64.4% respectively at relative density of 57.4%. 297

Figure 209 Example Pareto solution 42 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 34.6%, 100.4% and 47.8% respectively at relative density of 41.7%. 298

Figure 210 Supercell geometries for Pareto solutions; (a) solution 31; (b) solution 56; (c) solution 25; (d) solution 42 298

Figure 211 Composite plate FEA model mesh and boundary conditions 301

Figure 212 Composite plate FEA model layup stacking sequence 302

Figure 213 Network training performance evolution 305

Figure 214 Network H training history 307

Figure 215 Network H regression plots between predicted and actual known layup angles . 307

Figure 216 Baseline composite stress responses; example 1 (top) and example 2 (bottom) . 308

Figure 217 Comparison of requested and actual stress responses - Example 1 310

Figure 218 Comparison of requested and actual stress responses - Example 2 312

Figure 219 Schematic of predictor neural network model. Network architecture is figurative and simplified for illustrative purposes. Actual architectures used featured a greater number of neurons and hidden layers. Labels w and b represent weights and biases of network channels. 314

Figure 220 Training dataset design space - 2000 sample Latin hypercube test plan 315

Figure 221 (a) Shallow artificial neural network architecture; (b) Test mean squared error and regression vs number of hidden layer neurons; (c) Training time vs number of hidden layer neurons 317

Figure 222 ANN training state 318

Figure 223 ANN training performance 318

Figure 224 ANN test regression between known target stress values and output predictions 319

Figure 225 Error histogram for entire dataset of 200,000 observations..... 319

Figure 226 Absolute error between known stress values and values predicted by the ANN for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan. 320

Figure 227 Mean absolute error averaged across the composite plate for 500 test layups – ANN 321

Figure 228 Predictor network CNN schematic layout 323

Figure 229 Training history of network DNN-10 326

Figure 230 Training history of network CNN-12 326

Figure 231 Regression plots between predicted and actual normalised stress values for networks (a) DNN-10 and (b) CNN-12. Each plot features 30,000 datapoints generated from the test partition of the dataset. 327

Figure 232 Absolute error between known stress values and values predicted by DNN-10 for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan. 328

Figure 233 Absolute error between known stress values and values predicted by CNN-12 for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan. 328

Figure 234 Mean absolute error averaged across the composite plate for 500 test layups (a) DNN-10 and (b) CNN-12 329

Figure 235 Reduced mesh element centroid locations 330

Figure 236 Comparison between FEA and CNN predicted composite stress contours. Test sample numbers correspond to those in Figure 234. 333

Figure 237 Training history of best performing maximum stress predictor CNN, trained on a dataset of 200,000 observations 335

Figure 238 Regression between normalised predicted and actual known values for the best performing maximum stress predictor CNN 335

Figure 239 Absolute error between predicted and known maximum stress values for best performing CNN 336

Figure 240 Maximum Von Mises stress in the composite plate for the optimised layup sequences 339

Figure 241 Genetic Algorithm Test 2 optimisation history (left) and solution layup FEA stress response (right) 340

Figure 242 Particle Swarm Test 2 optimisation history (left) and solution layup FEA stress response (right) 340

Figure 243 Simulated Annealing Test 5 optimisation history (left) and solution layup FEA stress response (right) 340

Figure 244 Example geometry and boundary conditions for the parametric predictor model. Red line indicates load, green line is fully fixed support 343

Figure 245 Parametric predictor model training dataset design of experiments test plan. 2000 datapoints generated via Latin hypercube sampling. 344

Figure 246 Training history of best performing CNN model for parametric stress prediction 345

Figure 247 Regression between predicted and actual known normalised Von Mises stress values for parametric predictor CNN. Plot features 168,131 datapoints corresponding to separate test observations. 346

Figure 248 Mean absolute error between CNN predictions and FEA values of Von Mises stress averaged across each composite plate for 500 test cases of selected geometry, loading and layup combinations 347

Figure 249 Maximum stress parametric predictor CNN training history 348

Figure 250 Maximum stress parametric predictor CNN regression between predicted and actual normalised maximum Von Mises Stress 349

Figure 251 Maximum stress predictor CNN error histogram for test dataset of 30,000 observations 349

Figure 252 Pairwise plots of 50 test geometry and loading conditions chosen by Latin hypercube sampling 350

Figure 253 Exhaustive search test plan. 1000 layups chosen by Latin hypercube sampling. 351

Figure 254 Comparison of maximum Von Mises stress for optimised test layups against smallest and largest values found by exhaustive search 352

Figure 255 Optimised composite layup prediction consistency for genetic algorithm and particle swarm optimisation 353

List of Accompanying Materials

Dataset DOI: 10.5258/SOTON/PG/D166

Research Thesis: Declaration of Authorship

Print name: Peter Anthony Hawes

Title of thesis: Machine Learning for Tailored Multifunctional Properties of Lightweight Metamaterials

I declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:-

Hawes, P., Boccaccio, M. and Meo, M., 2025. Design of subwavelength wide bandwidth sound absorbers by inverse convolutional neural networks. Applied Acoustics, 231, p. 110543. DOI: 10.1016/j.apacoust.2025.110543

Signature: .....Date:...28/04/2026.....

Acknowledgements

Professor Michele Meo – Primary supervision, assistance with conceptualisation, editing, provision of resources for literature review

Dr Jie Yuan – Secondary supervision, advice and feedback

Dr Marco Boccaccio – Supervision, advice on methodology, sound absorption theory and training for acoustic impedance testing

Dr Gian Piero Malfense Fierro – Advice on unit cell layout for section 5.3

Dr Kostas Myronidis – Assistance with setup of additive manufacturing equipment and manufacture of test specimens

Stuart Findlow – Provided training on electromechanical testing machines and assistance with setup of machinery and strain measurement devices for mechanical testing

Permission to reproduce figures 1, 3 and 10 has been granted by the Copyright Clearance Center (CCC) on behalf of The Royal Society (U.K.)

Permission to reproduce figure 2 has been granted by the Copyright Clearance Center (CCC) on behalf of Springer Nature

Definitions and Abbreviations

AI.....Artificial Intelligence

AM .....Additive Manufacturing

NREL .....National Renewable Energy Laboratory

EV.....Electric Vehicle

PHEV .....Plug-in Hybrid Electric Vehicle

EPA .....Environmental Protection Agency

FEA .....Finite Element Analysis

TPMS .....Triply Periodic Minimal Surface

DTPMS.....Diamond Triply Periodic Minimal Surface

GTPMS.....Gyroid Triply Periodic Minimal Surface

PTPMS .....Primitive Triply Periodic Minimal Surface

IWPTPMS .....Isotropic Woodpile Triply Periodic Minimal Surface

NTPMS.....Neovius Triply Periodic Minimal Surface

BCC.....Body-centred Cubic

FCC .....Face-centred Cubic

GRF .....Gaussian Random Field

DfAM .....Design for Additive Manufacture

STL.....Stereolithography file format used in 3D printing

SIMP.....Solid Isotropic Material with Penalisation

RAMP.....Rational Approximation of Material Properties

ESO .....Evolutionary Structural Optimisation

BESO.....Bidirectional Evolutionary Structural Optimisation

GA .....Genetic Algorithm

DSM .....Direct Stiffness Method

UTM.....Upper Triangular Matrix

PSO.....Particle Swarm Optimisation

ANN .....Artificial Neural Network

DNN .....Deep Neural Network

ReLU .....Rectified Linear Unit

Adagrad .....Adaptive gradient algorithm

RMSprop .....Root Mean Squared propagation

Adam .....Adaptive moment estimation

Epoch.....When the entire training dataset has been passed forward and backward through the neural network once this is referred to as one epoch.

RMSE .....Root Mean Square Error

CNN .....Convolutional Neural Network

ME.....Manufacturable Element

FDM .....Fused Deposition Modelling

ABS .....Acrylonitrile Butadiene Styrene

PCA .....Principal Component Analysis

RVE.....Representative Volume Element

GAN .....Generative Adversarial Network

CAD .....Computer Aided Design

CNC.....Computer Numerical Control

SLS .....Selective Laser Sintering

TLP.....Transient Liquid Phase bonding

CFRP .....Carbon Fibre Reinforced Polymer

PP .....Polypropylene

SLA .....Stereolithography

UV .....Ultraviolet

SLM.....Selective Laser Melting

EBM .....Electron Beam Melting

FFF .....Fused Filament Fabrication

PVA.....Polyvinyl Alcohol

PLA .....Polylactic Acid

CCFRAM .....Continuous Carbon Fibre Reinforced Additive Manufacture

CCFRP .....Continuous Carbon Fibre Reinforced Polymer

DLW.....Direct Laser Writing

TPL.....Two-Photon Lithography

ALD.....Atomic Layer Deposition

SEA .....Specific Energy Absorption

FE .....Finite Element

HOBE .....Honeycomb Before Expansion

BMC .....Bulk Metallic Glass

NVH .....Noise, Vibration and Harshness

STL.....Sound Transmission Loss

HR .....Helmholtz Resonator

LMD .....Light Medium Duty

PEL.....Passenger Ear Level

PBN .....Pass By Noise

PU.....Polyurethane

HREN .....Helmholtz Resonator with Extended Neck

MPP .....Microperforated Panel

PPA.....Perforated Panel Absorber

HRIN .....Helmholtz Resonator with Inserted Neck

PET.....Polyethylene Terephthalate

AFR .....Airflow Resistivity

RAM.....Random-access Memory

GPU .....Graphics Processing Unit

CPU .....Central Processing Unit

ASTM.....American Society for Testing and Materials

PETG.....Polyethylene Terephthalate Glycol

Definitions and Abbreviations

QWR.....Quarter-wavelength Resonator

LVDT .....Linear Variable Displacement Transducer

MSE.....Mean Squared Error

APDL .....Ansys Parametric Design Language

SA.....Simulated Annealing

Chapter 1 Introduction

1.1 Background

Throughout history materials science has defined the limits of society's technology, setting the boundaries of what is achievable in component design in terms of mass, strength, cost, ease of manufacture and aesthetics. Eras such as the stone age, bronze age and iron age are defined by the materials at the core of the civilization's technological advancement. Improvements in materials are able to filter down to affect all areas of industry.

The modern day is an era of smart materials, a digital age defined through advances in computing power which has birthed the emergence of artificial intelligence (AI) algorithms that already affect our lives on a daily basis, influencing what products we buy, evaluating and selecting courses of action, proving mathematical theorems and helping to diagnose diseases [1]. The evolution of additive manufacturing (AM) into the mainstream allows for new unprecedented levels of collaboration between mankind and artificial intelligence in engineering product design. Furthermore, structural design architectures can now be applied not only at component level, but also at the material level. This allows for tailoring of cellular microstructures for a given application, leading to improved structural efficiency and unlocking previously unseen degrees of lightweighting and multifunctionality that exceed conventional monolithic materials and pre-existing composite hybrids, with the potential for large scale national and societal impact.

The laws of physics tell us that the lighter a vehicle is, the less energy is required to accelerate it. The US spends on average 26% of its total energy output on transportation, of which 90% is accounted for by petroleum products [2]. The US department of energy states "Advanced materials are essential for boosting the fuel economy of modern automobiles while maintaining safety and performance" [3]. They estimate a 10% reduction in vehicle weight can improve fuel economy of vehicles by 6-8%, projecting by 2030 a saving of 5 billion gallons of fuel each year in the US alone could be possible if only a quarter of their vehicles made use of advanced lightweighting materials. In addition to making more efficient use of depleting fuel reserves, UCS estimates for every 100kg of mass reduction there will be a 3-5% decrease in carbon emissions [4]. Scientists at the National Renewable Energy Laboratory (NREL) predict the importance of vehicle lightweighting for reducing costs and increasing lifespan of batteries in electric (EV) and plug-in hybrid (PHEV) vehicles as society switches to electricity as the primary energy source for road transport in the future [5]. It is of vital

importance that costs of electric powertrains reduce and EV range increases to become more competitive with conventional combustion engines. A study conducted by NREL found that a 150kg mass decrease resulted in a reduction in battery energy consumption of EVs by 5% on EPA drive cycles, by reducing rolling resistance and improving braking and powertrain efficiency. They found that the costs of lightweighting must be kept below $6/kg of mass eliminated to remain cost effective for mass market passenger vehicles, highlighting the need for cost effective lightweighting materials for a low carbon, sustainable future.

1.1.1 Material Property Space

Material properties are governed by performance trade-offs. For example the strength of a material is closely related to how lightweight it can be and typically has an adverse effect on toughness. Stiffness often comes at the cost of increased density and reduced energy dissipation in impacts. Material property relations can be visualised using pair-wise material property charts that represent slices of multi-dimensional material property space (Figure 1). Merit indices can then be used to mathematically define performance criteria, such as E/ρ, $E^{1/2}/\rho$ and $E^{1/3}/\rho$ to represent maximum stiffness at lowest density [6]. These index lines can then be plotted for use as guide lines for minimum mass design. Materials that lie above the line will produce stiffer and lighter components. All combinations of material properties will feature “white space” voids that remain unpopulated by currently existing materials [7]. In many cases these regions will always remain inaccessible for fundamental reasons. However in other spaces holes can in theory be filled, in some governing performance trade-offs may be reduced in extent or broken completely. Over the years many methods have been developed in attempts to fill holes in material property space. One approach is through manipulation of material chemistry to create new monolithic materials such as metal alloys, polymers and ceramics. Another is to manipulate microstructure through heat treatments and mechanical processing to control phase distribution and remove defects [8]. These approaches have been long established and largely exhausted over the years. An alternative method introduces design architecture at the material level, breaking performance trade-offs by defining material mechanical properties through geometry rather than chemical composition.

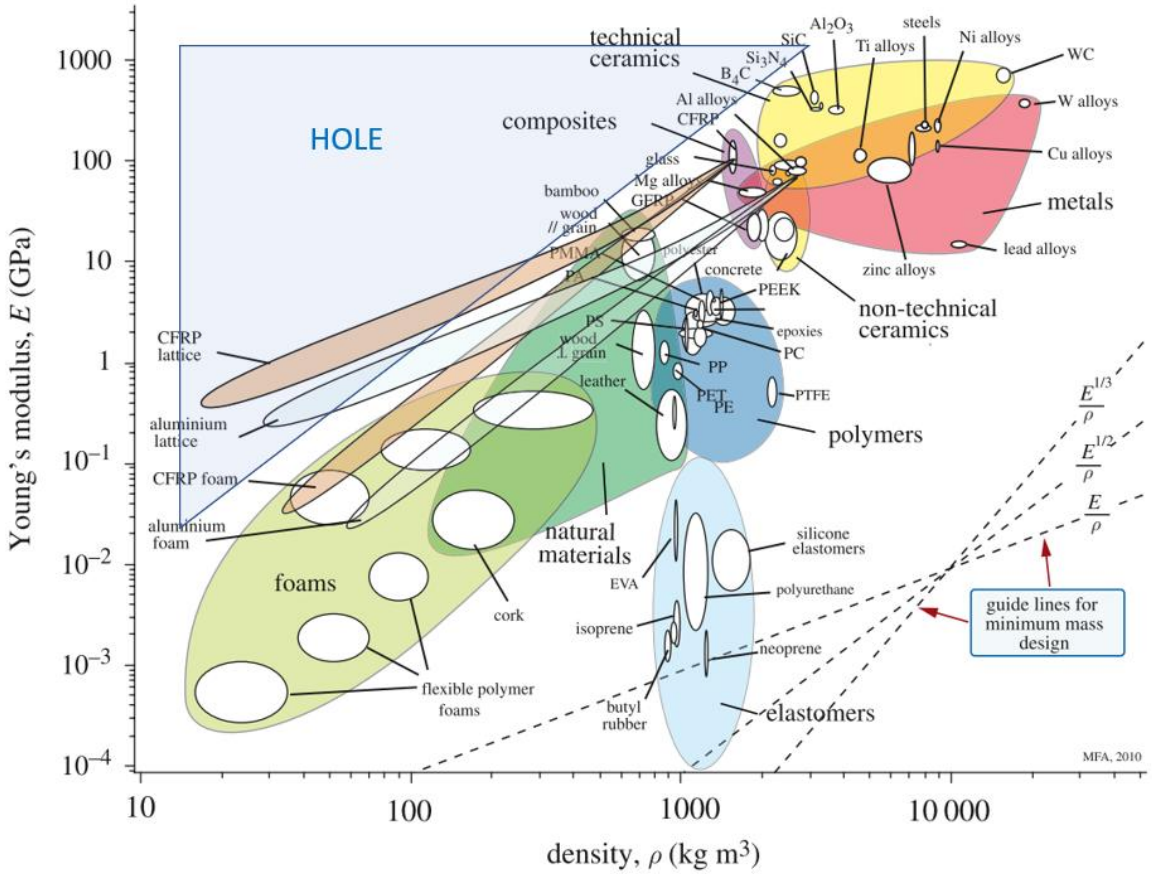


Figure 1 Modulus-density material properties space. Carbon fibre reinforced polymer (CFRP) and aluminium lattice structures are predicted to extend from the location of the legacy material to fill holes in the design space by defining material properties through geometry rather than chemical composition. Adapted from [8]

1.1.2 Architectured Cellular Materials

Architectured materials were first defined by Michael Ashby as hybrid materials that are combinations of two or more materials, or of materials and space which are assembled in configurations that enhance the performance of the combined substance, offering new attributes that cannot be found in the constituent parts alone [9]. Their existence is typically driven by a clear objective to optimise a certain set of properties for a given application. Nature is an excellent designer of architectured structures. The natural evolutionary process of variation and selection has been engineering structurally efficient cellular solid materials for millennia, made from interconnected networks of solid struts or plates forming the edges and faces of unit cells [10]. Leaves feature parallel skins reinforced by stiff ribs and separated by low density foam-like cellular cores in order to grow upwards towards the sun to capture light for photosynthesis without collapsing under their own weight. Scaling of tessellating cellular arrangements have formed strong, impact resistant bone structures that protect animals against

predators. The human skeleton is shaped by spongy, shock absorbing trabecular bone sandwiched between dense cortical outer layers, which grows by naturally depositing material along load paths in the directions of greatest stress (Figure 2). High porosity hollowed core geometries have developed to produce remarkable stiffness to weight characteristics in the bones of birds that has enabled flight.

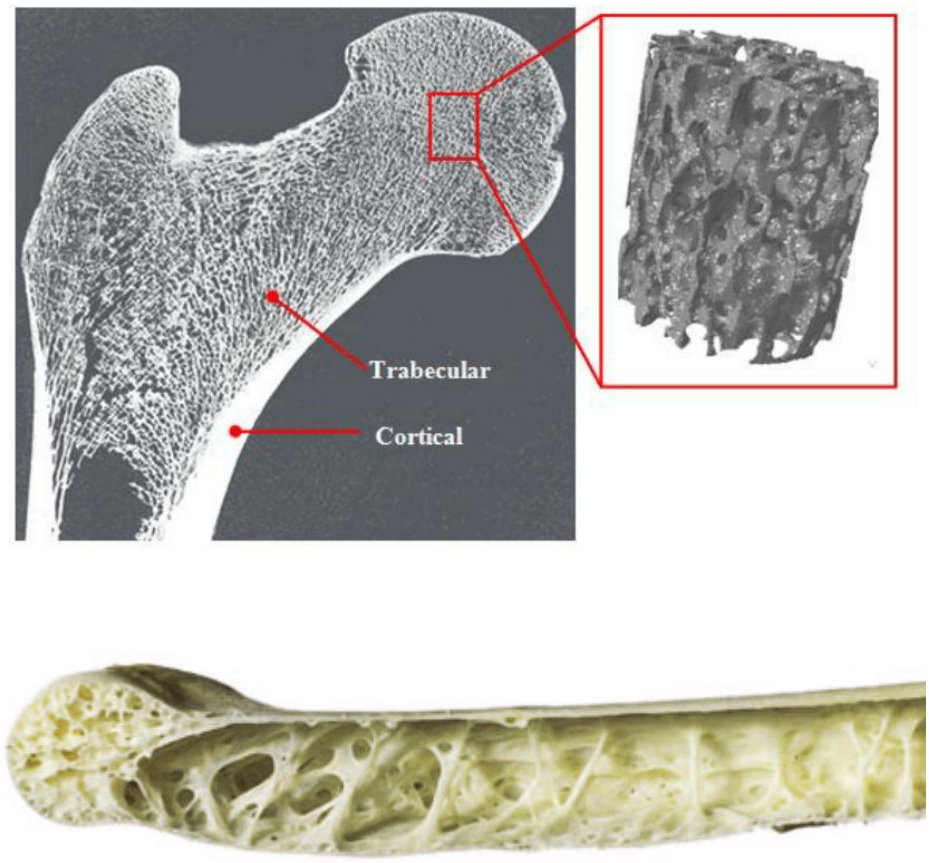


Figure 2 Cross section of human femur showing trabecular and cortical bone (top) [11] and cross section of a bird wing (bottom) [12]

The randomised, stochastic evolutionary approach to material design has the advantage of growing geometry adaptively to a specific loading scenario at all scales. The thicknesses and orientations of individual strut elements depend on the magnitude and direction of loading. Each strut naturally aligns in the direction of principal stress trajectories.

Relative density is defined as the ratio between the density of a cellular structure and the solid material from which it is made:

$$\bar{\rho} = \frac{\rho}{\rho_s} \tag{1.1}$$

A key problem with naturally occurring cellular materials is that many exhibit quadratic or greater scaling between relative density and stiffness or strength [13]. That is to say $E/E_s \propto (\rho/\rho_s)^n$ and $\sigma_y/\sigma_{ys} \propto (\rho/\rho_s)^n$ where E is Young's modulus and σ_y is yield strength and the scaling power n is greater than 2. This leads to a rapid drop off in mechanical performance at ultra-low density due to bending dominated behaviour of structural elements. A solution to this problem comes in the form of manmade metamaterials, materials created by design not by chance. Metamaterials are able to maintain near linear scaling between stiffness and density spanning over orders of magnitude through use of stretch-dominated rather than bend-dominated unit cell architecture, which carries loads axially along struts loaded in tension or compression. This has previously largely been achieved through use of periodic rather than stochastic lattice structures of uniform repeating tessellated unit cells. Such architectures like honeycomb, octet-truss or Kagomé lattices are an expedient way to maximise stretch-dominated behaviour independently of the loading scenario, however room exists for improving structural efficiency by organising strut topology corresponding to the actual acting principle tensile and compressive stresses. The most structurally efficient design will feature as many struts as possible oriented in the direction of localised load paths while guaranteeing structural stability. By using artificial intelligence to design periodic cellular structures it may be possible to gain the advantages of both natural and manmade approaches, achieving largely stretch-dominated architecture that breaks performance trade-offs and maintains strength and stiffness at ultra-low density, while enabling localised cell optimisations tailored to a specific application. Machine learning is able to simulate natural evolution through automation of the process of variation and selection, increasing the number of design iterations possible, producing and analysing thousands of alterations in the time taken for a human to produce a single design.

Lattices can be open or closed-cell. Closed-cell honeycombs while being lightweight can withstand large compressive strains at near constant stress, allowing them to absorb the kinetic energy of impacts without generating high peak stresses [14]. This makes them popular for use as core material in automotive composite sandwich panels and crash safety structures. However conventional honeycomb core material is highly anisotropic, exhibiting lower strength and stiffness in lateral directions. Its susceptibility to unique failure modes such as cell wall buckling limits improvements in scalability of material properties with reduced density. At ultra-low density open cell lattices are able to maintain uni-axial loading along the direction of their struts, exceeding strength of equivalent honeycombs. The open cell space can also be used to enable additional multifunctional features such as energy storage or allowing fluid flow for active cooling. Honeycomb cores are difficult to fit to

curved contours and require expensive time consuming machining processes. Additively manufactured lattices allow for net-shape composite core structures compatible with more complex geometries with graded, location specific material properties. Conventional engineered cellular materials have the drawback of inherent structural redundancies. For a given loading scenario all of the cellular structure will not necessarily be used to support the load and many individual struts will not be perfectly optimised for the part they have to contribute towards the overall bulk mechanical performance. If the loading conditions are known and well defined the material can be tailored to remove redundant features and perform local optimisations on cell architecture. Any features that add to the overall mass but don't significantly contribute towards satisfying strength and stiffness objectives, or offer multifunctionality for a given requirement can be removed. A key focus of this project is aimed at enabling tailored, application specific structural efficiency optimisations, using machine learning to expand the limitations of materials property space while activating additional levels of multifunctionality, simultaneously performing multi-objective optimisations of diverse physical properties which cannot be related through conventional theoretical means.

1.2 Research Aim and Objectives

This project aims to **develop machine learning based optimisation frameworks that allow for tailoring of material functionalities according to application-specific design requirements.** Neural network models are developed and proposed as approaches for prediction of geometry parameters of cellular architectures necessary to provide desired acoustic and mechanical responses. The developed workflows are then expanded to enable activation of additional multifunctional features through consideration of multiple performance objectives simultaneously. The versatility of the framework is demonstrated over a selection of case studies in which the mechanical properties, sound absorption performance and impact resistance of metamaterial and composite architectures are tuned according to practical design criteria. The target of the research is to maximise potential for lightweighting in aircraft and motor vehicles, leading to reduced fuel consumption and carbon emissions, improved refinement through augmented noise, vibration and harshness, and enhanced impact resistance for improved crash safety. This aim is achieved through completion of the following objective milestones:

1. Generation of datasets of material responses by analytical theory and numerical simulation. Sound absorption coefficient-frequency responses of metamaterial unit cell configurations are generated analytically. Mechanical stress-strain relationships are

generated via quasi-static finite element analysis, while dynamic simulations are performed to characterise specific energy absorption under high velocity impact. Parametric models are created to automate cycling through design points. For each case study a minimum dataset size is targeted with order of magnitude $1x10^3$ observations.

2. Physical testing and validation of analytical and numerical models used for prediction of material properties. Acoustic impedance tube testing is carried out under industry standard ASTM E-1050, while mechanical characterisations are performed following ASTM D695 and ASTM D638 to measure base material properties and assess the effectiveness of the solutions when subjected to real-world loading.

3. Establishment of a framework for development of inverse neural network architectures able to output geometry parameters needed for a given desired material response. The networks are required to possess sufficient prediction accuracy to allow fine tuning of a vector of material responses matched to a given input signal requested by the user. Where large training datasets ($1x10^5$ observations or greater) can be generated the coefficient of regression between network predictions and known values is targeted at greater than 0.9 and no less than 0.7. In cases where the dataset size is limited by the simulation method, the minimum regression is targeted at 0.7 and the network should demonstrate responsiveness to requested inputs able to perform the required function.

4. Application of the developed inverse neural network framework to optimise micro-perforated panel sound absorber metamaterials tailored to a given target frequency and absorption amplitude. Demonstration of the optimiser is carried out maximising absorption of individual unit cells at discrete target frequencies in the low-end regime (<1000Hz) in addition to providing the greatest absorption over the widest frequency bandwidth at subwavelength thicknesses. The model is used to design combinations of multiple unit cells in parallel targeting no less than 50% absorption between 500-1000Hz at overall material thickness less than 25mm.

5. Tuning of force-displacement structural responses of additively manufactured polymer metamaterials under uni-axial compression. Inverse neural networks are used to tune the radii of key struts within a truss-based lattice structure to satisfy a variety of mechanical performance targets to simulate practical design scenarios. The lattices are tailored to obtain the lightest solutions that maintain targeted safety factor limits

between 2-2.5 at design loads ranging from 0.5-5kN with stiffness that can be tuned to offer deflections between 0.5-1mm.

6. Successful multi-objective optimisations of sound absorption and mechanical properties in the same analysis. A metamaterial unit cell is designed to allow simultaneous tuning of acoustic and structural responses through a set of integrated geometry parameters. Structural responses from finite element analysis simulations are combined with analytical sound absorption-frequency responses in the same input signal. An inverse neural network is used to provide solutions with maximum absorption at lowest frequencies with tuneable stiffness and strength. Absorption amplitude is targeted greater than 90% at frequencies below 400Hz for thicknesses less than 25mm.

7. Synthesis of metamaterial architectures using artificial intelligence based methods to determine the layout of nodes and struts within the unit cell to maximise both static and dynamic structural performance simultaneously. Neural networks are trained from datasets collected from both quasi-static implicit simulations and high velocity impact simulations using explicit dynamics analysis. The specific stiffness, strength and energy absorption of the resulting designs are compared against equivalent pre-existing architectures formed from geometric primitives across a range of relative densities between 10-90% to analyse potential for lightweighting and improvements in high velocity impact resistance.

8. Proof of concept for machine learning optimisation of the layup stacking sequences of carbon fibre reinforced polymer laminates for minimum stress under a given loading, enabling weight reductions and strength enhancements in composite component design. An inverse neural network model is created that allows the stress contours through an example composite laminate plate to be dictated by providing the desired stress response as input. The network outputs the layup stacking angles necessary to provide greatest agreement with the request, allowing tailoring of load paths. An alternative model is explored working in the opposite direction, predicting the stress at a given location in the composite for a certain stacking sequence. This model is used to provide an early proof of concept that machine learning models may be used as a faster alternative to FEA in composite component design, able to rapidly predict the stress distribution in addition to providing the layup that minimises the maximum stress in the laminate.

1.3 Outline of the Thesis

The thesis structure is arranged over eight chapters as follows:

Chapter 1: Introduction

An introduction to the major topics of the research and background information including definition of metamaterials and explanation of the main issues in design of architected cellular materials. The motivation for the research and reasons why it is of importance are presented along with the potential benefits that may result from the outcome of the project. A demonstration is provided of how the project aligns with current key research priority areas of advanced materials and devices, low carbon futures and sustainability. The core aim of the project is outlined and the objectives that are accomplished to achieve this aim are introduced.

Chapter 2: Literature Review

An in depth review of the key literature relevant to the research area of the project. The main theory and governing equations of metamaterial design are presented. Previous historic works of substantial importance are reported and an analysis is made of the current state of the art of architected materials, identifying the major research gaps that will be filled by the project. An outline is provided of the major previous approaches for design optimisation of metamaterials, including both manually generated periodic unit cells and numerically generated stochastic geometries. The fundamentals of topology optimisation, evolutionary structural optimisation and machine learning are presented. The potential advantages of applying machine learning to the design of metamaterials are explored by analysis of previous works. An overview is given of the base materials and manufacturing methods available to produce complex architected lattice geometries, including mechanical assembly and additive manufacturing processes. The enhancements in material properties obtained in previous studies through the use of multi-scale hierarchy are provided. Methods are outlined that have been used to bestow inherent multifunctional features on architected metamaterials, with a focus on sound absorption and energy absorption for impact resistance.

Chapter 3: Design of Subwavelength Wide Bandwidth Sound Absorbers by Inverse Convolutional Neural Networks

Microperforated panel sound absorbers are tuned by an inverse convolutional neural network model to provide maximum absorption strength and bandwidth range. The materials are optimised to provide peak absorption at ultra-low frequencies to filter out road noise in automotive vehicles while maintaining subwavelength thicknesses. Designs are validated

experimentally by acoustic impedance testing. An assessment of network architecture, hyperparameters and dataset size requirements is carried out that is used to inform the studies in further chapters.

Chapter 4: Artificial Intelligence Based Optimisation of Metamaterials for Tailored Stress-Strain Response and Enhanced Structural Efficiency

An inverse machine learning approach is proposed for optimising the structural efficiency of lattice structures based on training deep neural networks on a database of stress-strain responses, enabling tailored mechanical performance. A convolutional network is trained on non-linear FEA responses to predict the required metamaterial geometry parameters to produce a desired stress-strain curve. The ability of the model to fine-tune structural properties is assessed and results are validated by experimental mechanical characterisations.

Chapter 5: Machine Learning Based Multi-objective Structural Optimisation of Lightweight Sound Absorber Metamaterials

The development of a deep learning model is presented to allow for simultaneous multi-objective optimisations of sound absorption and structural properties. An inverse convolutional neural network is trained to output the geometry parameters of a quarter-wavelength resonator labyrinth metamaterial to provide desired combinations of acoustic and mechanical properties, granting the user control for fine-tuning responses. An alternative approach is then examined that works in the opposite direction, in which the network outputs predictions of the sound absorption and structural performance which are then maximised by an evolutionary optimisation algorithm. The unit cell layout is then revised to limit overall material thicknesses and the inverse neural network is used to provide strongest absorption for the lowest frequencies at the reduced thickness.

Chapter 6: Multi-objective Optimisation of Lightweight Metamaterials for Impact Resistance by Machine Learning

The results from both quasi-static implicit finite element analyses and explicit high velocity impact simulations are combined to train convolutional neural network models with the ability to predict mechanical performance of metamaterial lattices. The performance objectives are mapped to the placement of nodes and thicknesses of struts within the unit cell. The predictor networks are combined with an evolutionary genetic algorithm to synthesise architectures with maximum stiffness, strength and energy absorption. The approach is proposed as a way to harness the advantages of artificial intelligence to create ultra-

lightweight architected materials for enhanced vehicle performance, reduced carbon emissions and improved crash safety.

Chapter 7: Layup Optimisation of Carbon Fibre Reinforced Polymer Composites by Machine Learning

A convolutional neural network is trained on the stress values in elements within a finite element analysis model of a carbon fibre composite plate under in plane loading to predict the angles of the layup stacking sequence required to produce desired stress responses. An inverse design method is used where instead of altering the stacking sequence and then evaluating the stress, target stress contours can be dictated at the outset and the model provides the layup necessary to satisfy the request, providing additional speed, performance and flexibility over conventional approaches. Examples of highly stressed composite plates are presented. The model is then used to reduce the stresses by up to 80% under the same loading whilst maintaining a given load path through manipulation of the existing stacking sequence, without adding any additional plies or extra mass. This is accomplished by inputting the desired reduced stress response contours to the deep learning model which then outputs the stacking angles, allowing for tailoring of the structural properties of the composite.

A predictor model is then created that works in the opposite direction, inputting the stacking sequence to the neural network and outputting the stresses in every element. The model is used to display a proof of concept that trained machine learning models may in future be used as a faster alternative to finite element analysis throughout all engineering design practices. The predictor neural network is combined with a stochastic evolutionary genetic algorithm in addition to particle swarm optimisation and simulated annealing to output the layup that minimises the maximum stress anywhere in the composite. A method is proposed in which the predictor is trained on results from a fully parametric FEA model. The parameters that are traditionally used to define an FEA model including geometry, boundary conditions and loading are supplied as input to the AI model which then performs the stress analysis in a fraction of the time in addition to returning the optimal composite stacking sequence necessary to provide the lowest maximum stress.

Chapter 8: Conclusion and Future Works

The key conclusions of the research are presented and related back to the aim of the thesis. The advantages and current limitations of the proposed machine learning approaches are highlighted. Suggestions are provided for potential future works that may continue to build on the frameworks established during the study.

Chapter 2 Literature Review

2.1 Design Optimisation of Metamaterials for Mechanical Performance

2.1.1 Bending Vs Stretch Dominated Structures

The deformation mechanism of a cellular solid structure can be either bending or stretching dominated. A stretch-dominated structure is expected to be approximately ten times as stiff and three times as strong as the bending-dominated equivalent at a relative density of $\bar{\rho} = 0.1$ [15]. A small representative volume element can be modelled as a set of pin-jointed struts. If the pin-joints are assumed able to freely rotate without friction, any load forces will be resisted purely by axial loading of the struts. The failure load then depends largely on the axial tensile or compressive strength of the struts. However if the rotational degree of freedom is removed and the pin joints are now fixed or welded in place, the load will now cause bending moments about the joints. The failure mechanism now relies on the rotational stiffness and strength at the nodes and the structure becomes bending-dominated. However on a macroscopic scale the material may exhibit periodic collapse mechanisms which do not produce a macroscopic strain, the bulk properties remain unaffected and the structure remains stretch-dominated. This behaviour is largely dictated by the degree of connectivity of the joints rather than the regularity of the microstructure [8].

Maxwell's criterion can be applied to determine whether a pin-jointed structure will be statically or kinematically determinate [16]. In a two-dimensional pin-jointed frame made up of b struts and j frictional joints, if $b - 2j + 3 = 0$ then the structure is statically determinate. In three dimensions this relation becomes $b - 3j + 6 = 0$. These conditions for rigidity were further generalised in three-dimensional space by Calladine to take into account states of self-equilibrated stress in the absence of external loading, and of mechanisms and is given by [17]:

$$b - 3j + 6 = s - m \tag{2.1}$$

where s is the number of states of self stress and m is the number of mechanisms. If $m > 0$ and the structure contains mechanisms it has no stiffness or strength, it simply collapses under load. Locking the joints in this case leads to bending-dominated behaviour. If $s = m = 0$ the frame is rigid and both statically and kinematically determinate. Its struts will become loaded axially in tension or compression under loading. Deshpande et al. determined that for a large framework the necessary, but not sufficient, condition for rigidity is Z = 4 in two dimensions and Z = 6 in three dimensions, where Z is the number of struts connected at each node [15].

Cellular lattice materials will resemble frameworks in the regime where the relative density is less than 0.2 [18]. Under this condition the relative density of a two-dimensional lattice scales directly with the thickness t and length l of the struts [8]:

$$\bar{\rho} = A\left(\frac{t}{l}\right) \tag{2.2}$$

The scaling factor A depends on the geometry. Hexagonal topology scales with A equal to $2/\sqrt{3}$. Triangular and Kagomé lattice structures scale with factors of $2\sqrt{3}$ and $\sqrt{3}$ respectively. In three dimensions the relationship depends on whether the lattice is open or closed cell. For open cell lattices relative density scales with $(t/l)^2$. Closed cell lattices with uniform cell faces of thickness t and side length l scale with t/l.

2.1.2 Manually Generated Lattice Unit Cell Configurations

The majority of pre-existing lattice unit cell architectures that are used in literature are generated manually from geometric primitive shapes. Examples of periodic lattice unit cell designs which are commonly used in lightweight composite sandwich structures for load bearing applications are displayed in Figure 3. As a sandwich structure is loaded in bending the outer and inner face skins will be loaded conversely in tension and compression, inducing shear loads in the core material. Closed cell honeycomb lattices have been composed from triangular, square and hexagonal unit cell arrangements and have been long considered to be the most structurally efficient way of resisting both bending and transverse shear loads.

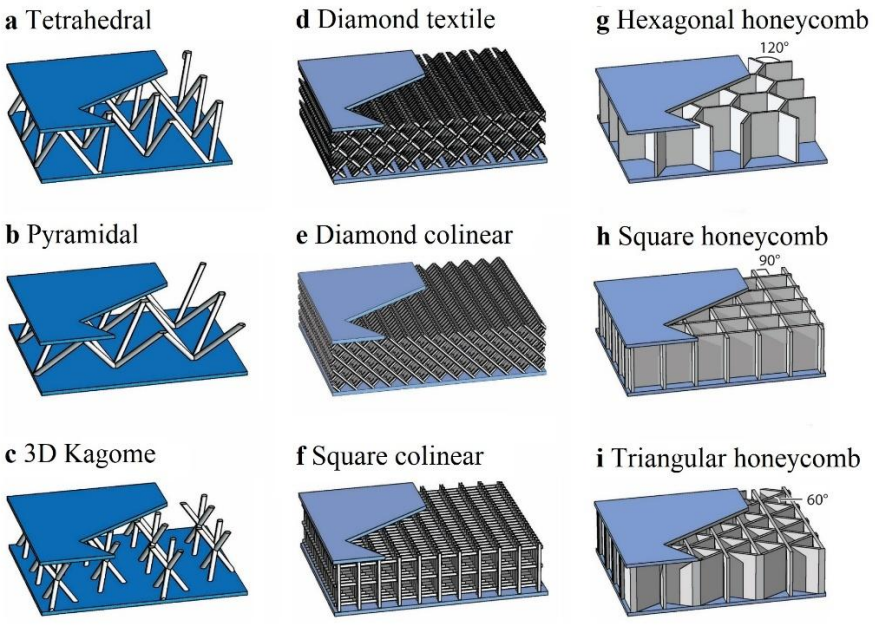


Figure 3 Examples of periodic lattice unit cell topologies commonly used as cores of composite sandwich panels. Adapted from [19]

Open cell tetrahedral truss cores were introduced in 2001 by Wicks and Hutchinson [20], who compared weights of optimised sandwich plates with truss cores against those with honeycombs, considering four competing failure modes of face sheet buckling, face sheet yielding, core web buckling and yielding under shear. A sequential quadratic programming algorithm was used to solve the multi-objective optimisation problem. Following the optimisation procedure the honeycomb sandwich required a significantly larger core thickness than the tetrahedral truss structures, eliminating it as a practical contender in the study. In order to make meaningful comparisons the core thickness, H_c was restricted and the outcome is displayed in Figure 4.

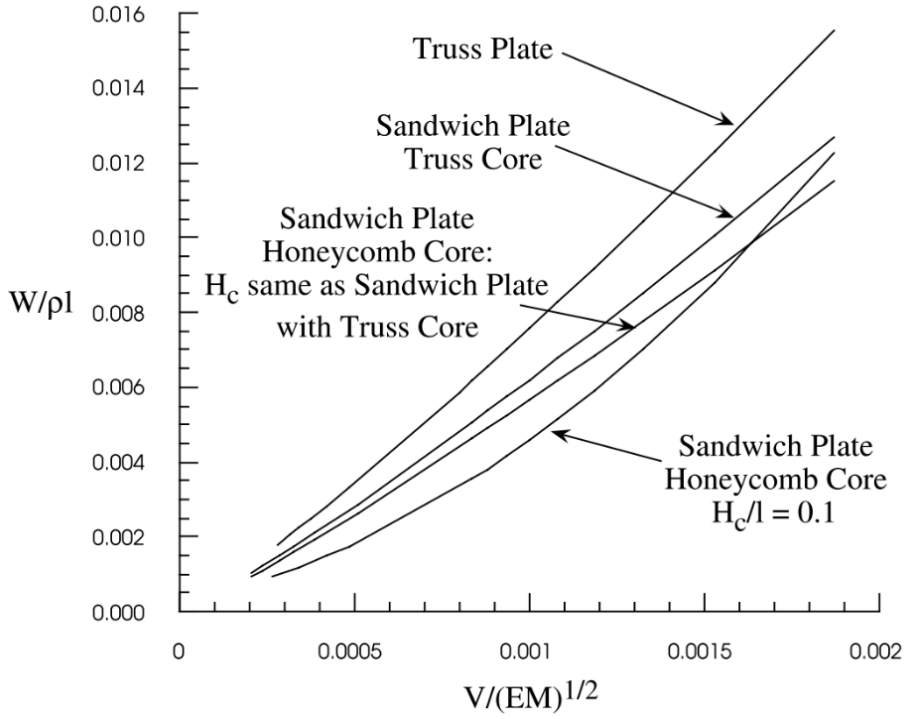


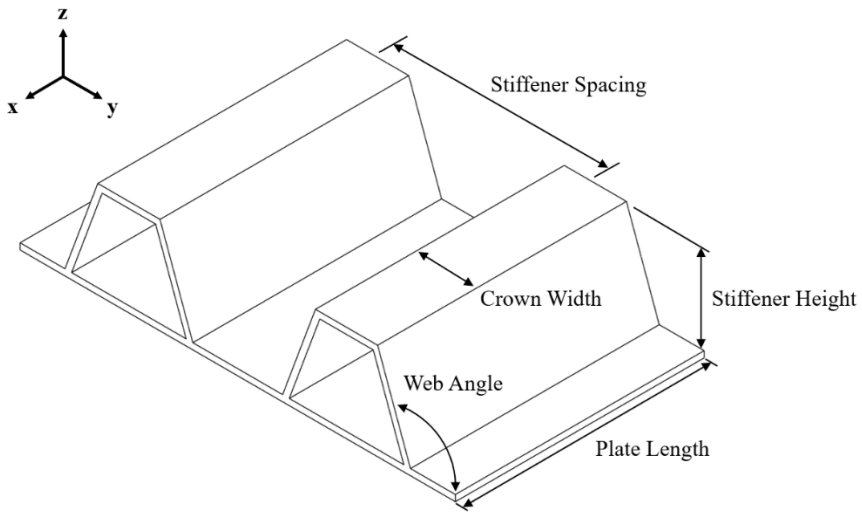
Figure 4 Weight per unit area against dimensionless loading parameter for comparison between optimised truss cores and honeycombs, all with $\sigma_y/E = 0.007$ and Poisson ratio = 1/3. Both truss core plates are fully optimised. Two cases of honeycombs are considered, one with core thicknesses, H_c fixed at $H_c/l = 0.1$ and one with core thickness taken equal to that of the sandwich plate with truss core at the same value of $V/(EM)^{1/2}$. From [20]

Results for tetrahedral truss core and honeycomb core were very similar, suggesting truss cores are equally effective for composite sandwich panels. Truss cores offered additional advantages in face to core bonding and the opportunities enabled by utilising open cell structures make them more suitable for multifunctional applications. These assertions were also supported by Wadley et al. [21] who demonstrated that truss cored sandwich panels are able to sustain loads at much lower relative densities than stochastic foams, while offering

superior peak strength to honeycombs at ultra-low density due to superior buckling resistance, at lower manufacturing cost. Wicks and Hutchinson warned that errors may be created from using the pin-jointed assumption, which can lead to an underestimate in buckling capacity, highlighting the need for accurate buckling analysis. Deshpande and Fleck found the pin-jointed assumption to provide adequate predictions for macroscopic effective stiffness and strength of tetrahedral cored sandwich beams [22]. Under experimental testing the deformation mechanisms were stretch-dominated in most cases. However failure under shear was mostly caused by plastic shearing of pins attaching the core to the outer face sheets, highlighting the need for stiff and strong joints that minimise nodal mass if lattice structures are to be effectively refined in the future. This requirement was also demonstrated by Chiras et al. who suggested a highly ductile parent material should be used, ideally greater than 20% under uniaxial tension, with low manufacturing induced porosities to achieve full load carrying potential [23]. This was found to be necessary to overcome high-hydrostatic stress concentrations forming at the nodes. Deshpande and Fleck measured deformation responses in 3-point bending and used the results to map the collapse mechanisms of truss core sandwich beams to the geometric parameters of the truss, allowing the selection of minimum weight designs for a given structural load under face sheet yielding, face wrinkling, core shear and indentation. The optimal designs were found to occur at the boundaries of competing collapse modes and outperformed metallic foam cores. During their experiments the strut aspect ratio was limited by investment casting manufacturing constraints to $\bar{\alpha} = 0.07$ which resulted in minimum weight designs 30% greater than those achieved by Wicks and Hutchinson's algorithm, which allowed aspect ratio to vary to find global minimum weight designs. They predicted further advances in manufacturing able to reduce strut aspect ratio would likely yield further substantial weight savings.

The performance of tetrahedral struts under compression loading was compared by Wicks and Hutchinson to hat-stiffened plates (Figure 5(a)), previously regarded as the most efficient design. The optimised tetrahedral core sandwich plates outperformed the hat-stiffened plates in terms of normalised weight per unit area under compression. The results are shown in Figure 5(b).

(a)



(b)

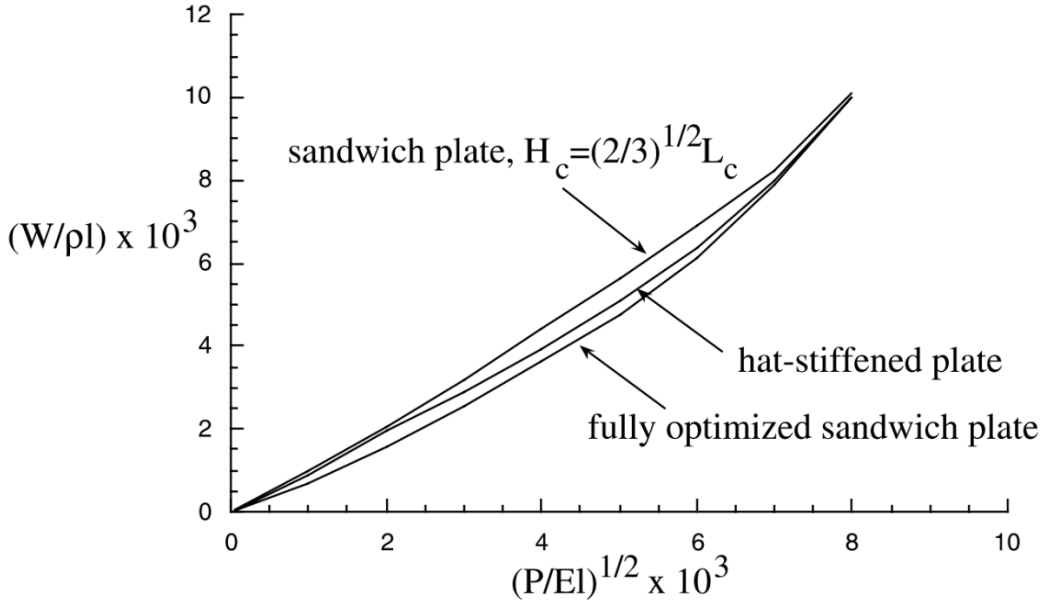


Figure 5 (a) Diagram of hat-stiffened plate; (b) Normalised weight per unit area as a function of the dimensionless load parameter for optimally designed simply-supported wide panels under compression ($\sigma_y/E = 0.007$ and Poisson ratio = 1/3). The upper and lower curves are for the sandwich plates with a truss core, constrained and unconstrained respectively. The middle curve is an optimised hat-stiffened single layer plate, which is regarded as one of the most efficient structures for this application. From [20]

A following study by Wicks and Hutchinson in 2003 compared tetrahedral truss core sandwich plates to honeycombs under compression and concluded that tetrahedral cores were as efficient as honeycombs under realistic minimum crushing strength conditions [24]. On relaxing this constraint honeycombs showed a slight weight advantage, but exceptionally low

crushing strength. The tetrahedral core on the other hand exhibited an inherent advantage in compression at low relative densities at which typical composite sandwich panel core applications occur. A close competition was highlighted between open cell lattice and closed cell honeycomb approaches for lightweighting. It was suggested that the major advantages of open cell hinge on the ease of manufacture, reduced vulnerability to delamination, low moisture ingress, and enablement of multifunctional capabilities. Kooistra, Deshpande and Wadley compared peak compressive strengths of aluminium tetrahedral lattice cores with competing open celled aluminium foams and closed cell honeycomb, and also found that tetrahedral cores are comparable to, and at low relative densities, superior to equivalent honeycombs [25]. The stress-strain responses exhibited a region of linear elastic loading followed by gradual core yielding up to the peak compressive strength. Further loading resulted in a period of strain softening followed by a stress plateau until the core densified, leading to greatly increased load resistance. In some instances only small amounts of strain softening was achieved after the peak strength, which is highly desirable for absorbing energy during impacts. The degree of softening was dependent on the relative density and could be altered by changing the metallurgical state of the parent material, allowing for tailoring of energy absorption properties for impact resistant applications by tuning of annealing and age-hardening treatments.

The pyramidal unit cell, featuring four rather than three trusses meeting at each face sheet node was studied by Wallach and Gibson, who carried out linear elastic truss analysis in MATLAB before confirming results by non-linear finite element analysis (FEA) in ABAQUS [26]. The elastic moduli, uniaxial compressive strength and shear strength were analysed as a function of relative density and aspect ratio of the unit cell. Good agreement was found between experimental data and theoretical predictions of Young's modulus and shear modulus, though in some cases the failure stress was underestimated in FEA as simple truss elements were used that only experience tension and compression and do not take into account any buckling that may have occurred. Comparisons with commercially available closed cell metallic foam at similar relative density showed the truss material to have superior properties in all cases. Shear moduli and uniaxial compressive strengths were particularly significantly higher. Zok et al. proposed a protocol for characterising the structural performance of pyramidal truss cores that combined experimental measurements, mechanism maps, FEA and analytical optimisation [27]. Failure mechanism domains were mapped using beam theory to allow selection of face sheet and core thicknesses to achieve optimal weight design, which occurred at the confluence of face sheet yielding, face sheet buckling and core member yielding domains. Responses of the core material were assessed under transverse

compression, in-plane shear and stretching and results were used to calibrate an orthotropic constitutive law which can be used for design calculations.

Chiras et al. observed substantial differences in responses of tetrahedral trusses loaded in tension compared to those in compression [23]. Trusses in tension continued to strain harden beyond yield, whereas those in compression experienced buckling after small plastic strains. This results in an asymmetric response across truss-cored composite sandwich panels when loaded in 3-point bending. In addition, early truss unit cell designs suffered from significant anisotropy, particularly the pyramidal configuration. Empirical testing and numerical simulations revealed the Young's modulus, yield strength and shear modulus vary significantly depending on the loading direction. Hyun et al. proposed the three-dimensional Kagomé unit cell design, inspired by a Japanese basket weave pattern to overcome these two key limitations of previous truss designs [28]. Kagomé structures were known to be highly structurally efficient [29], with elastic moduli approaching the optimal Hashin-Shtrikman upper bounds over a wide range of relative densities and offering excellent buckling resistance. The Kagomé has nodal connectivity of Z = 4 and has no strain-producing collapse mechanisms, it can only collapse by periodic mechanisms which do not produce a macroscopic strain. It is rigid in all directions. Kagomé architectures made from both copper and aluminium were compared against tetragonal lattices in non-linear FEA simulations of compression loading. Figure 6 compares the relationships found between non-dimensional compressive force and strain for the lattice core structures. Non-dimensional force was calculated as $F_S/\sigma_Y(\pi R_C^2)$, where F_S is the compressive load, σ_Y the yield strength of the base material and R_C the truss radius, while the vertical strain is taken as the displacement divided by the core thickness (height of the un-displaced lattice). For both designs the maximum compressive force was reached and followed by a period of softening. The Kagomé was able to reach significantly higher loads over greater displacements before softening. The peak loads coincided with the onset of plastic buckling and can be related directly to the buckling strength of the struts. Shear tests in three different orientations revealed isotropic material behaviour of the Kagomé core, exhibiting similar responses in all directions and maintaining isotropy beyond yielding. It was concluded that the Kagomé unit cell has greater load carrying capacity and is superior to tetrahedral trusses for use as ultra-lightweight core material in composite sandwich panels. Findings were confirmed and validated experimentally by Wang et al. [30].

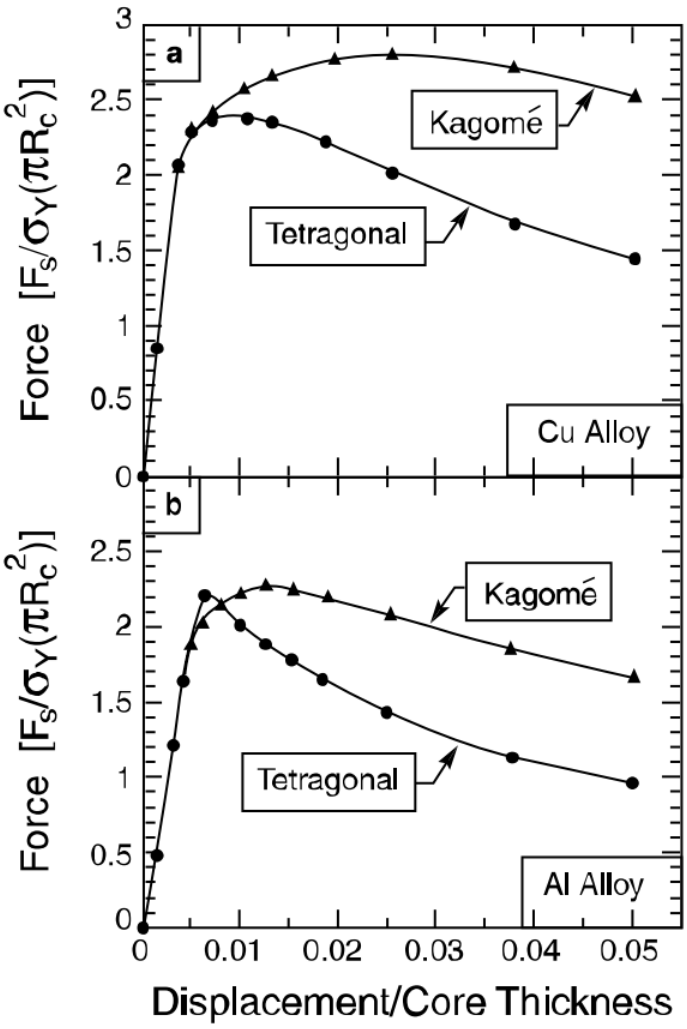


Figure 6 Non dimensional force/strain curves simulated in compression for tetragonal and Kagomé cores made from (a) Cu alloy and (b) Al alloy, where R_c is the truss radius [28].

Another alternative to tetrahedral and pyramidal unit cell designs is the octet-truss, proposed by Deshpande et al. [31], which features a symmetrical face-centred cubic (FCC) structure with nodal connectivity of Z=12 that provides near isotropic stretch-dominated material properties with a greater number of structural redundancies compared to the Kagomé. The unit cell design is displayed in Figure 7 and features a central octahedral cell combined with eight regular tetrahedrons at each face of the octahedron.

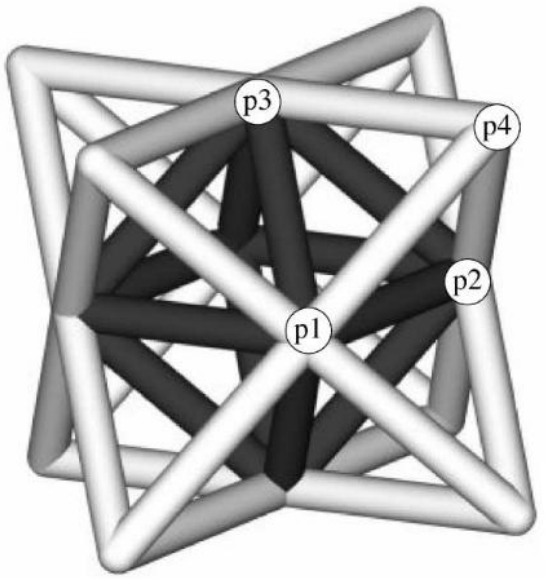


Figure 7 Structure of the octet-truss lattice unit cell. Darkened struts represent the octahedral core while the lighter struts connected at nodes p1-4 form one of eight tetrahedral cells [31].

The relative density of the octet-truss is given by:

$$\bar{\rho} = 6\sqrt{2\pi}\left(\frac{a}{l}\right)^2 \tag{2.3}$$

Where a and l are the radius and length of the struts. The relative density scales linearly with relative stiffness and strength, approaching the Hashin-Shtrikman upper bounds for stiffness and strength of isotropic voided materials. Hashin and Shtrikman defined the theoretical upper bounds of two-phase composite structures to achieve optimal effective elastic modulus for a given volume fraction [32,33]. In two dimensions the upper bounds on effective bulk modulus k_e and shear modulus G_e that are achievable by an isotropic cellular solid are given by

$$\frac{k_e}{k} \leq \frac{G\phi}{k(1-\phi)+G} \tag{2.4}$$

$$\frac{G_e}{G} \leq \frac{k\phi}{(k+2G)(1-\phi)+k} \tag{2.5}$$

where the material consists of solid phases of volume fraction ϕ, bulk modulus k and shear modulus G. In the low density regime where $\phi \longrightarrow 0$ the optimal effective modulii become linear functions of ϕ. Deshpande et al. manufactured octet-truss lattice material from LM25 aluminium alloy using a lost wax investment casting process for comparison against conventional metallic foams. Under uniaxial compression testing the stiffness and strength of the octet-truss exceeded metallic foams by a factor between 3 and 10 depending on the relative density, as shown in Figure 8.

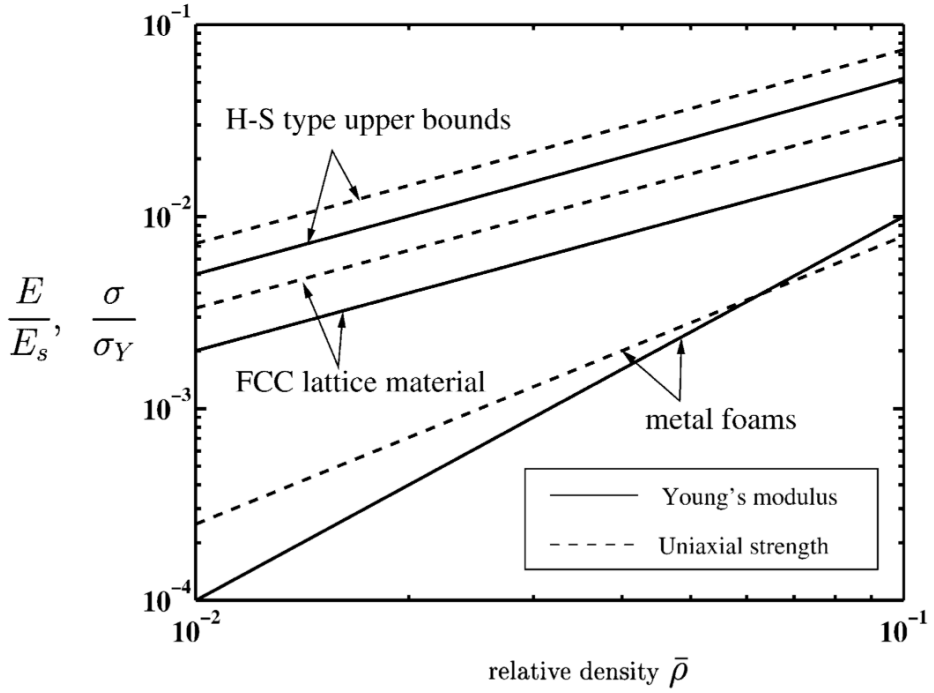


Figure 8 Comparison between stiffness and strength of the FCC octet-truss lattice, metallic foams and the Hashin-Shtrikman upper bounds [31].

The Euler elastic buckling strength of the octet-truss is given by:

$$P_E = \frac{n^2\pi^3 E_s a^4}{4l^2} \tag{2.6}$$

Where the factor n depends upon the rotational stiffness of the end nodes of the strut. If the struts are assumed pin jointed, the rotational stiffness is zero and $n = 1$. The struts will collapse by Euler buckling if P_E is less than the plastic yielding load $P_Y = \pi a^2 \sigma_Y$.

The macroscopic stress versus strain relationship in three dimensions was derived by Deshpande et al. The analytical solution can be found by applying the following:

$$\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_{23} \\ \varepsilon_{13} \\ \varepsilon_{12} \end{pmatrix} = \frac{1}{E_s\bar{\rho}} \begin{pmatrix} 6 & -2 & -1 & 0 & -2\sqrt{2} & 0 \\ -2 & 6 & -1 & 0 & 2\sqrt{2} & 0 \\ -1 & -1 & 5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 20 & 0 & 2\sqrt{2} \\ -2\sqrt{2} & 2\sqrt{2} & 0 & 0 & 20 & 0 \\ 0 & 0 & 0 & 2\sqrt{2} & 0 & 16 \end{pmatrix} \begin{pmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{pmatrix} \tag{2.7}$$

They simulated the material behaviour in FEA using Timoshenko beam elements under stress states of uniaxial tension and shear. Comparison between the analytical solution, FEA and experimental results under uniaxial tension demonstrated the effectiveness of the analytical and numerical approaches (Figure 9). Beyond yielding the lattice exhibited a hardening

response, followed by a softening corresponding to plastic buckling of the struts. The analytical and numerical responses showed overall good agreement with the experimental results, however in both cases the stiffness and strength are overestimated, suggesting conservative factors of safety and thorough experimental validation is advisable for practical design applications. Current theory only allows for periodic design and all struts are assumed to have the same values of sizing variables. Limited theory exists for location specific tailoring of struts and non-uniform, non-periodic designs. More recently Zhai and Zhang manufactured octet-truss lattice structures using spatial weaving of continuous glass and Kevlar fibre braids [34]. They found a strong link existed between strength and stiffness and the geometrical parameters of the unit cell design. The compressive modulus and strength of both glass fibre and Kevlar composite specimens rapidly increased with larger strut inclination angles and average strut diameter, but sharply decreased with strut length. The main failure mode observed was catastrophic failure of inclined struts that bear compressive load, while in-plane struts remained largely intact due to the fibres being loaded mainly in tension. This suggests that if it was possible to further optimise the structure in a non-periodic, non-recurring fashion it may be possible to ensure a statistically larger number of struts remain loaded optimally over a greater amount of time for a given loading. Instead of a series of straight beams that may be susceptible to failure if the loading becomes compressive in an inclined position, each individual strut could be tailored according to the specific localised loading conditions they experience, ensuring they remain largely in tension.

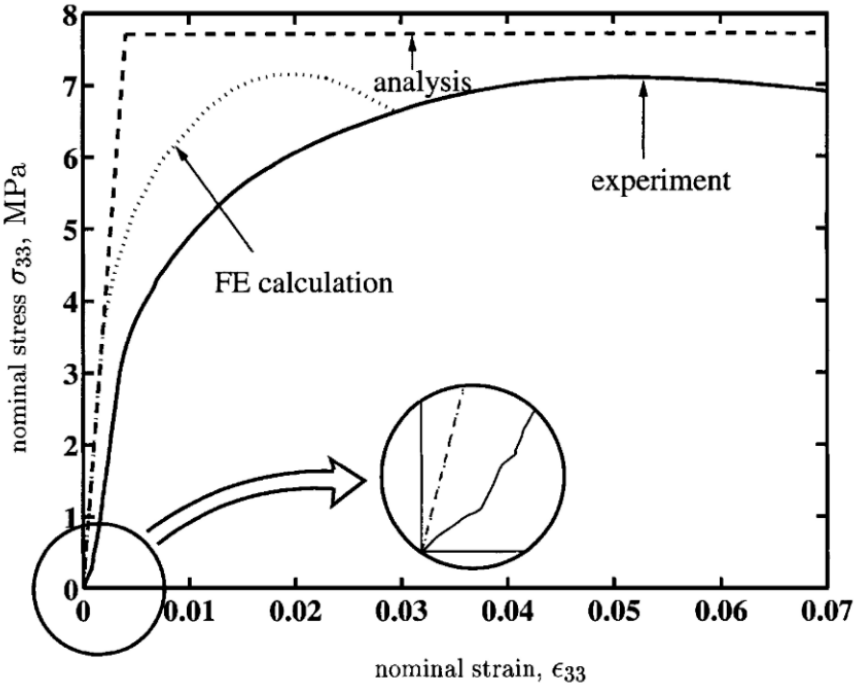


Figure 9 Comparison between analytical, numerical and experimental stress-strain responses for octet-truss lattice material cast from LM25 aluminium alloy [31].

Mechanical Performance

The nodal connectivity of a lattice unit cell is directly related to the scaling relationship between relative density and stiffness and strength [8]. Beam theory dictates the relations can be represented by the following power laws, taking the subscript S to refer to a property of the parent solid material:

$$\frac{E}{E_S} = B\bar{\rho}^b \quad and \quad \frac{\sigma_Y}{\sigma_{YS}} = C\bar{\rho}^c \tag{2.8}$$

Coefficients for the scaling laws can be found in Table 1 for standard unit cell topologies. Hexagonal honeycombs, commonly used as core material in composite sandwich panels, are susceptible to wall buckling due to bending dominated behaviour, with third order scaling of stiffness with $b = 3$, whereas both triangulated and Kagomé feature linear scaling with $b = 1$ and present stretch-dominated properties. Linear scaling is also observed in the strength of triangulated and Kagomé truss structures, whereas hexagonal honeycomb scales with a factor of 2. Higher nodal connectivity can generate additional structural redundancies and resistance to the effects of imperfections in the lattice, at the cost of reduced structural efficiency [35]. All pre-existing periodic lattices have fixed nodal connectivity. By applying additional tailoring at nodal locations using machine learning it may be possible to increase the level of structural redundancies at key regions of peak stress concentrations within a component, combining both material level and component level optimisations.

Table 1 Coefficients for relative density scaling laws of common lattice topologies [8].

Topology	A	B	b	C	c	D	d
Hexagonal	$2/\sqrt{3}$	$3/2$	3	$1/3$	2	0.90	2
Triangular	$2/\sqrt{3}$	$1/3$	1	$1/3$	1	0.61	1
Kagomé	$\sqrt{3}$	$1/3$	1	$1/2$	1	0.21	$1/2$

Theory on fracture resistance of truss lattice structures is currently less extensive than equivalent analytical solutions for predicting stiffness and strength and typically relies on the assumption that individual struts behave as elastic-brittle beams. The fracture toughness of a lattice may be defined as the value of the applied stress intensity factor K which causes the maximum local tensile stress at any point in the lattice to attain the value of the fracture strength of the constituent solid material from which the lattice is made [8]. Fleck and Qiu generalised the scaling law for fracture toughness K_{IC} by the following relationship [36]:

$$\frac{K_{IC}}{\sigma_{fS}\sqrt{l}} = D\bar{\rho}^d \tag{2.9}$$

Where σ_{fS} is the fracture strength of the solid and l is the cell size. The coefficients D and d can be found in **Error! Reference source not found.** for common lattices. For the Kagomé l attice $d = 0.5$, making its fracture toughness remarkably insensitive to changes in relative density. The Kagomé structure is able to maintain high levels of fracture resistance at ultra-low relative density due to localised buckling causing transformation of the lattice pattern in bands that radiate out from the crack tip. The result is a highly pronounced crack blunting effect that offers superior fracture toughness to other periodic unit cells.

An important consideration for the toughness of lattice metamaterials is defect tolerance. The damage tolerance of a lattice material and the parent solid are shown in Figure 10.

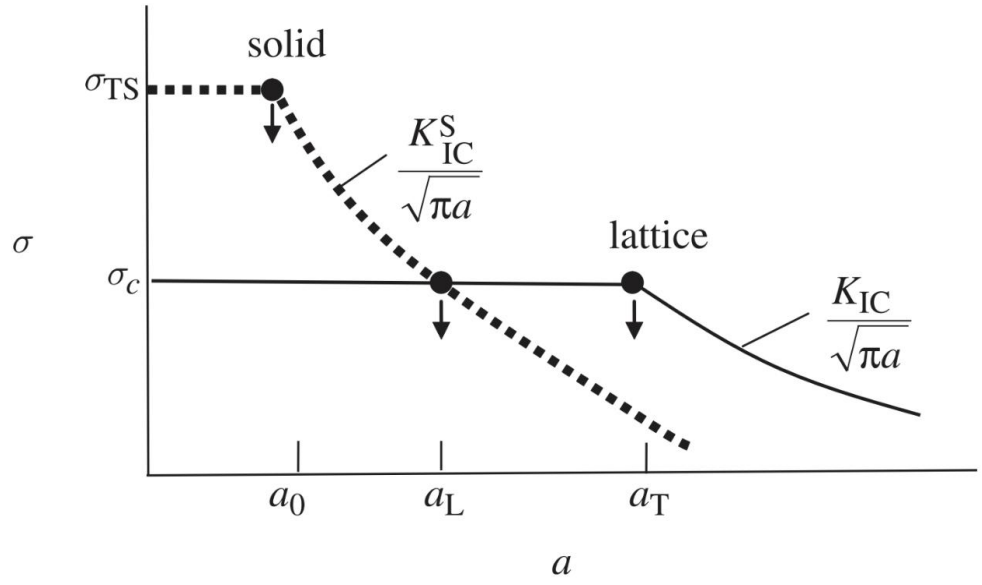


Figure 10 The damage tolerance of a lattice material and the parent solid [8].

When a lattice contains a sufficiently long internal defect of length $2a$ its tensile strength no longer depends on the fracture strength of the lattice material σ_c but instead can be related to the fracture toughness by:

$$\sigma = \frac{K_{IC}}{\sqrt{\pi a}} \tag{2.10}$$

The transition flaw size at which this occurs is given by:

$$a_T = \frac{1}{\pi}\left(\frac{K_{IC}}{\sigma_c}\right)^2 = \frac{1}{\pi}\left(\frac{D}{C}\right)^2 \bar{\rho}^{2(d-c)} l \tag{2.11}$$

The tensile strength of the lattice will exceed that of the parent material provided the flaw size is greater than:

$$a_L = \frac{1}{\pi}\left(\frac{{K_{IC}}^S}{\sigma_c}\right)^2 = C^{-2}\bar{\rho}^{-2c} a_0 \tag{2.12}$$

For an octet-truss a single missing strut can exceed the critical flaw size a_T. For the Kagomé the transition size varies inversely with relative density [36] and could be as large as multiple unit cells. However Kagomé structures are more sensitive to imperfections that result in perturbations of nodes from their ideal positions. The octet-truss has greater resistance to such defects due to higher structural redundancy from greater nodal connectivity of Z = 12 rather than Z = 4, though sacrifices structural efficiency as a result. A research gap exists for providing lattice metamaterials with ultra-high specific toughness while maintaining high stiffness and strength at ultra-low density. For this requirement to be filled in the future new novel lattice topologies are required. Possible approaches include numerical and machine learning based unit cell design optimisation, heterogeneous lattice topology and multi-scale hierarchical design.

2.1.3 Numerically Generated Lattice Structures

Instead of creating lattice architectures manually from configurations of geometric primitives, mathematical modelling and numerical optimisation approaches have been developed in attempts to improve on standard truss-based topologies. The resulting structures typically feature increased geometric complexity and can often only be created using additive manufacture (AM). Lattice structures can be generated using two different approaches, uniform or conformal (Figure 11). The lattice can either be a uniform distribution of periodic unit cells, each cell with the same size, shape and topology, or individual cells can be adapted to conform to the shape of the macroscopic design space boundaries [37]. Conformal lattices are more complex and time consuming to model, as they are usually created bespoke for a specific application, however they show significant advantages in strength and stiffness [38]. They are able to accommodate elaborate geometries that allow for greater design freedom, enabling highly optimised structural topologies and biomimetic shapes.

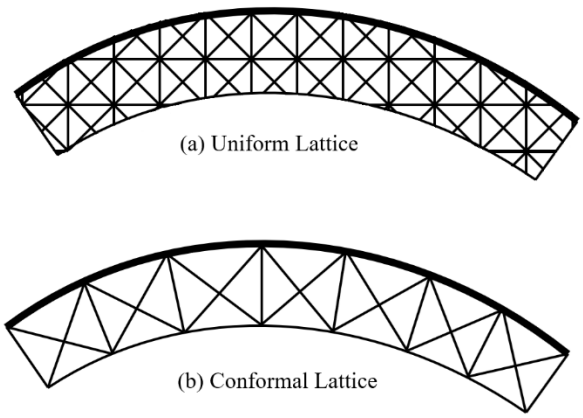


Figure 11 Uniform and conformal lattices.

Triply Periodic Minimal Surfaces

Hyperbolic geometries commonly occur in natural structures such as sea urchin splines, lipid bilayers and detergent films. A subset of hyperbolic surfaces, the minimal surface has attracted the attention of material scientists for use in architected lattice metamaterials. A minimal surface is a surface for which the mean curvature $H = (k_1 + k_2)/2$ is zero at every point, where k_1 and k_2 are the principal curvatures in two mutually orthogonal planes [39]. A surface can be defined as periodic if the following applies [40]:

$$\phi(\boldsymbol{r}) = \sum_{k=1}^{K} A_k \cos\left[\frac{2\pi(\boldsymbol{h}_k \cdot \boldsymbol{r})}{\lambda_k} + p_k\right] = C \tag{2.13}$$

Where $\boldsymbol{r}$ is the location vector in Euclidean space, $\boldsymbol{h}_k$ is the kth lattice vector in reciprocal space, A_k is the magnitude factor, λ_k is the wavelength of periods, p_k is the phase shift, and C is a constant.

A triply periodic minimal surface (TPMS) can be defined as a minimal surface that is infinite and periodic in all three directions of three-dimensional space. The coordinates of a TPMS can be calculated using the Weierstrass formula [40]:

$$\begin{cases} x = Re\int_{\omega_0}^{\omega_1} e^{i\theta}(1 - \omega^2)R(\omega)d\omega \\ y = Im\int_{\omega_0}^{\omega_1} e^{i\theta}(1 + \omega^2)R(\omega)d\omega \\ z = -Re\int_{\omega_0}^{\omega_1} e^{i\theta}(2\omega)R(\omega)d\omega \end{cases} \tag{2.14}$$

Where ω is a complex variable, θ is the Bonnet angle, and $R(\omega)$ is the Weierstrass function which varies for different surfaces. For example for TPMS primitive (P), diamond (D) and gyroid (G) surfaces, $R(\omega) = 1/\sqrt{1 - 14\omega^4 + \omega^8}$. Examples of common TPMS surfaces are shown in Figure 12.

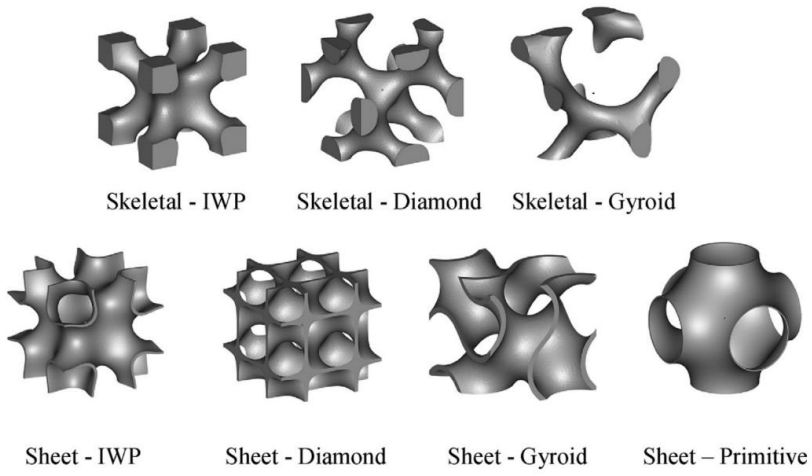


Figure 12 Triply periodic minimal surface unit cell configurations [41].

Studies of TPMS surfaces have shown that some configurations exhibit mechanical properties that can outperform strut-based lattices. Al-Ketan et al. compared truss unit cell based lattices with mathematically generated TPMS structures manufactured from maraging steel via powder bed fusion [41]. Both sheet and skeletal TPMS lattices were studied. Sheet-TPMS's are created by thickening the minimal surface, whereas skeletal-TPMS's are formed by solidifying the volume enclosed by the surface. Skeletal lattices exhibited bending-dominated behaviour and quadratic scaling of stiffness with relative density, while sheet based unit cells were almost purely stretch-dominated and showed superior mechanical performance to Kelvin, Gibson-Ashby and octet-truss topologies. In particular diamond (D) sheet TPMS configurations scaled with a power law exponent of n = 0.522, which suggests it is able to maintain high stiffness at significantly lower relative densities than octet-truss, which scaled at n = 1.23. Gyroid (G), primitive (P) and isotropic woodpile (IWP) configurations scaled with values between n = 1.15 and n = 1.31. The comparison of the scaling relationships found for Young's modulus are shown in Figure 13. This low sensitivity to reducing relative density indicates significant mechanical performance improvements may be achievable at ultra-low mass if future manufacturing advancements are able to enable smaller minimum feature sizes. This finding was supported by Abueidda et al. who highlighted that at low relative densities there is greater sensitivity to printing direction and defects are more likely to occur due to limitations of the resolution of the AM process [42]. They suggested fabricating at the micro or nano scale using methods such as atomic layer deposition, photopolymer waveguide and metal etching may open possibilities for highly lightweight and efficient metamaterials.

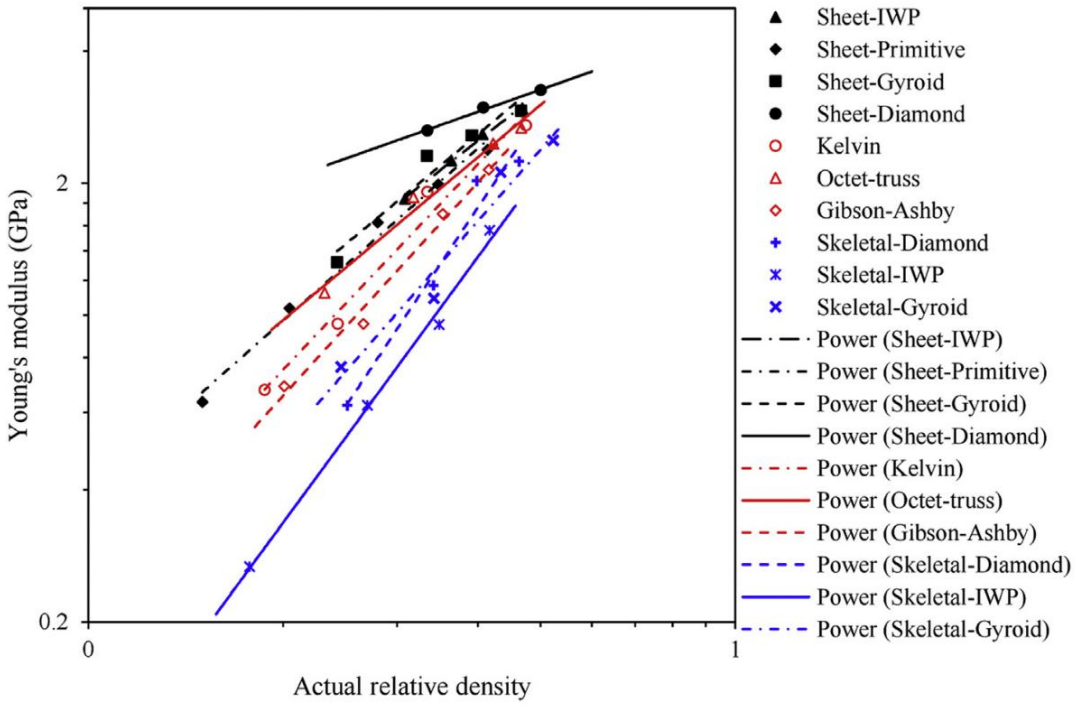


Figure 13 Scaling relationships between Young's modulus and relative density [41].

TPMS structures can exhibit greater toughness and a more efficient stiffness to weight ratio than truss lattices, showing particularly good suitability for energy absorbing applications. Under compressive loading many strut-based structures will exhibit an abrupt significant drop in stress levels immediately after the ultimate strength of the material is reached due to sudden failure of cell layers, often accompanied by the formation of shear bands. This can often be followed by an unstable series of periodic fluctuations in stress, which is undesirable for efficient energy absorption. Energy absorption per unit volume can be calculated by integrating the area under the stress-strain curve. Therefore the most efficient absorption occurs with a smooth gradual decline of stress after the initial peak, which transitions into a flat, progressively increasing stress plateau before an eventual inevitable rise due to densification of the material [43]. For truss-structures the lattice exhibits high Young's modulus when the loading is aligned with the direction of the struts, but low values in other directions. Hence the response depends on keeping the loading aligned favourably over as much of the structure as possible for the greatest amount of time, which is not always the case as the structure begins to buckle and shear. TPMS's can offer superior mechanical performance as the geometry continuously changes alignment across the unit cell. The cell walls are much thinner than the overall cell dimensions and tensile loading will also induce local compressive loading due to the complex geometries. At small wall thickness and low relative density strength is governed by buckling, while at higher thickness and relative density it is governed by plastic yielding. The transition point was found by Abueidda et al. to occur at approximately 2.5%, 1% and 2% relative density for primitive, IWP and neovius (N) configurations respectively [42]. The topology of TPMS's also transitions smoothly into adjacent cells, solving issues with nodal stress concentrations that occur at the joints between cells in strut lattices and allowing for a more efficient and progressive load transfer through the material.

TPMS surfaces are inherently free of self-intersections. They can be assembled into periodic lattice structures without producing meshing artifacts, requiring limited post processing which is highly advantageous for additive manufacture. This also presents opportunities for hybridisation of different unit cell configurations in the same lattice. The Schoen gyroid TPMS is particularly manufacturable thanks to an inherently self-supporting structure [44]. Generating unit cell topology mathematically allows for ease of control over spatial distributions of material density, enabling functionally graded lattice materials. Relative density can be easily varied by altering the thicknesses of TPMS sheets. Predicting the mechanical performance of graded density TPMS structures is currently challenging due to lack of theory. Xiaojie Fan et al. studied the manufacturability, mechanical properties and

energy absorption capacities of graded thickness gyroid, neovius and primitive TPMS lattices [45]. The heterogeneous graded thickness lattices offered reduced stress fluctuations, improving the stability of stress plateau regions while offering higher elastic moduli and cumulative energy absorption compared to constant thickness structures. At 30% volume fraction the Young's modulus of neovius and gyroid lattices were boosted by 94MPa and 136MPa respectively, while cumulative energy absorption increased by 24% and 15.4%.

Random Generation

Practically all existing theoretical mechanical property relations are for periodic topologies. Randomised numerically generated models typically depend on computational finite element analysis and experimental mechanical characterisation. On applying FEA to the study of randomised Voronoi tessellations and Gaussian random fields (GRF), Roberts and Garboczi found that large variances occur in the value of scaling exponent due to the stochastic nature of the material [46]. At low densities, the scaling relationships featured a more complex dependency on the material microstructure than for periodic unit cells, requiring consideration of larger scale, multi-cellular models.

Voronoi tessellations are a three-dimensional random open-cell arrangement with a mathematical definition similar to that of the physics of natural cellular foam formation [47]. Cell regions are defined through the random distribution of "seed points" in space. Placing solid walls along the faces of each cell will form a closed-cell Voronoi tessellation. If instead only the edges where the cells intersect are defined as solid, an open-cell Voronoi tessellation lattice is formed. The degree of order in the lattice can then be tuned by altering the sampling method for the seed points. More regularised tessellations can be generated using body-centred cubic (BCC) or face-centred cubic (FCC) arrangements, or purely randomised, stochastic formations can be synthesised by random distributions of seed coordinates. Delauncy or Voronoi-dual tessellations are formed by defining the cell edges a rod placed between two points in space which share a common face in the Voronoi model. Roberts and Garboczi used a Monte-Carlo algorithm to generate non-periodic, evenly spaced and isotropic arrangements of seed points provided by the centre points of equi-sized hard spheres packed in thermal equilibrium [46]. They proposed a "node-bond" model in which a bond is placed between each seed point and its nearest neighbours. The resulting structures are displayed in Figure 14. At low relative densities between 0.04 and 0.15, the node-bond model scaled with relative density with exponent n = 1.3, slightly higher than perfectly linear as the topology has no straight-through struts. However this did not lead to significant bending that might result in quadratic scaling as there was a significant average number of struts emanating from each

node ($\bar{Z} = 12$) to lock the nodal positions in space and provide the rigidity criterion necessary for largely stretch-dominated behaviour. At low nodal connectivity ($\bar{Z} = 5.5$) quadratic scaling began to occur. This highlights a major consideration that must be taken into account when designing stochastic cellular architectures. A key design trade-off exists between nodal connectivity and structural efficiency. It is desirable to remove redundant struts in regions that experience low stress concentrations for lightweighting and prevention of notch effects, however it is essential that high nodal connectivity must be maintained along key load paths. A research gap exists for achieving tailored stochastic lattices with tuned nodal connectivity that maximises structural efficiency without straying into quadratic scaling, bend-dominated behaviour in vital regions. Such behaviour was observed in Gaussian random field models, where excessive curvature of struts allowed for greater levels of bending deformation, leading to cubic scaling exponent of n = 3. Large variations in strut thickness of the GRF lattice resulted in reduced Young's modulus as the stiffness of struts with non-uniform cross sections is limited by the thinnest sections. The study focused on global connectivity and geometrical cell arrangements, rather than local characteristics such as cross section or curvature, emphasising the fact that these are difficult to model simultaneously due to limitations in computing power. A strong requirement exists for analytical models and experimental research of stochastically generated cellular materials at multiple length scales.

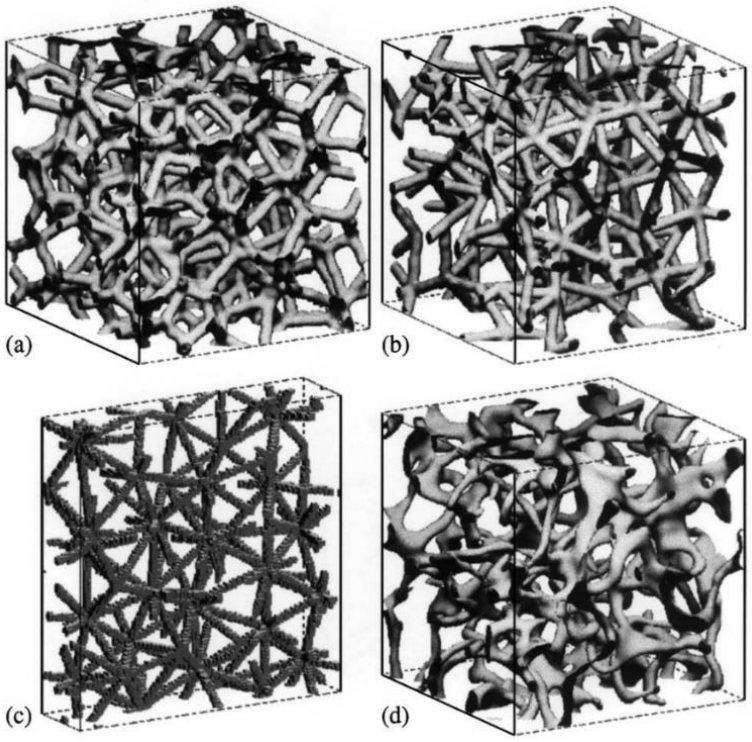


Figure 14 Random models of open cell solids. (a) Open-cell Voronoi tessellation with ρ/ρ_s = 0.15 (b) Node-bond model with $\bar{Z} = 5.5$ (c) Digital node-bond model with $\bar{Z} = 12$ (d) Open-cell GRF with $\rho/\rho_s \approx 0.1$. From [46]

Topology Optimisation

Topology optimisation is a mathematically driven analysis process for redistributing material within a design space in the most efficient way subject to a target objective function and a set of boundary constraints. Manually generated architectures, while having the advantage of implicit analytical solutions for predicting mechanical performance as a function of unit cell geometry, have the drawback of limiting the scope of any further optimisations for considering competing design trade-offs or multifunctional requirements for a specific application, as the topology has been locked by the initial choice of unit cell. Topology optimisation is able to change the connectivity of the structure as the design is iterated, eliminating dependence on standard unit cell architecture.

The framework for applying topology optimisation to architected material design was first envisioned by Ole Sigmund [48]. The designer choses the constituent base of the material, which may be a combination of legacy materials or of materials and void space. An objective function is then formulated, such as minimise mass while targeting maximum stiffness or strength. Constraints of the problem are defined, for example orthotropic or isotropic material behaviour, elastic symmetry or manufacturing restrictions such as minimum feature size. The density-based topology optimisation for achieving minimum compliance with a constrained volume can be formulated as follows [49]:

$$\begin{gathered} Find: \mathbf{x} \in R^n, \mathbf{x} = (x_1, x_2, \dots, x_n) \\ Minimise: c = f^T u \\ s.t.: K(x)u = f \\ \sum x_i v_i \leq V \\ 0 < \varepsilon < x_i \leq 1 \end{gathered} \tag{2.15}$$

Where **x** represents an n-D vector which contains the relative density of each element within the functional volume; c is the structural compliance calculated from the global external force vector f and a global displacement vector u; $K(x)$ is the global stiffness matrix; v_i is the volume for the ith element, and V is the upper limit of the total volume. The optimiser iterates towards convergence of a solution to the objective function that satisfies the design requirements. The design space is typically discretised into finite elements so that the strain energy and stress distributions can be mapped through the unit cell. Void space is then assigned to regions of lowest stress and material added at stress concentrations. The resulting

unit cell design can then be periodically repeated to achieve the desired macroscale bulk material properties. These bulk properties can be upscaled from unit cell level by the method of homogenisation, where local material properties are related to the design variables through consideration of representative volume elements [50]. Inverse homogenisation, introduced by Sigmund [51] involves minimising the difference between the homogenised material properties and the target material properties within the representative volume element.

Depending on the complexity of the problem and the number of interdependent design objectives, the amount of variables to be considered can make a topology optimisation study extremely computationally expensive. Parameterisation of the design space is able to simplify optimisation problems by reducing the analysis down to a limited number of key design variables, describing the possible lattice topologies as a compact vector of real numbers. A lattice unit cell can be described by four types of parameters [52]; A description of the simple lattice joint positions, a description of the position of the remaining joint sets, naturally described as the set of joint basis vectors, a description of the struts connecting the joints, and (optionally) a rotation that describes the orientation of the unit cell relative to a global coordinate system. Parameterisation can greatly facilitate the analysis of conformal lattices in which topology rotates to conform to the macroscale design shape. Additional design for additive manufacture (DfAM) constraints are usually necessary to create printable structures. Generated models must be “watertight”, free of meshing irregularities and surface artifacts such as cracks and self-intersecting polygons. In many cases the raw results require post modelling, smoothing and refinement to ensure the geometry can be efficiently meshed for further FEA validation and accurately converted to STL format for additive manufacture. Repairing STL meshes is often time consuming and resource intensive. To improve the reliability and manufacturability of topology optimisation results, interpolation methods such as solid isotropic material with penalisation (SIMP) and rational approximation of material properties (RAMP) have been developed to artificially reduce the complexity of homogenisation approaches and improve convergence rates [53]. In the level set method the structure is created by mapping a level set field onto the mechanical model, which is then iteratively updated by an optimiser. The topology of the lattice is defined by the domain in which the level set function takes positive values [54]. The process is well suited to the generation of stochastic geometry variations, however the geometry can only evolve from pre-existing boundaries and the effectiveness of the results depends greatly on making an initial starting guess. This is a common issue in topology optimisation. Large design spaces typically require gradient descent based optimisation algorithms to search for the best solution. Such algorithms use the method of moving asymptotes to converge at minima points [55].

Depending on the starting point the solver may end up converging on only a localised minima rather than the global optimal solution. An initial starting geometry is therefore usually required to obtain a guess close to the global solution which can then be further refined. This limits the amount of design freedom, a problem that may be overcome by the use of stochastic evolutionary algorithms.

Evolutionary Structural Optimisation Algorithms

Topology optimisation methods have been expanded with heuristic updating rules to simulate natural evolution. Evolutionary structural optimisation (ESO) was first proposed by Xie and Steven [56] to remove material in regions of a structure where it is not required, "evolving" towards an optimal structurally efficient solution. This requires a large initial volume of material that fills the entire design space from which mass is progressively removed. The process is very time consuming, usually involving many hundreds of FEA studies and the majority of the initial structure is redundant and gets replaced with empty space. The bidirectional evolutionary structural optimisation (BESO) algorithm was later suggested by Querin, Steven and Xie as an extension of ESO which in addition to removing material in regions of low stress, also adds material in high stress areas [57]. This allows the lattice to grow from a small initial structure with increased efficiency. However in some scenarios BESO can be limited by overly discrete density updates and lack of convergence due to difficulties defining appropriate stopping criteria to identify when the optimal solution has been reached. Mahdavi and Hanna synthesised lattice topology by evolving the microstructure using an evolutionary genetic algorithm (GA) combined with deterministic analysis of forces via the direct stiffness method (DSM) [58]. The volume of the design space is first divided into a grid of modular units. The connections between nodes are described by a graph of connections, which is then evolved by the GA under a range of tensile conditions. Chromosomes are constructed as a series of upper triangular matrices (UTM) that contain all of the nodal connections in binary format as a series of zeros and ones. If an entire row of the matrix becomes populated with zeros, that node is eliminated from the structure. An initial population of ones and zeroes is randomly assigned to the matrix and a crossover function is applied to swap values of a random selection of points. A mutation mechanism is then applied which selects a number of mutation points and flips the bits, replacing zeros with ones or vice versa, resulting in the breaking or forging of connections. The mutations are carried out over time to vary the population with random probability. A function is generated to assign a "fitness" value to each individual member of the population, considering a number of factors based on key design criteria, such as overall mass, maximum deflection, or DfAM considerations. Each factor is assigned a weight that allows tuning of the model. The

population is assessed according to an elitist approach that takes the top performing individuals of each generation and places them directly into the next generation. The nature of this approach ensures the average and maximum fitness levels can only ever increase with subsequent generations. The algorithm is run for thousands of generations until the fitness levels converge towards an optimal solution. Particle swarm optimisation (PSO) is an extension of genetic algorithms. It can be used to perform parametric topological optimisation of structures by simulating the movement of birds in a flock [59]. Individual "birds" are influenced by their own experiences and those of surrounding birds in their immediate vicinity, leading to a hive mind response to solving problems. This approach combines both local and global optimisations, allowing the solver to converge faster than a standard GA as localised birds cooperate with others to assist the overall effort of finding a global optimal solution. PSO also allows for additional secondary design objectives alongside mechanical performance targets during the topology synthesis, enabling opportunities for multifunctionality.

Currently existing methods for numerical generation of cellular architectures present challenges with the scalability of designs for use in large scale three-dimensional components. Results are often highly dependent on selection of initial starting geometry and require an initial guess to ensure good convergence. There are issues with the tunability of parameters during the design iteration process, many approaches lack the ability to handle multiple constraints and complex boundary conditions while also taking into account DfAM considerations such as minimum feature size, overhang angles, build direction and printing induced residual stress concentrations. Excessive amounts of sacrificial support structure is usually required to accommodate the complex geometries created, increasing build time and material wastage. More research is required to develop methods for applying topology optimisers across multiple length scales to take advantage of hierarchical architectures. Existing solvers require extensive prior FEA experience and knowledge of mathematical theory. An easy to use intuitive and fast optimisation framework, able to produce multi-scale lattice topology that can link material level with component level design would benefit large scale uptake of architected materials within industry. Use of established FEA design iteration methods within the optimisation process requires significant computing power and results in a bottleneck in design of optimised lattices. Consideration of alternative approaches such as machine learning and generative design is required to overcome this computational bottleneck.

2.2 Machine Learning

Machine learning is a subset of artificial intelligence that is able to learn features and tasks directly from sets of numerical, sequential or image data. A statistical model is trained by passing data through neural network architectures that are able to identify patterns that correlate between the input and the desired output. Once trained the model can be used to make predictions on new unseen data for tasks such as classification and regression. In a classification problem, the probabilities with which the input is likely to belong to a certain set of categories are established. Examples include image recognition, spam mail filtering and fake news detection. During regression, instead of outputting the probabilities to determine a class, the model outputs a set of continuous scalar values. This can be extremely useful in the design of metamaterials, as the outputs may be a set of geometry parameters of a unit cell that are mapped to the stress/strain response of desirable mechanical properties. Both classification and regression tasks belong to a subset of machine learning known as supervised learning [60]. Supervised learning is the most common branch of machine learning. Algorithms are designed to learn by example, similar to the way humans learn things. Current knowledge is applied to solve the problem. If the solution turns out to be incorrect the knowledge is modified and updated until the correct outputs can be obtained for future data. Unsupervised learning is typically used in exploratory data analysis. Unlike supervised learning it does not use data that can be labelled and instead relies on extracting data features to manifest underlying patterns. Reinforcement learning uses rewards and punishments as signals for positive and negative behaviour and learns from trial and error based on feedback from its own actions and experiences [61]. It is often used in robotics, business strategy planning and road traffic control. Machine learning is increasingly gaining popularity within society, as big data processing becomes more and more prevalent. Advancements in computing hardware and availability of streamlined open source software provide increasing opportunities for collaboration between humans and artificial intelligence in all areas of civilization. The potential advantages are particularly compelling in engineering design, where the ability of machine learning to solve multi-objective design problems with no pre-existing analytical theory make it an especially powerful tool.

2.2.1 Artificial Neural Networks

Artificial neural networks (ANN) form the basis of machine learning. They are in essence an attempt at bio-mimicry that aims to simulate the firing of neurons in the human brain. The fundamental structure of an ANN is shown in Figure 15. A nodal arrangement of neurons is connected in a series of layers. Early neural networks featured only input and output layers, known as single-layer neural networks. The downside of a single-layer network is it can only solve linearly separable problems [60]. Introducing a hidden layer in between the input and output results in a shallow multi-layer neural network. Input data signals are transmitted via the input layer to the sequence of hidden layers. The resulting outputs may contain real numbers between 0 and 1 corresponding to classification probabilities, or alternatively continuous regression values.

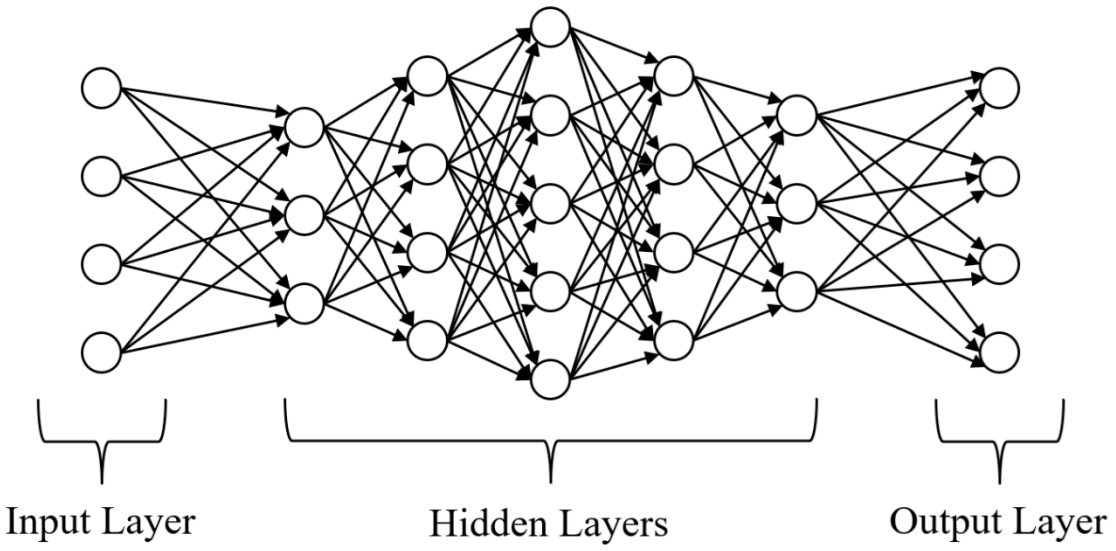


Figure 15 Artificial neural network architecture

2.2.2 Deep Learning

When a neural network has two or more hidden layers it becomes a deep neural network (DNN). The neurons are connected to neurons of the next layer through channels (Figure 16). Each channel is assigned a numerical value known as the weight w_i, an indication of how important that connection is. As the input data signals X_i are passed along the channels they are multiplied by the weights and the sum is then transferred to the neurons in the hidden layer. Each neuron in the hidden layer has an associated bias b_i which is added to the input sum. The weighted sum is then passed through a non-linear function known as the activation function σ, which decides on whether or not the neuron will fire and contribute towards the

next layer. The neuron with the highest value determines what value the output $\hat{y}$ takes. This propagation of information from input layer to output layer is known as forward propagation. The relationship between the input and output signals can be formulated as follows:

$$\hat{y} = \sigma(\sum_{i=1}^{n} X_i w_i + b_i) \tag{2.16}$$

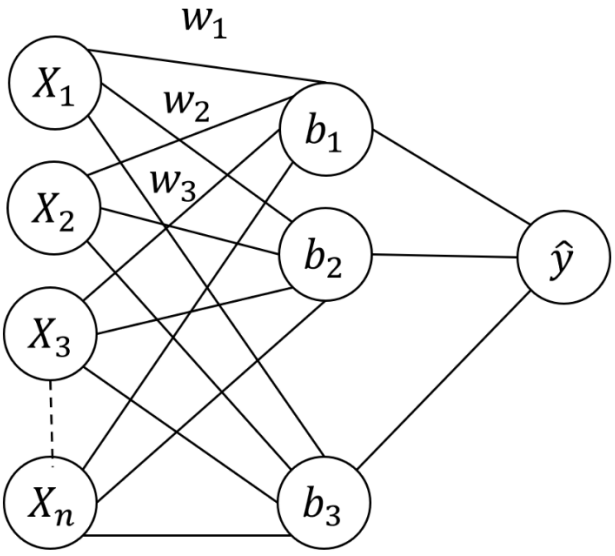


Figure 16 Forward propagation

The process of backwards propagation allows for the neural network to check its own performance by passing information from the output layer back to the input to see if the prediction was right or wrong. This is a powerful mechanism and is what allows the neural network to learn complex features of the dataset. The weights and biases are initially assigned randomly. A loss function quantifies the deviation between the predicted output and the expected output. An optimiser is then used to adjust the weights and biases to increase the accuracy of the predictions. As the dataset is iteratively fed through the network the training loss is continuously minimised and the network is able to make better and better predictions.

The most simple activation function is the step function, which activates only when its value is greater than zero. This provides a binary yes or no response, which can't model situations requiring probabilities or continuous scalar values. The linear function is non-binary and the activation is proportional to the input. However as the derivative is constant the gradient is independent of the input. During backpropagation adjustments wont depend on the input at all and the advantages of stacking layers is negated. Using a non-linear activation function allows the network to learn non-linear relationships. One of the most widely used is the biologically inspired sigmoid function (Figure 17(a)), given by $sig(x) = 1/(1 + e^{-x})$. Unlike linear activation, the output is always between 0 and 1. The function is smooth and differentiable everywhere. However as the input increases the signal begins to saturate and the gradient

becomes very small. Multiplication of the gradient of the activation function is required to calculate the loss function. If the gradient becomes very small the network will stop learning. This is known as the vanishing gradient problem [62]. The hyperbolic tangent function suffers from the same issue. It is a version of the sigmoid function rescaled between -1 and 1 instead of 0 and 1 (Figure 17(b)). The advantage of this is that the activation approximates to the identity function at the origin, increasing the convergence speed of gradient descent optimisation.

(a) (b)

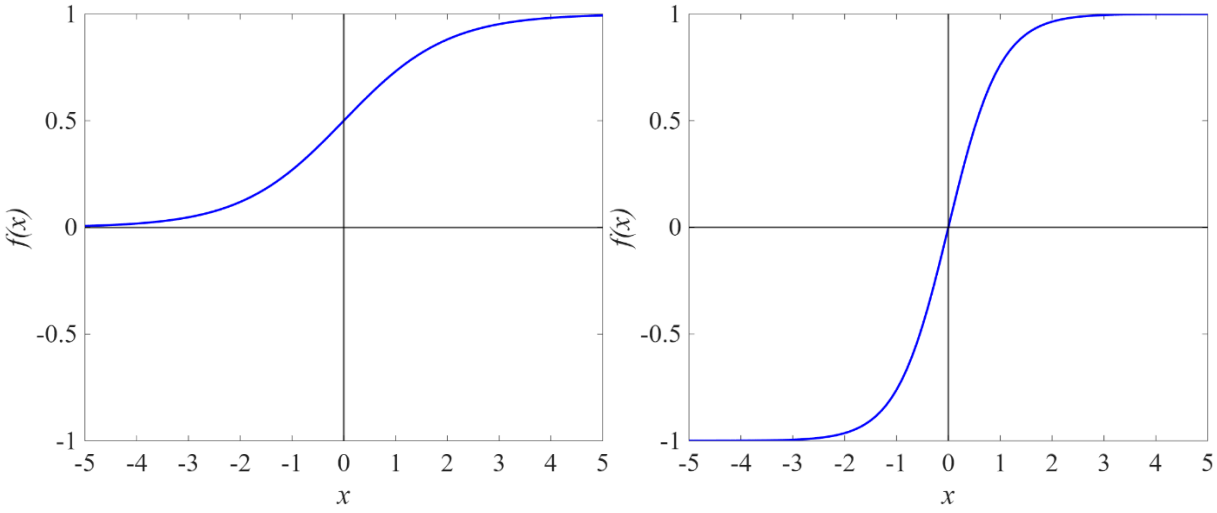


Figure 17 (a) Sigmoid activation function; (b) hyperbolic tangent activation function

The vanishing gradient problem limits how deep the neural network can be before the network stops learning. Sigmoid and hyperbolic tangent functions are limited to shallow networks with only a few layers. The rectified linear unit (ReLU) is a non-linear activation function that does not saturate as the input increases and does not suffer from vanishing gradient, making it suitable for deep networks (Figure 18(a)). It produces sparse activations making it computationally efficient, however its gradient is zero whenever the input is negative. This can result in dead neurons that always return zero, known as the dying ReLU problem. Leaky ReLU was introduced as a workaround to this problem, in which a slope, commonly set at a gradient of around 0.01 is used in the negative region (Figure 18(b)).

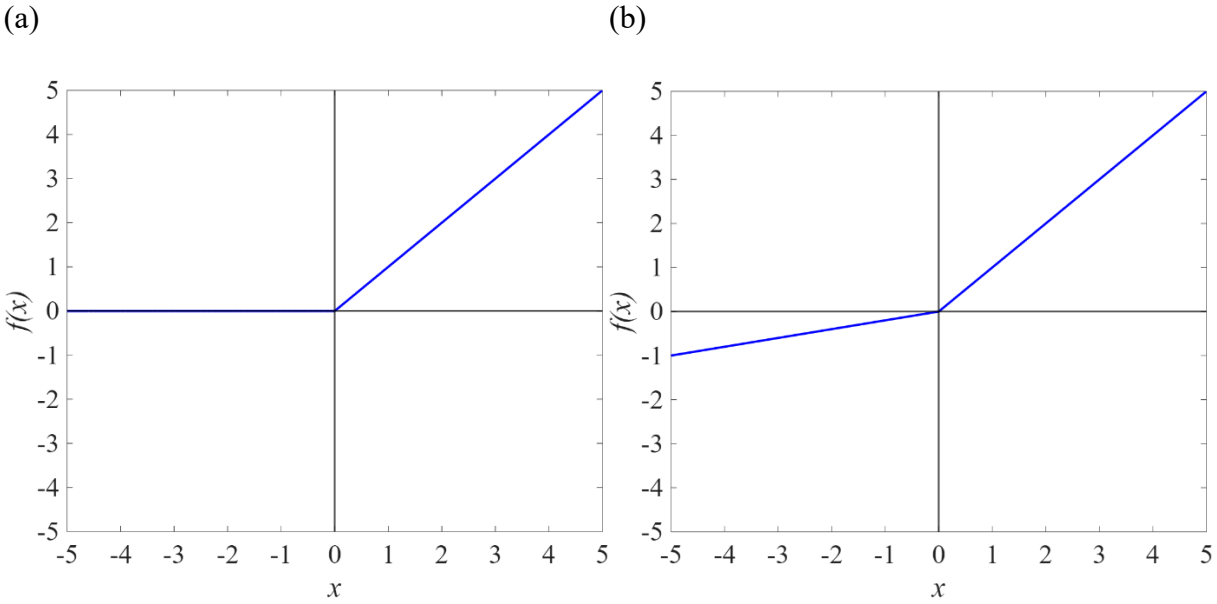


Figure 18 (a) ReLU activation function; (b) leaky ReLU activation function

The most commonly used optimiser for adjusting network parameters during training is gradient descent. This is an iterative algorithm that begins at a random point on the loss function and calculates the gradient as a vector of partial derivatives with respect to all independent variables. A new point is selected further down the slope and the algorithm progresses taking steps in the direction of the negative gradient until it reaches the lowest point of the function. Stochastic gradient descent uses a random subset of training examples rather than the entire dataset at once. It is often implemented with a momentum mechanism to prevent the updates overshooting and jumping wildly back and forth over the minima point, improving computational efficiency and speed of convergence. Momentum sends the steps in the general direction of the descending gradient, like rolling a ball into a bowl at an inclined angle (Figure 19). The ball will come to rest at the bottom of the bowl more smoothly than if it was dropped directly down from one side and allowed to oscillate back and forth.

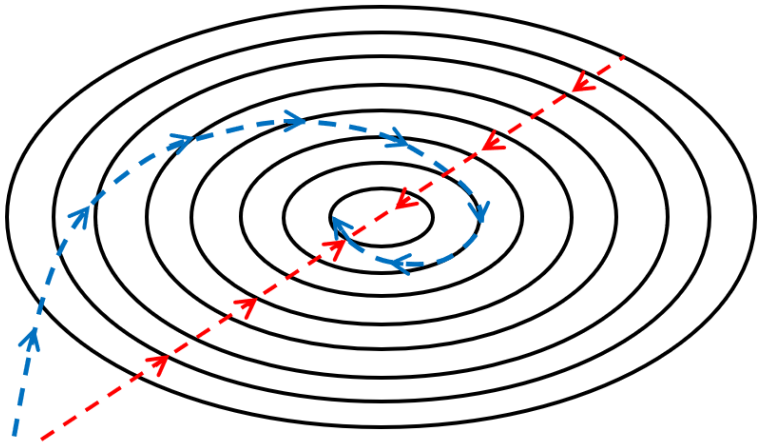


Figure 19 Stochastic gradient descent with momentum (blue) and without momentum (red)

Taking steps that are too large or too small can also skip over the optimal value and may cause the optimiser to converge on a local rather than global minimum. To avoid this a learning rate is used to scale the gradients and control how aggressively the algorithm changes the network weights. This is usually a small number, typically starting at around 0.001 and can be tuned by the process of learn rate annealing to ensure it is not so large as to overshoot the global minimum, but not too small to excessively slow down the convergence. A general rule of thumb is to divide the learning rate by 10 if the weights blow up or the network performs poorly. The learning rate may be either fixed for the duration of the training, or set to periodically decrease as the training progresses using a learning rate schedule. The schedule can be set by defining the period after which it drops and the factor with which it drops. A common problem with gradient descent is that the loss function gradient may change rapidly with some of the input parameters but slowly with others. This can cause the weights of some parameters to update wildly. To control this the gradients can be clipped by setting a maximum size on their values using a gradient threshold.

Even with an upper limit on the gradient values, if some values are at the maximum and others are near zero, it can be hard to choose the right learning rate. Smaller steps are desirable in the direction of rapid change, but larger steps are needed in the direction of slow change. A solution for this is to use different learning rates for each parameter. The adaptive gradient algorithm (Adagrad) adapts the learning rates for individual parameters depending on how frequently they are updated during training [63]. The more often a weight is updated, the lower the learning rate becomes. This is ideal for sparse datasets with missing entries, however the learning rates tend to get very small over time which limits speeds and efficiency. Root mean squared propagation (RMSprop) is a specialised version of Adagrad which accumulates gradients in a fixed window, keeping a history of their size which is used to scale the learning rate for each parameter. The adaptive moment estimation (Adam) optimiser scales the learning rate for each parameter similar to RMSprop, but also uses momentum to smooth out the updates [64]. This method is highly computationally efficient, fast, flexible and is well suited to problems involving large datasets, making it one of the most widely used optimisers for training deep neural networks.

Training deep neural networks to make accurate predictions often requires large amounts of data. Calculating the loss and gradient by passing the entire dataset through the network at once is computationally expensive. Because of this the dataset is usually split up into mini-batches. Each iteration of the training is then performed using a different mini-batch. When the entire dataset has been passed forward and backward through the neural network once this is referred to as one epoch. The batch size is set by the network designer depending on the

size of the dataset. The total number of batches is equivalent to the number of iterations required for one epoch. Multiple epochs are typically used to help the model generalise better over the training period. As the number of epochs increases, more parameters are adjusted, leading to a better performing network. However too many epochs can lead to overfitting, where the network overlearns complex details specific to the training data instead of the general patterns. The network then performs poorly on any new unseen data as a result. In general, the deeper the neural network, the more susceptible it is to overfitting as there are more hidden layers, hence more weight. Overfitting can be avoided by increasing the overall amount of data to improve generalisation, or by reducing the model size by limiting the number of parameters and capacity of the network. These solutions may not be practical if generating training data requires time consuming experimentation or the complexity of the problem requires many hidden layers. A common alternative approach is to separate the data into multiple datasets, one set for training which is seen by the network, and one for validation which is used to test the network's performance on unseen data. In addition to training and validation datasets a third "test" set is often used after the training is complete to evaluate the correlation between the network's predictions and known values. Network training performance can be measured by accuracy and loss. For classification problems mini-batch accuracy shows the percentage of images that the network classifies correctly, and should increase as the training progresses. For regression the root mean square error (RMSE) between the predicted values and the actual values is used instead. Mini-batch loss shows how far the network is from a perfect prediction totalled over the entire dataset, providing an indication of how confident the network is in the predictions it is making. Ideally RMSE should decrease towards zero and level off, while loss continues to decrease. The network predicts the correct data, then becomes better at making decisions, becoming more and more sure of its predictions as it passes the dataset multiple times over more epochs. The onset of overfitting can be identified by a rise in validation loss. If this occurs the training can be stopped to prevent the performance of the network dropping off with future iterations. In addition to early stopping, the distribution of weight values throughout the network can be regularised to constrain the complexity of the model. A cost associated with having large weights is added to the loss function. L_1 regularisation adds a cost with regards to the absolute value of the weights coefficient, or the L_1 norm of the weights. The penalised cost function is given by $\mathcal{L}_{l1}(\mathbf{x}) = \mathcal{L}(\mathbf{x}) + \lambda|W|$, where $|W|$ is the L_1 norm of the weights and λ is a user-defined value [62]. L_2 regularisation adds a cost with regards to the squared value of the weights coefficient, $\mathcal{L}_{l2}(\mathbf{x}) = \mathcal{L}(\mathbf{x}) + \lambda\|W\|^2$. In general the more data the network has to train on, the better the model will be at making predictions. Data augmentation can be used to artificially increase the size of the dataset by applying transformations to the existing set to

synthesise more data. For example images can be rotated, scaled or mirrored. Flipping, blurring or zooming an image can allow the network to find features that cannot be picked up by the human eye. Another method to prevent overfitting is neuron dropout [65]. A fully connected network layer will occupy most of the parameters and the neurons may begin to develop a co-dependency amongst each other during training, which curbs the individual power of each neuron and leads to overfitting. Dropout selects neurons at random at each iteration and deactivates them by setting their outputs to zero. As a result the training becomes more generalised, featuring a greater level of randomness which builds a more robust predictive model that is better suited to deal with new data.

The model parameters such as the weights and biases are variables that are internal to the neural network. They cannot be set manually, they are updated automatically during the training process. Their values define the skill of the model at making predictions and are saved as part of the trained network. Hyperparameters are external to the neural network and cannot be calculated directly from the dataset. Examples are the learning rate, batch size and number of epochs. These must be set manually and tuned to extract performance from the network. There is no clear-cut theory on how to find the optimal values of hyperparameters. Initial estimates can be made using values from similar studies in literature, then an extensive process of trial and error is usually required.

2.2.3 Convolutional Neural Networks

Convolutional neural networks (CNNs) are designed for applications where the input has an inherent 2-dimensional structure, for example an image data array. This makes them well suited for multifunctional metamaterial design applications where the input may be a 2-dimensional stress-strain characteristic, a sound absorption-frequency response or images from thermography scans. CNNs imitate the organisation of neurons in the visual cortex of the human brain, replicating the way in which humans process and recognise images. Unlike regular neural network layers, convolution layers do not use connection weights. Instead filters are used to scan the input image matrix to extract feature maps, which accentuate any unique features (Figure 20).

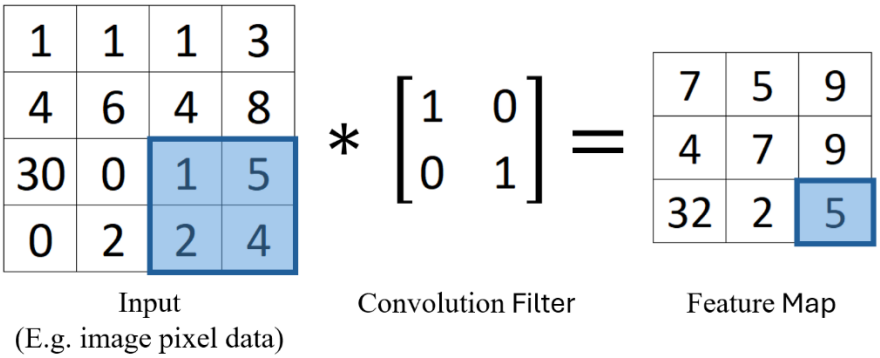


Figure 20 Example of convolution process

A bounding box moves across the data and performs a mathematical convolution operation on individual readings of parts of the image, it then sends the result to the corresponding neuron in the convolutional layer. This effectively reduces the number of parameters required and allows the network to be deeper without increasing computational requirements [62]. Mathematically, the convolution operation can be formulated as the dot product of the input function and a kernel function as follows:

$$(f * g)(i) = \sum_{j=1}^{m} g(j) \cdot f(i - j + m/2) \tag{2.17}$$

Pooling is used to produce a down-sampled matrix for each filter in the convolution layer, reducing the number of neurons necessary in subsequent layers whilst still retaining the most important information. This significantly reduces the computational load and prevents overfitting. Max pooling picks out the maximum value from the selected region of the data and places it in the corresponding cell of a reduced matrix. Minimum and mean pooling pick out the minimum and mean values respectively. The factor to which the down-sampling occurs is known as the stride. For example a stride of 2 will down-sample a vector with 12 dimensions to 6 dimensions. Convolution and pooling steps are repeated multiple times using the previous extracted features as inputs. This is typically followed by a number of fully connected hidden layers to extract a classification or regression result at the output of the network.

2.2.4 Potential for Applications of Deep Learning in Metamaterial Design

A major benefit of deep learning is that it is extremely flexible. It can be applied to virtually any multi-variable design problem as long as it is possible to generate a training dataset. Because of this the potential for applying deep learning to metamaterial design is vast. Once a model is trained it can produce extremely powerful predictions in minutes without ever seeing the governing theory and equations behind the physical phenomena being studied. It can be

used to solve problems for which there is no analytical solution, where there is only empirical data available through experimentation. Neural networks are able to extract complex features for describing non-linear relationships. They are able to fit predictive models to data that is otherwise impossible to analyse using standard polynomial regression techniques.

Tang et al. highlighted the ongoing requirement for combining DfAM considerations into metamaterial optimisations and used an artificial neural network (ANN) to predict the manufacturability of heterogeneous lattice structures that were optimised for stiffness using a BESO algorithm [49]. Many optimisation methods only consider DfAM constraints on strut thickness and minimum feature size, however other considerations such as inclination angle and overhang angles are likely to affect the quality and chance of successful manufacture considerably. The ANN method linked the design and manufacturing stages and can be used to automatically consider the effects of different geometrical parameters on the manufacturability of the optimised lattice, guaranteeing successful fabrication. The concept of "manufacturable element" (ME) was used which creates a predefined, parameterised decomposition of volumetric regions of the material. In this case the ME was defined as a lattice strut with its related geometry, material and manufacturing process data. The ME data was then input into a meta-model machine learning process to obtain the manufacturing constraints of each ME as the geometry parameters are varied. The neural network was trained on raw data of printing quality from fused deposition modelling (FDM) ABS samples of different ME configurations. Quality metrics used were the ratio of deflection of horizontal struts compared to their thicknesses, R_h, and the thickness deviation of slanted struts, R_s, used to measure the discrepancy ratio between actual and design thicknesses. R_h values greater than 25% and R_s values greater than 10% led to printing issues such as stringing. The targeted printing quality was assigned and a multilayer ANN with back-propagation was able to adjust the lattice design domain to satisfy these requirements.

Yang et al. demonstrated the potential of machine learning in accelerating material design optimisations by using convolutional neural networks to predict non-linear stress-strain responses of binary composites [66]. The high dimensionality of the stress-strain curves presented modelling issues for the CNN approach. The complexity of the problem was reduced by converting the non-linear responses into principal component analysis (PCA) coordinates before training, then converting the predictions back again afterwards. The input to the network was a binary matrix of 0's and 1's describing the locations in a 2-dimensional representative volume element (RVE) where the material should be soft (E=21MPa) and where it should be hard (E=2100MPa). The model was trained on FEA predictions from a crack phase field solver implemented in ABAQUS, able to simulate crack evolution through

the material. Crack phase field modelling allows simulation of complex fractures through the material including curvatures, branching and coalescence of crack pathways. A hybrid formulation method was used in which a crack phase field is coupled with the displacement field using the degradation of strain energy density. Crack nucleation and propagation was then modelled based on the critical energy release rate. This approach was highlighted as suitable for modelling of composites, offering more accurate fracture patterns and load-displacement responses compared to alternative commonly used anisotropic formulation. The CNN was able to approximate the non-linear behaviour well despite only using linear dimensionality reduction. Measurements of model performance based on predictions of elastic modulus, strength and toughness indicated a mean absolute error less than 10% of the range of values in the dataset, even though the network was only trained on 10^{-27}% of the possible microstructure configurations. The coefficients of determination indicated good correlation between predicted and known values. The R^2 values for modulus, strength and toughness were 0.888, 0.918 and 0.763 respectively. Liu et al. also developed a DNN method for non-linear modelling of heterogeneous materials, however they used a multi-scale approach featuring a collection of hierarchically connected building blocks to describe the structure-property relations across material length scales from individual phases through to the overall macroscopic material [67]. Four RVE configurations were simulated in LS-DYNA: uniform, circular matrix inclusions, amorphous geometry from the Cahn-Hilliard equation [68], and anisotropic with phases penetrated in one direction. 300 samples were generated for each configuration, 200 for training and 100 for validation. The neural network was able to achieve average errors below 0.7% for all configurations after 10,000 epochs and 9 hours of training.

Many previous works focus mainly on the prediction of material properties using neural networks. Few studies go beyond prediction to optimise specific properties for a particular application. Even fewer have gone as far as validating the resulting materials with full experimental mechanical characterisation. Large scale applications of machine learning optimised materials in industry are therefore currently still limited. One study that did validate results experimentally was conducted by Gu et al. who trained convolutional neural networks with FEA results from hundreds of thousands of microstructure designs of hierarchical composites for the purpose of maximising strength and toughness [69]. Each design was assigned a ranking according to the FEA predictions. The machine learning process was able to realise a coarse-grading approach, in which a collection of building blocks are condensed into a single unit cell, significantly reducing the number of parameters required and allowing accelerated design and analysis of the material without needing full microstructural material

characterisation data. The building blocks were made up of a binary composite arrangement of stiff and soft material regions (Figure 21).

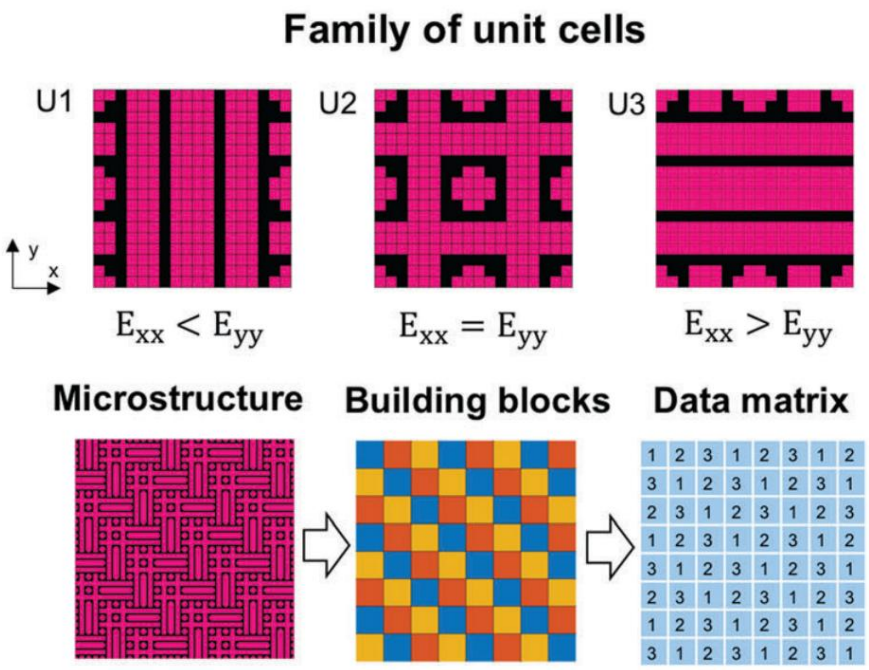


Figure 21 Unit cell configurations, pink and black colours referring to stiff and soft building blocks respectively (top) and microstructure consisting of an assemblage of unit cells that is then encoded (blue=U1, orange=U2, yellow=U3) and converted to a data matrix (bottom). From [69]

At the boundaries of the unit cells each type of element in one unit cell was connected to the same type of element in the adjacent cell to satisfy tessellation continuity. The degree of isotropy in the unit cells could be tuned by changing the orthogonal stiffness ratios. The assembled lattice was converted into a data matrix of building blocks by encoding the individual unit cells. This matrix representation was then used as input for training the CNN model. A dataset of 100,000 materials with randomly assigned unit cell arrangements (80% used for training and 20% for testing) was able to provide excellent predictions for normalised toughness and FEA performance rankings, despite representing less than 10^{-8}% of the entire design space of possible arrangements. Root mean square deviation was 0.2978 for training and 0.4996 for testing. The normalised toughness of the highest ranking designs extracted from the resulting model was approximately 13 times higher than the average value found in the original dataset that the network was trained on. Under experimental testing of AM samples the optimised composite offered performance considerably greater than that of its individual constituent base materials. It was approximately 25 times tougher and twice as strong as the stiff material and 40 times tougher and more than 100 times stronger than the soft material.

Previous studies that apply machine learning to composite materials typically deviate from conventional composite fibre/ply arrangements to binary configurations of discrete soft and stiff phases of material, distributed in a grid of blocks throughout the design space. This allows for hierarchical designs in which smaller block arrangements can be positioned to form higher level designs on a larger length scale. However this approach sacrifices the benefits of continuous fibre reinforcement found in traditional carbon fibre. One possible alternative may be to map stacking arrangements of carbon fibre plies to mechanical performance characteristics, however the size of the dataset required would likely make it very time consuming to generate. The binary composite could be expanded to feature more phases of varying stiffness. These stiff and soft regions could then be realised by multi-material additive manufacture. Instead of placing discrete blocks throughout the design space, the model could be used to determine the best arrangement for placement of continuous fibre strands within the microstructure. All of the above studies for predicting stress-strain behaviour considered 2-dimensional RVE problems under plane strain conditions only. The 2-dimensional discretisation approach greatly speeds up the generation of the training dataset using FEA. The model created by Gu et al. was able to analyse 100,000 microstructures in 5 days. The training process itself took up to 10hours and the predictive phase of the study took less than a minute to solve. Extension is required into 3-dimensional problems with more elaborate geometries, multiphysics effects and complex material laws such as those that govern carbon fibre reinforced composites.

2.2.5 Generative Design

Complex 3-dimensional geometries are likely to take months to create the number of data samples necessary for a neural network to make good predictions. If the dataset is limited in size, for example if it can only be generated through practical experimentation considerable simplification of the original design problem is often required. Generative design, first proposed in the 1970's by J. H. Frazer [70] does not require time consuming pre-generation of a training dataset. It instead works on a ground up approach that synthesises stochastic geometry by evaluating randomly generated iterations against a set of boundary conditions and design constraints using an adversarial network configuration, known as a generative adversarial network (GAN). A generator network proposes new designs, while a second discriminator network continuously evaluates them until an equilibrium between both networks is reached [71]. Any time constraints will depend more on the fidelity of the model, the number of cell divisions and the load cases considered. Structural generative design studies begin with the engineer defining a model genotype that comprises a design space, the

set of initial conditions and target parameters [72]. The AI analyses the structural load paths and performs a form synthesis process analogous to the way bone forms in the human body. Material is applied in regions of greatest stress in the same way calcium deposits build up in the human skeleton, as evidenced by astronauts losing bone strength when in zero gravity. The limits of the design space that can be explored are defined by an exploration envelope. The design parameters, initial values and limits are stored in a design table. The resulting phenotype designs are generated in a parametric CAD engine and evaluated using performance filters based on requirements such as strength, stiffness and manufacturability before further iteration and refinement. Thousands of iterations are able to be performed in the time it would take a human designer to make a single alteration. The generative solver presents the most suitable solutions to the engineer, who then uses their engineering experience to select the optimal result to take forward. When combined with additive manufacture, the design to manufacture time of a structural composite component could take a fraction of the time of manual design processes. New components could go from first concept to a final result fitted into an assembly in hours rather than days or weeks.

The first commercially available generative design solver was released by Autodesk in 2018 [73]. It functions using a method that can be described as “designing in the void”. Instead of only considering where material is required, the user may also define void spaces where material will not form. First key geometry regions that are essential in the final design, referred to as “preserve” geometry are modelled in CAD. Obstacle geometry is then added to the model to represent keep-out regions. An initial guess can be used by selecting starting geometry to reduce the design space, however unlike topology optimisation where this step is essential, in generative design this is optional. The solver can either remove material from an initial shape or it can grow freeform geometry to populate the design space entirely from scratch. Also unlike topology optimisation, manufacturing requirements can be set before the optimisation process begins. Production volume can be set so the solver knows if the result is a one off prototype or is to be mass produced. Multiple manufacturing methods such as CNC machining, casting or AM can be included in the same study and the model will produce separate solutions for each. Restrictions can be set for number of CNC axes available, tool direction, minimum tool diameter, casting draft angles, mould ejection direction, or in the case of AM, build direction, overhang angles and minimum wall thicknesses. The program is able to evaluate multiple base materials at the same time, and generates thousands more iterations than in a conventional topology optimisation study. This is made possible through a cloud based solver that links to multiple servers so that the model can be computed purely online, making the solutions extremely fast with low computational hardware requirements,

offering significant advantages over alternative genetic algorithm or particle swarm optimisation approaches.

A major research gap exists for the application of generative design for adding stochastic bio-inspired architecture at the material level. The majority of previous applications have been at component level, largely initial proof-of concept prototypes and bespoke low production volume products. Autodesk partnered with General Motors for the design of a prototype seatbelt bracket (Figure 22). Generative design was able to reduce an 8 part assembly down to a single component that weighed 40% less while being 20% stronger [74,75].



Figure 22 General Motors generative design seat belt bracket iterations [75].



Figure 23 BAC Mono generative design wheel [76].

Briggs Automotive Company (BAC) were able to reduce the mass of the wheels on the BAC Mono by 35% using generative (Figure 23) [76]. The solver was able to incorporate manufacturing constraints during the optimisation to ensure the resulting geometries could be produced by 5-axis CNC milling. Generative design is being used by Hyundai to facilitate design of their elevate walking car concept vehicle by allowing for lightweight strong legs

that are able to handle the torque requirements of the vehicle's electric motors [77]. In the aeronautical sector, Airbus created a bionic partition to separate the passenger cabin from the galley of the Airbus A320 [78]. Generative reduced the mass of the partition by 45% (30kg). Airbus estimated that the application of the design approach across their fleet of A320 planes could save up to 465,000 metric tonnes of CO2 emissions every year. In a previous study by the author generative design was applied to Formula Student suspension uprights, which were manufactured by selective laser sintering (SLS) of titanium Ti6Al-4V powder [79]. The resulting uprights were 34% lighter than the CNC equivalent machined from aluminium 7075-T6.

A potential barrier currently exists for large scale uptake and implementation in mass produced parts as results typically require thorough post validation using FEA and practical testing and often require extensive post modelling or further tuning of boundary conditions to achieve performance targets. Studies are often prone to material separation, in which one of the preserve regions becomes separated from the rest of the model as the geometry grows, leading to non-convergence. This issue can usually be solved by inserting additional loading into the initial conditions at the location of the preserve geometry in the direction of the problematic region to ensure formation of material. There are significant potential benefits for applying generative design at the material level for stochastic lattice cell architectures. Individual struts will be morphed into an optimised shape and thickness depending on the localised loading conditions. Many periodic lattice structures, particularly the octet-truss have high degrees of redundancy due to high nodal connectivity ($Z = 12$ in the case of the octet-truss). Generative optimisations can allow redundant struts to be removed, while maintaining the level of nodal connectivity required for rigidity and stretch-dominated behaviour along key load paths. Multiple loading scenarios are able to be satisfied at the same time and optimisations can be carried out over multiple length scales, from global component level, to local cell divisions, down to microstructural hierarchy orders. Completely freeform biomimetic designs are possible where locations of nodes within the lattice are no longer fixed and can be relocated to where they are needed depending on the loading conditions. The density of nodal locations can be greater in certain areas than others, allowing the designer to tune relative density and enable graded lattice structures for multifunctional energy absorption applications. Regions of solid material can still be incorporated as necessary by setting preserve geometries, gaining the advantages of both cellular lattice and bulk solid in the same architecture. The lattice will transition smoothly between solid and porous regions and density can progress to solid material at the outer layers if required, analogous to the formation of cancellous and trabecular bone structures.

2.3 Materials and Manufacturing Methods

2.3.1 Mechanical Assembly

Early lattice structures relied on traditional manufacturing processes such as investment casting, drawing and folding of perforated sheets, assembly of woven wires or slotting of beam elements [19]. Tetrahedral lattices can be created by perforating rolled steel sheets with a hexagonal punch before folding periodically in a die. This creates a significant amount of waste material and requires sufficient bonding to ensure nodal strength. Mechanical assembly of metallic lattices usually requires high-temperature brazing post processes such as transient liquid phase bonding (TLP) [80]. The forming processes are typically time consuming and labour-intensive. Limited geometrical complexity makes biomimetic and stochastic features near impossible to manufacture in this way. The methods are only compatible with a limited selection of periodic unit cells. If a new unit cell size or topology is required the entire process must be reworked with new machinery. Structures can only be created on the mesoscale, microstructures and hierarchical structures are not possible due to limits on minimum feature size. Investment casting is a cheap alternative that is able to produce complex geometries and can be scaled to large components. However it is challenging to achieve relative densities lower than 5% due to difficulty filling mould channels with small diameter. The need for high fluidity limits the possible selection of base materials. Results are also commonly susceptible to porosities and build defects.

More recent developments take advantage of composite materials to achieve superior mechanical properties. George, Deshpande and Wadley manufactured pyramidal truss cored lattices from carbon fibre reinforced polymer (CFRP) using a snap-fitting method, achieving relative densities in the range of 1-10% [81]. The snap-fit patterns were manufactured from 0/90° carbon fibre laminate sheets by water jet cutting (Figure 24). Pyramidal truss geometries were then assembled by snap fitting which occurred at each of the nodes. The lattice was then bonded into composite face sheets using an epoxy adhesive and curing cycle. The trusses were countersunk into cruciform-shaped slots cut into the face sheets. The resulting lattice never exhibited failure at the nodes, the main failure modes observed were Euler buckling of struts and delamination of the laminate. The snap-fit lattice was competitive with conventional honeycomb, however did not achieve the full potential of composite cores predicted by micro-buckling analysis due to delamination failure of the struts.

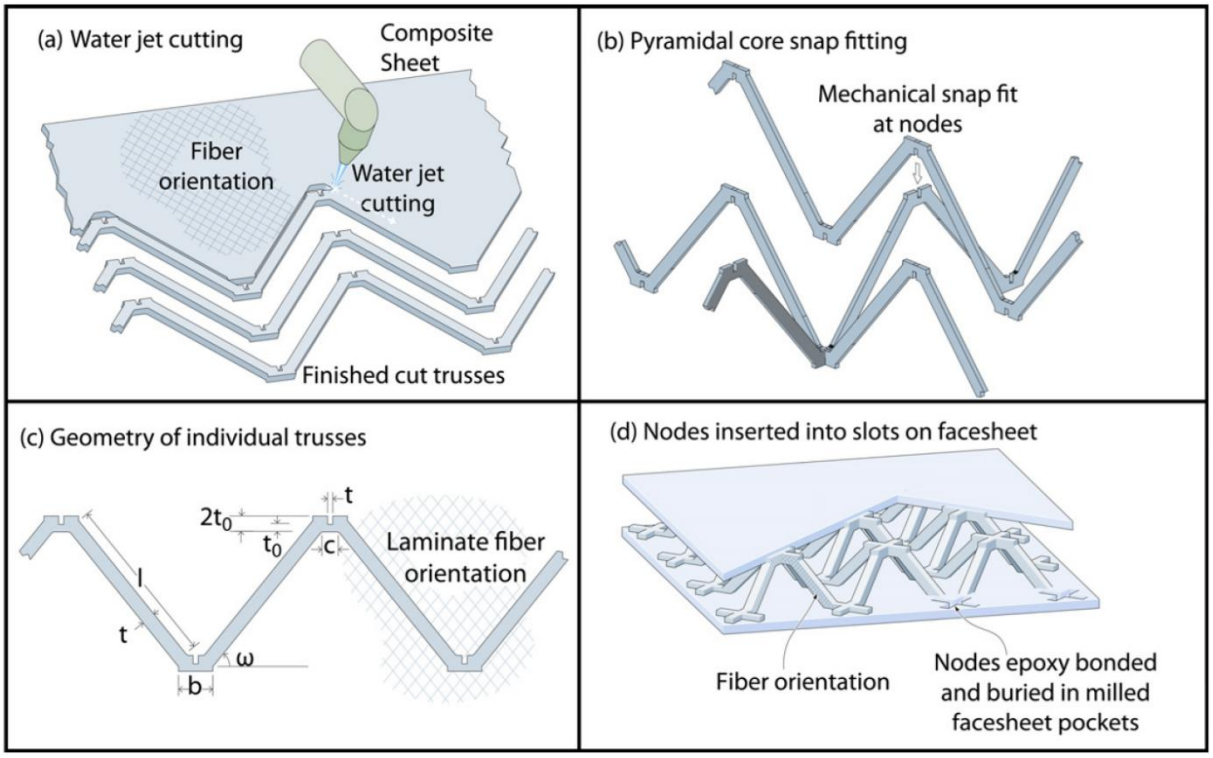


Figure 24 Pyramidal truss cored lattices manufactured by snap-fitting CFRP laminates. From [81]

Liu et al. designed octet-truss lattice structures with hollow cylindrical struts of unidirectional carbon fibre, which were able to outperform laminated strut-based counterparts [82]. The struts were joined by custom designed CFRP connectors fabricated via assembly of sub-components cut from a 0/90° laminate using a micro-water jet (Figure 25). A CNC mill was used to machine the tips of the connectors. The struts were then snap fitted to the connectors and bonded with a high strength epoxy. The failure mode was dominated by fibre fracture of the struts, except at low relative densities ($\bar{\rho} = 1.33\%$ and 0.85%) where the failure mode transitioned to Euler buckling.

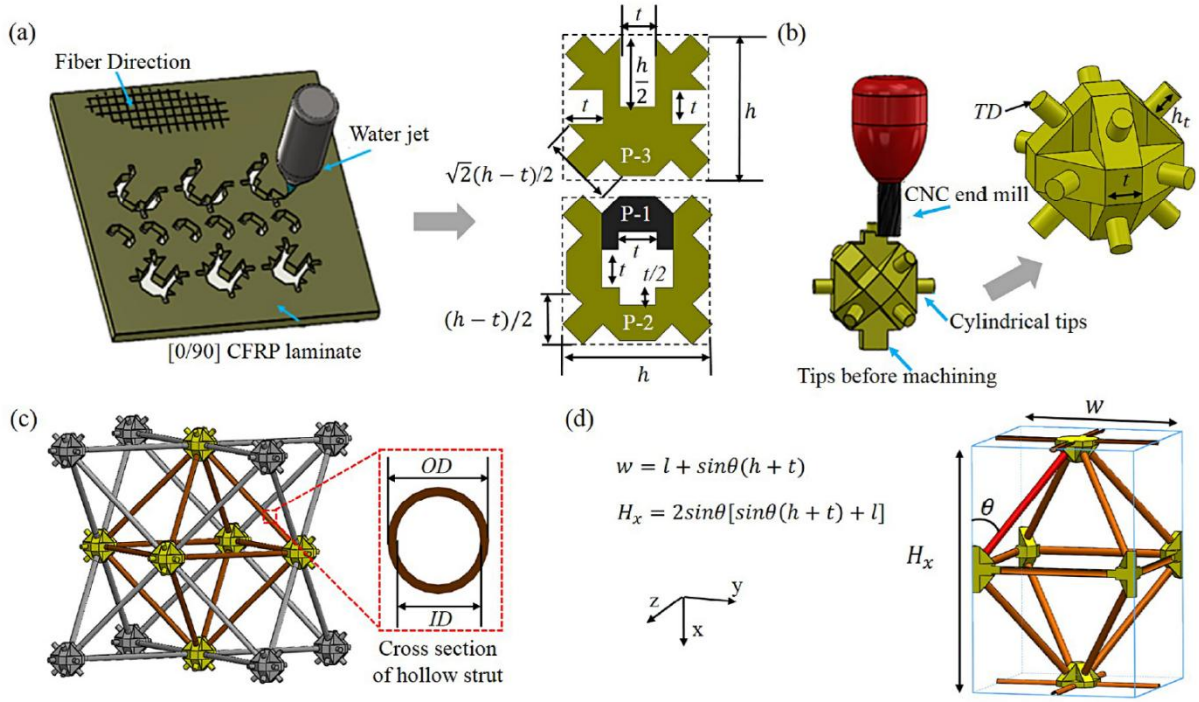


Figure 25 Illustration of the fabrication approach for CFRP cylindrical tube-based octet lattice structures. (a) A micro-water jet cuts sub-components from CFRP plate, (b) CNC milling of the cylindrical tips of joint connectors, (c) a single CFRP tube-based octet lattice structure, and (d) a single unit cell of the octet lattice structure. From [82]

Jiang et al. designed an interlocking orthogrid lattice of carbon fibre composite panels [83]. Prepreg carbon fibre strips were cured to form ribs of trapezoid cross section, with periodically spaced slots along their length. The strips were then interlocked to form an orthogonal grid and bonded to face sheets with epoxy resin. The structure showed particular strength under compression, competitive with even metallic lattices. High ductility was observed through a long progressive deformation plateau that offered greater specific energy absorption than alternative composite lattice configurations including pyramidal honeycomb, corrugated truss and octet-truss. The continuous layers of adhesive resin between the orthogrid and the face sheets offered greater resistance to shear compared to standard trusses that are only bonded at the nodes. However susceptibility to skin flexural failure was observed at low face sheet thicknesses.

Cheung and Gershenfeld reversibly assembled a “cuboct” lattice of vertex-connected octahedrons, inspired by the structure of the mineral perovskite, from planar components linked at the nodes by shear clips (Figure 26) [84]. Rather than cutting the components from a quasi-isotropic laminate, they were instead sliced from a wound fibre stock (Figure 26(e)) so as to align the orientation of the fibres in the direction of the struts. As shown in Figure 26(d), the reinforcing fibres were wound around the connecting holes and along the axis of the struts

to optimise load-bearing capacity. The lattice was assembled from cross-shaped components arranged so that each central connecting node met the strut ends of 4 adjacent components, as shown in Figure 26(a). A shear clip was then inserted through the four coincident connection holes to fasten the strut ends together. Connections were designed to be stiffer and stronger than the struts to maintain stretch dominated behaviour that scales according to the response of the struts rather than the response at the nodes. A Young's modulus of 12.3MPa was reported at a density of 7.2mg/cm^3, an order of magnitude greater than alternative ultralight materials of similar density. The modular lattice solution offers greater flexibility than snap fitting approaches, allowing for heterogeneous regions and localised tuning of material behaviour by changing the relative locations of more-rigid and less-rigid components. The components are able to be mass produced and assembled to fill any arbitrary shape and unit cell sizes can be scaled to match the length scale of the application.

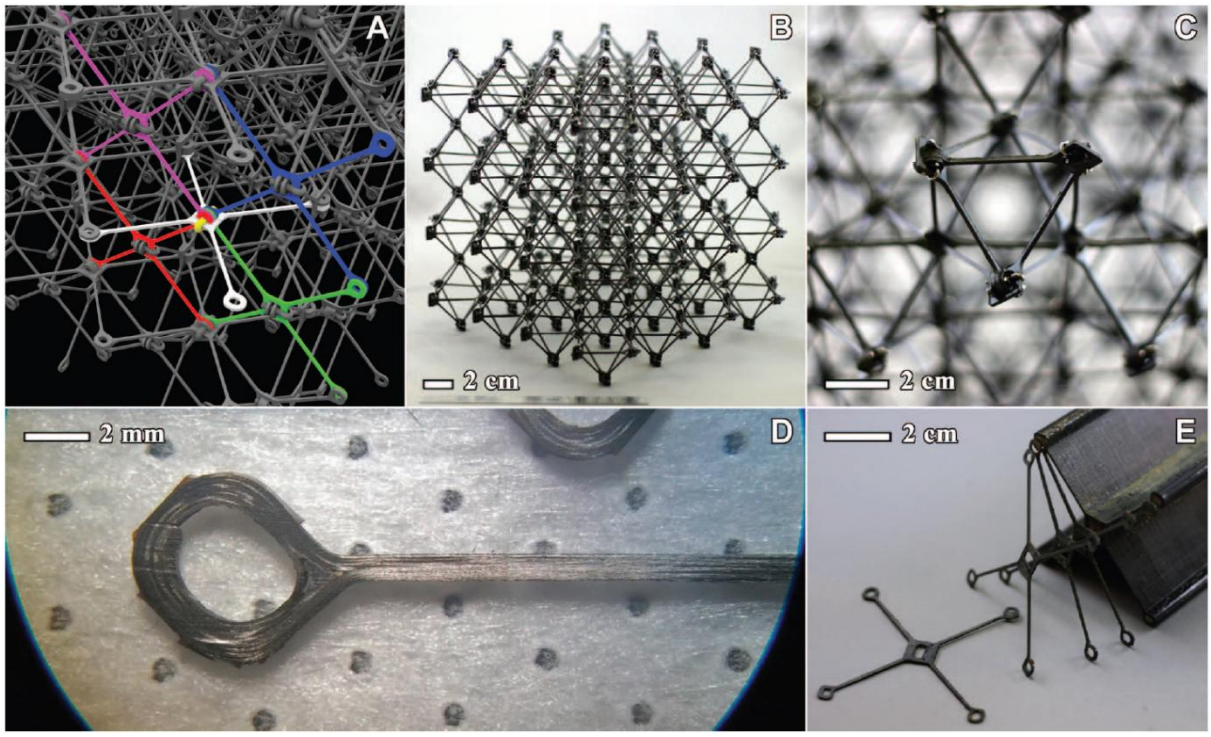


Figure 26 (a) "cuboct" lattice assembly; (b) carbon fibre–reinforced polymer composite cuboct lattice; (c) reversible joints; (d) micrograph showing fibre orientations; (e) slicing of components from wound fibre stock. From [84]

Zhao et al. designed a similar reversible mechanical interlocking method for assembling lattice structures, however they used non-planar shapes which helped avoid misalignment problems during assembly [85]. The non-planar cross-shaped parts were injection molded using long-glass-fibre-reinforced polypropylene thermoplastic. Overlapping square loops were fixed with chamfered pins at the nodes to assemble the cross-shaped components into an octahedral lattice. A handheld ultrasonic welding process was used to manually bond the

resulting lattice to face sheets cured from glass/PP prepreg, which significantly improved compressive stiffness. Due to asymmetry of the cross-shapes 120 alternate unit cell arrangements could be created by changing the assembly sequence, allowing the lattice to be redesigned for a specific loading scenario. Most unit cells showed higher specific stiffness than equivalent structures fabricated by planar cross-shaped parts made of CFRP.

The traditional Japanese art form of origami is gaining increased popularity with material scientists, inspiring assembly processes based on paper folding techniques. Ahn et al. combined direct-write assembly with a wet-folding technique based on origami [86]. A concentrated titanium hydride (TiH_2) ink was deposited using direct-writing to create complex 2D lattice patterns that were then folded, twisted or rolled into 3D origami structures. These patterns were then transformed to metallic and ceramic structures by thermal annealing. Yan et al. introduced an approach that exploits controlled mechanical buckling for reversible autonomic assembly of 3D origami structures, which can be applied at multiple length scales, from the mesoscale down to the nanoscale [87]. Placement of thickness variations within a substrate material was able to create engineered folding creases under compressive stress. This enabled active control over intermediate states as 2D patterns cut in the material transformed into complex 3D origami structural arrangements using buckling-induced deformation. The 3D configurations could be dynamically and reversibly tuned by loading and unloading of the substrate through stretching or bending. A diverse range of origami inspired geometries were demonstrated, including hierarchical forms. The concept can be expanded to produce kirigami lattice structures by placing patterns of cuts at appropriate locations. This can reduce the effects of stress concentrations and introduce unique deformation modes that are otherwise unobtainable.

Established mechanical assembly techniques suffer from increased manufacturing time and complexity, and commonly require structural compromises associated with weaknesses at joints. Braiding and weaving techniques can be used to arrange fibre tows to form continuous composite lattice structures. Maypole braiding uses a meshed gear arrangement to weave together continuous fibre tows in interlacing patterns to improve the strength-to-weight ratio of tubular shafts and beams [88]. However it requires complex and expensive 3D braiding machinery which usually takes up a lot of space. The approach is generally limited to low production volumes. There is also limited complexity of geometries and a lack of design flexibility, as changing the design of the lattice usually involves redesigning the machine hardware configuration. A research gap exists for lattice structures that can take advantage of the high strength-to-weight ratio of continuous fibre composites, combined with complex bio-mimetic stochastic geometries. Zhai and Zhang proposed a method of continuous spatial

weaving of scalable lattice structures from continuous fibre braids [34]. Octet-truss lattices were manufactured up to a meter-scale in overall size, with a minimum critical feature size of 350μm. The lattice was woven in three-dimensional space from braids of glass and Kevlar fibres (Figure 27). After weaving the structures were impregnated with low-viscosity epoxy resin and cured at a fixed temperature.

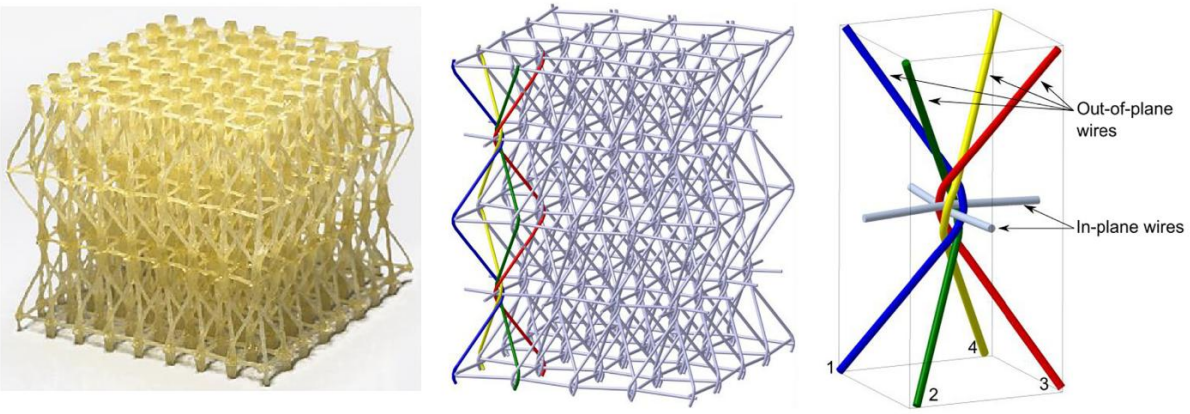


Figure 27 Glass fibre octet-truss lattice manufactured by continuous spatial weaving. Adapted from [34]

The material properties could be tuned by varying geometric parameters of the octet-truss, however location-specific tailoring is not possible. The approach offers good manufacturing scalability to larger components, however small unit cells are impractical. The method can be applied to other unit cell architectures and graded density structures are possible. However it only allows for periodic, recurring unit cell designs. Gaps between the braids due to insufficient winding lead to manufacturing defects. Kevlar fibres were observed to be more difficult to tighten, resulting in less straightness in the struts. Redundant resin may accumulate and remain within the structure post-curing, particularly in the centre of the nodes. Achieving good geometrical accuracy and homogeneous results proved difficult. High fibre misalignment was observed at nodal locations. In this case the design and manufacturing method are coupled, one depends on and is limited by the other. The combination of continuous fibre with additive manufacture on the other hand may be able to remove these limits rather than imposing them.

2.3.2 Additive Manufacture

Additive manufacturing has the potential to make lattice metamaterials more affordable, more practical and more widely readily available due to the availability of low-cost open-source hardware and software. Between 2001 and 2011 the average cost of additive manufacturing dropped by 51% [89]. Advances in additive manufacturing have enabled opportunities for metamaterials with complex 3D structures with non-periodic topologies, offering precision at the micro and even nano scales. In the AM process geometries are nested onto a build plate using a slicer software package to split the topology into a series of thin layers. Material is then laid down progressively to build the three-dimensional structure up layer-by-layer. The methods for depositing material and bonding subsequent layers vary between approaches. The oldest technique for polymeric AM is stereolithography (SLA). SLA was developed in 1986 and was the first commercially available AM technology [90]. It is a vat photopolymerisation process in which an ultraviolet laser is used to cure photopolymer resin. The laser scans over the print area, activating cross-linking of the polymer and solidifying the resin in the regions required. The platform is then moved and the process repeated to solidify the next layer. SLA has high manufacturing accuracy and offers resolutions that allow for sub-micron feature sizes. However compared to other approaches it is relatively slow. Large quantities of lattice material may take days to manufacture, making it impractical for high volume components. In most AM methods there is a trade-off between minimum resolution and manufacturing speed. A gap currently exists for multi-stage additive processes that are able to take pre-built micro-scale lattice building blocks and combine them at the macroscale to allow for practical production rates of multi-scale lattices. Photopolymer waveguide is able to produce very small features to obtain hierarchical microstructures and functionally graded materials. Lattices can be produced with ultra-low relative densities below 5%. Large layer sizes at the centimetre scale can be developed in short exposures of several minutes, orders of magnitude faster than standard layering methods, however complex arbitrary shapes with extensive overhang angles are not possible [91]. Projection microstereolithography (Figure 28) exposes photosensitive resin to UV light patterns projected via a reconfigurable digital photomask through a reduction lens, achieving minimum feature sizes less than one micron, albeit with overall build volume restrictions [13]. The resulting lattices can be combined with nanoscale coating methods such as electroless plating of nickel or copper and atomic layer deposition of ceramics to create ultra-low relative density lattices below 0.1%.

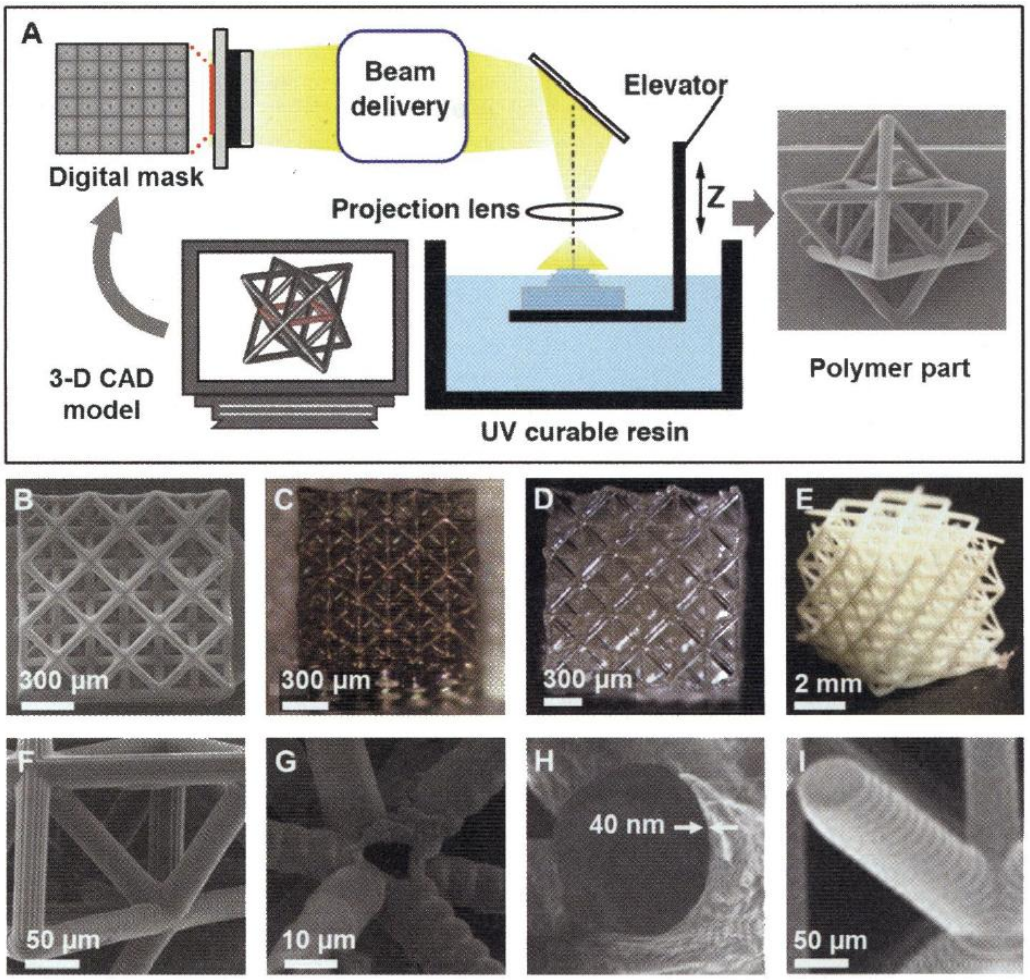


Figure 28 (a) Projection microstereolithography process; (b-e) fabricated octet-truss microlattices; (b) solid polymer hexanediol diacrylate (HDDA); (c) hollow-tube metallic nickel-phosphorus (Ni-P); (d) hollow-tube ceramic (alumina); (e) solid ceramic (alumina); (f-i) magnified views of the struts of the microlattices in (b) to (e), respectively. From [13]

Metallic structures can be additively manufactured via either direct energy deposition or powder bed processes, the predominant methods of which are selective laser sintering (SLS) and electron beam melting (EBM). In SLS the geometry of each layer is traced across a bed of metallic powder by a high powered laser, which sinters the particles together. The bed is then lowered and new powder is spread across the build area and the process is repeated. Minimum feature size depends on the size of the particles of the metal powders. Large unit cell sizes above approximately 5mm are difficult to fabricate due to the support structure requirements of large overhang areas. Unit cell designs must generally be largely self-supporting in nature due to difficulty in removing support structures within the lattice. SLS causes greater build-up of residual stresses due to a large thermal gradient between the melt pool and the lower layers, requiring more extensive post processing heat treatments to relieve. EBM uses an electron beam focused through electromagnetic lenses to melt the powder. It is slightly faster and the build chamber is preheated to higher temperatures which reduces thermal stresses and holds

the powder together better so less support structure is required. SLS produces better surface finishes and can achieve tighter geometric tolerances, between 0.05 to 0.1mm compared to 0.2mm for EBM. The downsides of both approaches are expensive hardware and material costs and long build times, though very high strength and stiffness base materials such as titanium Ti-6Al4V can be used to achieve lattice struts with high stiffness-to-weight ratio.

Fused deposition modelling (FDM), also known as fused filament fabrication (FFF) uses extrusion of melted filament through a heated nozzle to deposit successive layers of material onto a print bed [92]. The ability to tune layer thickness can make the process relatively fast. Hardware is inexpensive and widely available, and compatible with a wide range of recyclable thermoplastics. Multi-material FDM printers are able to combine multiple materials within the same lattice. Water-soluble PVA filaments can be used for support structures, which can be dissolved in a water bath post build. This allows for incredibly complex geometries deep within the lattice structure without requiring access for removal of support material. FDM is however prone to significant accumulation of residual stresses due to thermo-cyclic loading in subsequent layers [93]. It is highly temperature sensitive and susceptible to warping and stringing phenomena, difficulties with first layer separation and nozzle jamming that may result in build failure.

FDM can be used to produce fibre reinforced composites by combining the polymer matrix with reinforcement fibres. Thermoplastics combined with discontinuous chopped fibre reinforcements are currently commercially available and can provide an inexpensive boost in material properties, readily compatible with off-the shelf FDM printers. Quantitative investigations carried out by Hofstätter, T. et al. demonstrated a direct relationship exists between fibre content and both ultimate tensile strength and Young's modulus in carbon fibre reinforced PLA up to 30% fibre content [94]. They found chopped fibre reinforced FDM thermoplastics to be an environmentally friendly, cheap and fast method for flexible rapid prototyping, helping to significantly improve component lifespan and reduce crack propagation velocity [95]. However they found that material properties were significantly affected by fibre orientation and highly dependent on the flow characteristics through the nozzle and deformation on nozzle exit. They highlighted a key drawback in short carbon fibre reinforcement, in that increases in Young's modulus were limited by a detrimental effect on ultimate tensile strength, concluding that the fibre-matrix interface requires improvement for components in high strength applications. Ivey et al. found the same drop off in tensile strength occurred while increasing Young's modulus of FDM printed PLA from 3.37GPa to 5.68GPa through the use of short carbon fibre reinforcement [96]. However this offset was negated by annealing at 115°C for 2 hours to allow recrystallisation of the polymer and

removal of void content, improving Young's modulus further to 6.73GPa, approximately twice that of the unreinforced thermoplastic.

In recent years materials scientists have begun to experiment with continuous fibres in additive manufacturing with the goal of producing material properties competitive with conventional composite layup. Brooks, H. attempted a method in which continuous fibre tows are threaded through channels left in the component during the AM build [97]. Carbon fibre and Kevlar reinforcements were able to increase stiffness of a pulley, 3D printed in PLA, ABS and PA12, by approximately 200% in most cases. Kevlar reinforcement of ABS increased elongation to failure by 2660%. Reinforcement of the PLA pulley with carbon fibres caused a 762% increase in tensile strength over the unreinforced pulley, improving elongation to failure by 521%, as shown in Figure 29.

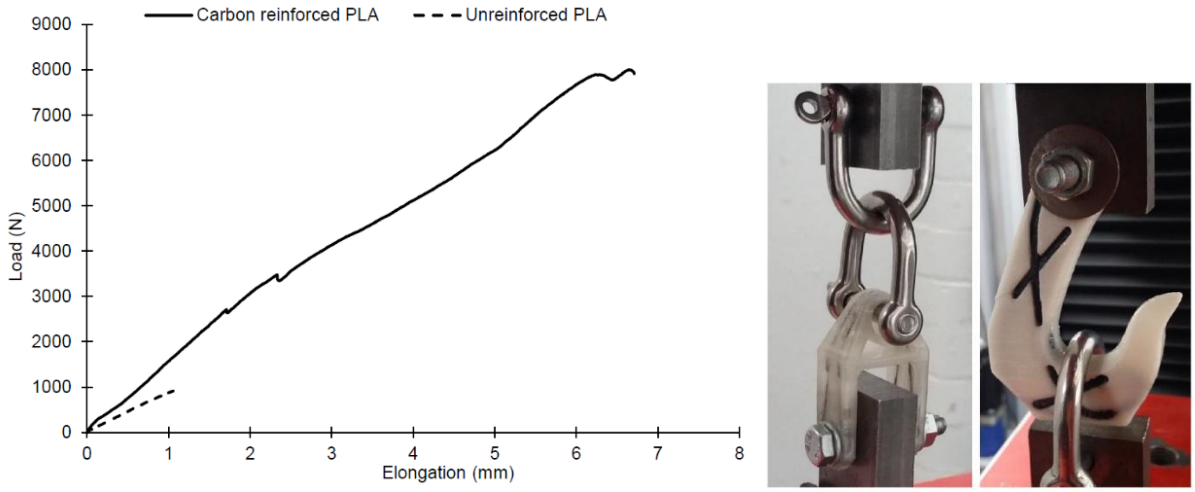


Figure 29 Load-elongation curves for PLA pulley housings (left) and test fixtures (right) [97].

At 174 and 178 kPa/(kg/m³), the specific strength that was achieved by reinforcing the PLA and PA12 pulleys with continuous carbon fibre exceeded additively manufactured aluminium Al-Si-10Mg and was comparable to 7075-T6 aerospace grade aluminium alloy, while able to be manufactured at a fraction of the time, at a cost an order of magnitude lower than metal AM.

In 2014 Ken-ichiro Mori et al. developed a die-less forming process for manufacturing carbon fibre reinforced ABS by sandwiching continuous carbon fibre bundles between AM build layers using an FDM printer [98]. They found that a thermal bonding process between the fibre and matrix was vital in improving the material's strength and by switching to an in-situ impregnation of the fibres directly into the FDM nozzle, the ultimate tensile strength of the resulting composite was increased by 76%. Masahito et al. successfully reinforced a PLA matrix with both carbon fibres and natural jute fibres using in-nozzle impregnation during

FDM [99]. The tensile modulus and strength achieved for carbon fibre reinforcement were 19.5 (± 2.08) GPa and 185.2 (± 24.6) MPa respectively, 599% and 435% improvements over standard PLA. Jute reinforcement resulted in properties of 5.11 (± 0.41) GPa and 57.1 (± 5.33) MPa, 157% and 134% that of unreinforced PLA. During FDM each layer cools rapidly as heat flows into the colder layers below, deteriorating the conditions for recrystallisation and setting up residual stresses between layers which results in poorer material properties. To address this issue Yamawaki et al. studied the effects of heat treatments on continuous carbon reinforced nylon formed through in-situ FDM [100]. On performing a hot-isostatic press treatment, holding a test piece at 270°C and 0.2MPa static pressure for 10 minutes, Young's modulus was increased from 53 to 91GPa, exceeding that of aluminium. Tensile strength was doubled from approximately 700 to 1400MPa. FDM is currently limited to thermoplastics, however the firm Continuous Composites has developed a process for continuous fibre composite additive manufacturing in which thermoset resin is cured instantaneously using UV light, allowing printing in free space without the need for support structures (Figure 30) [101]. In 2017 they exhibited an AM composite rudder featuring nichrome wire for de-icing on the leading edge and fibre optics built into the outer skin of the material able to detect changes in temperature, acceleration and stress.



Figure 30 Continuous Composites continuous fibre additive manufacturing. From [101]

Several key issues require investigation to allow continuous carbon fibre reinforced additively manufactured (CCFRAM) materials to become competitive with conventional composites before widespread investment and uptake in industry can take place. Typical AM components currently suffer from high levels of void content compared to conventional materials, limiting the maximum strength and stiffness. The resulting material is also likely to feature a large degree of anisotropy, the material properties will depend largely on the direction in which the fibres are laid throughout the component. There is also currently limited theory on the failure modes unique to continuous fibre composites such as fibre pull out, delamination between build layers and fibre buckling. Current commercially available CCFRAM machines are very

expensive, Markforged offers a machine capable of printing carbon reinforced nylon using fused filament fabrication (FFF) for $20,000 [102].

To explore the potential of CCFRAM for lightweighting a generative study was carried out by the author using material properties from Markforged to design a formula student suspension upright from continuous carbon fibre reinforced nylon. The resulting upright design is shown in Figure 31 below. The continuous carbon fibre reinforced component satisfying the same design space, with identical key geometry hardpoints and load cases was 81% lighter than the CNC aluminium equivalent. The properties quoted by Markforged of 60GPa Young's modulus and 800MPa tensile strength were on par with properties found in literature using modified FDM equipment, and below those achieved through hot-isostatic pressing. The manufacturing time of the continuous carbon fibre reinforced polymer (CCFRP) upright is expected to be 6-8 hours. The titanium upright was printed successfully in a 26 hour build [79]. Layup and curing of an equivalent conventional carbon fibre component would be expected to take 50+ hours, not including tooling & mould design and manufacture. The total material cost is expected to be £150-200. The titanium version cost £1050 per build including build prep, consumables (gas, filter), material costs of £450 per kg of Ti6Al4V alloy and £200 for wire EDM removal. Conventional carbon fibre layup including materials, moulds, surface prep, manual lamination, debulking, curing and pattern block machining would be expected to be in excess of £2000.

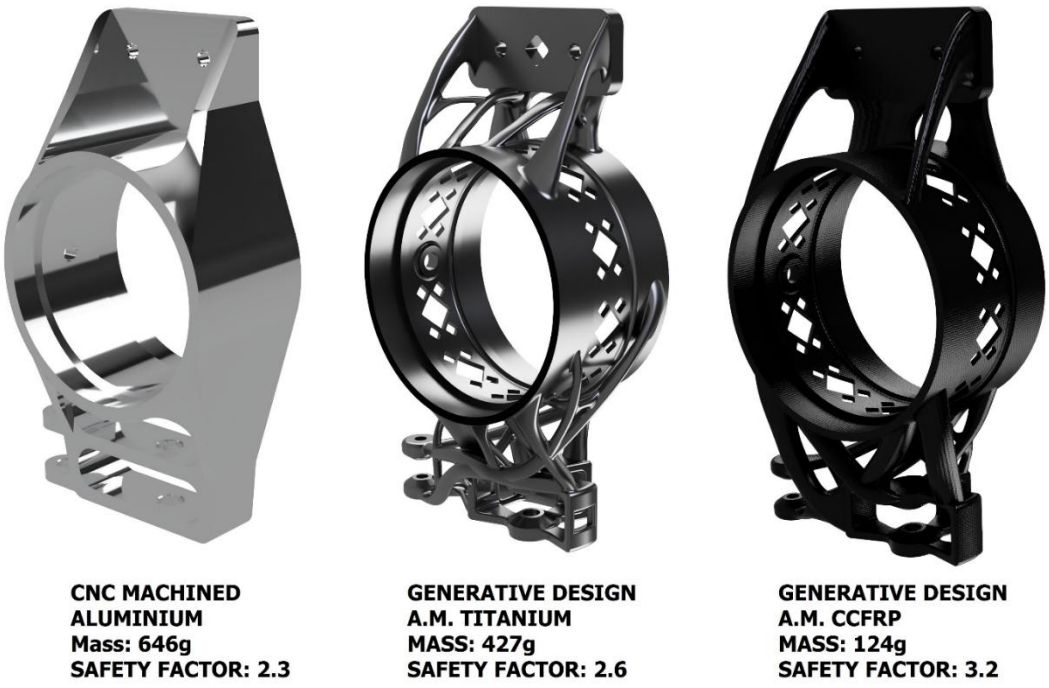


Figure 31 CCFRP demonstration component – Team Bath Racing 2020 Formula Student car rear suspension upright

2.4 Multi-scale Hierarchical Design

Many natural materials exhibit cellular structure on more than one length scale. Structural hierarchy is present when structural elements have their own structure at a smaller scale. In typical lattice structures there is a strong direct relationship between compressive strength and density. As density decreases strength reduces significantly, limiting the potential for lightweighting. Natural materials evolve by organising 3D cellular architecture over a range of length scales from nano to macrostructure, overcoming this drawback. Each hierarchy level contributes towards the bulk material properties. Fleck, Deshpande and Ashby derived the effects of combining topologies over multiple length scales [8]. For a hierarchical lattice with 2 levels, where level 1 exists at the coarse-scale and contains level 2 struts at the finer scale, the effective elastic modulus E_2 of the fine-scale struts scales with the modulus of the parent material E_S and the relative density of the fine-scale struts according to:

$$E_2 = B_2 \bar{\rho}_2^{b_2} E_S \tag{2.18}$$

If the volume fraction of material 2 within the lattice material 1 is $\bar{\rho}_1$, then the relative density of lattice material 1 to the solid is

$$\bar{\rho} = \bar{\rho}_1 \bar{\rho}_2 \tag{2.19}$$

The effective modulus of the coarse-scale lattice E_1 is given by

$$E_1 = B_1 \bar{\rho}_1^{b_1} E_2 = B_1 B_2 \bar{\rho}_1^{b_1} \bar{\rho}_2^{b_2} E_S \tag{2.20}$$

For the case where $b_1 = b_2$, this simplifies to

$$E_1 = B_1 B_2 \bar{\rho}^{b_1} E_S \tag{2.21}$$

A high stiffness material can therefore be achieved by combining stretch-dominated architecture at both length scales, where $b_1 = b_2 = 1$. If both levels contain bending-dominated structures, then $b_1 = b_2 = 2$ and the material will be highly compliant. A bird's wing contains stretching-bending hierarchy. At the coarse scale $b_1 = 1$ and the structure is stretch-dominated. At the finer scale $b_2 = 2$ and the microstructure is bending dominated. This allows for strength to be maintained at low density to enable flight, while the energy absorbent microstructure is able to resist impacts.

Material designers are able to replicate bio-organic hierarchy in man-made metamaterial structures. Zheng et al. presented scalable metallic hierarchical metamaterials with high strength, ultra-low density, high compressive and superelastic tensile behaviour [103]. The

materials were designed over length scales ranging across seven orders of magnitude, from centimetres at the macro-level down to nanometres at the smallest features. A successive reduction of feature sizes between each hierarchy was achieved with periodic recurring geometry made from hollow nano-tubes at the lowest scale (Figure 32). A network of hierarchical stretch-dominated octahedron-tetrahedron unit cells was manufactured by selective laser melting of liquid monomer combined with nanoscale post-processing to form a nickel-phosphorus alloy coating before chemical removal of the polymer core. The hierarchy was generated in a top-down approach. At each scale filaments were assembled along load paths defined by the architecture at the higher length scale. A hierarchical lattice generator script was developed to combine stretch-dominated and bend-dominated topologies. At the first order of hierarchy the material was optimised for stretch-dominated behaviour, while the second order was bending-dominated. The resulting metamaterial was highly deformable, by approximately 20% in tension and over 50% in compression, exceeding any previous lightweight metal foam or lattice. The hierarchy offered tuneable scaling between strength and density that depends on the 3D topologies at each design level. At ultra-low densities below 0.2% an unprecedented scaling existed between yield strength and relative density of exponent 1.3, far exceeding first-order lattices which scaled with powers above 2.5. At this ultra-low density regime, yield stress was two orders of magnitude higher than first-order periodic lattice alternatives.

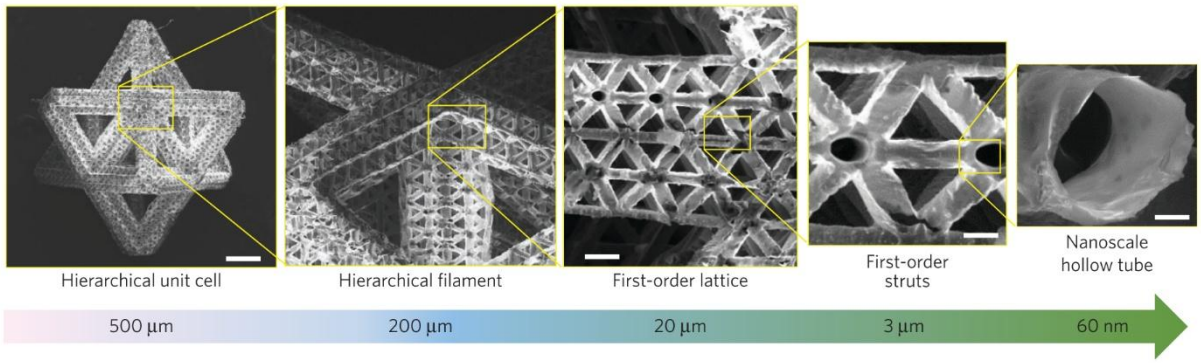


Figure 32 Scanning electron micrographs of structural hierarchy present in the unit cell of hollow nickel alloy metamaterial. From [103]

Jang et al. fabricated hierarchical octahedral unit cells by direct laser writing (DLW) of photopolymer using two-photon lithography (TPL) [104]. The structures were then coated in titanium nitride (TiN) particles using atomic layer deposition (ALD), before etching out of the polymer to create hollow ceramic struts of nano-scale thickness. The maximum tensile stress observed in the resulting unit cells was 1.75GPa, an order of magnitude higher than most standard brittle ceramics. Under compression testing the metamaterial did not fracture even

after repeated cycles. This was attributed to nano-scale size effects of the hollow struts. At the reduced length scale, the distribution of sizes of any flaws that exist in the material is greatly reduced, lowering the probability that a crack or void could be large enough to initiate failure. For a brittle material, the fracture strength is given by $\sigma^f = K_c/\sqrt{\pi a}$, where K_c is the fracture toughness and a is the initial flaw size [104]. The strength of a ceramic lattice is governed by pre-existing flaws in the weakest links. For such structures the fracture strength becomes $\sigma^f \propto (1/t)^{1/m}$, where m is the Weibull modulus and t the strut wall thickness. The smaller the wall thickness, the higher the fracture strength. For nano-scale wall thicknesses the hierarchical sizing effect greatly increases the strength of the weakest struts across the material, enabling high strength at ultra-low density.

A research gap exists for non-uniform, non-periodic designs that recreate randomly evolved naturally occurring biological hierarchy by applying stochastic geometry over multiple length scales. Most previous studies in literature that apply hierarchy make use of periodic unit cells formed from geometric primitives. Ajdari et al. studied 2-dimensional hexagonal cellular honeycombs featuring first and second order hierarchy levels, created by placing smaller hexagons at every vertex of the tessellation to form a fractal-like structure (Figure 33) [105].

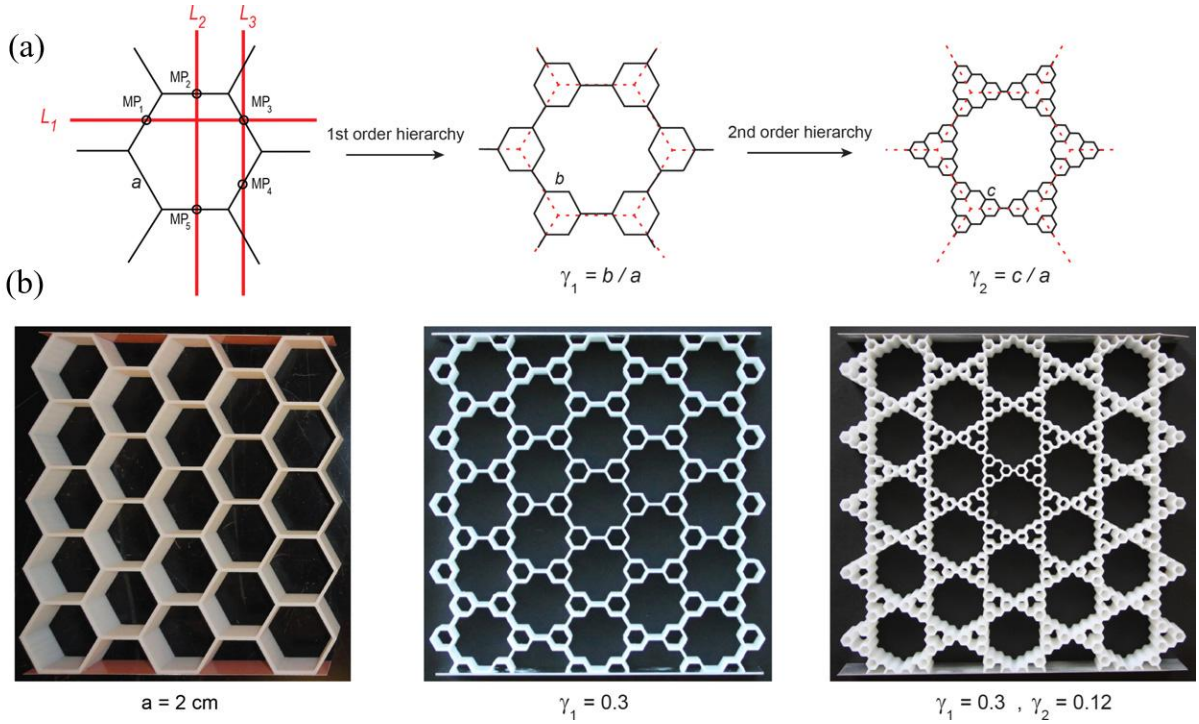


Figure 33 Hierarchical honeycombs; (a) Unit cell of the hierarchical honeycombs with regular structure and with 1st and 2nd order hierarchy; (b) Images of honeycombs with a = 2 cm fabricated using three-dimensional printing. From [105]

First and second order honeycombs were 2.0 and 3.5 times stiffer than regular honeycomb of the same mass. Increasing the level of hierarchy resulted in a broad range of elastic material behaviour that can be tailored by structural organisation of the hierarchy. The mechanical behaviour could be varied by altering the ratios of the introduced hexagonal edge length to the original hexagon's edge length at each hierarchy level. Poisson ratio was able to be tuned between 0.28 and 1.0 depending on the dimension ratios. They noted this could be extended to include the thicknesses of the cell walls to achieve greater levels of control, though at the cost of increased complexity of the problem. Such behaviour could potentially be mapped through a machine learning model.

Yang and Alam used hierarchical honeycomb lattices to mimic the structure of unidirectional fibre reinforced composites, creating a "pseudo-composite" material [106]. Regions with additional hierarchical subdivisions and greater stiffness act as the pseudo-reinforcement fibres, while less dense regions without hierarchy act as the pseudo-matrix. This allows for alignment of the hierarchically reinforced "fibre" regions in the direction of the loading, meaning the material response can be tailored to the given application. The strength of the honeycomb in the in-plane direction was increased as a result, something that standard honeycombs suffer with. They typically show good strength and stiffness in the out-of-plane direction only. The elastic modulus of the pseudo-composite metamaterial could be predicted using the Voigt model rule of mixtures from the modulus of the pseudo-fibre and pseudo-matrix. Models were subjected to tensile and compressive simulations in COMSOL Multiphysics and compared against a CFRP model at 57% fibre volume fraction. The honeycomb structures were able to mimic the overall stress-strain response curve shapes of unidirectional composites through the incorporation of the structural hierarchy, showing similar stages of mechanical behaviour. The fibre contribution coefficient, defined as $\eta_{FCC} = \frac{\sigma_{composite}}{\sigma_{fibre}} \times 100$ was greater than that of CFRP, indicating the pseudo-composite has more effective stiffness improvement from the combination of multiple phases, due to superior geometrical arrangement of the phases. Specific stiffness was improved compared to the standard honeycomb. At 1.13MNm/kg the best performing design was at the upper range of non-porous polymers and comparable with octet and octahedral lattices. The material showed considerable ductility and capacity for energy absorption. This was also observed by Sun et al. who studied the crashworthiness of vertex based hierarchical honeycombs in out-of-plane impacts using non-linear FEA [107]. First order hierarchy improved specific energy absorption (SEA) by 81.3% compared to regular honeycomb and second order hierarchy boosted SEA by 185.7%, while maintaining similar values of peak crushing force. SEA values were higher than most commonly used metallic foams and superior to the theoretical

limits of tubular-based structures, one of the most popular configurations for crash safety applications.

Hierarchy was applied by Tang and Yin in kirigami-based auxetic metamaterials [108]. Kirigami enabled great levels of compressibility by compaction of the cut-outs via buckling under load. In tension rigid rotation of the cut units lead to high stretchability, realising extreme levels of compressibility and stretchability within the same structure. When applied to re-entrant lattice geometry in silicone rubber, the second order hierarchical porous lattice achieved 81% compressibility, while exhibiting a zero Poisson's ratio without lateral expansion of the structure.

Wu et al. investigated implementing structural hierarchy in composite pyramidal truss lattices by combining the standard design with a second order corrugated-pyramidal structure within each of the struts of the first order unit cell (Figure 34) [109]. Mechanical performance was characterised under out-of-plane and transverse compression, shear and three-point bending. Failure maps were generated for sandwich panels with second order hierarchical lattice cores under compression, considering 5 competing failure modes: face sheet wrinkling and crushing of the second order trusses, core member buckling and crushing of the second order trusses, and macro Euler buckling of the first order pyramidal truss. At small core relative densities, compressive strength of second order structures was almost four times greater than the first order equivalent. However under shear the stiffness of the second order lattice was slightly lower than the first order of equal density, as the majority of load was carried by the face sheets of the second order structure and contribution of the struts was negligible. Second order hierarchy did show significantly greater shear strength though due to superior resistance to buckling of struts. The hierarchy was able to supress the macro-Euler buckling failure mode under in-plane compression.

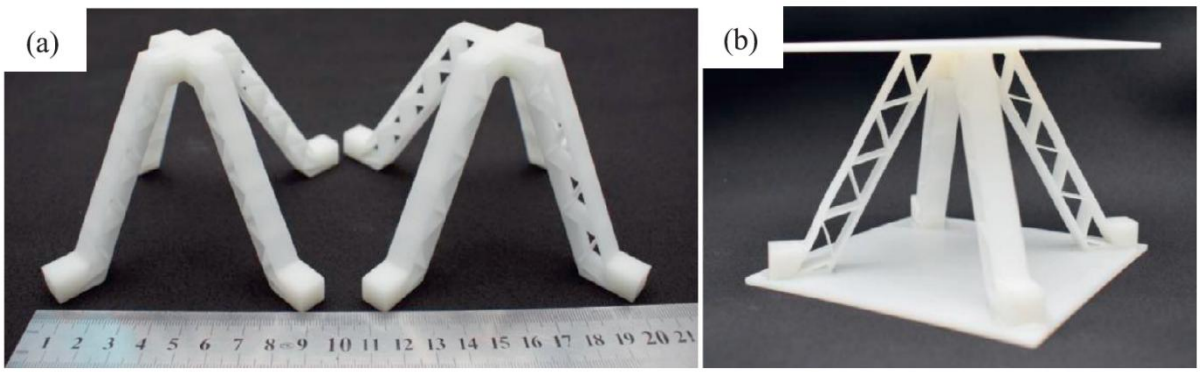


Figure 34 Second order hierarchical pyramidal truss structure; (a) without and (b) with face sheets. From [109]

The effects of geometrical parameters including truss aspect ratio and inclination angle at each level of hierarchical truss lattices were studied by Yin et al. [110]. The unit cell designs used were octet (Z=12, fully rigid), octahedron (Z=8, periodically rigid) and tetrakaidekahedronal (Z=4, non-rigid). By using hollow struts manufactured by SLS of nylon 12, the hierarchical octet-truss lattice was up to 46.1% stronger than the equivalent standard lattice. Mesoscopic and macroscopic configurations had equal effect on specific stiffness and failure strain. Macroscopic configuration was more important than mesoscopic for specific strength. The effects of hierarchy were analysed for both self-similar and hybrid lattices constructed by using different unit cell designs for each hierarchy level. Octet-octet, octahedron-octet and tetrakaidecahedron-octet combinations were studied, the former of each pair being applied at the macroscale and the latter at the mesoscale (Figure 35).

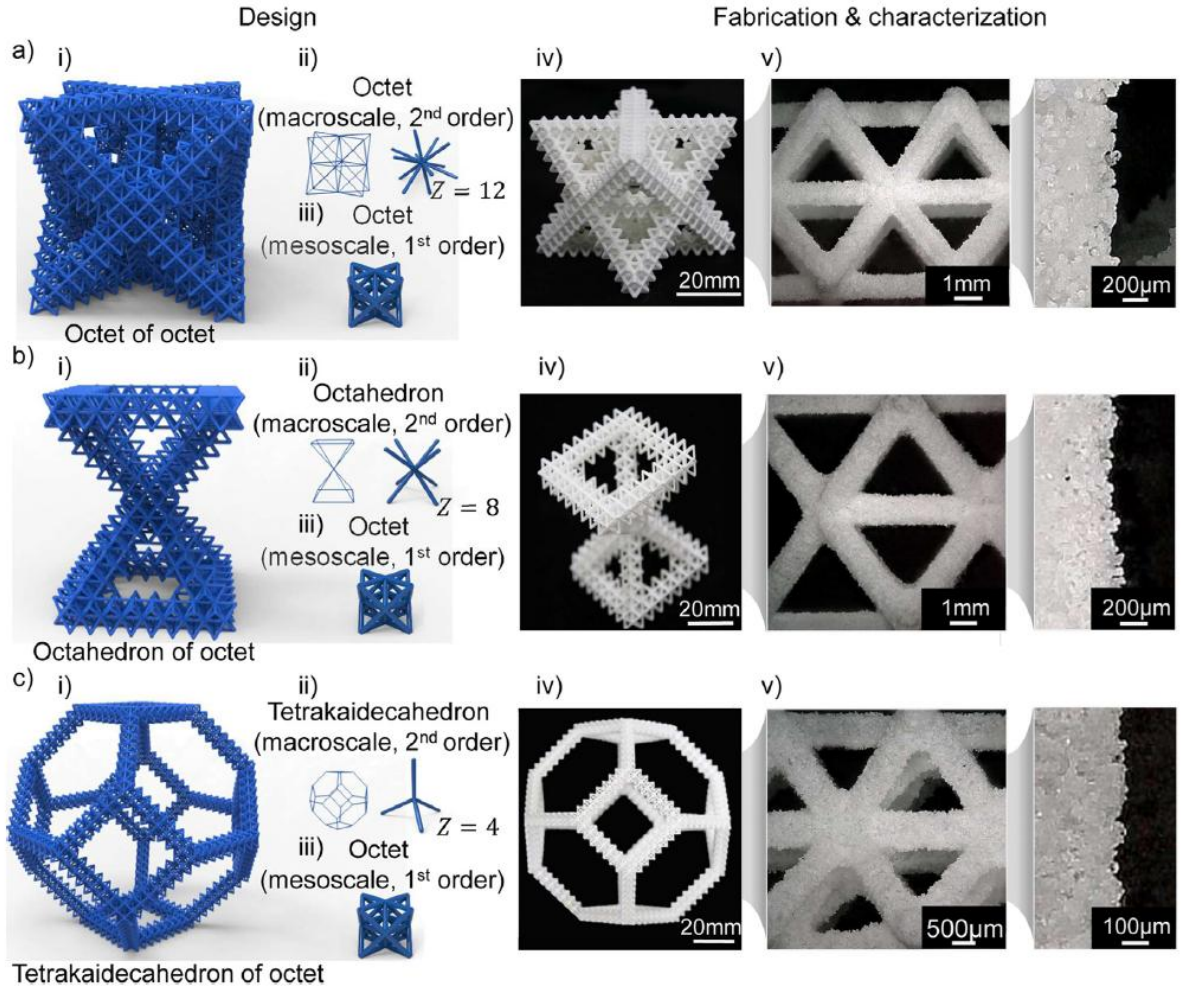


Figure 35 Hierarchical truss lattices [110]. (a) Octet of octet (OT –OT); (b) octahedron of octet (ON –OT) and (c) tetrakaidecahedron of octet (TN –OT). Figs. (i) are CAD models of the integral structures; (ii) are macroscopic configurations showing nodal connectivity Z; (iii) are mesoscopic lattice configurations; (iv) are hierarchical lattice materials fabricated by SLS 3D printing and (v) are the corresponding graphs showing microstructures at nodes and struts.

Specific strength increased as hierarchy was added and could be tuned by altering the hybrid combinations and strut thicknesses. FE analysis of hybrid configurations presented large variations in mechanical responses that are difficult to predict and purposefully design for without further multi-variable optimisation analysis. Second order hierarchy introduced large fluctuations in the stress-strain curves associated with layer-by-layer failures in the material. Third order hierarchy showed more homogeneous deformation, a sharp stress drop after the initial peak followed by a long stable stress plateau before densification. A wide range of mechanical responses were also observed by Zhang et al. who presented a framework for design and characterisation of hierarchical structures based on triply periodic minimal surfaces (Figure 36) [111].

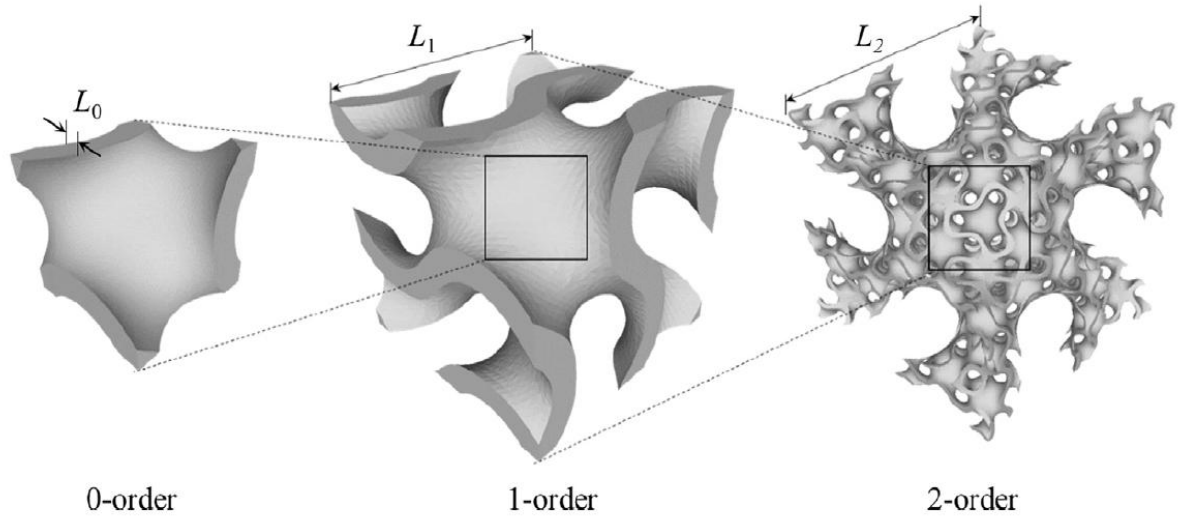


Figure 36 Hierarchical triply periodic minimal surface. From [111]

Hierarchies were generated by introducing smaller scale sheet TPMS structures within the walls of gyroid unit cells. Specimens were then manufactured by micro-selective laser melting (SLM) of stainless steel at relative densities between 0.031-0.071. Second order Gyroid lattices exhibited stable failure with a flat stress plateau and outperformed the single order equivalent in terms of strength under buckling at ultralow densities. However at moderate densities second order hierarchy showed reduced stiffness and strength compared to the single scale lattice. There is currently lack of clear theory for tuning hierarchy orders to obtain optimal material properties. Machine learning based approaches are likely to offer the advantage of being able to map the relationships between unit cell geometry parameters at each level to the mechanical responses without requiring underlying theoretical solutions. Being able to predict complex responses from combining stretching and bending dominated structures at different hierarchy levels using a deep neural network would be highly beneficial to enhancing the effectiveness of the next generation of architected cellular materials.

Structural hierarchy can also be used to enable and enhance multifunctionality in cellular materials. Shi and Akbarzadeh demonstrated the ability of hierarchy to improve figures of merit of ferroelectric metamaterials for sensing and energy harvesting applications [112]. First order body centred cube (BCC) unit cell configurations exhibited direct mode piezoelectric energy harvesting figure of merit (FOM_{33}) more than 50 times higher than the base ferroelectric material. Adding second order hierarchy using a hybrid octet-truss/BCC arrangement increased this by a further 50.7%. Hybrid combinations of different unit cell designs for each hierarchy level (Figure 37) outperformed self-similar equivalent structures. The hybrid second order octet-truss/BCC design presented FOM_{33} value 15.8% higher than that which used only BCC and 17.4% greater than those using only octet-truss at all levels. Furthermore, at ultra-low densities second order hierarchy was able to augment hydrostatic figure of merit, current responsivity, voltage responsivity and pyroelectric energy harvesting figure of merit by 89.8%, 118.3%, 222.2% and 603.5% respectively.

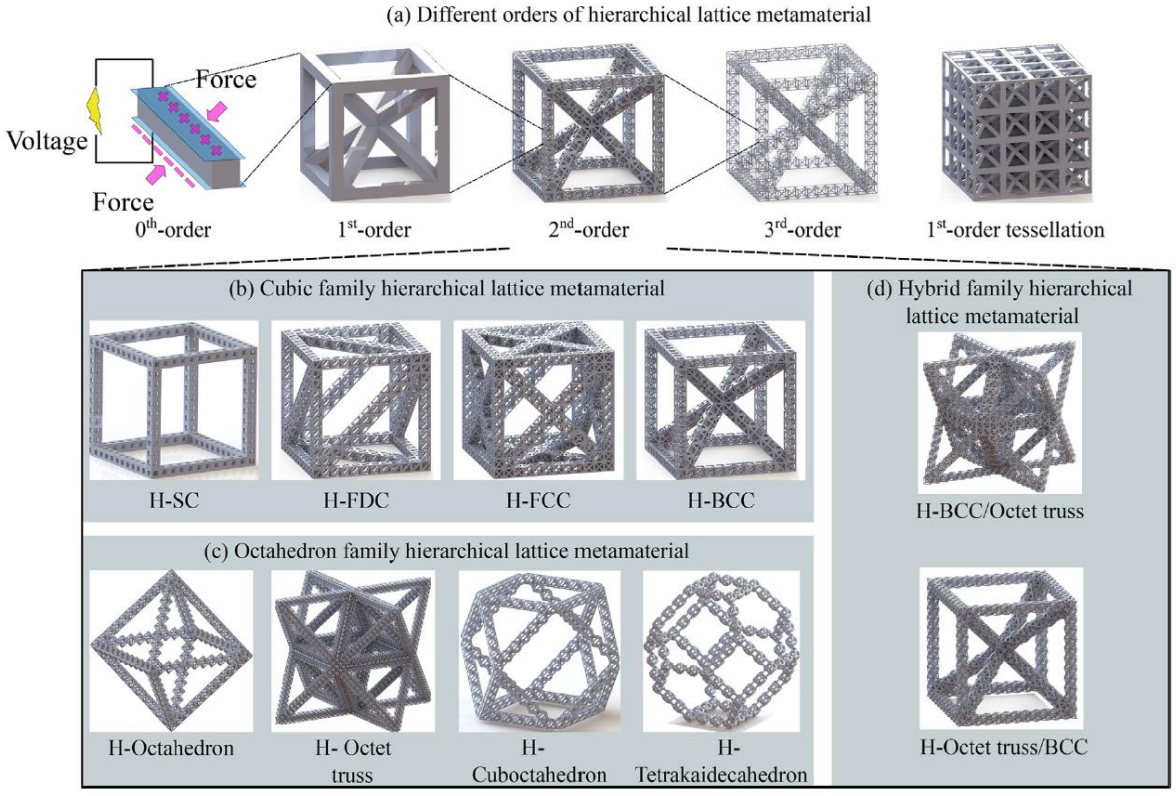


Figure 37 Hierarchical metamaterial configurations; (a) orders of hierarchy; (b) cubic family (simple cube, face diagonal cube, face centred cube, body centred cube); (c) octahedron family; (d) hybrid family. From [112]

2.5 Multifunctionality

2.5.1 Energy Absorption and Impact Resistance

Structural hierarchy can be an effective method for enhancing energy absorption of lattice metamaterials. This was demonstrated by Al Nashar and Sutradhar who showed that introducing hierarchy levels increases the number of struts within the lattice that fail via the mode of buckling, which facilitates energy absorption mechanisms [113]. FE analysis indicated the energy absorption capacity of a second order octet-truss lattice could be increased by between four to five times under low strain rate (20mm/s) compressive loading compared to the first order equivalent at the same density. The orientation of the second order lattice cells with the horizontal XY plane significantly affected the mechanical response. The energy absorption could be further boosted by 20% by orientating the cells at a 45° inclination. Truncated octahedron, dodecahedron and octet-truss unit cells were compared. The hierarchical truncated octahedron provided the greatest energy absorption efficiency under low loading, whereas at higher loads the octet-truss was superior due to its greater nodal connectivity and higher stiffness. The highest Von Mises stress concentrations occurred at very few locations. Less than 5% of the total lattice struts exhibited the highest stress. Of the rest the majority experienced only 10-15% of the maximum stress, indicating large amounts of material density within the lattice remained redundant. Further localised optimisations using machine learning at each hierarchy level may further augment strength-to-weight ratio by removing redundancies. The effect of cross-sectional geometry of the second order struts on the energy absorption performance was examined, comparing solid square, solid circular, hollow circular and hollow hexagonal cross sections. While solid struts offered greater stiffness, the volumetric energy absorption efficiency was highest for hollow trusses. The reduced stiffness facilitates greater levels of buckling within the lattice as it deforms, which absorbs more energy. Square cross sections offered greater stiffness than circular. Machine learning optimisations may be able to produce irregular stochastic cross sections for further localised stiffness augmentation. A drawback of the study is that it was limited to low strain rates. Further investigation into dynamic behaviour under high velocity impacts is required to evaluate effectiveness of hierarchical structures for vehicle crashworthiness.

An and Fan created energy absorption structures by combining aluminium foam with a hierarchical arrangement of cylindrical CFRP tubes for reinforcement (Figure 38) [114]. This approach was inspired by the natural structure of luffa sponge. Luffa features a cylindrical macrostructure consisting of micro-porous networks, separated by macro-pore regions. The struts are also hollowed at the microscale, resulting in four levels of hierarchy. The macropore

regions have a stiff inner surface that strengthens and maintains the shape of the structure during deformation. This biological hierarchy arrangement was replicated by inserting 25 thin-walled CFRP tubes into an aluminium foam cylinder to mimic the stiff layers around the macro-pores of the sponge. Under quasi-static axial compression the hybrid hierarchical material samples showed increases in yield force and mean crushing force between 40-73% over standard aluminium foam, depending on the density. The deformation of the CFRP tubes was stabilised by the surrounding foam, leading to a flatter, more progressive stress plateau and increased specific energy absorption. The SEA was comprised of the linear summation of the contributions from the aluminium foam and CFRP tubes plus an additional interaction effect that enhanced the hybrid material so that the energy absorption was greater than the sum of its parts. SEA was significantly improved, by a maximum ratio over 70% compared to standard foam, exceeding many metallic truss lattices.

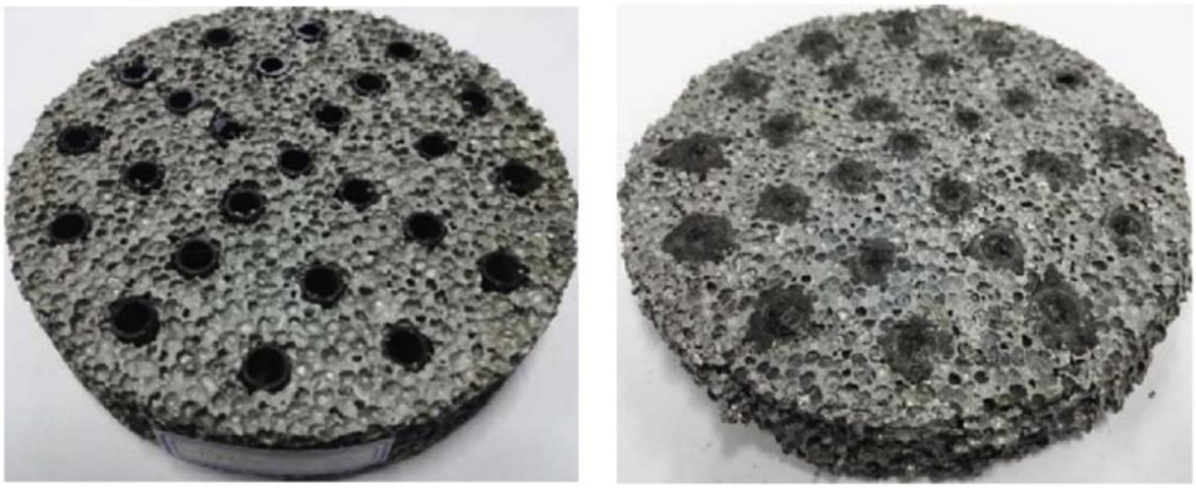


Figure 38 CFRP tube reinforced hierarchical foam metamaterial before (left) and after crushing (right). From [114]

The effectiveness of CFRP tubes for crash resistance was also demonstrated by Li et al. [115]. Samples were manufactured by pipe-rolling of carbon fibre pre-preg to form tubes that were then bundled together in concentric layers (Figure 39). After each layer a sheet of 0° pre-preg was wound around the bundles to create a nested hierarchy. CFRP tubes effectively dissipate energy by interlaminar crack propagation and shear-bending coupled fracture of fibres. Typically long and thick single-walled CFRP tubes with large diameter fail abruptly with poor energy absorption compared to short thin-walled tubes, which are highly efficient absorbers. Hierarchical structure was able to stabilise the global crushing behaviour of long and large diameter tubes, resulting in improved energy absorption. Quasi-static compression testing was carried out on a variety of tube diameters and geometry configurations. Crushing behaviour was considerably more progressive than equivalent single-walled tubes of the same

dimension and weight, improving specific energy absorption by as much as 20 times. Mean crushing force exceeded half of the peak force and the deformation under compression exceeded 50%. Tubes absorbed energy through progressive axial splitting and curling initiating from the location of the applied load. As the outer tubes began to split the inner tubes kept absorbing energy. Best performing samples were able to withstand peak crushing force in excess of 700kN before failure initiation, while absorbing over 40kJ of energy before densification.

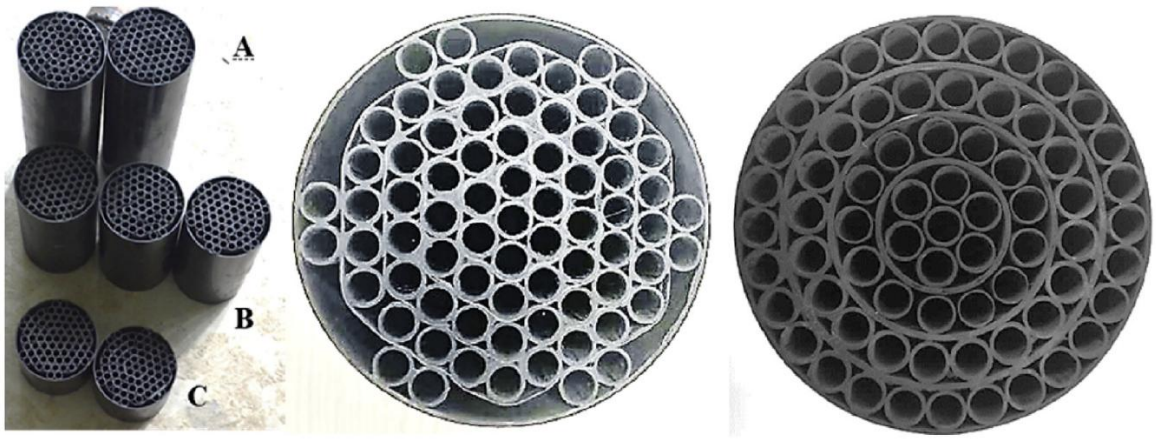


Figure 39 Hierarchical CFRP tube structures. From [115]

Most crash resistant lattices dissipate energy by crushing their structure from a large volume to a small volume. This means the maximum capacity for energy absorption is limited by their initial starting volume. Alan et al. developed smart aluminium honeycombs for rapid expansion on impact as a deployable crash safety device, increasing their volume by up to 75 times when a crash sensor is triggered [116]. Specimens were prepared from aluminium honeycomb before expansion (HOBE), an intermediate form in the process of manufacturing standard hexagonal honeycomb. Rapid expansion of the honeycomb was achieved within 16ms with no loss of integrity of the structure. The system was able to activate up to 5kJ of energy absorption with only 1.3kJ of energy input required for deployment, two orders of magnitude less than that required to inflate a passenger airbag.

Auxetic lattice structures possess negative Poisson's ratio, meaning they compress rather than expand in the lateral directions under uni-axial compression. This can make them highly efficient energy absorbers. A novel tubular lattice was recently introduced by Jiang et al. featuring sinusoidal ligaments able to switch from bending to stretching dominated mechanical responses while exhibiting auxetic behaviour and tuneable Poisson's ratio (Figure 40) [117]. Samples were manufactured using polyjet printing of flexible FLX95595-DM material. Uniaxial tensile testing revealed a J-shaped stress-strain response and strain-dependent auxetic behaviour. The sinusoidal ligaments were loaded axially, inducing a

stretch-bending deformation. As the ligaments deformed, the loading was transferred to circumferential ligaments and the response became bending-dominated. Further strain caused the ligaments to straighten entirely, resulting in a transition to purely stretch-dominated behaviour. By arranging the lattice heterogeneously it could be made to deform into a target shape under load. Poisson's ratio could be adjusted between -0.7 to 0 by tailoring the geometry parameters of the sinusoidal ligaments, such as beam thickness and depth to delay the transition from bending to stretch dominated response. Further multi-parameter inverse optimisation analysis is required to design for specific desired auxetic responses for practical application.

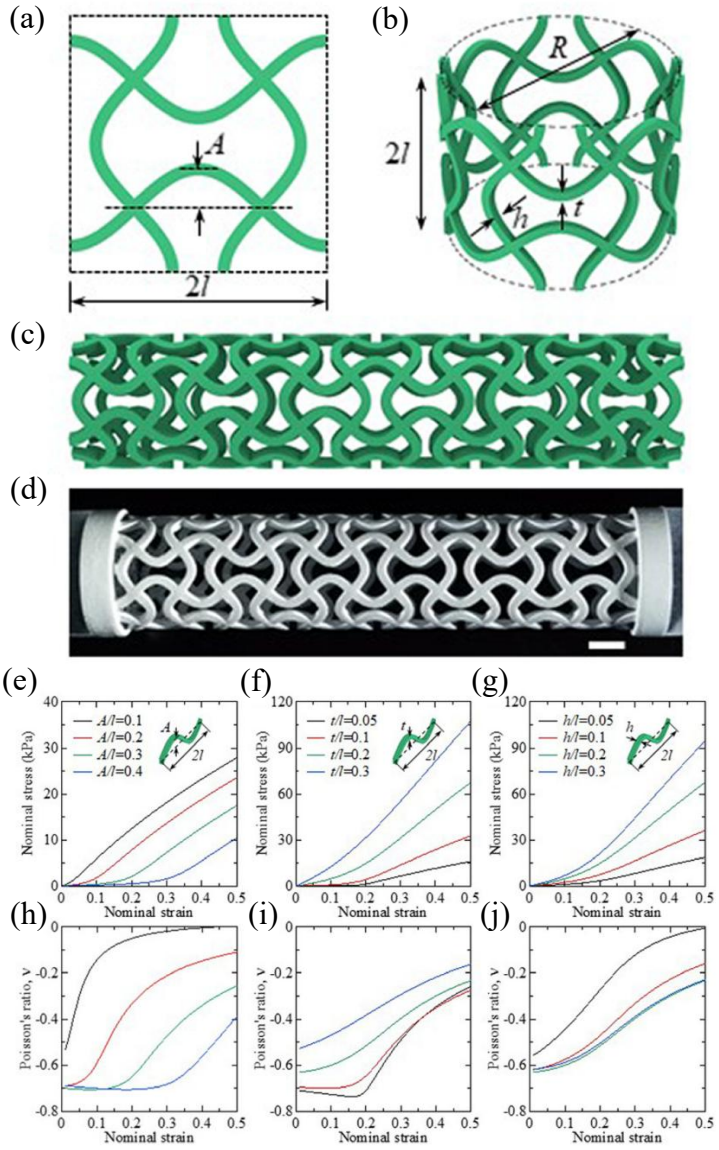


Figure 40 Tubular auxetic metamaterial design; (a) 2D planar unit cell design; (b) 3D RVE of four unit cells arranged along the circumference: (c) tubular lattice formed from 6 RVEs; (d) 3D-printed sample; (e) and (h) Effect of wave amplitude A/l on stress and Poisson's ratio, respectively. Here, t/l=h/l=0.1; (f) and (i) Effect of varying beam thickness t/l with A/l=0.3 and h/l=0.2; (g) and (j) Effect of varying beam depth h/l with A/l=0.3 and t/l=0.2. From [117]

Topology optimisation can be applied at the material level not only to improve structural efficiency of cellular materials, but also to target improvements in energy absorption capacity. Le et al. were able to tailor energy wave propagation through a two-phase composite by topology optimisation to focus, redirect and disperse energy [118]. At each point in the material the parameters of the laminate were assigned and the material properties were calculated from homogenisation theory. These properties were then used in transient finite element analysis to simulate the wave propagations through the material. Cost and constraint functions were optimised using gradient descent based algorithms to find the optimal laminate orientations and volume fractions to assign either stiff or compliant phases at each location. The process was able to funnel an impact load, uniformly distributed over a 266.67mm wide design space to a small concentrated region 6.67x16.67mm in size at the bottom centre of the material. Energy was also successfully redirected within a 600x100mm design domain from an impact occurring over 100mm of the width of the top centre surface to very specific locations: a 15x5mm region centred 42.5mm left of the bottom right corner; a 5x5mm region centred 42.5mm above the bottom right corner; a 5x10mm region centred 150mm left of the top right corner; and a 5x5mm region centred 47.5mm above and 102.5mm left of the bottom right corner. This was achieved by altering the objective functions to maximise normalised total energy over the target regions. A largely compliant region was assigned at the impact point to allow energy into the material. It was then channelled by stiff phases to direct energy towards the desired region. Such designs can be used to divert energy from vital areas and eliminate problematic stress concentrations. By setting the objective functions to minimise energy concentrations, the impact could be prevented from propagating to the bottom of the design space and was instead made to disperse gradually as it travelled through the material. In this case stiff microstructures were assigned to the top and bottom regions to either trap energy near the top surface or reflect it back from the bottom surface. The study was purely based on numerical FE analysis however and lacked experimental validation of results. Nakshatrala and Tortorelli presented a topology optimisation framework for management of dynamic energy propagation in rate-independent elastoplastic materials [119]. A discrete adjoint sensitivity analysis was combined with an optimality criteria based approach to update the design variables to create structures that effectively dissipate energy from impacts. Two material phases, one stiff and the other compliant were arranged within a 2-dimensional 600x100mm design space under plane strain conditions to create a binary composite. A triangular pulse of peak load 250MPa was applied over 10μs to the top face of the material. The model was able to efficiently minimise the energy in five discrete target regions located throughout the design space. Large concentrations of stiff material were assigned to the top surface to lower the input energy. Stiff regions were also assigned above and around the target

areas to divert energy around the key regions as it travelled through the material. The framework was extended to 3-dimensions in a 800x800x100mm design space, however the material was not manufactured and results were from numerical simulation only. Asadpoure et al. optimised multiphase architected materials consisting of combinations of stiff material, dissipative material and void space for enhanced vibration damping at highest stiffness and lowest density [120]. Effective stiffness was evaluated using homogenisation theory and the Bloch-Floquet theorem was used to estimate damping capacity. Each element in the 2-dimensional design space was initially assigned a mixture of all three phases. Interpolations and penalisation was then used to converge on near binary composite microstructures by deriving design sensitivities for gradient-based optimisers. The microstructures were successfully optimised to achieve superior figure of merit values for combined stiffness, density and effective damping for target frequencies between 1-200kHz. The best results were achieved for high target frequencies at low target densities. Performance was considerably improved by the inclusion of void spaces. The figures of merit decreased significantly when the study was limited to fully dense two-phase microstructures. Carstensen et al. designed the unit cell of a 2-dimensional periodic cellular lattice using topology optimisation to maximise energy absorption capacity [121]. A method of moving asymptotes gradient based optimiser along with SIMP penalisation was used to determine if each element of the design space should contain a void or solid material. Heaviside projection was used to limit the minimum feature size to ensure manufacturability from bulk metallic glass (BMC) base material. On considering plastic material behaviour by Von Mises yield function with linear isotropic hardening the energy dissipation improved by 10% compared to a fully elastic model. Unit cells were generated from square design domains of side length 1mm at volume fractions between 10-25%, with square symmetry conditions applied. The stress-strain responses of lattices of 6x6 unit cells were simulated by FE analysis. The energy absorption was almost 10 times greater than hexagonal honeycomb at the same volume fraction, as shown in Figure 41. Under real-world uniaxial in-plane compression testing with quasi-static loading rate, the optimised lattice exhibited energy absorption approximately 38% greater than the honeycomb equivalent. The new design buckled layer-by-layer, absorbing energy progressively as each layer collapsed, whereas the honeycomb failed by formation of shear bands.

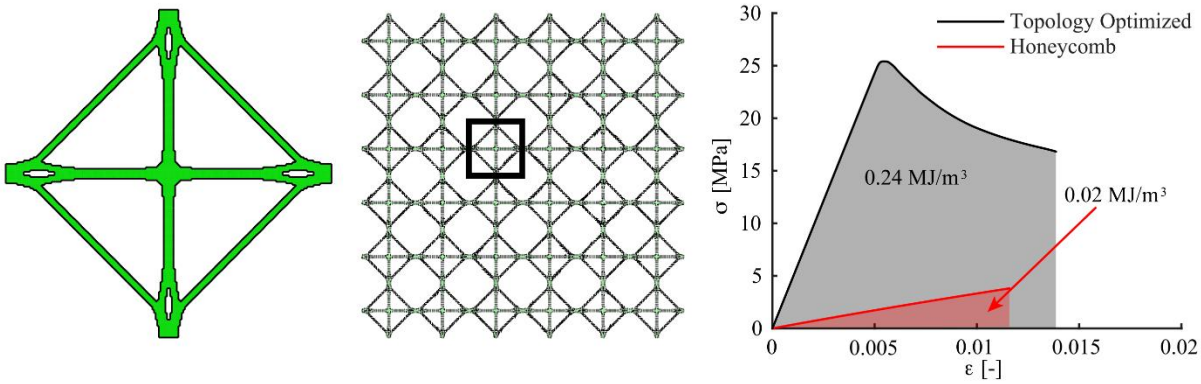


Figure 41 Unit cell optimised for maximum energy absorption and volume fraction of 12.5% (left), 6x6 lattice sample (centre) and stress-strain response with energy absorption comparison (right). From [121]

In uniform lattices energy absorption is often impaired by the formation of shear bands within the material. Maskery et al. eliminated formation of shear bands through the use of graded density regions, making the deformation mechanism more stable and predictable, which is advantageous for design of crash-safety structures [122]. Relative density was set at six levels of a body-centred-cubic lattice by altering the strut thicknesses, getting progressively more dense towards the base of the structure (Figure 42). The resulting functionally graded material was fabricated by selective laser melting of Al-Si10-Mg and subjected to quasi-static compression. Energy absorption prior to densification was increased by 10% from 5.7 ± 0.2MJ/m^3 to 6.3 ± 0.2MJ/m^3 compared to a uniform lattice of the same average density. This exceeds aluminium foams with similar volume fraction, which have achieved energy absorption between approximately 3.1 – 4.5MJ/m^3. The same thickness was used for each of the graded layers and changes at the layer boundaries were abrupt. In further study these thicknesses could be allowed to vary. Generative geometries are likely to be able to blend between layers with smooth transitions. Alternating density bands may be used to channel crack propagation to bypass key areas and eliminate undesirable failure modes. The study highlighted the limitations of Gibson-Ashby based analysis for predicting densification strain and cumulative energy absorption of graded density structures, suggesting further empirical models are required for functionally graded energy absorbers. Machine learning approaches, able to optimise designs based on datasets gathered by numerical FE analysis, are therefore highly desirable for designing a new generation of lightweight impact structures.

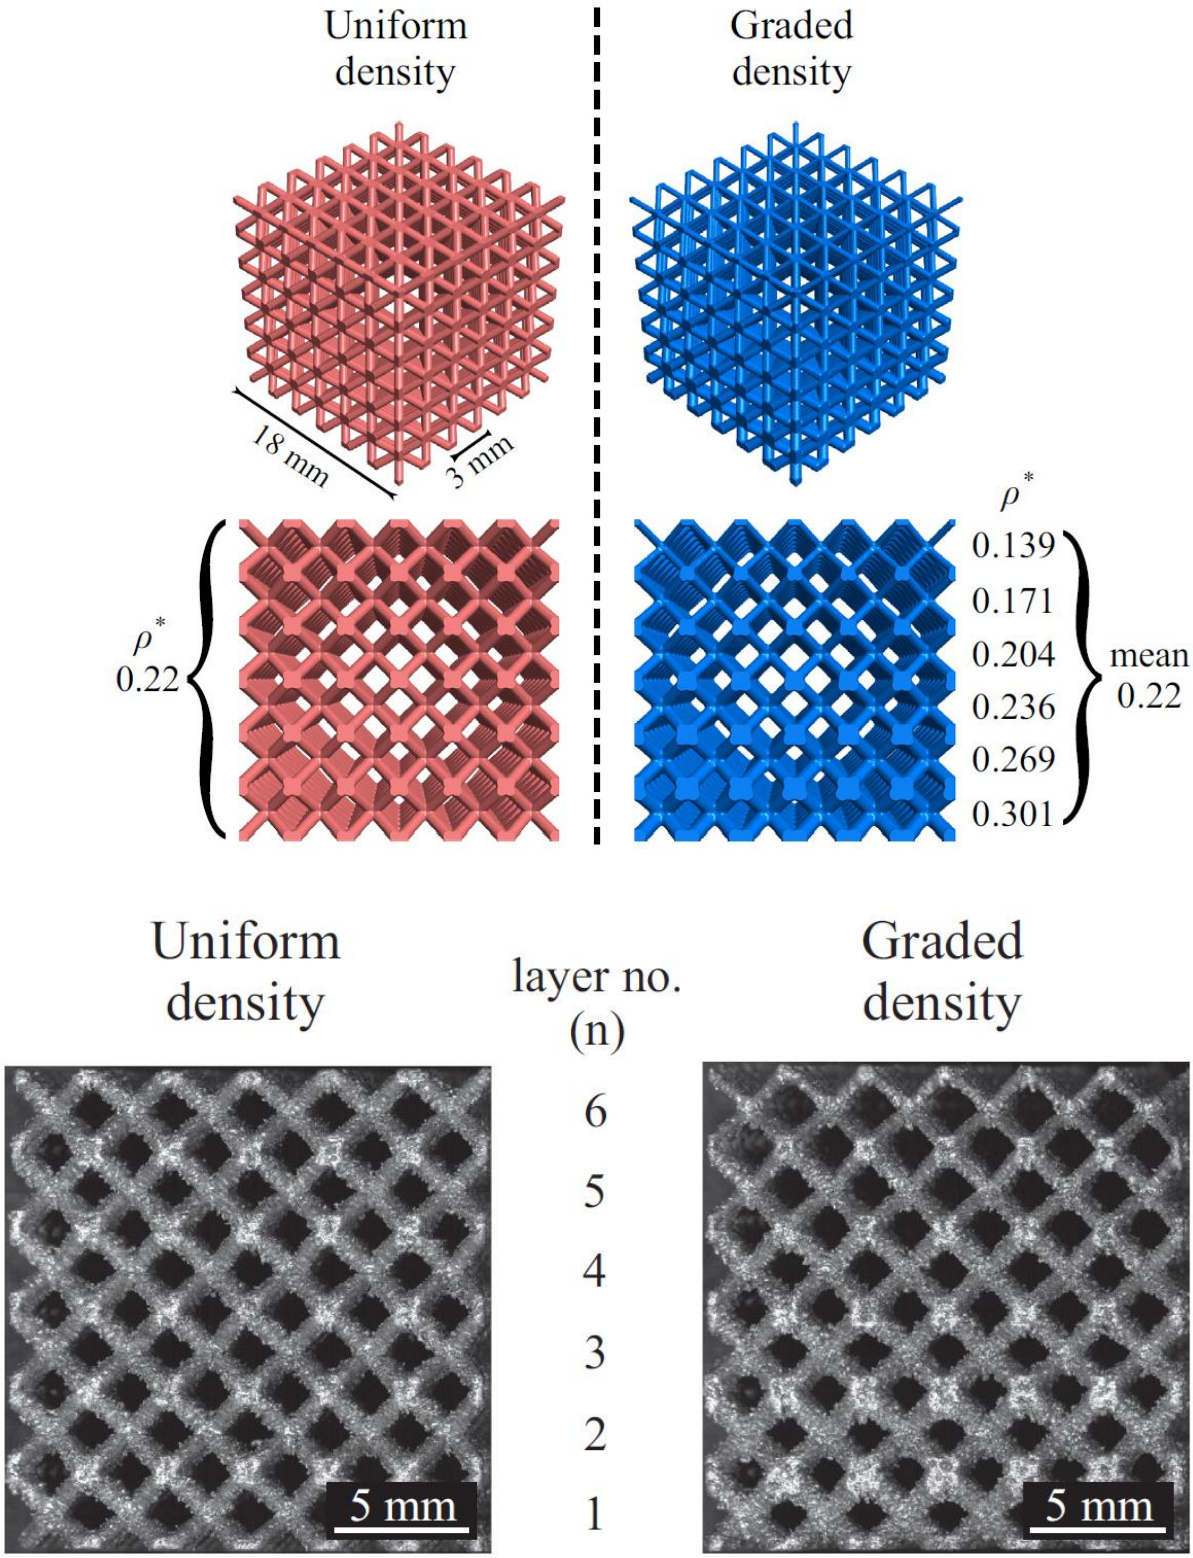


Figure 42 Uniform and graded density lattices. From [122]

2.5.2 Sound Absorption

There is increasing demand for noise reduction and refinement in passenger vehicles. Excessive noise, vibration and harshness (NVH) can cause significant discomfort, headache, fatigue and stress. Prolonged exposure can cause long term health effects such as hearing damage and hypertension. The human ear can detect sound in the frequency range of 20Hz to 20kHz. The main sources of vehicle cabin noise are road inputs, exhaust and engine vibrations which cause propagation of high amplitude sound waves at low frequencies. Low-frequency sound in the range of 0-2000Hz is particularly effective at propagating over large distances and through solid material with little energy loss. Sound absorption effects typically occur over a finite frequency range. Industrial sound absorbers are therefore usually a combination of multiple devices, each designed to cover a certain bandwidth. This generally makes broadband attenuation solutions large and heavy and impractical for use in aircraft and motor vehicles. The main parameters used to evaluate the performance of sound absorbers are coefficient of absorption and sound transmission loss (STL). The coefficient of absorption (α) is the ratio of sound energy absorbed by the material to the incident energy and takes value between 0 and 1, where 1 represents total absorption. Transmission loss is measured in decibels and is defined as the ratio of sound power incident to the sound power transmitted:

$$TL = 10\log_{10}(W_{incident}/W_{transmitted}) \tag{2.22}$$

Traditional sound absorption materials include glass fibrous material, mineral wool and porous foams, which have inherently weak attenuation at low frequencies. Classical absorption materials generally require a minimum thickness of at least one-quarter of the operating wavelength to achieve full absorption [123]. With this assumption, to absorb sound waves travelling at 343m/s in the range of 500-1000Hz would require minimum thicknesses between at least 85.75-171.5mm, many materials would require even greater thicknesses. A significant requirement exists for metamaterials that are able to provide ultra-low frequency absorption at sub-wavelength thickness.

The Helmholtz resonator (HR) was invented in 1860, originally as a method for determining the frequencies present in musical notes [124]. It consists of an enclosed resonance cavity of volume V with a small neck opening of length L and cross section area S. The volume of air in the chamber resonates at the frequency given by

$$f = \frac{c}{2\pi}\sqrt{\frac{S}{LV}} \tag{2.23}$$

where C is the speed of sound. Air flow velocity increases at the cavity opening when resonance occurs. This causes the frictional force due to viscosity to increase, facilitating the dissipation of the sound wave energy at the resonant frequency. Kasliwal et al. observed poor noise attenuation at low frequencies in the standard exhausts of LMD buses and designed an absorption system based on a Helmholtz resonator device [125]. Two cylindrical HRs of neck diameter 10 and 16mm and length 109 and 69.7mm were connected to the tailpipe to target frequencies of 80Hz and 160Hz. This achieved noise reductions in both passenger ear level (PEL) and pass by noise (PBN) testing without any significant drop in exhaust system back pressure. PEL tests showed maximum attenuation of 3.3dB at 2600RPM and 2dB at idle. At the tailpipe orifice 7.5dB reduction was recorded at 2600RPM and 4.9dB at idle. However due to the size of the resonators there were limitations on the number of frequencies that could be targeted at once. HRs are able to provide absorption coefficients exceeding 90% at ultra-low frequencies, however they are only effective over narrow frequency bands and require large cavity volumes that are not suitable for applications with spatial limitations. A strategy that has been used to reduce the volume requirements is the introduction of a neck that extends into the resonator cavity. Multiple HRs can then be combined to improve absorption range by superposition. The target frequencies can be tuned by varying the lengths of the necks without changing the back cavity depth. Gorain and Padmanabhan constructed a unit cell based metamaterial featuring an array of Helmholtz resonators with intruded necks embedded within polyurethane (PU) foam [126]. A hollow cylinder of inner diameter 34mm and height 18mm with one end open and the other end with a central hole of 4mm diameter offered a natural frequency of 415Hz. By intrusion of hollow tubular necks of varying length, the resonant frequency was reduced and the absorption could be tuned to frequencies between 210 to 415Hz. The HRs at the lowest (210Hz) and highest (415Hz) ends of the frequency range were kept constant, then intermediate HRs were iterated by trial and error to maximise the absorption over the frequency range. An array of 10 HRs was needed to maintain absorption coefficient between at least 0.4-0.5 over the range of 210 to 415Hz. Guo and Zhang studied checkerboard arrangements of alternating Helmholtz resonators with varied length extended necks (HREN), displayed in Figure 43 [127]. They found that the absorption coupling between resonators becomes weak if two adjacent resonators are dissimilar. Combination of HRs with peaks at 777.8Hz and 987.0Hz resulted in a dual-band absorption response at each of the target frequencies. Once the frequencies were closer together (839.8Hz and 909.4Hz) the coupling effect strengthened and lead to wide-bandwidth absorption. Absorption peaks greater than 90% were maintained over a range of 71.5Hz from 847.2Hz to 918.7Hz, 1.63 and 1.47 times wider than the respective individual HRENs.

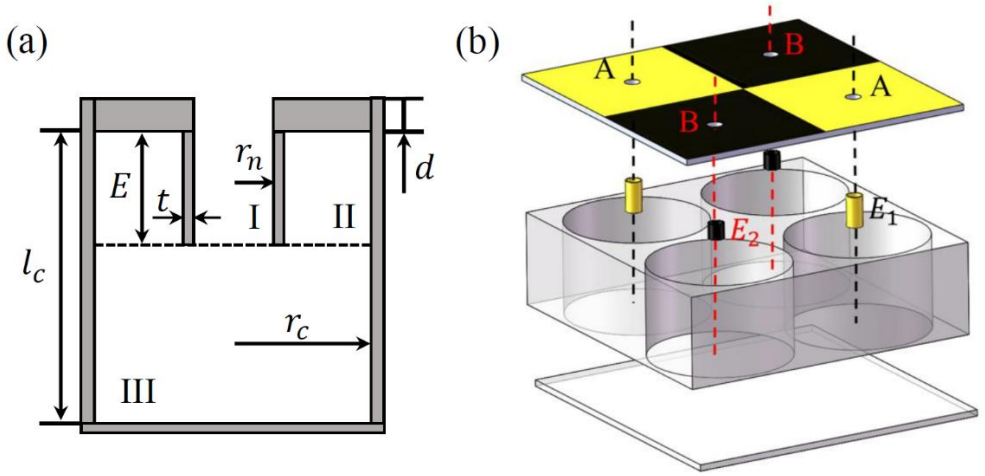


Figure 43 (a) Schematic of Helmholtz resonator with extended neck; (b) checkerboard arrangement of alternating resonators A and B with different length necks. From [127]

Baek et al. were able to reduce the size of a HR by 40% while also improving transmission loss by combining the design with a fractal-based internal structure (Figure 44), which achieves broadband sound insulation at low thickness through a channel extension effect [128].

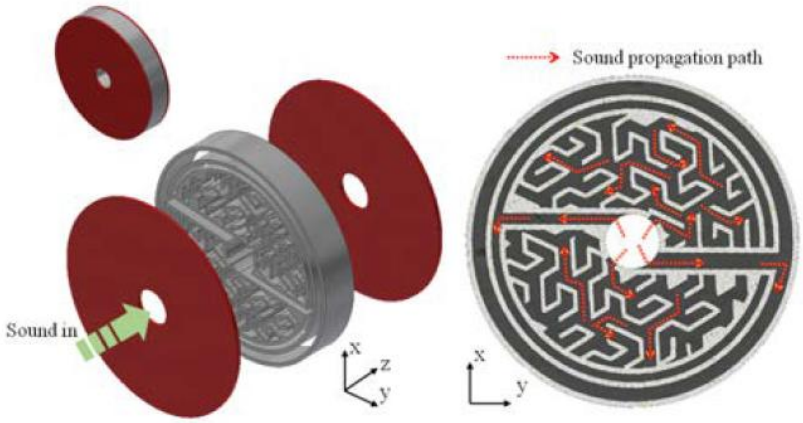
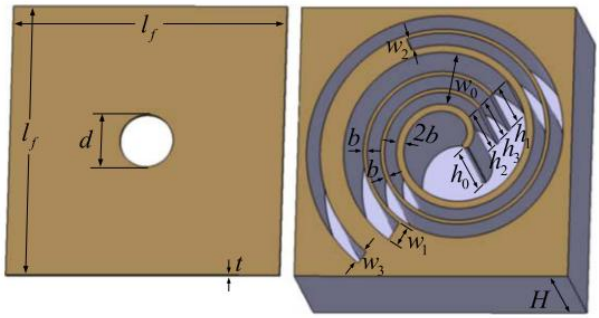


Figure 44 Hybrid HR and fractal structure metamaterial design. From [128]

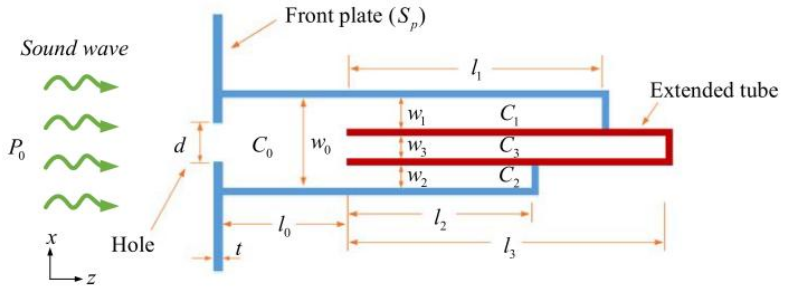
Mean sound transmission loss of 16.34dB was attained over the range of 600 to 1700Hz by hybridisation of an arrangement of 3 concentric HRs tuned for 600-1000Hz with a space-filling fractal structure designed for high absorption in the 1000-1700Hz range. A fundamental problem exists in designs featuring superposition of multiple absorbers. Analytical methods are able to find geometry parameters for a given target frequency, but knowing which frequencies to target to obtain an optimal combined response that provides maximum absorption across the greatest range at the lowest possible thickness is a very complex multi-variable design problem, for which there are limited approaches to solve that usually rely on a degree of intuition and trial and error.

An effective method for reducing the thickness requirements of resonator absorbers is by spiralling of the back cavity. This elongates the path of the sound waves through the resonator, shifting to lower resonant frequencies while limiting the increase of overall thickness. HRs with spiral extended necks were designed by Guo et al. who achieved total absorption at 320Hz with a thickness of only 13.5mm, approximately $1/80^{th}$ of the wavelength [129]. However broadband range was narrow, reaching absorption coefficient above 50% over only approximately 50Hz of range. A checkerboard arrangement of 4 HRs with spiral necks was able to improve performance, exhibiting average absorption of 90% over 360-410Hz at the same thickness. Cai et al. created a metamaterial featuring a coplanar arrangement of spiral tubes, the inner sections functioning as Helmholtz resonators and opening sections serving as the neck, which achieved 90% absorption at 400Hz at an overall thickness of 17mm [130]. On simulating the air flow velocity distribution at resonant state it was discovered that the opening of the tube dominated the damping behaviour and the rest of the tube functions as the back cavity tank of a HR. Replacing the spiral tube with an enlarged hollow chamber in all regions except the opening section shifted the absorption peak lower to 250Hz without changing the sample thickness. However again broadband range was considerably limited, reaching 50% absorption for less than approximately 50Hz range. This is a recurring problem with HR and spiral cavity based structures.

Han, Ji and Qiu were able to achieve broadband sound absorption greater than 70% across a bandwidth of 600Hz, between approximately 600-1200Hz, from a 64x64mm square metamaterial of 50mm thickness [131]. Samples consisted of a checkerboard arrangement of unit cells, each cell containing multiple coiled coplanar absorbers (Figure 45). Each absorber consisted of a perforated plated and coiled air chamber with an extended tube inserted at the end. The resulting structure divides the airflow into 3 separate coiled tubes, allowing for 3 absorption peaks at different frequencies that can be tuned by adjustment of the geometry parameters of each tube. By combining unit cells with different target frequencies in parallel, the absorption responses were superimposed to create the characteristic shown in Figure 46. They noted shifts in absorption peaks that deviate from the individual target frequencies as a result of the coupling between absorbers arranged in parallel. Due to the mutual couplings a trade-off exists between achieving strong absorption or large bandwidth.



(a) The detailed internal structure of the individual absorber of the absorbing panel.



(b) The schematic of a straight chamber with an extended tube inserted in it in theoretical analysis.

Figure 45 3D view of the individual absorber and its theoretical model. From [131]

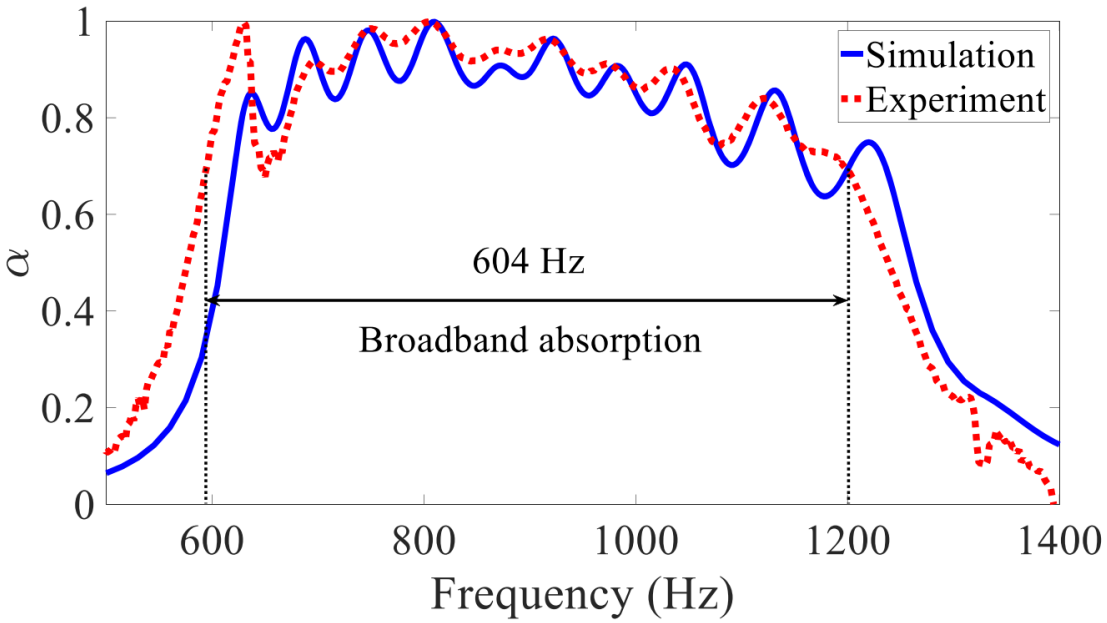


Figure 46 Absorption response of parallel coiled coplanar absorbers. From [131]

In 1974, Maa Dah-You proposed the revolutionary idea of reducing the scale of perforated panel sound absorbers to the submillimetre level to tune acoustic resistance [132]. Resulting microperforated panel (MPP) sound absorbers are able to effectively dissipate energy of incident sound waves by the effect of viscous boundary layer interactions that occur due to the reduced perforation size. A typical MPP-based metamaterial consists of a perforated panel placed in the path of the incoming wave before a contained back cavity air chamber. The micro-scale perforations offer greater viscous losses, resulting in wider absorption bandwidth compared to HR and spiral structures, up to 3 or 4 octaves and excellent performance at low frequencies. The absorption characteristics are highly predictable and can be related directly to the perforation diameter d, perforation ratio σ, panel thickness t and back cavity depth D. However the back cavity depth requirements can make it challenging to combine multiple absorbers at once while maintaining practical thicknesses, especially for automotive and aeronautical applications.

MPP arrangements can be combined either in series or in parallel. Tan and Ripin demonstrated a wide bandwidth absorption response using a double-layered series configuration [133]. A simulated annealing optimisation solver was used with the objective function of maximising the mean sound absorption coefficient over the target frequency range of 200-1000Hz. 50% absorption was realised over a range of approximately 1000Hz from 300 to 1300Hz, however the overall thickness was 60mm due to stacking of the back cavities. Sakagami et al. calculated the theoretical absorption coefficient responses from combinations of two different MPPs combined in parallel [134]. On allowing the back cavity depths of the two MPPs to vary ($D_1 = 50mm, D_2 = 150mm$) and keeping all other parameters constant ($d_1 = d_2 = 0.4mm, \sigma_1 = \sigma_2 = 0.75\%, t_1 = t_2 = 0.4mm$), the analytical solution predicted 50% absorption ranging between 250-1000Hz, with a peak of approximately 80%. The peak occurred part way between the peaks of the individual MPPs and at reduced amplitude. Li et al. studied combinations of four perforated panel absorbers (PPA) arranged in parallel [135]. A drawback of the MPP is that often to achieve wide bandwidth absorption at low frequency an extreme number of perforations, often in the hundreds of thousands per square metre may be required. To allow for ease of manufacture the perforation diameters were limited to above 1.5mm. On comparison between a PPA with perforations of diameter 1.5mm and a submillimetre MPP of perforation diameter 0.7mm with the same ratio of $t/d = 1.3$, the PPA exhibited higher absorption peak which shifted to lower frequency due to increase of acoustic mass with increased aperture. However the submillimetre panel offered greater bandwidth, hence the requirement for parallel combination at perforation diameters above 1mm. Quadruple parallel PPA arrangements optimised by simulated annealing of the geometry

parameters improved the bandwidth of a single MPP designed for the 120-250Hz range, increasing the 70% absorption threshold by approximately 55Hz at the low-end and 45Hz at the high-end, with the same depth restriction of $D = 100mm$. Multiple designs were successfully produced to offer absorption bandwidths of approximately 200Hz within the 100-900Hz regime for a variety of back cavity depths that ranged between 35-100mm, as shown in Figure 47.

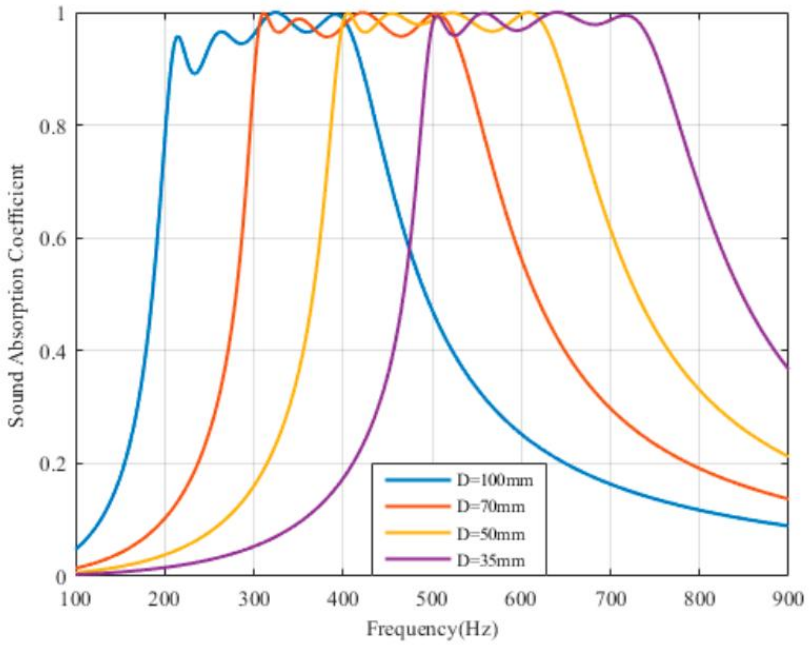


Figure 47 Optimisation of quadruple parallel arranged PPAs with different back cavity depths. From [135]

Choice of MPP configuration depends greatly on the application. If large thicknesses are acceptable series arrangement is likely to give stronger and more predictable absorption performance across a wider bandwidth. Parallel arrangements can effectively limit the overall thickness, though the absorption peaks in the superposed response are often reduced. Mahesh and Mini combined multiple Helmholtz resonators with inserted necks (HRIN) with integrated MPPs, studying both series and parallel arrangements using analytical modelling and FEA numerical simulations [136]. Arranging a single MPP in series in front of four parallel HRINS resulted in 55% absorption over the frequency regime of 263-403Hz, at an overall thickness of 110mm. On placing the MPP and HRIN array in parallel, thickness was reduced to 56mm and the proposed structure offered 65% absorption at a higher frequency range, from 318-800Hz, though these results were not validated through real-world testing.

In recent years, studies have begun to explore the potential for machine learning approaches in managing sound attenuation. Huang et al. predicted the sound quality of vehicle interiors using convolutional neural networks [137]. A training dataset was gathered from subjective

evaluations of vehicle noise over a variety of road surfaces. These sound quality ratings were then mapped to the time and frequency domain noise characteristics. The model was successfully able to classify the recordings of vehicle interior noise, rating them by how severely they affect the human participants to an accuracy of 97.03% and a loss of 0.167. The CNN approach was able to mathematically predict with near certainty an abstract and subjective feeling of how badly a human is annoyed by particular noises. Iannace et al. applied machine learning to model the absorption coefficient responses of broom fibres under impedance testing [138]. An artificial neural network with a single hidden layer was used with inputs of the material thickness, diameter, resistivity, porosity and the frequency of interest. These were mapped to a single output containing the absorption coefficient at the corresponding frequency. A dataset of 8172 observations (70% used for training and 30% for validation) was able to make predictions with a Pearson's correlation coefficient of 0.989, considerably better than a generalised linear model, which only provided correlation of 0.571 under testing on known data. While the model showed excellent predictive power, the ability to make inverse optimisation, outputting the material parameters required to maximise absorption over a given frequency range, is limited by having the frequency as one of the inputs. Jeon et al. used an artificial neural network to predict the sound absorption characteristics of multi-layered fibrous materials [139]. Samples consisted of four layers of fibrous polyester (PET), a material commonly used for mid to high range absorption. Samples provided at least approximately 80% absorption across the 3000-6000Hz range. Densities and airflow resistivities (AFR) were varied between layers. 256 different combinations were used to generate a training dataset via impedance testing. The inputs to the ANN model were the thicknesses, densities and AFRs at each of the layers. The output was the absorption coefficient for 56 frequencies at 100Hz intervals in the range of 500 to 6000Hz. The model provided predictions with error less than 1.67% and demonstrated that if the ANN is trained on experimental data, it can be a more effective method for accurately estimating the actual real-world absorption response than the analytical transfer matrix solution, which gave average error of 2.46%. The downside being the time consuming process of dataset generation. The model showed accurate predictions that can be used in the design of high-frequency sound absorbers. A research gap remains for ultra-low frequency absorption (at less than 1000Hz) at lowest possible thickness. Many studies only work in one direction, taking material parameters as input and outputting predictions of the absorption response. A network able to work in the inverse direction, in which a 2-dimensional absorption coefficient – frequency response is input and the material parameters required to obtain the response are output, is likely to be able to make powerful optimisations for design of sound absorption metamaterials. However as this increases the dimensionality of the network inputs, a deeper

network and convolution is likely to be required to extract features of the required level of complexity.

Sun et al. applied an inverse optimisation method using a deep learning model trained on the analytical absorption responses of a two-order Helmholtz resonator [140]. The inverse approach allows for on-demand design optimisations in which specific frequencies can be targeted. The exact desired absorption-frequency characteristic is input and the network returns the geometry parameters of the metamaterial required. Mahesh et al. designed a deep neural network for inverse design of the geometry of an acoustic absorber for low-frequency absorption, based on 4 HRs and a single MPP arranged in parallel [141]. Frequency responses were obtained analytically using electro-acoustic theory to create a dataset for training. The geometry parameters predicted by the DNN to achieve given responses were fed back into the analytical model. Desired input and actual achieved absorption characteristics showed good agreement with low average error below 3%. The parallel MPP-HR combination optimised for the target frequency range of 340-630Hz exhibited excellent broadband characteristics, reaching quasi-perfect absorption of more than 90% across the range of 347-630Hz at a subwavelength total thickness of 70mm. The study used a data-driven approach as opposed to evolutionary, randomised approaches such as genetic algorithm and particle swarm optimisation. These methods are highly computationally expensive and time consuming due to the large number of iterations required and often do not converge on an optimal solution. On comparison with an alternative GA solution, the deep learning method provided superior accuracy. The maximum error in the geometry parameters predicted to optimise the absorber in the 0-700Hz frequency range was 3.1%, whereas the GA provided error of 12%. The strategy can be extended for many acoustic structures, anything for which a dataset can be generated using theoretical or experimental means. In this case no experimental validations were carried out, further real-world testing of resulting metamaterial designs is required to evaluate potential for practical industry applications. With this method designs are still restricted by selection of base geometry configuration. Many parameters had to be kept constant to prevent the problem from becoming too computationally expensive, with complexity beyond the scope of the neural network. The parameters allowed to vary have a limited range and have to be carefully preselected. Previous studies only focus on sound absorption. Designs could be extended to feature structural efficiency optimisations for enhanced lightness and strength to allow for multifunctional load bearing applications. If the sound absorption functionality can be combined into the same metamaterial that forms the structure of a component, low-frequency sound attenuation can be achieved without adding any additional material thickness.

2.6 Conclusion

The review of current literature has revealed that a vast range of potential applications of machine learning exist for design of architected materials. The variety of metamaterial base configurations that have been developed offer a broad range of possible functionalities. As such in order to evaluate the effectiveness of the machine learning based optimisation approach, the following research is presented across 5 different case studies which have been selected with a focus on addressing several key research gaps identified in existing literature. A major goal is to enhance the tunability of the metamaterial that allows for fine tuning and tailoring of performance criteria for bespoke real-world applications, using deep learning to create a link between the material level geometry and the component level structural performance. Many previous works using machine learning focus mainly only on the prediction of material properties, inputting the unit cell geometry and outputting the material response. This project also explores the potential for machine learning in inverse design, in which a desired performance characteristic may be supplied to an optimiser which then outputs the geometry required to provide a certain stress/strain curve or energy absorption response that can then be tuned to satisfy design requirements. In many previous works structural optimisations are simplified to 2 dimensions, where an optimiser assigns stiff or flexible material to certain locations within a grid. Fewer studies have attempted to synthesize geometry in 3D. Most current theory for predicting structural properties of truss lattice structures may only be applied to periodic unit cells where all struts share a common radius. Heterogenous and numerically formed lattices typically depend on empirical FEA based optimisations that suffer from computational bottlenecks. A key advantage of deep learning is that it allows for mapping of variables in a way that circumvents the need for theoretical solutions. Any kind of multi-variable design problem can be modelled provided a dataset of examples can be produced for the AI to train on. This allows multiple competing objectives to be optimised simultaneously even if there is no underlying theory to link them physically, allowing for activation of additional multifunctional features during the optimisation. The following chapters demonstrate this by optimising parametric unit cell designs for combinations of multiple properties. An inverse model is first developed to tailor sound absorber metamaterials, as acoustic characteristics can be rapidly predicted by theory. This is followed by an equivalent model for structural properties trained on numerical simulations. Following research adds simultaneous tuning of both structural and acoustic characteristics. Both static and high velocity dynamic structural responses are then optimised at the same time. To showcase the versatility of the solvers, the network architectures are applied to optimise the layup of composite materials for minimal stress design using transfer learning.

Chapter 3 Design of Subwavelength Wide Bandwidth Sound Absorbers by Inverse Convolutional Neural Networks

This study focuses on the application of machine learning for optimising metamaterials for enhanced broadband sound absorption at minimal subwavelength thicknesses. The microperforated panel was chosen as the base metamaterial configuration. A workflow is proposed that uses an inverse convolutional neural network model to provide on-demand design optimisations. The desired absorption-frequency characteristic is provided as input for the network, which is trained to return the geometry parameters of the metamaterial required to satisfy the request. The model can then be used to target specific frequencies and maximise absorption over a given frequency range. The MPP was selected as it offers wide absorption bandwidth at low frequencies and low thicknesses can be achieved using multiple panels arranged in parallel. MPPs are simple to manufacture with scalability for mass production and can be machined using inexpensive CNC router hardware. They can be easily combined with additively manufactured lattice structures for multifunctional load bearing applications. The geometry can be described by a discrete set of parameters that are suitable for use as machine learning outputs. Absorption responses are highly predictable and can be determined analytically, allowing for datasets to be generated rapidly. Hundreds of thousands of examples can be created within seconds. The problem has similar dimensionality and complexity to mapping of non-linear stress-strain responses to geometry parameters of a lattice unit cell. This allows for network architecture, hyperparameters and dataset requirements to be analysed before application to structural efficiency optimisations that rely on FEA or experimentation, which may take weeks to generate a dataset. The number of training samples and the amount of training time required to produce reliable predictions is assessed. The limits of the predictive power of the CNN model and its ability to compensate for different base materials and manufacturing methods are explored. The model is then expanded to predict a more complex arrangement of multiple parallel MPPs for practical application aimed at reducing low-frequency vehicle road noise between 500-1000Hz.

3.1 Dataset Generation

A training dataset of MPP absorption characteristics was generated using an analytical electro-acoustic model. The metric used to evaluate the performance of the MPP was the coefficient of absorption (α), the ratio of sound energy absorbed by the material to the incident energy. The coefficient takes value between 0 and 1, where 1 represents total absorption. For calculation the MPP is considered as an array of short tubular cavities of diameter d and length t (the panel thickness), separated at a distance D from a rigid wall by an enclosed back cavity. The arrangement can be represented by an electro-acoustic equivalent circuit that consists of the acoustic impedances of the panel Z_p and back cavity Z_D combined in series. This is shown in Figure 48.

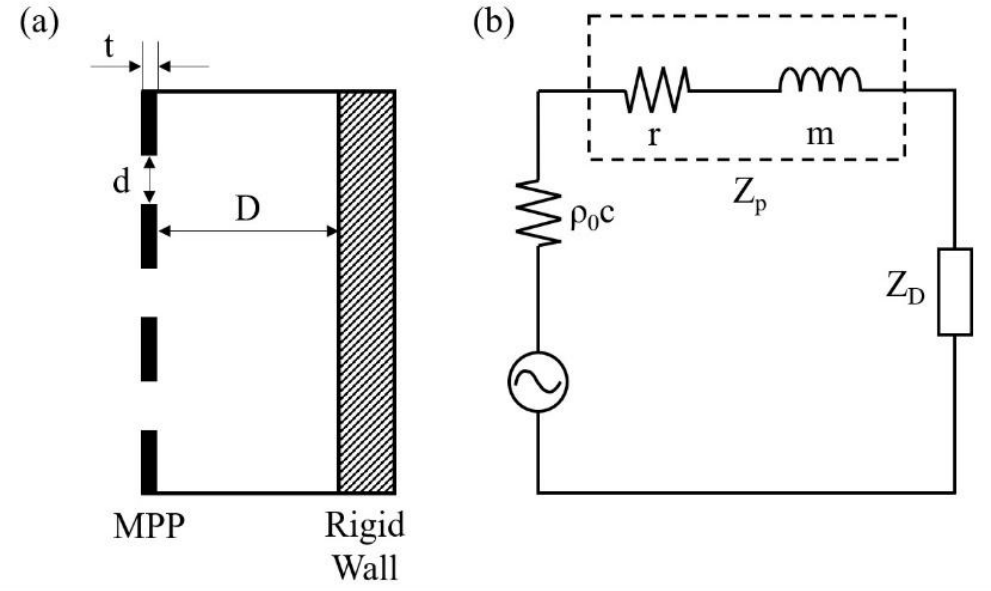


Figure 48 (a) Model of MPP absorber; (b) electro-acoustic equivalent circuit

The global impedance of the system can be found from the summation of the panel and back cavity impedances:

$$Z_{MPP} = Z_p + Z_D \tag{3.1}$$

The propagation of sound waves within the perforations can be modelled by the equation of aerial motion in a short tube [142]:

$$j\omega\rho_0 - \frac{\eta}{r_1}\frac{\partial}{\partial r_1}\left(r_1\frac{\partial}{\partial r_1}u\right) = \frac{\Delta p}{t} \tag{3.2}$$

where Δp is the sound pressure difference between the ends of the perforation, u is the particle velocity, ρ_0 is the air density, η is the air viscosity coefficient, r_1 is the radius vector of cylindrical coordinates within the tube, and $\omega = 2\pi f$, where f is the frequency of interest.

Solving the equation allows for calculation of the specific acoustic impedance of the perforation, which is given by the ratio of Δp to the average velocity over the cross section $\bar{u}$ [143]:

$$Z_i = \frac{\Delta p}{\bar{u}} = j\omega\rho_0 t\left[1 - \frac{2}{k\sqrt{-j}}\frac{J_1(k\sqrt{-j})}{J_0(k\sqrt{-j})}\right]^{-1} \tag{3.3}$$

where $k = d\sqrt{\omega\rho_0/4\eta}$, d is the perforation diameter, J_1 and J_0 are the Bessel functions of first and zeroeth order respectively. For an MPP k takes values between 1 and 10. An approximation that is valid for all k values is given by:

$$Z_i = \frac{32\eta t}{d^2}\left(1 + \frac{k^2}{32}\right)^{1/2} + j\omega\rho_0 t\left(1 + \left(9 + \frac{k^2}{2}\right)^{-1/2}\right) \tag{3.4}$$

The acoustic impedance of the entire panel can then be found by dividing by the perforation ratio σ:

$$Z_p = Z_i/(\sigma\rho_0 c) = r + jx_m = r + j\omega m \tag{3.5}$$

Where the relative acoustic resistance is given by

$$r = \frac{32\eta t}{\sigma\rho_0 c d^2}k_r, \tag{3.6}$$

the resistance coefficient is

$$k_r = \left[1 + \frac{k^2}{32}\right]^{1/2} + \frac{\sqrt{2}}{32}k\frac{d}{t}, \tag{3.7}$$

the mass reactance is given by

$$\omega m = \frac{\omega t}{\sigma c}k_m, \tag{3.8}$$

and the mass reactance coefficient is

$$k_m = 1 + \left[1 + \frac{k^2}{2}\right]^{-1/2} + 0.85\frac{d}{t} \tag{3.9}$$

The absorption coefficient of the total MPP absorber system is a function of the acoustic resistance and mass reactance of the panel and the relative acoustic reactance of the air in the back cavity as follows:

$$\alpha = \frac{4r}{(1+r)^2 + (\omega m - \cot(\omega D/c))^2} \tag{3.10}$$

where $-\cot(\omega D/c)$ is the relative acoustic reactance of air in the back cavity and $Z_D = -j\cot\omega D/c$ [144].

A MATLAB function was created to evaluate the MPP geometry parameters of perforation diameter, panel thickness, perforation ratio and back cavity depth as input and output the analytical absorption coefficient value at frequencies between 0 to 1000Hz, sampling at intervals of every 10Hz. This was then incorporated into a script to generate a uniform random distribution of geometry parameter variations and store the resulting dataset of absorption characteristics within a cell array matrix. Values of the constants were taken as the properties of air at room temperature: viscosity η = 1.79x10^{-5}kgm^{-1}s^{-1}, density ρ_0 = 1.2kgm^{-3}, speed of sound c = 343ms^{-1}. The design space of the dataset was defined by selecting the limits of the ranges within which the parameters were allowed to vary. The panel thickness t could take values between 0.1 and 5mm, while the perforation ratio σ was varied between 0.1 and 15%. The back cavity depths D were chosen to target overall sample thickness competitive with the lowest values found in literature, between 10 and 60mm. Perforation diameter d was varied between 0.1 and 2mm to allow for sub millimetre scale perforations down to the smallest size that can be manufactured by CNC drilling operations for strongest possible absorption, and also larger non-micro perforation configurations to allow tuning for ease of manufacture. The final script was able to generate datasets of hundreds of thousands of randomised MPP absorption responses on demand near instantaneously.

3.2 Convolutional Neural Network Architecture

A custom CNN architecture was developed bespoke for the MPP inverse design problem. A process of trial and error involved training and reiterating the network structure and hyperparameters over several weeks. The architecture was designed in MATLAB using the deep network designer plugin. The network was then trained using the deep learning toolbox to perform convolution operations. A MATLAB script was written to pre-process the dataset into training, validation and testing sets and evaluate the resulting network performance against known data before iteration of the design. The final network structure is shown in Figure 49. It features an input layer that passes sequence data of absorption responses to a series of hidden layers before outputting predictions of the MPP geometry parameters through a regression layer. The number of convolution layers was increased progressively to deepen the network until additional gains in predictive performance became negligible. At 14 and 16 convolution layers regression began to decrease, 12 layers were found to be the optimal depth to extract the features of required complexity without overfitting. The convolution layers were all 1-dimensional with filter size of 3, stride of 1 and no padding. The numbers of convolution filters in each layer were arranged hierarchically, starting small and gradually extracting more and more complexity as the data is passed deeper into the network. The first layer had 8 filters. Subsequent layers increased by 8 more each time, up until 96 filters in the twelfth layer. Each convolution layer used ReLU activation and was followed by a layer-wise normalisation to accelerate the training and reduce sensitivity to network initialisation. Max pooling down sampling with pool size of 3 and stride of 2 was initially applied to all convolution layers, which greatly improved the training time from 110sec/epoch to 62sec/epoch. Over 100 epochs of training this saved 1hr 20min. Removing max pooling down sampling from the first 3 layers then boosted regression values at the cost of 4.6sec/epoch, or 7min 40sec over 100 epochs. The best trade-off between performance and training time came from applying max pooling only on the layers with the largest numbers of filters. The 12 convolution layers were followed by a global average pooling layer, which condenses all of the feature maps into a single vector. This flattened feature vector was then fed into a traditional DNN architecture of 9 dense fully connected hidden layers. Each of the fully connected layers had batch-wise normalisation and ReLU activation. The number of neurons began at 1024 and was reduced by half with each following layer until the penultimate layer which featured 8 neurons. The final fully connected layer was set to the same number of neurons as geometry parameters of the MPP under study, then each of these was linked to a separate activation in the regression output layer. The number of activations and learnable weights and biases at each stage of the network are displayed in Table 2.

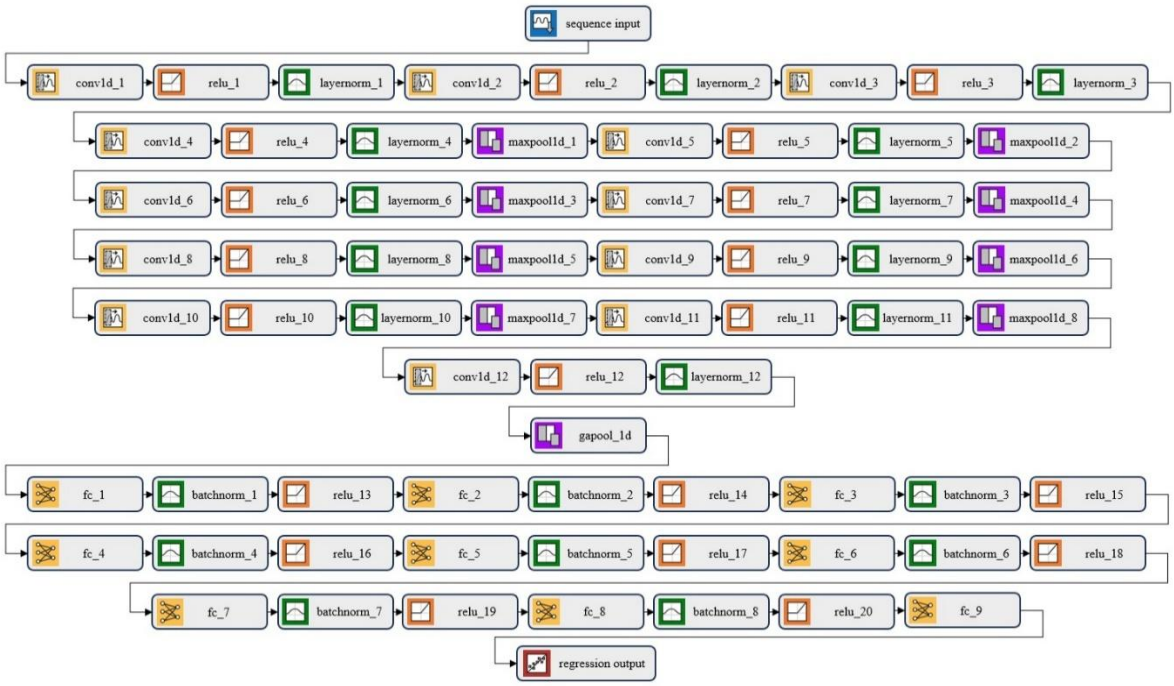


Figure 49 Convolutional neural network architecture. Block colours correspond to specific layer types as follows: blue = sequence input, red = regression output, yellow = convolution and fully connected, orange = ReLU activation, green = normalisation, purple = pooling

Table 2 Activations and learnable parameters at each stage of the convolutional neural network

No.	Layer Type	Activations	Weights	Biases	Total Learnables
1	Sequence Input	1	-	-	0
2	1D Convolution	8	3 x 1 x 8	1 x 8	32
3	1D Convolution	16	3 x 8 x 16	1 x 16	400
4	1D Convolution	24	3 x 16 x 24	1 x 24	1176
5	1D Convolution	32	3 x 24 x 32	1 x 32	2336
6	1D Convolution	40	3 x 32 x 40	1 x 40	3880
7	1D Convolution	48	3 x 40 x 48	1 x 48	5808
8	1D Convolution	56	3 x 48 x 56	1 x 56	8120
9	1D Convolution	64	3 x 56 x 64	1 x 64	10816
10	1D Convolution	72	3 x 64 x 72	1 x 72	13896
11	1D Convolution	80	3 x 72 x 80	1 x 80	17360
12	1D Convolution	88	3 x 80 x 88	1 x 88	21208
13	1D Convolution	96	3 x 88 x 96	1 x 96	25440
14	Fully Connected	1024	1024 x 96	1024 x 1	99328
15	Fully Connected	512	512 x 1024	512 x 1	524800
16	Fully Connected	256	256 x 512	256 x 1	131328
17	Fully Connected	128	128 x 256	128 x 1	32896
18	Fully Connected	64	64 x 128	64 x 1	8256
19	Fully Connected	32	32 x 64	32 x 1	2080
20	Fully Connected	16	16 x 32	16 x 1	528
21	Fully Connected	8	8 x 16	8 x 1	136
22	Fully Connected	3	3 x 8	3 x 1	27
23	Regression Output	3	-	-	0

3.3 Training Performance Analysis

The network was trained on computer hardware featuring an intel i7-10700K processor at 3.8GHz and 32GB RAM. The MATLAB Parallel Computing Toolbox was used to accelerate the training process on an Nvidia RTX 3060 GPU, which allowed for large depth of convolution layers. The datasets were partitioned, with 70% used for training, 15% for validation and 15% for testing. The effect of the number of samples on the training performance was established by training the network for 50 epochs on datasets ranging between 100 and 500,000 observations. The learning rate was set at 0.001 and the mini batches were shuffled after every epoch. The Adam algorithm was used to update the learnable network parameters. Mini batch sizes of 10, 100 and 1000 were used for datasets of less than 1000 samples, 1000 to 5000 samples and 5000 to 500,000 samples respectively. Training times ranged between 42 seconds and 80 minutes. The final validation root mean square error (RMSE) between the network predictions and known values is plotted along with the validation loss in Figure 50. In each case the test set of absorption responses were fed into the trained networks to evaluate them on unseen data. The resulting predictions for the MPP perforation diameter d, perforation ratio σ and back cavity depth D required to produce the absorption responses were compared against the actual values. The coefficients of regression between predicted and actual values that were obtained at each dataset size are displayed in Figure 51.

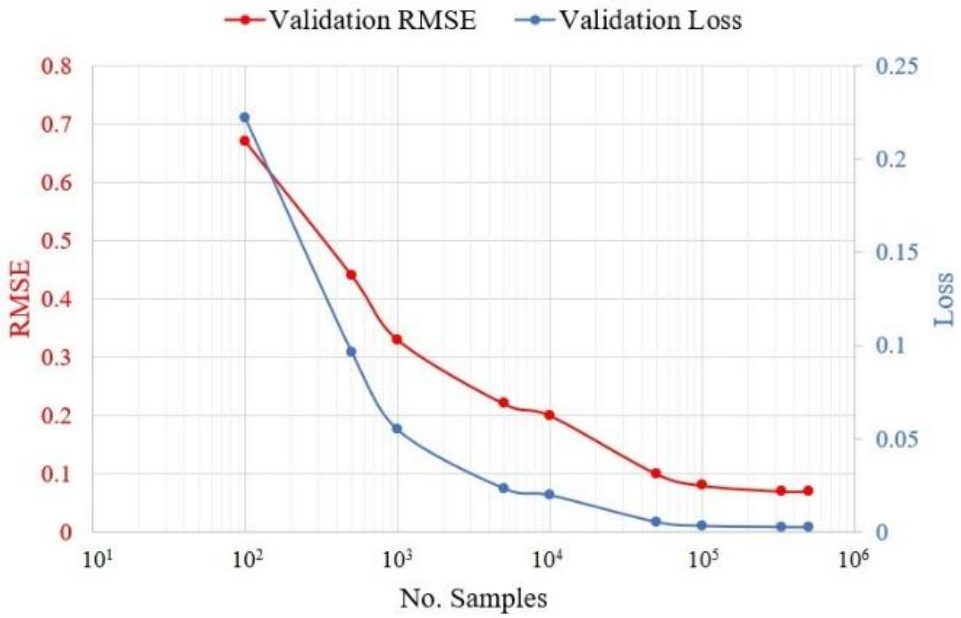


Figure 50 Effect of logarithmically increasing dataset size on validation RMSE and loss

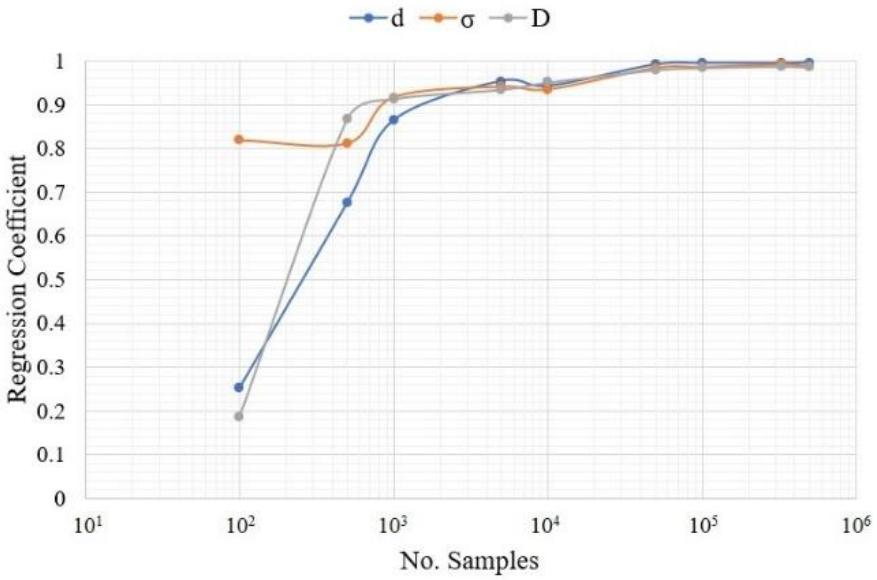


Figure 51 Effect of logarithmically increasing dataset size on regression coefficient

The results indicate datasets of at least 1000 samples are able to predict the geometry parameters at above 80% regression. To achieve 90% regression across all three parameters a dataset of 2000-5000 samples is required. The performance continues to improve until approximately 100,000 samples, where the regression, validation RMSE and validation loss all begin to plateau. At 500,000 samples there is negligible improvement to justify the additional computational costs. A dataset size of 330,000 samples was chosen as a compromise between training time requirements and prediction accuracy. At this size the network provided near perfect predictions and trained over 50 epochs in 53min 28sec. For inverse optimisation predictions for the final MPP designs the training was extended to 100 epochs to give maximum possible performance within a practical timeframe, training within approximately 1hour 45min.

A trade-off existed between the number of output parameters that it was possible to predict and the training performance. Therefore the geometry parameters to vary and those to keep constant must be carefully selected to achieve the best results. On attempting to optimise all four parameters of the MPP at once, the predictions for the panel thickness *t* suffered whereas the other three were predicted well. Table 3 shows the predictions from the best performing CNN.

Table 3 Summary of best performing DNN and CNN architectures

Case	Architecture		Training Options		Training Performance			Regression	
DNN, 4 geometry parameters predicted	Type	DNN	Training Algorithm	Adam	RMSE	Training	3.96	Parameter	R Value
			Mini Batch Size	1000		Validation	4.29		
	Hidden Layer Neurons	[801, 512, 256, 128, 64, 32, 16, 8, 4]	No. Epochs	500	Loss	Training	7.831	d	0.9021
			Initial Learning Rate	0.0001		Validation	9.205	σ	0.8001
	Activation	ReLU	Validation Frequency	50	Training Time	2hours 3min 57sec		t	0.6519
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	115000		D	0.9682
CNN, 4 geometry parameters predicted, trained for 10 epochs	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.29	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.29		
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96]	No. Epochs	10	Loss	Training	0.043	d	0.9747
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 4]	Initial Learning Rate	0.001		Validation	0.042	σ	0.8233
	Activation	ReLU	Validation Frequency	231	Training Time	10min 48sec		t	0.6810
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	2310		D	0.9646
CNN, 4 geometry parameters predicted, trained for 100 epochs	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.28	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.28		
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96]	No. Epochs	100	Loss	Training	0.040	d	0.9828
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 4]	Initial Learning Rate	0.001		Validation	0.038	σ	0.8325
	Activation	ReLU	Validation Frequency	231	Training Time	1hour 46min 18sec		t	0.6997
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	23100		D	0.9779
Final CNN, 3 geometry parameters predicted	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.05	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.07		
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96]	No. Epochs	100	Loss	Training	0.001	d	0.9966
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 3]	Initial Learning Rate	0.001		Validation	0.002	σ	0.9922
	Activation	ReLU	Validation Frequency	231	Training Time	1hour 44min 36sec		t	-
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	23100		D	0.9873

On setting the thickness constant at t = 1mm, the network offered approximately 99% regression across the other 3 variables. In practice it is desirable to fix the thickness to integer values in any case for ease of manufacture. The thickness of stock material for manufacturing of an MPP is readily available at 1mm and lower values are likely to make the material too flexible. Increasing thickness above 1mm is more likely to reduce than improve the strength of the absorption. In experiments at t = 2mm there was a drop-off in performance that required larger back cavity depths and smaller perforations to compensate. At t = 0.5mm excessive warping and flexure of additively manufactured polymeric samples was observed. It

was decided to lock the thickness at 1mm and gain better overall predictions as a result. Table 3 also demonstrates the benefits of convolution. The CNN model is compared against a standard DNN with the convolution layers removed. The convolutional network was able to produce superior training metrics and regression values using 50 times fewer epochs, in a fraction of the time of the DNN. In all cases during pre-processing all of the geometry parameters were scaled between 0 and 1. The network output values between 0 and 1 which were then rescaled back again between the original limits of each parameter. This allowed multiple parameters across different length scales to be combined in the same analysis. Without this normalisation process the predictions for the perforation diameter declined sharply from regression of approximately 0.99 to 0.67 because the values were too small compared to the other parameters. The rescaling equation is as follows:

For rescaling data X to an arbitrary interval $[a\ b]$:

$$X_{rescaled} = a + \left[\frac{X - min_x}{max_x - min_x}\right](b - a) \tag{3.11}$$

The training history of the final CNN model for predicting 3 geometry parameters to give desired MPP absorption responses in the 0-1000Hz range is shown in Figure 52. The training error and loss decrease rapidly during the first epoch. After approximately 20 epochs the training stabilises. At this point further training has little effect on the validation RMSE and validation loss, though continuing to 100 epochs gradually improves the regression statistics. Small spikes are observed in training RMSE, however there is good agreement between training and validation results and no indication of overfitting. Figure 53 plots the geometry parameters predicted by the final network against the actual known values for the 48,000 samples of test data which remained unseen by the network during training. The line of Y = T represents a perfect prediction, equivalent to a regression coefficient of 1. The network performed well on the unseen data, showing strong generalisation and offering regression coefficient greater than 0.98 for all 3 parameters. The chances of an erroneous prediction are particularly unlikely for perforation diameter. Predictions of back cavity depth showed some outliers however the network gave strong performance overall.

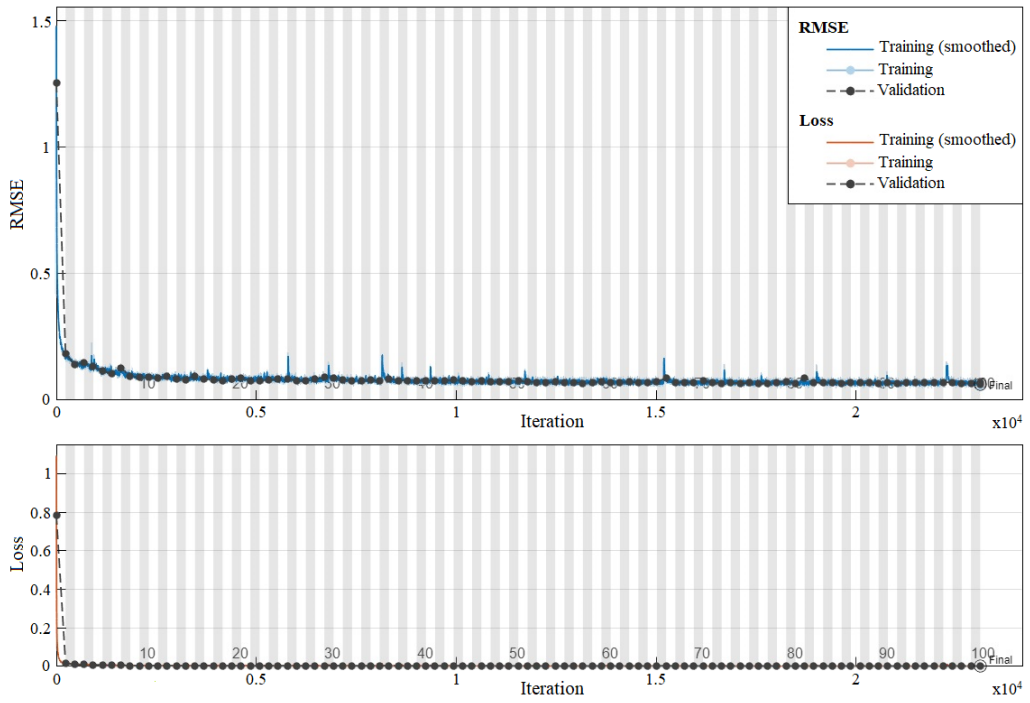


Figure 52 Training history of final CNN architecture over 100 epochs. Black markers indicate validation points

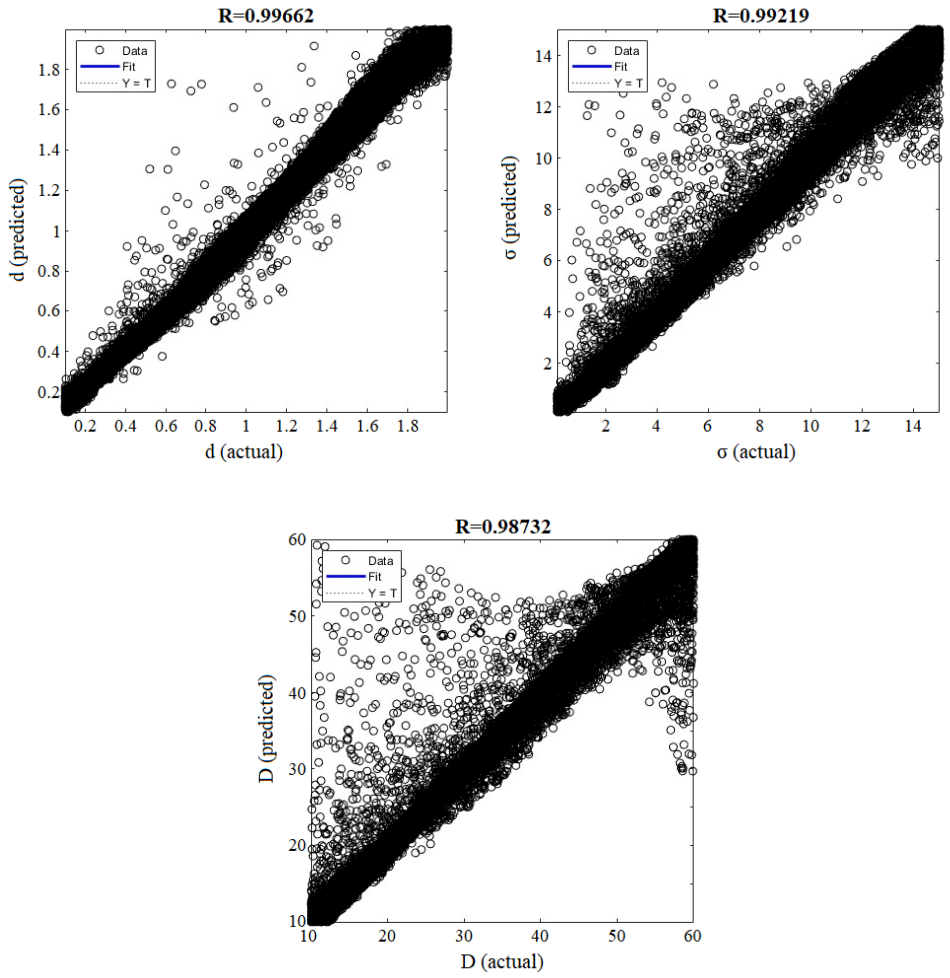


Figure 53 Regression plots of predicted vs actual MPP geometry parameters of 48,000 examples tested on unseen absorption responses

3.4 Optimisation and Experimental Validation

In order to optimise the MPP to provide maximum absorption at a specific target frequency, the desired absorption coefficient characteristic was represented as a spike input. This simplifies the absorption curve to the 3 most important points so that the amplitude, location and bandwidth of the response can be requested from the deep learning model. In addition, inputting a step function allows the absorption to be maximised over a given frequency range. The target absorption responses were input to the CNN, which returned predictions of the geometry parameters necessary to achieve them. These parameters were then fed back into the analytical solution to estimate the actual absorption response that the network had designed in response to the request. The proposed solutions were then validated experimentally on an impedance tube test rig (Figure 54) built to industry standard ASTM E-1050 [145].

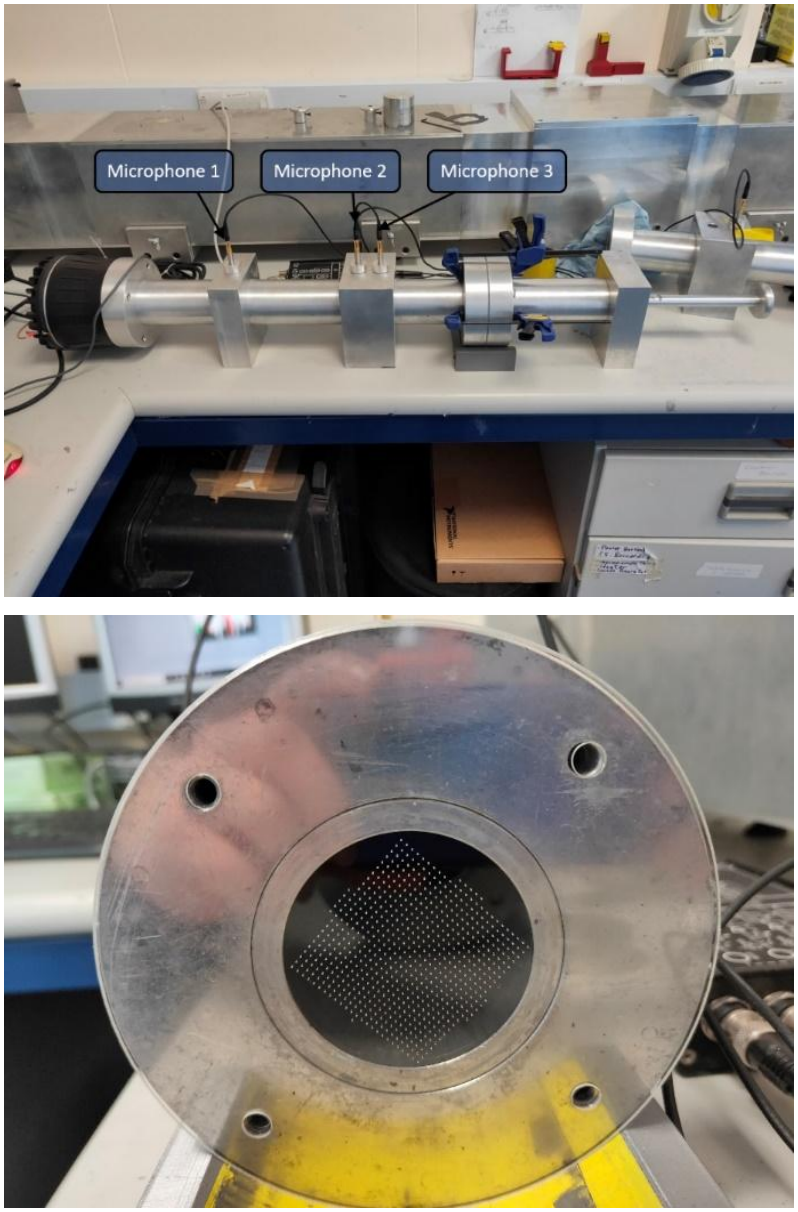


Figure 54 Impedance test rig (top) and MPP sample (bottom)

The test rig featured an aluminium tube of constant circular cross section of internal diameter 50.8mm and 15mm wall thickness. A broadband acoustic signal was generated from a loudspeaker attached at one end, propagating a plane standing wave into the tube to simulate a worst-case scenario where the wave meets the surface of the specimen at normal incidence. The test samples were placed at the other end of the tube 500mm from the loudspeaker and sandwiched in place by bolting another flanged tube which contained a plunger. The plunger was used to set the back cavity depth D, measured by a digital micrometer from the open end. Three microphones positioned upstream of the test sample were used to measure the incident and reflected sound pressure levels. Microphones 1 and 3, positioned 200mm apart were used to capture the low-frequency response between 85-700Hz, whereas microphones 2 and 3 captured the high-frequency response between 250-3000Hz and were positioned 30mm apart. Microphone 1 was mounted 160mm from the loudspeaker and microphone 3 was 140mm from the test sample. The high and low-frequency responses were combined to calculate the absorption coefficient across the range of 85-3000Hz. The measurements were post processed by a transfer function method implemented in a LabView interface according to ASTM E-2611 [146].

The complex reflection coefficient was calculated by [147,148]:

$$R(f) = \frac{H - e^{-jks}}{e^{jks} - H} e^{j2k(l+s)}, \tag{3.12}$$

where H is the acoustic transfer function, k is the wave number, s the spacing between the two active microphones, and l the distance from the sample to microphone 3. The absorption coefficient was then given by [149,150]:

$$\alpha(f) = 1 - |R(f)|^2 \tag{3.13}$$

The test samples were manufactured from discs cut from extruded acrylic and drilled on a CNC router. The number of perforations required was calculated as:

$$N = \frac{A_{active}}{a}\left(\frac{\sigma}{1+\sigma}\right), \tag{3.14}$$

where A_{active} is the total active area of the panel, a is the area of a single perforation and σ the perforation ratio. The perforations were arranged in square lattices and separated at a distance recommended by Maa theory [143] of

$$b = d\sqrt{\pi/4\sigma}, \tag{3.15}$$

where d is the perforation diameter.

Figure 55 displays example responses of MPPs optimised by the CNN to maximise absorption over the ranges of 200-800Hz and 700-1000Hz using step functions as input. The network shows strong optimisation power for maximising low-frequency absorption. Some fine-tuning of the requested input is sometimes required. It is important to ask for reasonable responses otherwise the network returns a flat absorption curve, particularly when the step input is located only at the edges of the design space. For example on asking for perfect absorption from 0 to 50Hz with a back cavity of 2mm, the network will indicate that this is not possible and the range and maximum absorption can then be updated to relax the request and try again. This allows the limits of what is physically possible and what the network is capable of achieving to be probed. There is general agreement between the analytical predictions and the experimental measurements, though the design for 700-1000Hz shows accurate range but a reduced absorption magnitude.

(a) (b)

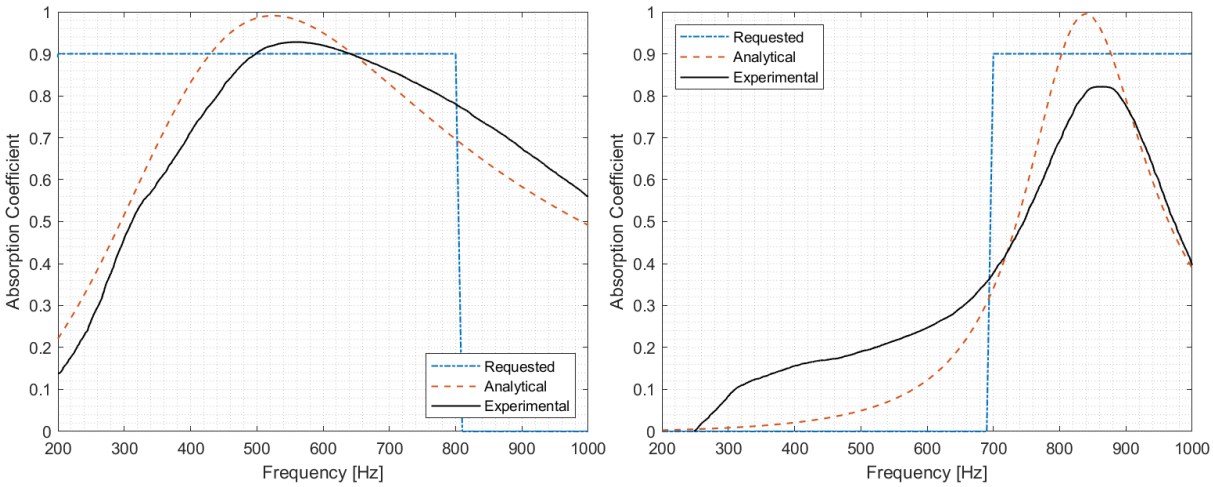


Figure 55 Requested absorption responses vs actual responses returned from the CNN model and experimental validations of MPPs optimised for (a) 200 to 800Hz, where d=0.35mm, σ=1.0594%, D=57.7mm and (b) 700 to 1000Hz, where d=0.9mm, σ=0.42171%, D=9mm. In both cases panel thickness is kept constant at t=1mm

Using spike inputs to request the desired absorption provided more predictable responses that were more accurately reflected in the experimental validations. In most cases both the absorption peak amplitude and bandwidth range found via experimentation exceeded that of the analytical responses predicted by the inverse CNN design method. Furthermore the exact range can be fine-tuned, widening or narrowing the response while maintaining the same peak frequency. Figure 56 demonstrates this ability on an MPP targeted at 800Hz. The spike input can be used to control the range, which in turn tailors the back cavity depth and overall thickness of the metamaterial. The range can simply be reduced until the required thickness is

achieved, resulting in the maximum absorption over the greatest range centered on any frequency of target. Alternatively the design space of the training dataset can be altered, limiting the back cavity ranges and training only on MPPs of a certain overall thickness, as explored in Section 3.6.

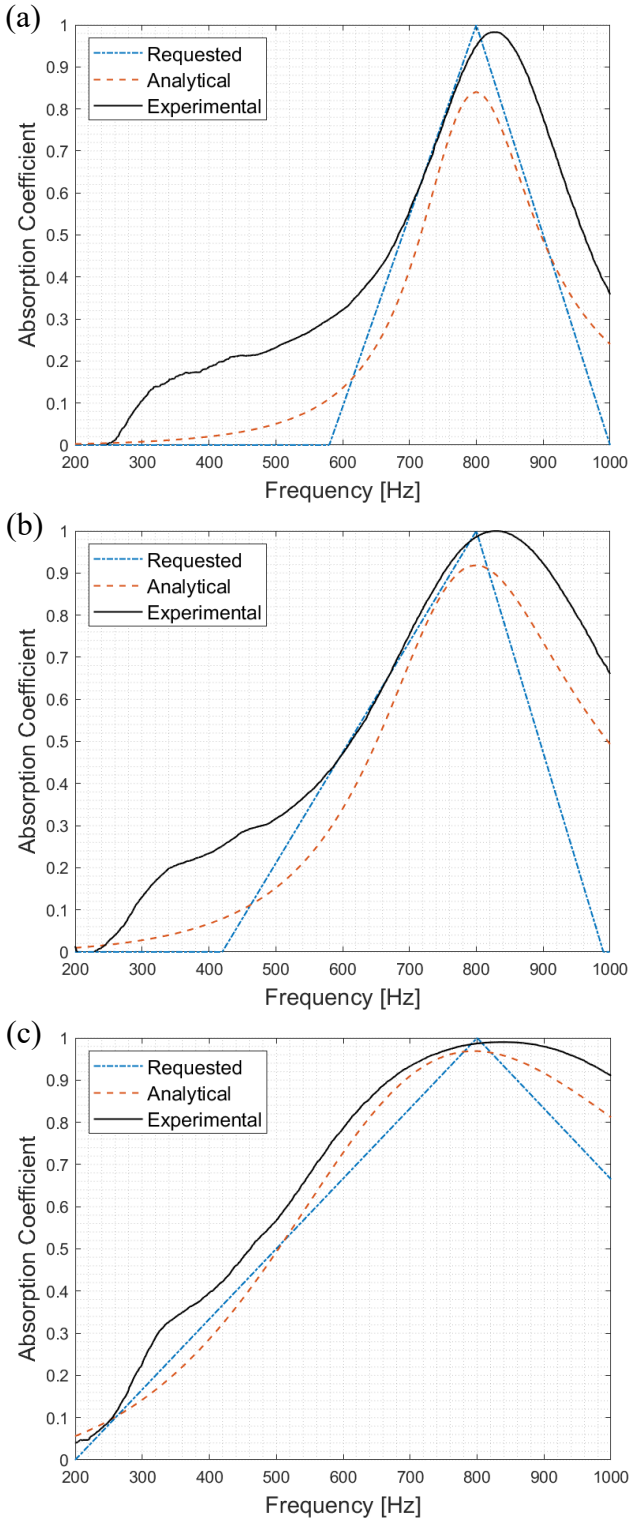


Figure 56 Requested spike inputs vs actual returned analytical solutions and experimental validations for MPP designs targeted to give maximum absorption at 800Hz. (a) narrow bandwidth requested, d=1.25mm, σ=0.60355%, D=13mm, t=1mm (b) medium bandwidth, d=0.75mm, σ=0.80806%, D=20mm, t=1mm (c) broad bandwidth, d=0.4mm, σ=1.6089%, D=40mm, t=1mm

Figure 57 displays the ability of the CNN model to produce MPP designs at frequencies between 400 and 800Hz by using spike inputs to request each target frequency. Once trained the network was able to return all predictions in a matter of minutes. All peaks occurred within 94% of the target frequencies. Small shifts of between 20 and 40Hz to slightly higher frequencies were observed in all cases, though the absorption coefficient at the exact target frequency was above 0.95 for all except the lower frequency MPP at 400Hz, which gave coefficient of 0.8.

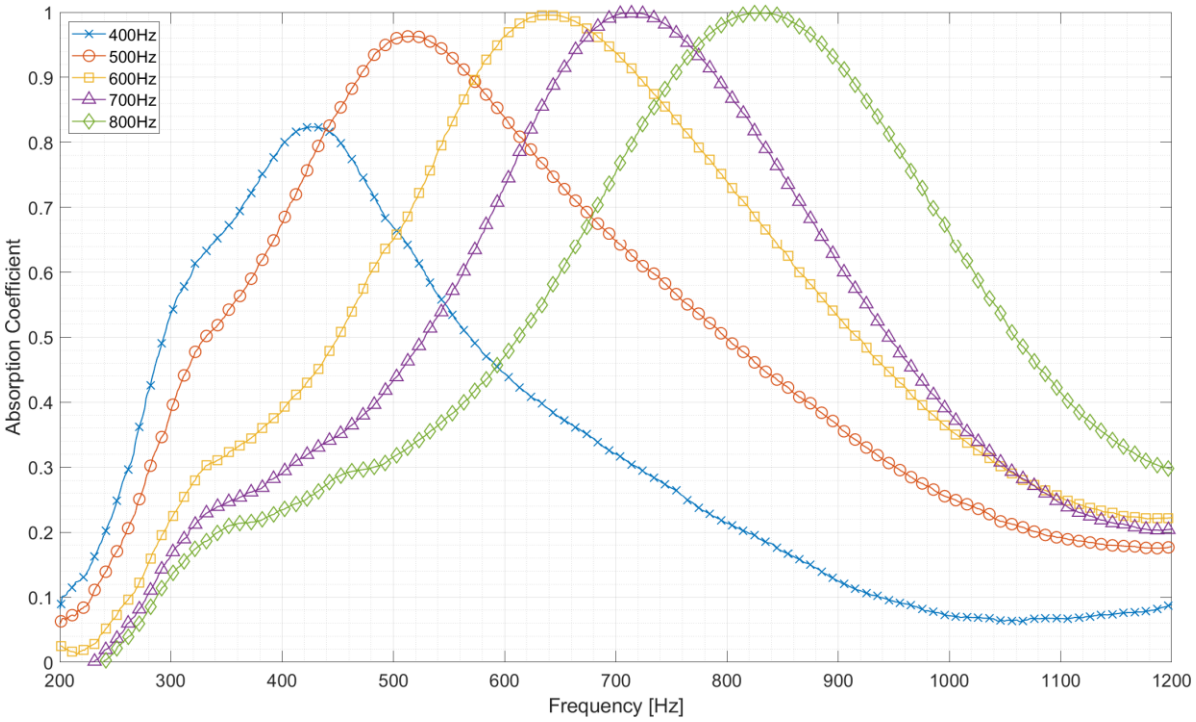


Figure 57 Experimental absorption responses of MPPs optimised by the CNN to target specific frequencies between 400-800Hz

Results of the best performing MPPs optimised by the CNN to provide the greatest absorption over the widest range in the low-frequency 0 to 1000Hz regime using the spike and step input approaches are shown in Figure 58. The spike input provided at least 70% absorption over a range of 923Hz, from 548 to 1471Hz, at a back cavity depth of 40mm. The step input result was obtained by requesting perfect absorption coefficient of 1 between 200 and 1000Hz. The CNN returned an MPP design that measured no less than 70% absorption over 1000Hz of range, from 470 to 1470Hz, at a back cavity of 56mm.

(a) (b)

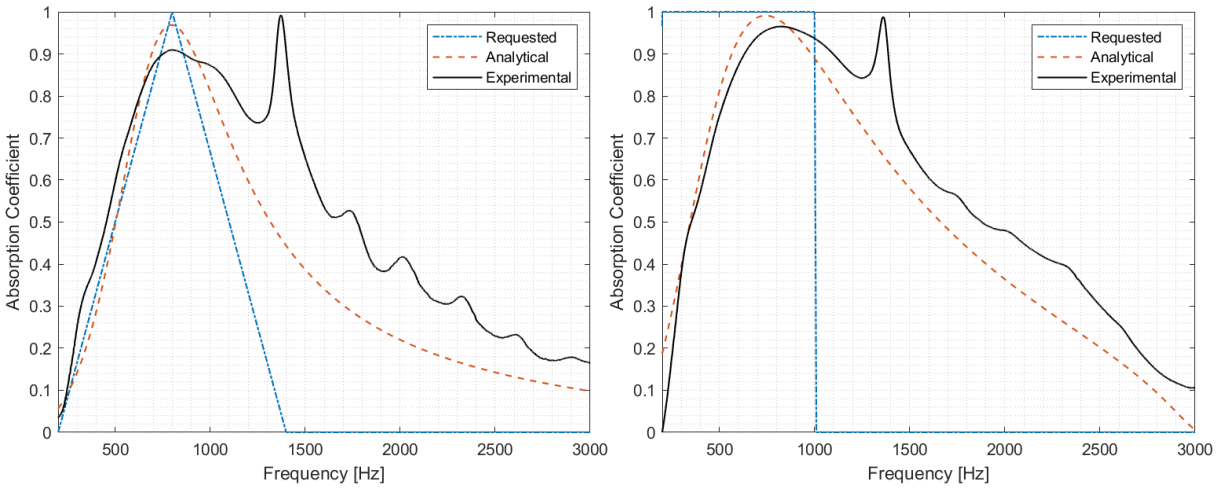


Figure 58 Absorption responses of MPPs optimised by the CNN to maximise absorption between 0-1000Hz using (a) spike inputs, d=0.4mm, σ=1.6089%, D=40mm, t=1mm and (b) step inputs, d=0.25mm, σ=2.3325%, D=56mm, t=1mm

3.5 Compensation for Materials and Manufacturing

As the CNN model is able to return the geometry parameters for a desired absorption characteristic near instantaneously, this grants the ability to accurately and rapidly fine-tune the response to compensate for discrepancies between the analytical and experimental absorption curves caused by the materials and manufacturing methods used to create the test samples. During the workflow between an analytical solution and the sound measured by a microphone in real-world testing many sources of error are introduced at various stages. General inaccuracies are inherent in matching the analytical results to available material sizes and drill bit diameters, typically the diameter must be rounded to the nearest available drill bit and the perforation ratio adjusted to give an integer whole number of perforations. The impedance tube is unlikely to perfectly replicate pure plane wave conditions and may give better results when the perforations are more tightly spaced in the centre of the sample. The assumption of normal incidence used in the analytical theory may be a source of error as some diffusion of sound waves is unavoidable and is expected to have occurred within the tube. Measurements depend on the quality and calibration of the microphones and correct mounting distances, which are affected by the geometric tolerances of the rig manufacture. The tube is not perfectly airtight and energy losses from leakages are to be expected. The air density and air viscosity coefficient may not match theoretical values in practice and results are likely to vary with room temperature. Manually setting the back cavity depth may have resulted in errors associated with the precision and orientation of the micrometer, which may not have

been perfectly level. The sample perforations may not have been perfectly perpendicular to the panel surface due to flexure in the samples as they were drilled or slight misalignments of the drill bit in the chuck, meaning the perforation cavity walls are not entirely aligned with the direction of the wave within the tube. The stock material for the samples was extruded acrylic. The thicknesses would have been dependent on the tolerances of the extrusion process. Acrylic samples manufactured by manual drilling on a pillar drill, laser cutting and machining on a CNC router were compared under optical microscope inspections using a Keyence VHX-6000. Example perforations from each process are shown in Figure 59, all having a nominal target diameter of 750µm.

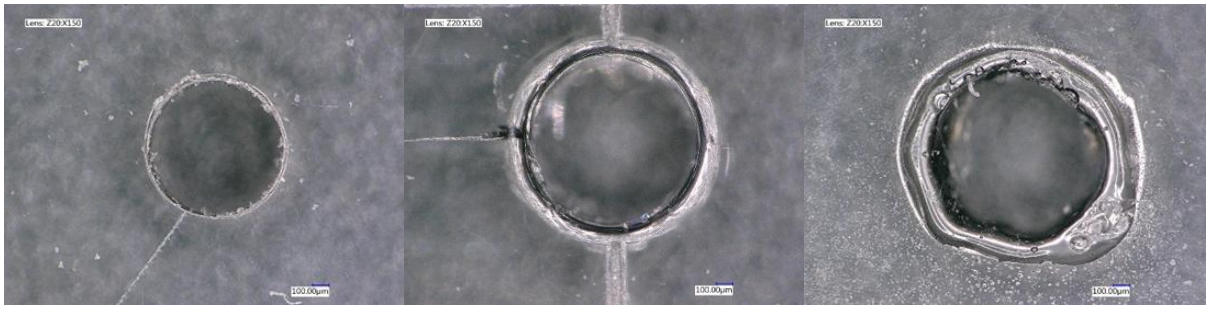


Figure 59 Optical microscope scans of perforations manufactured by CNC (left), pillar drill (middle) and laser cutting (right). All scans are at the same scale of 150x magnification. Nominal perforation diameter in all cases is 750µm

The perforations of the manually drilled samples were considerably oversized due to swarf accumulation, melting of material and inconsistent plunge rate. The minimum resolution of the laser cutting process was highly dependent on the laser power. In all cases excessive melting of acrylic caused the perforations to be oversized. The scans were all taken from MPP samples of the same design optimised by the CNN model to provide peak absorption at 800Hz. The resulting absorption characteristics are compared for each manufacturing method in Figure 60. The CNC sample provided an absorption peak at 823Hz, whereas the laser cut and pillar drill versions are shifted to considerably higher frequencies, 989Hz and 1086Hz respectively. The actual perforation diameters were measured from microscope scans of the samples to determine the errors of each manufacturing method. The diameters were averaged across all of the perforations across each sample.

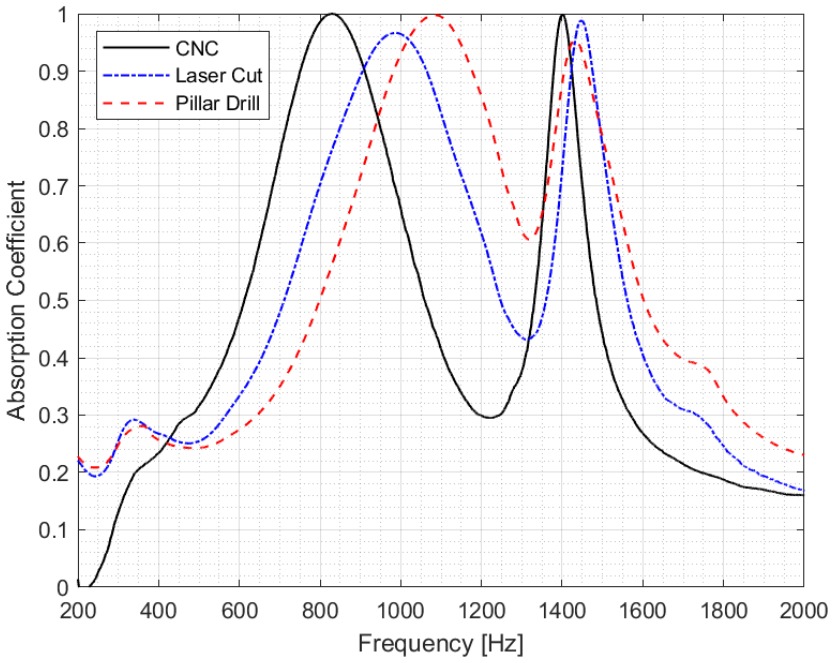


Figure 60 Manufacturing method comparison for MPP designed for 800Hz; d=0.75mm, σ=0.80806%, D=20mm, t=1mm, no. perforations=37

Each sample featured 37 perforations in total of nominal target diameter 750µm. An automatic scanning method was used to measure the circularity of the perforations. The perforated regions in each scan were extracted according to the brightness values in the image. The degree of circularity was then calculated as $4\pi A/P^2$, where A is the total area of pixels within the perforation and P is the perimeter of the extracted region. The results, displayed in Table 4 demonstrate substantial errors as a result of the manufacturing process for both laser cutting and manual pillar drilling. Laser cut samples showed poor circularity due to melting of the acrylic. CNC provided the best manufacturing tolerances, average perforation diameter was accurate to within 36µm which explains the improved absorption performance.

Table 4 Measurements of manufacturing accuracy

	Laser Cut		**Pillar Drill**		**CNC**	
	Diameter (µm)	Circularity	Diameter (µm)	Circularity	Diameter (µm)	Circularity
Average	1088.68	0.85	1025.38	0.95	767.03	0.96
Standard Deviation	23.05	0.06	77.46	0.05	7.22	0.02
Max	1155.16	0.91	1166.18	0.99	786.06	1.00
Min	1036.78	0.70	832.53	0.74	753.01	0.92
Range	118.38	0.21	333.65	0.26	33.05	0.08
Target Nominal Diameter	750.00	-	750.00	-	750.00	-
Upper Tolerance	405.16	-	416.18	-	36.06	-
Lower Tolerance	286.78	-	82.53	-	3.01	-
Count	37					

By benefiting from the speed and flexibility of the inverse CNN approach it was possible to successfully decouple the absorption characteristic from the material or manufacturing method used, producing performance at the desired frequency irrespective of the manufacturing error. Once an initial design is optimised for the given target frequency and validated experimentally it can be continuously refined, producing multiple iterations at intervals around the target within minutes without any additional calculations. Figure 61 shows the effects of using the CNN model to compensate for the overshoots observed in the laser cut and manually pillar drilled samples targeted at 800Hz. The absorption peaks of the laser cut and pillar drill samples now occur at 850Hz and 790Hz respectively. The pillar drilled MPP was able to produce better results after compensation, likely due to the superior circularity of the perforations. In this case the CNN model effectively eliminates the influence of the previous manufacturing error.

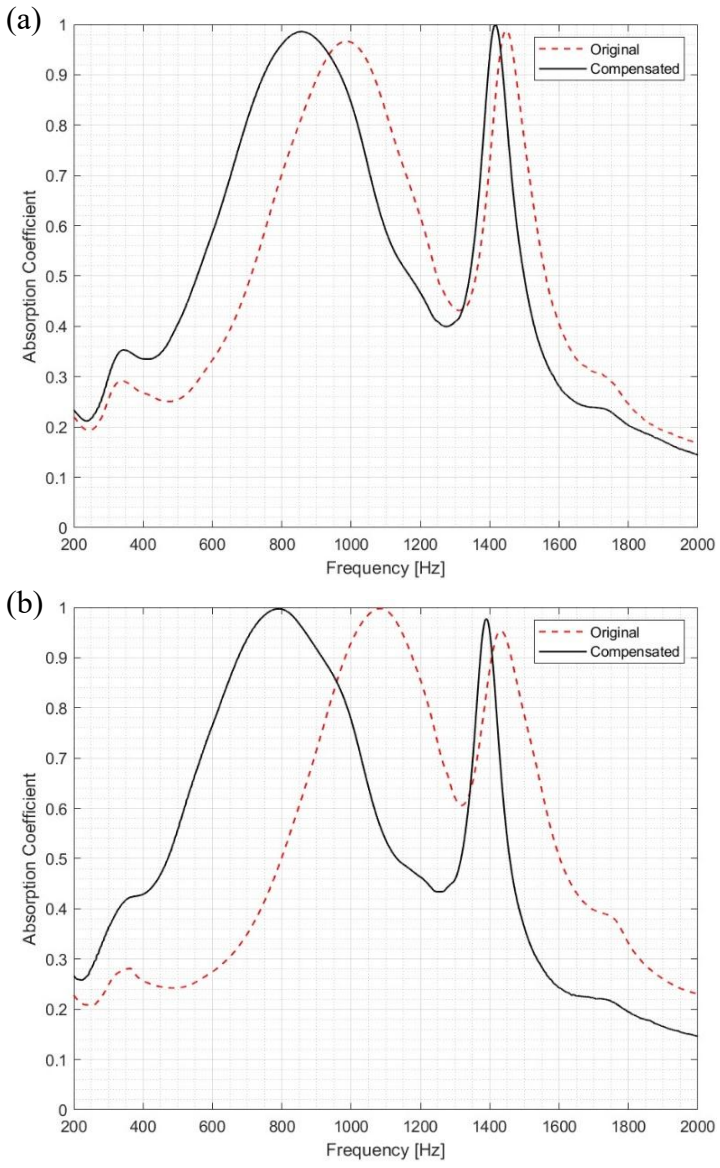


Figure 61 Original absorption response vs response compensated using CNN model for (a) laser cut and (b) pillar drill experimental samples

Figure 62 shows similar attempts at compensating for the use of additive manufacturing and different base materials for the same MPP designed for 800Hz. Fused Deposition Modelling (FDM) 3D printed samples showed inconsistent thickness and poor surface finish due to the hatch pattern of alternating layers as the sample is built. The responses featured a distinctive double peak characteristic. One peak occurs due to the natural vibration mode of the panel membrane, while the other is associated with the viscous boundary layer effects within the tube. In all cases a strong peak occurred at a much higher frequency than the one targeted. The CNN was able to shift this peak from 1005Hz to 826Hz and 1091Hz to 807Hz for FDM nylon and PolyEthylene Terephthalate Glycol (PETG) samples respectively, though peak bandwidth was sacrificed.

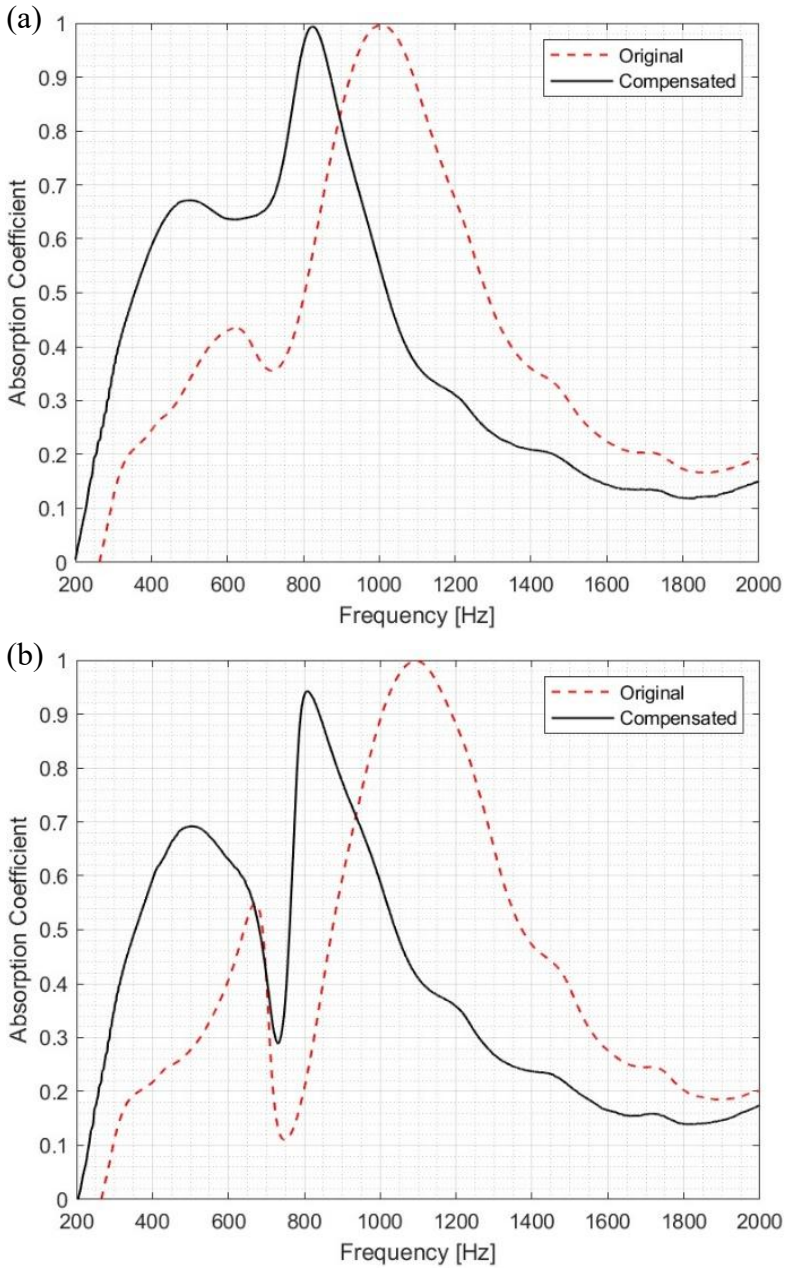


Figure 62 Original vs compensated absorption responses for FDM (a) nylon and (b) PETG MPP samples designed to maximise absorption at 800Hz

3.6 Design of Subwavelength Wideband Absorber for Industry Application

The inverse CNN optimisation method was applied to design a sound absorbing MPP based metamaterial for use in automotive/energy applications to filter out low-frequency noise. The target design specification was to achieve 50% absorption between 500-1000Hz at a maximum of 25mm overall thickness. To attain such a low thickness the model was applied to parallel arrangements of multiple MPPs. The schematic and electro-equivalent circuit of a triple parallel MPP configuration is displayed in Figure 63.

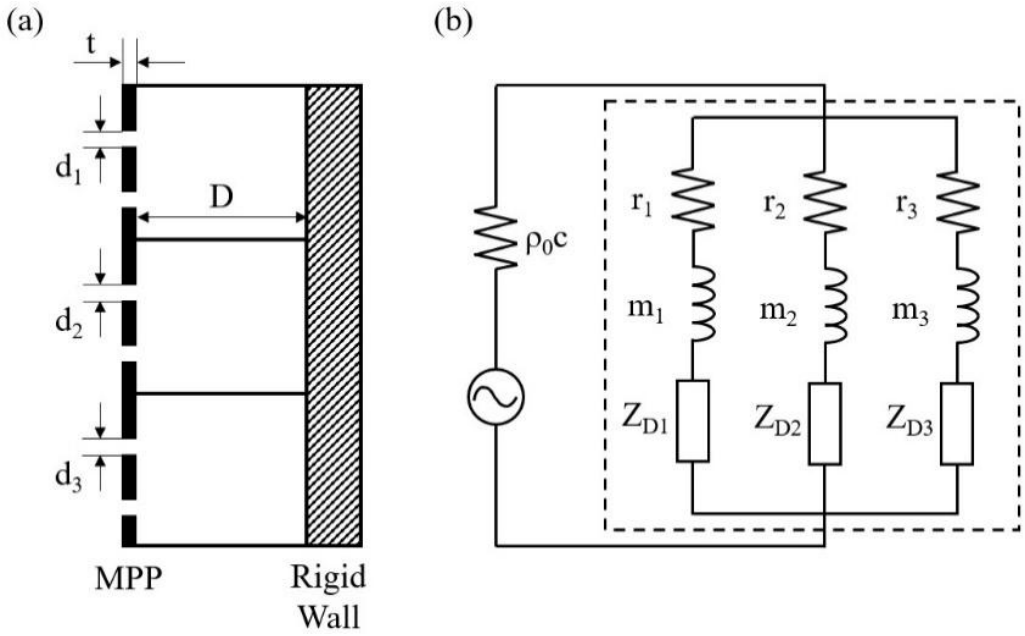


Figure 63 (a) Model of triple parallel MPP (b) electro-acoustic equivalent circuit

The total acoustic impedance of the three MPPs arranged in parallel is given by [135]:

$$Z = \left(\sum_{i=1}^{3} \frac{\phi_i}{Z_i}\right)^{-1} \tag{3.16}$$

The coefficient of sound absorption is a function of the real and imaginary parts of the total acoustic impedance, given by [151]:

$$\alpha = \frac{4Real(Z)}{(1+Real(Z))^2+(Imag(Z))^2} \tag{3.17}$$

The same CNN model was retrained to limit the value of back cavity depth by only training on MPPs where $D < 25mm$. The remaining parameter limits were kept the same as before and panel thickness was kept constant at $t = 1mm$. A solution using only a single MPP at overall thickness of 21.6mm was found by using step inputs to request 70% absorption between 350Hz and 1000Hz and was validated experimentally as shown in Figure 64. The

design provided 50% absorption between 548Hz and 1016Hz. To improve the low-end response the optimisation process was repeated with double and triple parallel MPP arrangements. In these cases the back cavity depth was now kept constant at $D = 22mm$ to allow additional depth for a rigid back wall and the remaining two geometry parameters, d and σ were predicted by the network for each MPP individually before parallel combination. Only predicting two parameters further improved the training performance. As shown in Figure 65 the CNN provided regression above 99% across both parameters, using the same dataset sizes and segmentation as previously.

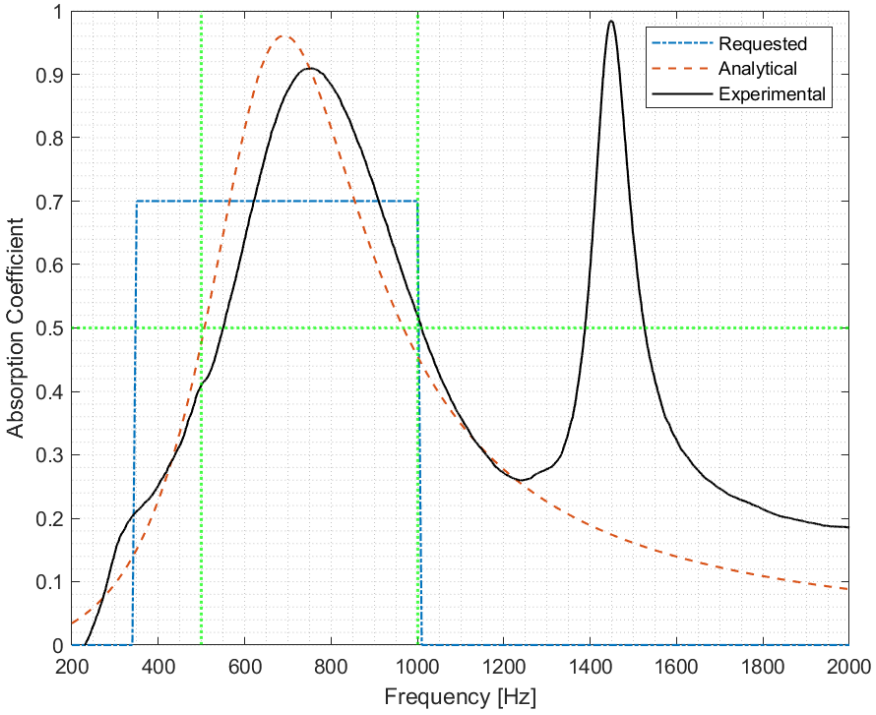


Figure 64 Single MPP solution; d=0.45mm, σ=0.58803, D=20.6mm, t=1mm

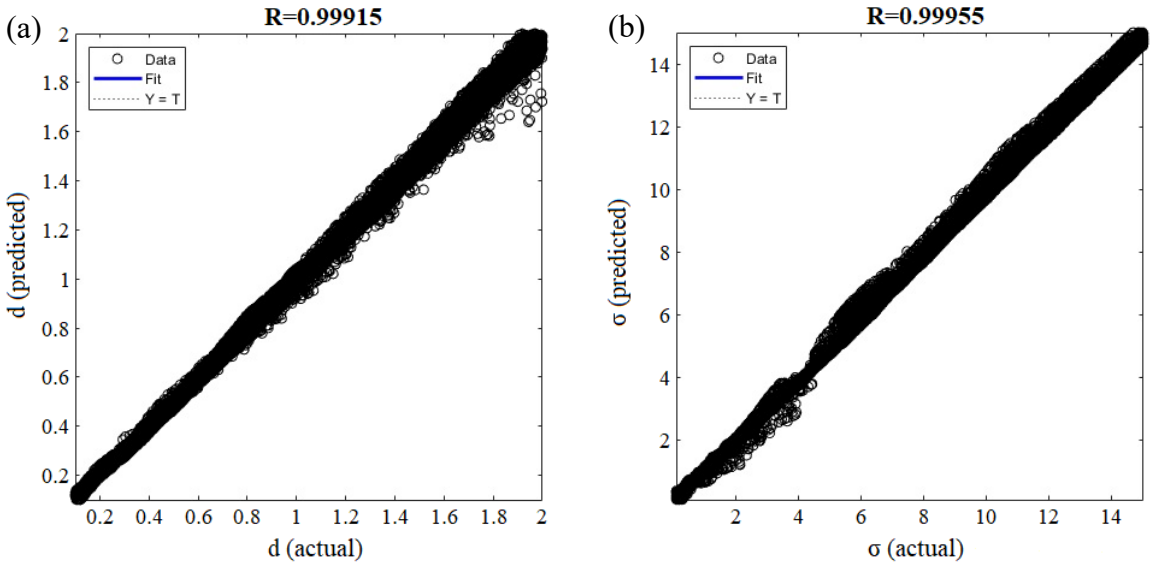


Figure 65 CNN regression values for (a) perforation diameter and (b) perforation ratio

Back cavity dividers were CNC machined from solid blocks of acrylic, splitting the cavities into equal parts. The thickness of the back wall was 1mm, giving an overall sample thickness of 24mm. To eliminate errors in panel thickness t associated with machining tolerances the perforated panels were drilled from discs of extruded acrylic and fixed to the back cavity sections with adhesive. An example triple parallel MPP test sample set in the test rig is shown in Figure 67. The experimental results for double and triple parallel configurations optimised by the CNN model are compared against the analytical predictions in Figure 66. The double MPP boosted the high-end 50% absorption threshold to 1122Hz, though gave the same low-end threshold as the single MPP of 548Hz. The triple parallel combination provided the best low-end bandwidth, giving 50% absorption between 518-1027Hz at overall thickness of 24mm.

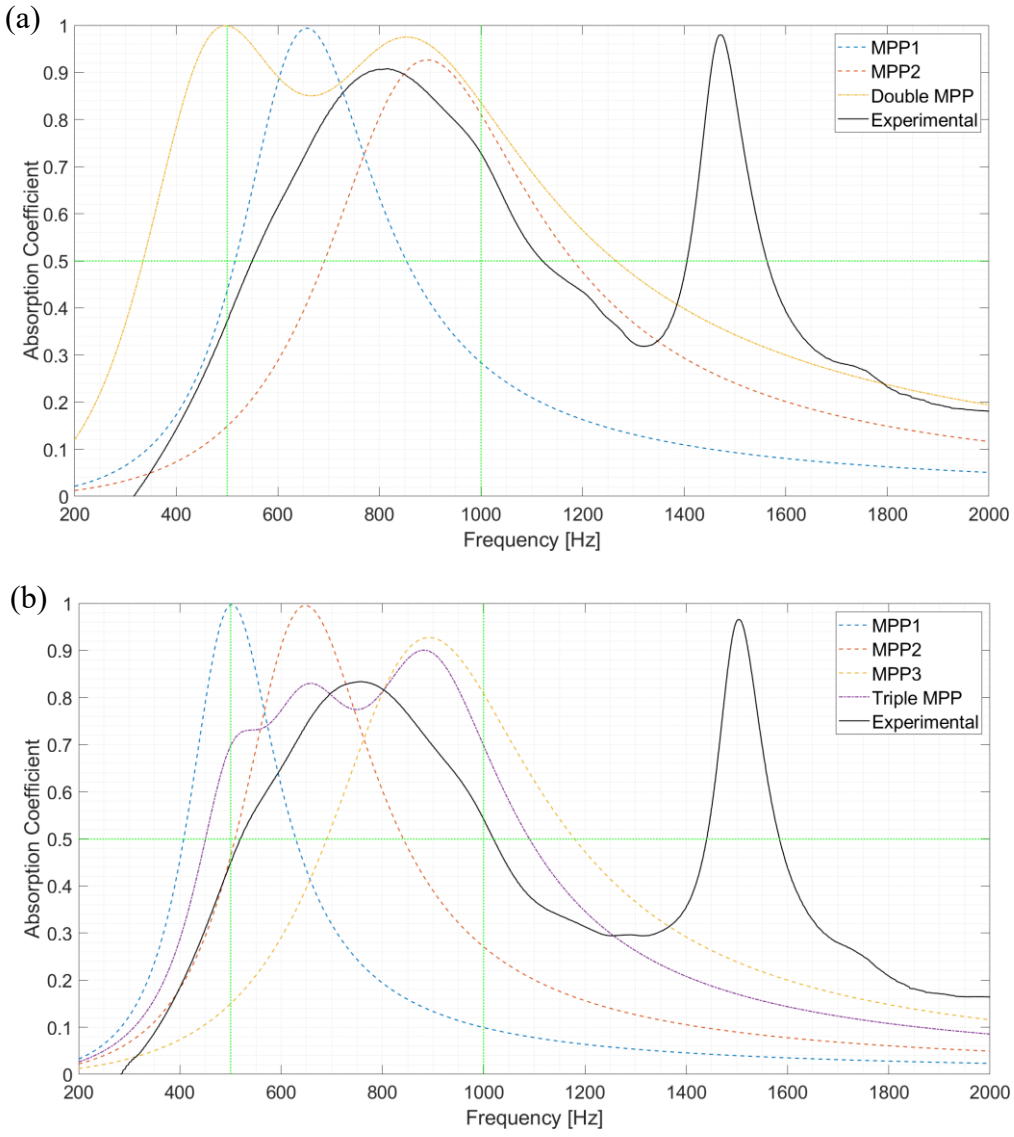


Figure 66 Experimental vs analytical absorption responses for parallel MPP configurations optimised via CNN model (a) double parallel MPP; d1=0.65mm, d2=0.6mm, σ1=0.5723, σ2=1.0655, D1=D2=22mm, t1=t2=1mm and (b) triple parallel MPP; d1=0.75mm, d2=0.65mm, d3=0.6mm, σ1=0.3571, σ2=0.5727, σ3=1.0655, D1=D2=D3=22mm, t1=t2=t3=1mm

The triple parallel MPP design was selected and validated experimentally at a larger sample size of 185mm square cross section. For the larger samples the MPPs were arranged in a hexagonal honeycomb tessellation with wall thickness of 1mm between adjacent cells, as shown in Figure 67.

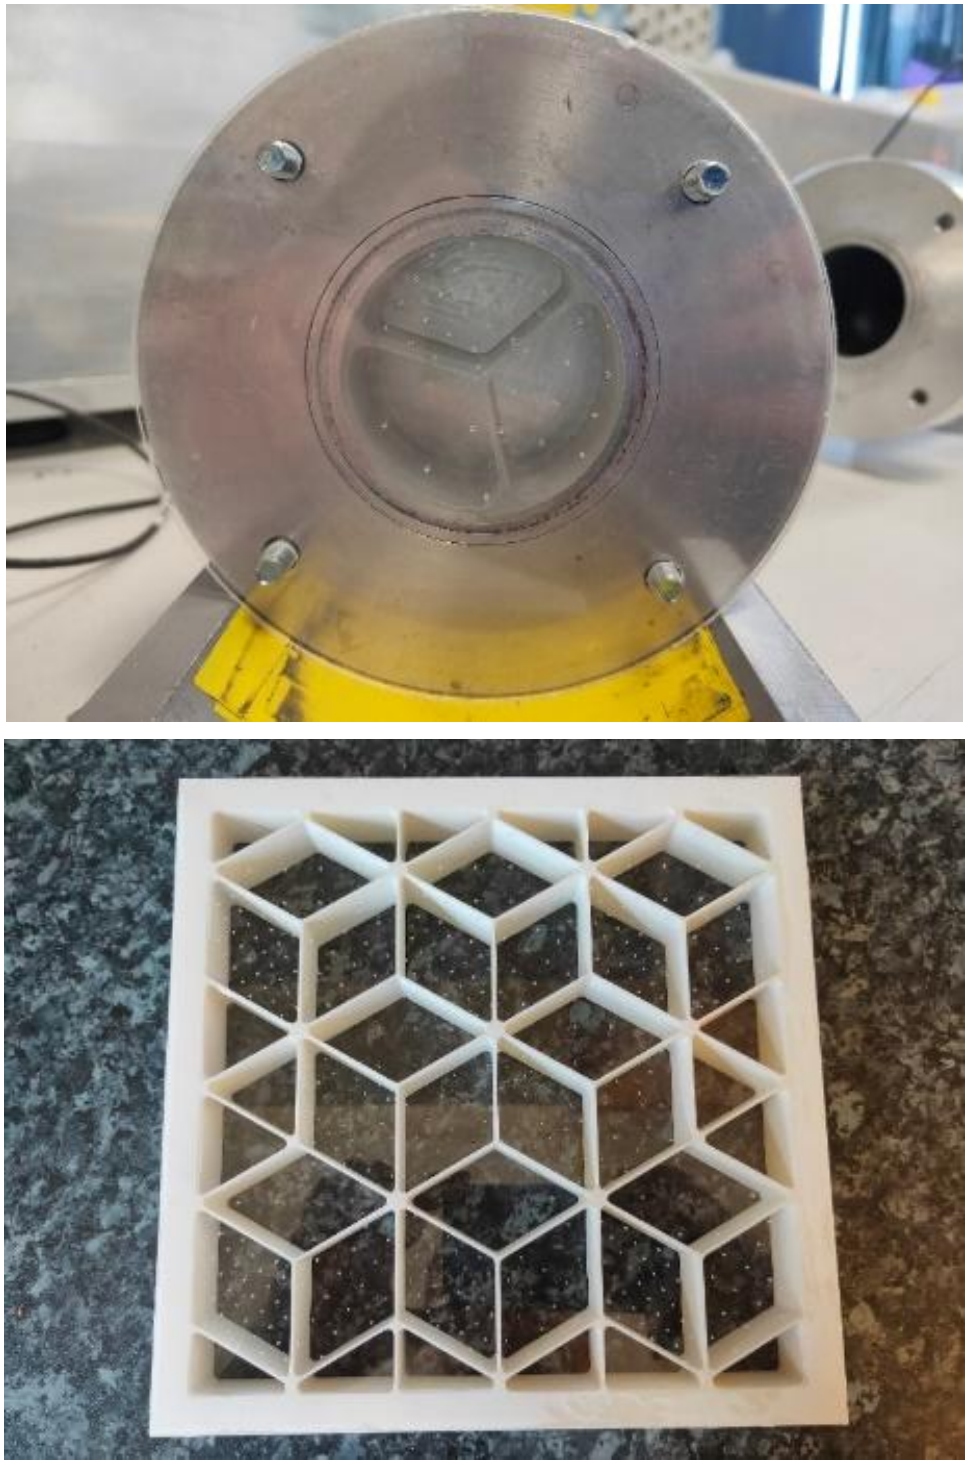


Figure 67 Triple parallel MPP test samples for 50.8mm diameter circular impedance tube (top) and 185mm square tube (bottom)

Figure 68 displays the experimental vs analytical absorption responses for the final triple parallel MPP design, tested on an impedance tube of 185mm square internal cross section. The measurement range of this tube was limited by the increased dimensions. The absorption could be reliably calculated between 200-800Hz. After 845Hz the signal began to return to zero before full desaturation beyond 1050Hz therefore the measured high-end threshold is likely a worst case scenario. Validations were made across a range of back cavity materials. In each case a perforation panel drilled from extruded acrylic using CNC was fixed to the back cavity with adhesive. All results fell within acceptable limits of the initial target specification. The peak absorption amplitudes exceeded those of the smaller samples. The parallel superposition provided experimental bandwidth less than that predicted by theory. In most cases the actual recorded peaks occurred at points midway between adjacent peaks in the analytical response. The best performing sample was manufactured via FDM from eSUN PLA+, a reinforced PLA filament, which provided 50% absorption between 506Hz and 1032Hz. The lowest bandwidth was recorded for CNC machined acrylic back cavities, though the absorption was maintained above 50% from 521Hz to 1010Hz. For these samples the back cavities were open, the rigid back wall was set by a square plunger held against the sample. The overall thickness of the metamaterial was therefore only 23mm, the combination of panel thickness and back cavity depth.

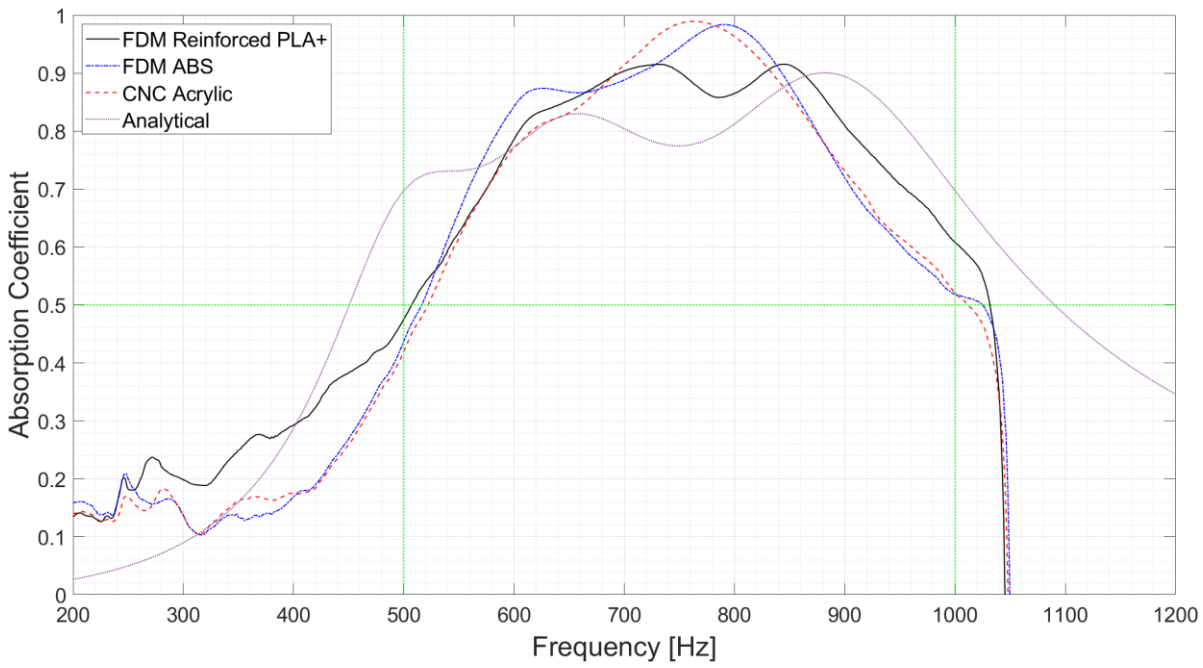


Figure 68 Experimental validation of 185mm triple MPP samples for a variety of back cavity materials

3.7 Design of Experiments Study of Dataset Size Requirements

A critical factor in the use of deep learning to optimise metamaterials is the dataset size requirements, particularly when the dataset can only be generated using FEA or experimental testing. If initial results are unsatisfactory, the model can be improved by increasing the amount of training data until the required level of performance is met. Alternatively the complexity of the problem can be reduced. For example if instead of a full stress-strain response only the elastic modulus, yield stress and ultimate tensile strength are predicted the inputs and outputs of the network would now have the same dimensionality and can likely be mapped using a standard DNN with only a few hidden layers. This may also provide the option for training on purely experimental data, with sample numbers below several hundred. The potential for reducing the dataset size requirements for further sections of the project was evaluated using a design of experiments approach. The CNN model for optimising the MPP sound absorbers was retrained on datasets generated using Latin hypercube, Sobol sequence and Halton sequence sampling in addition to a full factorial test plan featuring 6 levels. Table 5 shows the discrepancy, minimum and maximum point-to-point deviation values for examples of the space filling designs with 1300 samples. Pairwise projection plots for each design can be found in Appendix A.

Table 5 Discrepancy and deviation values for design of experiments test plans of 1300 samples

	Latin Hypercube	Sobol Sequence	Halton Sequence
Discrepancy	2.0433	1.4180	1.0319
Min Point-Point Deviation	0.0673	0.1131	0.1109
Max Point-Point Deviation	3.6907	3.5624	3.5558

Discrepancy is defined as a measure of the amount of deviation from the average point density. Latin hypercube offered the highest value, indicating a better spread of datapoints throughout the design space may be possible when the number of observations is limited. The minimum deviation between datapoint locations was also reduced compared to Sobol and Halton sequences, which showed similar deviation values to each other. This is desirable to limit the amount of empty space in the design, however Latin hypercube also showed the largest maximum point-to-point deviation.

The same convolutional network evaluated in Section 3.3 was trained for 50 epochs on datasets of reduced sizes ranging between 100 and 5000 samples. In each case 70% of the

data was used for training, with 15% used for validation and 15% for testing. The regression coefficients between predicted and known values of the perforation diameter, perforation ratio and back cavity depth were calculated for each test dataset. The mean regression across all three geometry parameters is plotted against the size of the dataset for each sampling method in Figure 69. Plots of validation error, loss and regression for all geometry parameters for each test run can be found in Appendix B.

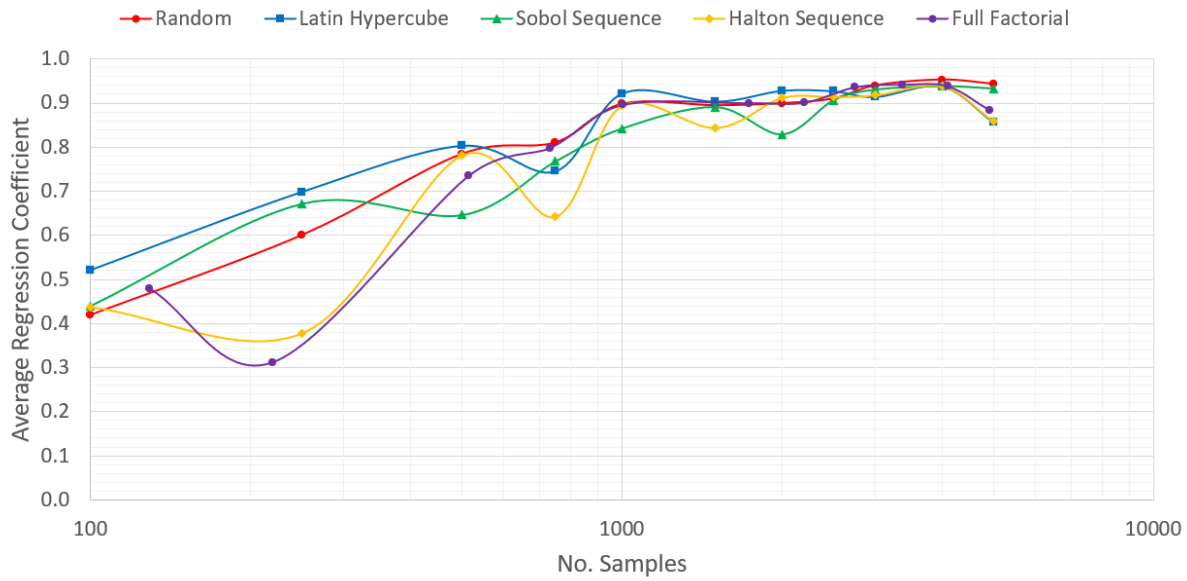


Figure 69 Mean regression vs dataset size for design of experiments test plans

In general increasing the number of samples in the dataset resulted in improved regression. However some variation was observed as peaks and dips as the dataset size increased. In some cases these steps in performance are likely to be a result of increasing the mini batch size to accommodate larger datasets. Typically a trade-off exists when selecting mini-batch size between performance and training speed. Datasets between 100 and 1000 samples were passed through the network in mini-batches of 10 samples at a time. From 1000 samples this was stepped up to batches of 100. At 5000 samples the mini-batch size was increased again to 1000 samples. Latin hypercube sampling showed the best performance at low dataset sizes, giving the best regression results for all datasets below 1000 samples with the exception of an outlier at 750 samples. At this point the majority of other methods also showed a dip in performance, which then improved significantly once the mini-batch size was increased for the 1000 sample datasets. The advantage of Latin hypercube sampling began to drop off for larger datasets. Above 3000 samples there was less difference between approaches and both the randomised and space filling designs were able to offer regression above 90% across all geometry parameters. The randomised distribution of datapoints in fact slightly outperformed the design of experiments approaches. However at 5000 samples there was a drop in

regression across all methods as a result of the increase in mini-batch size. It may be possible to improved performance further by increasing the number of epochs and training for a greater amount of time. On training the network on full factorial datasets it became able to produce predictions only at intervals corresponding to the test plan levels. Due to this it is not recommended for problems where the output parameters contain continuous data.

For large datasets with greater than 100,000 samples there is expected to be little difference compared to the randomised distribution of datapoints, as there is likely to be very little empty space in the designs regardless of the sampling method. However at low sample numbers typical of those that may be practical to generate via FEA, the design of experiments sampling was able to improve the predictive power of the network. The sampling approach that offered the best network performance at limited dataset sizes was Latin hypercube, which is selected as the preferred test plan for following works. The minimum number of samples required to produce predictions within 90% regression of known values across all parameters is expected to be within the range of 1000 to 2000 samples for problems of similar complexity.

3.8 Conclusion

This study introduces a machine learning framework for designing sound absorber metamaterials to achieve specific absorption characteristics at subwavelength thicknesses. The dataset of absorption responses was generated analytically from a micro-perforated panel (MPP) unit cell, which was then used to train a convolutional neural network (CNN). This CNN predicted the necessary geometry parameters for desired low-frequency absorption (0-1000Hz) with over 98% accuracy. The model allowed for precise control over the absorption bandwidth and amplitude at target frequencies. It demonstrated at least 70% absorption across the frequency ranges of 548-1471Hz and 470-1470Hz with single MPPs of 41mm and 57mm thickness, respectively. Additionally, an arrangement of three MPPs in parallel provided 50% absorption between 506-1032Hz at a thickness of 23mm. This machine learning approach significantly enhances the speed and flexibility of design compared to traditional methods. If discrepancies arise during real-world testing, the CNN can quickly adjust the design. The framework also considers manufacturing constraints by including only feasible parameters in the training dataset. Optimisation is achieved by setting limits on the design space and retraining the model. Future work may explore alternative geometries and multifunctional designs, incorporating structural, thermal, or aerodynamic properties for multi-objective optimisations. The development of efficient sound absorber metamaterials can significantly improve noise management across various industries, such as the automotive and aerospace sectors, where this technology can lead to quieter and more comfortable vehicles and aircraft, enhancing passenger experience and reducing noise pollution. By streamlining the design process and reducing material weight and costs, this approach can lead also to more lightweight and cost-effective production of sound absorbers.

Chapter 4 Artificial Intelligence Based Optimisation of Metamaterials for Tailored Stress-Strain Response and Enhanced Structural Efficiency

The majority of industrial engineering design practices currently apply structural optimisations at the macro level only, defining the distribution of mass through a component by altering its overall shape to satisfy design criteria. At the material level the micro and mesoscale structures typically remain limited to selections between traditional monolithic bulk materials. Even more advanced composite components generally rely on the use of pre-existing periodic lattice core materials with fixed geometries. This leaves potential for additional structural efficiency enhancements and tailoring of mechanical properties for specific loading conditions through application of architected cellular metamaterials and artificial intelligence based design optimisation. This section of the research aims to tailor material structure for a given real-world loading scenario according to engineering design constraints, setting the mechanical properties by material level architecture rather than overall component level geometry. Historically engineers have developed components by iterating geometry and then using software tools to analyse the design in terms of its mechanical performance. This study proposes a new approach that works in the opposite direction. The desired mechanical properties are requested by the designer initially and the software tools then return the required architecture, allowing for previously unobtainable levels of flexibility and design freedom. These works present an overview of the inverse design workflow to allow for ease of application of similar AI optimisation models within industry, demonstrating the potential advantages, assessing the time requirements and highlighting issues that may occur. The approach is applied to a case study of a lattice compression block under uniaxial loading for proof of concept. Properties of additively manufactured base materials are measured experimentally and used to build a non-linear FEA model to generate a dataset of sample mechanical responses. An inverse convolutional neural network is then trained to predict the geometry of an octet-truss framed lattice unit cell for a given stiffness and strength requirement. The ability of the network to tailor mechanical responses is assessed and the results are validated by real-world testing. Success of the model provides evidence of the benefits that may be gained by harnessing the power of artificial intelligence to tailor mechanical performance in engineering design. In future works design parameters and the developed network architecture may be substituted for application to full structural components.

4.1 Base Material Properties Characterisation

The metamaterial lattice samples were additively manufactured using fused deposition modelling (FDM). Additively manufactured materials are susceptible to non-isotropic behaviour and material data from manufacturer's datasheets must be validated experimentally to ensure an accurate FEA simulation. Mechanical properties are likely to vary depending on the FDM equipment and printing parameters such as nozzle, bed and chamber temperature, layer height and build speed. Properties may vary between different filament batches. Nylon filaments are particularly sensitive to humidity and moisture content, as they are hygroscopic and absorb moisture from the air at room temperature. The FDM process may cause build-up of residual stresses between layers as heat transfers from the melt pool at the nozzle tip to cooler layers below that have already solidified. Significant thermal gradients can build up from depositing geometries of large thicknesses, leading to warping. In most cases manufacturers do not provide full non-linear material data so experimental measurements were required to build the non-linear FEA model for predicting behaviour of the metamaterial lattice beyond yield.

4.1.1 Experimental Setup

Sample Materials

Material data was collected for comparison between 5 different FDM materials; ColorFabb XT-CF20 copolyester with 20% carbon fibre content, Esun ePA nylon, Esun ePA-CF nylon with chopped carbon fibre reinforcement, Fiberthree F3 PA-CF Pro carbon fibre reinforced nylon and Polymaker CoPA nylon. Mechanical characterisations were performed under uniaxial tension and compression according to ASTM standards D638 and D695 respectively. Tensile specimens were of the type I dog bone geometry displayed in Figure 70. The specimens were built in the z-direction at a layer height of 0.15mm, with the infill lines alternating between ±45° with respect to the direction of loading, as shown in Figure 71. This was selected to provide an approximation of isotropic material behaviour for comparison and selection of the best performing material, however was not used in the FEA model.

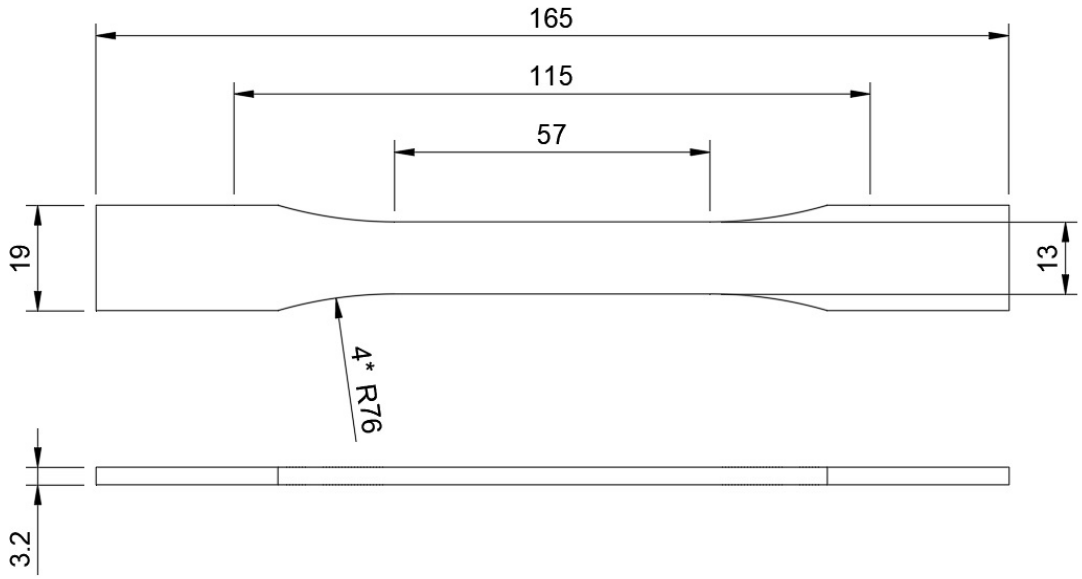


Figure 70 Tensile test coupon geometry. All dimensions in mm

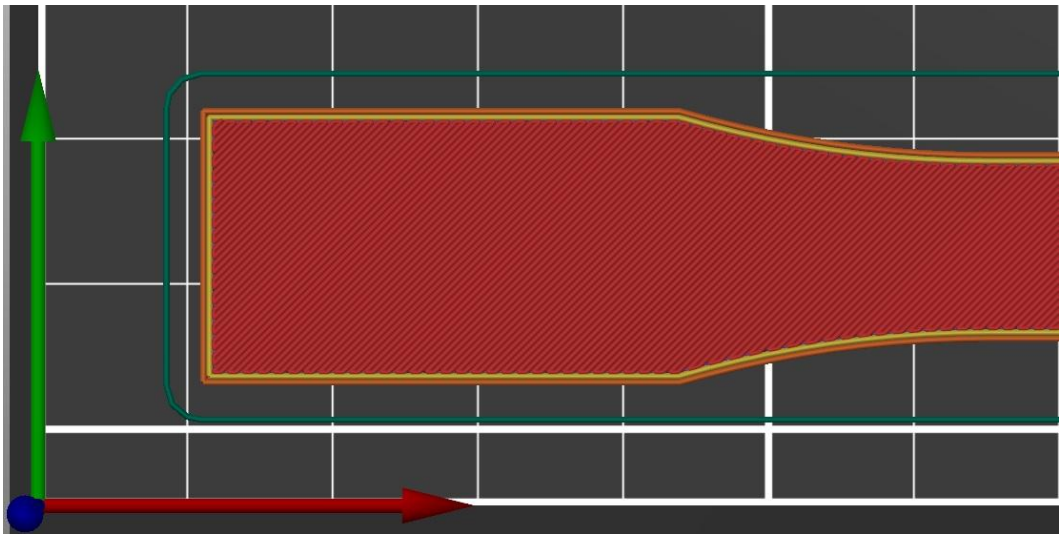


Figure 71 Tensile test coupon layer orientation

For compression tests cylindrical coupons of ½ inch diameter (12.7mm) and 1 inch height (25.4mm) were built at 0.15mm layer height in the z-direction. In addition, for the best performing material cuboid specimens of ½ inch width and depth and 1 inch height were built at 0, 45 and 90° to the loading direction. The FEA model was validated using each set of material data and the effects of build direction on the material properties were assessed. For both tensile and compression in each case tests were repeated for 5 specimens and a mean taken of the results. Tensile coupons for each sample of 5 were printed in the same build, nested in line at a spacing of 25mm to allow for variations in material properties across the build plate due to temperature, whereas compression coupons were positioned at the corners and centre of a 100 x 100mm square, as shown in Figure 72.

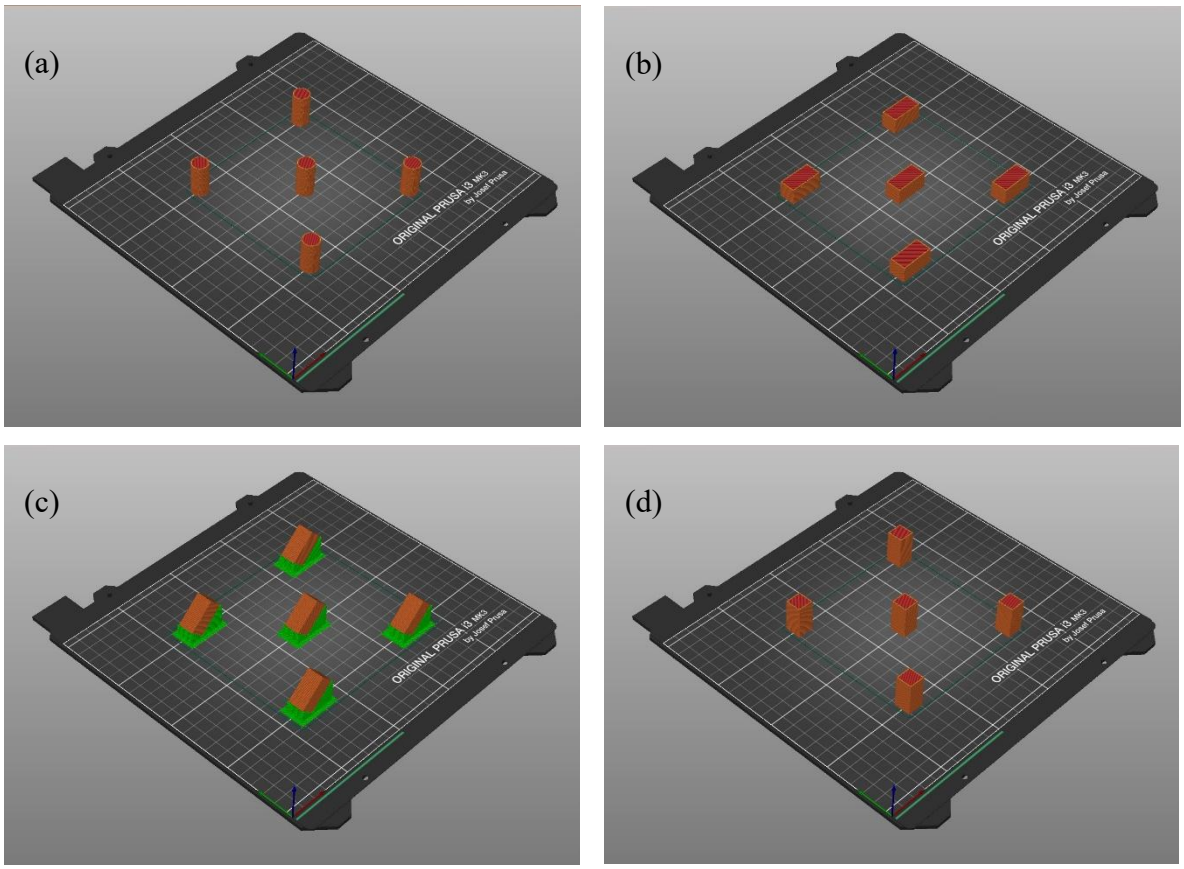


Figure 72 Compression coupon nested build plates (a) 90° cylindrical (b) 0° cuboid (c) 45° cuboid (d) 90° cuboid

Apparatus

The specimens were printed on a Prusa i3 MK3S machine that was modified for the project with an aftermarket enclosure, temperature controller and filtration system from PrusaBox to allow printing in nylon. The setup used is shown in Figure 73. The standard 0.4mm brass nozzle was upgraded to a stainless steel Prusa Nozzle X to allow for printing in composite filaments. A multi material unit with filament buffer was fitted and incorporated into the enclosure for printing with water soluble support structures to allow for more complex lattice geometries. The build g-code files were generated using PrusaSlicer with custom filament profiles tuned using test prints and trial and error. The print settings for each material can be found in Table 6. For the nylon filaments an Eibos Cyclopes filament drier was used to dry the filament spools at 70°C for 24 hours before printing. The filament was also held at 70°C throughout each build by printing directly from the drier. After printing the nylon spools were double vacuum bagged with silica gel desiccant to prevent moisture ingress. The specimens were immediately post treated by annealing at 70°C for 2 hours to slow down recrystallisation and limit residual stresses. They were then conditioned at room temperature and humidity for a minimum of 48 hours before testing.

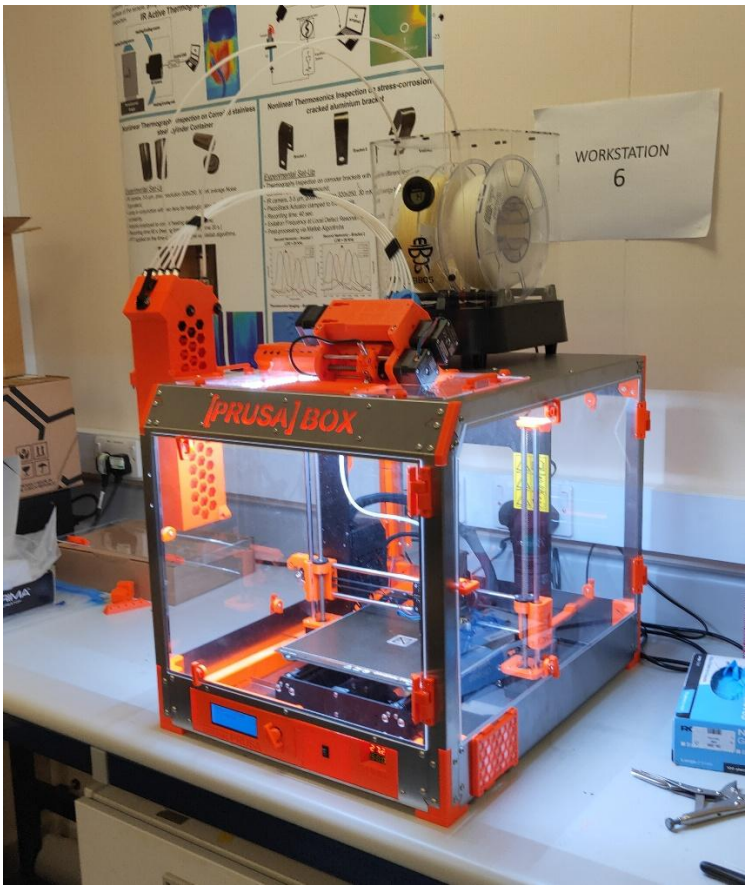


Figure 73 FDM equipment used for test specimen manufacture

Table 6 FDM slicer settings

	Parameter		Polymaker CoPA	Fiberthree F3 PA-CF Pro	Esun ePA	Esun ePA - CF	ColorFabb XT-CF20
Print Settings	Layer Height (mm)		0.15	0.15	0.15	0.15	0.15
	Minimum Perimeters		2	2	2	2	2
	Infill Density (%)		100	100	100	100	100
	Infill Pattern		Rectilinear	Rectilinear	Rectilinear	Rectilinear	Rectilinear
	Print Speed (mm/s)	Perimeters	45	45	45	45	45
		External	25	25	25	25	25
		Infill	80	80	80	80	80
		Bridges	25	25	25	25	25
Temperature	Nozzle Temperature (°C)	First Layer	265	285	270	270	260
		Other Layers	265	285	270	270	260
	Bed Temperature (°C)	First Layer	70	90	85	85	90
		Other Layers	70	90	90	90	90
Cooling	Fan Speed (%)	Min	0	0	30	30	30
		Max	0	0	60	60	50
	Bridges Fan Speed (%)		30	30	60	60	50
Retraction	Length (mm)		2	2	2	2	disabled
	Speed (mm/s)		40	40	40	40	disabled
Build Surface	Prusa Build Plate Used		PA Nylon Sheet	PA Nylon Sheet	PA Nylon Sheet	PA Nylon Sheet	Smooth Sheet

Mechanical testing was carried out on an Instron 3369 dual column electromechanical tension and compression system with a 50kN load cell, shown in Figure 74. To measure the Young's modulus a clip on contact extensometer was used. A set of digital callipers was used to measure the as-printed sample geometry and set the gauge length of the extensometer.



Figure 74 Instron 3369 uniaxial tensile (left) and compressive (right) experimental setup

Procedure – Tensile

Firstly the gauge width and thickness of the narrow section of the dog bone specimens were measured using the digital callipers to the nearest 0.01mm and the cross-sectional area was calculated. Measurements were taken at three different locations for both width and thickness and mean values were calculated. The sample measurements can be found in Appendix C. The specimens were marked to indicate positions for placement of the jaws and extensometer and tightened into the grips of the machine with the distance between the jaws set at 115mm. The gauge length of the extensometer was set at 50mm using the digital callipers and temporarily locked by holding the thumb buttons while it was attached. The markers on the specimen were used to ensure the extensometer was centred symmetrically on the sample while the digital callipers set the gauge length with greater accuracy before attachment to the specimen. A small preload of 5N was then applied at a speed of 0.1mm/min to eliminate the effects of any bending or warping of the specimen. The extensometer and load cell were then

rebalanced and zeroed before loading of the specimen at a constant rate of 5mm/min. The loading force and displacement were recorded digitally at a rate of 50 datapoints per second. The extensometer was removed from the specimen on reaching the limit of its measurement range and then the test was continued until fracture.

Procedure - Compressive

The average length and diameter (or length, width and thickness) of each specimen were measured to the nearest 0.01mm. These measurements can be found in Appendix D. The specimen was placed between the compression plates with its central axis aligned below the centre of the compression tool. The crosshead was jogged until it was a paper width thickness (approximately 0.1mm) above the top of the specimen and the machine was zeroed. For each successive sample the crosshead was automatically moved back to the identical starting position using the machine's return function. The load was balanced and then a compressive load was ramped on at a constant rate of 1.3mm/min until either the machine reached 80% of the maximum capacity of the load cell or rupture of the specimen occurred. The complete load-deformation curve was recorded at a rate of 50 datapoints per second.

4.1.2 FDM Material Comparison

Tension

For each specimen the force-displacement data was converted to a stress-strain characteristic using the measured cross-sectional area and gauge length. For each material the 5 specimens were then averaged to produce a mean response curve. Extensometer readings were used to calculate the Young's modulus and yield strength approximated by the 0.2% proof stress. Load cell readings were used to calculate the ultimate tensile strength, stress at fracture and percentage elongation at fracture. Full results for all specimens can be found in Appendix E. Comparisons of the average stress-strain responses recorded by the extensometer and load cell for every material are displayed below.

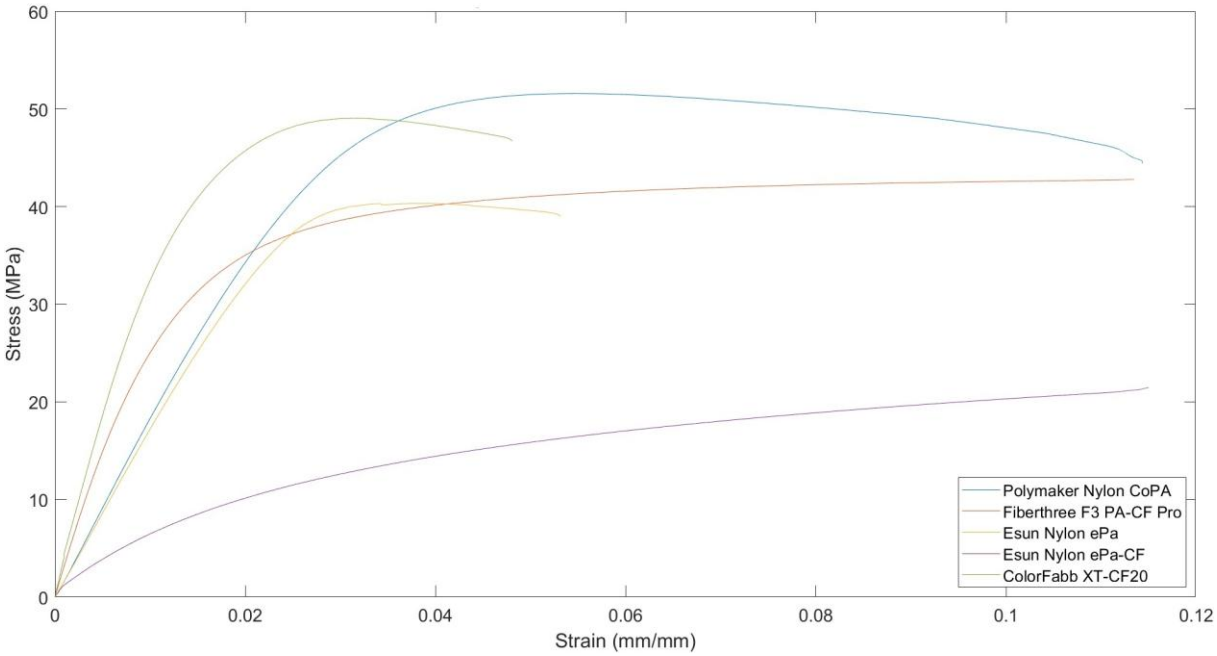


Figure 75 Material comparison - Tension ASTM D638 - Extensometer

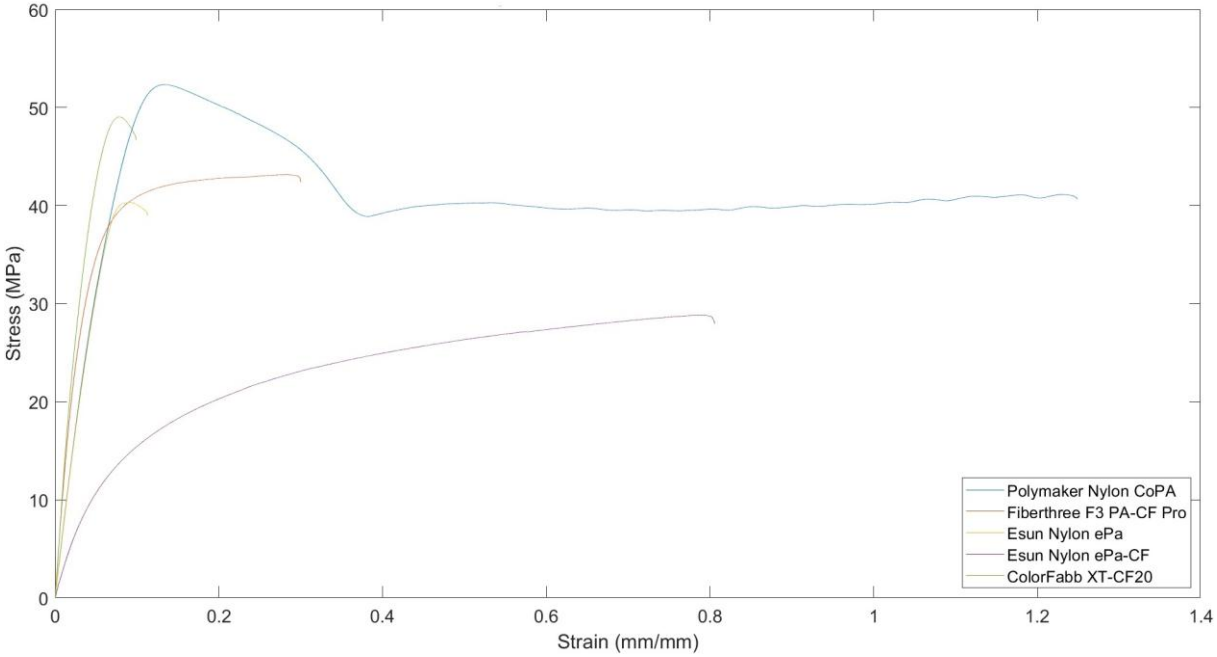


Figure 76 Material comparison - Tension ASTM D638 - Load cell

Compression

The mean stress-strain response was calculated for each material and compensation was carried out to correct for the toe artifacts generated by the take up of slack at the start of each test as the initial small gap between the sample and compression plates was closed. The linear region of the curve was continued through to the zero-stress axis to find the corrected zero strain point. The data was translated and the elastic modulus of the linear region and tangent modulus of the post-yield region were calculated in addition to the yield strength as the 0.2% proof stress. For samples that ruptured the ultimate strength, stress at fracture and percentage elongation at fracture were also recorded. Results for every specimen can be found in Appendix F and the mean responses are compared for every material in Figure 77.

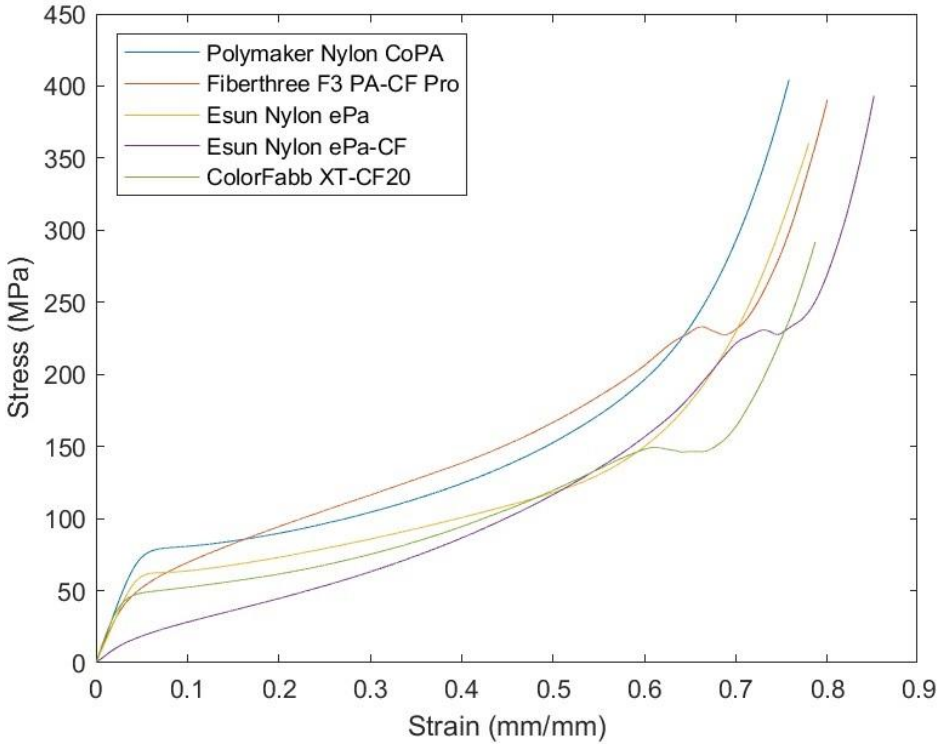


Figure 77 Material comparison - Compression ASTM D695

Material Properties Comparison

Table 8 demonstrates the importance of carrying out mechanical testing to measure material properties using the exact hardware and build settings that will be used to print the metamaterials rather than using values from the manufacturer's datasheets. In all cases strength and stiffness were below the quoted typical values. It should be noted that alternate testing methods were used by the manufacturers, as shown in Table 9. Mechanical properties are likely to vary between different test coupon geometries and at different load rates. Polymaker CoPA gave the closest values to the experimental data for tensile strength and

Young's modulus, though the ductility was greatly higher than expected. The nylon samples showed significant variations in percentage elongation at fracture from expected values. This is likely because the ductility of nylon can be strongly affected by the amount of moisture content absorbed from the air during conditioning. ColorFabb XT-CF20, being a non-hygroscopic material showed considerably lower deviation from the expected value.

Table 7 FDM tensile and compressive material properties comparison

		Polymaker CoPA	Fiberthree F3 PA-CF Pro	Esun ePA	Esun ePA - CF	ColorFabb XT-CF20
Tension	Young's Modulus (MPa)	1785.89	2835.23	1655.98	576.96	3470.66
	Yield Strength (MPa)	39.71	27.12	36.36	8.84	36.93
	Ultimate Strength (MPa)	51.58	43.09	40.30	28.81	49.05
	Stress At Fracture (MPa)	40.65	42.41	38.97	27.96	46.71
	% Elongation At Fracture	124.85	30.00	5.32	80.57	4.81
Compression	Young's Modulus (MPa)	1705.92	1501.67	1434.65	447.79	1832.32
	Yield Strength (MPa)	63.21	34.81	52.38	11.89	32.70
	Ultimate Strength (MPa)	N/A	233.11	N/A	230.45	149.39
	Stress At Fracture (MPa)	N/A	227.43	N/A	227.93	146.05
	% Elongation At Fracture	N/A	69.33	N/A	74.26	64.23

Table 8 Comparison of tensile material properties with manufacturer datasheet values

		Polymaker CoPA	Fiberthree F3 PA-CF Pro	Esun ePA	Esun ePA-CF	ColorFabb XT-CF20
Ultimate Tensile Strength (MPa)	Measured Mean Value	51.58	43.09	40.30	28.81	49.05
	Manufacturer Value	66.20	85.00	52.45	140.00	76.00
	% Deviation	22.08	49.31	23.16	79.42	35.46
Young's Modulus (MPa)	Measured Mean Value	1785.89	2835.23	1655.98	576.96	3470.66
	Manufacturer Value	2223.00	10500.00	N/A	N/A	N/A
	% Deviation	19.66	73.00	N/A	N/A	N/A
Elongation At Fracture (%)	Measured Mean Value	124.85	30.00	5.32	80.57	4.81
	Manufacturer Value	9.90	0.98	175.32	10.61	7.50
	% Deviation	1161.11	2961.22	96.97	659.38	35.87

Table 9 Manufacturer testing methods

Material	Testing Method
Polymaker CoPA	ASTM D638 (ISO 527, GB/T 1040)
Fiberthree F3 PA-CF Pro	DIN EN ISO 527-2 Typ 1A
Esun ePA	GB/T 1040
Esun ePA – CF	GB/T 1040
ColorFabb XT-CF20	ISO 527

Of the materials compared in Table 7 Polymaker CoPA showed the greatest strength in both tension and compression and the highest compressive stiffness. Under tension the composite ColorFabb and Fiberthree filaments with carbon fibre reinforcement showed the greatest stiffness. However the composite filaments all failed by eventual rupture in compression, whereas the nylon filaments without reinforcement continued to deform until densification. The Esun ePA-CF showed highly non-linear material properties and reduced strength and stiffness compared to the unreinforced ePA, though offered an increase in ductility of 75.25% strain before failure during tension. Polymaker CoPA provided the highest overall ductility under tension, showing an extended necking region where the material cross section continued to narrow without fracture. The response curve levels off into a plateau that remains stable across a displacement equivalent to approximately 80% of the original length. Polymaker CoPA was selected as the most suitable material for use in the metamaterial lattice compression block optimisations. In addition to high strength and stiffness it showed good printability with reduced warping and stringing compared to alternative nylons.

Build Direction Comparison

Polymaker CoPA specimens with square cross sections were built at 0, 45 and 90° to the loading direction and tested to assess the effects of build direction on the mechanical properties under compression. The stress-strain responses are compared against the previous sample with circular cross sections in Figure 78. The range of deviations in stiffness and strength between the different build directions are displayed in Table 10.

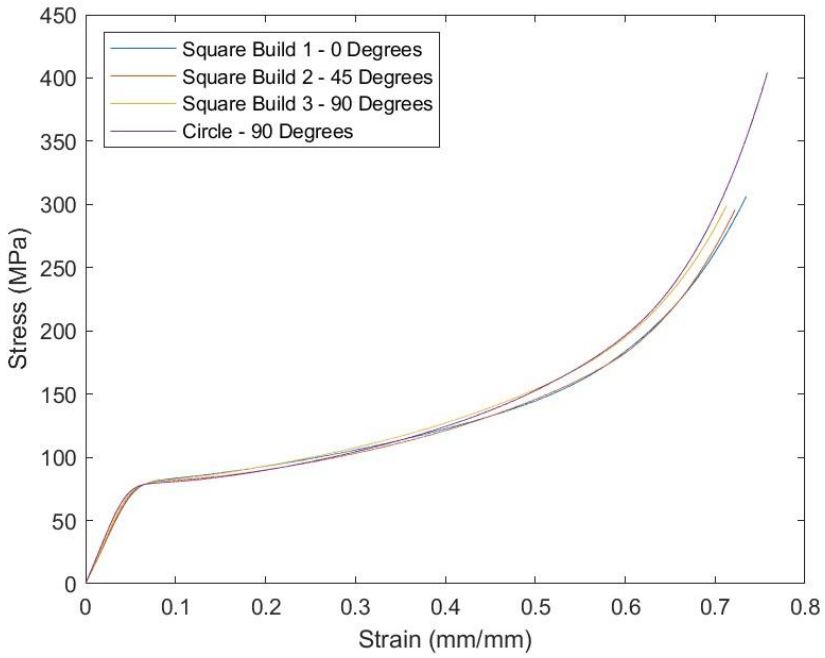


Figure 78 Build direction comparison of Polymaker CoPA compression samples

Table 10 Deviations in stiffness and strength with build direction of Polymaker CoPA compression samples

	Build 1 - 0°	Build 2 - 45°	Build 3 - 90°	Mean	Standard Deviation	Max	Min	Range
Young's Modulus (MPa)	1485.14	1501.07	1570.46	1518.89	45.37	1570.46	1485.14	85.32
Yield Strength (MPa)	66.55	67.41	66.31	66.76	0.58	67.41	66.31	1.10

The sample built at 90° to the loading direction showed the greatest compressive stiffness, likely as the FDM layers are forced together as the material deforms whereas 0 and 45° orientations may cause a small increase in relative motion between layers. However the largest difference in stiffness remained below 6% between the different build directions, indicating an assumption of near-isotropic behaviour is likely to be valid for the purposes of the non-linear FEA model. Yield strength remained consistent between build directions to within 98% of the maximum value. The engineering stress-strain responses for all build directions were converted to true stress-strain characteristics for use in the FEA model and the amount of plastic deformation at each point was calculated using the following formulae:

$$\sigma = \sigma_{eng}\left(1 + \varepsilon_{eng}\right) \tag{4.1}$$

$$\varepsilon = \ln\left(1 + \varepsilon_{eng}\right) \tag{4.2}$$

$$\varepsilon_{plastic} = \varepsilon - \sigma/E \tag{4.3}$$

The non-linear region of the curve of true stress against logarithmic strain was extracted and down sampled to 100 datapoints for input into an ANSYS non-linear isotropic material model. ANSYS requires the Young's modulus and yield strength to be input separately and the plastic region material data to be provided with the first datapoint occurring at the yield point. On transitioning past yield the solver uses the plastic region data as a lookup table to predict the non-linear mechanical response. Plots of the true stress against logarithmic strain curves that were used for direct input into ANSYS for each build direction can be found in Appendix G.

4.1.3 AM Manufacturability Tests

In order to determine the limits of the design space for dataset generation, a series of test prints were carried out to find the minimum strut radius that can be successfully printed using the chosen FDM material. A 20x20x20mm unit cell of the octet-framed truss lattice was printed at a layer height of 0.15mm, starting at a strut radius of 1mm. Successive test prints were then carried out reducing the strut radius by 0.1mm each time until the build failed. The

experiments were repeated for a variety of nozzle and bed temperatures, fan and print speed settings. The settings that gave the results with the best quality for each material are those recorded in Section 4.1.1. At very small strut radii a greater degree of stringing was observed. This depended on how wet the filament was. Print quality was significantly improved by drying the filament before printing. Build failure was largely due to deflection of the struts away from the nozzle during print moves so that new material was deposited into empty space. The prints typically failed completely while attempting to bridge the horizontal sections of the central tetrahedrons. However successful builds were possible down to 0.5mm without need for support structures. At lower radii soluble support material is recommended. A minimum strut radius of 0.5mm was set for the design space of the deep learning model to ensure any optimised results would be manufacturable.

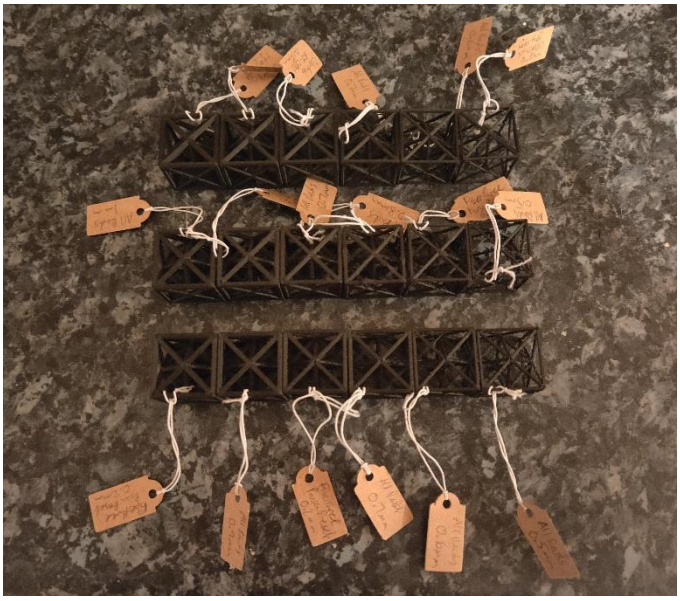


Figure 79 Minimum strut radius test prints

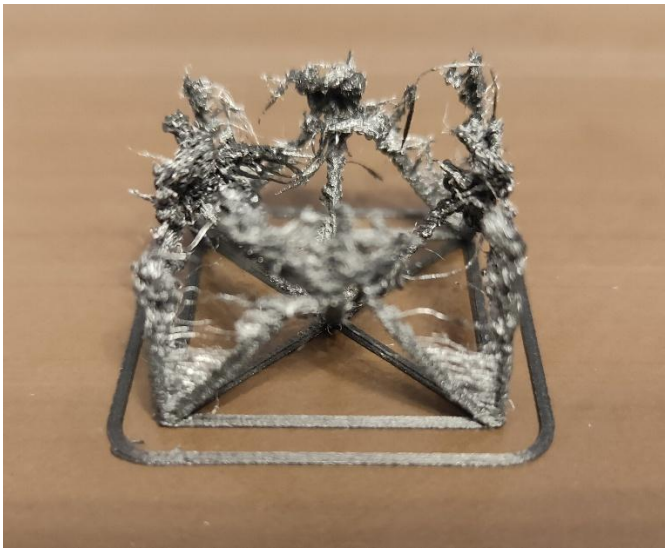


Figure 80 Example of a failed build at strut radius below 0.5mm

4.2 FEA Dataset Generation

4.2.1 Non-linear FEA Model

For generation of datasets of non-linear mechanical responses of the lattice compression blocks several parametric FEA models were created in ANSYS to assess the best approach. The first modelled the lattice geometry with 3-dimensional solid elements. The solid model was extended to include a non-linear buckling analysis, taking into account the deformation mode shapes due to Euler buckling of the struts. The lattice was then remodelled as line bodies using a slender beam assumption for the struts and non-linear stabilisation was applied to aid convergence and improve the robustness of the model as the geometry was updated across the range of parameters in the dataset. For all models a non-linear isotropic material model was used with the compressive Polymaker CoPA material data obtained in section 4.1. Simulations were repeated with the material data measured from the coupons built at each orientation to examine the influence on the lattice response. A 20x20x60mm octet-framed truss lattice featuring 3 unit cells of 20x20x20mm was created in ANSYS design modeller. Parameters r1, r2 and r3 were assigned to 3 key strut radii; the central octahedron, the connecting tetrahedrons, and the outer frame as shown in Figure 81. Each radius was allowed to vary between 0.5 and 3mm. The parametric CAD model implementation in ANSYS allowed these parameters to be updated from a table of design points and the analysis repeated automatically to generate the full dataset of responses. The setup of the solid and line body models are displayed below.

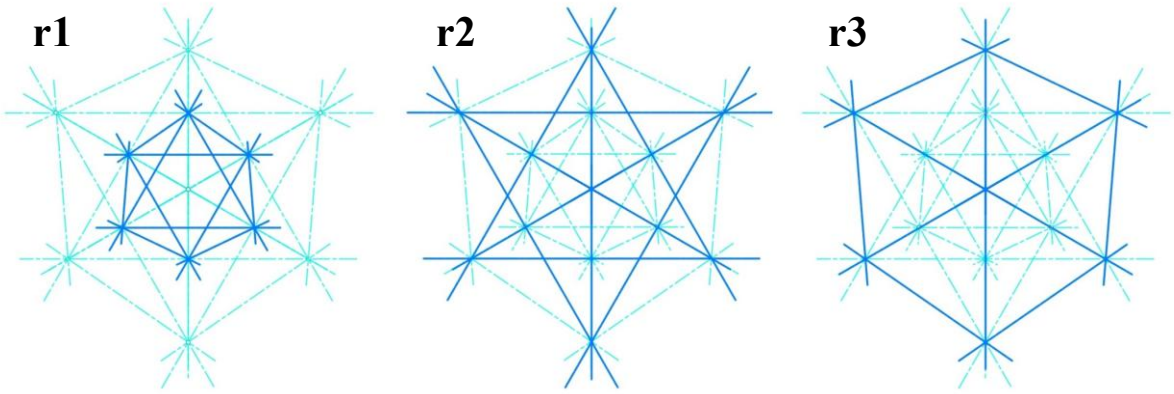


Figure 81 Lattice design optimisation parameters

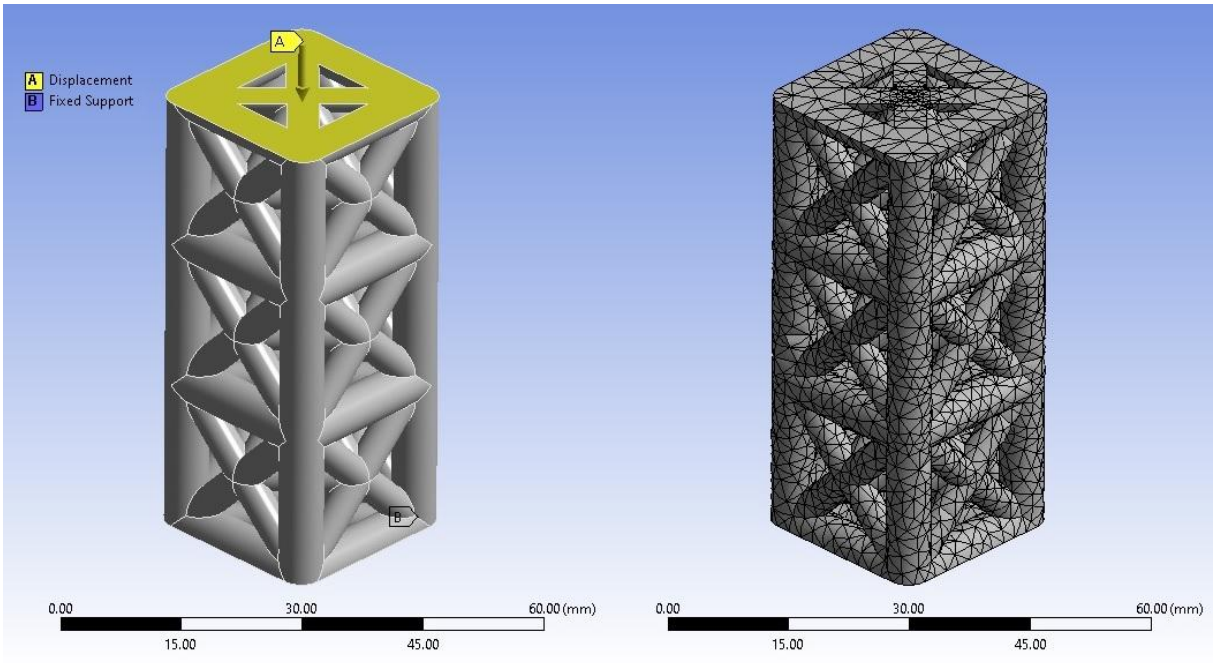


Figure 82 Solid body FEA model example boundary conditions (left) and mesh (right)

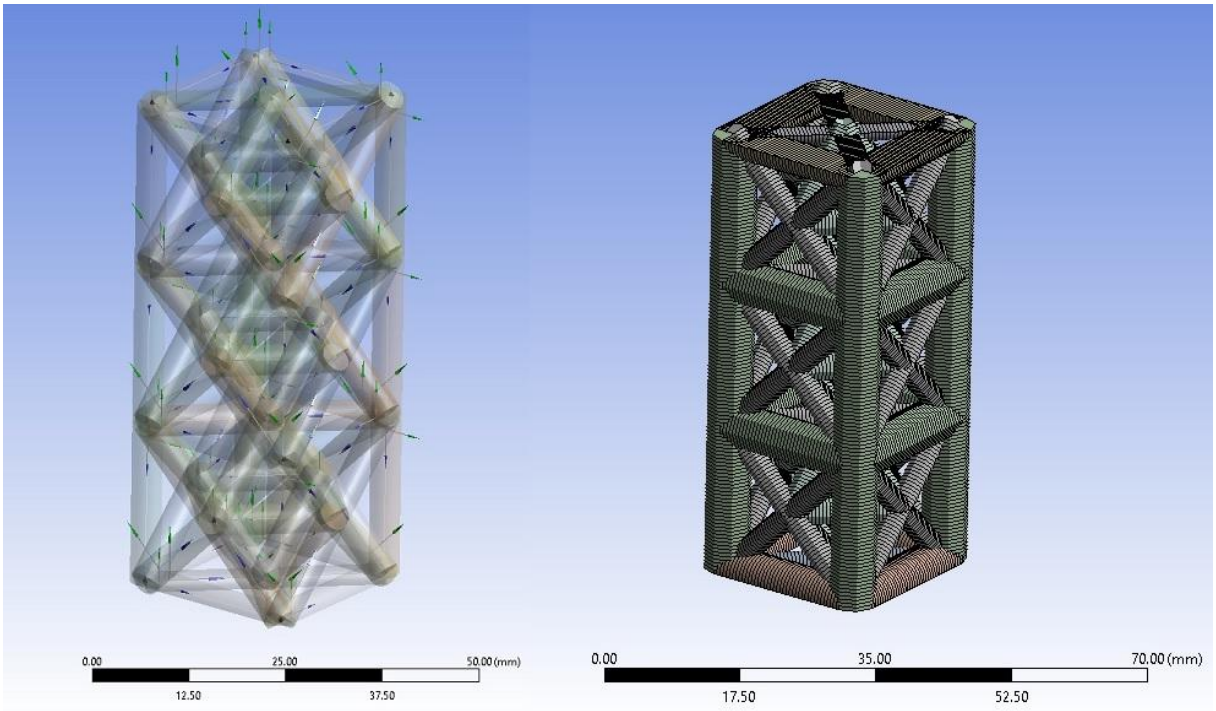


Figure 83 Line body FEA model example geometry (left) and mesh (right)

A fixed support was applied to the bottom face of the geometry and a 10mm displacement load was applied to the top in the negative z-direction, ramped on in 10 equally sized load steps of 1mm each. The solver was set to initially use 30 substeps, with a maximum of 100 and a minimum of 10. The programme then carried out a Newton-Raphson iteration process using the lowest number of substeps required for convergence of the internal forces within the structure. Large deflections were enabled to allow simulation of the non-linear post yield behaviour. The reaction force in the z-direction was probed at the fixed boundary condition. To generate the entire force-displacement curve for automatic output into the table of design points the probes were repeated at time intervals equivalent to 0.1mm intervals of displacement between 0 and 10mm. The solid body geometry was meshed with an initial element size of 2mm. On reducing the element size further the difference in the response curves was negligible and the additional time necessary to run the simulation was significant. On an example study of a lattice with all struts set to 2mm radius, reducing the default element size from 2mm to 1mm resulted in an increase in the number of elements from 13,766 to 61,100 and an increase in solve time by a factor of 4 from 2min 35sec to 11min 2sec. Over the course of generation of a 2000 sample dataset this would be equivalent to an increase from 3.6 days to 15.3 days, with larger datasets taking months rather than weeks. This additional time was unacceptable considering the change in response was minimal. A comparison between the response curves can be found in Appendix H. For the line body model the optimal trade-off between solve time and model fidelity was found at 0.5mm element divisions. For future studies an alternative approach is to use an adaptive mesh, however this would greatly increase the solve time as a mesh convergence study is carried out automatically for each design point. This may be practical for further works where the dataset is generated using high performance computing. For the current study running the solver on an intel i7-10700K processor at 3.8GHz with 32GB RAM was sufficient, returning datasets of 1392 samples for the solid body model and 4297 samples for the line body model within 2 weeks.

In addition to the solid body model a contact analysis was undertaken to compare the effects of modelling the compression plates of the Instron machine. The test setup is shown in Figure 84. The interface between the plates and the lattice sample was first modelled as bonded contacts and then the study was repeated with frictional contacts. A coefficient of friction of 0.4 was used to represent the resistance between nylon and steel. To study the effects of including buckling analysis in the simulation the solid body model was subjected to a unit load and an eigenvalue buckling study was used to extract the first buckling mode shape in addition to the load multiplier necessary to initiate buckling within the lattice, equivalent to

the Euler buckling load. The deformed mode shape of the solution was then used as the initial geometry for the full non-linear analysis, loading up to 10mm displacement with the same boundary conditions and solver settings as before. The advantage of this method is that a displacement scaling factor is used to set the amount that the lattice should be perturbed from its initial shape due to buckling at the start of the simulation. This allows the FEA model to be fine-tuned to produce more realistic responses as close as possible to the real-world experimental results. The project schematic of the buckling model implemented in ANSYS is displayed in Figure 85, with an example of perturbed starting geometry generated by the eigenvalue buckling analysis shown in Figure 86.

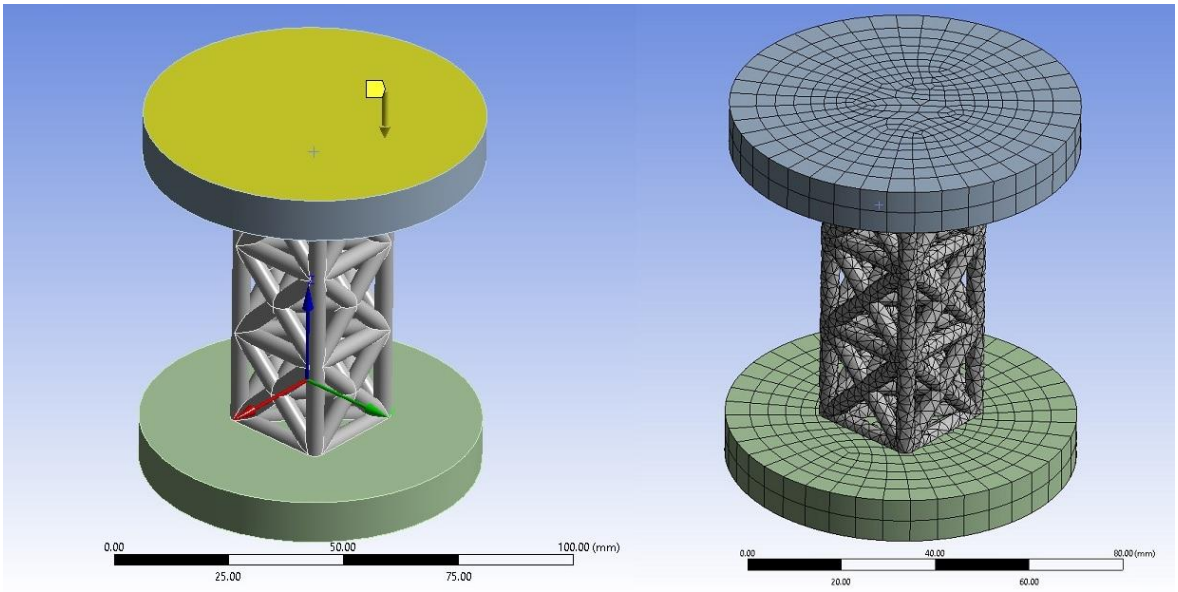


Figure 84 Frictional contact FEA model example boundary conditions (left) and mesh (right)

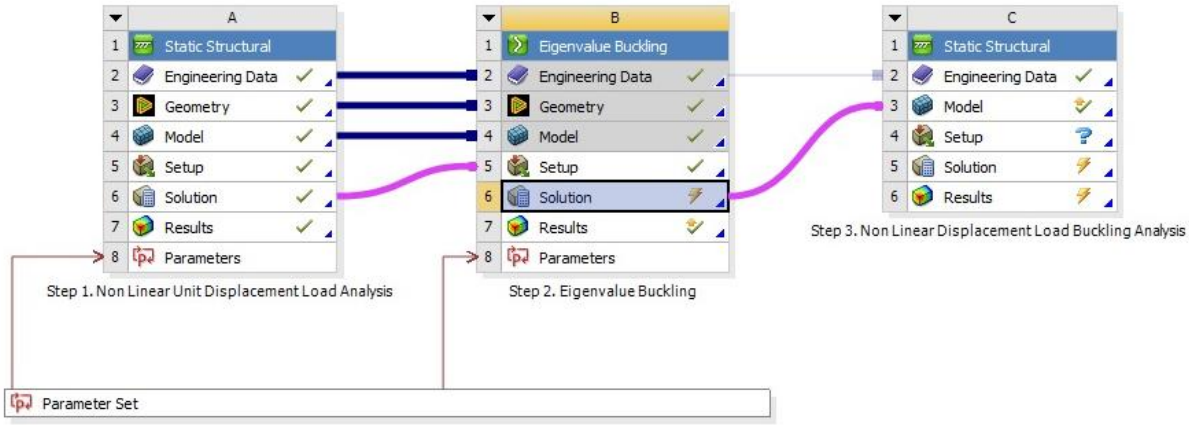


Figure 85 Non-linear buckling model project schematic implemented in ANSYS

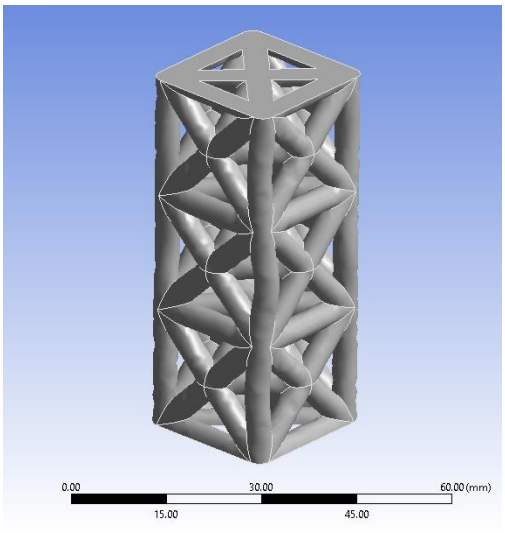


Figure 86 Example of perturbed starting geometry in non-linear buckling analysis

4.2.2 Experimental Validation

The non-linear FEA models were validated by Instron testing of lattice compression block samples with all strut radii set at 1mm, 2mm and 3mm. For each sample 3 specimens were tested for comparison against the FEA results. The same procedure for compressive testing as outlined in Section 4.1.1 was applied, loading the specimens at a rate of 1.3mm/min until fracture. The mechanical responses were compensated for toe artifacts as before and the deformation mode shape of the lattice was recorded using a digital camera. The experimental test set up is shown in Figure 87.



Figure 87 Lattice compression testing experimental setup

The mean results averaged across the 3 specimens for each experiment are plotted in Figure 88, which demonstrates the limits of the design space envelope that can be explored in the study. The results indicated a large amount of variation in the design space, suggesting a successfully trained model would likely be highly tuneable, able to provide responses from 500N up to 15-20kN failure load. Figure 89 shows the resulting geometries may have masses anywhere between 1.5 and 35grams. The FEA results generated using material models from data obtained at each build direction are compared against the experimental results in Figures 90, 91 and 93 for the solid body model and Figures 94-96 for the line body model below. Any differences between build directions remains small and for the dataset generation and all further results the material data collected from the circular coupons built at 90° to the loading direction is used. In all cases the standard solid body model without contact or buckling analysis overestimates stiffness and strength, likely as it does not account for bending occurring below the yield stress as a result of Euler buckling of the struts. However the overall deformation mode shape is captured with good accuracy, as shown in Figure 92. The line body model meshed with beam elements offers predictions closer to experimental values, particularly within the mid-range of the design space. At 2mm strut radius the stiffness, yield load and post buckling behaviour are all modelled with greater accuracy. When the strut radius is reduced to 1mm the FEA predictions show the same stiffness between solid and beam models, however the beam elements correct an overshoot in failure load, predicting the yield point of the lattice material within the range of the true experimental values. However the results suggest the beam element model may become less accurate at higher strut radius where the slender beam assumption becomes invalid. At 3mm radius the stiffness becomes underestimated, while the yield load and slope of the post buckling region remains in relative agreement with the real-world values. This must be taken into account while tuning the responses using the deep learning model. If the requests lie close to the limits of the design space additional compensations may need to be applied and thorough post validation of the optimised geometry is necessary.

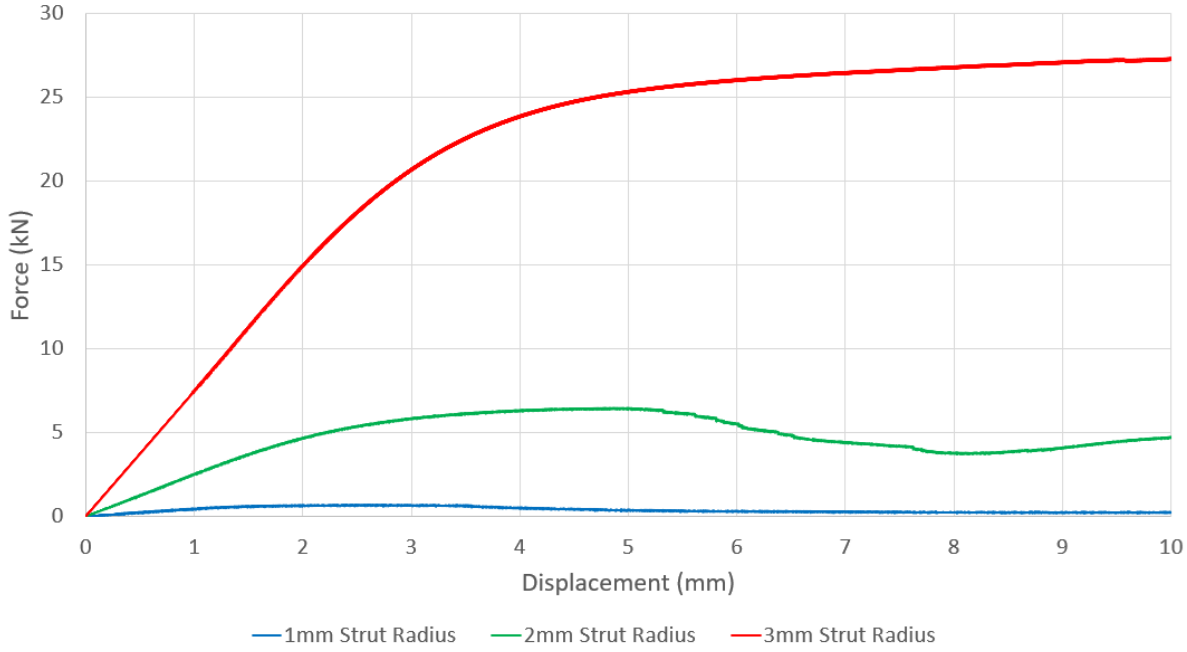


Figure 88 Mean experimental responses of lattice blocks loaded in uniaxial compression

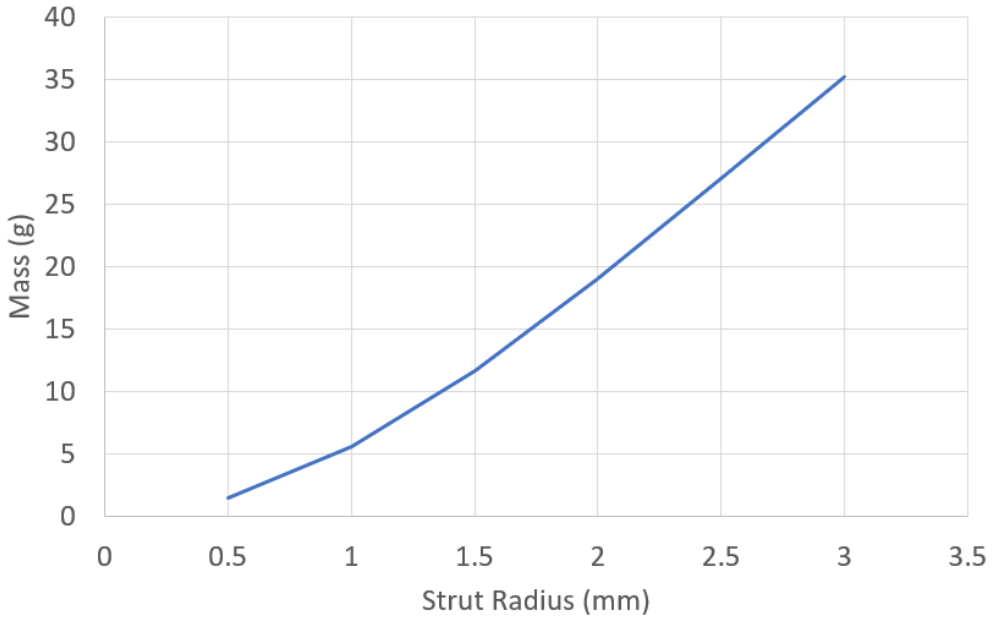


Figure 89 Mass vs strut radius of 20x20x60mm octet-framed truss compression block

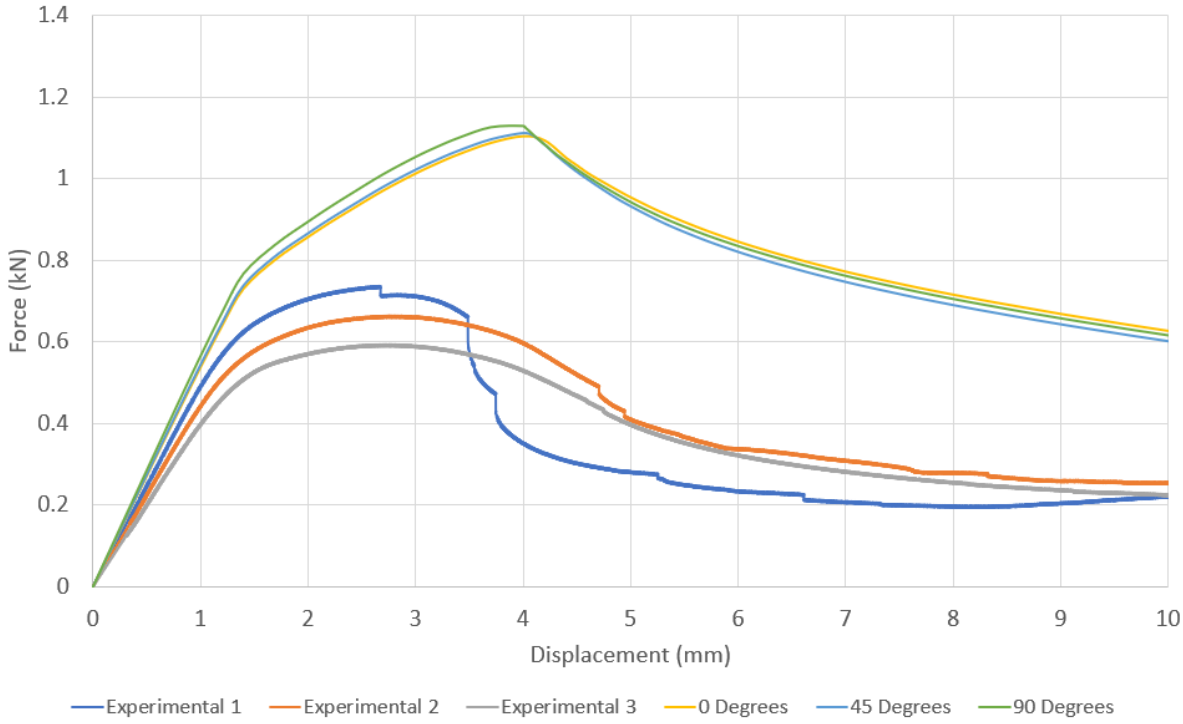


Figure 90 Experimental response vs solid body FEA model of lattice compression block - 1mm strut radius

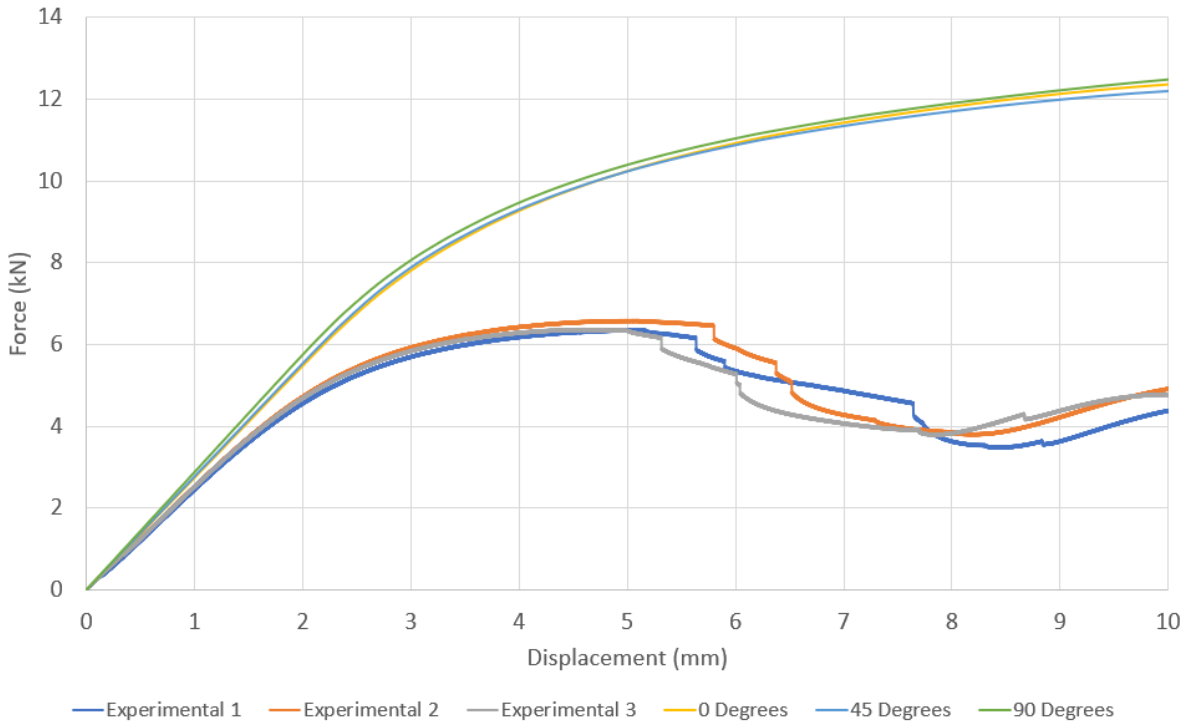


Figure 91 Experimental response vs solid body FEA model of lattice compression block - 2mm strut radius

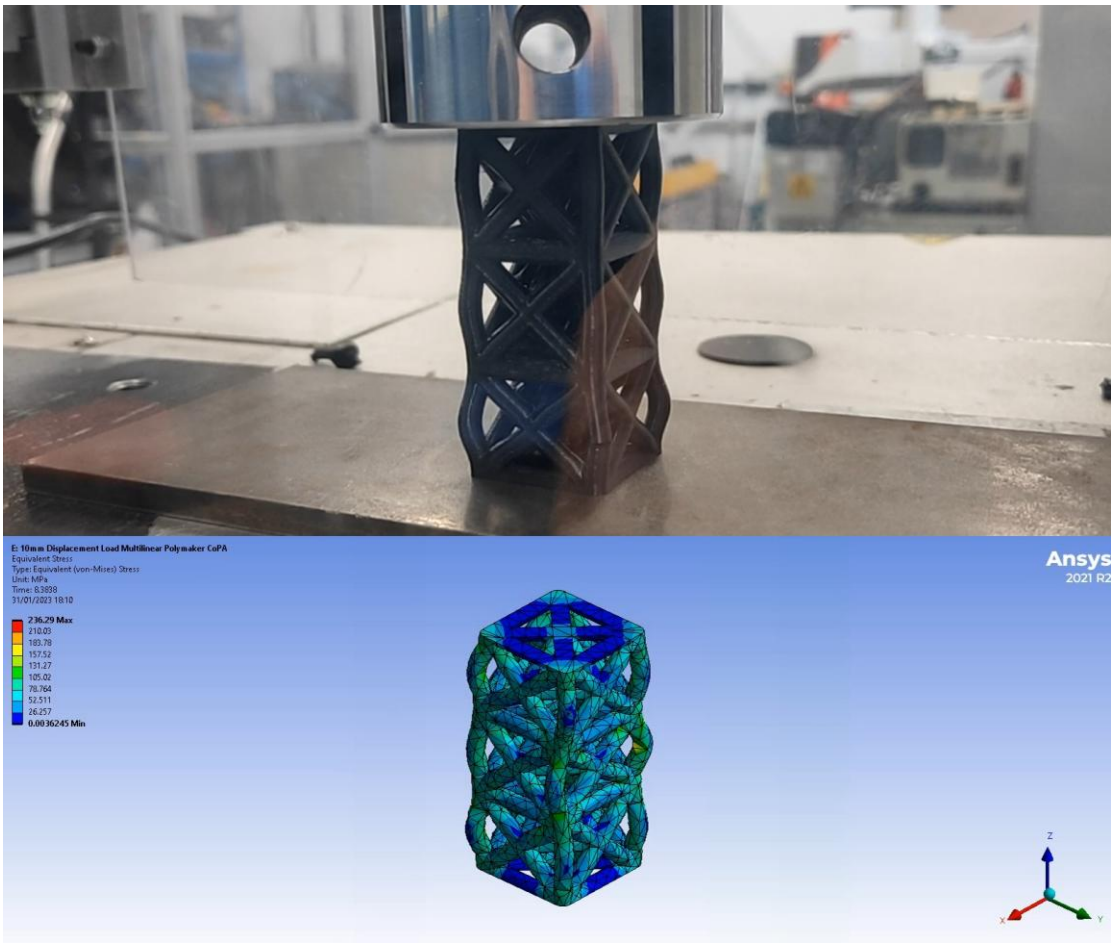


Figure 92 Experimental vs solid body FEA deformation mode shape for lattice compression block - 2mm strut radius

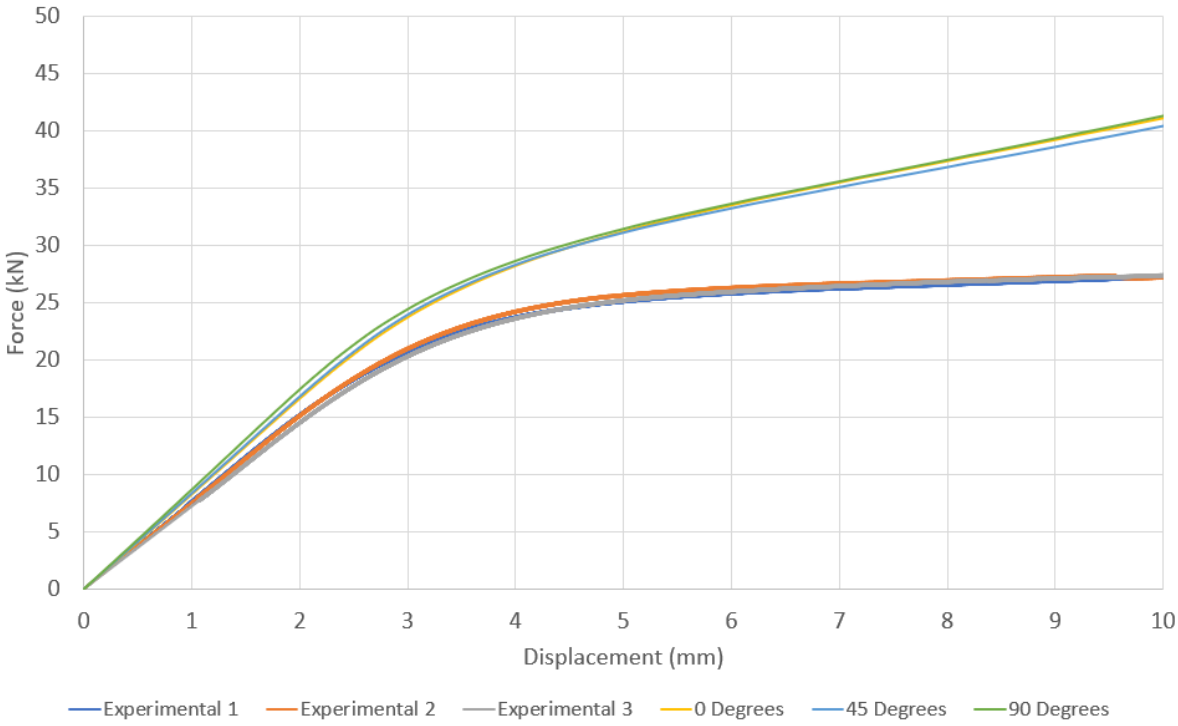


Figure 93 Experimental response vs solid body FEA model of lattice compression block - 3mm strut radius

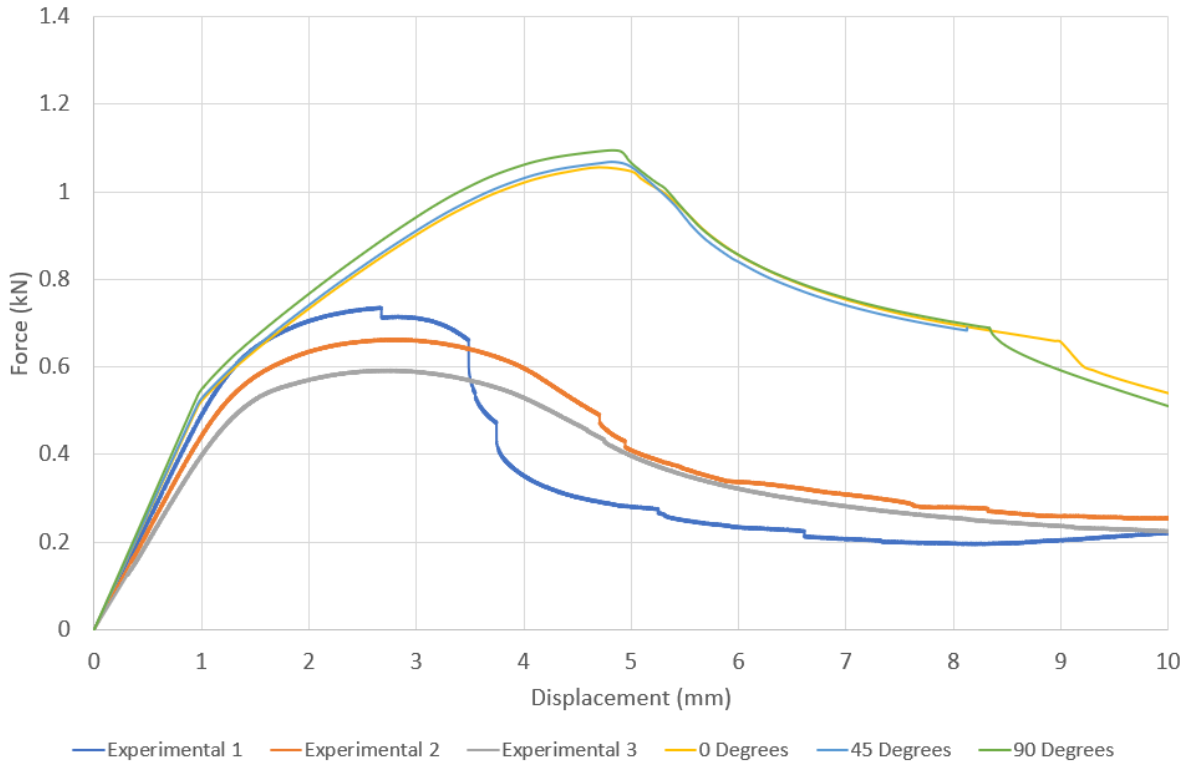


Figure 94 Experimental response vs line body FEA model of lattice compression block - 1mm strut radius

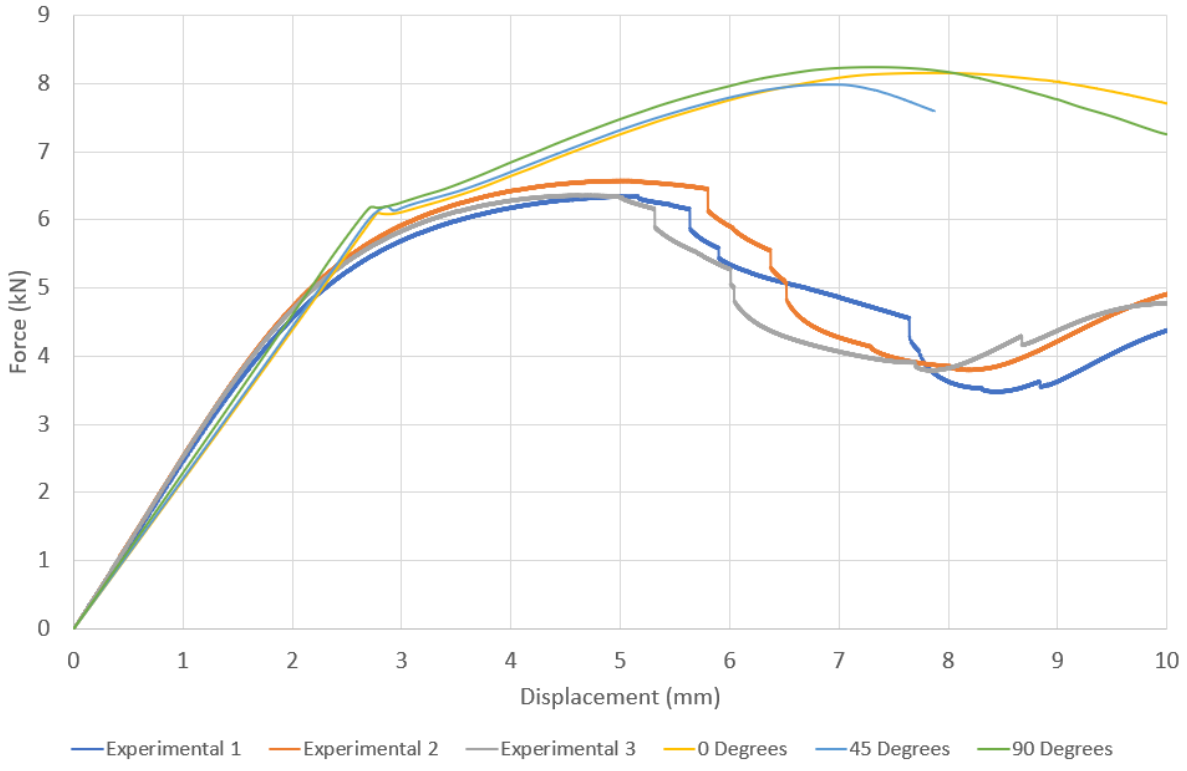


Figure 95 Experimental response vs line body FEA model of lattice compression block - 2mm strut radius

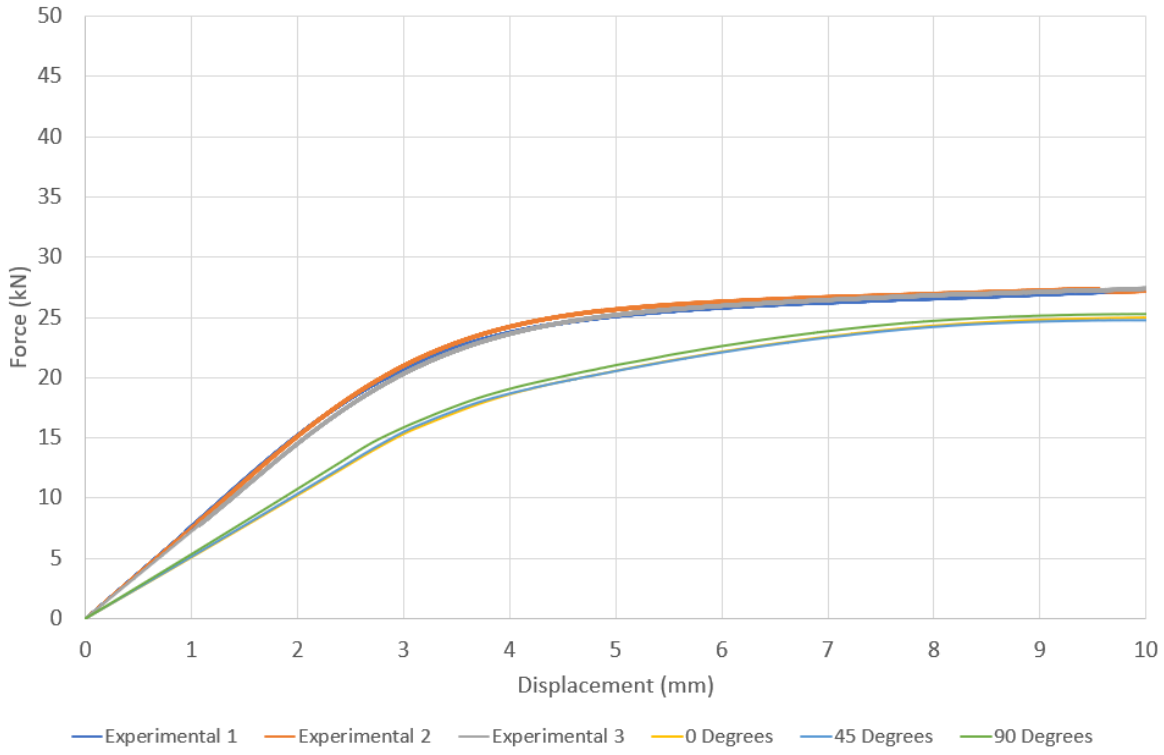


Figure 96 Experimental response vs line body FEA model of lattice compression block - 3mm strut radius

Figure 97 shows the effects of applying contacts to a solid body model of a lattice block with 2mm strut radius. The results from including the compression plate geometry and modelling the frictional interface with the test specimen is compared against fully bonded contacts and the original non-contact model. Applying bonded contacts to the analysis affected the post yield region only, while the stiffness and failure load remained in agreement with the non-contact model. On including friction the accuracy of the predicted stiffness improves. However the change is small and remains within 7% of the non-contact analysis. In this case the difference is unlikely to justify the additional time required for modelling contact effects over the course of generating an entire dataset of responses. As shown in Table 11, the frictional model increases the solve time by a factor of 4. Contacts may be introduced to aid distortion of elements undergoing highly plastic deformation in future studies. However as elastic modulus and yield strength were of greater focus than post yield behaviour in this case the contacts were removed from the simulations in favour of the reduced solve time.

Table 11 Non-linear FEA contact model solution times for 2mm strut radius lattice example

Contact Model	**No. Nodes**	**No. Elements**	**Solve Time**
None	25708	13766	2min 35sec
Bonded	38541	17474	3min 44sec
Frictional, μ=0.4	38541	17474	10min 44sec

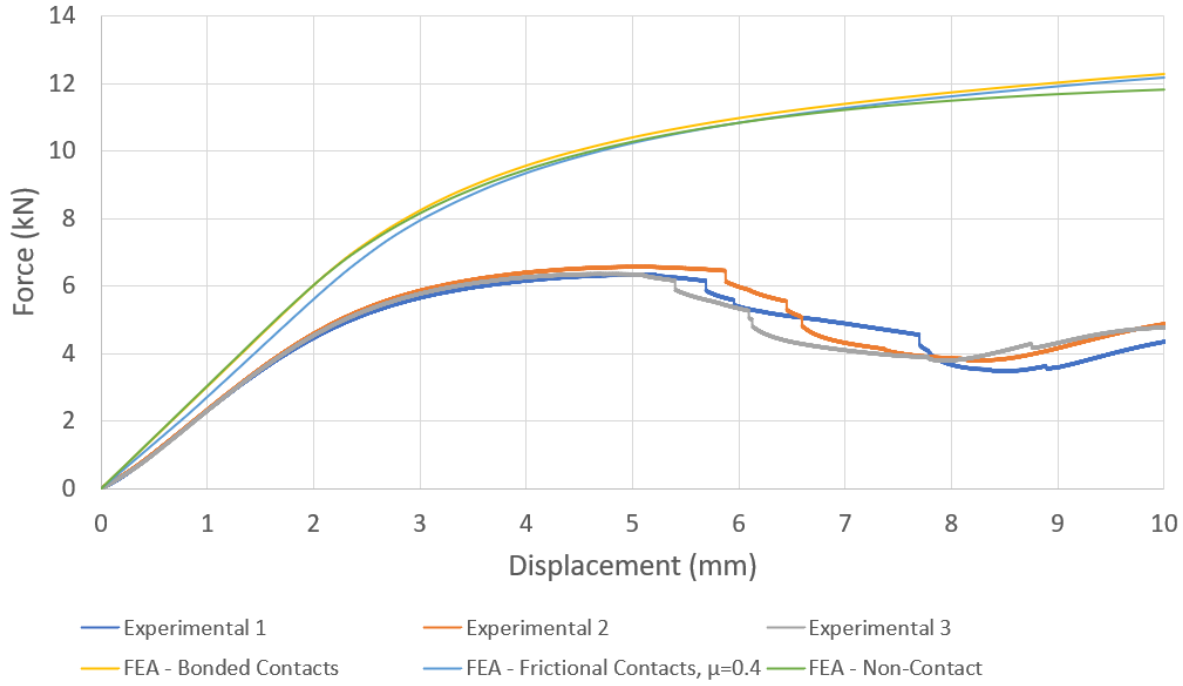


Figure 97 Experimental vs solid body FEA response of 2mm strut radius lattice with and without contacts

The results of applying non-linear buckling analysis to the solid body model without contacts are compared against the beam element model in Figure 98 for the 2mm strut radius lattice. The first 6 mode shapes extracted by the eigenvalue buckling analysis can be found in Appendix I. The simulations were repeated varying the perturbation between scale factors of 1-8. On increasing the buckling deformation scale factor the solid body FEA model can be tuned to more closely represent the experimental stiffness in its predictions. However the disadvantage of this approach is that the scaling factor must be updated as the geometry changes as the analysis is repeated over the entire dataset. Additional experimentation is required to create a suitable regression model for predicting the optimal value for a given geometry. In any case overshoots in the post-yield region remain and the beam element model proves more accurate overall as well as being more practical to implement over a large number of samples. The buckling analysis was therefore removed from further simulations for dataset generation in favour of the beam and standard solid element models. The solid body model with buckling analysis may however be useful for predicting mechanical performance of more complex geometries in future works.

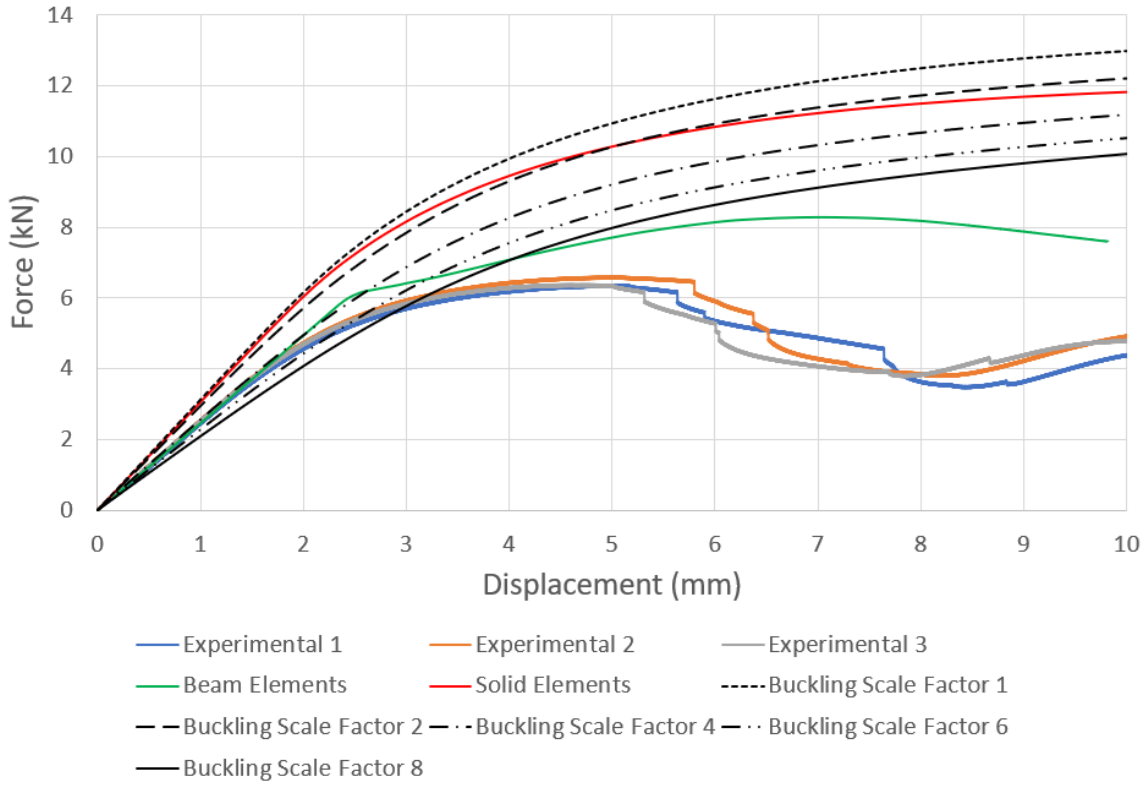


Figure 98 Experimental vs FEA responses for beam element and solid element model with buckling for 2mm strut radius lattice

The beam element FEA model was selected as the preferred approach for dataset generation. However this method resulted in reduced convergence rates when generating responses of design points from the design of experiments test plan. Particularly in cases where the outer frame radius was small the force convergence process would occasionally fail to complete. Buckling would initiate early and an extended simulation of post buckling behaviour over large displacements would be necessary, which is typically difficult to model using FEA. The solid body model offered on average 93% convergence, while the initial beam model gave 81%. To aid convergence non-linear stabilisation was used to apply small damping forces to the beam elements. The process is analogous to suspending the model in a fluid and then controlling the viscosity. Two methods were used to control the stabilisation forces and are compared against the non-stabilised solution for lattice blocks at 1, 2 and 3mm strut radius in Figures 99-101 below. The first method defines a single damping factor which is applied across all elements equally. The second is energy based and applies a dissipation ratio that is calculated for each element as the work done by stabilisation forces to the element potential energy. This method allows the stabilisation factor to vary between elements and is well suited for local instabilities such as those caused by extended periods of plasticity. In all cases use of a single damping factor affected the linear region as well as the non-linear region of the response, reducing the predicted stiffness and is not recommended. The optimal approach was

found at an energy based dissipation ratio of 0.01 which was used for generation of the final dataset. At this value the mechanical response does not deviate from the original unstabilised beam element prediction until simulation of post yield behaviour is initiated. At 1mm strut radius the stabilised post yield response is closer to experimental values. At the lower end of the design space predictions are expected to be enhanced by the stabilisation in addition to improvements in convergence without compromising other design points elsewhere in the dataset with greater strut thicknesses. At 2mm radius there is no deviation in response until approximately 6.5mm displacement, after initial fracture is seen in the experimental results. The energy dissipation result at 3mm radius indicates stabilisation forces are negligible, the curve lying directly on the original beam element prediction.

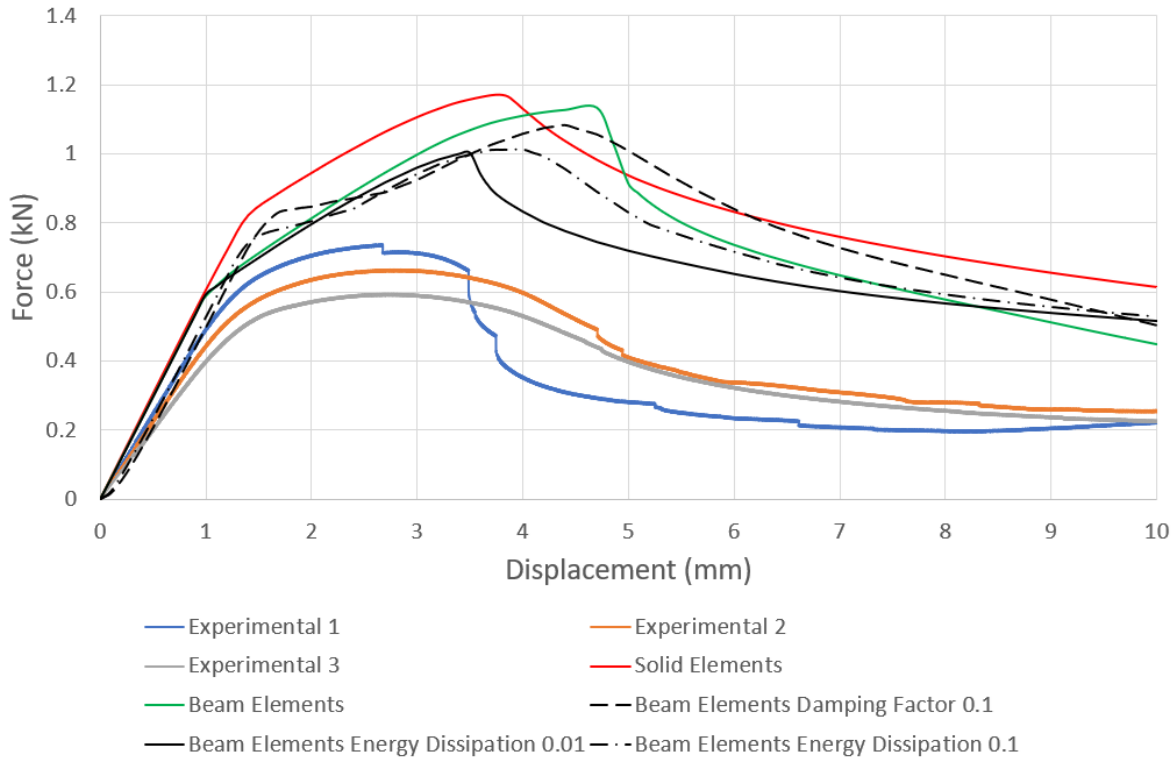


Figure 99 Non-linear stabilisation of lattice compression block with 1mm strut radius

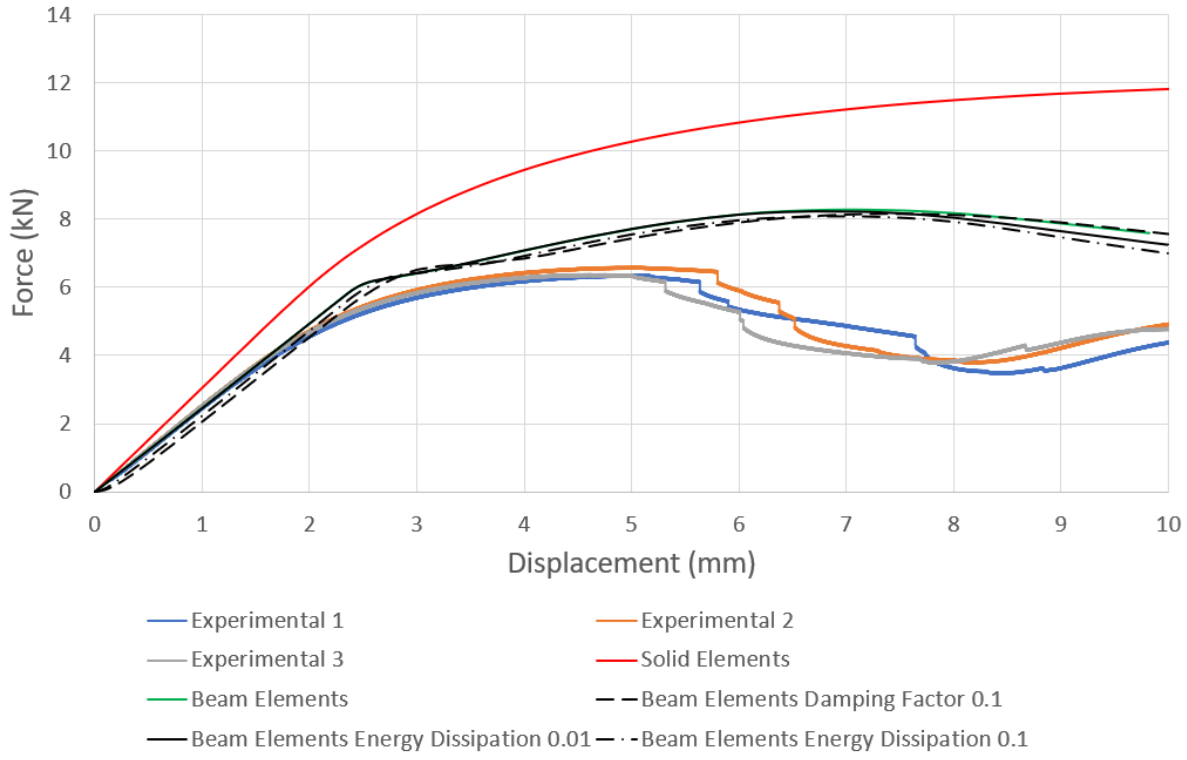


Figure 100 Non-linear stabilisation of lattice compression block with 2mm strut radius

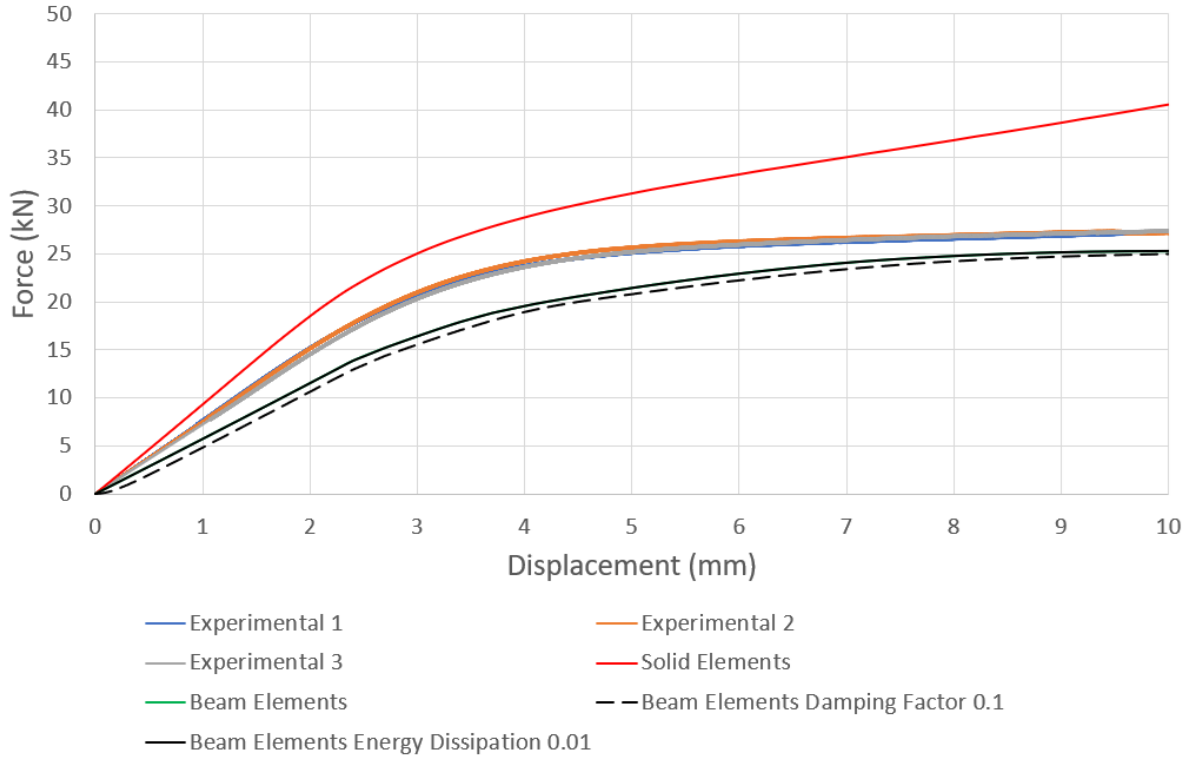


Figure 101 Non-linear stabilisation of lattice compression block with 3mm strut radius

On applying the stabilisation process 75% of previously non-converged design points gave successful force convergence. Figure 102 shows an example of the effects on a design point which had originally failed to converge halfway through the simulated displacement. The overall convergence rate of the beam element model was improved to 86%, providing a dataset of 4297 mechanical responses from an initial design of experiments test plan of 5000 design points generated using Latin hypercube sampling. A pairwise plot of the design space populated by the successfully converged results is shown in Figure 103. Generally satisfactory dispersal of datapoints is observed that matches the quality of datasets of similar quantity from the previous sound absorption study. Largest gaps appear where radius r3 is very low and r2 is between 1 and 2mm, indicating lack of convergence due to early buckling of the outer frame. However these regions are likely to be outside of those targeted by the network. As the frame is likely to be key to compressive stiffness desired tuned responses are expected to fall where r3 is greater than 1mm. The final model may experience some overshoots in the post-yield regions, though the validation results indicate the simulation is able to provide predictions of elastic modulus and yield load with good agreement to experimental values and may be used to train a convolutional neural network to tailor desired strength and stiffness with sufficient accuracy.

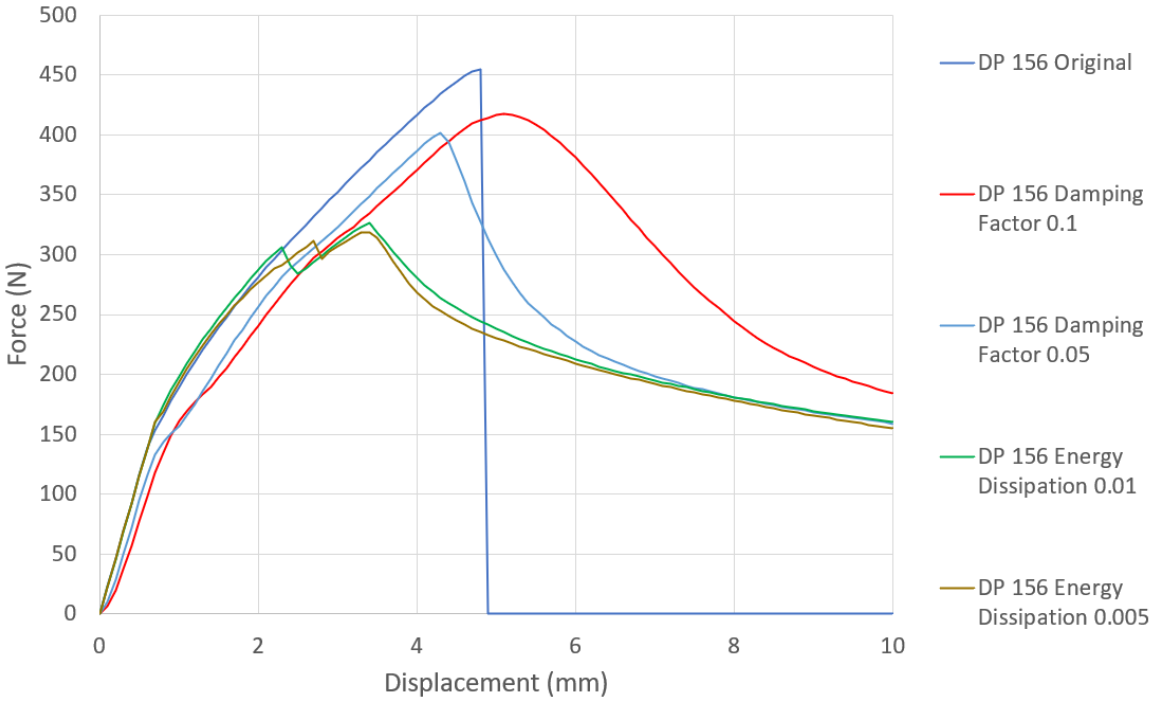


Figure 102 Non-linear stabilisation on example of a previously unconverged design point at r1=0.986, r2=0.516 and r3=0.508mm

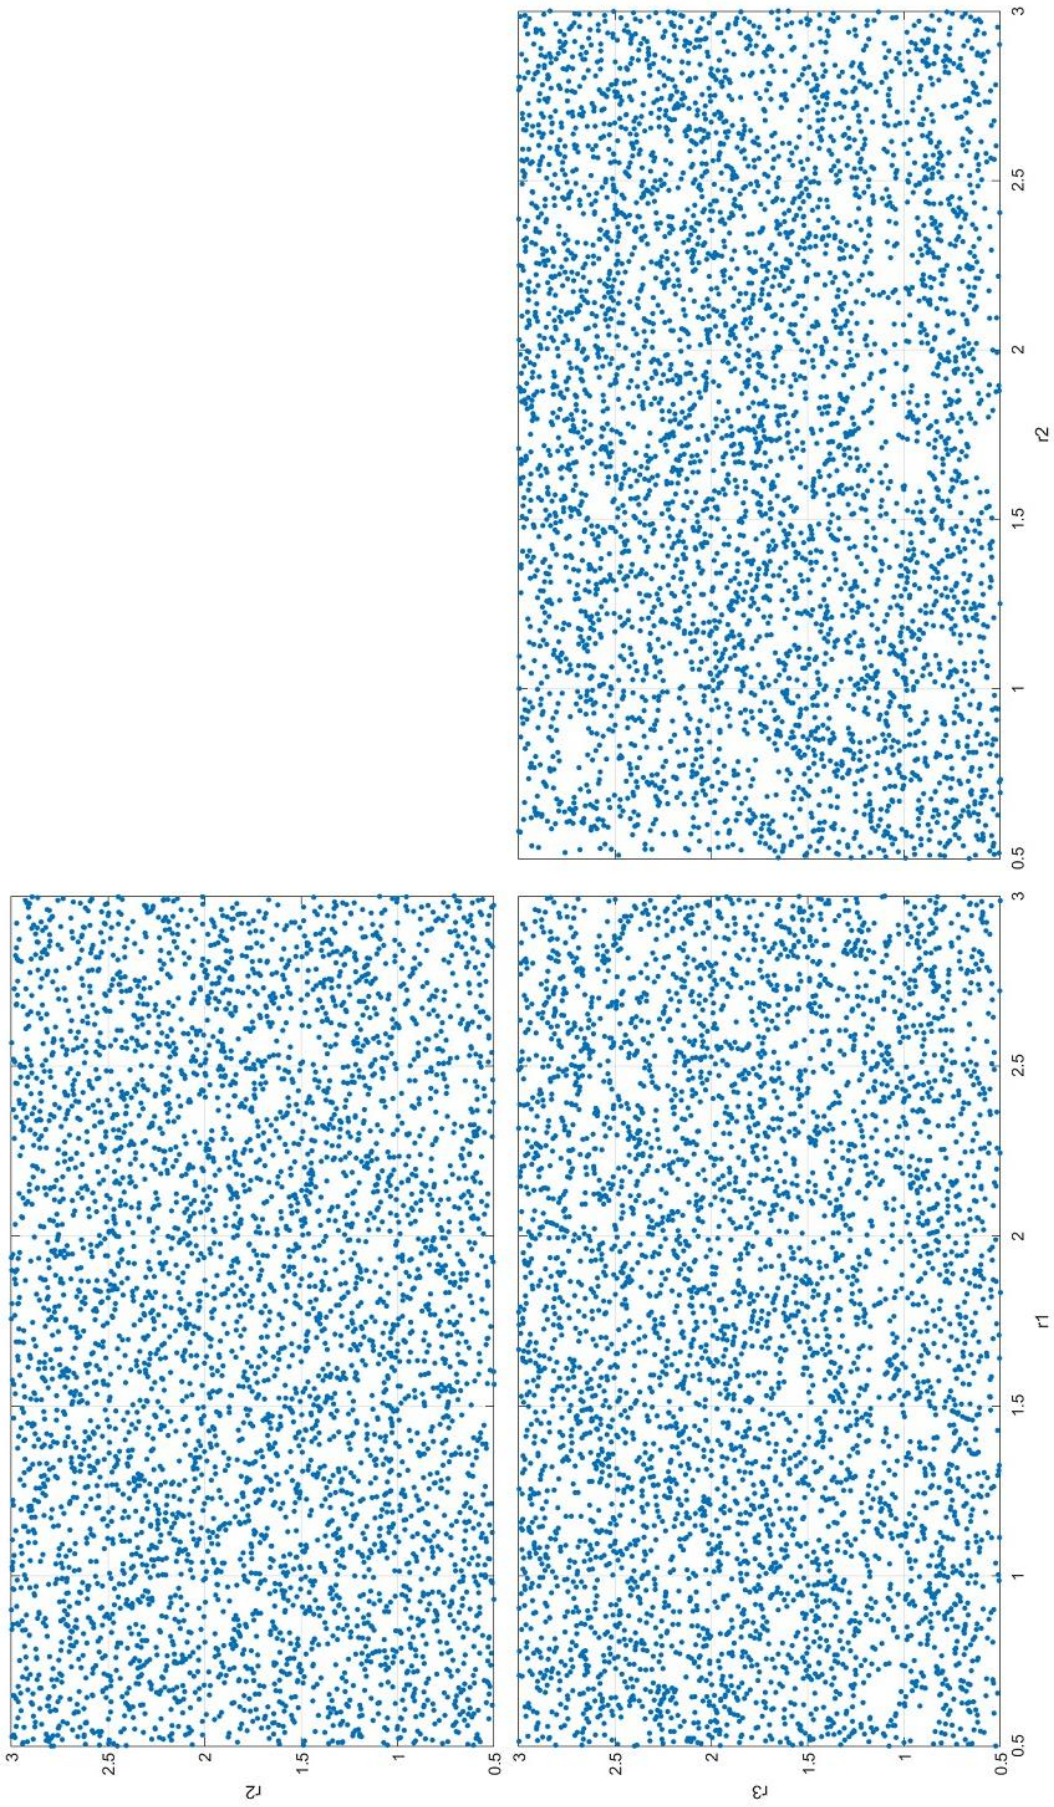


Figure 103 Pairwise plot of design space of final dataset of 4297 samples generated using the beam element FEA model with non-linear stabilisation

4.3 Convolutional Neural Network Training

The convolutional neural network architecture described in section 3.2 was repurposed using transfer learning to predict the lattice geometry required for a given mechanical response. The network model was implemented in MATLAB with a custom script used to perform the data formatting and set hyperparameters. The inputs to the network were now 100 force values recorded at 0.1mm displacement intervals between 0 and 10mm. The force data was formatted for sequence input into the first layer of the network as a matrix of cell arrays, with each array containing a vector of force values corresponding to a separate observation. This was mapped to a regression output which predicted the strut radius parameters r1, r2 and r3. For validation checks during training the true radius values were formatted as a matrix with each column representing a different parameter and each row a different observation. To aid network training performance the force values were normalised before training by rescaling to values between 0 and 1 corresponding to the minimum and maximum values found anywhere in the original dataset. Likewise the strut radius values were normalised between 0 and 1, scaled according to the minimum and maximum values in the table of design points. The network predictions were then rescaled afterwards to assess network performance against known values. The 4297 observations from the beam element FEA model were partitioned with 70% used for training the network, 15% for validation during training and 15% kept unseen by the network during training for testing the performance of its predictions afterwards. An additional dataset of 1392 samples generated by the solid body FEA model was used to train a separate network with identical architecture to evaluate its ability to tailor responses under circumstances where the quantity and fidelity of the data within the training dataset may be limited. This dataset was partitioned with the same ratio and the results used to assess if the network was able to compensate for the overshoots in the FEA predictions. The networks were each trained for 1000 epochs, passing the data through in mini batches of 100 samples at a time. The solid element network trained in 33min 17sec while the beam element network took 1hr 39min 26sec due to the increased dataset size. The training results are summarised in Table 12 and Figures 104-107 below show the training history and regression performance against known values from the test datasets.

Table 12 CNN training results for lattice mechanical response inverse design model

	Network Architecture		Training Options		Training Performance			Regression	
Solid Elements	**Type**	CNN	**Training Algorithm**	Adam	**RMSE**	Training	0.10		
	Convolution Filter Size	3	**Mini Batch Size**	100		Validation	0.15	**Parameter**	**R Value**
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96]	**No. Epochs**	1000	**Loss**	Training	0.0053		
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 3]	**Initial Learning Rate**	0.001		Validation	0.0234	**r1**	0.90693
	Activation	ReLU	**Validation Frequency**	1	**Training Time**	33min 17sec		**r2**	0.98728
	Normalisation	Yes	**Shuffle**	Every Epoch	**No. Iterations**	9000		**r3**	0.98593
Beam Elements	**Type**	CNN	**Training Algorithm**	Adam	**RMSE**	Training	0.07		
	Convolution Filter Size	3	**Mini Batch Size**	100		Validation	0.12	**Parameter**	**R Value**
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96]	**No. Epochs**	1000	**Loss**	Training	0.0027		
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 3]	**Initial Learning Rate**	0.001		Validation	0.0101	**r1**	0.92434
	Activation	ReLU	**Validation Frequency**	1	**Training Time**	1hr 48min 56sec		**r2**	0.98616
	Normalisation	Yes	**Shuffle**	Every Epoch	**No. Iterations**	30000		**r3**	0.99002

Training on a larger dataset for the beam element model provided performance improvements across training root mean square error and training loss. A clear difference can be seen in the training history plots in Figures 104 and 106 as a gap appears between the training and validation RMSE for the solid element network, indicating underfitting due to insufficient quantity of data. This gap closes for the beam element network and small increases in regression values are seen, suggesting improved predictive power without evidence of overfitting. Figures 108 and 109 on page 188 demonstrate an example of overfitting occurring on extending the training of the solid element model to 3000 epochs. In this case the validation error begins to increase and the curve rises away from the training error after a large number of iterations have passed. This leads to a drop-off in regression as the network has over learnt the characteristics of the training data to the point that it performs worse on new unseen test data. The networks trained for 1000 epochs on the other hand show good generalisation and sufficient performance for the purposes of the study. In both cases r1 showed the largest number of outliers. However both networks gave strong predictions overall, providing strut radii that agreed with known values to above 90% regression across all 3 parameters.

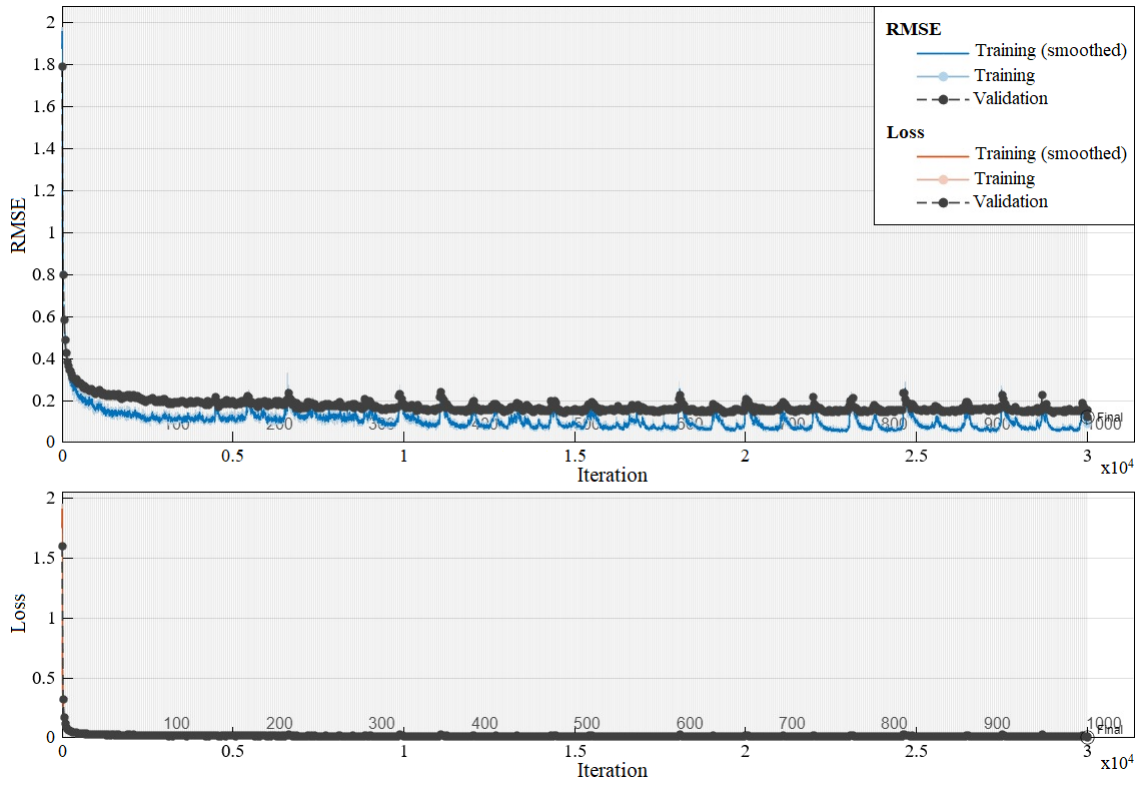


Figure 104 Training history of CNN for predicting lattice mechanical responses based on beam element FEA model

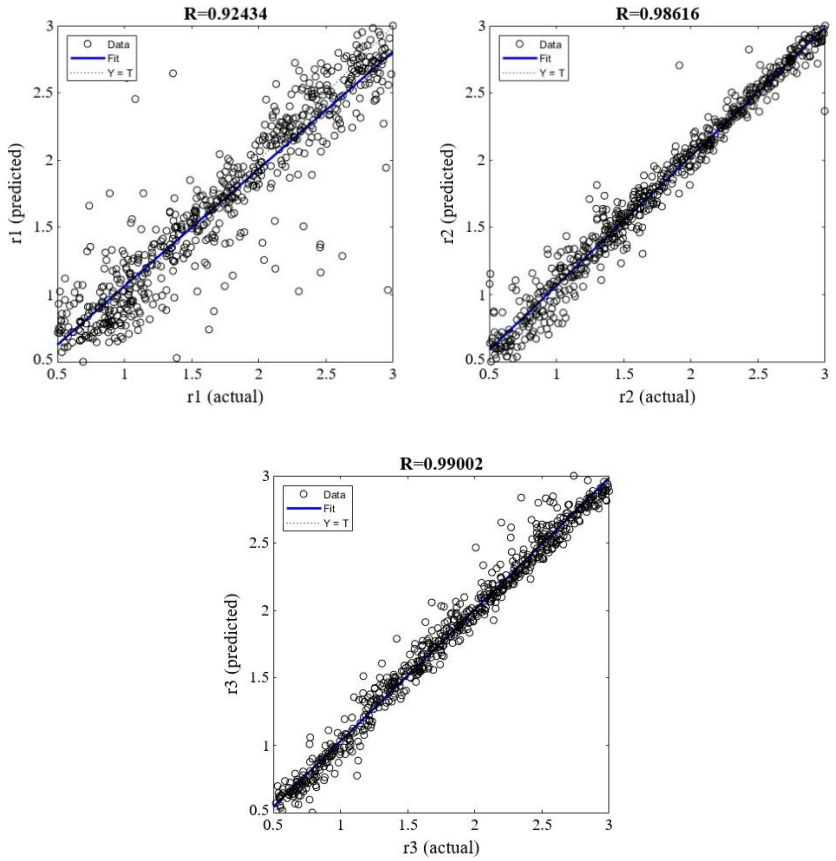


Figure 105 Regression plots of predicted vs actual lattice geometry parameters for network based on beam element FEA

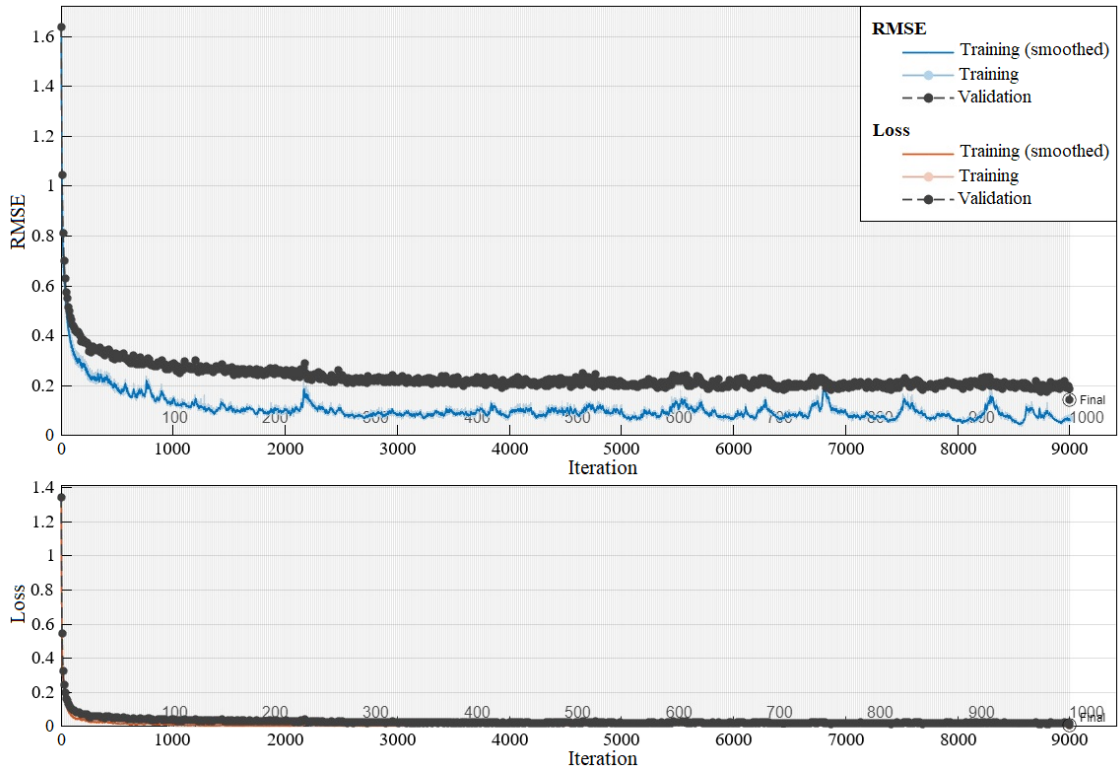


Figure 106 Training history of CNN for predicting lattice mechanical responses based on solid element FEA model

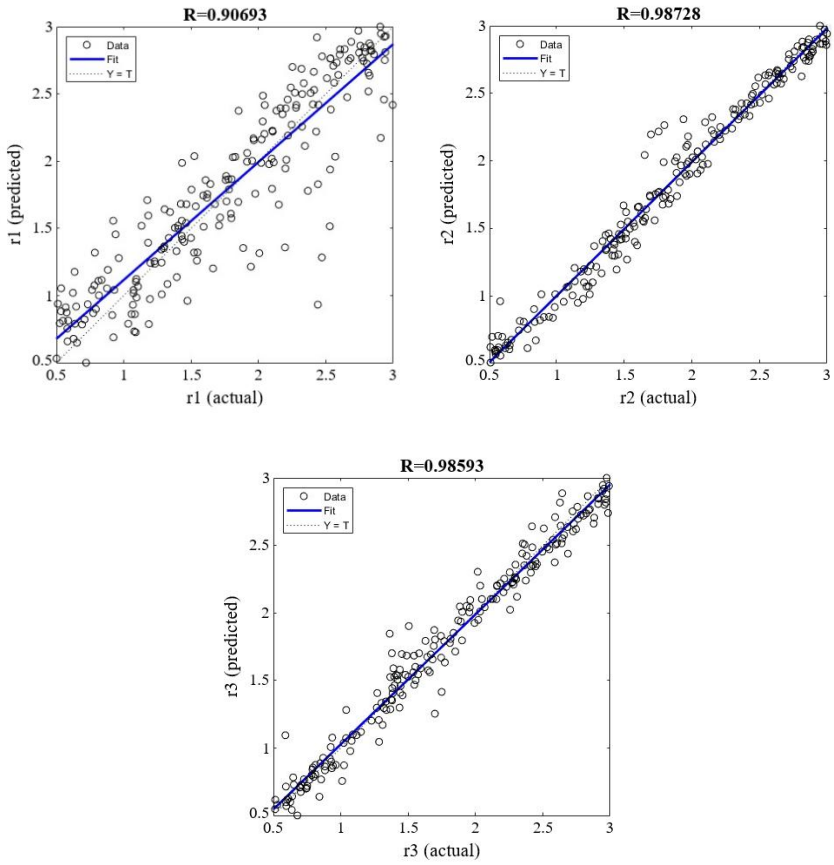


Figure 107 Regression plots of predicted vs actual lattice geometry parameters for network based on solid element FEA

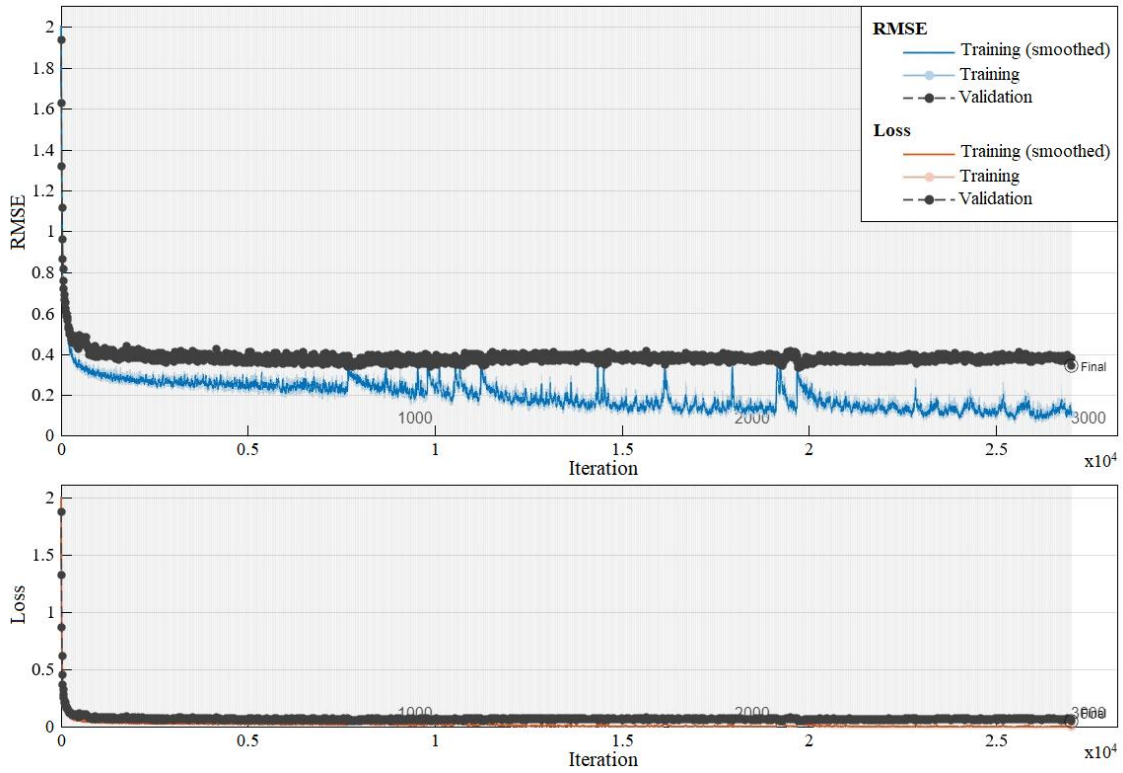


Figure 108 Training history of example network based on solid body FEA model with overfitting due to extended training time

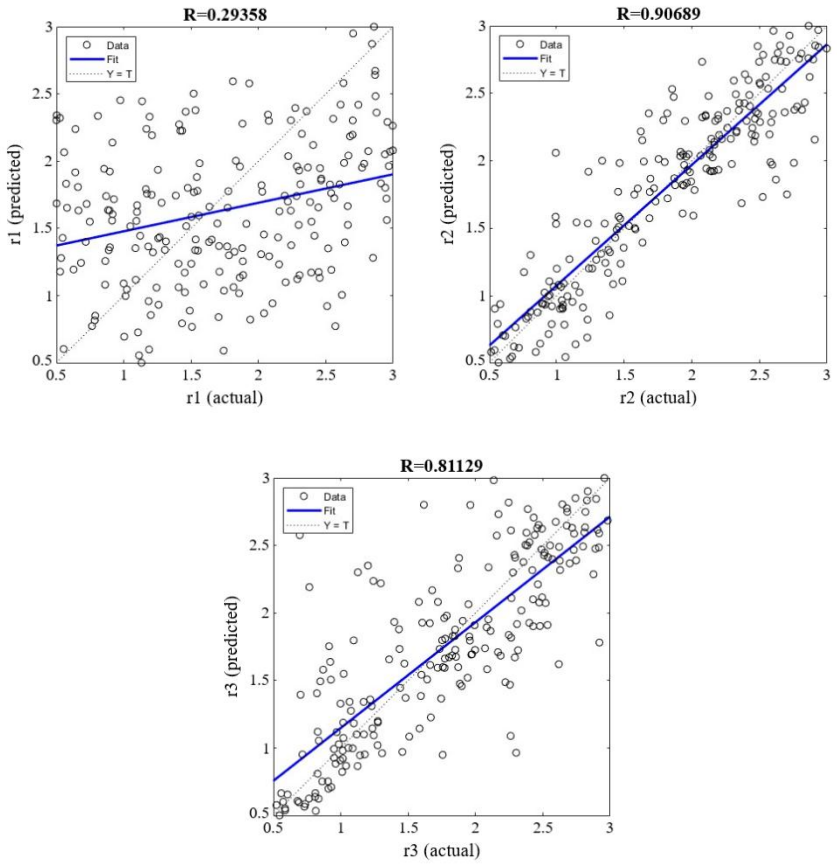


Figure 109 Regression plots for example network based on solid body FEA model with overfitting due to extended training time

4.4 Tailored Mechanical Responses

The trained networks were used to tailor mechanical responses to satisfy stiffness and strength design criteria, simulating objectives that may be targeted to meet a component design specification within industry. A design load and safety factor was set along with a target maximum deflection at the design load. The neural network was able to set the slope of the response and then the failure load could be increased incrementally until the safety factor was met, but by as small a margin as possible to maximise structural efficiency and ensure only the minimum mass necessary to satisfy the design criteria was used. To set the position of the force-displacement curve the requested responses were represented as a bilinear characteristic, setting the slope, failure load and post-yield tangent modulus. The requested forces were defined at intervals of 0.1mm displacement to match the format of the training data and input into the trained networks which returned the radii of the struts necessary to provide the response. These geometry parameters were supplied to the FEA model that was used to generate the initial dataset to give a numerical prediction which was then post validated under experimental testing using the same method as described in section 4.2.2. The networks showed high responsiveness to requested inputs, offering a large amount of control for tuning of the mechanical characteristics. There was consistently strong agreement between the request and the response curve resulting from FEA of the provided solution geometry, which could be fine-tuned through small adjustments. Example solutions from the network trained on the beam element FEA model are provided in Table 13 for a range of design specifications. Lattice geometries were tuned to target 1.5mm displacement at a design load of 5kN with a safety factor of 2, 1mm at 2kN with a safety factor of 2.5, and 0.5mm at 500N with a safety factor of 2. In Figures 110-113 below the bilinear requested responses are compared against the FEA predictions and experimental validations for each solution.

Table 13 Summary of solutions for tailored mechanical responses of lattice compression block using CNN trained with beam element model

		Solution 1 (5kN)	Solution 2 (2kN)	Solution 3 (500N)	Solution 4 (500N)
Design Specification	Design Load (kN)	5	2	0.5	0.5
	Target Maximum Displacement At Design Load (mm)	1.5	1	0.5	0.5
	Design Safety Factor	2	2.5	2	2
	Target Failure Load (kN)	10	5	1	1
Solution Geometry	Strut Radius (Predicted), r1 (mm)	2.433	2.475	0.864	1.522
	Strut Radius (Predicted), r2 (mm)	1.145	1.140	0.833	1.066
	Strut Radius (Predicted), r3 (mm)	2.947	1.936	1.506	1.265
	Mass (g)	22.725	16.358	6.371	8.545
	Volume - Solid (mm^3)	20290.187	14605.631	5688.111	7629.361
Numerical Prediction	Predicted Failure Load – FEA (kN)	10.664	5.468	2.512	2.341
Experimental Validation	Minimum Actual Failure Load – Experimental (kN)	10.102	5.003	1.295	1.329
	Actual Safety Factor	2.020	2.501	2.590	2.659
	Actual Displacement At Design Load (mm)	1.193	0.839	0.537	0.565

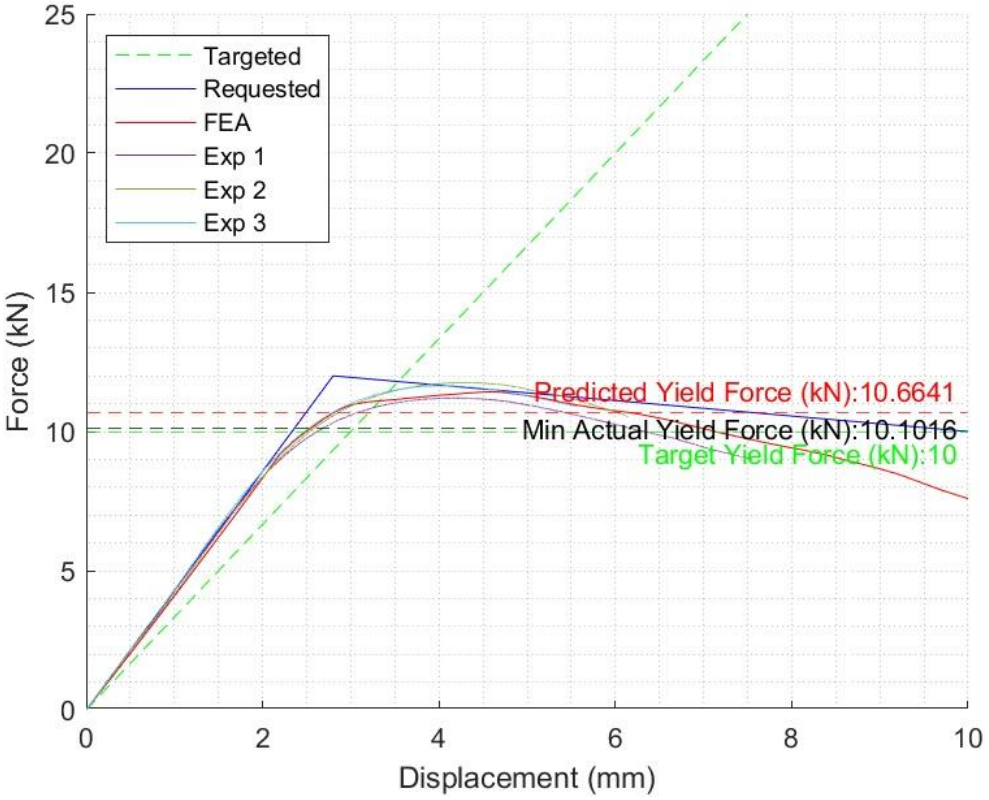


Figure 110 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 1 targeting 5kN design load

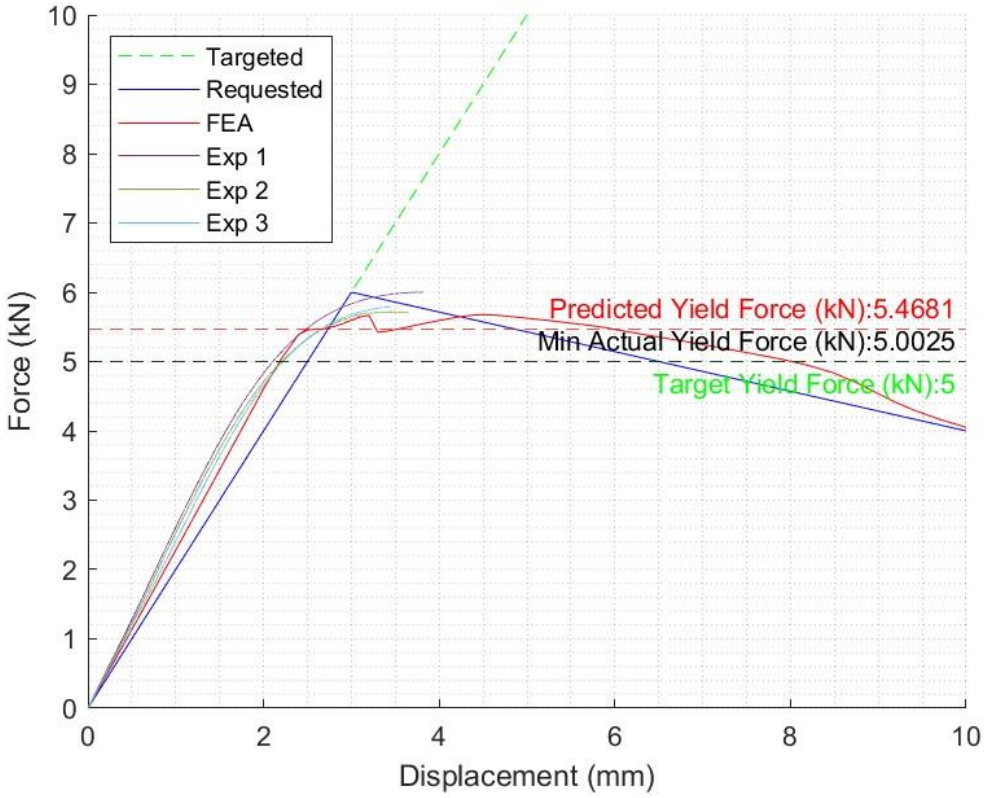


Figure 111 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 2 targeting 2kN design load

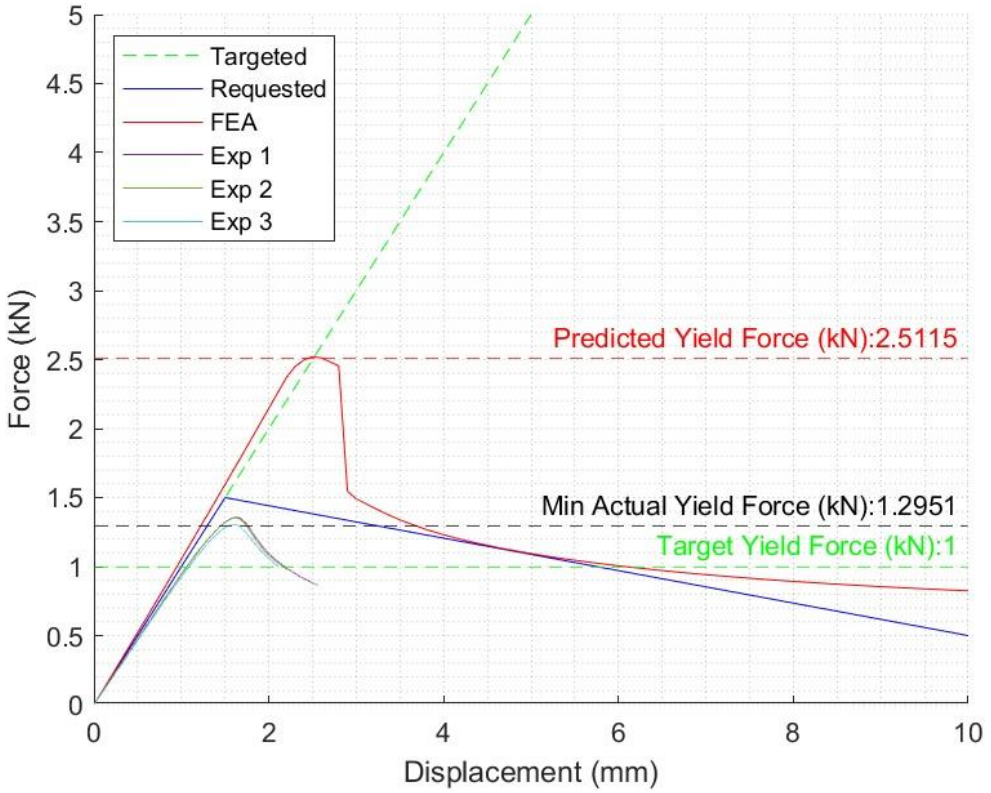


Figure 112 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 3 targeting 500N design load

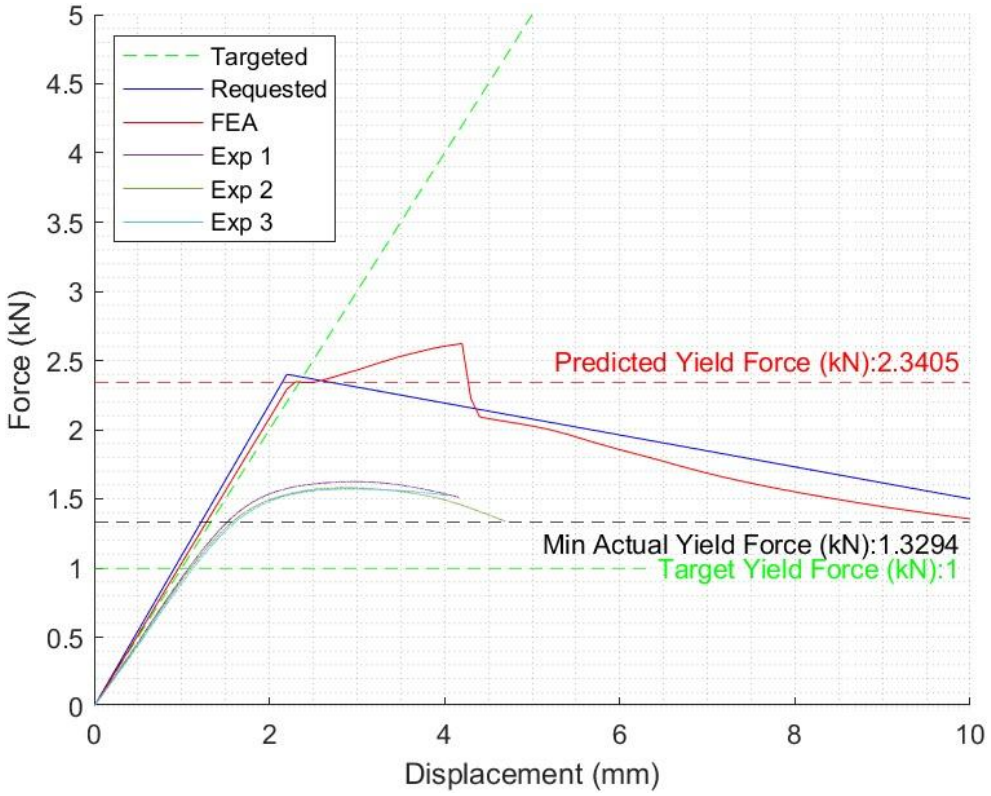


Figure 113 Requested mechanical responses vs FEA predictions and actual experimental validations for solution 4 targeting 500N design load

The experimental validations of the solutions at 5kN and 2kN design loads closely followed the FEA predictions. At 500N design load the geometries became susceptible to excessive bending and failure by buckling prior to the expected yield position due to the small strut radii. This caused overshoots in the predictions that required additional compensation. However the network model was still able to provide solutions to meet the target strength requirement. In addition on increasing the area below the requested response curve the specific energy absorption of the lattice could be altered while maintaining the targeted design specification. The bilinear request showed sufficient accuracy for tailoring non-linear behaviour. In each case an initial assumption of the second slope was followed by a process of trial and error facilitated by the speed of the solver, which returned the response to best match the requested input. By fine-tuning the slope of the post-yield region the mean volumetric energy density was increased by a factor of 2.7 at the cost of 2.174g additional mass. The experimental energy absorption values for each specimen tested are displayed in Table 14.

Table 14 Experimental energy absorption values measured for lattice compression block solutions targeting 500N design load

	Solution 3		Solution 4	
	Specific Energy Absorption (J/kg)	Energy Density (kJ/m^3)	Specific Energy Absorption (J/kg)	Energy Density (kJ/m^3)
Specimen 1	348.89	70.22	608.46	164.25
Specimen 2	299.36	60.25	678.51	183.16
Specimen 3	275.08	55.36	567.11	153.09
Mean	307.78	61.95	618.03	166.84

As predicted during validation of the non-linear FEA the network trained on the solid element model exhibited large overshoots in the post-yield regions when the solution geometries were subjected to real-world testing. However the network gave good predictions for stiffness and was actually more responsive to requested inputs. The slope of the non-linear regions of the FEA predictions from the returned geometry more closely followed the requested tangent modulus. This is likely because the line body model typically featured more peaks and troughs as it included the effects of beam buckling and larger plastic deformations so there was a greater degree of variation in the dataset. The solid body model was still able to provide experimental solutions at the required performance targets that were competitive with the beam element model, though additional steps were required to compensate for the overestimated failure load. Tables 15 and 16 below summarise the design iterations that were necessary to produce solutions at the same targeted design criteria as before for the 5kN and 2kN design load examples respectively. In each case the initial iteration was validated experimentally to quantify the deviation from the expected value. The speed of the solver then allowed multiple results to be produced around the target within minutes. The most time

consuming part of the process was the FEA studies which each took approximately 10-15 minutes to solve, the AI returned the geometry predictions instantaneously on making a request. For the 5kN design load target a successful solution was found on the third iteration, increasing the safety factor by 4.3% and reducing the deflection at 5kN by 0.1449mm compared to the beam model equivalent, though at 0.833g additional mass. For the 2kN target the solid body model produced a slightly improved solution overall, improving deflection by 0.1457mm at the design load and increasing safety factor by 3.8% at 97.6% of the mass of the beam element solution. However more extensive additional fine-tuning and validation iterations were required whereas the beam element model produced experimental results at the correct targets the first time for both design specifications. The requested, numerically predicted and experimentally validated mechanical characteristics of each iteration are displayed in Figures 114-121 below.

Table 15 Summary of solution iterations to satisfy design criteria at 5kN design load target using CNN trained with solid element model

		5kN Iteration 1	5kN Iteration 2	5kN Iteration 3
Design Specification	Design Load (kN)	5	5	5
	Target Maximum Displacement At Design Load (mm)	1.5	1.5	1.5
	Design Safety Factor	2	2	2
	Target Failure Load (kN)	10	10	10
Solution Geometry	Strut Radius (Predicted), r1 (mm)	0.912	0.954	1.946
	Strut Radius (Predicted), r2 (mm)	1.710	1.899	2.033
	Strut Radius (Predicted), r3 (mm)	2.949	2.961	2.813
	Mass (g)	19.213	20.511	23.559
	Volume - Solid (mm^3)	17154.609	18312.991	21034.637
Numerical Prediction	Predicted Failure Load – FEA (kN)	10.952	11.806	12.912
Experimental Validation	Minimum Actual Failure Load – Experimental (kN)	7.888	8.996	10.533
	Actual Safety Factor	1.578	1.799	2.107
	Actual Displacement At Design Load (mm)	1.237	1.169	1.048

Table 16 Summary of solution iterations to satisfy design criteria at 2kN design load target using CNN trained with solid element model

		2kN Iteration 1	2kN Iteration 2	2kN Iteration 3	2kN Iteration 4	2kN Iteration 5
Design Specification	Design Load (kN)	2	2	2	2	2
	Target Maximum Displacement At Design Load (mm)	1	1	1	1	1
	Design Safety Factor	2.5	2.5	2.5	2.5	2.5
	Target Failure Load (kN)	5	5	5	5	5
Solution Geometry	Strut Radius (Predicted), r1 (mm)	1.289	1.139	1.053	1.822	1.908
	Strut Radius (Predicted), r2 (mm)	0.921	0.981	0.931	0.745	0.743
	Strut Radius (Predicted), r3 (mm)	2.164	2.266	2.461	2.449	2.544
	Mass (g)	11.414	11.758	12.573	14.909	15.963
	Volume - Solid (mm^3)	10190.900	10497.874	11226.102	13311.388	14252.410
Numerical Prediction	Predicted Failure Load – FEA (kN)	5.289	5.822	6.616	7.138	7.776
Experimental Validation	Minimum Actual Failure Load – Experimental (kN)	3.655	3.756	4.456	4.803	5.195
	Actual Safety Factor	1.828	1.878	2.228	2.402	2.597
	Actual Displacement At Design Load (mm)	0.960	0.930	0.774	0.773	0.693

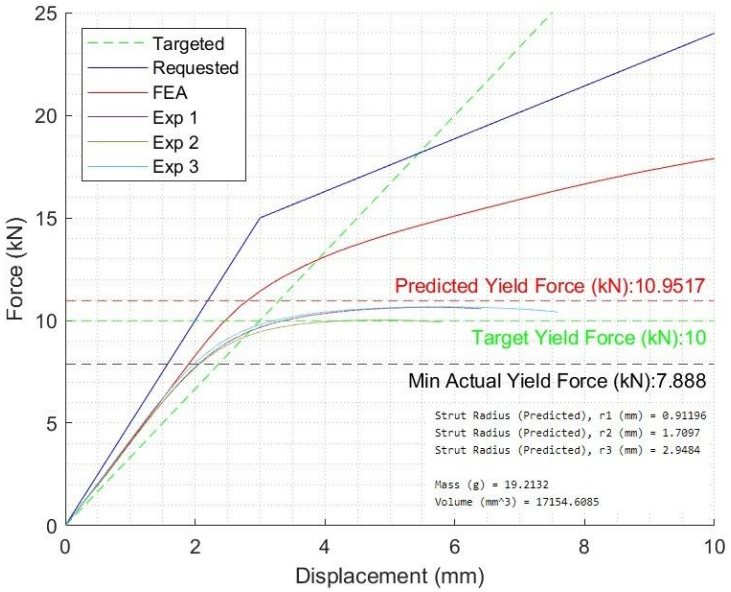


Figure 114 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 1

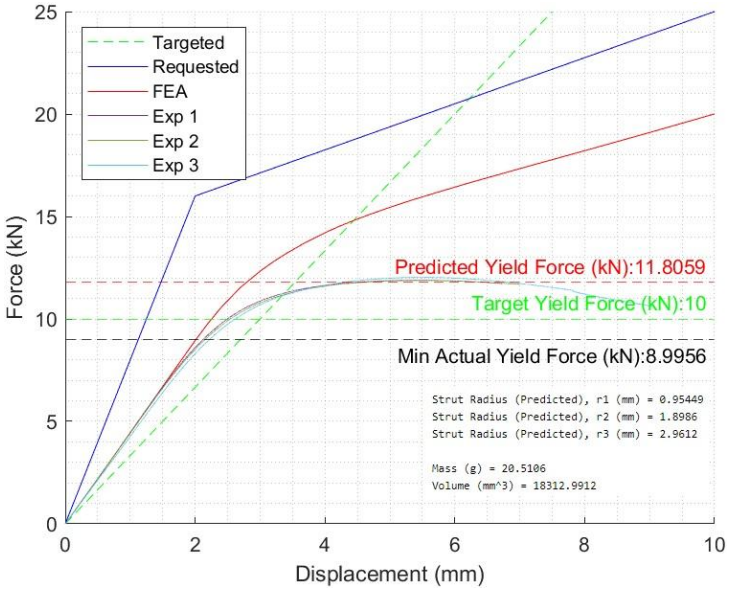


Figure 115 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 2

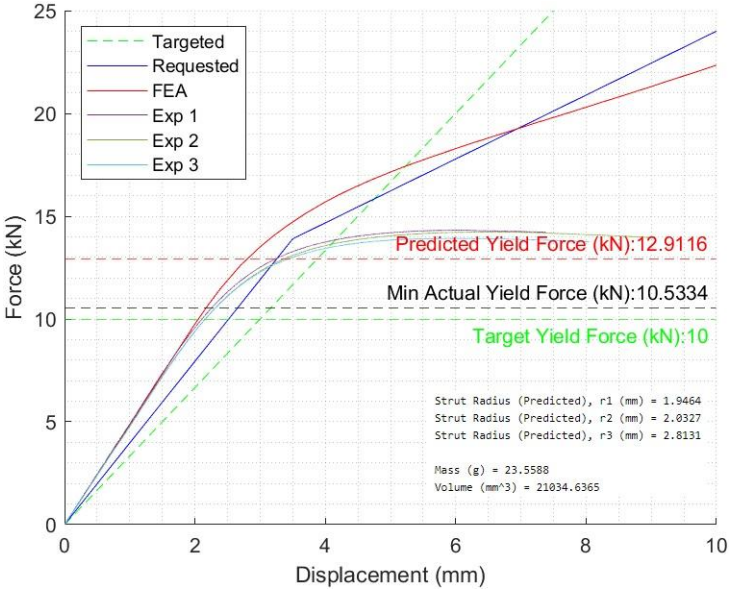


Figure 116 Mechanical response tuning result for CNN trained on solid body model - 5kN Iteration 3

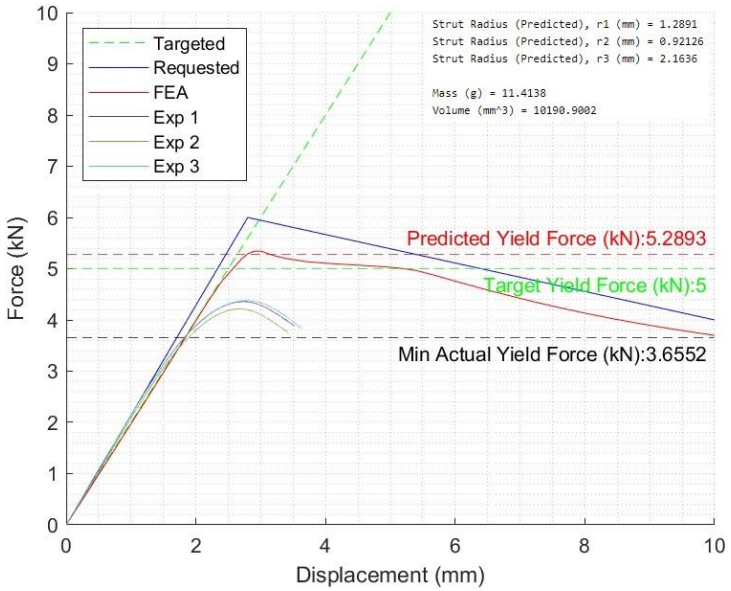


Figure 117 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 1

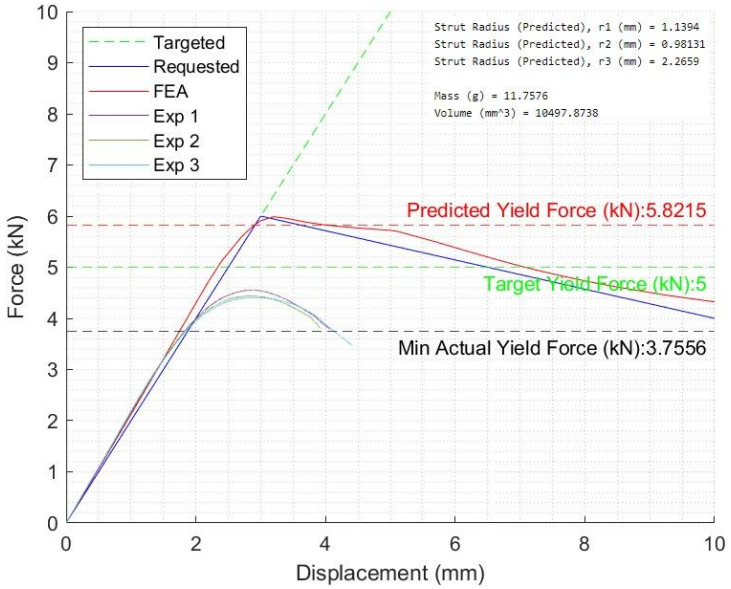


Figure 118 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 2

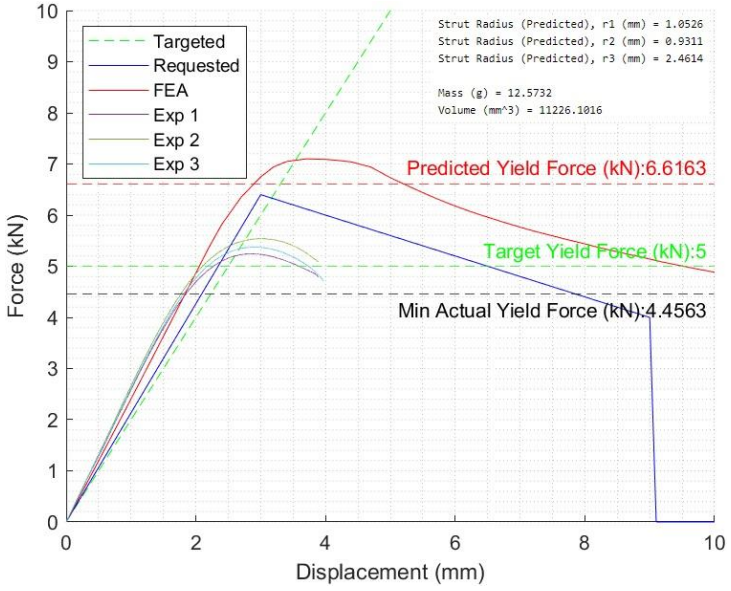


Figure 119 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 3

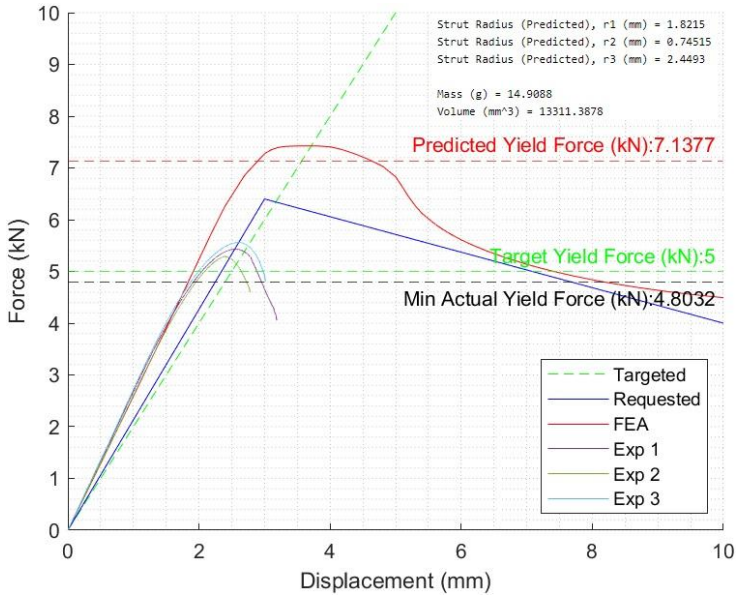


Figure 120 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 4

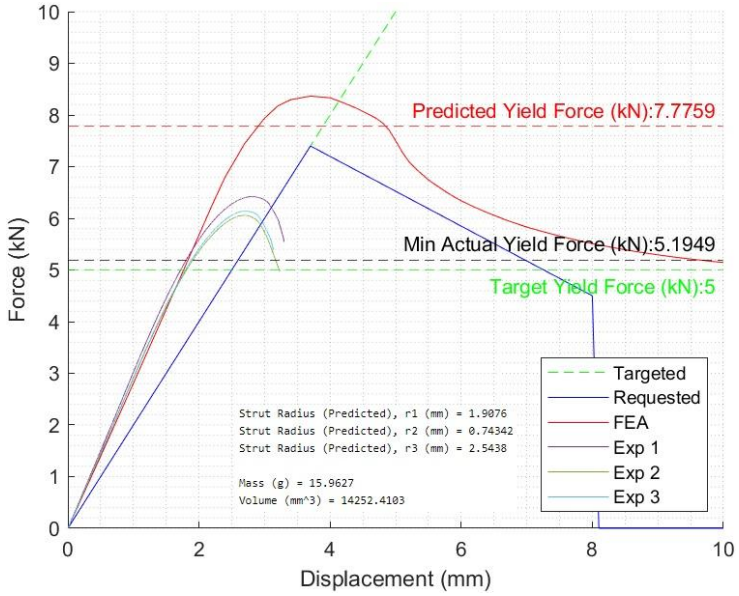


Figure 121 Mechanical response tuning result for CNN trained on solid body model - 2kN Iteration 5

4.5 Conclusion

The AI assisted inverse design process allows for tailoring of metamaterial mechanical properties so that a structurally efficient distribution of material can be assigned to satisfy design criteria under specific loading conditions, offering lightweighting advantages that can be applied across engineering design practices. During conventional workflows the designer creates a geometry which is then analysed by software tools to quantify the structural response before iteration. The proposed inverse design process instead allows the user to define the target structural response from the offset and the software tools then return the geometry to satisfy the request. The user makes an initial assumption of the most appropriate bilinear response curve to represent the desired characteristic. The geometry prediction is then returned by the neural network. This geometry is analysed through FEA simulation and the obtained response compared to the initial request. The user is free to make any updates as necessary while the machine learning model matches the mechanical responses as closely as possible to the requested inputs. Trade-offs between stiffness, strength and energy absorption can be tuned according to the given application. The stiffness can be set and then the failure load altered to get below a certain mass requirement or alternatively the failure load set by a safety factor and the stiffness tuned. Additional area under the force-displacement curve can be requested to improve specific energy absorption for enhanced impact resistance. After validation results can be easily updated using the same model to balance deflection and safety factor with freedom to make alterations to the initial design specification if desired. The ability to dictate mechanical responses at component level through tailoring material scale geometry offers significant advantages in design flexibility over traditional methods. Designs can be iterated and mechanical characteristics updated to satisfy new design criteria without changing the overall macro shape of the component. This may offer advantages in aerodynamics and heat transfer, where the shape may be set to achieve optimum flow rates and heat dissipation while the structural properties are set by the metamaterial architecture. The inverse convolutional neural network offers significant speed advantages over traditional topology optimisers. The solver returns the required geometry solutions to meet a specific response instantaneously on making a request. The designer is able to compensate for any deviations between numerical predictions and experimental results within a practical timeframe. If the initial result requires adjustment when tested the solution can be reiterated, printed and tested to meet design requirements within a matter of hours. Fine-tuning the stiffness and strength until it only just satisfies the requirement is likely to give the most lightweight solution. However combination of a CNN working in the opposite direction to predict the mechanical responses with an evolutionary optimisation algorithm to minimise

mass may improve results further and will be explored in following studies. The downside of this approach however is likely to be reduced design flexibility. It may be possible to improve the tunability of the inverse model further by using multilinear curves to request the responses. This is likely to allow optimisations of problems featuring curves with more complex features which could be particularly useful for tailoring non-linear plastic behaviour and results from dynamic simulations. The outlined workflow is likely to be applicable to other loading situations and full component designs, as long as a dataset can be generated of sufficient size. For larger FEA models high performance computing may be needed or alternatively the network can be simplified to predict only the stiffness and strength as scalar parameters rather than the full response curve, which is likely to allow for smaller datasets and studies based on experimental data only. Results from the solid body model suggest the inverse design CNN workflow may still be able to produce successful results when the fidelity of the FEA model is limited. As long as enough information is provided for the network to identify the underlying patterns, the mechanical response can still be tuned with good accuracy.

Chapter 5 Machine Learning Based Multi-objective Structural Optimisation of Lightweight Sound Absorber Metamaterials

Multi-objective optimisation problems present particular challenges when competing performance objectives are not possible to relate through conventional theoretical means. In many industrial, energy and transportation based engineering applications the demand for lightweight, strong and stiff materials is accompanied by a requirement for compact noise reduction that remains effective at low frequencies. Low-frequency wide bandwidth sound absorbers are typically combinations of multiple absorption devices, each designed to target a different frequency range. Commonly used absorbers feature layers of high thickness fibrous material and cellular foam structures with low stiffness and strength. Increasing demand for lightweighting and refinement in passenger vehicles creates a significant requirement for absorbers at sub-wavelength thicknesses with multifunctional load bearing capabilities that are inexpensive and can be manufactured rapidly. If both sound absorption and structural properties can be tuned in the same metamaterial, multiple material layers can be combined into a single more lightweight and compact solution, leading to improvements in fuel consumption. This presents a multi-objective optimisation problem in which material unit cell geometry must be designed in such a way that it may be tuned for both sound absorption and structural efficiency simultaneously.

While there are many applications of metamaterials that have been engineered for desired sound absorption or structural properties, far fewer have been developed with simultaneous integration of both functionalities. Such architectures may be designed based on Helmholtz resonance or cavity resonance absorption mechanisms to provide predictable sound absorption responses. Activation of Helmholtz resonance requires the formation of resonator cavities through the material which are preceded by narrow necks or pores. Cavity resonance effects may be initiated by engineering void regions of known cross section, such as hollowed struts or labyrinth channels. The strength of the absorption typically varies according to the area of the cross section, while the length of the cavity may be used to tune the natural frequency. Designs therefore often rely on an accurate calculation of the effective length of the sound wave as it passes through the material. Combining structural efficiency with optimal sound absorption in a single design remains a challenge, as these objectives typically conflict. High sound absorption often reduces strength and stiffness.

In order to achieve multifunctionality lattice architectures are often decoupled into separate sound absorbing and structural elements to define clear resonance mechanisms. This allows the configurations to be tuned independently of each other and combined as multiple layers in series. Decoupling structural and sound absorption mechanisms often allows greater freedom for impedance matching. However a key downside is that large overall thicknesses are usually required to achieve satisfactory structural performance and strong low-frequency sound absorption. For acoustic layers sound absorption is usually given the highest priority which compromises structural properties and the overall design is less likely to feature the most optimal compromise between the competing objectives. Integrating acoustic and structural architectures may reduce material thickness however greater consideration of the performance trade-offs between the competing attributes is necessary.

Previous studies typically investigate the effects of varying one geometry parameter while all others remain constant. There are currently limited examples where the optimal combination of parameters has been obtained by multi-objective optimisation due to a lack of theory able to link diverse physical phenomena. A research gap remains for repeatable design workflows that allow for simultaneous optimisation of sound and structural properties to obtain real world solutions that maximise the performance objectives of both considerations at the same time. Materials with strong noise attenuation at low frequency ranges below 2000Hz are of significant use in motor vehicles and aircraft where low-frequency noise is commonplace and space is highly limited. Practical sound absorber materials at thicknesses below 5mm are currently particularly limited and usually result in highly compromised absorption performance.

The following study proposes a process for simultaneous inverse tailoring of both sound absorption and structural performance in the same neural network model. Desired mechanical characteristics and absorption responses are supplied as input to a convolutional neural network. The network is trained to output the geometry parameters of a metamaterial lattice to satisfy the request. The trained network is then used to produce unit cell designs that satisfy a selection of performance criteria to simulate practical design scenarios. The model is then inverted to predict the structural performance for a given geometry. The resulting network is combined with a multi-objective genetic algorithm to seek a Pareto front of globally optimal designs with maximum relative specific stiffness, strength and sound absorption, for practical situations where the exact design criteria are unknown. The original inverse model is then applied to an alternate unit cell, designed to target highest sound absorption performance at lowest possible overall material thickness, while also offering tuneable structural properties.

5.1 Inverse Tailoring of Combined Sound Absorption and Structural Performance

To allow the neural network to tune both sound absorption and structural properties simultaneously, a combined vector containing both the absorption coefficient-frequency response and the force-displacement characteristic of a representative unit cell of the material was supplied as the network input. The network was then taught to output key geometry parameters of the unit cell for a given input request by training on a dataset of known examples. The process began with selection of an appropriate unit cell able to provide tunable sound absorption that can be reliably predicted by analytical theory, while at the same time allowing for a controllable combination of solid material and void space for structural optimisations. A parametric FEA model was created for the chosen unit cell in which the geometry parameters could be updated from a design of experiments test plan and the analysis repeated automatically to populate the dataset of structural responses. For each observation the FEA responses were combined with the corresponding analytical sound absorption characteristic for the given geometry parameters. The combined vectors were then formatted for input into the neural network, splitting the data into three separate partitions for training, validation and testing. The training partition is used to update the weights and biases of the network, while the validation set is used to carry out performance checks of the network predictions against known values at preselected intervals during the training. The test dataset is kept unseen by the network during training and is used to evaluate the network performance after the training is complete. The network architecture and hyperparameters are updated until the network provides satisfactory correlation between the predicted and known geometries. The trained network is then used to request desirable performance characteristics and fine-tune to satisfy practical design criteria. The results are validated experimentally and the network can then be used to compensate for any discrepancies between the real-world measurements and the numerical or analytical solutions.

5.1.1 Metamaterial Unit Cell Configuration

To allow for tunability of the sound absorption characteristics, the unit cell was designed to incorporate resonator chambers in which the natural frequency and acoustic impedance can be related directly to the geometry parameters of the cross section and the length of the resonator cavity. Two fundamental configurations were tested. The first consisted of closed cell honeycomb in which a Helmholtz resonator (HR) with extended neck is incorporated into each cell of the honeycomb in the layout shown in Figure 122. The wall thickness, cell height

and apothem of the honeycomb can be adjusted to tune the structural properties, where the apothem is given by half the width between parallel sides of each hexagon. The sound absorption can be predicted analytically using the length of the resonator neck l_N, the length of the resonator cavity l_C, the radius of the neck r_N and the effective radius of the cavity r_C, where the effective radius is that necessary to give a circular section with the equivalent cross-sectional area.

(a) (b)

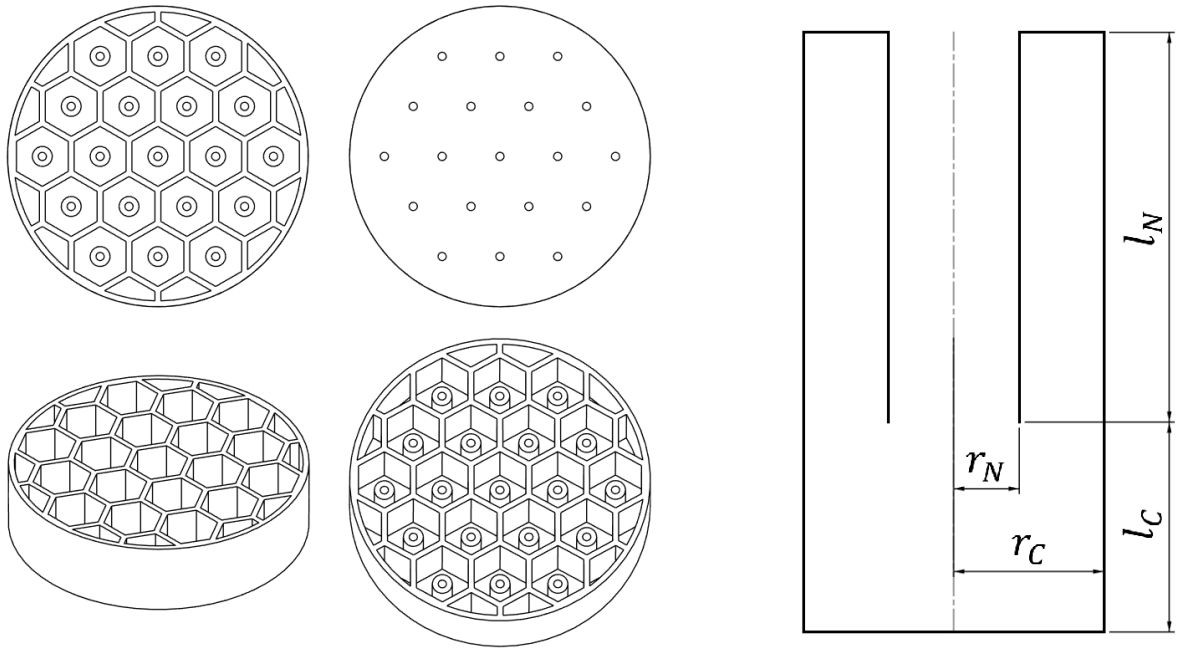


Figure 122 (a) Helmholtz resonator based honeycomb metamaterial configuration top, bottom, isometric and 3D views; (b) Section view of an individual Helmholtz resonator cavity with extended neck.

The second configuration featured a labyrinth of quarter-wavelength resonator (QWR) channels arranged transverse to the direction of the inlet. An initial spiral layout was tested and the design altered to the concentric circular and hexagonal arrangements shown in Figure 123 to increase the maximum channel length available for a given unit cell size, allowing for the absorption peak to reach lower frequencies. For these designs both the absorption and structural properties were set by the wall thickness t, the channel width w, and the channel height h. The amplitude of the absorption peak depends on the cross-sectional area of the channels, while the frequency at which the absorption peak occurs depends on the length of the channel, l. The natural frequency f at which the peak occurs is inversely proportional to the length of the resonator according to $f = \frac{c}{4l}$, where c is the speed of sound.

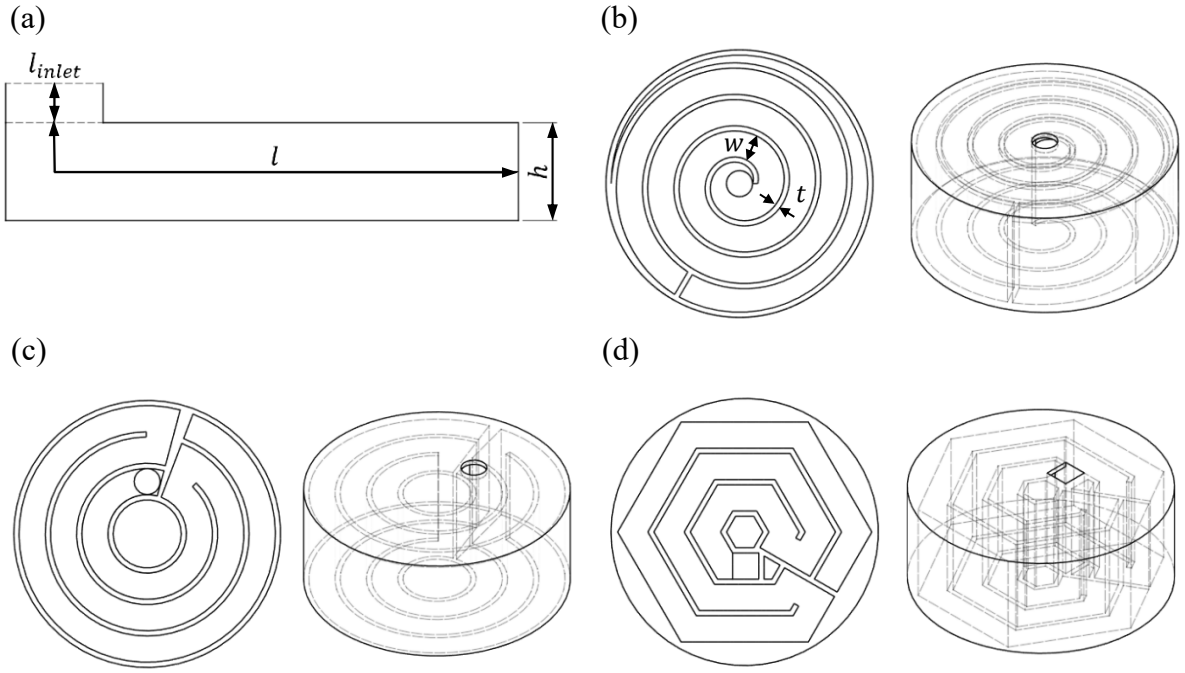


Figure 123 (a) Diagram of straight QWR with transverse inlet. Thickness of the front wall, l_{inlet} is considered separately to the length of the resonator cavity, l in analytical solution; (b) Equivalent QWR coiled into an Archimedes spiral; (c) Concentric circular labyrinth design; (d) Concentric hexagonal labyrinth design.

For both configurations the ability to combine multiple absorbers in parallel for increased broadband range was examined. The HR layout was divided by assigning different resonator parameters at certain regions within the unit cell, altering the cavity length where necessary by introducing rigid back wall terminations. Each individual resonator was allocated the same proportion of the overall active cross-sectional area and the assignments kept symmetrical with as many lines of symmetry as possible. The QWR labyrinths were divided into discrete sections, each with a separate inlet. Examples of the parallel geometry layouts are displayed in Figure 124 and Figure 125.

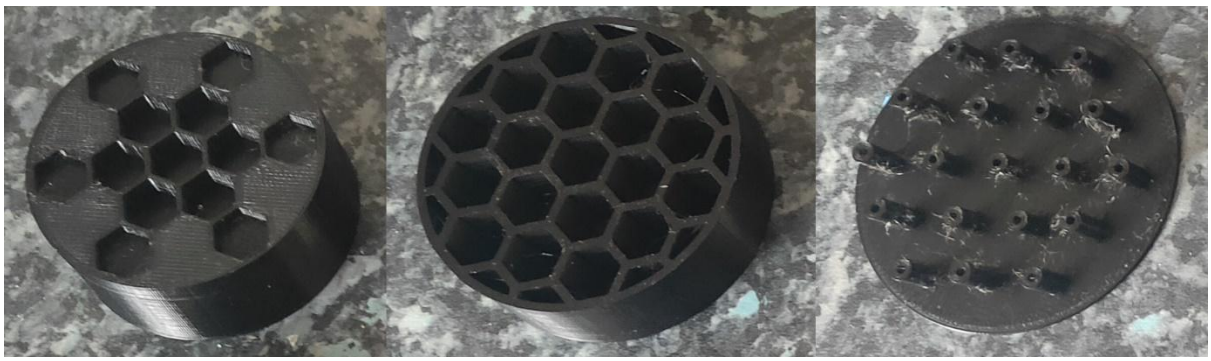


Figure 124 Triple parallel Helmholtz resonator specimen additively manufactured in two halves to allow for different neck and cavity lengths. Complete sample bonded with epoxy adhesive (left), back wall terminations are introduced at 3 different heights to allow for different cavity lengths. Specimen was assembled by joining honeycomb back section (middle) and front plate with inlet necks (right).

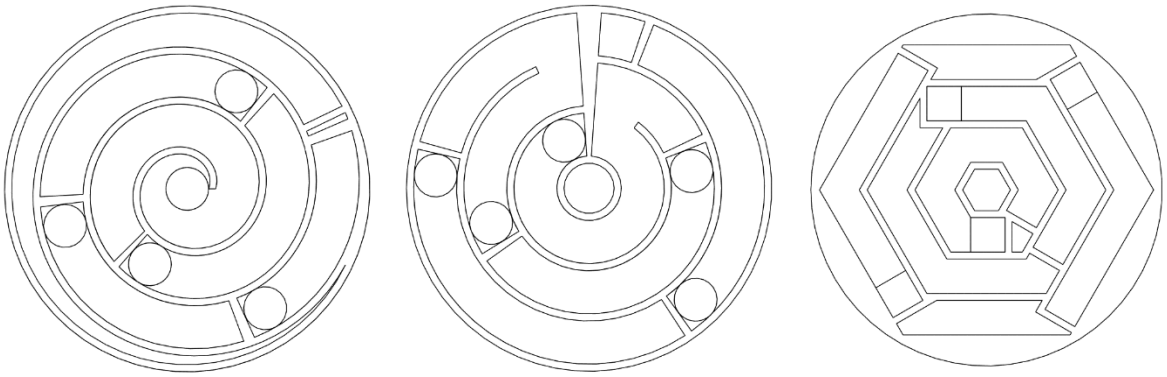


Figure 125 Parallel quarter-wavelength resonator configurations; spiral (left), concentric circular (middle) and concentric hexagonal (right).

Both the HR and QWR configurations offer analytical solutions for theoretical prediction of the sound absorption coefficient α at a given frequency. The absorption coefficient is defined as the ratio between the sound energy absorbed by a material and the incident energy. The value is determined from the acoustic impedances of the media that the sound wave passes through, in this case the resonator and the ambient air.

The acoustic impedance of a Helmholtz resonator with extended neck is given by [152]

$$Z_{HR} = Z_{RN} + Z_{RC} \tag{5.1}$$

where Z_{RN} is the impedance of the resonator neck and Z_{RC} the impedance of the resonator cavity.

The impedance of the Helmholtz resonator may also be expressed as a function of the normalised specific acoustic resistance, R_{HR} and the normalised specific reactance, X_{HR} by the following:

$$Z_{HR} = z_0(R_{HR} + jX_{HR}) \tag{5.2}$$

$$R_{HR} = \frac{\sqrt{8\eta\rho\omega}((l_N/2r_N)+1)}{\varepsilon_T \rho c} \tag{5.3}$$

$$X_{HR} = \frac{1}{\varepsilon_T} k(l_N + \delta) - \frac{1}{k l_{RC}} \tag{5.4}$$

where ρ is the density of the air, η is the dynamic viscosity of the air, c is the soundwave velocity, k the acoustic wavenumber and $\omega = 2\pi f$ is the angular frequency.

The perforation ratio, ε_T is given by

$$\varepsilon_T = \frac{S_n}{A_0} \tag{5.5}$$

where S_n is the cross-sectional area of the resonator neck and A_0 the cross-sectional area of the cavity.

A modified resonator cavity length, l_{RC} is used to account for the effective length of the neck, l_e where

$$l_{RC} = l_N + l_C - \varepsilon_T l_e \tag{5.6}$$

and

$$l_e = l_N + \delta$$

$$\text{with end correction } \delta = 1.7 r_N \tag{5.7}$$

The length of the resonator neck l_N, the length of the resonator cavity l_C, the radius of the neck r_N and the radius of the cavity r_C may then be used to calculate the absorption coefficient α at frequency f using

$$\alpha = 1 - \left|\frac{Z_{HR} - z_0}{Z_{HR} + z_0}\right|^2 \tag{5.8}$$

where $z_0 = \rho c$ is the specific acoustic impedance of the air.

For multiple HRs combined in parallel the acoustic impedance of each individual resonator is calculated using equations 5.2-5.7. The total impedance of the parallel resonator array can then be calculated from the individual impedances using

$$Z_{Total} = \left(\sum_{i=1}^{n} \frac{\phi_n}{Z_n}\right)^{-1} \tag{5.9}$$

where n is the number of parallel resonators and ϕ_n is the area ratio of each resonator to the total active area of the unit cell. The absorption coefficient α can be calculated from the total impedance of the array using

$$\alpha = 1 - \left|\frac{Z_{Total} - z_0}{Z_{Total} + z_0}\right|^2 \tag{5.10}$$

For a quarter-wavelength resonator tube with rigid back wall termination arranged perpendicular to the inlet hole, the surface acoustic impedance at the point immediately after the junction between the cavity and the inlet opening is given by [153-155]

$$Z(1) = -jZ_{c1}\cot(k_{c1}l_1) \tag{5.11}$$

where $Z_{c1} = \rho_{eff}c_{eff}$ is the effective impedance and $k_{c1} = \omega/c_{eff}$ is the effective complex wave number of the tube.

The effective density ρ_{eff} includes inertial and viscous contributions and is given by

$$\rho_{eff} = \frac{k_c Z_c}{\omega} = \frac{k_0 Z_0}{\omega}\frac{1}{\psi_v} = \frac{\rho}{\psi_v} \tag{5.12}$$

The effective speed of sound c_{eff} is given by

$$c_{eff} = \frac{\omega}{k_c} \tag{5.13}$$

where $k_0 = \omega/c_0$ is the wave number of the air with angular frequency ω and speed of sound c_0, $z_0 = \rho_0 c_0$ is the specific acoustic impedance of the air.

The complex wavenumber k_c is given by

$$k_c^2 = k_0^2\left(\frac{\gamma-(\gamma-1)\psi_h}{\psi_v}\right) \tag{5.14}$$

and the specific impedance Z_c in the cavity is

$$Z_c^2 = \frac{Z_0^2}{\psi_v(\gamma-(\gamma-1)\psi_h)} \tag{5.15}$$

where γ is the specific heat ratio.

The functions ψ_v and ψ_h correspond to viscous and thermal components obtained through linearised Navier-Stokes modelling and can be evaluated using:

$$\psi_i = k_i^2\sum_{m=0}^{\infty}\left[(\alpha_m m')^{-2}\left(1-\frac{\tan(\alpha_m W/2)}{\alpha_m W/2}\right)+(\beta_m m')^{-2}\left(1-\frac{\tan(\beta_m H/2)}{\beta_m H/2}\right)\right] \tag{5.16}$$

where the index i is replaced by subscripts v and h for viscous and thermal fields respectively, $\alpha_m = \sqrt{k_i^2-(2m'/W)}$, $\beta_m = \sqrt{k_i^2-(2m'/H)}$ and $m' = (m+1/2)\pi$.

The viscous and thermal wavenumbers for the system are calculated by

$$k_v^2 = -j\omega\frac{\rho}{\mu} \tag{5.17}$$

$$k_h^2 = -j\omega\frac{\rho C_p}{k} \tag{5.18}$$

where μ is the dynamic viscosity, k the thermal conductivity and C_p the specific heat at constant pressure of air.

The surface acoustic impedance at the inlet hole is found using

$$Z(1_a) = Z_{c1_{aJCA}}\frac{-jZ(1)\cot\left(k_{c1_{aJCA}}l_{inlet}\right) + Z_{c1_{aJCA}}}{Z(1) - jZ_{c1_{aJCA}}\cot\left(k_{c1_{aJCA}}l_{inlet}\right)} \tag{5.19}$$

where $Z_{c1_{aJCA}} = \rho_{JCA}c_{JCA}$ is the effective impedance of the inlet, $k_{c1_{aJCA}} = \omega/c_{JCA}$ the effective wave number of the inlet and $c_{JCA} = \sqrt{K_{JCA}/\rho_{JCA}}$ the effective speed of sound.

The effective mass density ρ_{JCA} and bulk modulus K_{JCA} are calculated by

$$\rho_{JCA} = \frac{\alpha_\infty\rho_0}{\phi} + \frac{\sigma}{i\omega}\sqrt{1 + \frac{4j\alpha_\infty^2\eta\rho_0\omega}{\sigma^2\Lambda^2\phi^2}} \tag{5.20}$$

$$K_{JCA} = \frac{\gamma P_0/\phi}{\gamma - (\gamma-1)\left(1 + \frac{8\eta}{j\rho_0\omega N_p\Lambda'^2}\sqrt{1 + \frac{j\rho_0\omega N_p\Lambda'^2}{16\eta}}\right)^{-1}} \tag{5.21}$$

where η is the dynamic viscosity and N_p Prandtl's number.

The airflow resistivity of the inlet is given by

$$\sigma = \left[\left(\frac{2l_{inlet}}{r_{eq}} + 1\right)\frac{\alpha_s R_s}{\phi}\frac{1}{l_{inlet}}\right]\frac{1}{\phi_{sec}} \tag{5.22}$$

where the specific airflow resistance $R_s = \frac{1}{2}\sqrt{2\eta\omega\rho_0}$ and $\phi_{sec} = A_1/A_{uc}$ is the ratio between the cross-sectional area of the resonator channel and the overall unit cell. The factor α_s is determined empirically and is equal to a value of approximately 4 for an inlet with sharp edges. The parameters Λ and Λ' can be set equal to the equivalent radius of the inlet hole r_{eq} and ϕ is the tortuosity, which is set to 1 for circular inlets.

The total input surface acoustic impedance of the resonator is found by

$$Z_{in} = [Z(1_a)/Z_0]/\phi_{sec} \tag{5.23}$$

The absorption coefficient value α at frequency f may then be calculated using the resonator geometry parameters of the length of the inlet l_{inlet}, the length of the cavity l_1 and the radius of the inlet/cavity r_{eq}. For cavities of square or rectangular cross section the radius of the equivalent circle with equal area was calculated using the width W and height H.

The value of the absorption coefficient is determined from the complex input impedance using

$$\alpha = \frac{4Re(Z_{in})}{|Z_{in}|^2 + 2Re(Z_{in}) + 1} \tag{5.24}$$

In the case where the QWR labyrinth is separated into multiple resonator chambers arranged in parallel, with separate inlets, varied lengths and identical cross sections, the total combined input impedance of the parallel array can be found from the individual impedances using

$$Z_{Total_{in}} = \left(\sum_{i=1}^{n} \frac{\phi_{sec}/n}{Z(n_a)/Z_0}\right)^{-1} \tag{5.25}$$

where n is the number of resonators in parallel and $Z(n_a)$ is the surface acoustic impedance at the inlet hole for the n^{th} resonator. The absorption coefficient of the array can be calculated from the total input impedance using

$$\alpha = \frac{4Re\left(Z_{Total_{in}}\right)}{\left|Z_{Total_{in}}\right|^2 + 2Re\left(Z_{Total_{in}}\right) + 1} \tag{5.26}$$

To validate the analytical solutions examples of the HR and QWR designs were additively manufactured using fused-deposition modelling (FDM) on a Prusa i3 MK3S printer from PLA filament. The use of FDM allowed the cavities to be fully sealed so that the inlet holes were the only path for air to enter the resonators. The honeycomb or labyrinth geometry was built in the z-direction with the inlet face flat against the build plate. The back walls were sealed over with bridging manoeuvres, printing the entire specimen in one piece as demonstrated in Figure 126. This was the case with all samples except the parallel HR arrangements which were printed in 2 halves and then sealed with epoxy adhesive as shown in Figure 124. All samples were circular in overall cross section with an outer diameter of 50.8mm. In all cases the front and back walls were both 1mm thick.



Figure 126 Spiral QWR open specimen (left) and same specimen sealed over during the FDM process (right)

The samples were tested on an impedance tube built to ASTM E-1050 [145], featuring an aluminium tube with a loudspeaker mounted at one end and a series of microphones positioned to measure the incident and reflected sound pressure levels at set distances along the tube, upstream and downstream of the specimen. The test rig is shown in Figure 127. The specimen was mounted in the rear half of the tube with the front face flush to the opening of the tube. A solid back wall termination was provided by an aluminium plunger with a rubber o-ring to seal against leakages. The two halves of the tube were then sealed together and clipped in place before propagation of a plane standing wave down the length of the tube. Table 17 contains the geometry parameters of the microphone placement. As shown in Figure 128, two microphone configurations were used simultaneously to cover two overlapping frequency ranges and extend the frequency bandwidth of the measurements. For the lower frequency response a larger separation gap was needed between microphones. Microphones 1 and 3 were used to measure the 85-700Hz range, while microphones 2 and 3 captured the high-frequency response between 250-3000Hz and results were combined in post-processing to cover the entire frequency range between 85-3000Hz. The acoustic signal was controlled by a LabView interface and the measurements post-processed using a transfer function method according to ASTM E-2611 [146]. The ratio of sound energy reflected back from the specimen was calculated by evaluating the complex reflection coefficient [147,148]:

$$R(f) = \frac{H - e^{-jks}}{e^{jks} - H} e^{j2k(l+s)}, \tag{5.27}$$

where H is the acoustic transfer function, k is the wave number, s the spacing between the two active microphones, and l the distance from the sample to microphone 3. The absorption coefficient at frequency f was then found using [149,150]:

$$\alpha(f) = 1 - |R(f)|^2 \tag{5.28}$$

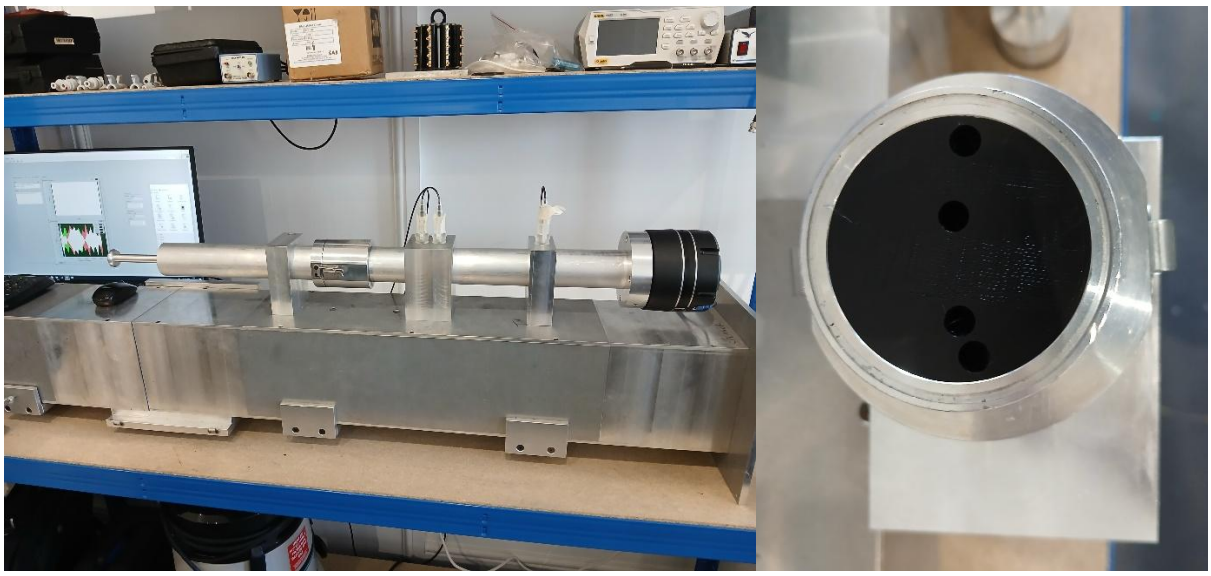


Figure 127 Impedance tube test rig (left) and parallel QWR specimen mounted in the tube (right)

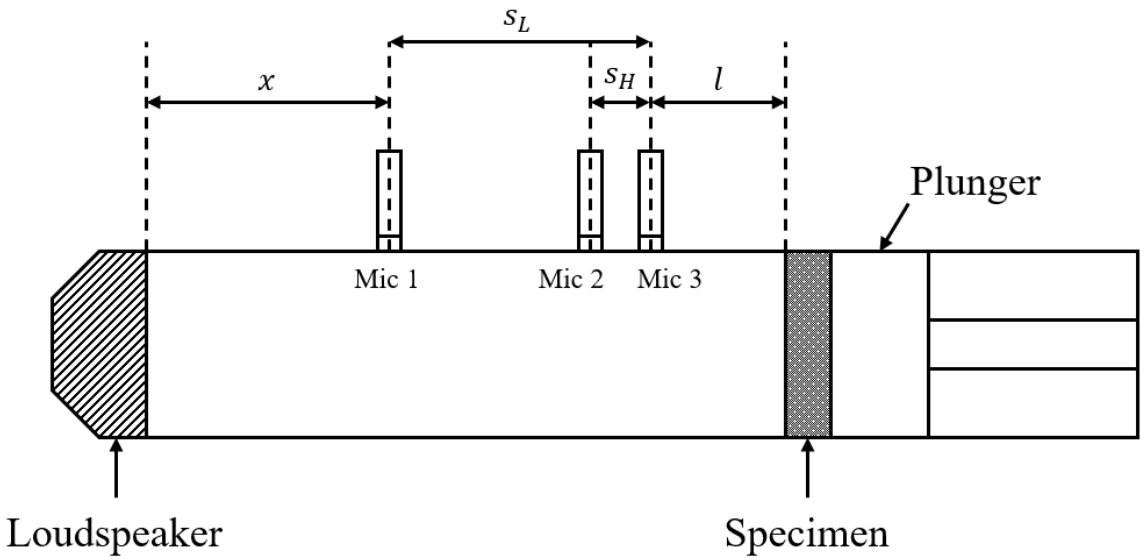


Figure 128 Diagram of impedance tube with microphone placement configuration

Table 17 Impedance tube geometry parameters

Geometry Parameter	Distance (mm)
l	140
x	160
s_L	200
s_H	30

A comparison between analytical and experimental absorption responses using the single Helmholtz resonator-based unit cell is displayed in Figure 129 for examples targeted at 100Hz intervals between 400-800Hz. In each case the apothem of the honeycomb was set at 4.4mm. This divided the samples into 19 complete identical resonators with no inactive cells to maximise the active area for the available sample size. The radius of the resonator neck, r_n was locked at 0.7mm. Below 0.6mm inaccuracies were observed with the diameter and circularity of the perforations due to the FDM process, as compression from the nozzle designed to aid first layer adhesion caused melted polymer to encroach on the perforations. For all examples the cavity length, l_c was 5mm and the neck length, l_n was varied to set the frequency of the absorption peak.

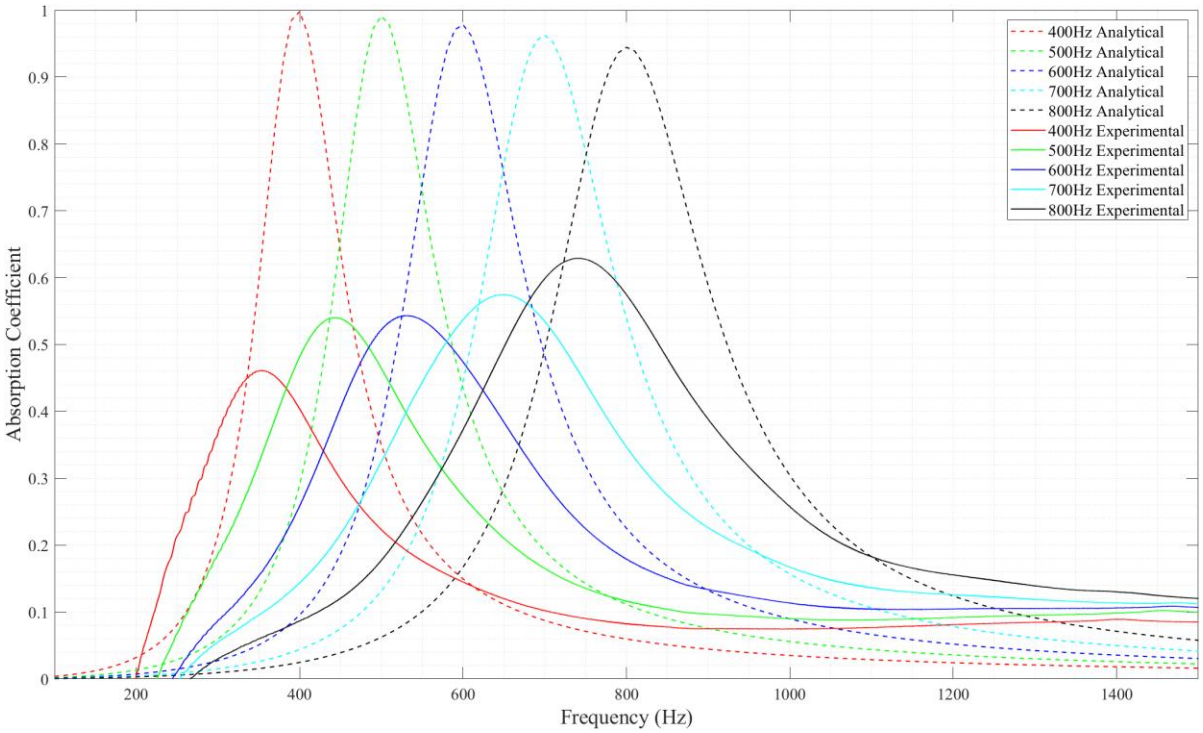


Figure 129 Comparison between analytical and experimental absorption coefficient responses for single Helmholtz resonator-based honeycomb samples manufactured via FDM. In all cases r_n=0.7mm, A_0=67.065mm^2, l_c=5mm. For 400, 500, 600, 700 and 800Hz specimens l_n=18, 13.7, 11, 9 and 7.5mm respectively.

In all cases the magnitude of the experimental absorption peak was below expected values. The strongest experimental peak offered 63% absorption, while the analytical solution predicted all peaks to exceed 94%. In addition all responses were shifted to lower frequencies than expected, suggesting inaccuracies in the manufactured neck lengths. Significant improvements were observed from manufacturing specimens via stereolithography (SLA) on an Elegoo Saturn 3 Ultra from 8K photopolymer resin. The example targeted at 800Hz provided a peak of 86% absorption at 790Hz, printing from an identical .stl geometry file. A comparison between FDM and SLA specimens is displayed in Figure 130.

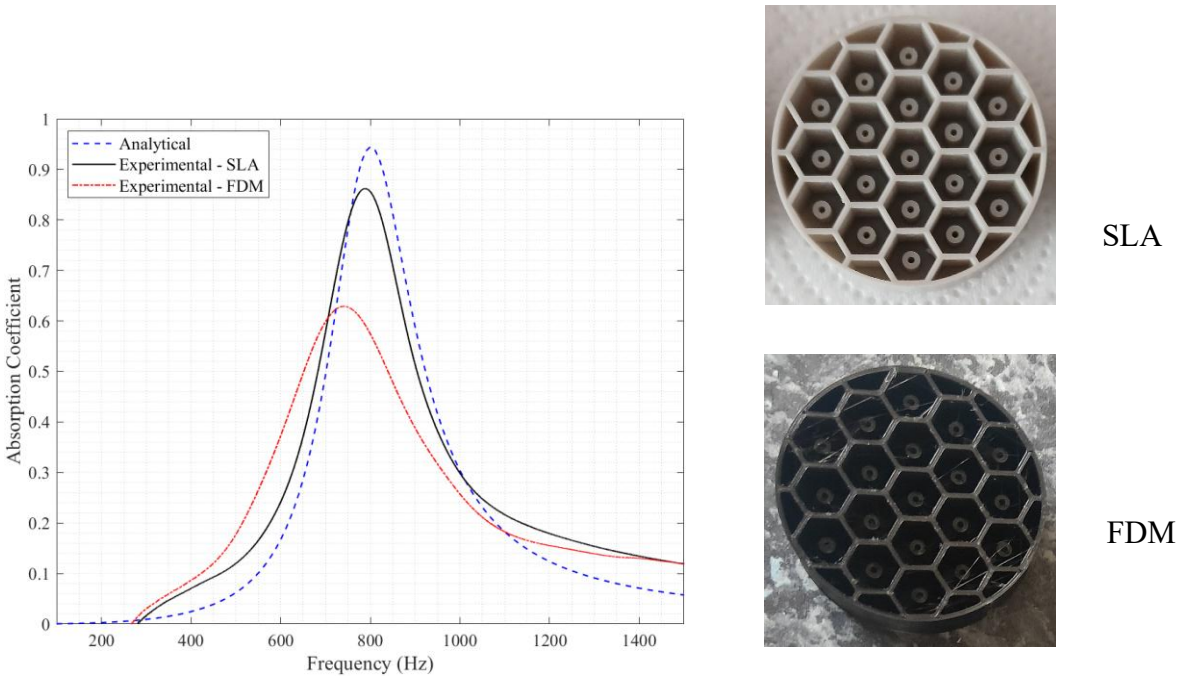


Figure 130 Comparison between specimens manufactured via FDM and SLA for HR targeted at 800Hz

The absorption response of a parallel arrangement of 6 Helmholtz resonators designed to target the 500-1000Hz frequency range is shown in Figure 131(a). Each cell had a cross-sectional area of A_0 = 25mm^2, neck radius r_n = 1mm and cavity length l_c = 5mm. The specimen was 42mm in overall thickness and printed via FDM. An equivalent version of the same resonators arranged perpendicular to the direction of the wave within the tube was also designed and manufactured for comparison to determine if the transverse arrangement could be used to reduce the size of the specimen. The design used an arrangement of concentric hexagonal channels containing extended resonator necks with periodic support material. The printed specimen, shown in Figure 132 had an overall thickness of 6mm. The specimen arranged in parallel with the wave direction showed a strong initial absorption peak, exceeding 80%, however with reduced range compared to the analytical solution. The

increased amplitude compared to previous single HR specimens may have resulted from a larger neck radius leading to a more accurate FDM print. Transverse arrangement showed an overall characteristic in agreement with the parallel version, though at a weaker magnitude, offering peak absorption below 30%.

(a) (b)

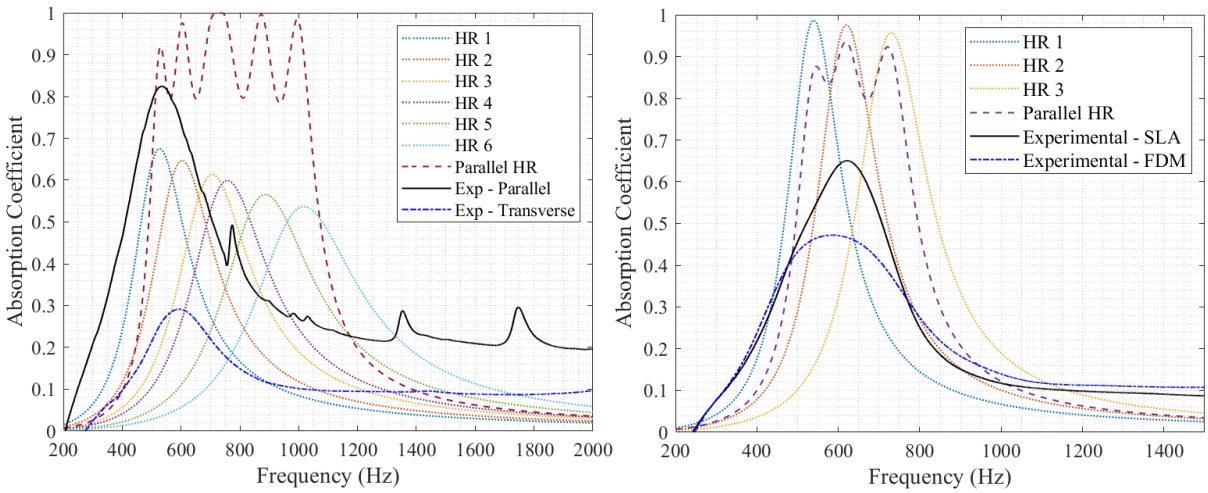


Figure 131 Experimental vs analytical absorption responses for parallel HR specimens; (a) 6 resonators, A_0=25mm^2, r_n=1mm, l_c=5mm, l_{n1}=36mm, l_{n2}=31mm, l_{n3}=26mm, l_{n4}=24mm, l_{n5}=20mm, l_{n6}=17mm; (b) 3 resonators, A_0=67.065mm^2, r_n=0.7mm, l_c=5mm, l_{n1}=12.5mm, l_{n2}=10.5mm, l_{n3}=8.5mm.

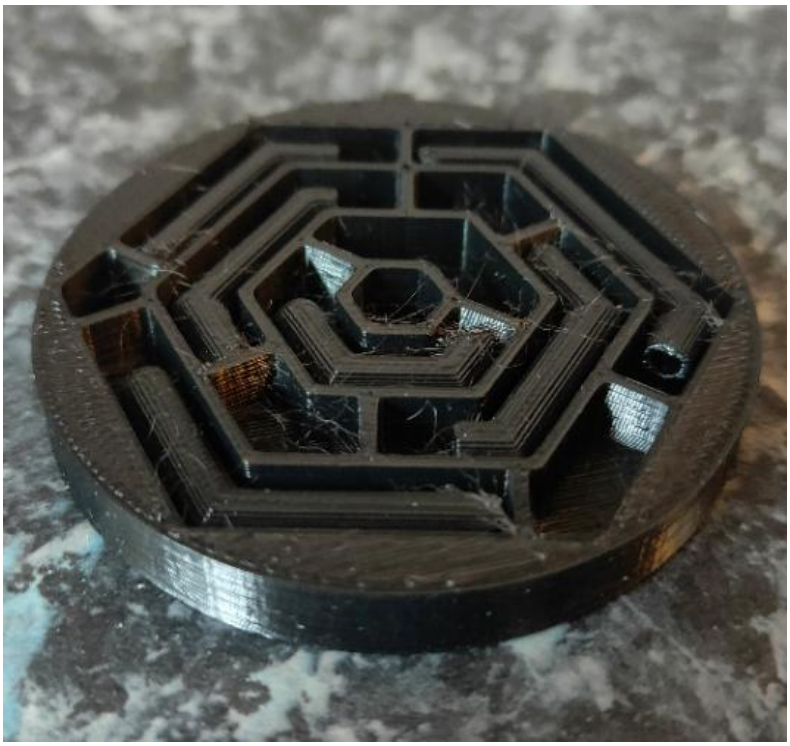


Figure 132 Parallel configuration of 6 Helmholtz resonators arranged transverse to the sound wave direction.

With 6 HRs in parallel the lower frequency resonators dominated in the experimental response despite being allocated the same amount of active area as the others. Absorption results for a design featuring 3 parallel HRs oriented parallel to the wave direction are displayed in Figure 131(b) for specimens manufactured using both FDM and SLA. A more symmetrical response was obtained centered closer to the midpoint of the analytical solution. However the amplitude was weaker than expected for both manufacturing approaches, giving maximum absorption of 47 and 65% for FDM and SLA specimens respectively.

Experimental absorption responses for the spiraled QWR design manufactured by FDM are shown in Figure 133. For each example the channel width, w and height, h were both set to 10mm to ensure a cross-sectional area of A_0 = 100mm^2, with an overall specimen thickness of 12mm. The wall thickness, t was 1mm leading to a maximum possible spiral length within the 50.8mm diameter sample of 638.2° at l = 94.77mm. Specimens were tested at the maximum possible channel length and hence lowest peak frequency for the given channel width, in addition to examples where a 1mm thick back wall termination was introduced to shorten the spiral to 585° (l = 75.75mm) and 540° (l = 60.60mm) to examine ability to tune for a given target frequency. Corresponding analytical solutions provided absorption peaks at 890, 1110 and 1410Hz. Experimental absorption peaks were all within 20Hz of expected values, occurring at 908, 1130 and 1415Hz. All specimens provided peaks stronger than those predicted, reaching absorption coefficient values of 0.99, 0.98 and 0.95 respectively.

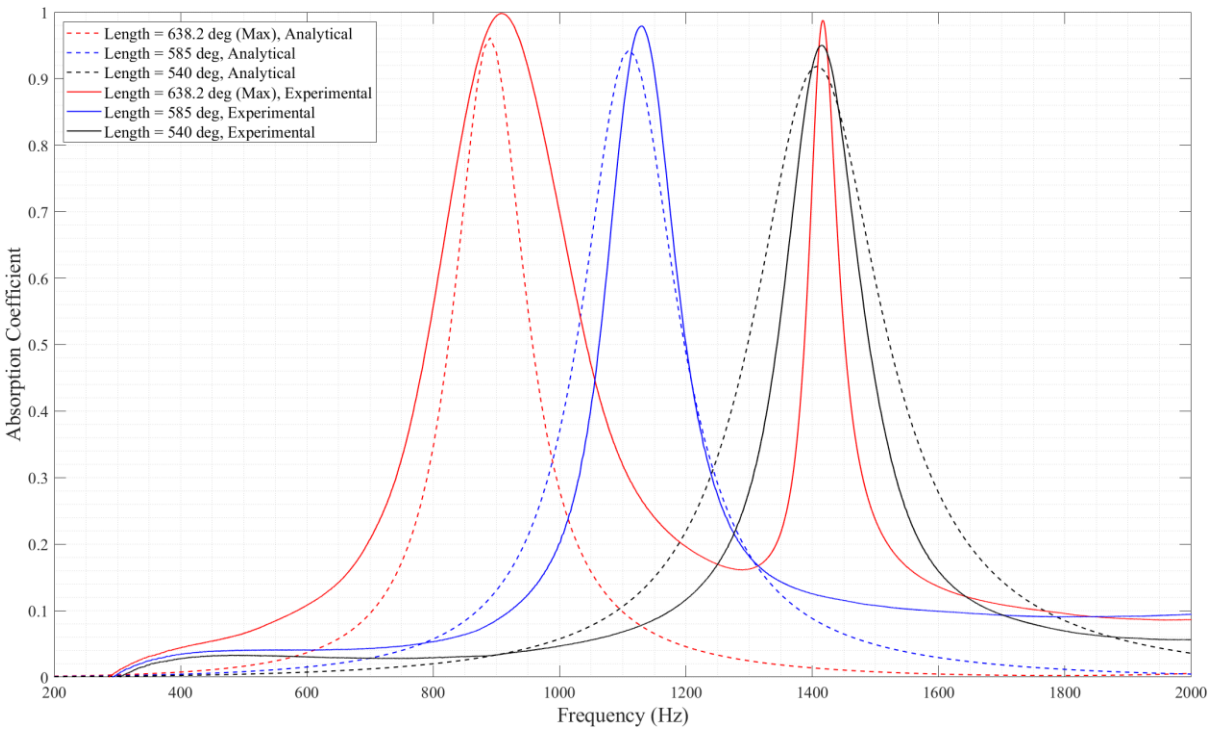


Figure 133 Comparison of analytical and experimental absorption coefficient responses for spiral QWR specimens manufactured via FDM. In all cases w=h=10mm, t=1mm, $l_{inlet} = 1mm$ and overall specimen thickness = 12mm.

Spiral quarter-wavelength resonator samples with rectangular channel cross sections were manufactured via FDM to evaluate the effects of aspect ratio on the absorption performance. Theory suggests the analytical solution is likely to be more accurate the closer the channels are to a 1:1 square section. For combination with structural optimisations the width and height of the channels are to be set as input parameters that may vary. To ensure this does not compromise absorption, tests were carried out in which the aspect ratio was varied at intervals between 1:1 and 1:6 while the overall cross-sectional area of the channels remained fixed at approximately 100mm^2. For each test sample the maximum possible spiral length was used for the given channel width, resulting in absorption peaks at the minimum frequency achievable at the given ratio for the unit cell diameter of 50.8mm. The locations and amplitudes of the experimental peaks are compared against the analytical predictions in Table 18. Peaks exceeding 90% absorption were maintained for all samples down to a ratio of 1:2. The lowest aspect ratio to maintain at least 70% absorption was 1:4, which provided a peak of 71.66% at lowest frequency of 321.36Hz with channel width 5mm and height 20mm. At overall material thicknesses greater than 20mm multiple levels of stacked resonator channels are recommended to ensure sufficient aspect ratio across the cross section of each individual level.

Table 18 Absorption peak locations for spiral quarter-wavelength resonator samples with varied aspect ratio

Aspect Ratio	Channel Width (mm)	Channel Height (mm)	Cross Sectional Area (mm^2)	Maximum Spiral Length (°)	Maximum Channel Length (mm)	Analytical Max Absorption		Experimental Max Absorption	
						Frequency (Hz)	Absorption Coefficient	Frequency (Hz)	Absorption Coefficient
1.00	10	10	100	638.2	93.22	900	0.9577	919.83	0.9943
0.82	9	11	99	720.0	114.26	740	0.9694	750.45	0.9959
0.67	8	12	96	820.0	139.14	610	0.9659	616.27	0.9592
0.50	7	14	98	945.0	169.52	500	0.9827	505.83	0.9261
0.35	6	17	102	1105.7	211.75	400	0.9947	404.95	0.8953
0.25	5	20	100	1319.6	266.83	320	0.9645	321.36	0.7166
0.20	4.5	22	99	1456.4	301.04	280	0.9938	284.85	0.5791
0.16	4	25	100	1620.0	346.86	240	0.9704	240.19	0.5152

The labyrinth arrangements featuring concentric circles and concentric hexagons were tested at the minimum aspect ratio limit, with channel width of 5mm and height 20mm for comparison against the spiraled equivalent at the maximum channel length able to fit in the unit cell. For the hexagonal design the maximum length was 270.13mm, while the circular layout allowed for an increased length of 297.89mm. Both labyrinth designs outperformed the spiral layout, offering greater absorption at a lower minimum frequency. The concentric hexagon design gave 79.1% absorption at 317.7Hz, while the circular design provided 98.6% absorption at 282.6Hz. The experimental absorption characteristics for the labyrinth configurations are compared against analytical predictions in Figure 134.

(a) (b)

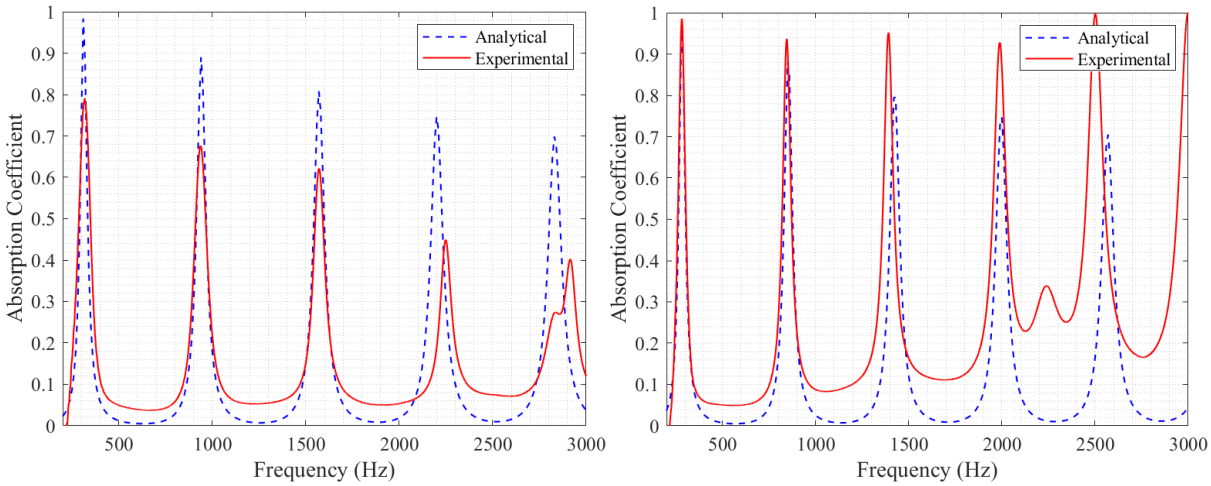


Figure 134 Comparison of analytical and experimental absorption coefficient responses for labyrinth QWR specimens; (a) hexagonal, w=5mm, h=20mm, l=270.13mm, l_{inlet}=1mm; (b) circular, w=5mm, h=20mm, l=297.89mm, l_{inlet}=1mm.

The concentric circular labyrinth layout was divided into arrangements of multiple QWRs combined in parallel to assess the strength of the absorption and predictability of the responses. The absorption-frequency characteristics that resulted from splitting the labyrinth into 4 and 5 resonator cavities are shown in Figure 135(a) and (b) respectively. Both cases used an aspect ratio of 0.35 for the cavity cross sections, with a width, w of 6mm and height, h of 17mm. The front and rear walls were both 1mm thick, giving an overall material thickness of 19mm. The experimental measurements were compared against the analytical solutions of the individual resonators in addition to the predicted combined parallel response. Both absorbers offered peaks greater than 90% and provided ranges of absorption that exceeded the analytical prediction for resonators combined in parallel. Arranging 4 QWRs in parallel gave experimental measurements of no less than 50% absorption between 896-1265Hz, while for 5 parallel resonators the 50% absorption threshold was 1007-1515Hz.

(a) (b)

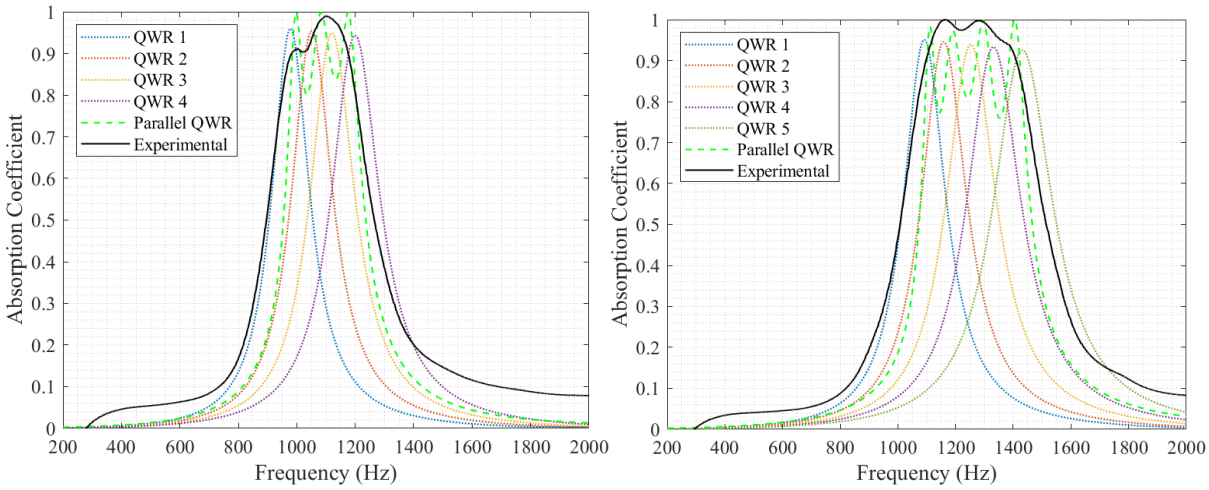


Figure 135 Analytical vs experimental absorption responses for labyrinth QWRs arranged in parallel; (a) 4 resonators, w=6mm, h=17mm, l_1=86mm, l_2=80mm, l_3=75mm, l_4=70mm, l_{inlet}=1mm; (b) 5 resonators, w=6mm, h=17mm, l_1=77mm, l_2=72.5mm, l_3=67mm, l_4=63mm, l_5=58.5mm, l_{inlet}=1mm.

Of the HR and QWR configurations the HR design is likely to provide strongest structural performance, however high absorption peaks can only be assured at discrete target frequencies with limited range. Broadband responses from parallel combinations of multiple Helmholtz resonators were less predictable with the honeycomb layout, often providing insufficient absorption performance for practical design purposes. The strength of the absorption improved on manufacturing samples using SLA rather than FDM, switching from a 50x50µm print resolution to 19x24µm. The improvements were likely due to the finer resolution, fewer build defects and a more homogeneous base material. The QWR layouts gave stronger and more reliable absorption peaks at restricted unit cell sizes for FDM samples, likely because they can be manufactured with greater accuracy via FDM due to the increased dimensions of the perforations. QWR specimens showed absorption peaks above 90% and broader responses that offered better agreement with analytical predictions. The responses were highly controllable and could be fine-tuned by adjusting the length of each individual resonator. Parallel QWR combinations also showed improved performance over parallel HRs that more consistently agreed with or exceeded analytical predictions. An inherent limitation of the HR is a tight absorption range for each individual absorber, whereas the QWR offered wider bandwidth in the combined parallel responses at lower frequencies with fewer absorbers needed. The channels can be arranged transverse to the direction of the inlet to achieve considerably lower overall material thicknesses without reducing the absorption performance. The concentric circular labyrinth design offered the greatest maximum channel length, consistently providing the strongest absorption peaks at the lowest minimum frequencies and was chosen as the base configuration for the multi-objective optimisations.

5.1.2 Dataset Generation

The dataset used to train the deep learning network was created from a set of sound absorption-frequency responses calculated using the analytical solution and a set of structural force-displacement responses determined by a non-linear FEA model. The absorption amplitude of the chosen QWR unit cell depended on the cross-sectional area of the resonator channel, while the frequency at which the absorption peak occurred depended on the length of the channel. The models were designed with the goal of finding the geometry that provides the strongest first harmonic absorption peak at the lowest possible frequency for a given relative specific stiffness and relative specific strength requirement. The location of the first harmonic is set by the maximum resonator channel length that can fit within the unit cell. If the first harmonic can be optimised, by theory the QWR absorber will be able to achieve the same or greater amplitude at any frequency above the initial result by introducing a back wall termination to shorten the length of the resonator channel, maximising the sound absorption performance for the given structural target at the frequency of interest. In addition, achieving a first harmonic at the lowest possible frequency leaves the greatest possible channel length available for dividing the labyrinth into multiple resonators, leading to the broadest possible absorption range in parallel QWR designs. A MATLAB script was written to output the analytical sound absorption response automatically for each design point in the dataset. The script featured a function that calculated the maximum length, l_{max} of the QWR labyrinth cavity for a given width, w, wall thickness, t, and unit cell apothem size, a. The absorption coefficient was then calculated from l_{max}, w, a, the channel height h and the inlet length l_{inlet} using equations 5.11-5.24 at 20Hz intervals between 0-1000Hz.

At each observation in the dataset the equivalent structural responses for the given geometry were generated using a parametric non-linear FEA model developed in ANSYS Mechanical, loading a single representative unit cell of the metamaterial under uni-axial compression. Non-linear material data was measured experimentally via electro-mechanical testing on an Instron 3369 with a 50kN load cell according to standards ASTM D695 [156] and ASTM D638 [157]. The base material was ColorFabb XT-CF20, a carbon fibre reinforced polymer filament. For each test 5 specimens were manufactured via FDM on a Prusa i3 MK3S. The specimens were post annealed for 2 hours at 70°C and then conditioned for 2 weeks at room temperature and humidity prior to testing. All coupons were built in the z-direction with a layer height of 0.15mm and 100% infill. The infill lines for the tensile coupons alternated between ±45° to the direction of loading, while the layers for the compression coupons were orientated at 90° to the loading. The tensile coupons were type 1 dogleg while the

compression coupons were cylindrical with diameter 12.7mm and height 25.4mm. The as-built dimensions of the test coupons were measured using a digital micrometer. The tensile specimens were loaded at a rate of 5mm/min until fracture and the strain measured by a digital extensometer with a gauge length of 50mm. The compression specimens were loaded at 1.3mm/min until rupture and a linear variable displacement transducer (LVDT) was used for the strain measurement, held in place between the compression platens using an additively manufactured jig. In both cases the stress-strain relations were averaged by taking the mean across the 5 specimens. The measured structural properties in uniaxial tension and compression are summarised in Table 19.

Table 19 Measured mechanical properties of ColorFabb XT-CF20 FDM base material

Testing Standard	Tension ASTM D638	Compression ASTM D695
Young's Modulus (MPa)	3470.66	1832.32
Yield Strength (MPa)	36.93	32.70
Ultimate Strength (MPa)	49.05	149.39
Stress At Fracture (MPa)	46.71	146.05
% Elongation At Fracture	4.81	64.23

For the FEA model the compression material data was converted from engineering stress-strain to true stress-strain and the plastic strain separated from the elastic strain using the formulae:

$$\sigma = \sigma_{eng}(1 + \varepsilon_{eng}), \varepsilon = \ln(1 + \varepsilon_{eng}), \varepsilon_{plastic} = \varepsilon - \sigma/E \tag{5.29}$$

The true stress is plotted against the logarithmic strain in Figure 136. The relationship between the true stress and the plastic strain was then down sampled to 100 datapoints and used as input for a non-linear isotropic material model in ANSYS.

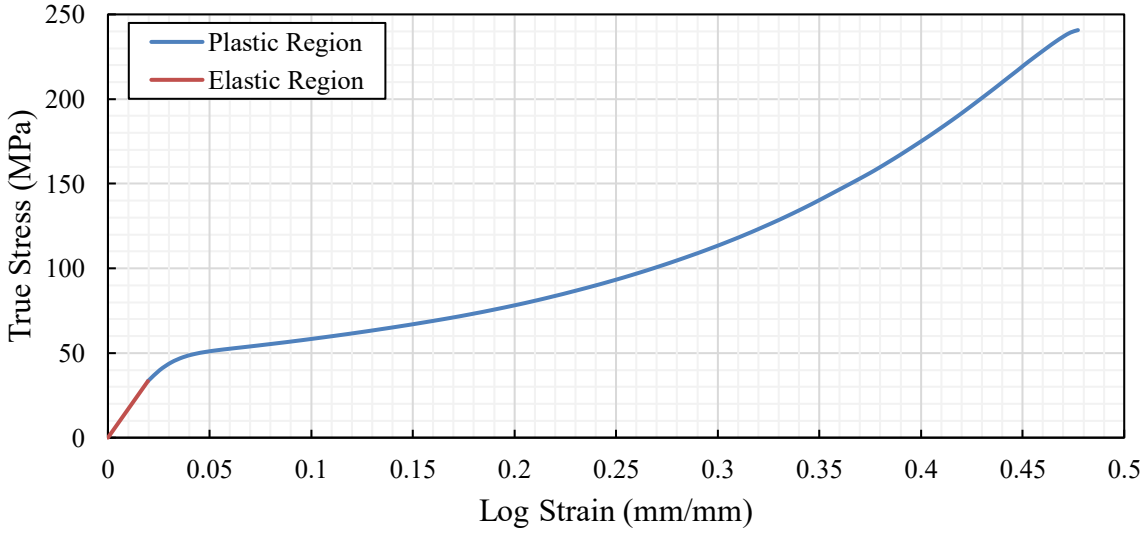


Figure 136 ColorFabb XT-CF20 base material uni-axial compression test data

The metamaterial unit cell geometry featured the concentric circular labyrinth layout shown in Figure 123(c) modified to allow hexagonal tessellation between cells to form a closed-cell honeycomb lattice. The apothem of the hexagonal unit cell was set at 25.4mm to allow testing of an equivalent specimen for sound absorption on the 50.8mm diameter impedance tube. The geometry was created to allow for parametric updates, ensuring the model remained valid at any point within the design space by progressively generating the wall bodies of the interior labyrinth outwards from the centre. The centre circle remained 5mm in diameter for all design points. The locations of the remaining concentric circles were calculated from the channel width, w and wall thickness, t. Any inactive geometry that fell outside the 50.8mm diameter outer limits was then removed using a Boolean operation. The remaining interior walls were then combined with the outer circular and hexagonal walls to form a single solid body. Finally front and back walls were added to seal the unit cell. The thicknesses of the front and back walls were kept constant at 1mm. A cutaway example unit cell is shown in Figure 137(a). The geometry was meshed with 3-dimensional solid elements of quadratic order. A fully fixed support constraint was applied to the bottom face to prevent free-body motion and a displacement load applied to the top face in the negative z-direction as shown in Figure 137(b). The load was ramped on linearly from 0-2mm over 2 load steps, each of size 1mm. The step controls were set to initialise the solve with 30 substeps and use the lowest number of substeps necessary for convergence, between a minimum of 10 and a maximum of 100. Large deflections were enabled to allow calculation of post-yield plastic deformation. The reaction force was probed at the location of the fixed boundary condition. The geometry parameters were provided as a table of design points and the model automatically cycled through each non-linear stress analysis, outputting the reaction force values at displacement intervals between 0-2mm, sampling every 0.04mm.

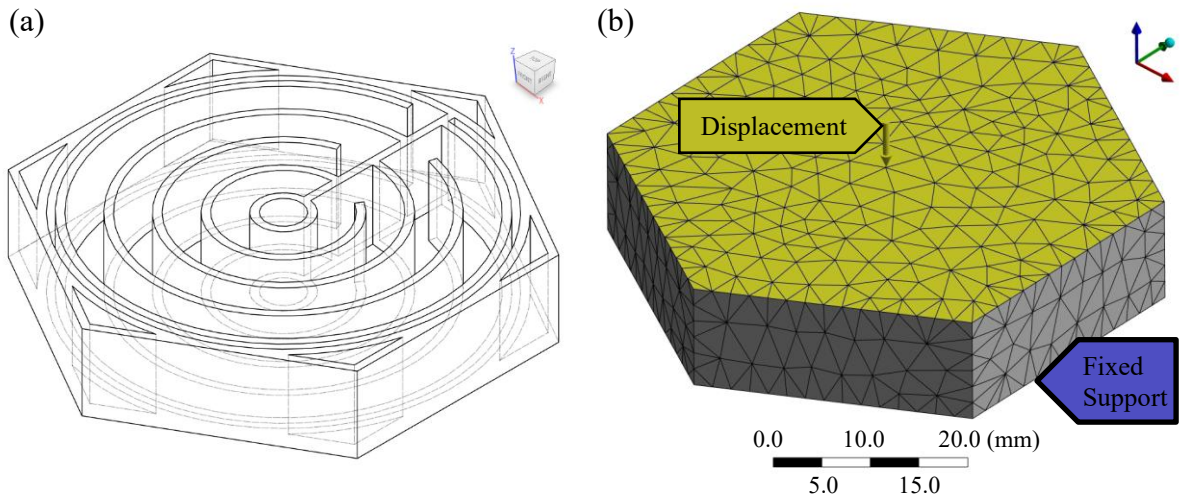


Figure 137 (a) Cutaway example unit cell, t=1mm, w=4mm, h=10mm; (b) FEA model mesh and boundary conditions

To determine a suitable element size a mesh convergence study was carried out, repeating the simulations at element sizes between 0.5-3mm for the lowest volume, greatest volume, weakest and strongest possible geometries that may appear within the design space. The geometry parameters for each of these cases is shown in Table 20. The relationship between the number of elements and the maximum force recorded during the analysis is plotted for each test case in Figure 138. All plots show evidence of convergence as the element size approaches the minimum value in the study. Models with coarser meshes all lead to overestimates of yield strength, the coarser the mesh the greater the overestimation. The greatest discrepancies in the results for the coarser meshes were observed in the non-linear sections of the force-displacement responses. Differences in slope remained negligible within the linear region in all test cases. A 1mm mesh gave maximum force values to a mean error of 3.54% when compared against the minimum values in the study. Using 2mm elements the mean error increased to 17.46% for the maximum force values. These values occurred at a displacement load of 2mm, at the end of the non-linear region in the responses. However the yield force values for the 2mm mesh remained consistent to within an average error of 1.45% when compared against equivalent results using 0.5mm elements. In practice the key values of interest for this study were the stiffness and yield strength. As such a 2mm mesh was selected due to a significant improvement in solve time. Figure 139 shows how the element size caused the solve time to vary for each test run carried out at the limits of the design space. Each simulation was carried out on hardware featuring an Intel i7-10700K processor at 3.8GHz with 32GB RAM and an Nvidia RTX 3060 GPU. Based on the mean solve time the 0.5mm mesh would be expected to take at least 50 days to carry out 4500 simulations.

Increasing the element size to 2mm was predicted to allow creation of an equivalent dataset within 48 hours of continuous runtime. In practice the simulations were run for approximately 8 hours a day over the course of 6 days to create a dataset of 4248 observations. 4500 FEA runs were carried out in total with a convergence rate of 94.4%.

Table 20 Geometry parameters for mesh convergence study test cases

	t(mm)	w(mm)	h(mm)
Lowest Volume	1	10	10
Greatest Volume	4	4	25
Strongest	4	4	10
Weakest	1	10	25

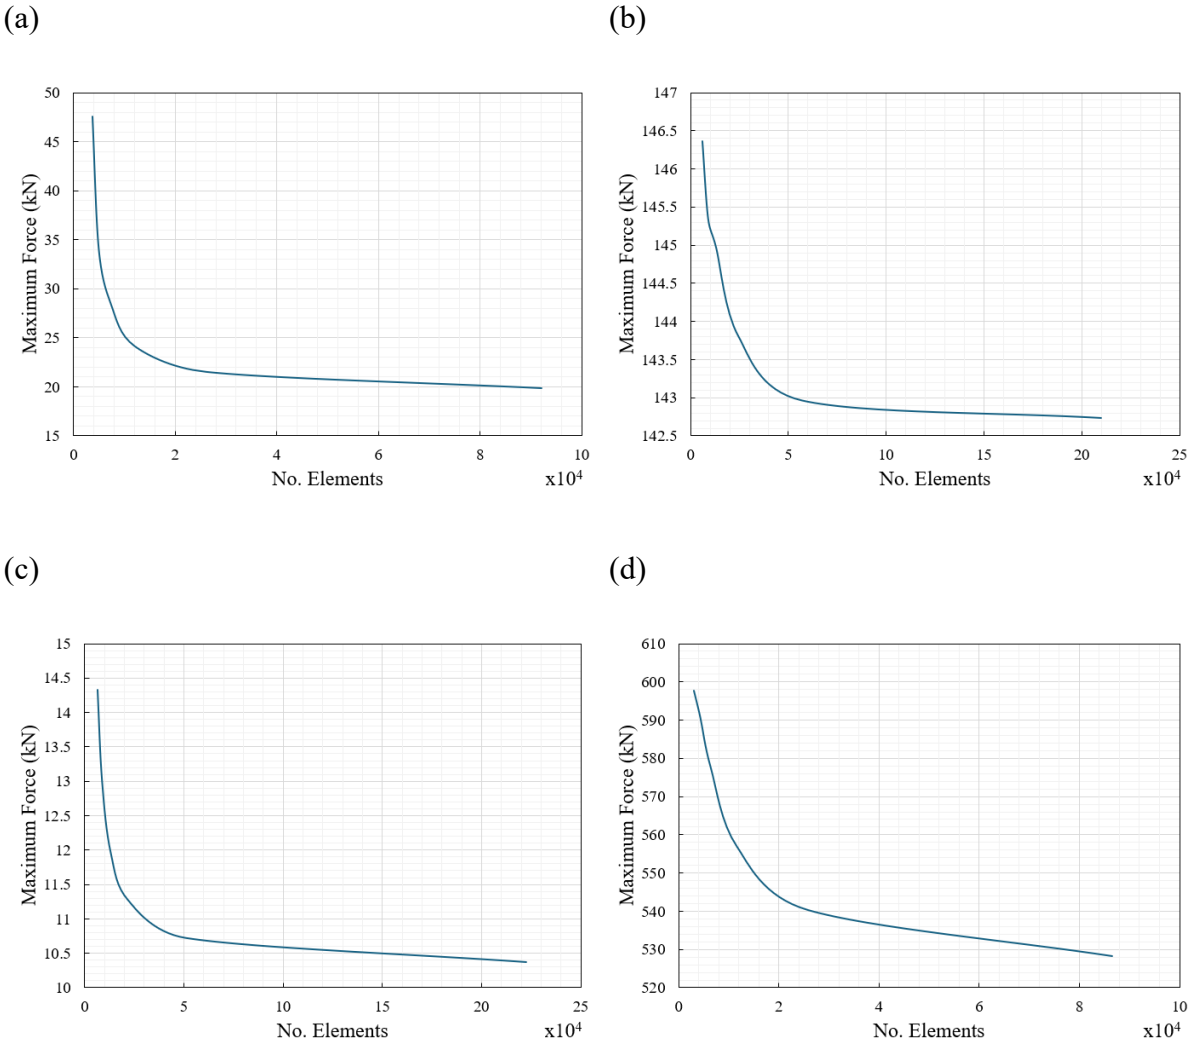


Figure 138 Mesh convergence study for (a) lowest volume, (b) greatest volume, (c) weakest, (d) strongest unit cell designs possible within the chosen design space limits

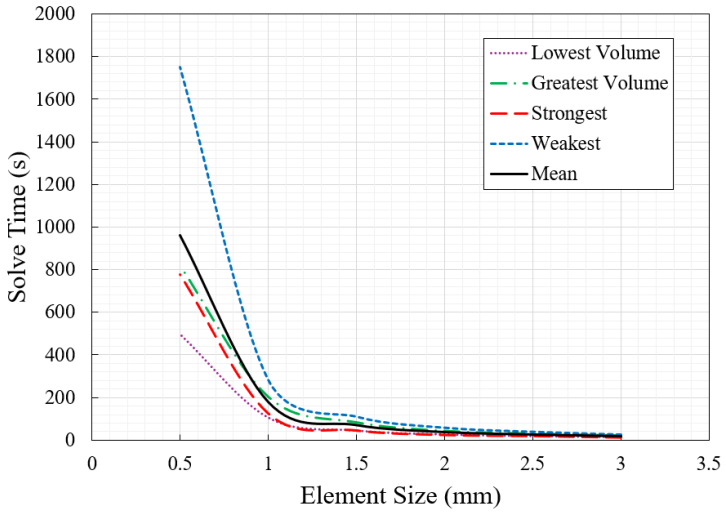


Figure 139 FEA mesh element size vs simulation solve time

The geometry parameters for each design point in the dataset were chosen according to a design of experiments test plan of 4500 observations using Latin hypercube sampling. The parameter ranges were selected to target overall material thicknesses desirable for use in passenger vehicles. The resonator channel dimensions were also set so as to feature adequate cross-sectional area to allow for strong sound absorption at low frequencies. The wall thicknesses were varied to allow for a large range of possible structural properties that can be fine-tuned. The chosen limits of the design space for each parameter are shown in Table 21. The population density of the design space is illustrated in Figure 140, which contains a pair-wise plot of the design space test plan. The possible range of structural responses expected within the design space is demonstrated in Figure 141, populated from the FEA results of a full factorial design of experiments test plan with 3 levels based on the chosen parameter limits. The study predicted the recorded yield force may vary between 16-73kN. This data was used to inform selection of an appropriate capacity load cell, ensuring the entire range of possible results could be tested during experimental post validation of any optimised solutions. At each observation in the dataset a vector of 51 datapoints was assembled from the FEA force-displacement structural characteristic and 51 datapoints from the analytical sound absorption-frequency response. The forces, F were all normalised to take values between 0 and 1, rescaled from the minimum and maximum values that appeared anywhere in the dataset according to the following formula.

For rescaling vector, F with original minimum and maximum limits F_{min} and F_{max} to an equivalent vector normalised to take values between the limits of a and b, where a is the lower limit and b the upper limit:

$$F_{rescaled} = a + \left[\frac{F - F_{min}}{F_{max} - F_{min}}\right](b - a) \tag{5.30}$$

The normalised force vector was combined with the sound absorption coefficient vector to produce a single vector of 102 entries, all taking values between 0 and 1. Entries 1-51 were populated by the raw absorption vector (already between 0 and 1) and 52-102 the normalised force vector. This combined sequence vector was then supplied as input for the convolutional neural network model. An example of the combined vector is shown in Figure 142.

Table 21 Parameter limits for chosen design space

Parameter	t(mm)	w(mm)	h(mm)
Lowest Value Limit	1	4	10
Greatest Value Limit	4	10	25

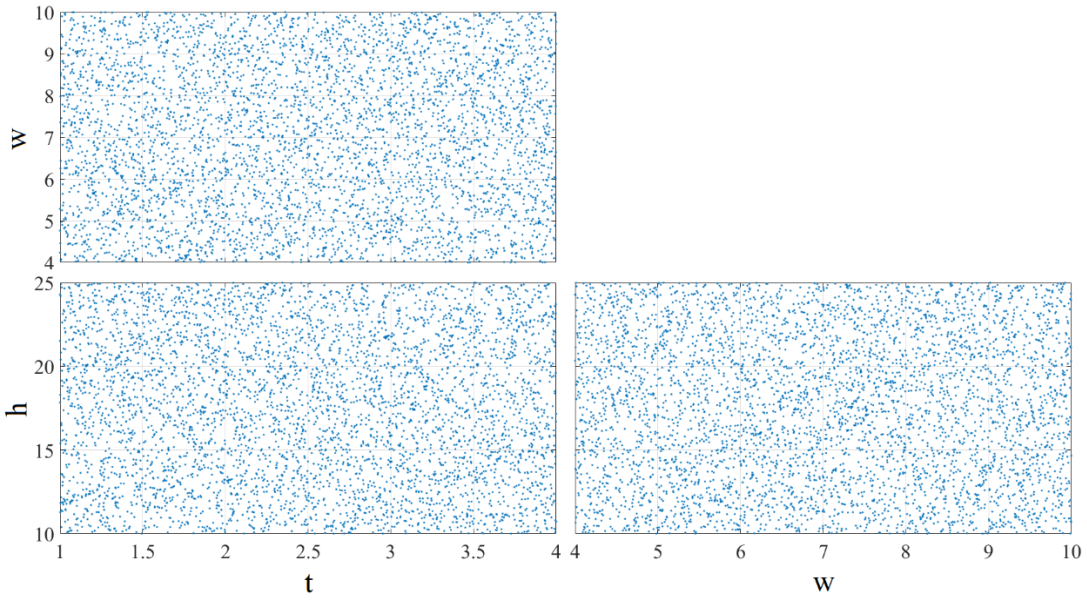


Figure 140 Pairwise plot of Latin hypercube design of experiments test plan of 4500 design points

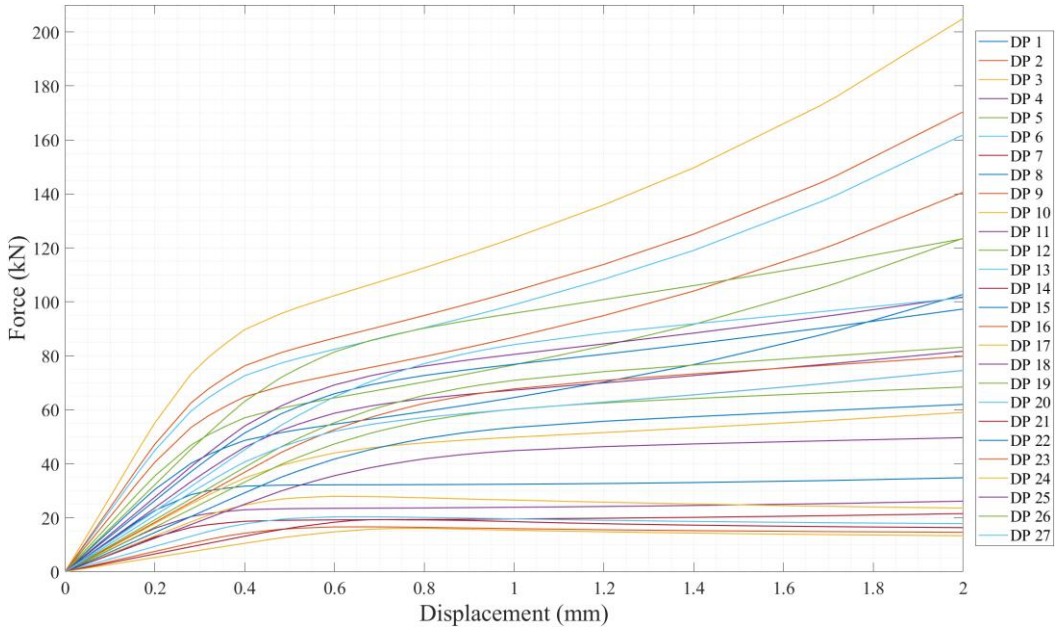
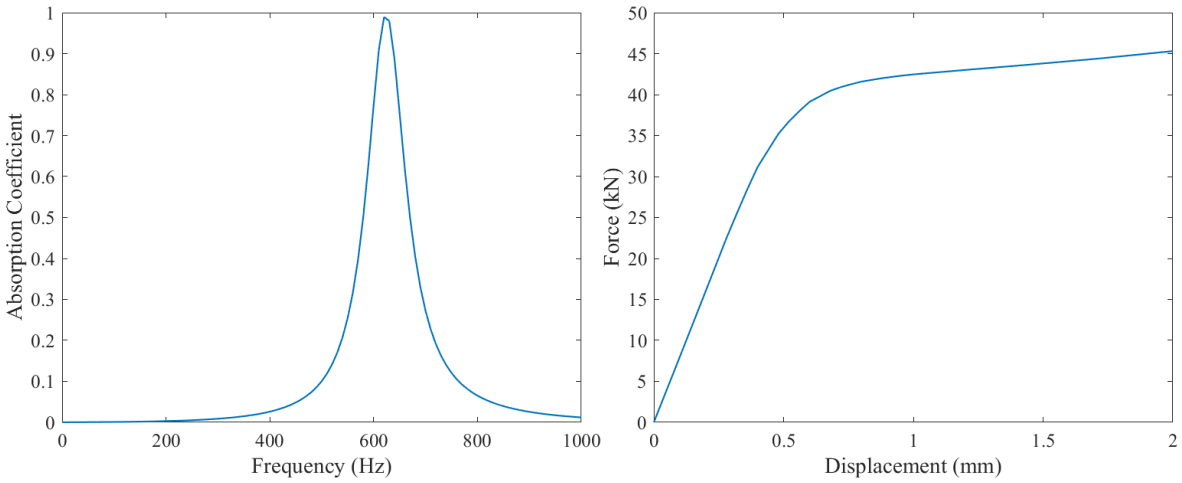


Figure 141 Range of possible FEA responses for chosen design space limits. 27 design points were selected from a full factorial design of experiments test plan with 3 levels based on the limits of the 3 geometry parameters.

(a) (b)



(c)

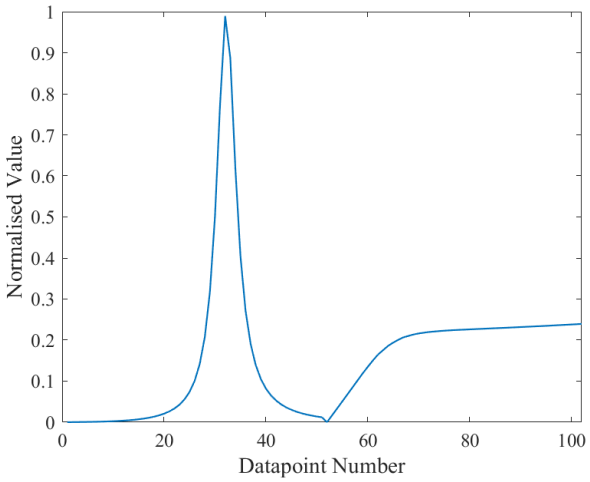


Figure 142 Example dataset design point observation signal used as input for the convolutional neural network; (a) analytical sound absorption-frequency response; (b) FEA force-displacement structural response for equivalent geometry; (c) final sequence vector generated from combination of raw sound absorption and normalised structural characteristics

5.1.3 Inverse Convolutional Neural Network Architecture

The CNN architecture is illustrated in Figure 143. The combined input signals were formatted as a 4248x1 column matrix of cell arrays containing the entire dataset. Each cell contained a 1x102 row vector corresponding to a single design point. These were provided to the network through the initial sequence input layer. The signals were then passed through a sequence of 5 1-dimensional convolutional layers, each with ReLU activation and layer-wise normalisation. The convolutional layers all had filters of size 3, stride of 1 and no padding. The filters were used to scan each input sequence, performing convolution operations to extract a series of complex feature maps which are passed through subsequent layers. The first layer featured 8 filters and the number of filters was doubled for each of the following layers to continually increase the complexity until 128 filters in the 5^{th} convolution layer. The resulting feature maps were then flattened into a single vector using global average pooling. This vector was sent through a series of fully connected hidden layers with batch-wise normalisation and ReLU activation. The first contained 512 neurons, then the number of neurons was halved with subsequent layers until 8 neurons in the 7^{th} fully connected layer. The 8^{th} and final fully connected layer featured 3 neurons corresponding to the 3 output geometry parameters, which were connected to a regression output. The weights and biases of each layer were updated as the signals were sent through the network using the Adam adaptive learning rate with momentum gradient descent optimisation algorithm. The objective of the algorithm was to minimise the root mean square error (RMSE) between the network predictions and known values of the output parameters from the original dataset. The number of learnable parameters at each layer within the network is summarised in Table 22. The network was developed within MATLAB using the deep learning toolbox and deep network designer plugin.

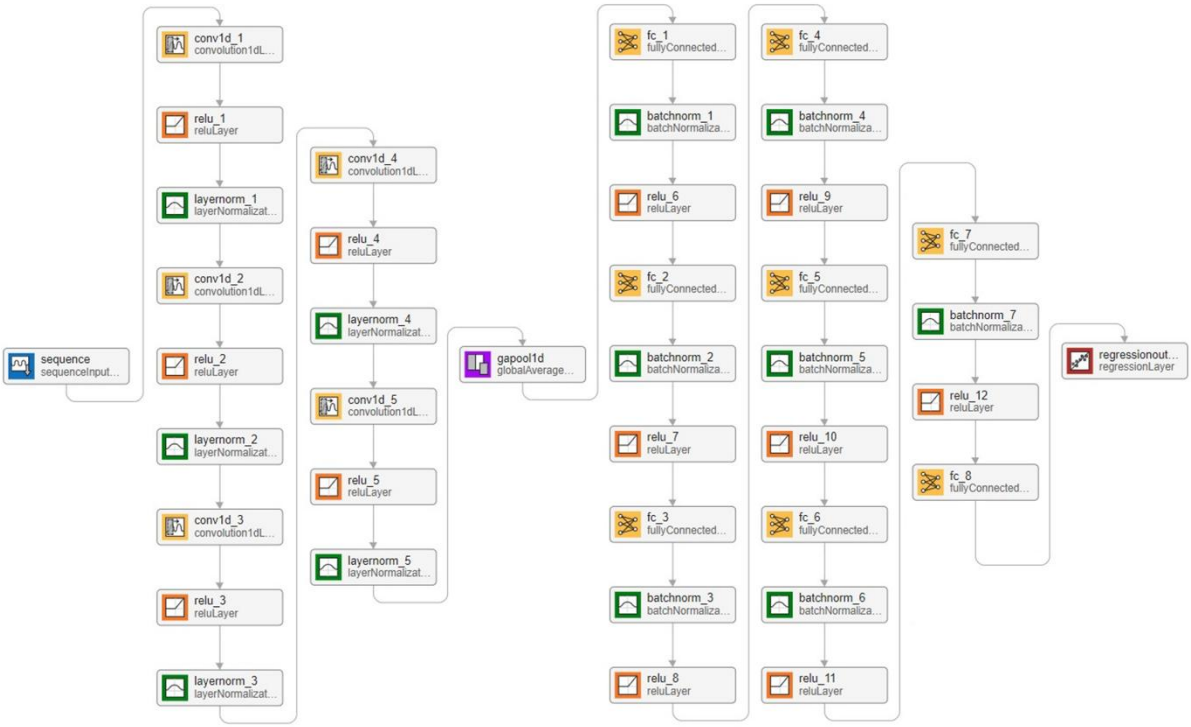


Figure 143 Inverse convolutional neural network architecture schematic

Table 22 Inverse convolutional neural network layers

No.	Name	Type	Activations	Learnable Parameters
1	sequence	Sequence Input	1 x 1 x 1	-
2	conv1d 1	8 Filter 1-D Convolution	8 x 1 x 1	Weights 3 x 1 x 8, Bias 1 x 8
3	relu 1	ReLU Activation	8 x 1 x 1	-
4	layernorm 1	Layer Normalisation	8 x 1 x 1	Offset 8 x 1, Scale 8 x 1
5	conv1d 2	16 Filter 1-D Convolution	16 x 1 x 1	Weights 3 x 8 x 16, Bias 1 x 16
6	relu 2	ReLU Activation	16 x 1 x 1	-
7	layernorm 2	Layer Normalisation	16 x 1 x 1	Offset 16 x 1, Scale 16 x 1
8	conv1d 3	32 Filter 1-D Convolution	32 x 1 x 1	Weights 3 x 16 x 32, Bias 1 x 32
9	relu 3	ReLU Activation	32 x 1 x 1	-
10	layernorm 3	Layer Normalisation	32 x 1 x 1	Offset 32 x 1, Scale 32 x 1
11	conv1d 4	64 Filter 1-D Convolution	64 x 1 x 1	Weights 3 x 32 x 64, Bias 1 x 64
12	relu 4	ReLU Activation	64 x 1 x 1	-
13	layernorm 4	Layer Normalisation	64 x 1 x 1	Offset 64 x 1, Scale 64 x 1
14	conv1d 5	128 Filter 1-D Convolution	128 x 1 x 1	Weights 3 x 64 x 128, Bias 1 x 128
15	relu 5	ReLU Activation	128 x 1 x 1	-
16	layernorm 5	Layer Normalisation	128 x 1 x 1	Offset 128 x 1, Scale 128 x 1
17	gapool1d	1-D Global Average Pooling	128 x 1 x 1	-
18	fc 1	512 Neuron Fully Connected	512 x 1	Weights 512 x 128, Bias 512 x 1
19	batchnorm 1	Batch Normalisation	512 x 1	Offset 512 x 1, Scale 512 x 1
20	relu 6	ReLU Activation	512 x 1	-
21	fc 2	256 Neuron Fully Connected	256 x 1	Weights 256 x 512, Bias 256 x 1
22	batchnorm 2	Batch Normalisation	256 x 1	Offset 256 x 1, Scale 256 x 1
23	relu 7	ReLU Activation	256 x 1	-
24	fc 3	128 Neuron Fully Connected	128 x 1	Weights 128 x 256, Bias 128 x 1
25	batchnorm 3	Batch Normalisation	128 x 1	Offset 128 x 1, Scale 128 x 1
26	relu 8	ReLU Activation	128 x 1	-
27	fc 4	64 Neuron Fully Connected	64 x 1	Weights 64 x 128, Bias 64 x 1
28	batchnorm 4	Batch Normalisation	64 x 1	Offset 64 x 1, Scale 64 x 1
29	relu 9	ReLU Activation	64 x 1	-
30	fc 5	32 Neuron Fully Connected	32 x 1	Weights 32 x 64, Bias 32 x 1
31	batchnorm 5	Batch Normalisation	32 x 1	Offset 32 x 1, Scale 32 x 1
32	relu 10	ReLU Activation	32 x 1	-
33	fc 6	16 Neuron Fully Connected	16 x 1	Weights 16 x 32, Bias 16 x 1
34	batchnorm 6	Batch Normalisation	16 x 1	Offset 16 x 1, Scale 16 x 1
35	relu 11	ReLU Activation	16 x 1	-
36	fc 7	8 Neuron Fully Connected	8 x 1	Weights 8 x 16, Bias 8 x 1
37	batchnorm 7	Batch Normalisation	8 x 1	Offset 8 x 1, Scale 8 x 1
38	relu 12	ReLU Activation	8 x 1	-
39	fc 8	3 Neuron Fully Connected	3 x 1	Weights 3 x 8, Bias 3 x 1
40	regressionoutput	Regression Output	3 x 1	-

5.1.4 Network Training Performance

The network outputs were contained in a 4248x3 matrix of scalar values normalised between 0 and 1. The network provided predictions for each geometry parameter between 0 and 1 which were then rescaled according to the minimum and maximum values of the parameters in the original dataset. The dataset was partitioned with 70% used to train the network, 15% for validation checks during training, and 15% kept unseen by the network for testing of the predictions against known values after the training was complete. The network was trained with GPU acceleration on an Nvidia RTX 3060 through use of the MATLAB parallel computing toolbox on the same hardware used to generate the FEA dataset, as detailed in section 5.1.2. The dataset was passed through the network 1000 times, over the course of 1000 epochs in mini batches of 100 observations at a time. Each epoch is defined as the period during which the entire dataset is sent through the network once. There were 29 training iterations per epoch. At the end of every epoch a validation check was carried out and the mini batches were shuffled. The initial learning rate was set at 0.001. The training progress throughout each iteration is plotted in Figure 144. A sharp drop in both RMSE and loss is seen over the first 10 epochs as the network initialises. A gap is observed between the training and validation error values. This may be due to limitations on the dataset size. Increasing the number of observations may lead to a fuller representation of the available design space and further improve the network performance where necessary. The validation error begins to level off after approximately 100 epochs, while the training error continues to decrease at a gradual rate as the training progresses. There were occasional small spikes in both error readings. However there were no sustained periods of increases in validation error throughout the training, suggesting there is no evidence of overfitting. An early stopping strategy was not implemented as allowing the training to progress past 100 epochs lead to improved regression performance when the trained network was tested on unseen data.

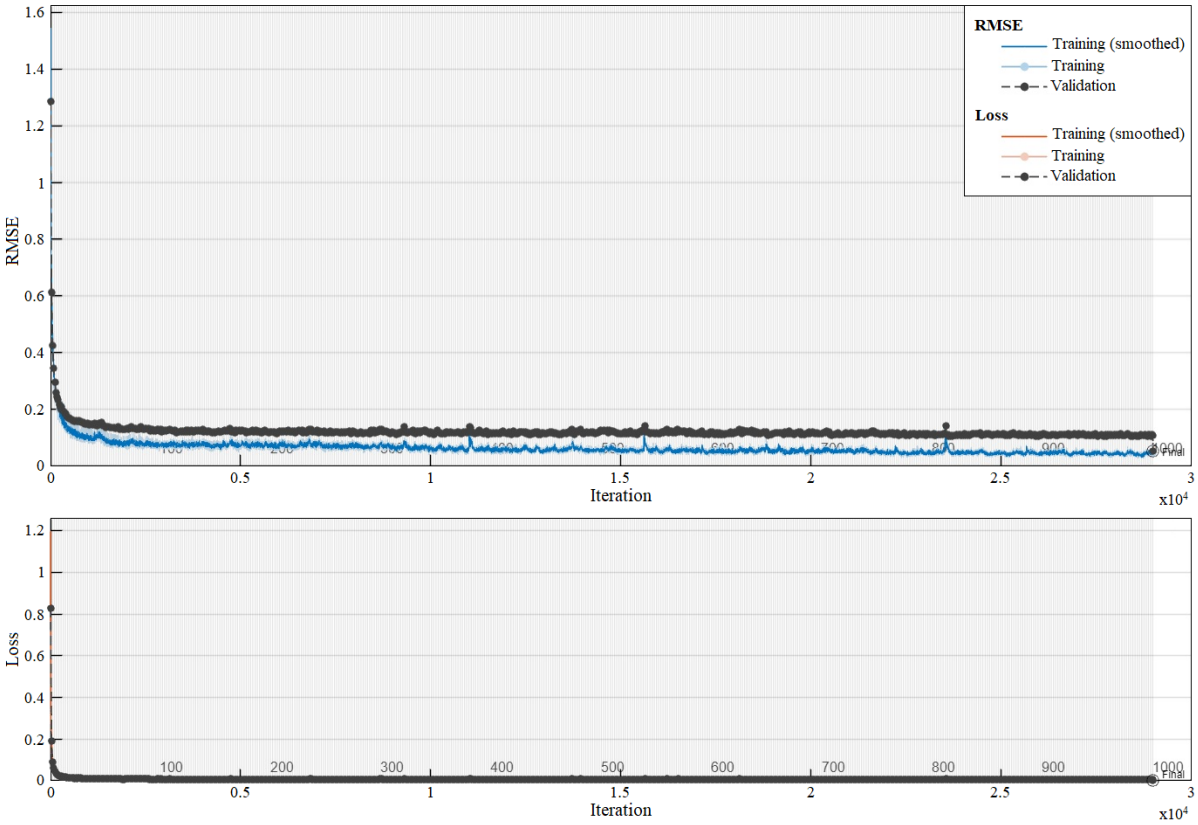


Figure 144 Training history of convolutional neural network

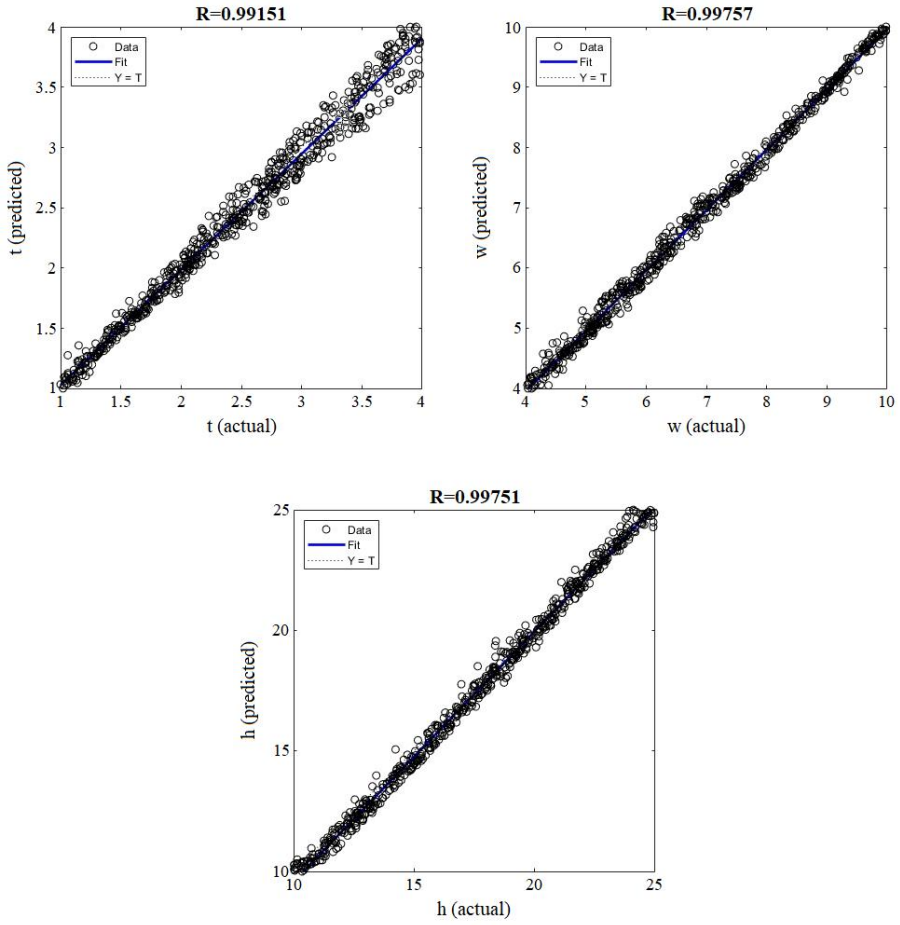


Figure 145 Regression between inverse network predictions and known unit cell geometry parameters

The regression between the network predictions and actual known values for the test dataset partition of 637 datapoints is displayed in Figure 145 for each of the 3 variable geometry parameters. Strong positive correlation is observed in each case with few significant outliers. Predictions of the wall thickness, t featured the greatest discrepancies from known values, however the overall regression coefficient remained above acceptable limits. The performance of the network trained on the combined input signal is compared in Table 23 against an equivalent network trained on only the force-displacement signal for the same dataset, sampled at 102 equally spaced intervals between 0-2mm displacement. In each case the dataset partitions, network architecture, hyperparameters and training options were kept identical. Both networks performed well on the unseen data, indicating they are well generalised and able to adapt to new inputs. The test coefficient of regression exceeded 0.91 for the structural-only model. The greater the value the better the performance, where a value of 1 indicates perfect predictions. The regression notably improved for the combined model, offering values above 0.99 across all geometry parameters. Reductions are also observed in the final training and validation root mean square error and loss values. This suggests the predictions are more accurate and the network has greater confidence in its outputs for the combined model, while the amount of time spent training was relatively similar to within 15 seconds. The AI model is likely able to more effectively identify the metamaterial unit cell required for a given input for the combined model as more information has been provided about the desired solution to help narrow the search of the design space. The predictive performance may have improved as a result of providing the network with a more complete picture of the problem being solved, despite the complexity increasing through consideration of multiple competing objectives.

Table 23 Summary of training performance and test regression for inverse CNNs trained for multi-objective and single-objective optimisation

Input Signal	**Training Performance**			**Regression**	
Sound Absorption-Frequency & Force-Displacement	**RMSE**	**Training**	0.0484	**Parameter**	**R Value**
		Validation	0.1092		
	Loss	**Training**	0.0012		
		Validation	0.0060	**t**	0.99151
	Training Time	23min 17sec		**w**	0.99757
	No. Iterations	29000		**h**	0.99751
Force-Displacement Only	**RMSE**	**Training**	0.0862	**Parameter**	**R Value**
		Validation	0.1490		
	Loss	**Training**	0.0037		
		Validation	0.0111	**t**	0.98294
	Training Time	23min 32sec		**w**	0.91019
	No. Iterations	29000		**h**	0.99454

5.1.5 Inverse Optimisation and Experimental Validation

The trained network was used to simultaneously optimise the sound absorption and structural properties of the metamaterial unit cell. The first harmonic absorption peak was tuned to achieve the lowest possible frequency at the maximum amplitude by representing the requested absorption coefficient-frequency characteristic as a spike input centred on the desired frequency. The targeted structural response was set using a bi-linear input to represent the force-displacement curve. The solution geometry parameters returned by the CNN were run back through the analytical sound absorption and the non-linear FEA models to compare the responses against the requested inputs. Normalised relative values taken from the representative unit cell were used to allow for evaluation of the stiffness and strength of the metamaterial irrespective of the size of the test specimen and number of unit cells within the lattice. The force-displacement reponses were converted to stress-strain using the cross-sectional area and original height of the specimen. Young's modulus was calculated from the slope of the linear region and the yield strength was taken as the 0.2% proof stress. The density of the representative unit cell was calculated as the mass divided by the total volume including void spaces. The relative density of the equivalent metamaterial lattice consisting of any number of repeating cells was then calculated as the density of the single unit cell divided by the density of the solid material within that cell. The relative specific stress of each specimen was found by dividing the stress vector by the relative density. A script was written in MATLAB to carry out these calculations automatically. The force-displacement requests for the individual unit cell could then be updated to fine-tune the relative specific stiffness and strength values of the full lattice material.

Example specimens optimised by the inverse CNN were manufactured by FDM from ColorFabb XT-CF20 filament for experimental validation of the material properties. Sound absorption-frequency responses were measured using the test setup detailed in section 5.1.1, while structural properties were evaluated under uni-axial compressive testing on an Instron 6025 with 100kN load cell. The structural specimens were printed with the same hexagonal geometry that was used in the parametric FEA model, as shown in Figure 146. Sound absorption specimens featured the equivalent geometry with the layout featured in Figure 123(c), with identical internal walls but with circular outer walls and the hexagonal features removed to fit the impedance tube of 50.8mm diameter. The layout of the Instron test rig is displayed in Figure 147. The as-built dimensions of the samples were measured with a digital micrometer. They were then set on a raised steel platform and the compression platen lowered to a paper's width (0.1mm) above the specimen before loading at a rate of 1.3mm/min until rupture. The force values were recorded by the load cell every 0.02 seconds and the strain

measured by an LVDT held in an additively manufactured jig bolted to the base of the machine. The slope of the linear region was extracted and used to perform a toe compensation to remove the initial period prior to the platen meeting the surface of the specimen in line with the process outlined in ASTM standard D695 Annex A1.



Figure 146 Manufactured compression test specimens

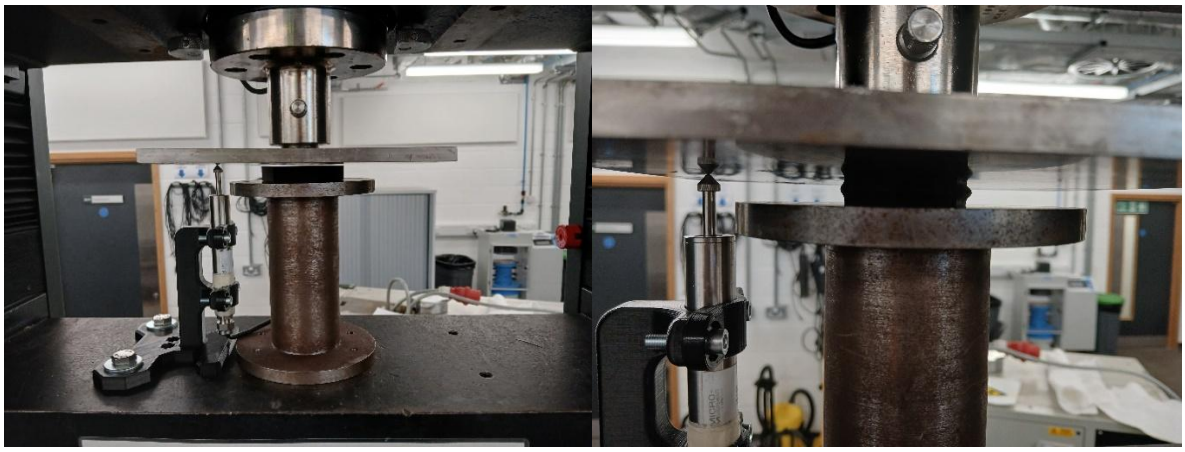


Figure 147 Test rig setup for experimental mechanical characterisation of optimised test samples

The ability of the CNN to perform inverse optimisations of both sound absorption and structural properties at the same time is demonstrated below for a series of examples tuned for a range of performance targets. In each case the required relative specific stiffness and strength was set before tuning of the first harmonic absorption peak by shifting the spike input to gradually lower frequencies until a solution was found at the lowest frequency to maintain a peak amplitude above 90% absorption for the given structural target. Table 24 compares the requested first harmonic absorption peak amplitudes and frequencies against the analytical predictions and measured experimental values. The structural performance targets, FEA predictions and experimental mechanical properties for the same examples are provided in Table 25. The corresponding geometry parameters for each solution returned by the CNN are displayed in Table 26. The requested spike and bi-linear inputs for each example are compared against the predicted analytical or FEA responses and the experimental responses in Figure 148. The geometries of the solution unit cells are displayed in Appendix J. While tuning each solution the original FEA model used to generate the dataset was used to supply an initial coarse prediction. An additional FEA model featuring a finer mesh of 0.5mm elements was then used to compensate for discrepancies with the experimental results. This model featured a non-linear buckling analysis to take into account the potential for premature yielding due to Euler buckling of the cell walls. The structure of the analysis is shown in Figure 149. To begin the simulation subjected the geometry to a unit load to extract the first buckling mode shape and the load factor required to initiate buckling. A perturbed version of the starting geometry was then generated in the shape of the first buckling mode, examples of which can be seen in Figure 150. The extent to which the geometry was perturbed was set by a scale factor that could be updated to fine-tune the model according to the experimental measurements. These resulting models featured higher fidelity, were in closer agreement with the real-world situations and were used to fine-tune the optimised metamaterials to create a final result. For each solution a first iteration was tuned using the coarse model and then validated experimentally. The simulation was then repeated using the buckling model to find the perturbation factor to give greatest accuracy with the experimental response. The requested inputs were then redefined iteratively, running the geometry returned by the solver through the tuned buckling model until the stiffness and strength targets were met. The results were then revalidated through a final round of experimental testing. This approach combining both coarse and higher fidelity simulations both speeds up the process of dataset generation and also reduces the likelihood of having to repeatedly reiterate, remanufacture and retest the solutions to compensate for discrepancies between the FEA and experimental measurements. Both the fine buckling model FEA solutions and the initial coarse solutions are included for each example for comparison against the final experimental responses in Figure 148.

Table 24 First harmonic sound absorption performance of example optimised metamaterials

Example	Requested Targets		Sound Absorption (Analytical)		Sound Absorption (Experimental)	
	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)
A	1	340	0.9666	330	0.9877	318.45
B	1	400	0.9419	380	0.9849	373.06
C	0.97	320	0.9740	340	0.9776	326.19
D	0.97	320	0.9695	330	0.9680	303.06
E	1	320	0.9694	330	0.9653	311.71
F	1	360	0.9703	350	0.9470	331.02
G	1	360	0.9131	360	0.9511	355.06
H	1	360	0.9565	350	0.9771	340.47

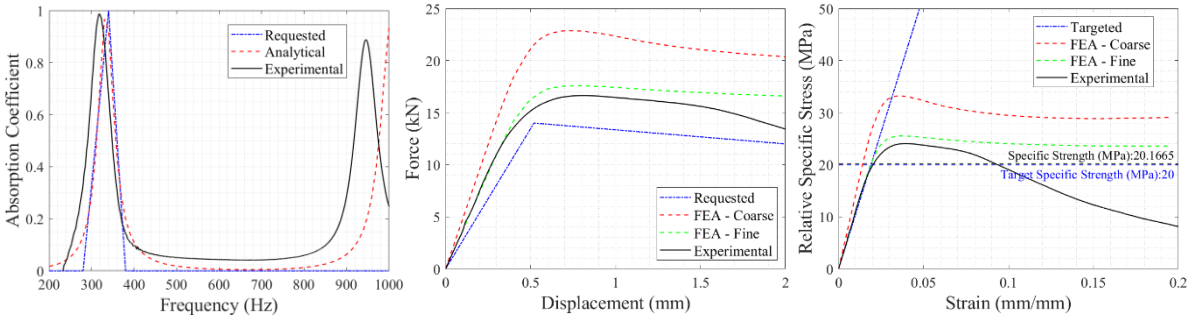
Table 25 Mechanical properties of example optimised metamaterials

Example	Requested Targets		Structural Properties (FEA – Buckling Model)		Structural Properties (Experimental)	
	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)
A	1050	20	1095.94	22.73	1089.66	20.17
B	1200	23	1232.72	25.53	1228.63	23.93
C	1250	23	1275.26	25.37	1276.60	23.80
D	1300	24	1316.61	28.57	1311.79	24.07
E	1400	25	1419.39	30.30	1419.80	25.47
F	1500	26	1513.43	30.15	1555.90	26.20
G	1550	33	1551.69	35.67	1554.36	33.64
H	1600	34	1593.71	36.52	1599.91	34.01

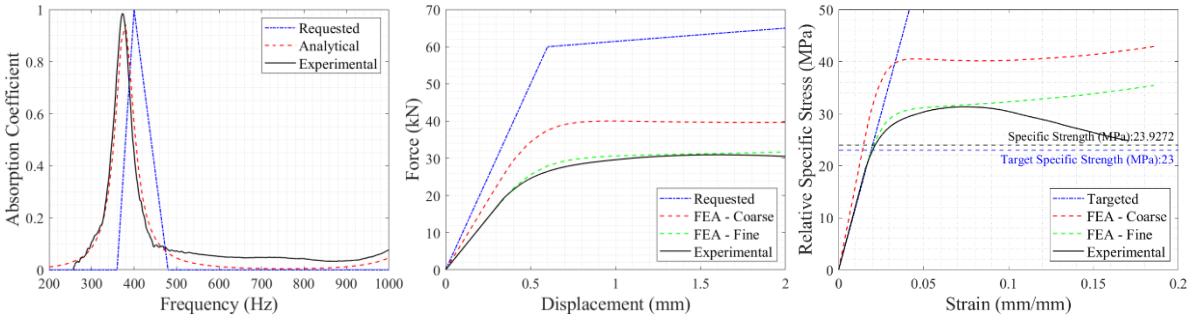
Table 26 Geometry parameters of example optimised metamaterials

Example	Solution Unit Cell Geometry Parameter					
	t (mm)	w (mm)	h (mm)	Mass (g)	Solid Volume (mm^3)	Relative Density
A	0.98	6.19	18.56	13.18	14169.45	0.3083
B	1.52	4.44	19.55	19.79	21277.38	0.4418
C	1.07	5.81	18.43	15.33	16483.19	0.3610
D	1.49	5.61	22.47	18.84	20248.81	0.3703
E	1.35	5.73	20.49	16.85	18111.28	0.3604
F	1.11	5.51	19.23	15.98	17180.11	0.3621
G	2.15	4.07	19.83	26.41	28383.96	0.5817
H	2.10	4.47	22.03	26.08	28037.10	0.5221

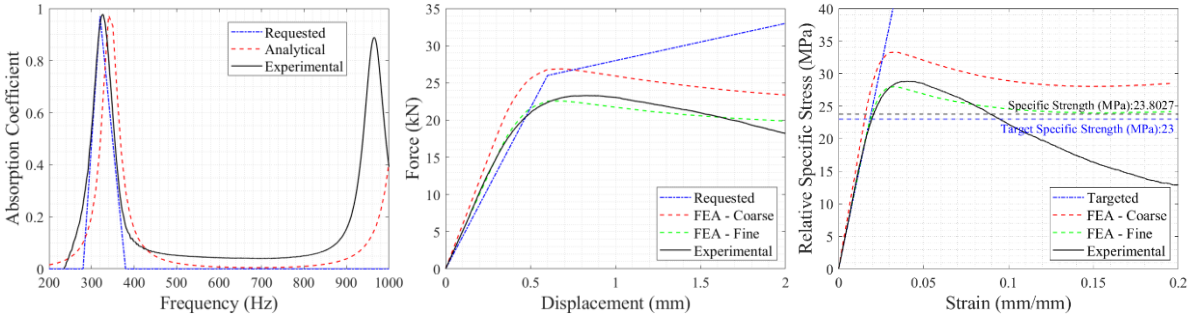
Example A



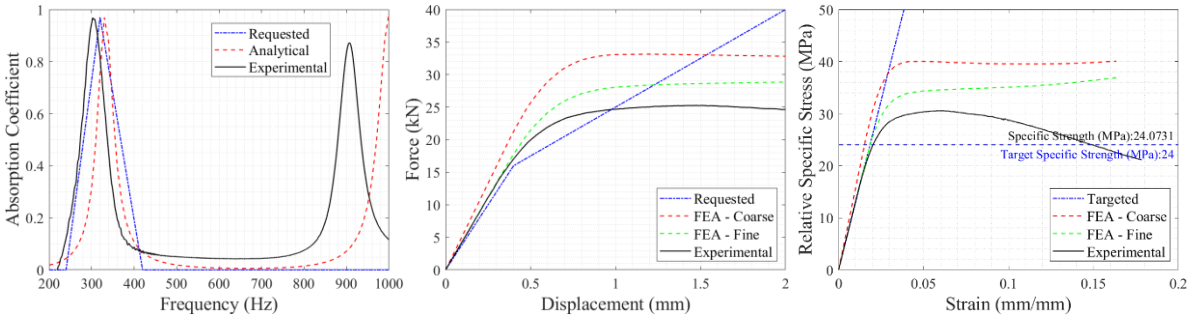
Example B



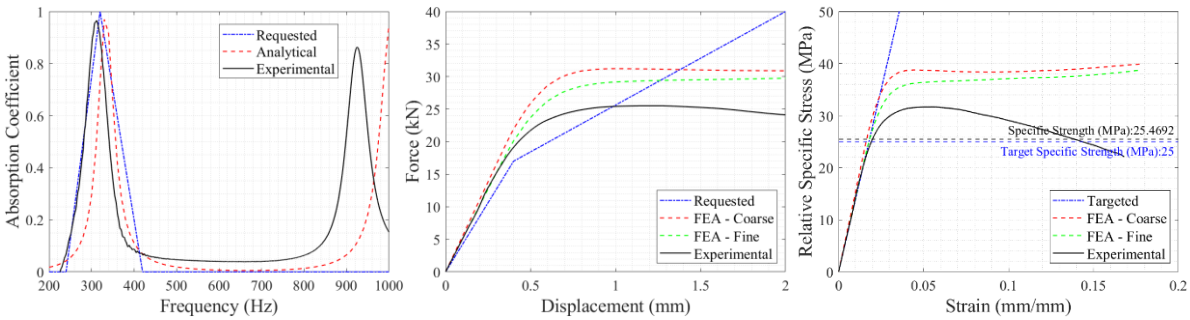
Example C



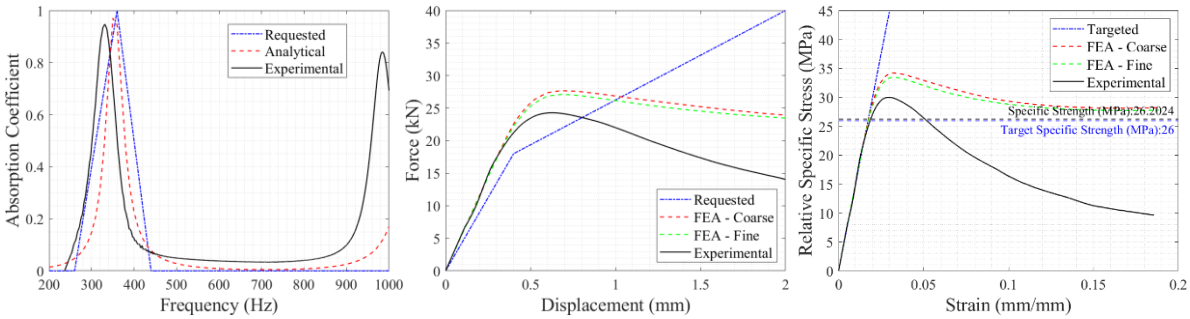
Example D



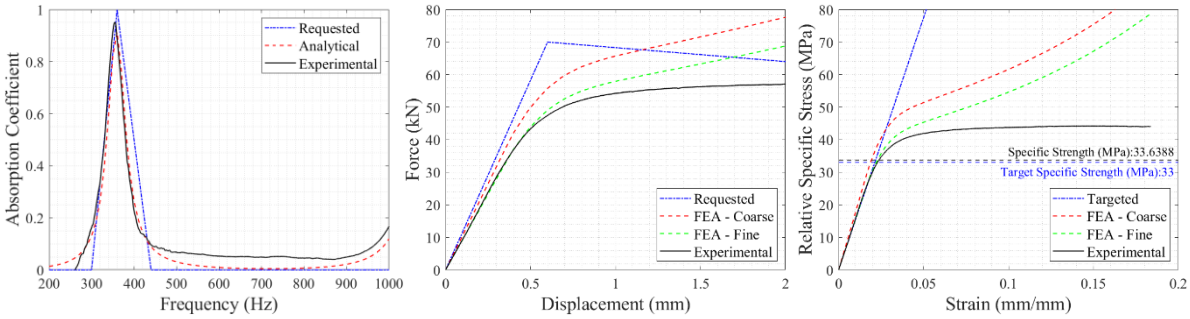
Example E



Example F



Example G



Example H

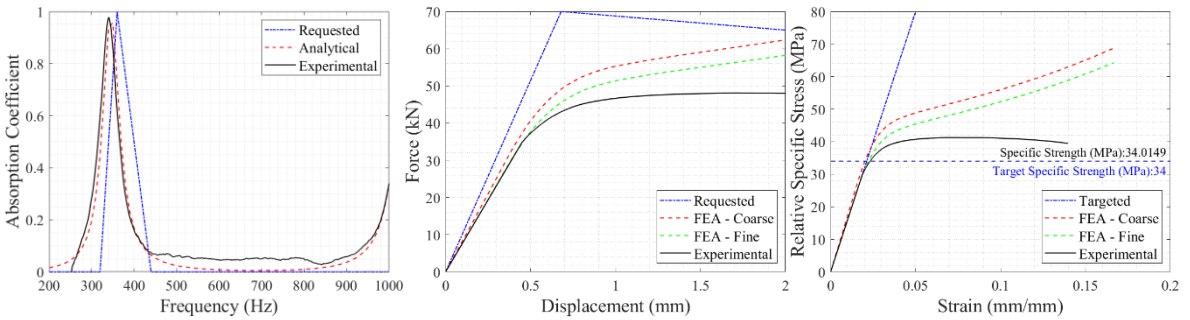


Figure 148 Sound absorption responses and mechanical characteristics of example optimised metamaterials

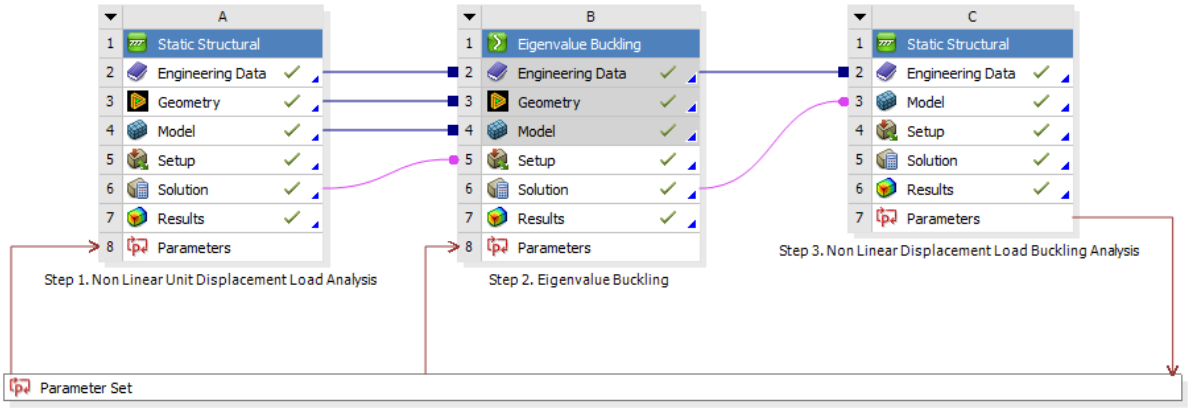


Figure 149 Non-linear buckling analysis FEA model schematic

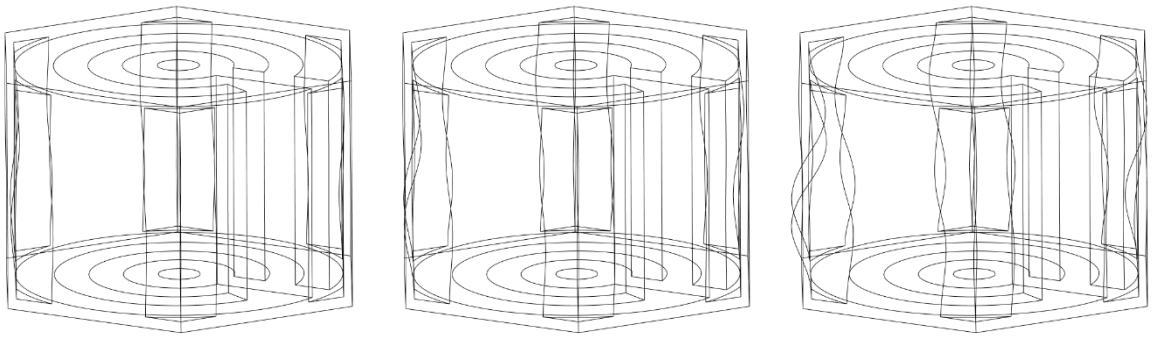


Figure 150 Example geometry perturbed according to the first buckling mode shape using a scale factor of 1 (left), 2 (middle) and 4 (right)

The inverse model was able to tune the relative structural properties of the metamaterial to satisfy performance targets between 1050-1600MPa specific stiffness and 20-34MPa specific strength. The specific stiffness values predicted by the non-linear buckling FEA simulations gave mean error of 0.59% when compared with experimental measurements. Specific strength values showed overshoots in all cases with mean value of 2.94MPa. However the CNN offered the ability to rapidly fine-tune the model to minimise discrepancies with the target values by updating the input signal. After compensation the solutions provided experimental relative specific strength that always exceeded the specification while remaining on average within 1.66% of the target value. The experimental relative specific stiffness of the optimised solutions gave mean error of 1.83% in comparison to the target. Agreement was observed between the requested bi-linear force-displacement inputs and the experimental response curves. Example B showed the greatest discrepancy, possibly as a solution closer to the request may not be possible. In this case the AI model still shows an attempt to maximise the area under the curve to fit the bi-linear request, suggesting an exact input may not always be necessary and the system may be used to probe the limits of what is possible. In addition to performing the structural optimisations, the trained CNN was able to simultaneously maintain

a first harmonic sound absorption peak above 90% at frequencies that were below 400Hz in every example. The experimental peaks occurred on mean average at locations within 16.67Hz of the targeted frequencies. By theory the optimised metamaterials may be tuned to provide similar absorption magnitude at any frequency above the initial result while maintaining the structural properties by introducing a back wall termination to shorten the length of the resonator channel. The natural frequency, f at which the peak occurs will be inversely proportional to the length, l according to $f = \frac{c}{4l}$, where c is the speed of sound [154]. Examples C-E demonstrate the ability of the model to lock the minimum absorption peak frequency to a given value while tuning the structural performance. The target minimum absorption frequency was locked at 320Hz. The optimisations were then performed assuming a minimum requirement of 23MPa relative specific strength is required to satisfy yield safety limits. On setting these criteria the stiffness could then be tuned by updating the input, fine-tuning the deflection to satisfy a given mass requirement for lightweighting purposes. Examples F-H demonstrate a similar scenario with the minimum frequency target locked at 360Hz.

To ensure sound absorption performance of the QWR labyrinth geometry layout remains valid at larger specimen sizes with multiple repeating unit cells, example A was manufactured for a larger impedance tube of 185x185mm square cross section. Due to the increased size of the cross section, the larger tube measurements had an upper frequency limit of 800Hz. The specimen geometry featured a honeycomb lattice with 7 complete active unit cells. A cutaway view of the printed test specimen is shown in Figure 151(b). The absorption response of the 185mm square specimen is compared against the equivalent measurement taken from a single unit cell on the 50.8mm diameter tube in Figure 151(a). The 185mm specimen provided a first harmonic peak at 315.33Hz, within 3.12Hz of the smaller version. A reduction was observed in the peak amplitude and broadband range, however the peak absorption was maintained above 90% and the response remained suitable for practical design purposes. Validation at larger specimen sizes is recommended in cases where the lattice must be conformed to fit the shape of a component. The reduction may likely have been caused due to the honeycomb tessellation featuring a greater amount of inactive area that contained incomplete unit cells as a result of fitting the geometry to the new specimen size.

(a) (b)

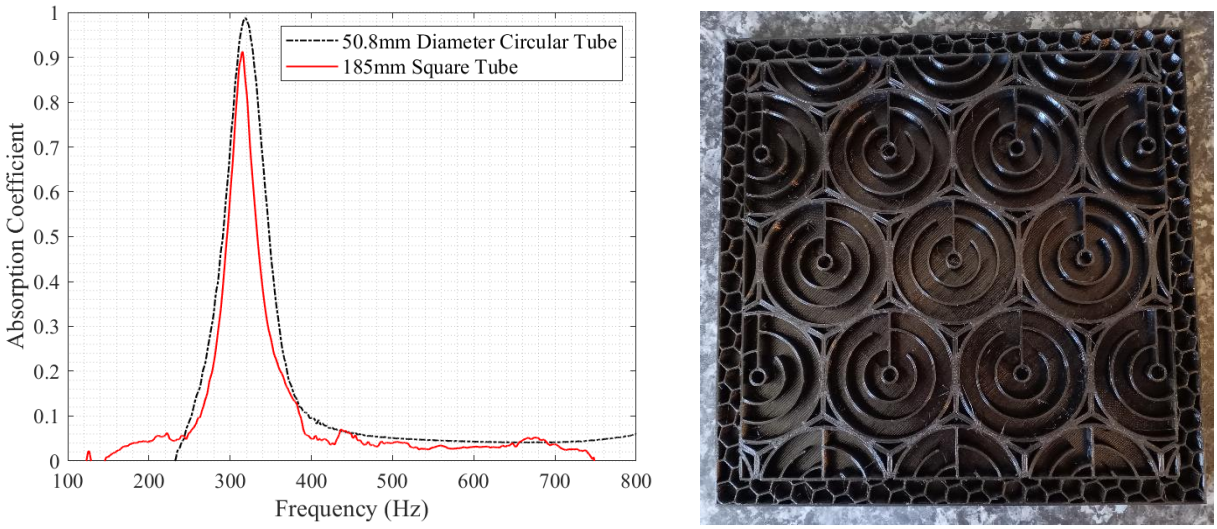


Figure 151 (a) Comparison of absorption responses for example A recorded on an impedance tube with 185mm square cross section vs 50.8mm diameter circular tube; (b) cutaway view of 185mm square test specimen

By allowing the first harmonic absorption peak to be tuned to the lowest frequency for a given structural design specification, the CNN provides a solution with the longest resonator channel that is able to satisfy the performance criteria. The channel can then be divided into a combination of multiple absorbers arranged in parallel for greatest overall broadband range, while maintaining the desired structural properties. The lower the initial first harmonic peak, the greater the maximum available channel length and the greater the achievable absorption range for parallel designs. In Figure 152 the analytical and experimental absorption responses are shown that resulted from dividing the solution geometry from example A into arrangements of 4 and 5 resonator channels acting in parallel. The layouts of each parallel design are displayed in Figure 153. The arrangement of 4 parallel resonators targeted no less than 50% absorption between 900-1200Hz and achieved a measured range of 857.4-1256.3Hz at an overall material thickness of 20.56mm. The unit cell was divided into 5 resonators to target a 50% absorption threshold of 1000-1500Hz. The manufactured specimen measured a range of 982.5-1542.1Hz at the same overall thickness of 20.56mm. The inverse CNN optimisation approach has been shown to offer control for fine-tuning of sound absorption and normalised relative specific structural properties of a metamaterial lattice. In addition the process offers potential for tailoring performance according to a given practical application. The input structural response may be substituted with that of a component under specific predetermined design loads and the network retrained without any additional complexity. Due to the dataset size requirements it is recommended to train on FEA models with as fine a mesh as can be practically generated within the timeframe available. The trained model can

then be used to compensate for any discrepancies in the experimental results. The system may be used to ensure the minimal amount of mass is used to satisfy stiffness and strength requirements for real-world applications that require multifunctional tunable sound absorption properties.

(a) (b)

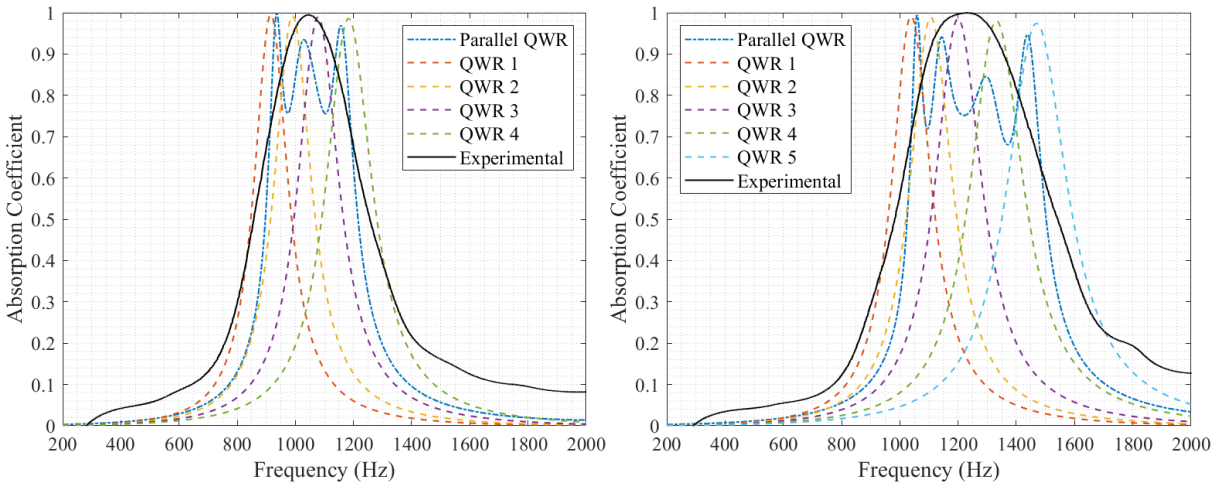


Figure 152 Analytical vs experimental absorption responses for parallel QWR specimens based on optimised example A, t=0.98mm, w=6.19mm, h=18.56mm, l_{inlet}=1mm ; (a) 4 resonators, l_1=92mm, l_2=85mm, l_3=78mm, l_4=71mm; (b) 5 resonators, l_1=81mm, l_2=76mm, l_3=70mm, l_4=63mm, l_5=57mm

(a) (b)

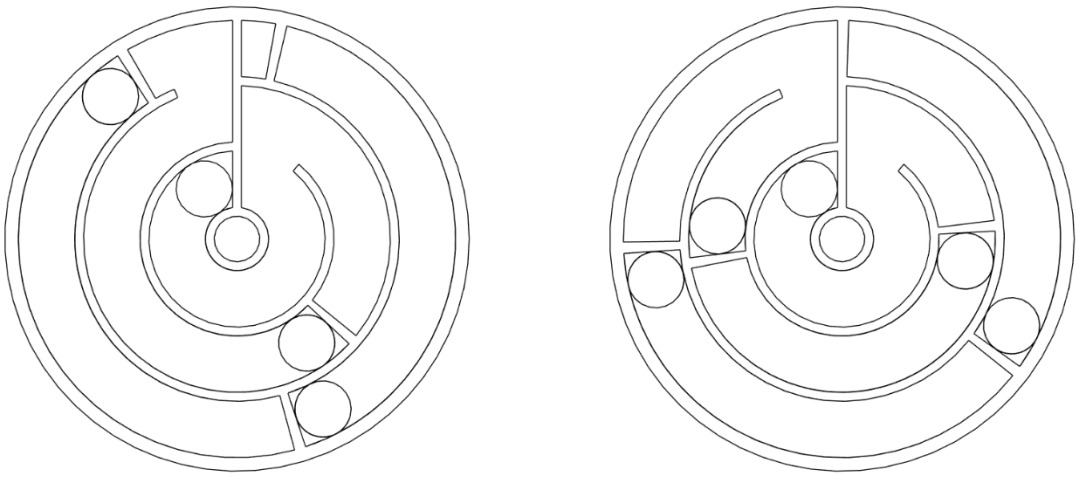


Figure 153 Geometry layout for parallel QWR specimens; (a) 4 parallel resonators; (b) 5 parallel resonators

5.2 Multi-objective Optimisation by Predictor Neural Network

While the inverse optimisation process is able to tune a metamaterial according to defined performance targets, an alternative approach is proposed for use when the specific desired properties are unknown and an overall best compromise between competing objectives is required. In this situation the optimal design points that exist at the limits of the performance trade-off between sound absorption and structural material property relations may be extracted through the use of a neural network that works in the opposite direction. Instead of outputting the geometry for a target set of material properties, the geometry parameters of the unit cell are provided as input and the network is used to perform rapid predictions of the performance objective values. The predictor network is combined with a multi-objective genetic algorithm to extract a Pareto front of optimal solutions. For each of these solutions one performance objective cannot be improved further without reducing another and a point of best compromise is reached. For multiple competing objectives multiple solutions will be found, all of which exist at an optimal point along the front. The user may then implement criteria to narrow the search, for example discarding any material that gives absorption below 90%, then select the solution that is most suitable for their application from those that remain. As a trained neural network is able to provide predictions faster than equivalent FEA simulations, it allows the objective values to be calculated for a large population of design points at each generation of the genetic algorithm within a reduced timeframe that would otherwise not be possible. In this section of the study networks are trained to predict the relative specific stiffness and specific strength of the QWR labyrinth unit cell. The models are used as the objective function of the multi-objective genetic algorithm, along with predictions of the sound absorption amplitude and frequency of the first harmonic peak. The algorithm is set to maximise the stiffness and strength of the material, while minimising the first harmonic frequency and maximising the first harmonic amplitude. The performance objectives are now represented as scalar values rather than a vector containing the full response signal. In addition, separate networks can be trained for each objective and in this case networks are only needed for structural objectives as sound absorption predictions may be carried out rapidly through the use of analytical theory. This significantly reduces the complexity required of the individual networks, allowing for use of a shallow ANN instead of a CNN.

5.2.1 Predictor Neural Network Model

The same dataset of 4248 observations that was used for the inverse model was reformatted for training the predictor network. A MATLAB script was written to calculate the relative specific stiffness and relative specific strength from the FEA structural responses for each design point. The corresponding geometry parameters were contained in a 4248x3 matrix and supplied as input to the network, with the outputs in individual 4248x1 column vectors. The network was trained using a Bayesian regularisation algorithm to minimise the mean squared error (MSE) between the predicted and known output value. The dataset was partitioned with 80% of the data used for training. The remaining 20% was kept unseen by the network during training and used for testing the trained network against known values. A validation partition was not necessary as the algorithm features internal validation checks based on Bayesian error, which depends on the size of the weights between neurons in addition to the accuracy of the predictions. The network was a shallow ANN featuring an input layer of 3 neurons corresponding to the 3 geometry parameters, followed by a single fully connected hidden layer with sigmoid activation. The hidden layer was connected to a linear output layer with a single neuron containing the value of either the relative specific stiffness or strength. A schematic of the network architecture is displayed in Figure 154.

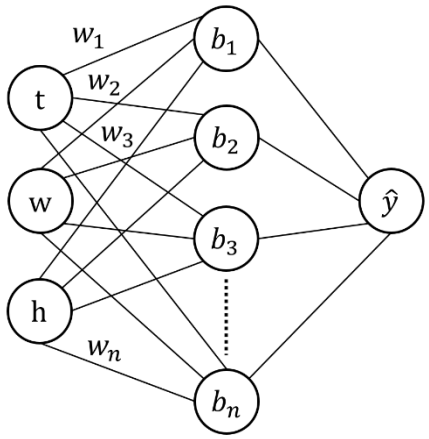


Figure 154 Schematic of predictor ANN

The network was trained until one of several stopping conditions was met. The training would end if the maximum number of epochs exceeded 1000; if further improvements in the predictions become negligible and the performance gradient fell below $1x10^{-7}$; or if the adaptive variable mu, which is increased in increments by the algorithm until the MSE improves raises to a maximum value of $1x10^{10}$. Individual networks were trained to predict values of the relative specific stiffness and relative specific strength separately for a given unit cell geometry. In both cases the training was repeated to determine the number of neurons in

the hidden layer necessary to achieve a target of 99.9% regression between the predictions and known values for the test partition of 850 observations. The test MSE and regression are plotted against the number of neurons in Figure 155. The performance improves for both models as the size of the network is increased. For the specific strength network the improvements begin to plateau beyond 150 neurons for both MSE and regression. At 200 neurons the networks provided relative specific stiffness and strength predictions that offered regression of 99.99% and 99.94% respectively and corresponding MSE of 3.20MPa and 0.0195MPa. This architecture was selected for use in the final predictor models and trained in 1 minute 56 seconds on the same hardware used to train the inverse model.

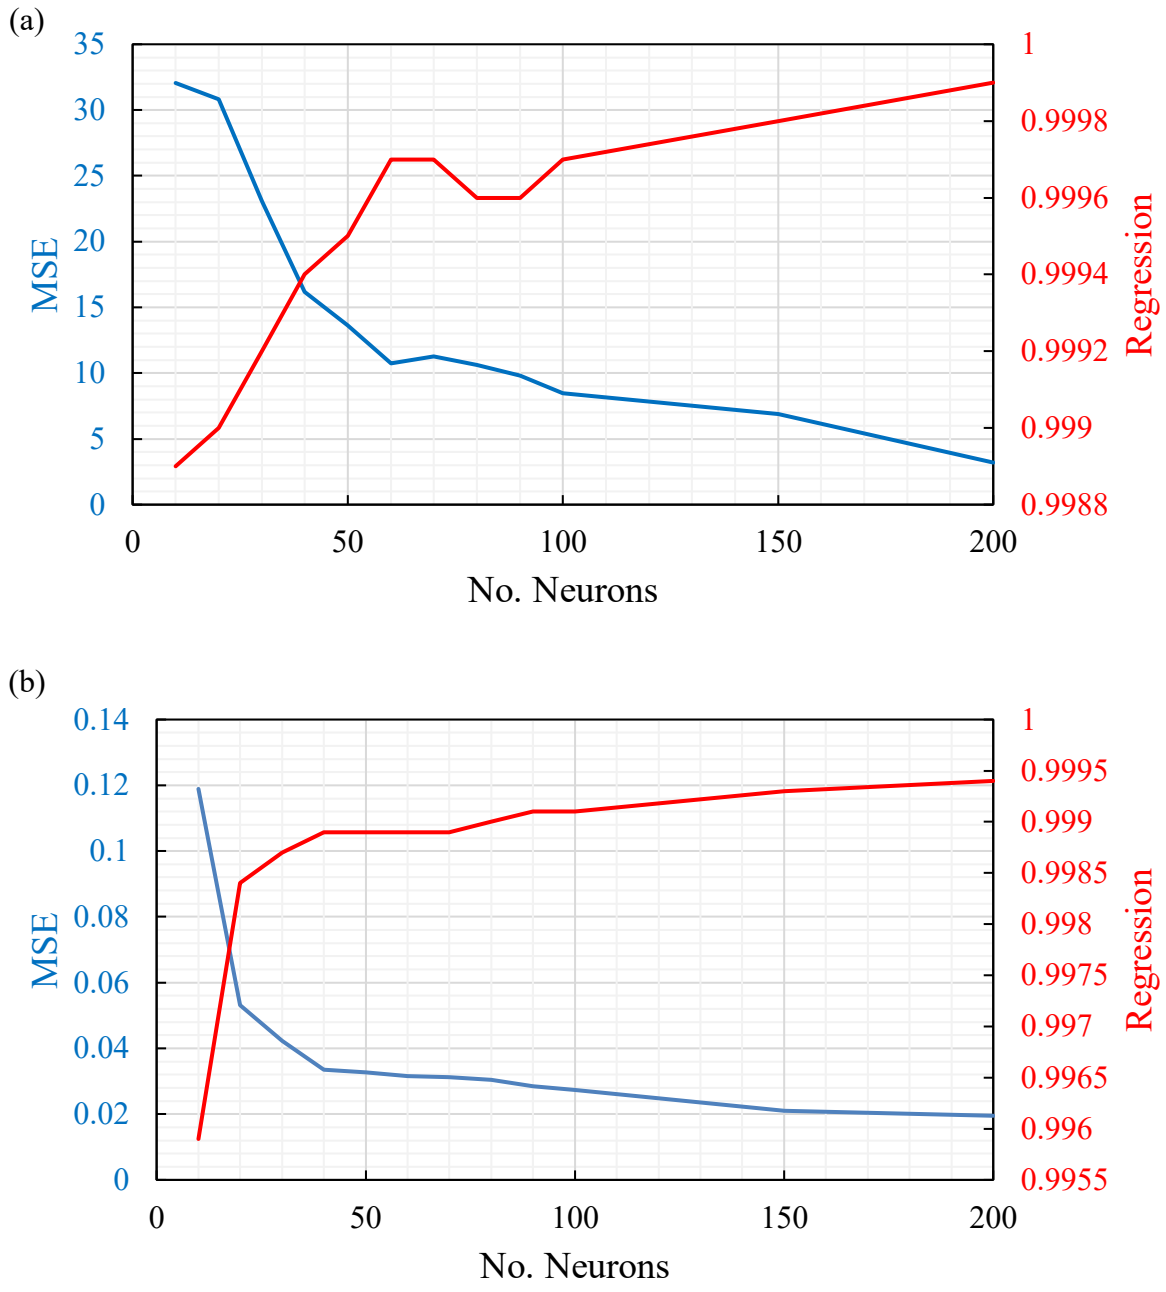


Figure 155 Test mean squared error and regression vs number of hidden layer neurons for networks trained to predict (a) relative specific stiffness and (b) relative specific strength

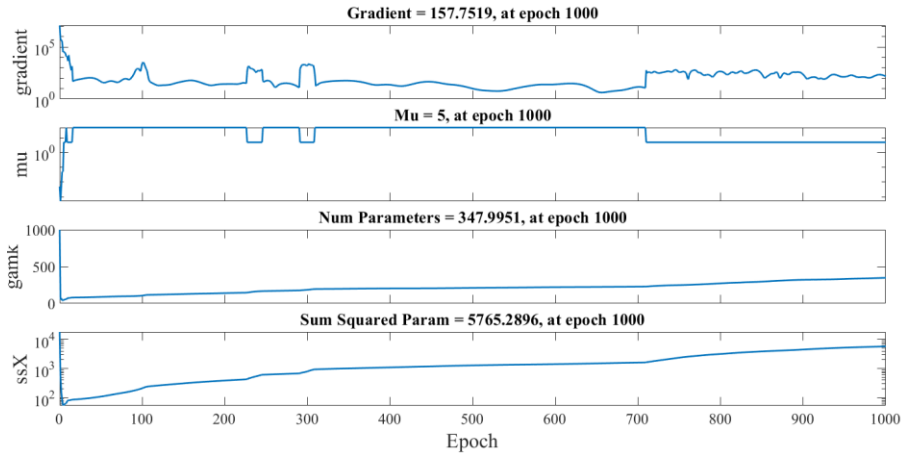
5.2.2 Training and Test Performance

The training progress and the performance of the network predictions are summarised in Table 27. Both networks stopped as a result of reaching the maximum number of epochs. The training state at every epoch is displayed in Figure 156. The mean squared error measured from values taken from the training and test dataset partitions is plotted in Figure 157. A high performance gradient is observed at the start of the training, which is reflected by a sharp decrease in MSE. During this period the mu variable is increased rapidly as the network initiates. Following this the mu variable remains consistent with occasional step changes. The number of network parameters continues to increase as the network attempts to extract additional levels of complexity, represented by the size of the weight and bias vector, gamk and the sum of the squared weight and bias vector, ssX. This results in further gradual reductions in MSE. The network trained for relative specific strength showed greater agreement between training and test MSE values throughout, while stiffness predictions began to separate after approximately 300 epochs. This may be due to a greater range in possible output stiffness values within the dataset (1176.4-1922.9MPa) compared to the range of strength values (24.1-45.2MPa). However test error continued to decrease with no evidence of overfitting and both networks offered test regression exceeding 99.9%.

Table 27 Training progress and test performance of predictor neural networks

Output	**Training Progress**				**Prediction Performance**			
Relative Specific Stiffness (MPa)	**Unit**	**Initial Value**	**Stopped Value**	**Target Value**		**Observations**	**MSE**	**R**
	Epoch	0	1000	1000				
	Elapsed Time	-	00:01:56	_				
	Performance	9.03E+06	2.10E+00	0	**Training**	3398	2.0999	0.9999
	Gradient	1.23E+07	1.58E+02	1.00E-07				
	Mu	0.005	5.00E+00	1.00E+10				
	Effective # Param	1.00E+03	3.48E+02	0	**Test**	850	3.1979	0.9999
	Sum Squared Param	1.80E+04	5.77E+03	0				
Relative Specific Strength (MPa)	**Unit**	**Initial Value**	**Stopped Value**	**Target Value**		**Observations**	**MSE**	**R**
	Epoch	0	1000	1000				
	Elapsed Time	-	00:00:25	_				
	Performance	2.41E+03	1.70E-02	0	**Training**	3398	0.0170	0.9994
	Gradient	3.30E+03	6.48E-03	1.00E-07				
	Mu	0.005	5.00E+00	1.00E+10				
	Effective # Param	4.01E+02	2.92E+02	0	**Test**	850	0.0195	0.9994
	Sum Squared Param	3.93E+03	7.37E+03	0				

(a)



(b)

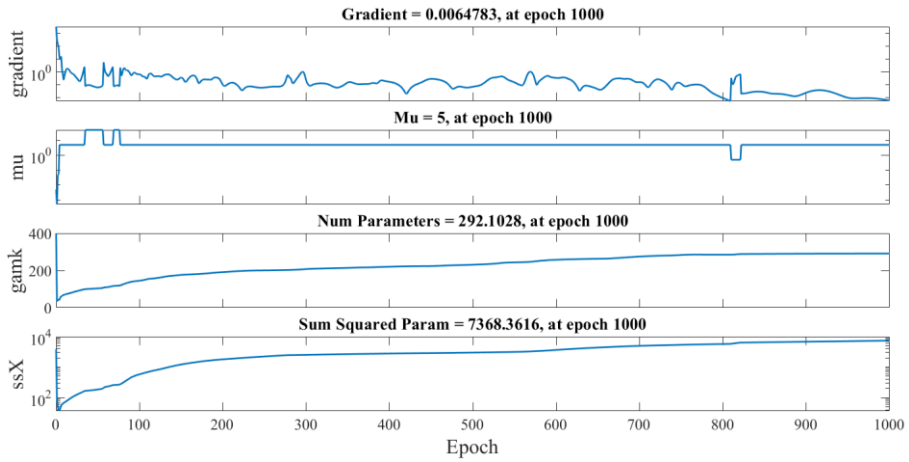


Figure 156 Training state history of neural networks for predicting; (a) relative specific stiffness; (b) relative specific strength

(a) (b)

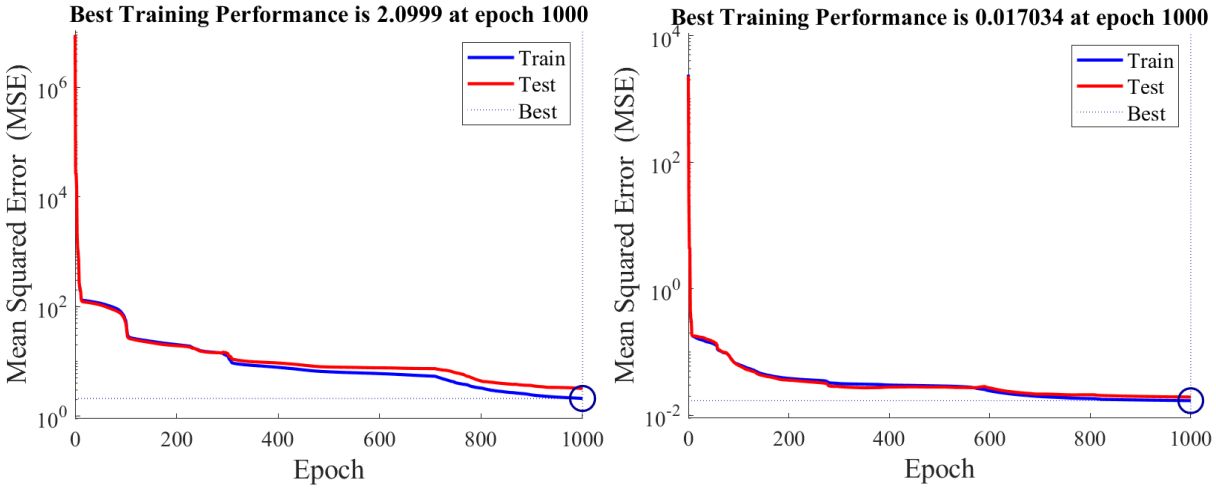


Figure 157 Training performance history of neural networks for predicting; (a) relative specific stiffness; (b) relative specific strength

The correlation between the network predictions and the known target values of relative specific stiffness and strength are plotted in Figure 158 for the test dataset partition of 850 observations. Both networks showed strong predictive performance suitable for use in further multi-objective optimisations. Relative specific stiffness was predicted to a mean test error of 1.79MPa and specific strength 0.14MPa. The distribution of error values across the entire dataset is shown in Figure 159. In 91.3% of the instances the stiffness error fell below 1.5MPa. The strength predictions showed a 91.7% chance of the error remaining below 0.17MPa.

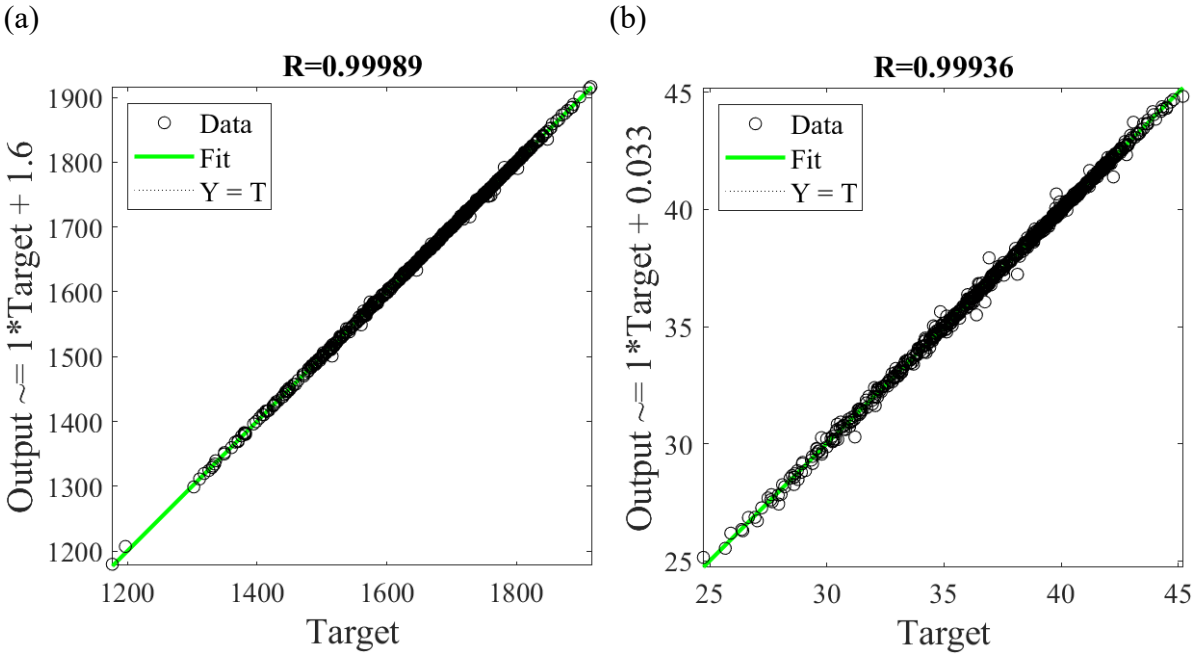


Figure 158 Test regression between known target values and output predictions of (a) relative specific stiffness; (b) relative specific strength

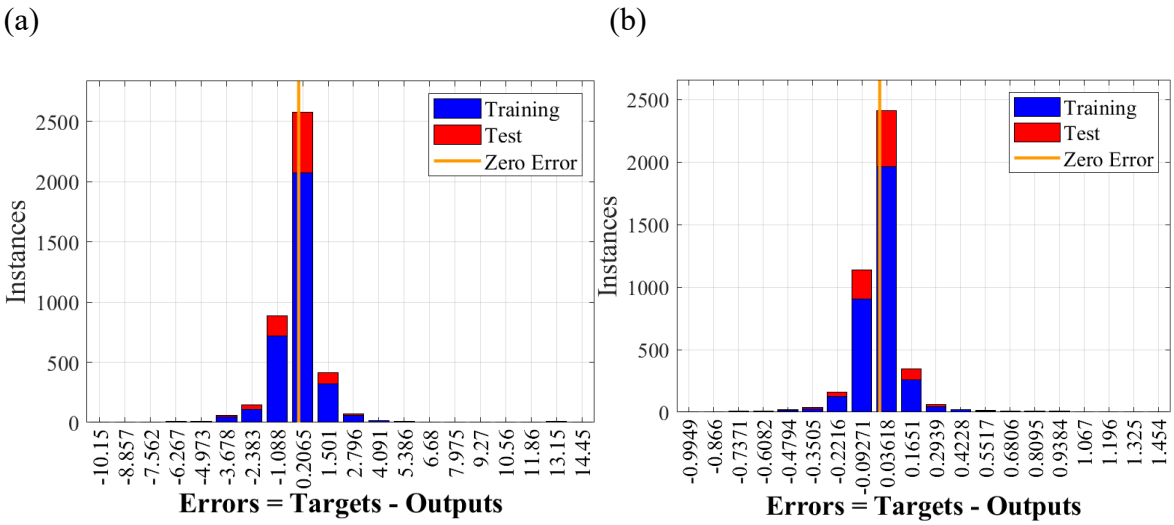


Figure 159 Histograms of error between known target values and output predictions of (a) relative specific stiffness; (b) relative specific strength

5.2.3 Multi-objective Genetic Algorithm

The predictor networks were combined with a multi-objective genetic algorithm using a script written in MATLAB with the optimisation toolbox. The objective function of the algorithm consisted of 4 objectives:

1. Maximise relative specific stiffness
2. Maximise relative specific strength
3. Maximise first harmonic peak sound absorption coefficient amplitude
4. Minimise first harmonic peak frequency

Each objective was represented as a function of the geometry parameters of the unit cell; the wall thickness, t, the channel width, w, and the channel height, h subject to the constraints:

- $1mm \leq t \leq 4mm$
- $4mm \leq w \leq 10mm$
- $10mm \leq h \leq 25mm$

To set the objectives the script called the trained networks to predict the relative specific stiffness and strength values in MPa. The first harmonic peak amplitude and frequency were predicted using the analytical solution outlined in section 5.1.1. The length of the resonator cavity was calculated automatically from the wall thickness, channel width and unit cell size. A function was then called to determine the absorption coefficient-frequency response between 0-3000Hz and extract the locations and magnitudes of all the peaks in the response using the findpeaks command. The peak occurring at the lowest frequency was then returned as the objective value. If no peaks were found the function would instead output the point at which the absorption took maximum value to prevent convergence issues. By default the GA solver worked by minimising all of the outputs provided from the objective function script. In order to maximise objectives the values were multiplied by -1 as necessary before output and then minimised by the algorithm. The GA selected an initial population of 50 possible unit cell designs using a quasi-random Sobol sequence creation function. Every member of the population was ranked according to its fitness values and the top 5% were designated as elite members for guaranteed survival into the following generation. At each stage of the algorithm a set of individuals were chosen to act as crossover parents using tournament selection. A scattered crossover function was used to select genes at random from the parents. The children that resulted from the crossover made up 80% of the non-elite members at every generation. The remaining 20% were created by random mutation to promote genetic diversity. The algorithm was set to continue to iterate until any changes in the average spread of the Pareto solutions fell below a function tolerance of $1x10^{-6}$, up to a maximum of 300

generations or 50 stall generations. The solver had no time limit and the fitness values were able to take any value down to negative infinity.

The algorithm converged after 69.05 seconds and 110 generations, at which point the average change in Pareto spread satisfied the tolerance limit. The mean Pareto spread at every generation is shown in Figure 160. A large initial spread was followed by fluctuations between values exceeding 0.9 and significantly reduced values below 0.3. These fluctuations became less frequent as the algorithm progressed before convergence on a final Pareto front of 18 solutions with average spread of 0.21. A pairwise plot of the performance objectives predicted by the neural networks and analytical theory for every solution along the Pareto front is displayed in Figure 161. The front provided a range of predicted structural properties between 1337.7-1936.8MPa relative specific stiffness and 28.8-45.5MPa relative specific strength, slightly exceeding the upper limits found within the original training dataset (1176.4-1922.9MPa and 24.1-45.2MPa). A positive linear correlation is observed between stiffness and strength objectives, which may be expected as the unit cells become both stronger and stiffer with the addition of material. Across the range of structural responses a strong trend is observed favouring solutions with low first harmonic frequency. 14 of the 18 solutions are predicted a first harmonic peak below 600Hz. 3 of the remaining 4 are in the upper quartile for relative specific strength and no solutions show peaks above 1500Hz. The majority of solutions provided absorption amplitude of at least 80%. 3 solutions fell below 70% absorption however all 3 provided high relative specific structural properties that exceeded 1890MPa stiffness and 44MPa strength. Of all the solutions, 1/3 provided an absorption amplitude greater than 99%. These trends towards higher structural properties and greater absorption at lower frequencies are also illustrated in Figure 162, which displays histogram plots of the Pareto front solutions for each design objective. 88.9% of the solutions predicted relative specific stiffness above 1500MPa, the largest number of which fell between 1740-1860MPa. 3 solutions were predicted to feature relative specific strength values larger than the greatest values found in the original dataset. 2 of these resulted from the solver converging on the same solution and featured identical geometries. All other points along the front were distinct. The highest number of solutions were between 35-45MPa relative specific strength.

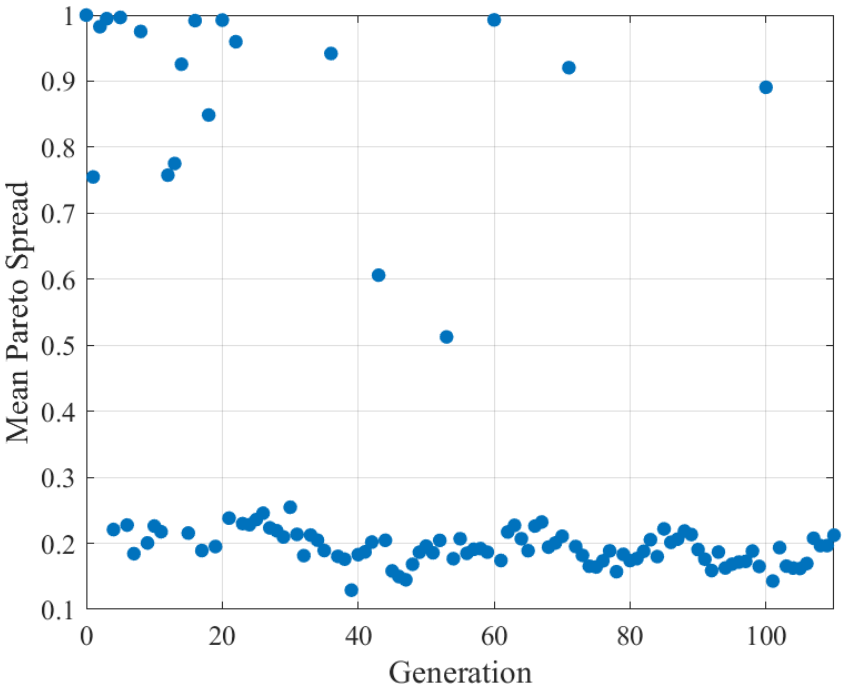


Figure 160 Multi-objective genetic algorithm optimiser history

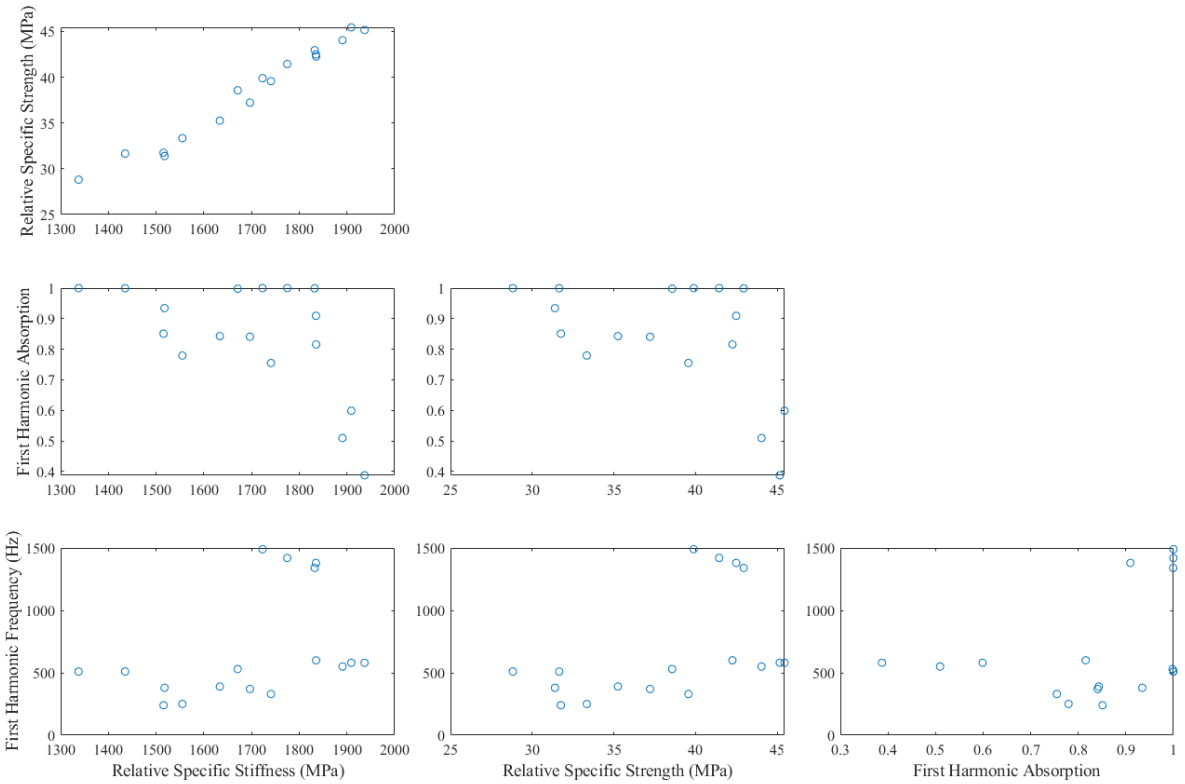


Figure 161 Pareto front of multi-objective GA solutions

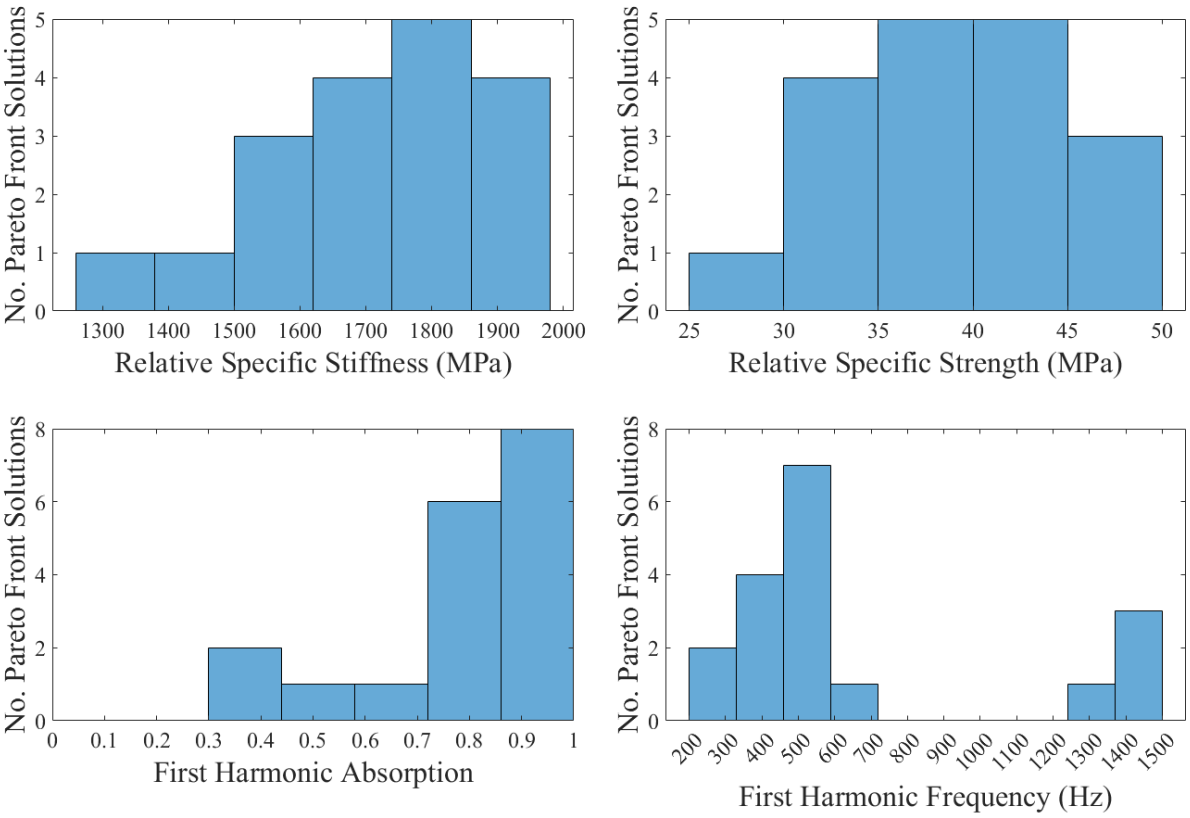


Figure 162 Performance objective histograms of Pareto front solutions

5.2.4 Experimental Validation

Each of the Pareto front solutions was manufactured via FDM and tested for both sound absorption and structural properties using the experimental procedures outlined in section 5.1.5. The analytical first harmonic peak values are compared against those obtained by impedance testing, and the experimental structural properties are compared against those obtained using the non-linear buckling FEA model in Table 28. All experimental sound absorption peaks fell within 20.77Hz of the natural frequencies predicted by theory. The greatest error in the absorption amplitude was observed in examples with resonator channel cross-sectional area below 50mm^2; solutions 4, 7 and 17. The analytical theory predicted the peak absorption coefficient of the remaining solutions to a mean error of 0.029. The largest overall error was 0.129, recorded for solutions 4 & 7, likely as they featured the lowest channel area found anywhere in the Pareto front of 41mm^2. However this solution provided the strongest overall unit cell with an experimental value of 42.76MPa. The sound absorption performance was compromised as a result of maximising the strength and the test specimen recorded the lowest peak absorption coefficient in the study of 0.5285. The predicted sound absorption and mechanical responses are compared against the experimental measurements in Figure 163. Solutions 2 and 16 showed the greatest discrepancies between experimental stiffness and values predicted by FEA. In both cases the experimental measurement exceeded

the prediction. The remaining Pareto front solutions offered relative specific stiffness values to a mean error of 12.57MPa, 0.60% of the maximum experimental value. The buckling FEA model predicted the relative specific strength of the Pareto front solutions to a mean error of 4.12MPa. Solution 2 provided the best combined structural performance, recording the greatest specific stiffness of 2111.77MPa while maintaining a specific strength of 39.27MPa, 91.8% of the maximum value in the study. In addition the unit cell is able to offer 90% absorption at frequencies above 1191Hz. The minimum first harmonic response is plotted along with the measured structural properties in Figure 164. The Pareto front may be narrowed down further by selecting only solutions that achieve a target first harmonic peak amplitude or frequency for use in applications where low-frequency sound absorption is of heightened importance. Solution 18 shows the lowest first harmonic frequency (350Hz) to achieve approximately 90% or greater absorption amplitude, as displayed in Figure 165. The overall lowest first harmonic frequency was recorded for solution 5, which measured a natural frequency of 232.55Hz. As shown in Figure 166, the resonator was able to maintain 79.99% absorption at this frequency at the cost of reduced structural performance. Plots of the measured properties for all Pareto solutions can be found in Appendix K.

Table 28 Sound absorption and structural properties of GA optimised Pareto front solutions

Solution No.	Solution Parameters			Sound Absorption (Analytical)		Sound Absorption (Experimental)		Structural (FEA – Buckling Model)		Structural (Experimental)	
	t (mm)	w (mm)	h (mm)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)
1	3.02	9.31	14.38	0.9970	1250	0.9988	1270.77	1719.54	39.65	1746.73	34.33
2	3.92	9.90	13.39	0.9978	1190	0.9893	1191.21	1827.87	42.59	2111.77	39.27
3	3.82	4.07	13.53	0.6144	520	0.5564	520.32	1761.90	41.93	1757.61	38.62
4	3.74	4.07	10.08	0.3994	530	0.5285	543.65	1504.62	35.05	1492.01	42.76
5	1.15	4.28	18.45	0.8514	240	0.7999	232.55	1378.03	28.46	1395.38	25.34
6	1.27	9.37	11.73	0.9999	510	0.9974	518.41	916.85	19.83	897.41	20.73
7	3.74	4.07	10.08	0.3994	530	0.5285	543.65	1504.62	35.05	1492.01	42.76
8	2.16	7.56	14.95	0.9994	500	0.9999	506.64	1667.47	38.22	1692.21	34.61
9	3.91	9.22	10.60	0.9177	1270	0.9993	1269.26	1747.28	40.23	1739.14	39.61
10	1.53	8.98	12.47	0.9994	510	0.9913	513.98	1117.67	24.58	1123.42	22.31
11	1.57	4.08	17.65	0.8494	350	0.7814	354.86	1331.95	28.44	1353.63	24.09
12	3.42	4.29	17.17	0.8314	500	0.8239	512.37	1832.22	42.28	1848.69	37.67
13	2.55	4.17	15.34	0.7638	310	0.7660	312.01	1381.70	31.45	1370.62	33.61
14	1.25	4.08	15.51	0.7958	240	0.7433	235.67	1198.63	25.31	1195.74	22.25
15	1.84	4.13	17.46	0.8431	330	0.7786	339.57	1409.35	30.84	1406.57	25.35
16	3.43	9.56	13.79	0.9976	1230	0.9937	1239.18	1771.01	41.30	1933.50	35.81
17	3.99	4.57	10.58	0.5210	510	0.6288	510.76	1279.25	29.75	1286.15	39.63
18	1.15	4.78	18.20	0.9287	350	0.8996	351.74	1143.10	23.35	1135.81	24.56

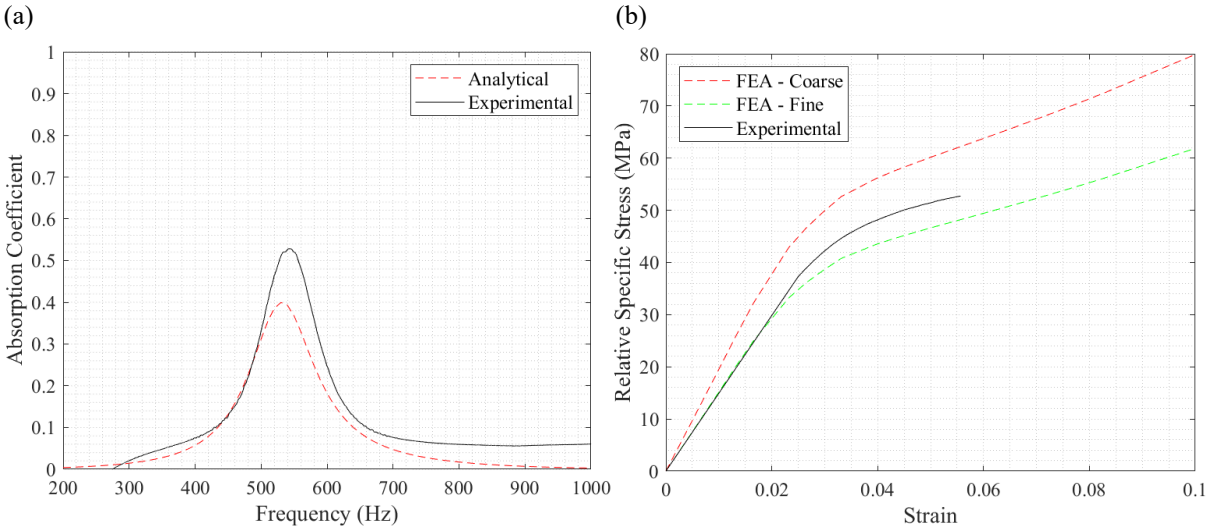


Figure 163 Pareto solution 4 material properties; (a) sound absorption response; (b) mechanical response

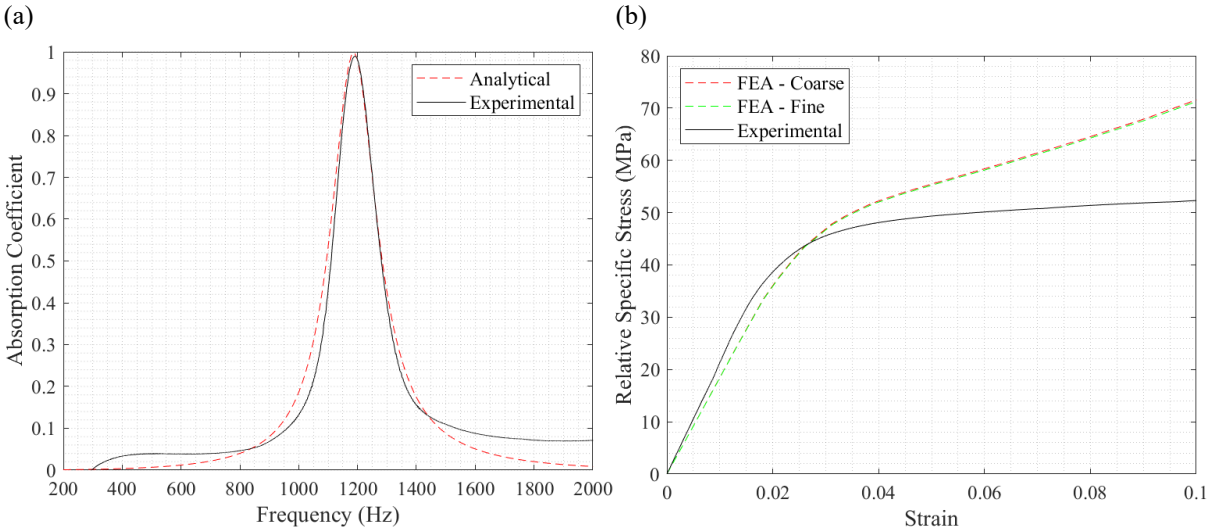


Figure 164 Pareto solution 2 material properties; (a) sound absorption response; (b) mechanical response

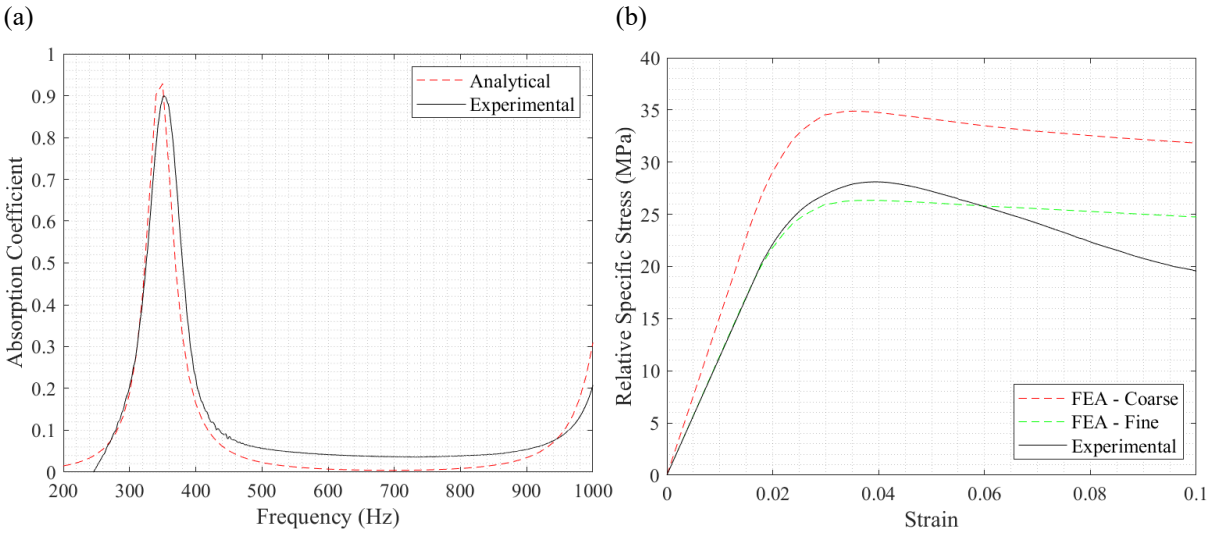


Figure 165 Pareto solution 18 material properties; (a) sound absorption response; (b) mechanical response

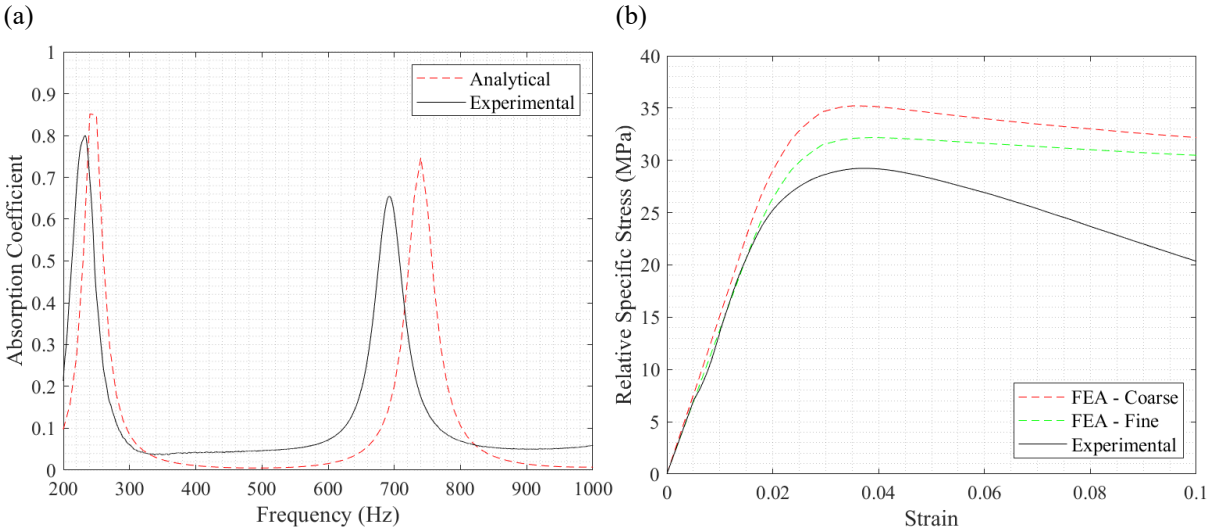


Figure 166 Pareto solution 5 material properties; (a) sound absorption response; (b) mechanical response

To compare the experimental structural performance of the Pareto front against equivalent solutions that may be found using an exhaustive search, the design points that provided the greatest values of relative specific stiffness and strength anywhere in the original dataset were extracted and manufactured for mechanical characterisation. The mechanical properties predicted by both the FEA model used to train the neural network and the revised buckling FEA model are compared against experimental measurements in Table 29. As was the case with the inverse optimiser, the coarse FEA model used to generate the dataset overpredicted the maximum stiffness and strength values while the buckling model offered greater agreement with the experimental values. In the case of the predictor direction optimiser, the network is only required to analyse the influence of the geometry parameters on increases or decreases in the structural properties. Provided the neural network has sufficient information to extract the underlying patterns and trends suitable for maximising the performance objectives, the accuracy of the exact values output by the FEA may be of less importance. The predictor network may therefore offer means to carry out rapid optimisations when the fidelity or size of the available dataset is limited, however post-validation of results is necessary before selection of individual solutions from the Pareto front to meet design criteria. Through exhaustive search of the dataset of 4248 potential unit cell designs, the maximum experimental relative specific stiffness and strength values found were 1392.15MPa and 37.82MPa respectively. Of the Pareto front extracted by the GA combined with predictor neural network, 61.1% of the solutions gave experimental stiffness above the stiffest unit cell in the exhaustive search while also simultaneously considering the first harmonic absorption performance. One third of the Pareto front provided strength greater than the strongest unit cell in the original dataset despite the trade-off with competing sound absorption objectives, as displayed in Figure 167. The use of machine learning offers the ability to rapidly predict

the material properties of a population of possible unit cell designs across multiple generations, enabling genetic optimisations of metamaterials that would not otherwise be possible. Furthermore the neural networks may be trained on empirical datasets obtained by simulation or experimental means for which there is limited or no theoretical solution. The approach presents an effective method for performing multi-objective optimisations across a diverse range of competing design objectives which may be based on unrelated underlying physical phenomena which cannot be linked analytically through other means.

Table 29 Structural properties of solution unit cells determined via exhaustive search

Performance Objective	Geometry Parameters			FEA Prediction - Original Dataset Model (MPa)	FEA Prediction - Buckling Model (MPa)	Experimental Measurement (MPa)
	t	w	h			
Max Relative Specific Stiffness	3.76	4.23	10.48	1922.93	1446.25	1392.15
Max Relative Specific Strength	3.73	4.00	13.87	45.20	39.32	37.82

(a) (b)

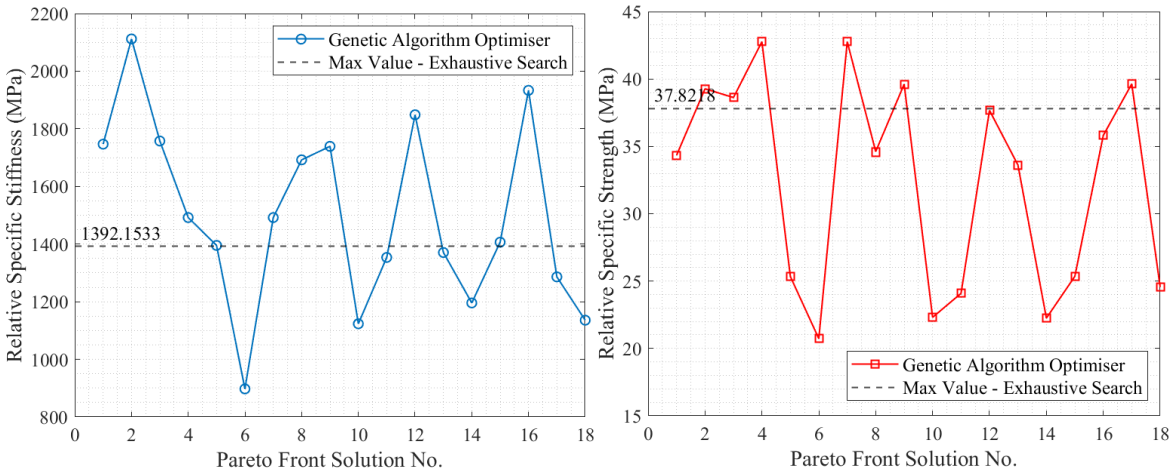


Figure 167 Comparison between experimental structural properties of genetic algorithm optimised Pareto front solutions and equivalent maximum values found by exhaustive search of the dataset of 4248 design points for (a) relative specific stiffness; (b) relative specific strength

5.3 Reducing Overall Material Thickness

The requirement for sound attenuation is accompanied by strict spatial limitations in many engineering sectors where low-frequency noise is commonplace such as automotive and aeronautical. In these applications interior packaging and lightweighting requirements are of significant importance, as is a need for tunability of structural properties to meet design criteria. There is significant demand for strong low-frequency sound absorption at minimal material thicknesses. In order to study the potential of the machine learning approach for satisfying this demand, the inverse CNN model was revised with an alternative parametric unit cell design to target QWR based metamaterials with overall material thicknesses below 6mm. At ultra-low material thicknesses the aspect ratio of the resonator cross section is expected to become considerably reduced. Table 18 suggests this is likely to cause a reduction in absorption amplitude at the lowest resonant frequencies achievable for a given channel length. A solution is proposed to overcome this issue by combining multiple identical QWR labyrinth paths in parallel to increase the effective total cross-sectional area. A new dataset is generated using the revised geometry and updated parameter limits. The inverse CNN is then retrained and the network is used to tailor both sound absorption and structural properties simultaneously while outputting only solutions that have thicknesses below the target value.

5.3.1 Dataset Generation

The unit cell layout was divided into 5 equally sized labyrinth segments that all shared a common inlet at the centre. The revised geometry is displayed in Figure 168.

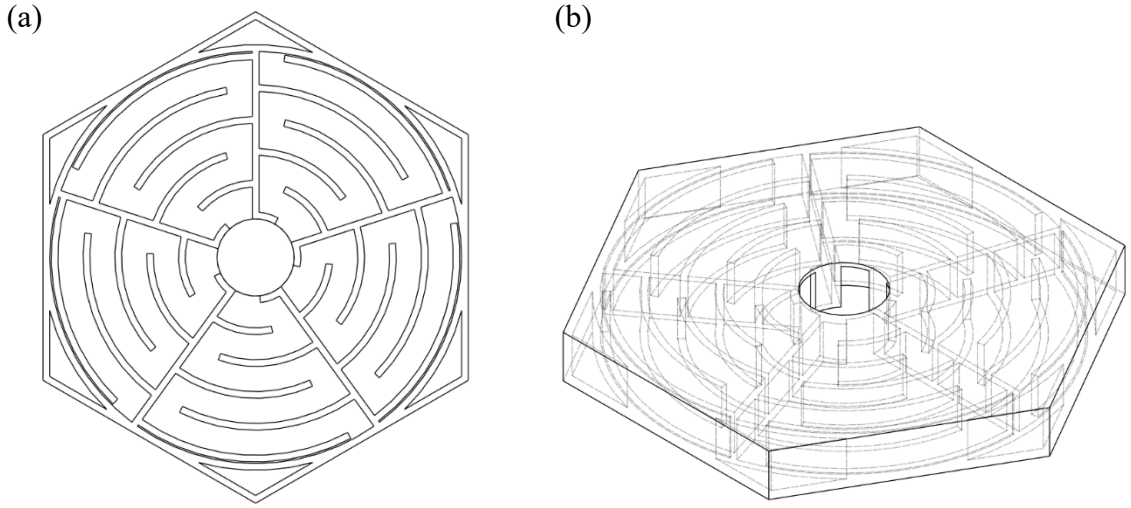


Figure 168 Revised metamaterial unit cell geometry for reduced overall thickness; (a) cutaway view (b) 3D view

A parametric CAD model was developed to allow for automated updates of the geometry parameters. The unit cell featured the same honeycomb tessellation with apothem of size 25.4mm as previously. The parameters allowed to vary remained the wall thickness, t, channel width, w, and channel height, h. The ranges of the limits of each parameter were adjusted to the values shown in Table 30.

Table 30 Revised geometry parameter limits for reduced thickness model dataset

Parameter	t(mm)	w(mm)	h(mm)
Lowest Value Limit	0.8	3	1
Greatest Value Limit	3	7	5

The maximum channel height was set at 5mm and the thicknesses of the front and back walls were both locked at 0.5mm to ensure maximum overall thickness of 6mm. The radius of the central inlet varied depending on the geometry parameters according to $r_{inlet} = \frac{5(w+t+x)}{2\pi}$, where x remained constant at 2mm. The entrance to each individual resonator was assigned an arc length of w mm. The concentric circular walls were created working outwards from the inlet. The interior radius of the n^{th} wall was calculated using $r_n = r_{inlet} + n(t + w)$. The angles of the gaps in the labyrinth walls were updated automatically to ensure an arc length of w mm according to $\theta_n = 180w/\pi r_n$. For larger values of w, any walls that fell outside the unit cell were subtracted using a Boolean operation and the remaining geometry combined with the outer circular and hexagonal walls of thickness t mm.

The script used to generate the sound absorption dataset was updated to automatically calculate the maximum channel length of each segment of the revised geometry. The analytical responses were calculated taking into account the increased effective overall area by multiplying the area by the number of resonators. To determine a suitable range of frequencies for the neural network to target, specimens were manufactured at the longest possible channel length that may occur for the chosen parameter limits. This provided the lowest possible first harmonic peak frequency for the unit cell design. The greatest channel length was found when the largest number of complete concentric circular channels fit entirely within the 50.8mm diameter. This occurred at $t = 0.9mm$ and $w = 3mm$, giving maximum length of each individual resonator of 87.45mm. The experimental absorption responses are compared against the analytical predictions in Figure 169 using varied channel heights to provide overall material thicknesses of 3mm, 4mm and 5mm. The amplitudes of the experimental absorption peaks were 0.51 at 966.8Hz, 0.66 at 972.3Hz and 0.91 at 963.2Hz

respectively. In each case the experimental values exceeded the analytical predictions, which all occurred at 960Hz at absorption coefficients of 0.22, 0.44 and 0.64. The results represent the lowest frequency sound absorbers that may be obtained for the given geometry layout if consideration of structural properties is not also required.

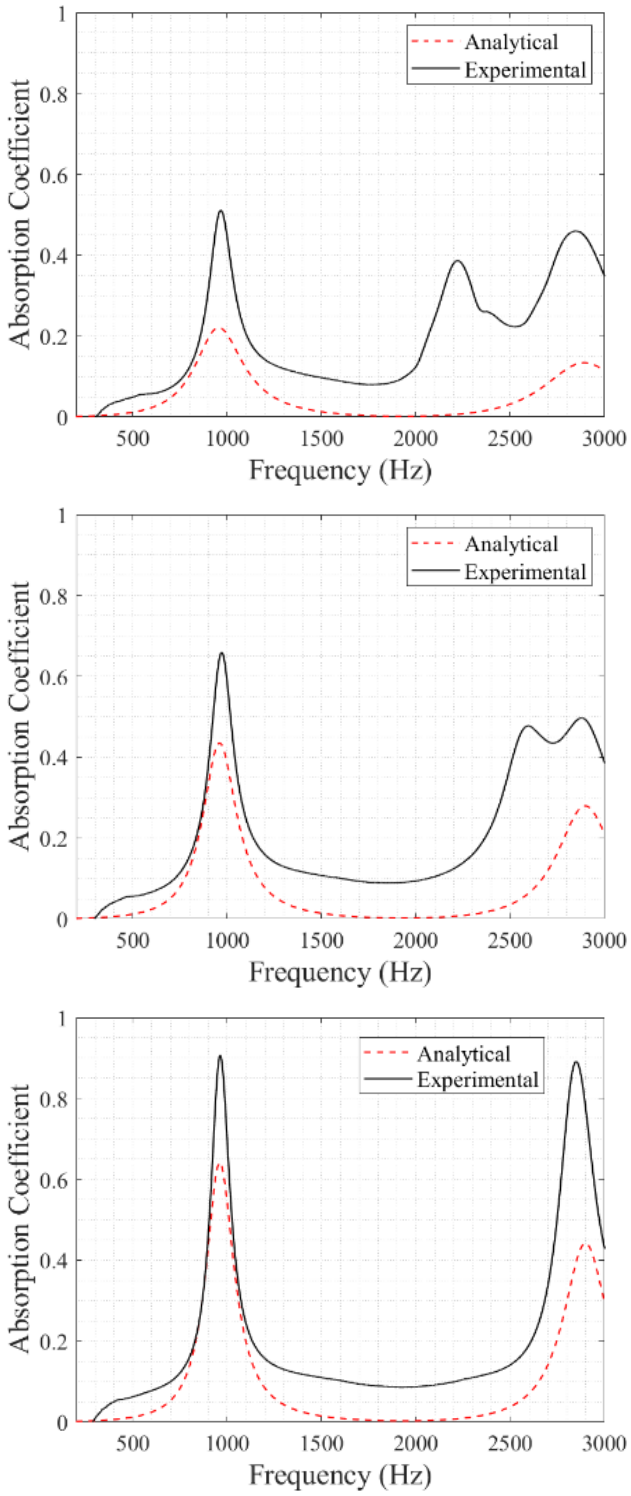


Figure 169 Comparison of analytical and experimental sound absorption responses for unit cell revised to reduce overall thickness at greatest possible channel length, where t=0.9mm and w=3mm; (a) h=2mm, overall thickness=3mm; (b) h=3mm, overall thickness=4mm; (c) h=4mm, overall thickness=5mm

A target frequency range of 500-1500Hz was selected for inverse tuning of sound absorption alongside structural characteristics to allow for the lowest possible first harmonic peaks. The analytical absorption coefficient values were calculated at intervals of 20Hz, providing 51 datapoints for each observation within the dataset. This was combined with 51 datapoints along the equivalent force-displacement curve to create the input signal sequence vector of 102 entries. The non-linear FEA model outlined in section 5.1.2 was updated with the revised parametric geometry, meshed with 2mm solid elements. The front face was fully fixed in all degrees of freedom and a compressive displacement load applied to the back face in the negative z-direction as shown in Figure 170. The displacement load was ramped on linearly up to 0.5mm across 5 load steps of 0.1mm. The reaction force was extracted at the fixed boundary constraint every 0.01mm. The range of possible mechanical characteristics within the dataset is displayed in Figure 171, which features the FEA responses for a full factorial test plan with 3 levels based on the chosen geometry parameter limits. A dataset of 5000 design points was selected using Latin hypercube sampling, of which 4985 gave converged FEA solutions and were used to train the inverse CNN.

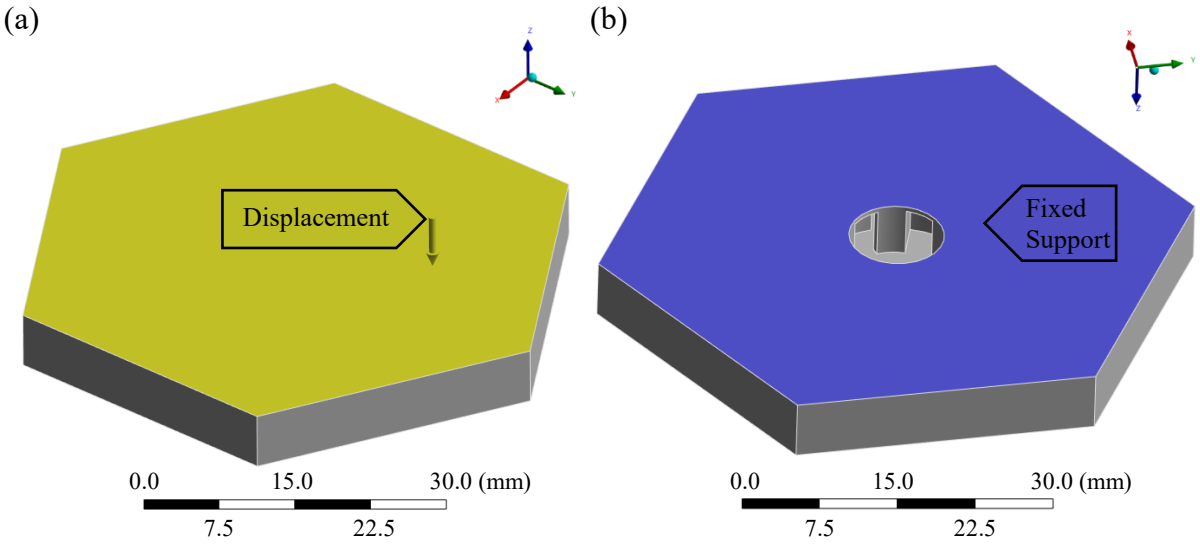


Figure 170 FEA model loading and boundary conditions for unit cell revised to reduce overall thickness

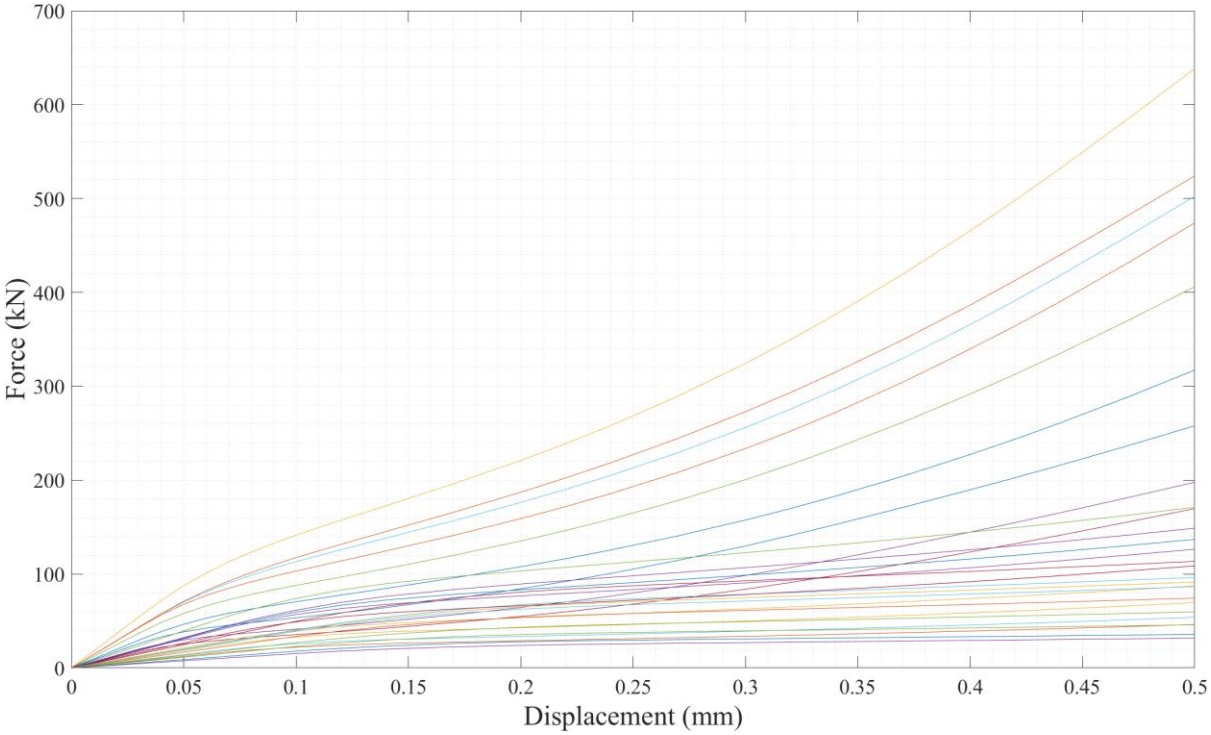


Figure 171 Range of possible FEA responses for reduced thickness model design space. Design points were selected from a full factorial test plan with 3 levels based on the chosen geometry parameter limits.

5.3.2 Inverse Convolutional Neural Network

The architecture detailed in section 5.1.3 was used to train an inverse CNN using the same dataset partition ratio, hyperparameters and training options as previously. The RMSE and loss are plotted at each training and validation iteration in Figure 172. Greater agreement is observed between training and validation RMSE values in comparison to the previous inverse model, likely due to the increased dataset size and the reduction in the ranges of the output parameters. The performance values after the final iteration are displayed in Table 31. A small discrepancy remained between the training and validation performance. The trend in the training history suggests this would be unlikely to improve if the network was trained for a greater number of epochs. However predictions exceeded 99% regression for all 3 geometry parameters when compared against known values for the test data partition of 748 observations. The correlations between the predicted geometry parameters output by the network and the actual known values in the test dataset are shown in Figure 173. Outliers remain a possibility in rare cases though the network showed strong overall predictive performance suitable for inverse tailoring of the material properties.

Table 31 Summary of training performance and test regression for inverse CNN trained to predict the geometry parameters of the metamaterial unit cell designed for reduced thickness

Training Performance			Regression	
RMSE	**Training**	0.0591	**Parameter**	**R Value**
	Validation	0.0938		
Loss	**Training**	0.0017		
	Validation	0.0044	**t**	0.99367
Training Time	30min 20sec		**w**	0.99337
No. Iterations	34000		**h**	0.99455

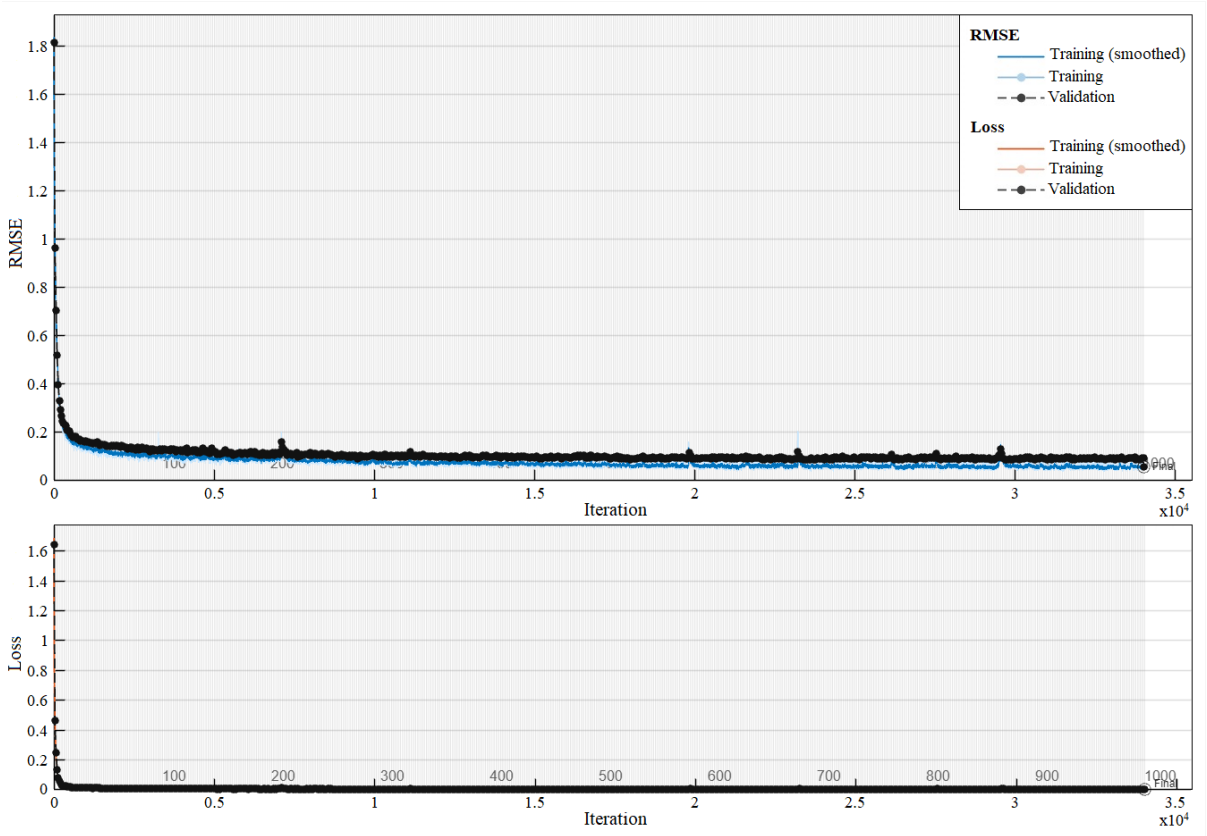


Figure 172 Training history of inverse convolutional neural network for reduced thickness model

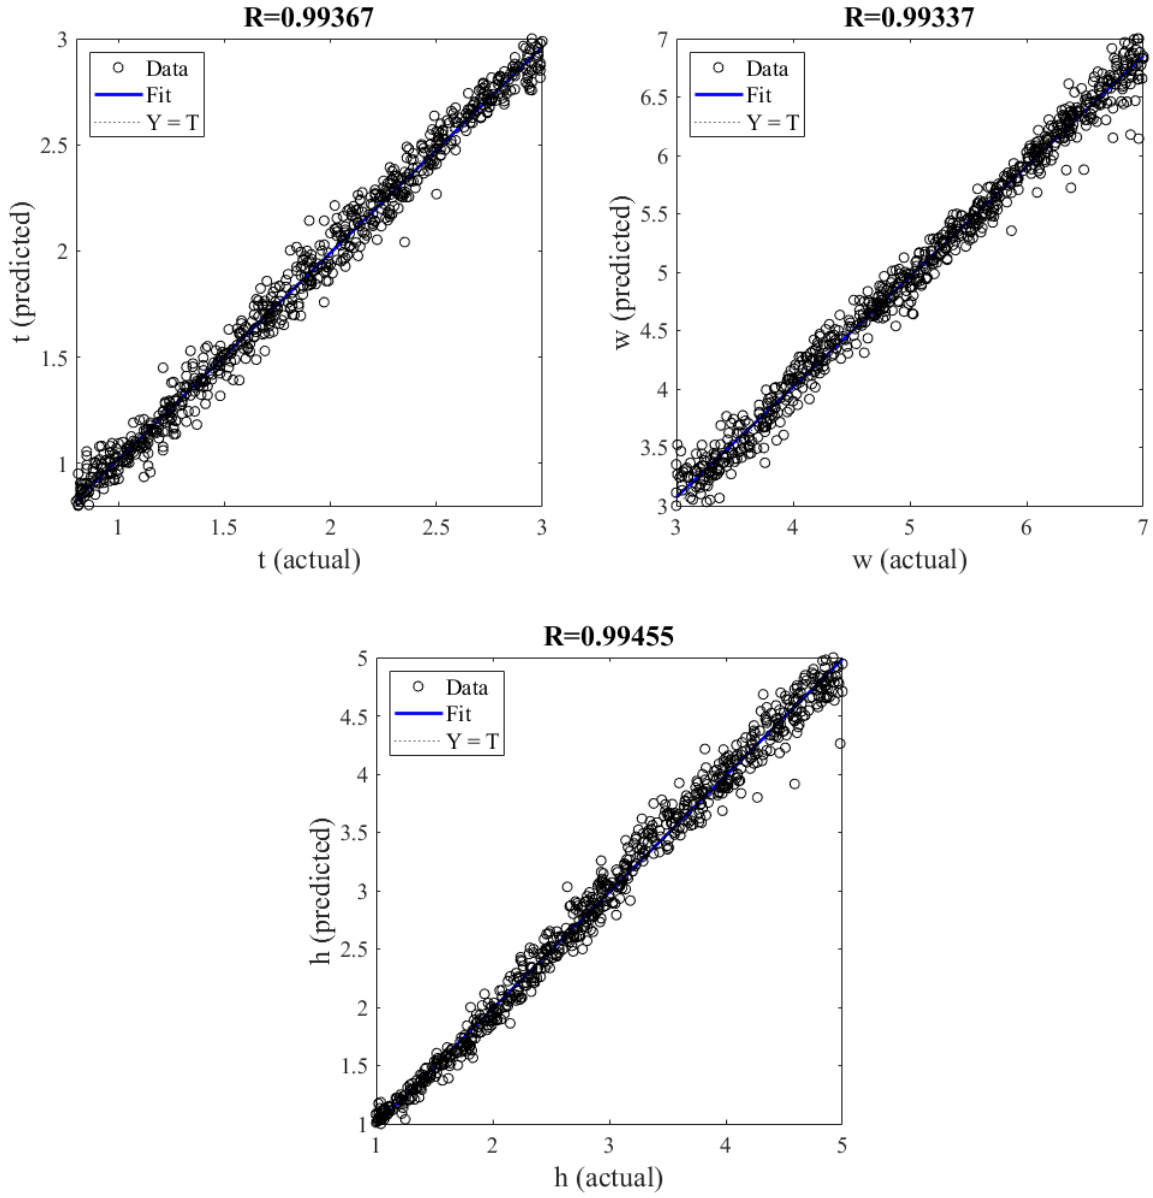


Figure 173 Regression between inverse network predictions and known unit cell geometry parameters for metamaterials with reduced overall thicknesses limited to below 6mm

5.3.3 Inverse Optimisation and Experimental Validation at Reduced Thickness

The trained network was used to tune sound absorption and structural properties of unit cells below 6mm overall thickness by requesting spike and bi-linear inputs. A set of optimised example specimens were manufactured via FDM and tested for acoustic impedance and compressive mechanical characteristics according to the methods detailed in section 5.1.5. Examples of the sound absorption and structural test specimens are shown in Figure 174. The specimens were printed with the inlet face flat downwards against the build plate. The FDM process successfully bridged across the labyrinth channels and the inlet gap to create the back wall, sealing the resonators without requiring sacrificial support structures. The requested first harmonic sound absorption peaks are compared against the analytical and experimental measurements for the solutions returned by the CNN in Table 32. Table 33 contains the equivalent structural properties for each example. The solution geometry parameters returned by the network in each case are shown in Table 34. The input signals provided to the network are compared against the analytical or FEA predictions and the experimental characteristics of every unit cell in Figure 175. The CAD geometries for each solution can be found in Appendix L. All examples were targeted to provide a first harmonic absorption peak at 1200Hz. Examples I-L demonstrate the ability of the network to lock onto a specific target frequency and then allow the overall thickness to be fine-tuned according to the design requirements by updating the amplitude of the absorption. Examples M-O show additional tuning of structural properties while maintaining at least 90% absorption at 1200Hz. The CNN was able to provide solutions with material thicknesses between 3.64-5.66mm that gave experimental absorption peak amplitudes between 84.3 - 99.9%. In all cases the experimental absorption exceeded analytical predictions. The measured peaks occurred within 20Hz of the target frequency for all except the thinnest specimen, which gave 86.6% absorption at 1237.7Hz at 3.64mm overall thickness. By theory the optimised unit cells are expected to provide similar absorption at any frequency above the target by shortening the resonator cavity with an additional back wall termination, tuning the length to the frequency of interest. Structural properties were targeted between 900-1700MPa relative specific stiffness and 20-28MPa relative specific strength. All solutions exceeded the target stiffness. The difference between the experimental measurements and the target stiffness values showed mean average of 19.88MPa. The non-linear FEA buckling model showed mean error of 2.11% when comparing the predicted stiffness against the experimental measurement. The FEA strength predictions showed overshoots in all cases, with mean value of 6.91MPa. Solution J fell below the target strength by 0.08MPa during experimental testing as a result of the overshoot, highlighting the importance of post-validation of the optimised solutions. For the remaining

tests the network provided a degree of control for fine-tuning the material properties that allowed for compensation of the discrepancy to satisfy the performance criteria. All remaining specimens gave experimental strength above the target value by on average 0.60MPa.

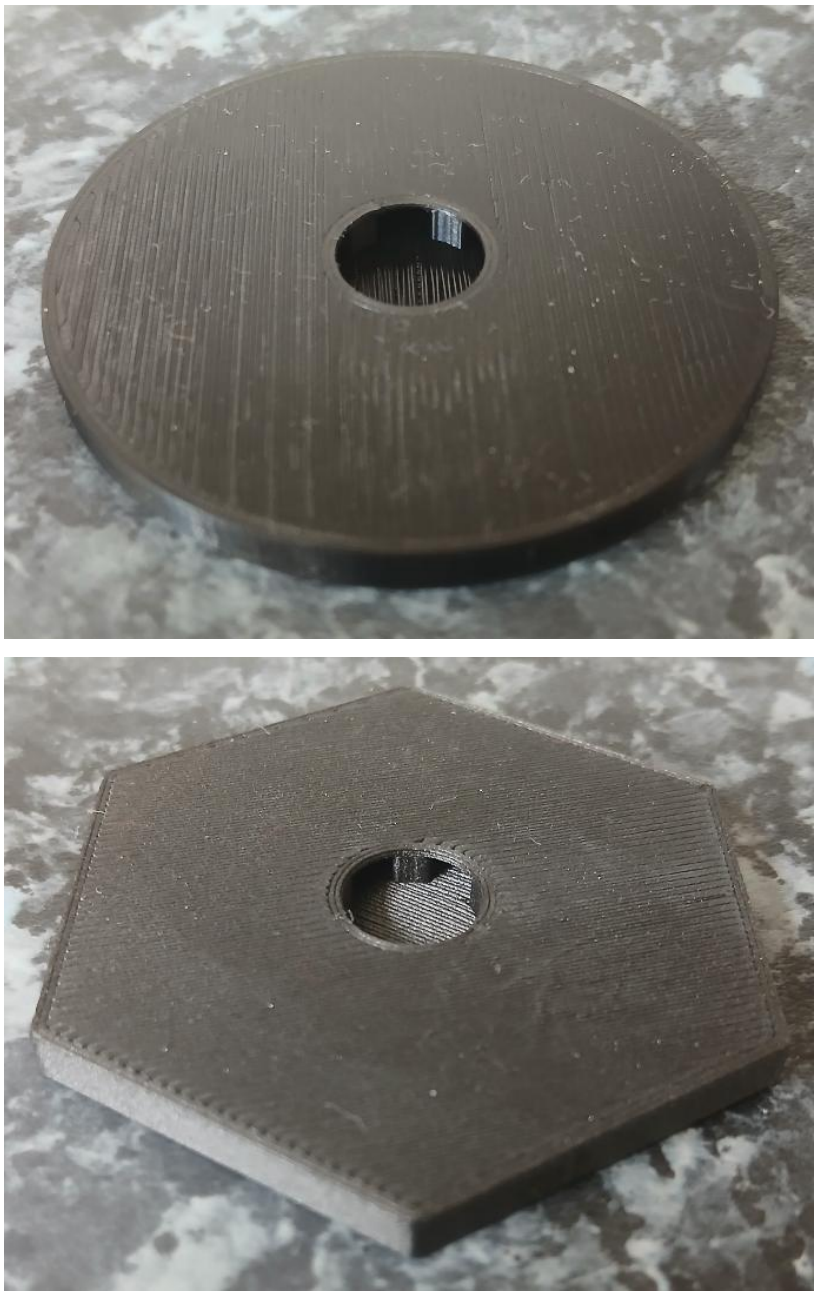


Figure 174 Example manufactured specimens for acoustic impedance testing (top) and mechanical compression testing (bottom)

Table 32 First harmonic sound absorption performance of example metamaterials optimised for overall material thicknesses below 6mm

Example	Requested Targets		Sound Absorption (Analytical)		Sound Absorption (Experimental)	
	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)	First Harmonic Absorption Coefficient	First Harmonic Frequency (Hz)
I	0.5	1200	0.4648	1200	0.8660	1237.67
J	0.5	1200	0.5148	1200	0.8425	1213.94
K	0.6	1200	0.5775	1210	0.8542	1192.51
L	0.5	1200	0.6166	1220	0.9447	1180.04
M	0.9	1200	0.8434	1210	0.9999	1205.29
N	0.9	1200	0.8071	1210	0.9957	1193.72
O	0.6	1200	0.8318	1200	0.9872	1179.64

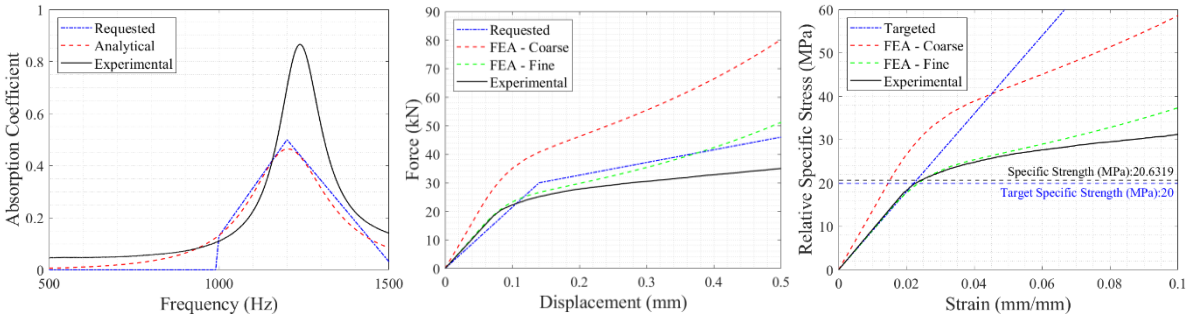
Table 33 Mechanical properties of example metamaterials optimised for overall material thicknesses below 6mm

Example	Requested Targets		Structural Properties (FEA - Buckling)		Structural Properties (Experimental)	
	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)
I	900	20	901.83	20.81	928.16	20.63
J	950	22	954.33	22.32	961.28	21.92
K	1550	25	1541.68	36.14	1557.44	25.97
L	1700	28	1695.53	39.84	1728.14	28.27
M	1350	24	1336.37	31.33	1369.56	24.83
N	1550	26	1543.97	36.12	1560.79	26.74
O	1600	26	1555.24	36.32	1633.80	26.18

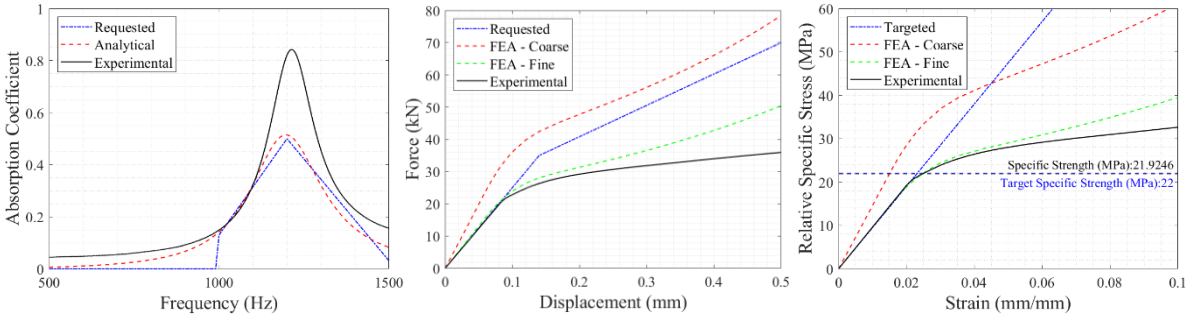
Table 34 Geometry parameters of example metamaterials optimised for overall material thicknesses below 6mm

Example	Solution Unit Cell Geometry Parameters						
	t (mm)	w (mm)	h (mm)	Overall Thickness (mm)	Mass (g)	Solid Volume (mm^3)	Relative Density
I	0.93	3.71	2.64	3.64	3.59	3860.19	0.4749
J	1.14	3.54	2.97	3.97	3.97	4269.14	0.4815
K	1.36	3.28	3.50	4.50	4.69	5038.51	0.5013
L	1.63	2.98	4.08	5.08	5.62	6042.47	0.5323
M	0.97	3.68	4.46	5.46	4.72	5075.62	0.4160
N	1.16	3.48	4.47	5.47	5.07	5445.31	0.4458
O	1.23	3.46	4.66	5.66	5.22	5607.98	0.4436

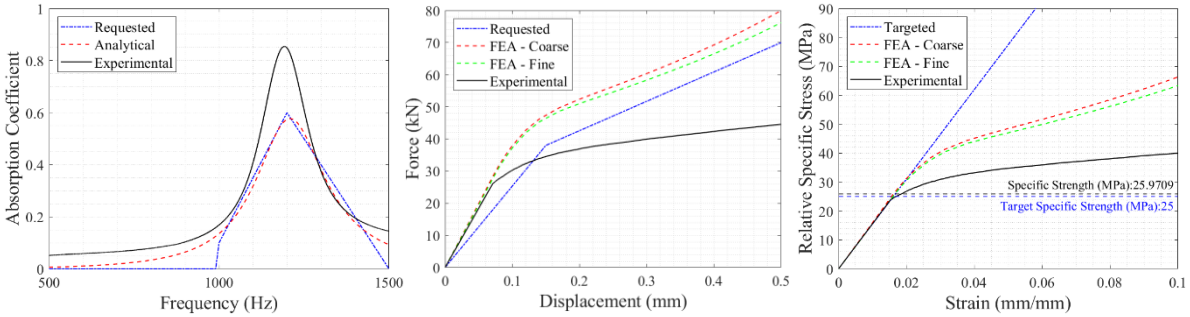
Example I



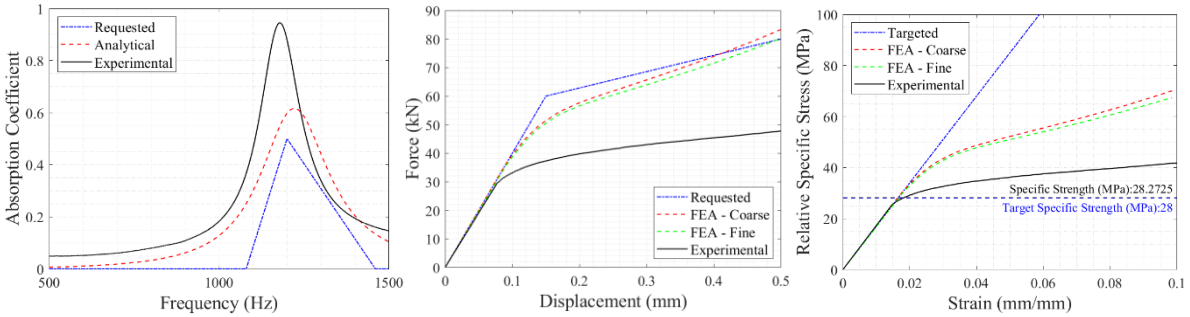
Example J



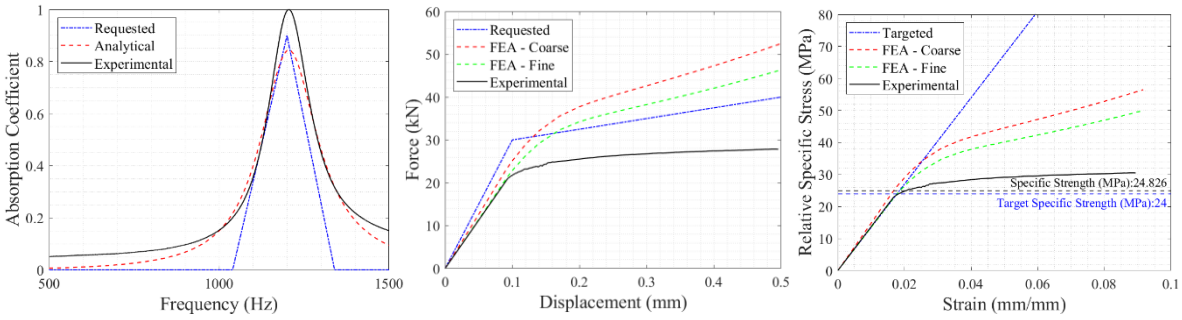
Example K



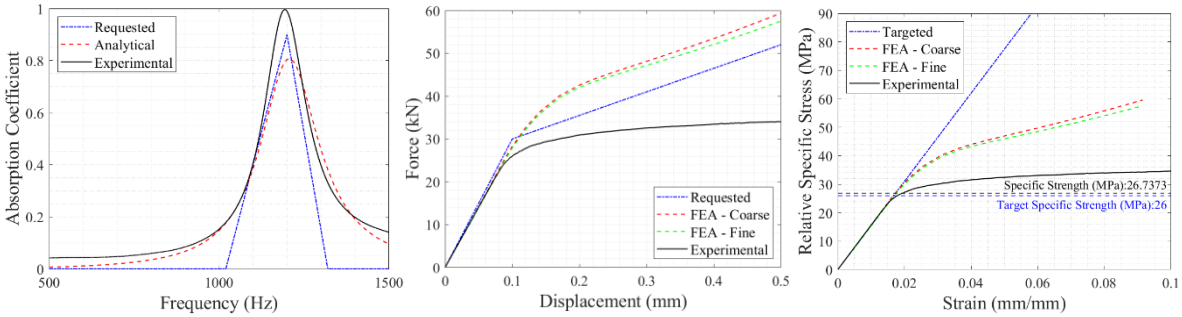
Example L



Example M



Example N



Example O

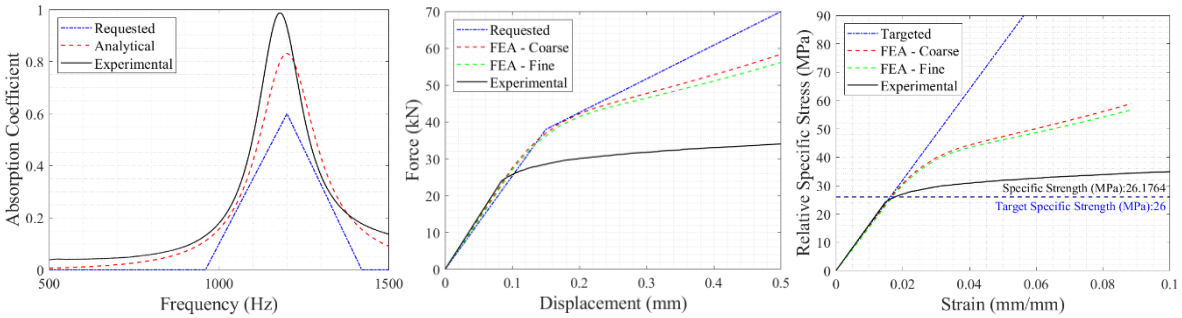


Figure 175 Sound absorption responses and mechanical characteristics of example metamaterials optimised for overall material thicknesses below 6mm

Additional acoustic tests were carried out to evaluate the best performing solutions for sound transmission loss. Transmission loss is defined as a logarithmic ratio of the incident and transmitted sound pressure levels, P_1 and P_2, as the sound wave passes through the specimen. It is quantified in decibels as $T_L = 20\log(P_2/P_1)$. The values were calculated from experimental measurements using the transfer matrix method and an alternative impedance tube setup, shown in Figure 176. 4 microphones were mounted into the tube at the separation distances shown in Table 35. The plunger was removed and an alternative tube assembly mounted to allow the sound wave to continue down the tube after passing through the specimen. Two microphones were located downstream and two upstream of the specimen. To

minimise reflections at the end of the tube an anechoic chamber was fitted filled with acoustic foam. Of the specimens example O gave the strongest performance in transmission loss, recording no less than 40dB between 225.2-2068.4Hz and 30dB between 193.5-2793.6Hz, as shown in Figure 177(a), at an overall thickness of 5.66mm. Example I offered the strongest low-frequency transmission loss at the lowest thickness, providing 28.9dB between 207.6-1633.5Hz at overall material thickness of 3.64mm. The measured response is shown in Figure 177(b). The inverse CNN approach is applicable to any problem where the target performance characteristic may be represented by a 2D response curve. In future works networks may be trained on datasets of transmission loss characteristics. The transmission loss may then be optimised alongside absorption coefficient in the same analysis.

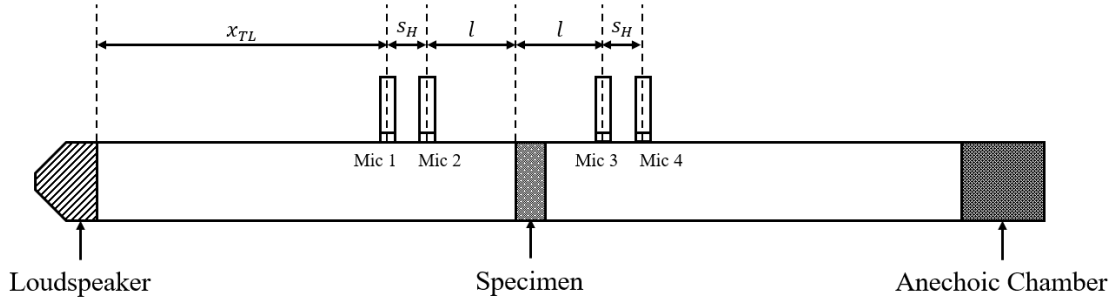


Figure 176 Impedance tube layout for measurement of transmission loss

Table 35 Geometry parameters for 4 microphone test setup to measure transmission loss

Geometry Parameter	Distance (mm)
l	140
x_{TL}	330
s_H	30

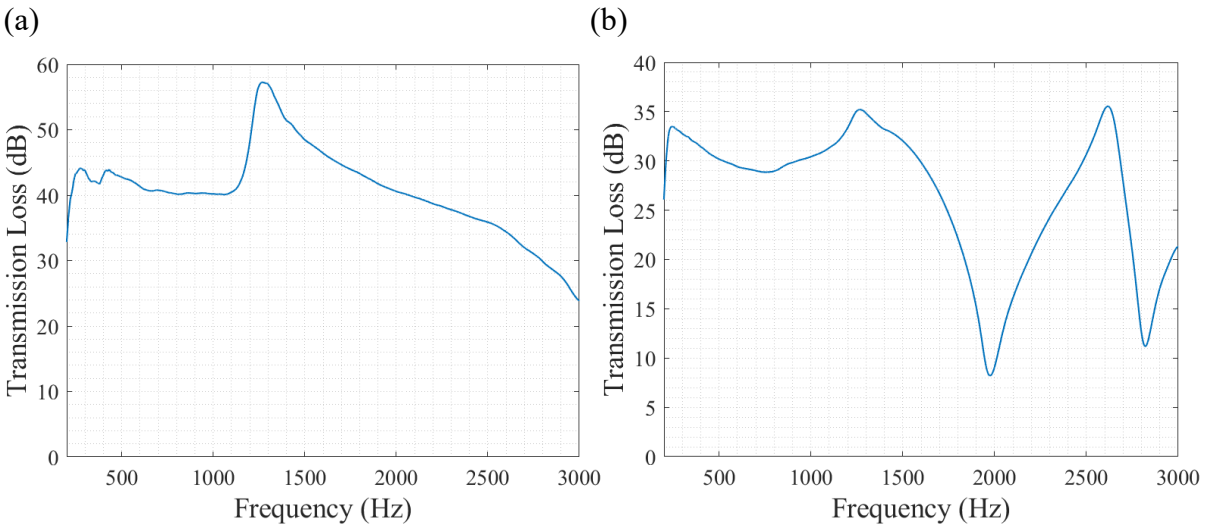


Figure 177 Experimental transmission loss measurements for reduced thickness model test specimens; (a) example O, t=1.23mm, w=3.46mm, h=4.66mm, overall thickness=5.66mm; (b) example I, t=0.93mm, w=3.71mm, h=2.64mm, overall thickness=3.64mm

5.4 Conclusion

This study presented a machine learning based approach for simultaneous optimisation of multiple competing performance objectives to facilitate new degrees of multifunctionality in design of architected metamaterials. A convolutional neural network was trained to predict the geometry parameters required to obtain desired sound absorption and mechanical properties of a parametric lattice material based on quarter-wavelength resonator labyrinths from datasets obtained by theory and finite element analysis. The network was able to output predictions of the necessary unit cell geometry for given input signals that combined the low-frequency absorption responses and compressive force-displacement characteristics to a level of accuracy that exceeded 99% regression when compared against known values. The model offers advantages in material performance and design flexibility over conventional alternatives. Requested material properties can be continuously updated to meet changing design criteria without repeating the analysis from scratch. Parameter limits may be updated and the network retrained to satisfy limits on overall material thickness according to the spatial limitations of a given application. Networks may be trained on experimental data for which there is no theoretical solution. Physical properties that have previously been seen as unrelated can be linked and considered in the same analysis, permitting new levels of multifunctionality and optimisation of multiple competing design criteria simultaneously. This ability was demonstrated by tailoring additively manufactured polymer unit cells for a range of relative mechanical properties targeted between 1050-1600MPa specific stiffness and 20-34MPa specific strength. All optimisations simultaneously provided tunable absorption peaks that exceeded 90% at frequencies below 400Hz. In addition the model could lock the absorption peak to a specific target frequency while fine-tuning the structural performance to achieve desired stiffness and strength at minimum mass. This allows the resonators to have the longest channel length for a given set of structural properties, which can be divided into multiple absorbers acting in parallel for greatest bandwidth. The network designed arrangements of parallel resonators that measured no less than 50% absorption between the ranges of 857.4-1256.3Hz and 982.5-1542.1Hz at an overall material thickness of 20.56mm. The process may be used to produce positive results when the size of the dataset is limited or the training simulations require coarse levels of fidelity to produce the quantity of data required. In these situations post validation of optimised solutions is of significant importance, however the network offers a large degree of control for fine-tuning responses that can be used to compensate for discrepancies between predicted properties and real-world measurements. The requirement for large quantities of training data remains a barrier that in future may require advancements in computing power or development of machine learning

networks that are able to carry out physics simulations autonomously to expedite dataset generation. Where necessary, training on a coarse dataset and post-compensation is recommended to ensure the process can be carried out within practical timeframes. This becomes of particular importance when multiple simulations are required to predict a range of performance targets. The inverse CNN showed an increase in training performance when adding additional objectives rather than a decrease, as the solution it is seeking becomes better defined. Further studies should be carried out to observe the ability of the model to process additional competing objectives such as impact resistance, bending characteristics, thermal conductivity, electrical properties and composite failure criteria. Reversing the model and working in the predictor direction to output material properties for a given geometry significantly reduces the complexity of the neural network. This may allow for additional objective functions and an increased number of geometry parameters. The existing geometry is locked to an initial parametric base configuration. Future works may allow the AI to set the layout of the unit cell depending on the requested input. The inverse CNN showed consistent predictive performance when the process was repeated on a revised unit cell with reduced thickness, providing evidence that the workflow is repeatable on problems of similar complexity. Increasing the dataset size and reducing the scope of the problem by restricting the limits of the dataset enhanced the generalisation of the network, leading to improved performance during validation checks. Additional unit cells were optimised to target relative structural properties between 900-1700MPa specific stiffness and 20-28MPa specific strength. The resulting materials are tuneable between 84.3-99.9% absorption at frequencies from 1200Hz and above at overall material thicknesses between 3.64-5.66mm. The AI driven approach allows for simultaneous optimisation of sound absorption and mechanical characteristics aimed at providing strongest absorption at lowest frequency and lowest overall material thickness while satisfying stiffness and strength criteria. This enables lightweight multifunctional metamaterials tailored for a given design purpose with potential for a wide range of applications across engineering industries.

Chapter 6 Multi-objective Optimisation of Lightweight Metamaterials for Impact Resistance by Machine Learning

In aircraft and motor vehicles lightweight structures that feature optimal specific stiffness and strength while maintaining high levels of crashworthiness are essential to prevent critical failure of components and mitigate injury to passengers in the event of high velocity impacts. During a collision the rate of change in momentum the vehicle exerts on the occupants before coming to rest is the product of force and time. Crash structures typically feature deformable regions engineered to crush to extend the period of time over which the impact takes place and reduce the forces acting on the passengers of the vehicle. This presents a performance trade-off between stiffness, strength and impact resistance. The following study aims to tune this trade-off using machine learning to optimise specific energy absorption under high velocity impacts while maximising static structural properties for improvements in lightweighting over pre-existing truss-based lattice structures. If components that are primarily structural in purpose can be lightened while also possessing increased levels of energy absorption, new generations of passenger vehicles may benefit from improved performance and crash safety as a result. The neural network optimisation workflow from previous chapters is extended to allow the AI to set the layout of the unit cell. Parametric geometry is designed based on an arrangement of nodes and struts, in which the node locations and strut radii are allowed to vary. Convolutional neural networks are trained to predict the relative specific stiffness and strength under quasi-static uniaxial compression from datasets of implicit FEA simulations. Additional networks are trained to predict the specific energy absorption of the same geometry from high velocity explicit dynamic simulations. The predictor networks are then combined using a multi-objective genetic algorithm to maximise the specific stiffness, strength and energy absorption simultaneously. Each solution along the resulting Pareto front of optimal solutions is compared against the equivalent octet-truss at the same relative density, one of the best performing pre-existing geometries found in literature, to assess the potential for use of machine learning to enhance mechanical properties and impact resistance for a new generation of AI synthesised metamaterials.

6.1 Parametric Unit Cell Configuration

The unit cell design featured 18 nodes. The first 14 nodes made up the outer frame and their locations remained fixed to ensure continuity with surrounding cells. The layout of the outer frame is shown in Figure 178(a), while Table 36 contains the coordinates of each node. Nodes 15-18 were internal nodes that were allowed to vary in position anywhere within the design space of a 20x20x20mm unit cell.

(a) (b)

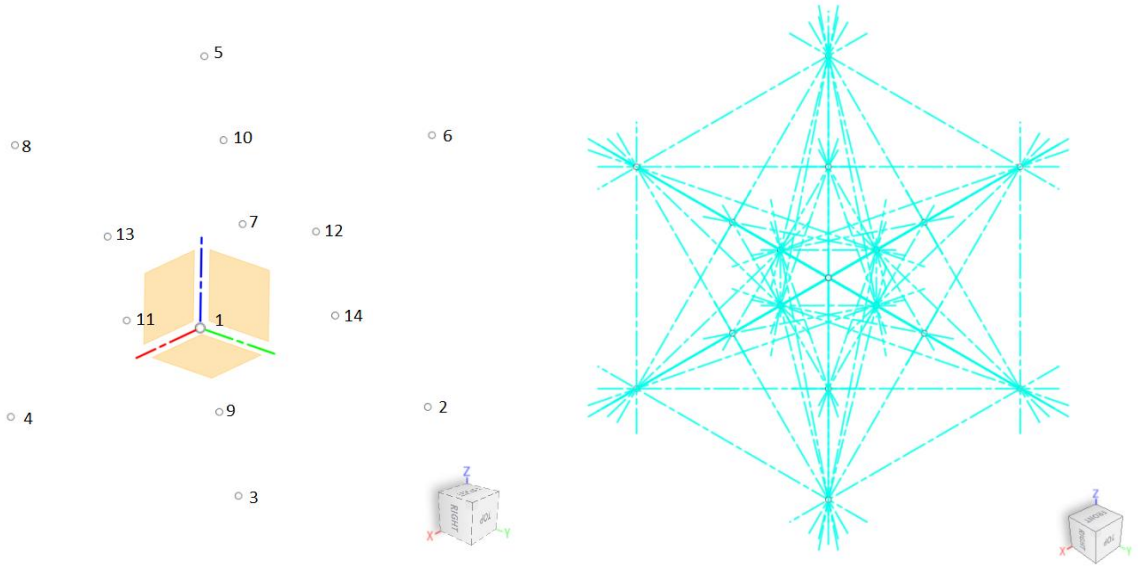


Figure 178 Parametric unit cell layout; (a) outer frame nodes with fixed positions; (b) strut connections between nodes at arbitrary default positions

Table 36 Node location matrix

	n1	n2	n3	n4	n5	n6	n7	n8	n9	n10	n11	n12	n13	n14	n15	n16	n17	n18
X (mm)	0	0	20	20	0	0	20	20	10	10	20	0	10	10	n15X	n16X	n17X	n18X
Y (mm)	0	20	20	0	0	20	20	0	10	10	10	10	0	20	n15Y	n16Y	n17Y	n18Y
Z (mm)	0	0	0	0	20	20	20	20	0	20	10	10	10	10	n15Z	n16Z	n17Z	n18Z

Fixed
CNN Parameter

Table 37 shows the connections between every node. Nodes 1-8 were connected by struts to form a cubic outer frame reinforced by cross braces joining nodes 9-14. Each internal node was then connected to every other node in the unit cell. Figure 178(b) illustrates the placement of struts between nodes, with the internal nodes at arbitrary default positions.

Table 37 Node connectivity matrix of strut radius parameters

	n1	n2	n3	n4	n5	n6	n7	n8	n9	n10	n11	n12	n13	n14	n15	n16	n17	n18
n1	-	r1.2	r1.3	r1.4	r1.5	r1.6	r1.7	r1.8	r1.9	r1.10	r1.11	r1.12	r1.13	r1.14	r1.15	r1.16	r1.17	r1.18
n2	r2.1	-	r2.3	r2.4	r2.5	r2.6	r2.7	r2.8	r2.9	r2.10	r2.11	r2.12	r2.13	r2.14	r2.15	r2.16	r2.17	r2.18
n3	0	r3.2	-	r3.4	r3.5	r3.6	r3.7	r3.8	r3.9	r3.10	r3.11	r3.12	r3.13	r3.14	r3.15	r3.16	r3.17	r3.18
n4	r4.1	0	r4.3	-	r4.5	r4.6	r4.7	r4.8	r4.9	r4.10	r4.11	r4.12	r4.13	r4.14	r4.15	r4.16	r4.17	r4.18
n5	r5.1	0	0	0	-	r5.6	r5.7	r5.8	r5.9	r5.10	r5.11	r5.12	r5.13	r5.14	r5.15	r5.16	r5.17	r5.18
n6	0	r6.2	0	0	r6.5	-	r6.7	r6.8	r6.9	r6.10	r6.11	r6.12	r6.13	r6.14	r6.15	r6.16	r6.17	r6.18
n7	0	0	r7.3	0	0	r7.6	-	r7.8	r7.9	r7.10	r7.11	r7.12	r7.13	r7.14	r7.15	r7.16	r7.17	r7.18
n8	0	0	0	r8.4	r8.5	0	r8.7	-	r8.9	r8.10	r8.11	r8.12	r8.13	r8.14	r8.15	r8.16	r8.17	r8.18
n9	r9.1	r9.2	r9.3	r9.4	0	0	0	0	-	r9.10	r9.11	r9.12	r9.13	r9.14	r9.15	r9.16	r9.17	r9.18
n10	0	0	0	0	r10.5	r10.6	r10.7	r10.8	0	-	r10.11	r10.12	r10.13	r10.14	r10.15	r10.16	r10.17	r10.18
n11	0	0	r11.3	r11.4	0	0	r11.7	r11.8	0	0	-	r11.12	r11.13	r11.14	r11.15	r11.16	r11.17	r11.18
n12	r12.1	r12.2	0	0	r12.5	r12.6	0	0	0	0	0	-	r12.13	r12.14	r12.15	r12.16	r12.17	r12.18
n13	r13.1	0	0	r13.4	r13.5	0	0	r13.8	0	0	0	0	-	r13.14	r13.15	r13.16	r13.17	r13.18
n14	0	r14.2	r14.3	0	0	r14.6	r14.7	0	0	0	0	0	0	-	r14.15	r14.16	r14.17	r14.18
n15	r15.1	r15.2	r15.3	r15.4	r15.5	r15.6	r15.7	r15.8	r15.9	r15.10	r15.11	r15.12	r15.13	r15.14	-	r15.16	r15.17	r15.18
n16	r16.1	r16.2	r16.3	r16.4	r16.5	r16.6	r16.7	r16.8	r16.9	r16.10	r16.11	r16.12	r16.13	r16.14	r16.15	-	r16.17	r16.18
n17	r17.1	r17.2	r17.3	r17.4	r17.5	r17.6	r17.7	r17.8	r17.9	r17.10	r17.11	r17.12	r17.13	r17.14	r17.15	r17.16	-	r17.18
n18	r18.1	r18.2	r18.3	r18.4	r18.5	r18.6	r18.7	r18.8	r18.9	r18.10	r18.11	r18.12	r18.13	r18.14	r18.15	r18.16	r18.17	-

Locked To Zero
Group R1 (Frame)
Group R2 (Crosses)
Group R3 (Internal to external nodes)
Group R4 (Internal to internal nodes)
Duplicates

The geometry parameters used as input for the neural network were the x, y and z coordinates of the 4 internal nodes, in addition to key strut radii grouped as follows:

- R1: Outer frame (12 struts)
- R2: Cross braces (24 struts)
- R3: Links between internal nodes and external nodes (56 struts)
- R4: Links between internal nodes and other internal nodes (6 struts)

To calculate values of specific stiffness, strength and energy absorption that remain applicable to larger lattices featuring combinations of any number of tessellated unit cells, a periodic boundary condition assumption was applied when carrying out FEA simulations. On its own a single unit cell of the proposed geometry is likely to be asymmetrical. To ensure validity of this assumption the geometry was formed into a representative 2x2 symmetric supercell of dimensions 40x40x20mm by mirroring a single cell in both the XZ and YZ planes. The struts of the resulting geometry that were assigned to each radius parameter are identified in Figure 179.

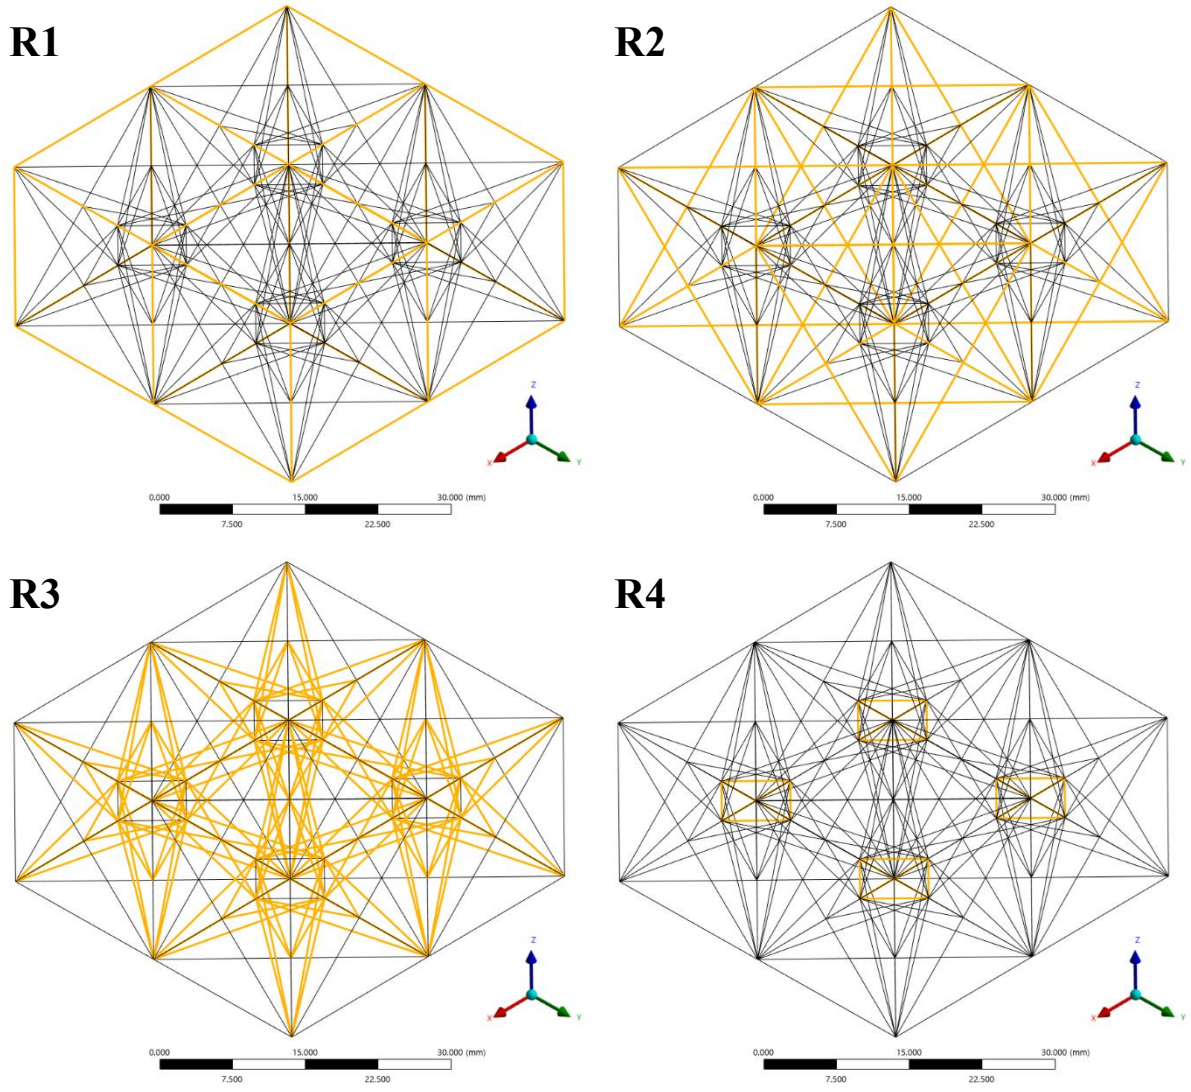


Figure 179 Strut radius geometry parameters assigned to 2x2 symmetric supercell

The CAD geometry was imported into ANSYS workbench using a 2 way associative interface to allow for automatic updates of the geometry parameters during each FEA run. The 4 strut radius parameters along with the 12 location parameters gave 16 total inputs for the neural network. The node locations took values between 0-20mm. Strut radii were limited between 0.5-3mm to ensure the resulting lattice structures would be manufacturable via FDM.

6.2 Dataset Generation

6.2.1 Static FEA Model

An implicit non-linear FEA model was created in ANSYS mechanical for measurement of static structural responses of the metamaterial supercell under uniaxial compression. The parametric geometry was represented as a line body model with user defined cross sections to set the radius of each strut. The model was meshed with beam elements of size 1.5mm. The material properties for FDM Polymaker CoPA nylon recorded in section 4.1 were assigned to the elements of the lattice using a multilinear isotropic hardening model to simulate post yield plastic deformation. The non-linear data was recorded on an Instron 3369 with 50kN load cell under standard ASTM D695 [156] from the cylindrical compression coupons 12.7mm in diameter and 25.4mm in height, built at 90° to the direction of the loading. The specimens were loaded at 1.3mm/min with strain measurements taken by an LVDT. A total of 5 specimens were tested and a mean taken of the results. The material had mean Young's modulus of 1705.92MPa, yield strength 63.21MPa, Poisson ratio of 0.35 and density of 1120kg/m^3. The non-linear material data was converted to a relation of true stress vs logarithmic strain and the plastic strain was separated from the elastic strain for use in the multilinear hardening model using the formulae:

$$\sigma = \sigma_{eng}(1 + \varepsilon_{eng}), \varepsilon = \ln(1 + \varepsilon_{eng}), \varepsilon_{plastic} = \varepsilon - \sigma/E \tag{6.1}$$

The non-linear data input into ANSYS featured 100 datapoints and is shown in Figure 180.

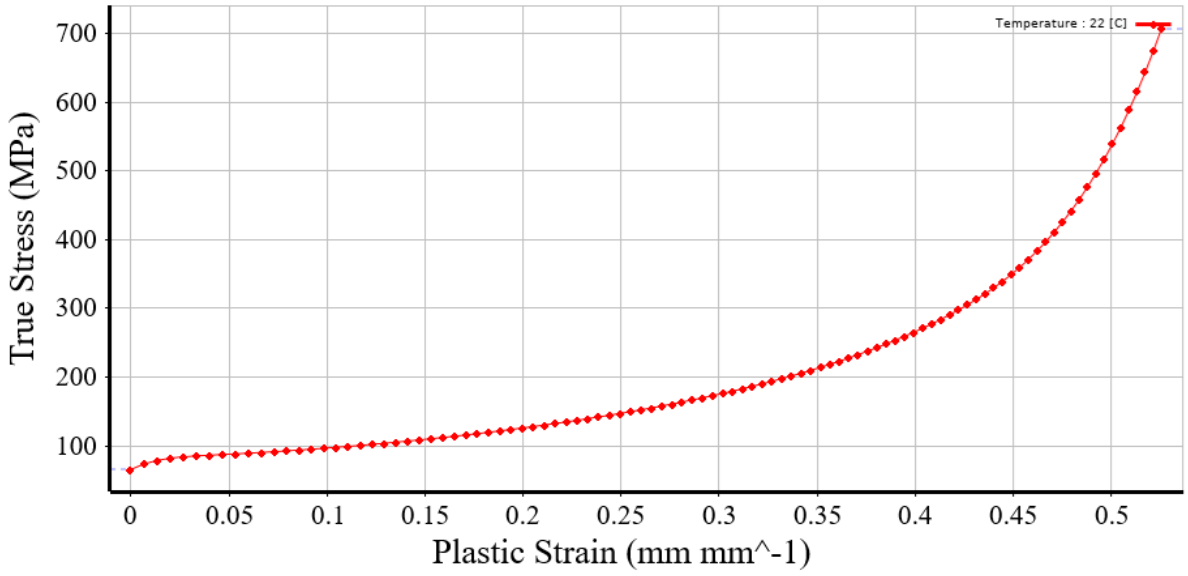


Figure 180 Polymaker CoPA nylon non-linear material data used in static ANSYS model

The beam elements at the outer limits of the supercell were assigned semi-circular cross sections to represent cut faces of the struts at the periodic boundary condition, as shown in Figure 181. A displacement load of 1.5mm was applied to the upper surface of the supercell in the negative z-direction, ramped on linearly over two load steps. A fully fixed support was applied to the bottom surface.

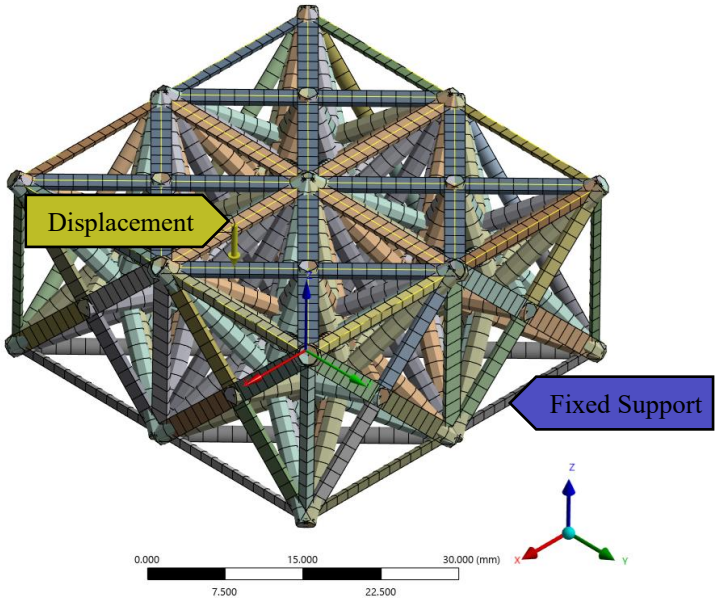


Figure 181 Static FEA model mesh and boundary conditions. Example shown has arbritary default positions and all strut radii set to 1mm

The solver type was set to Newton-Raphson with automatic time stepping enabled. The initial number of substeps was 30 and the solver was allowed to adjust the number of timesteps automatically between a minimum of 10 and maximum of 100 until convergence. Large deflections were enabled in the solver options to allow for modelling of post-yield behaviour using a lookup table populated from the plastic material data. To aid convergence energy based non-linear stabilisation was enabled with energy dissipation ratio of 0.01. The reaction force was probed at the fixed support boundary condition at intervals of every 0.03mm between 0-1.5mm displacement load to populate a force-displacement response of 51 datapoints. Each response was converted to stress-strain using the area and original height of the supercell. The Young's modulus was then extracted from the linear region using an automated MATLAB script and the yield stress was taken as the 0.2% proof stress. For every design point in the dataset the relative density was calculated by dividing the density of the supercell by the density of the solid material. Relative specific stiffness and strength were then calculated to allow the material properties to be evaluated irrespective of the number of unit cells in the lattice by dividing the Young's modulus and yield stress by the relative density. These values were then used as outputs for training the neural networks.

6.2.2 High Velocity Dynamic FEA Model

To measure the specific energy absorption of each supercell a high velocity explicit dynamics simulation was developed in ANSYS using the AUTODYN solver. The same material data used for the static analysis was adapted to use a bilinear isotropic hardening law by approximating the previous non-linear material data using a constant tangent modulus of 191.75MPa. The value was extracted by a least-squares linear regression of the raw engineering stress-strain data in the plastic region, as shown in Figure 182.

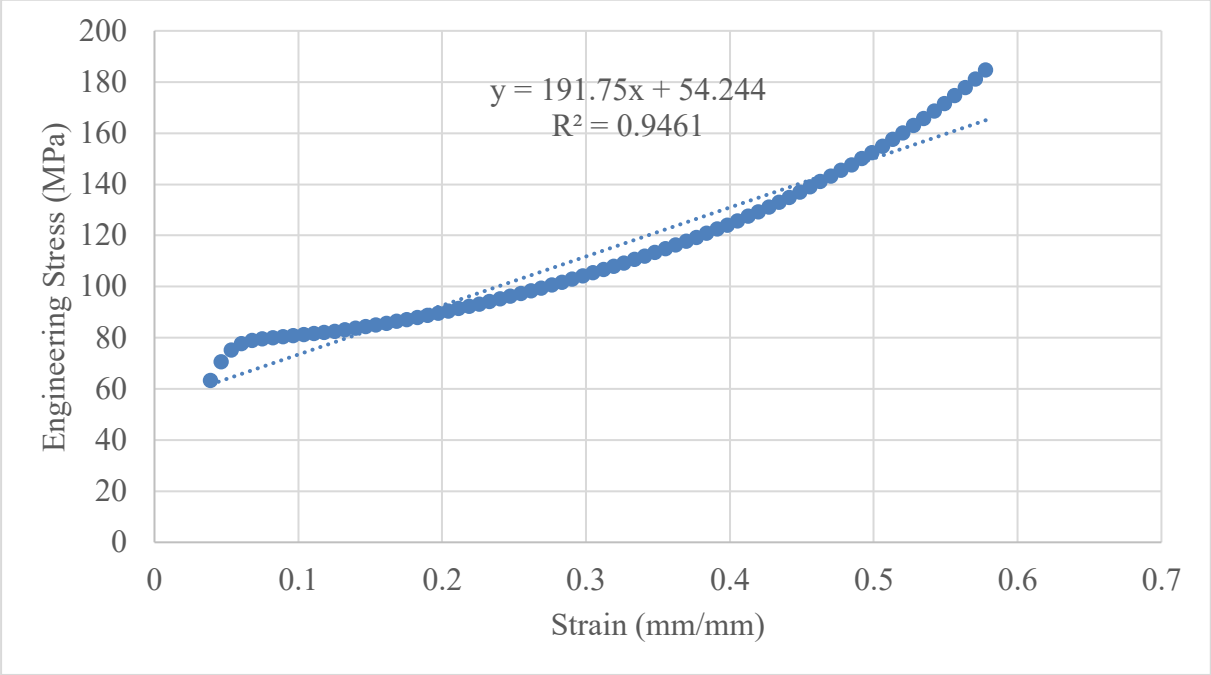


Figure 182 Tangent modulus extraction by least squares regression

The dynamic simulations represented the lattice supercell using a solid body CAD geometry meshed with solid elements of size 1.5mm. Initial simulations were carried out applying a displacement load directly to the top surface of the supercell while the bottom remained fixed. This resulted in unrealistic deformation as the interior of the lattice occasionally passed outside the cell limits, as shown in Figure 185. To prevent this, a rigid steel impactor of size 60x60x2mm was modelled and meshed with solid elements of size 2mm to simulate interactions between the struts and the external projectile. The steel was assigned Young's modulus of 200GPa, yield strength 250MPa, Poisson ratio of 0.3 and density of 7850kg/m^3. The impactor was placed with an initial separation distance of 0.15mm from the top face of the lattice. A displacement load of 10mm was applied to the impactor body (Figure 183), ramped on linearly over a period of 0.1 milliseconds to simulate a constant velocity of 100m/s. A fully fixed support was applied to the bottom face of the lattice supercell, as displayed in Figure 184.

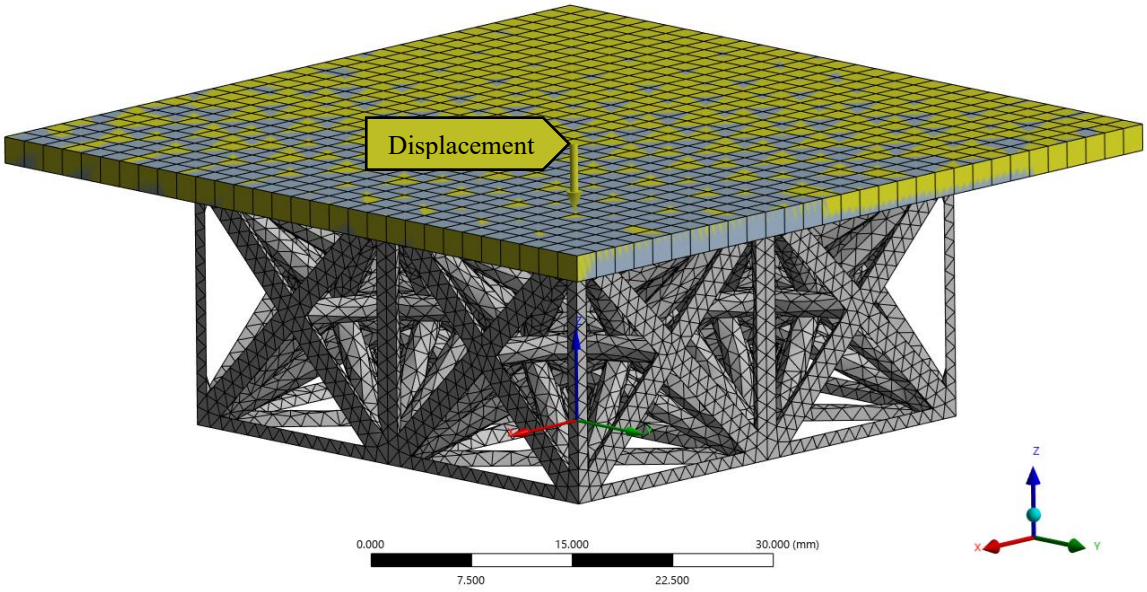


Figure 183 High velocity explicit dynamic FEA model mesh and loading. Example shown has arbritary default positions and all strut radii set to 1mm.

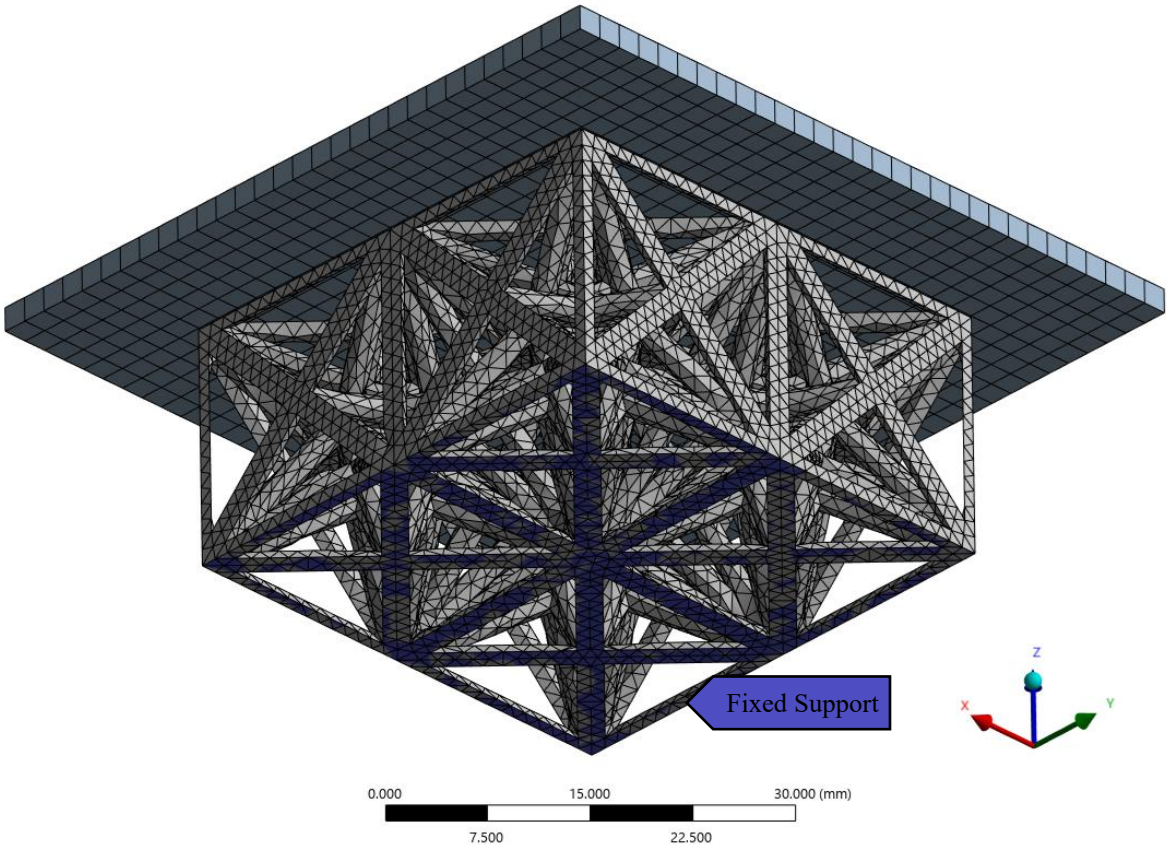


Figure 184 High velocity explicit dynamic FEA model mesh and fixed support. Example shown has arbritary default positions and all strut radii set to 1mm.

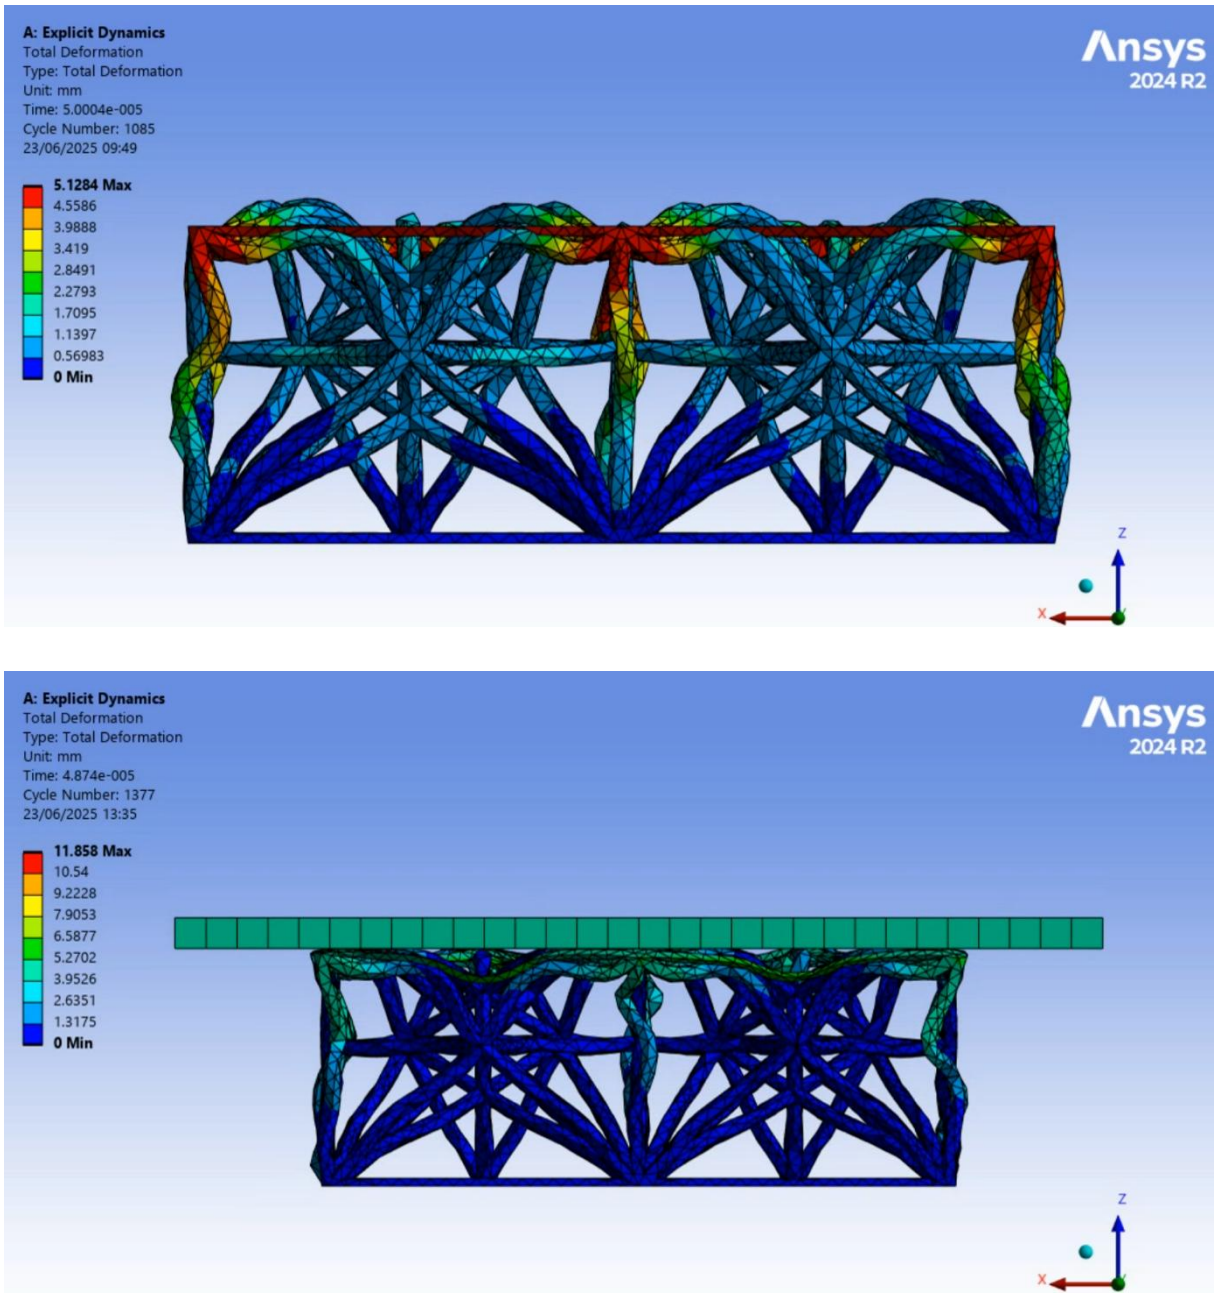


Figure 185 Comparison of high velocity dynamic impact simulation without (top) and with (bottom) steel impactor modelled. Examples use the same geometry parameters with arbitrary default positions and all strut radii set at 0.5mm

Frictional contacts were modelled between the impactor and lattice with static and dynamic coefficients of friction of 0.2 and 0.1 respectively. Self-intersections of the lattice as the struts interact with each other were modelled using a frictional body interaction with trajectory contact detection and penalty based formulation. The end time of the simulation was set at 0.1 milliseconds. Flanagan-Belytschko hourglass damping was used with stiffness coefficient equal to 0.1. Element erosion was enabled to model damage of the lattice. Element deletion was set to occur either on reaching a geometric strain limit of 1.5 or on initiation of material failure at a strain of 0.53. The reaction force was probed at the fixed support boundary

condition every 2 microseconds to build up a force-displacement response for each simulation of 51 datapoints at displacement intervals of 0.2mm between 0-10mm displacement of the impactor. The response was converted to a stress-strain characteristic. The energy density, also known as energy absorbed per unit volume was then calculated by integration of the area under the stress-strain curve. The specific energy absorption (SEA) was calculated for each simulation using the following formulae [158]:

$$SEA = \frac{\Delta W}{m} \tag{6.2}$$

where ΔW is work done, or total energy absorbed and m is the mass of the supercell given by

$$m = \rho_s \bar{\rho} V_{sc} \tag{6.3}$$

where ρ_s is the density of the solid, $\bar{\rho}$ is the relative density of the lattice and V_{sc} is the volume of the supercell.

The energy absorbed per unit mass, or specific energy absorption is given by

$$SEA = \frac{\Delta W}{\rho_s \bar{\rho} V_{sc}} = \frac{u}{\rho_s \bar{\rho}} \tag{6.4}$$

where u is the energy absorbed per unit volume or energy density.

To populate the dataset a design of experiments test plan consisting of 6000 geometry combinations was generated using Latin hypercube sampling. Each design point was first run through the static FEA model, in batches of 1000 observations. The static simulations were run on computer hardware with an Intel i7-10700K processor at 3.8GHz with 32GB RAM and an Nvidia RTX 4070 GPU. The mean solve time for the static model was 4min 2sec per run, which yielded a convergence rate of 76.1% with 4568 returned observations. The converged design points were then sent through the explicit dynamic model. Of these 70.6% were returned, largely due to meshing errors caused by complex curvatures where the faces of the struts were sliced at the outer boundary of the supercell. The dynamic model was run on an Intel i9-14900K processor at 3.2GHz with 64GB RAM and an Nvidia RTX 5070 GPU, taking on average 2min 24sec for each solve. The final dataset consisted of 3226 observations and provided coverage of the design space comparable to datasets generated in previous chapters. The relative specific stiffness and strength took values between the ranges of 161.51-1049.74MPa and 2.46-47.76MPa respectively. The range of specific energy absorption found in the dataset was 3.20-45.05kJ/kg, while the masses of the supercell were between 6.1-42.5grams.

6.3 Convolutional Neural Network Training

The geometry parameters for each design point within the dataset were formatted into a 3226x1 cell array. Each cell contained a 1x16 row vector of the node locations and strut radii to be supplied as sequence input to a convolutional neural network. The CNN architecture from section 5.1.3 was modified to provide a single regression output value. The 8th fully connected layer was changed from 3 to 4 neurons connected to an additional final 9th fully connected layer with a single neuron using batchwise normalisation and ReLU activation. The remaining layers and activations remained consistent with those displayed in Figure 143 and Table 22. As the network now worked in the predictor direction, a separate CNN was trained for each performance objective to maximise the prediction accuracy. The predictions could then be called individually by the genetic algorithm. All input and output parameters were normalised between 0 and 1 before training, scaled according to the minimum and maximum values for each parameter in the original dataset. The parameter ranges were then saved along with each trained network so any predictions could be rescaled back again as necessary. The normalised data was partitioned with 70% for training, 15% for validation and 15% for testing. Each network was trained on the same hardware used to run the explicit dynamic simulations, with GPU acceleration enabled on an Nvidia RTX 5070. For every case the Adam optimiser was used to minimise the Root Mean Square Error (RMSE) between the network predictions and known values from the dataset, with an initial learning rate of 0.001. The training was run over 1000 epochs and the data was passed through the network in mini batches of size 100. Separate networks were trained to predict the relative specific stiffness, relative specific strength, specific energy absorption and the mass of the lattice supercell. The training histories of each network are shown in Figure 186. In all cases a sharp drop in RMSE and loss is observed at the start of training as the network initialises. Gaps remain between the training and validation error over approximately the first 200 epochs, likely due to limitations on the size of the dataset. These gaps close towards the end of the training, indicating improvements in the network predictions as the training progressed. No evidence of overfitting is observed, suggesting the final networks were sufficiently well generalised to offer strong performance on unseen data. This assertion was supported during testing of the trained network, where the correlation between the predictions and known values of the network outputs showed coefficients of regression that exceeded 96% across all parameters. Figure 187 displays the regression between the network predictions and known values from the test dataset partition of 484 observations, while a summary of the training performance is provided in Table 38.

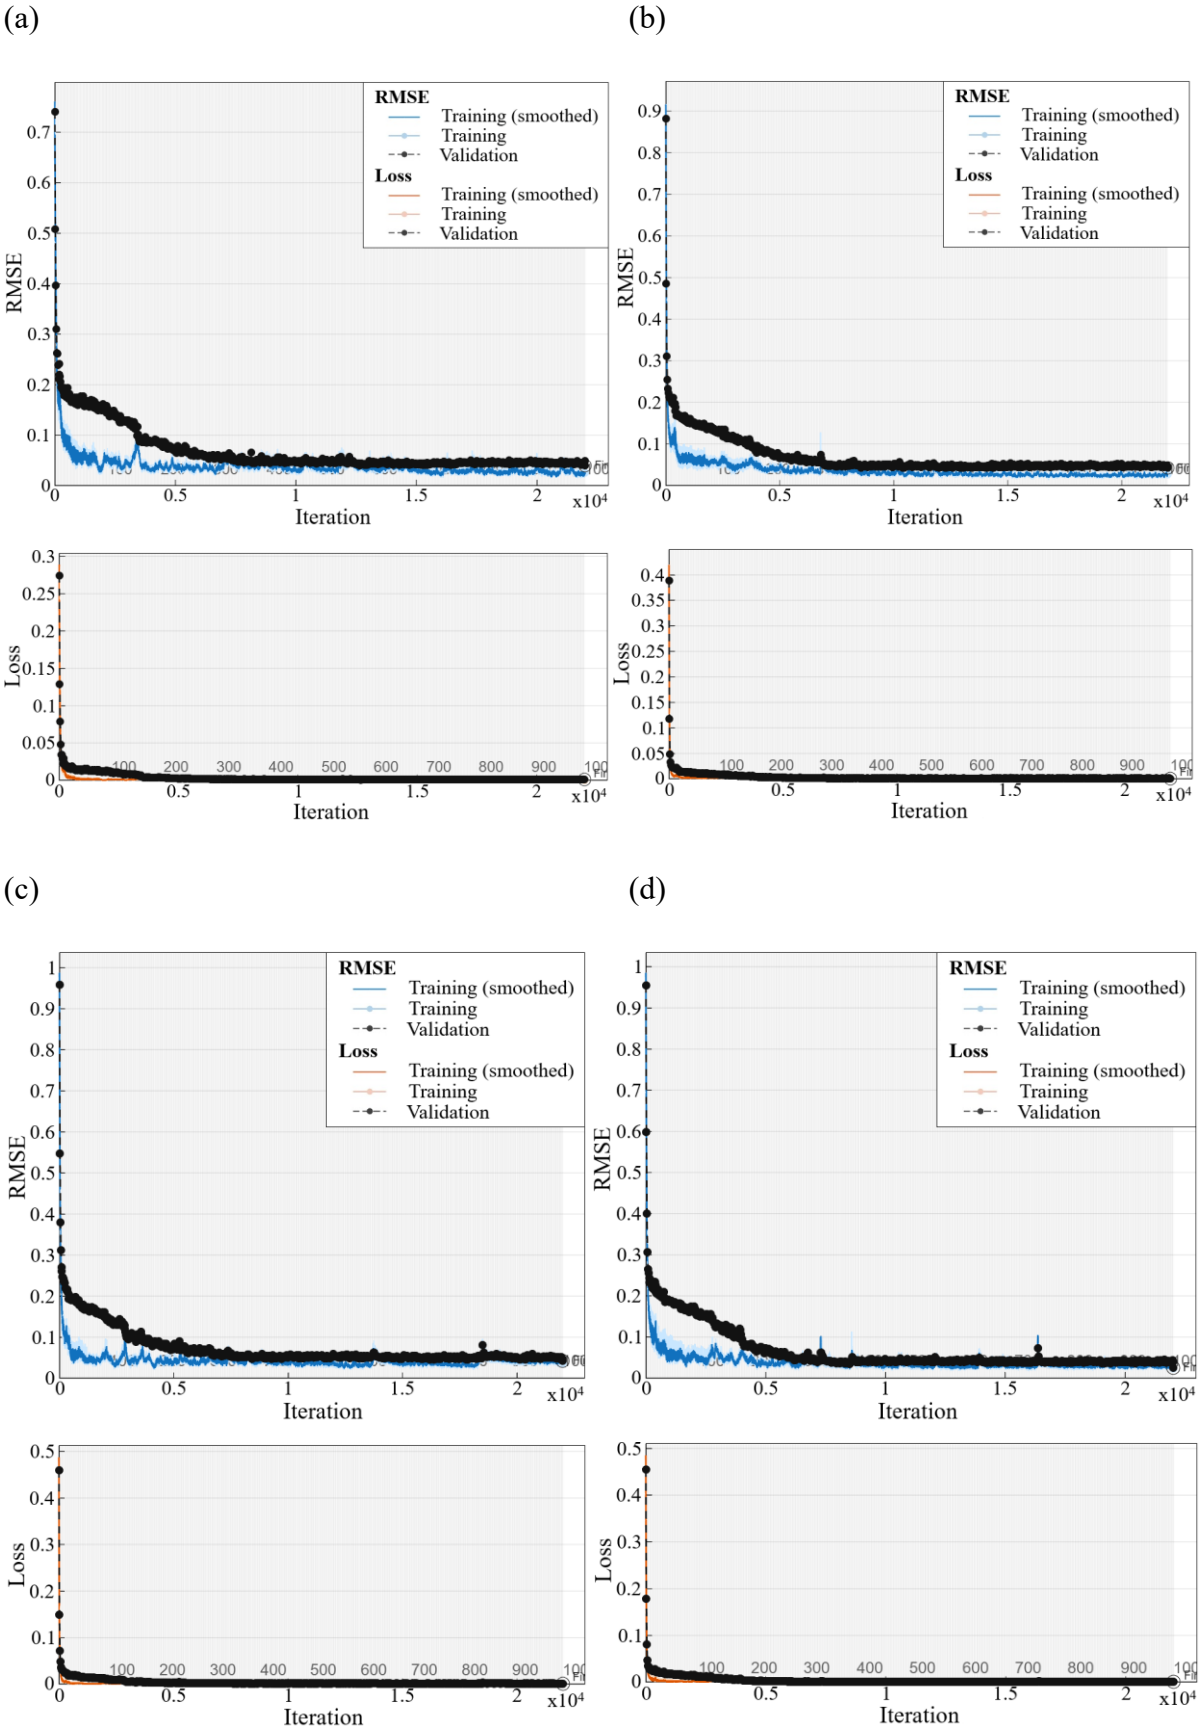


Figure 186 Training histories for CNNs to predict (a) relative specific stiffness, (b) relative specific strength, (c) specific energy absorption, (d) Mass

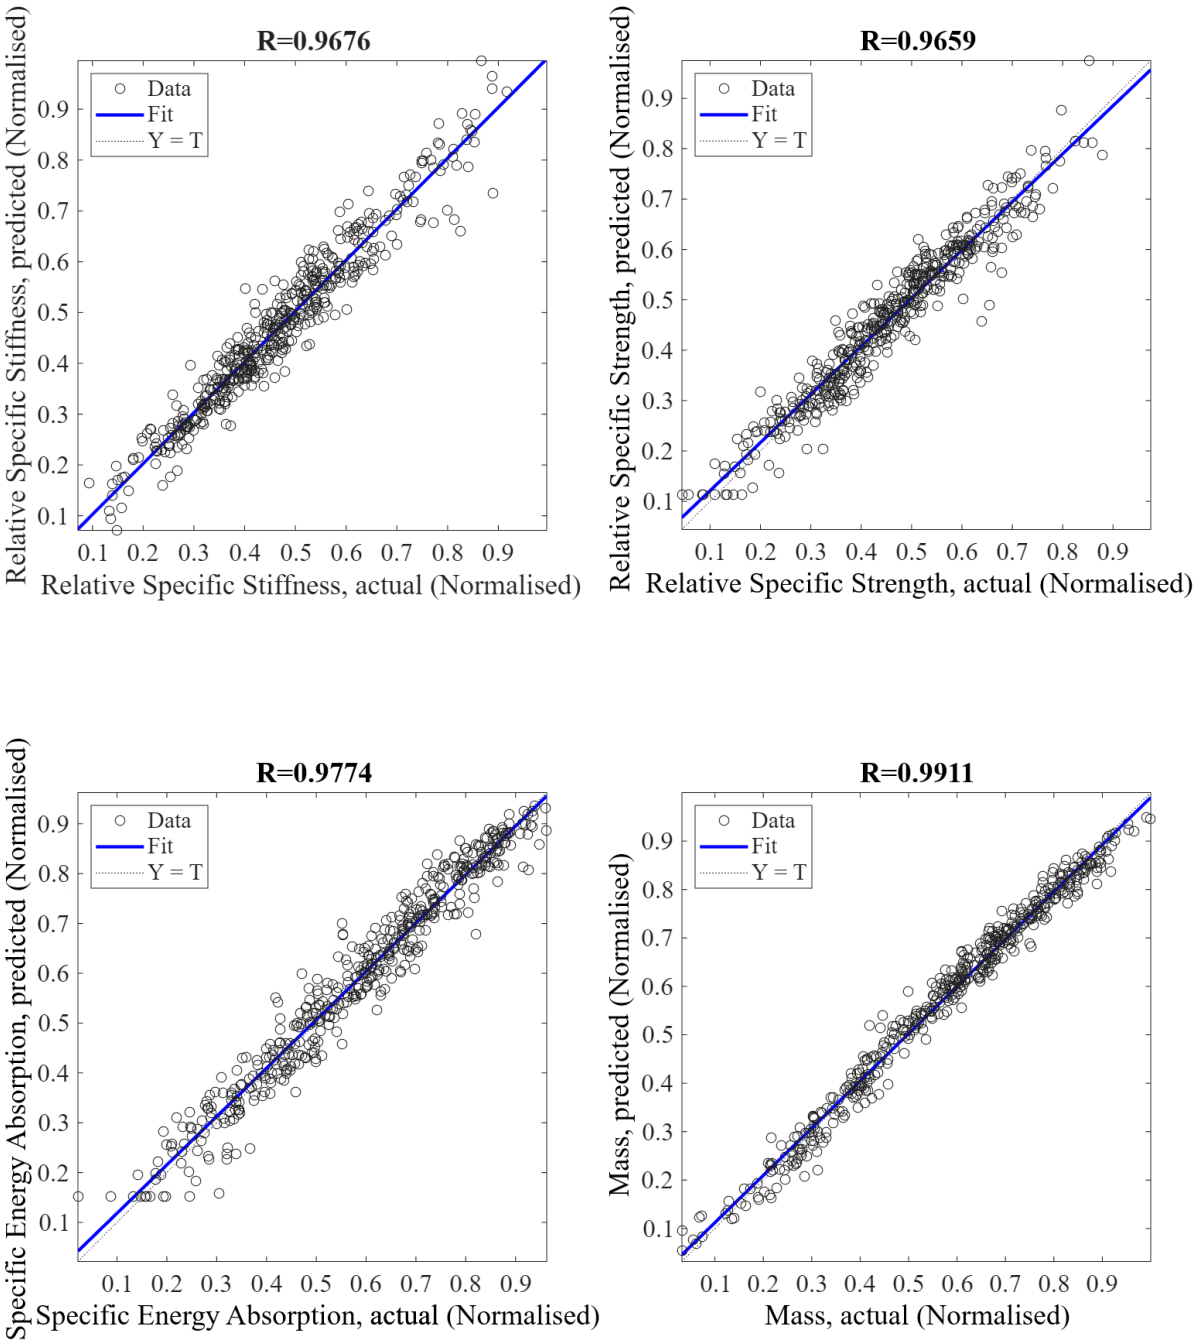


Figure 187 Regression between convolutional neural network predictions and known values for test dataset of 484 observations

Table 38 Summary of training performance and test regression for predictor CNNs

Output	Architecture		Training Options		Training Performance			Regression	
Relative Specific Stiffness	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0253	R Value	0.9676
	Convolution Filter Size	3	Mini Batch Size	100		Validation	0.0488		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.0003		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.0012		
	Activation	ReLU	Validation Frequency	22	Training Time	22min 18sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	22000			
Relative Specific Strength	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0263	R Value	0.9659
	Convolution Filter Size	3	Mini Batch Size	100		Validation	0.0461		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.0003		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.0011		
	Activation	ReLU	Validation Frequency	22	Training Time	24min 21sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	22000			
Specific Energy Absorption	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0361	R Value	0.9774
	Convolution Filter Size	3	Mini Batch Size	100		Validation	0.0474		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.0007		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.0011		
	Activation	ReLU	Validation Frequency	22	Training Time	24min 4sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	22000			
Mass	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0236	R Value	0.9911
	Convolution Filter Size	3	Mini Batch Size	100		Validation	0.0405		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.0003		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.0008		
	Activation	ReLU	Validation Frequency	22	Training Time	23min 55sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	22000			

6.4 Multi-objective Genetic Algorithm Optimisation

A multi-objective genetic algorithm (GA) was developed using MATLAB's optimisation toolbox. The objective function was set using a custom script to call each of the trained networks to provide predictions for each objective value across the entire population at every generation of the algorithm. The objectives of the script were set using combinations of the following:

1. Maximise relative specific stiffness
2. Maximise relative specific strength
3. Maximise specific energy absorption
4. Minimise mass

The algorithm provided 16 outputs corresponding to the geometry parameters of the lattice supercell. Constraints were placed on the algorithm to limit the values of the node location coordinates to between 0-20mm and the strut radii between 0.5-3mm, in agreement with the limits of the original training dataset. The GA generated a population of 200 geometry designs at each generation. The initial population was created using a quasi-random Sobol sequence. At every generation 10 elite members with the highest fitness values were guaranteed survival to the next generation. The remaining population consisted of 80% crossover children and 20% random mutations. Crossover parents were chosen by tournament selection according to their fitness values. A scattered crossover function was used to select genes from each parent. The optimiser had no time limit and there was no limit on the fitness values. The algorithm was allowed to continue until the average change in the spread of the Pareto solutions fell below a function tolerance of $1x10^{-6}$ or until a maximum of 1600 generations or 50 stall generations, at which point the stopping criteria would be triggered. For each optimisation the GA output a Pareto front of optimal solutions. The values of the solution geometry parameters were then run back through the static and dynamic FEA models to measure the properties of the lattice at each solution design point.

6.4.1 Structural Properties Optimisation

Initially the genetic algorithm was run with 2 objectives to maximise the relative specific stiffness and strength in order to seek solutions with the greatest overall structural properties. The algorithm took 1hr 25min 18sec to run. Stopping was triggered by the average change in the Pareto spread falling below the tolerance value. The mean Pareto spread and distance between individuals in the population at each generation are shown in Figure 188.

Fluctuations are observed in the Pareto spread before settling on a final value of 0.167 after 243 generations. The distance between individuals drops sharply over the first 10 generations before levelling out as the algorithm approaches convergence. The final generation had mean distance of 1.372.

(a) (b)

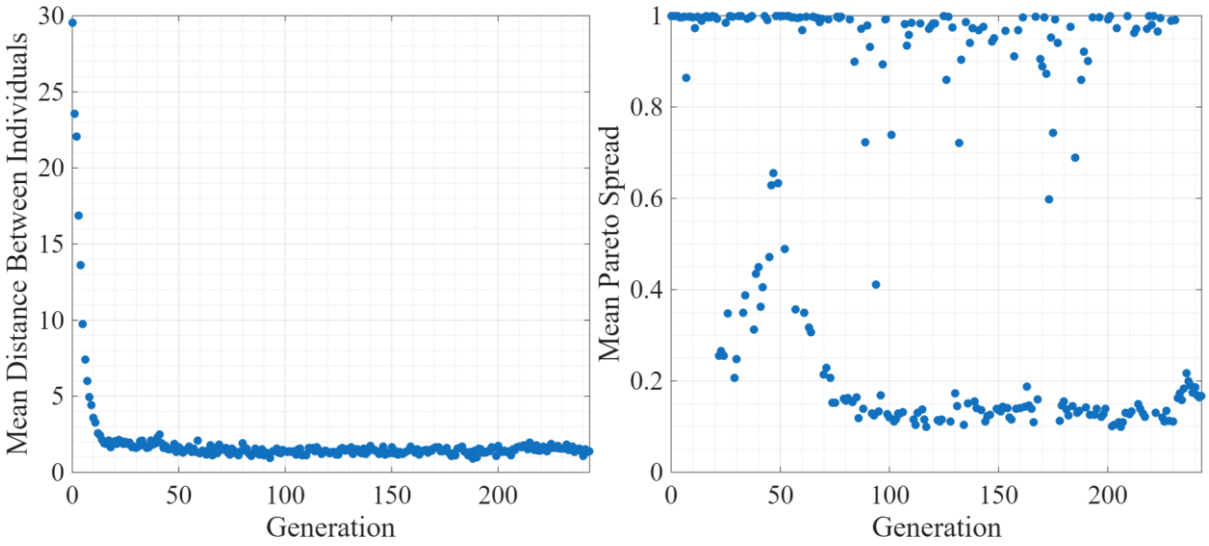


Figure 188 Optimiser history of multi-objective genetic algorithm set to maximise relative specific stiffness and strength; (a) mean distance between individuals; (b) mean Pareto spread

The Pareto front of solutions returned by the GA are displayed in Figure 189. The geometry parameters for each solution can be found in Appendix M. Each design point along the front represents an optimal solution for the multi-objective analysis, in which one performance objective cannot be further improved without reducing another. The GA provided 70 solutions with relative specific stiffness and strength between 1251.82-1305.75MPa and 57.75-62.82MPa respectively. All solutions therefore gave structural performance greater than the maximum values found in the original training dataset of 161.51-1049.74MPa and 2.46-47.76MPa. The stiffest and strongest solutions were 24.4% stiffer and 31.5% stronger than the greatest values in the dataset, which was equivalent to an exhaustive search of 3226 possible designs.

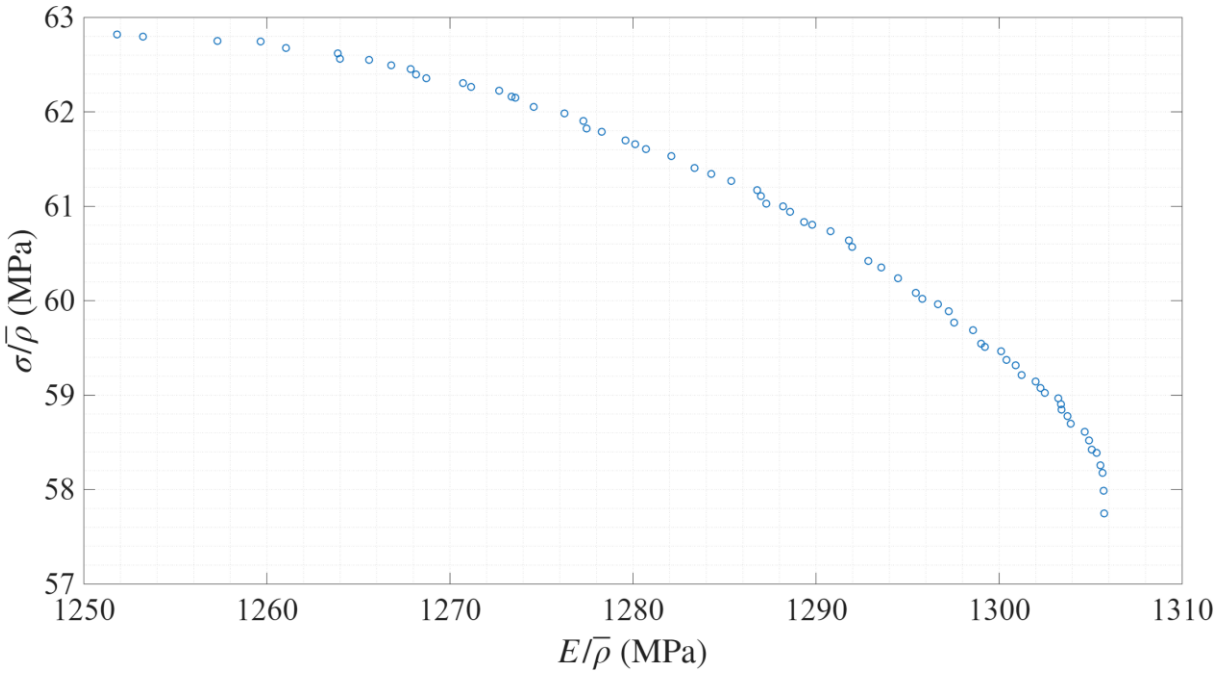


Figure 189 Pareto front of solutions from GA set to optimise relative specific stiffness and strength simultaneously

For each solution the mechanical performance was compared against the equivalent fixed radius octet-truss lattice at the same relative density. Equivalent FEA models to those outlined in section 6.2 were recreated substituting in the 40x40x20mm octet-truss supercell geometry shown in Figure 190, formed from a 2x2 arrangement of 20x20x20mm unit cells. All struts in the octet-truss unit cell had the same radius and the relative density was matched to the value of the AI unit cell by changing the radius value. In all cases the base material properties, mesh settings and loading conditions were kept identical to those used previously. Figure 191 plots comparisons of the stiffness and strength of the AI generated Pareto front solutions against the equivalent octet-truss. In each case the Young's modulus and yield strength of the supercell were normalised by dividing by the stiffness or strength of the parent solid base material. Figure 192 shows the percentage increase in performance. All designs optimised by combination of machine learning and genetic algorithm outperformed the octet-truss equivalent, increasing normalised stiffness by between 101.6-123.7% and strength by 63.1-88.8% for relative densities between 73.8-78.9%.

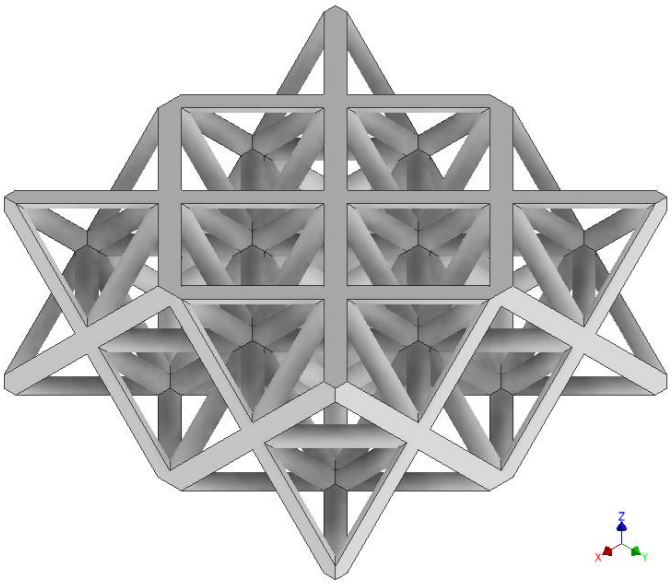


Figure 190 Octet-truss supercell geometry

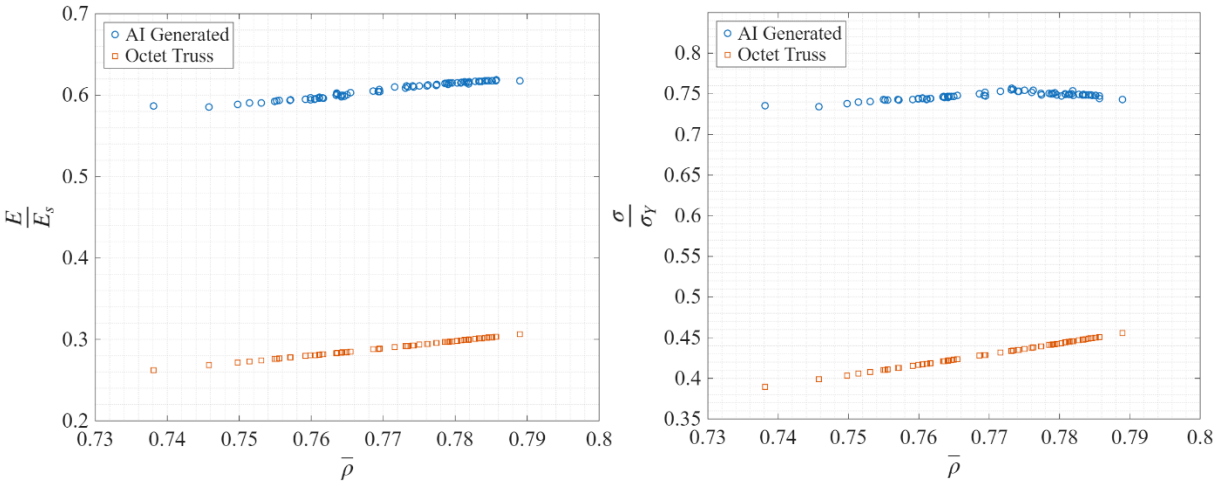


Figure 191 Comparison of Pareto front solutions against equivalent octet-truss at same relative density for normalised stiffness (left) and normalised strength (right)

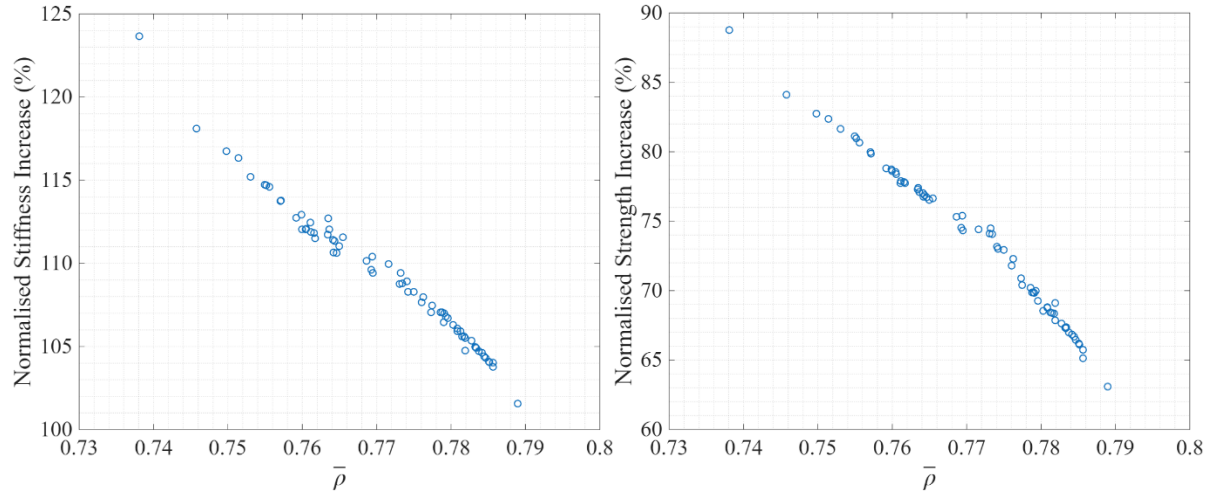


Figure 192 Percentage increase in performance from AI generated unit cell compared to equivalent octet-truss for normalised stiffness (left) and normalised strength (right)

To increase the range of relative densities in the solutions output by the GA the optimisation was repeated by including a third objective to minimise the mass of the supercell in addition to maximising relative specific stiffness and strength. The optimisation history of the revised GA is shown in Figure 193. The algorithm ran for 1hr 19min 43sec until stopping due to the average change in spread of the Pareto solutions falling below the function tolerance. A greater amount of variation is observed in the mean distance and Pareto spread as a result of increasing the number of objectives, however both plots show evidence of convergence. A trend emerges during which the dispersion of the population across the possible design space decreases over the first 50 generations before gradually levelling off as the GA hones in on a converged solution after 183 generations. The final set of Pareto solutions had mean spread of 0.136. The mean distance between individuals in the final population was 6.269.

(a) (b)

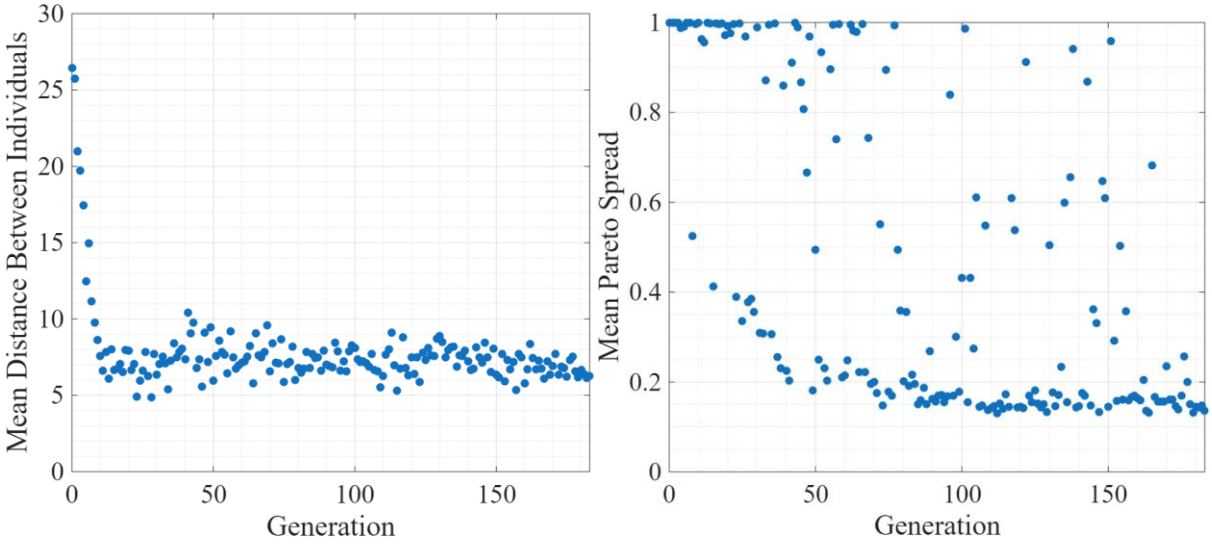


Figure 193 Optimiser history of multi-objective genetic algorithm set to minimise mass in addition to maximising relative specific stiffness and strength; (a) mean distance between individuals; (b) mean Pareto spread

The pairwise Pareto front of the GA solutions is shown in Figure 194, while the geometry parameters of each solution are provided in Appendix N. By including minimisation of mass in the objective function a front of optimal solutions was provided over an increased range of relative densities between 15.3-78.2%. The masses of the supercell ranged between 6.01-36.34 grams. Histograms of the performance objectives, displayed in Figure 195, show a trend towards favouring supercell designs with either high stiffness that exceeds the range of the training dataset, or low stiffness designs with reduced mass. Of the solutions output by the GA, 79% were in the lower half of the range of masses found in the original dataset (6.1-42.5g). Regions are observed in the Pareto front at highest mass where increasing mass

further no longer significantly improves the performance objectives. This may be due to the network predictions that were used to calculate the fitness values in the GA becoming less accurate towards the edges of the design space. In these regions the original training dataset was likely more sparsely populated due to convergence issues caused by excessively high forces during periods of large plastic deformation of the thickest radius struts. However the solutions showed strong performance when compared to both the original training dataset and equivalent octet-truss unit cells. The relative specific stiffness and strength of the solutions varied between 450.29-1423.68MPa and 9.40-62.03MPa, providing structures 35.6% stiffer and 29.9% stronger than the greatest values found in the original dataset.

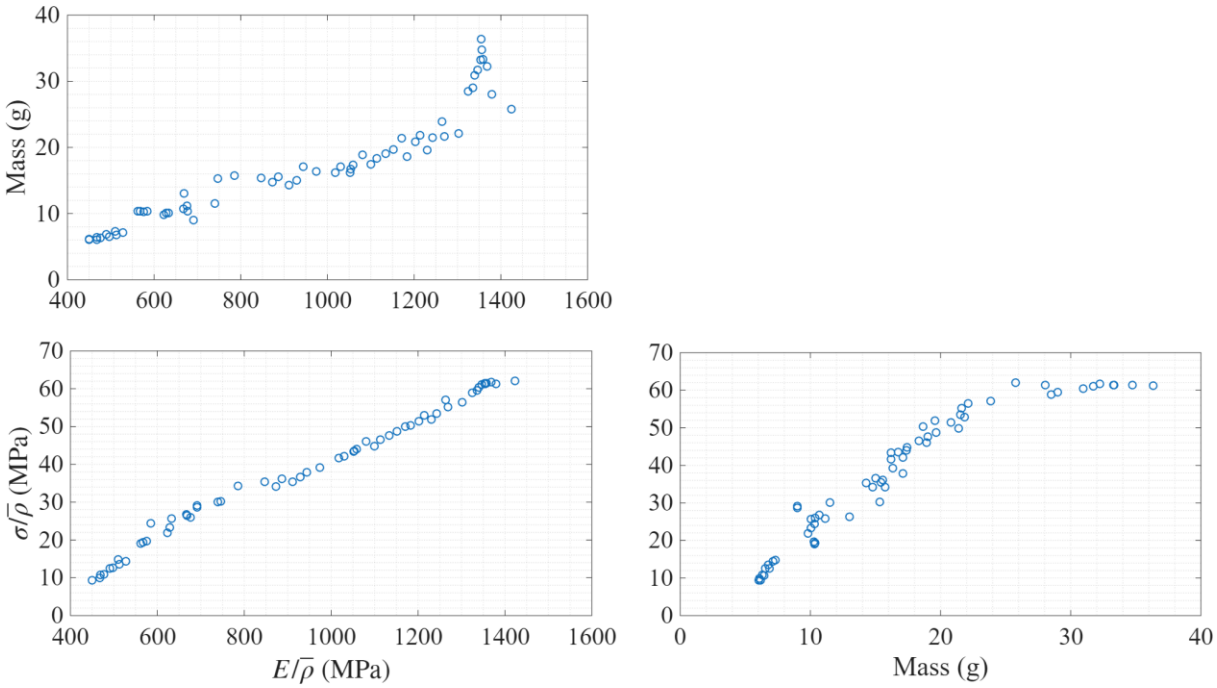


Figure 194 Pareto front of solutions from GA set to minimise mass while simultaneously maximising relative specific stiffness and strength

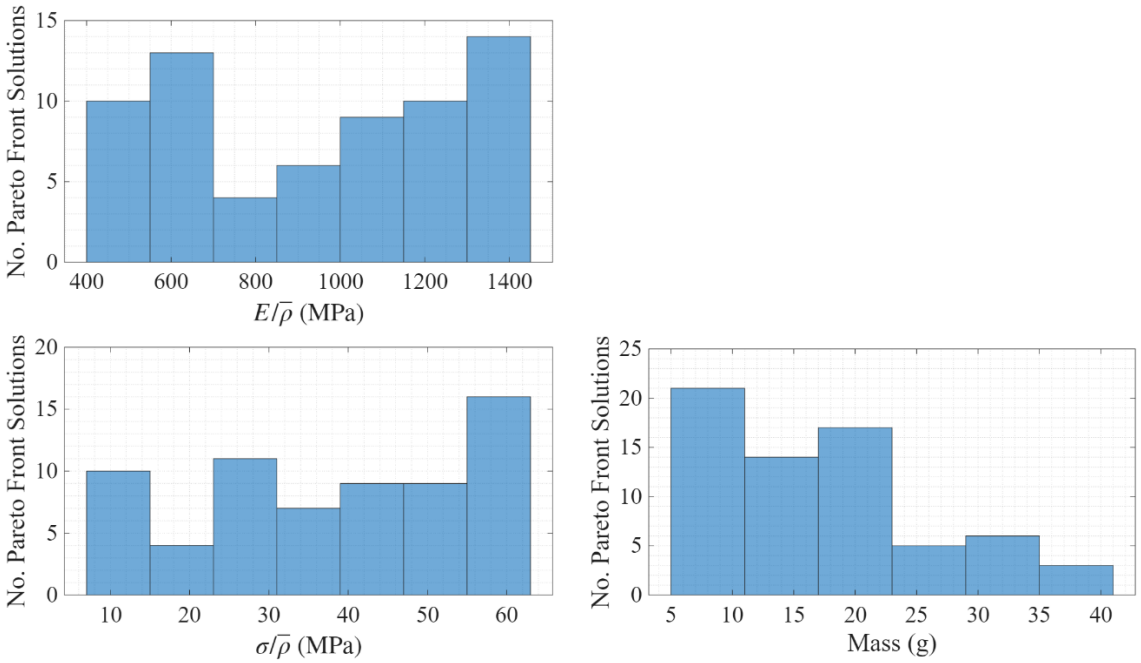


Figure 195 Performance objective histograms for GA set to minimise mass while simultaneously maximising relative specific stiffness and strength

All of the GA optimised solutions along the Pareto front outperformed the normalised stiffness of the equivalent octet-truss at the same relative density, as shown in Figure 196. The normalised stiffness was increased by between 32.2-186.3% at relative densities between 15.3-78.2%. The optimiser also provided normalised strength up to 126.7% greater than the octet-truss equivalent at relative densities between 21.7-78.2% (Figure 197). At relative densities below 20% the octet-truss was stronger than the AI generated unit cell, however in this region stiffness and strength are expected to be less than 5% of the Young's modulus and yield strength of the base material and specimens are likely to be impractical to manufacture via FDM.

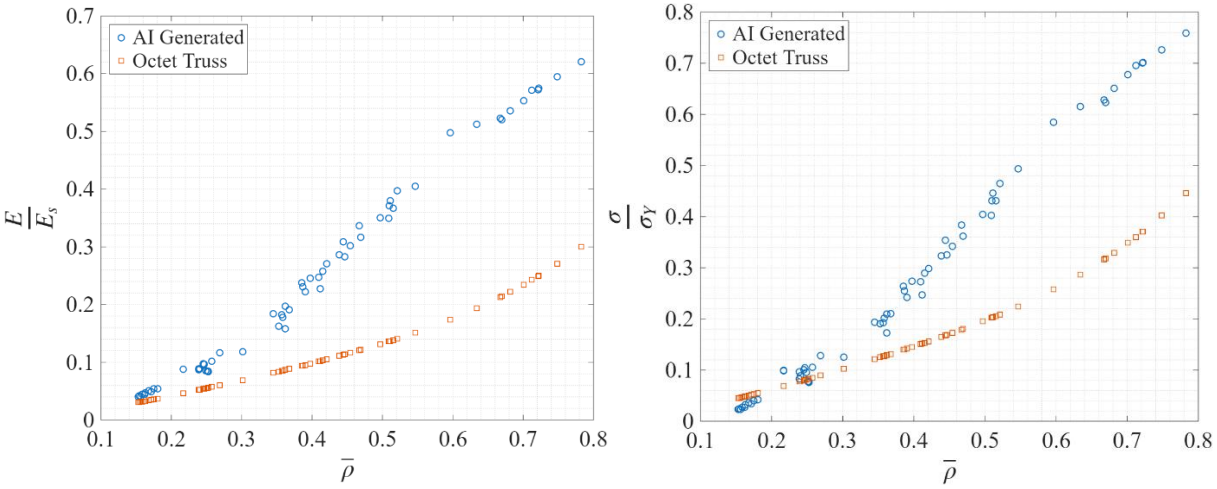


Figure 196 Comparison of revised Pareto front solutions against equivalent octet-truss at same relative density for normalised stiffness (left) and normalised strength (right)

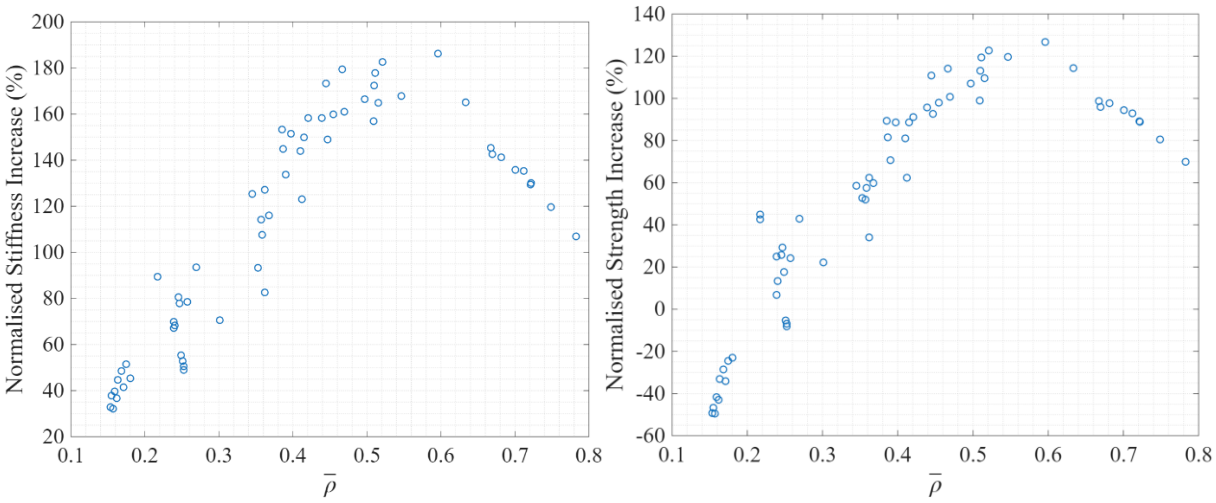
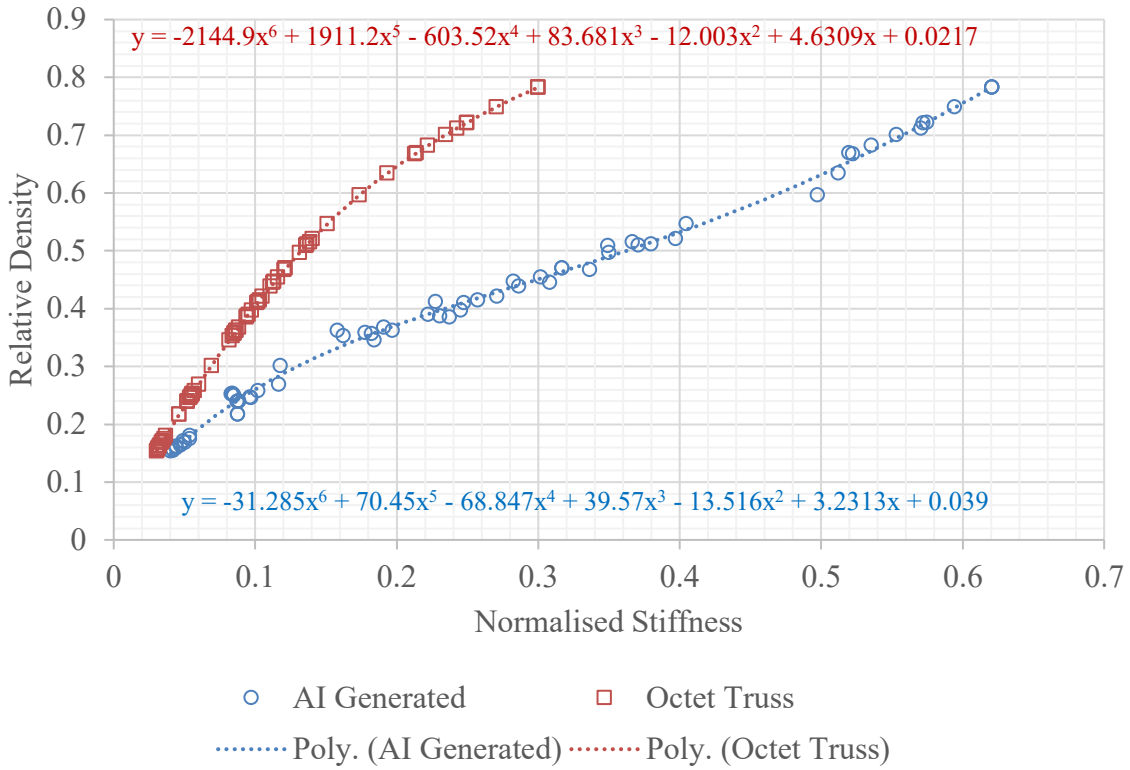


Figure 197 Percentage increase in performance from revised AI generated unit cell compared to equivalent octet-truss for normalised stiffness (left) and normalised strength (right)

In order to predict the expected mass savings that may result from using the CNN and GA based optimisation approach, a least squares polynomial regression model was fitted to the relationships between the relative density and normalised stiffness and strength (Figure 198). At each point along the octet-truss curves in Figure 196, the relative density of the AI generated unit cell required to give the same normalised stiffness or strength value was calculated from the polynomial and the mass of the supercell needed was found from the regression model shown in Figure 199. The percentage reduction in mass between the proposed optimisation workflow and the octet-truss alternative for a given stiffness or strength requirement is estimated in Figure 200. The approach is estimated to provide mass savings between 25.8-52.2% between relative densities of 15.3-78.2% for the same stiffness requirement. The AI generated metamaterial is estimated to provide strength equivalent to the octet-truss with a lattice that is up to 45.2% lighter between relative densities of 21.7-78.2%.

(a)



(b)

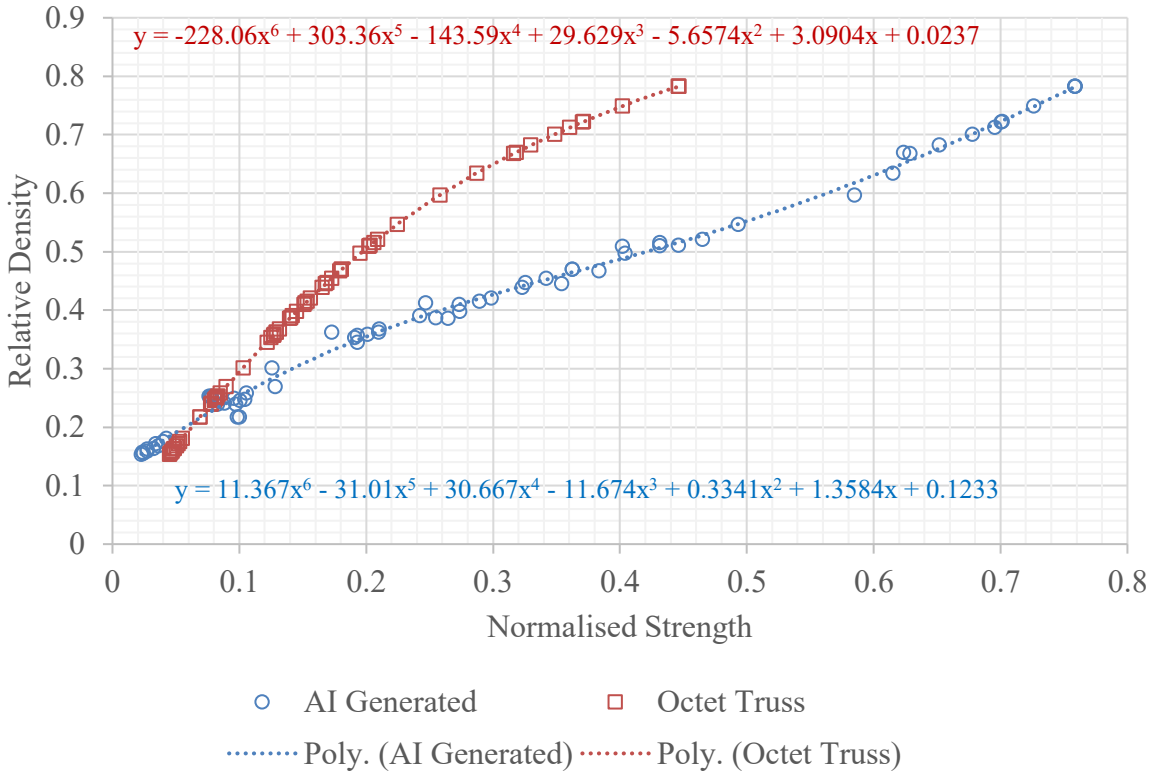


Figure 198 Polynomial regression models for prediction of relationships between performance objectives and relative density; (a) normalised stiffness; (b) normalised strength

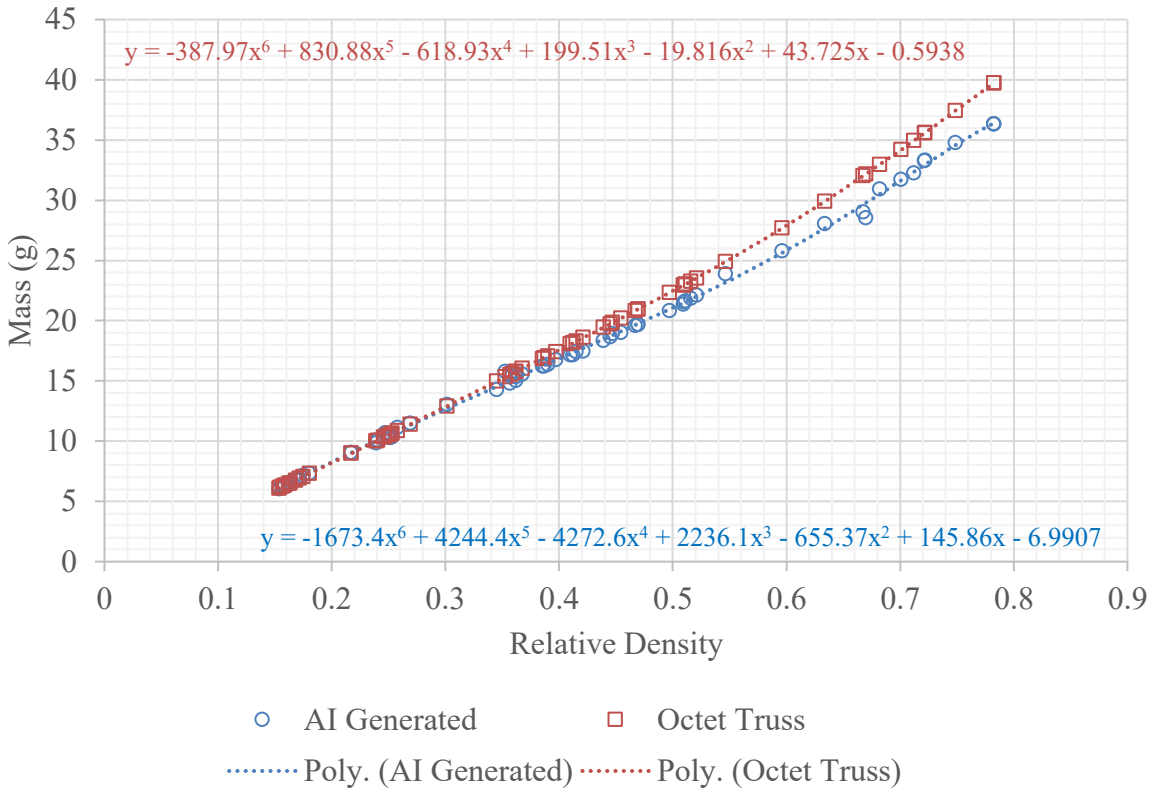


Figure 199 Polynomial regression models for prediction of relationships between supercell mass and relative density

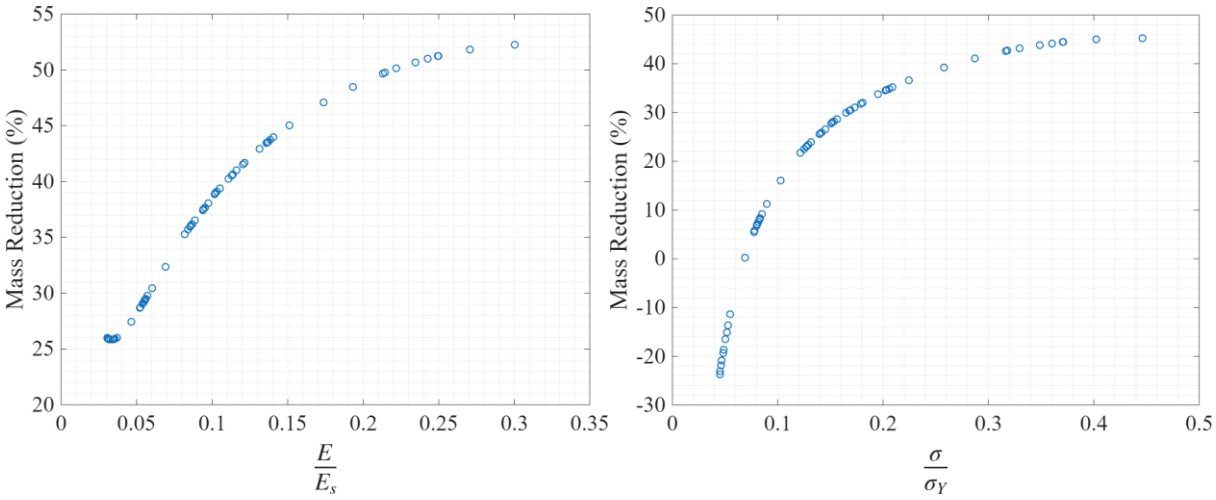


Figure 200 Estimated percentage mass reduction between revised AI generated unit cell and equivalent octet-truss for a given normalised stiffness (left) and normalised strength (right)

6.4.2 Specific Energy Absorption Optimisation

The multi-objective genetic algorithm was revised to include all 4 objectives, maximising specific energy absorption alongside relative specific stiffness and strength while minimising mass. The optimisation history is shown in Figure 201. The algorithm converged on a Pareto front after 1hr 49min 11sec and 195 generations with mean spread of 0.060 and mean distance between individuals of 12.637. Pairwise plots of the Pareto front solutions are displayed in Figure 202, while Appendix O contains the output parameters.

(a) (b)

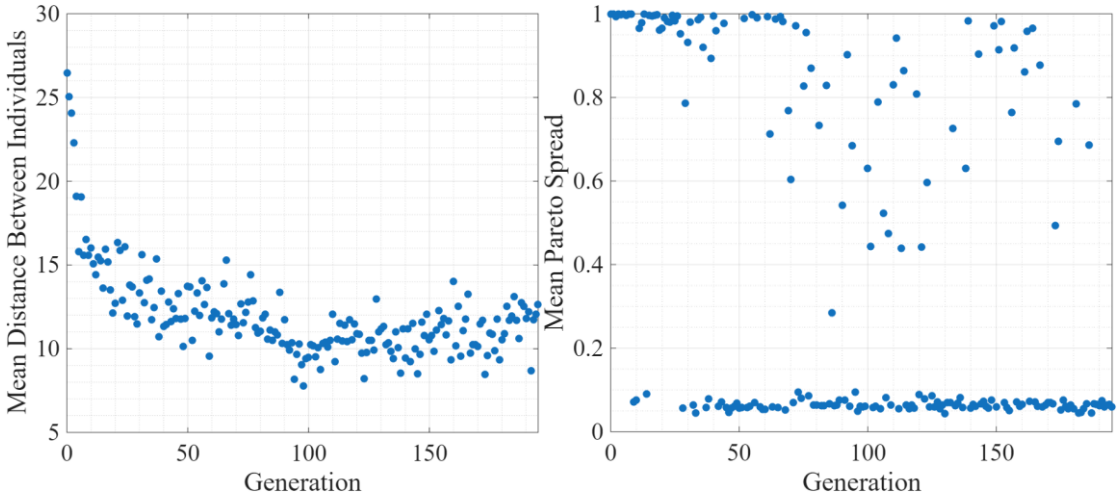


Figure 201 Optimiser history of multi-objective genetic algorithm set to maximise specific energy absorption in addition to relative specific stiffness and strength while minimising mass; (a) mean distance between individuals; (b) mean Pareto spread

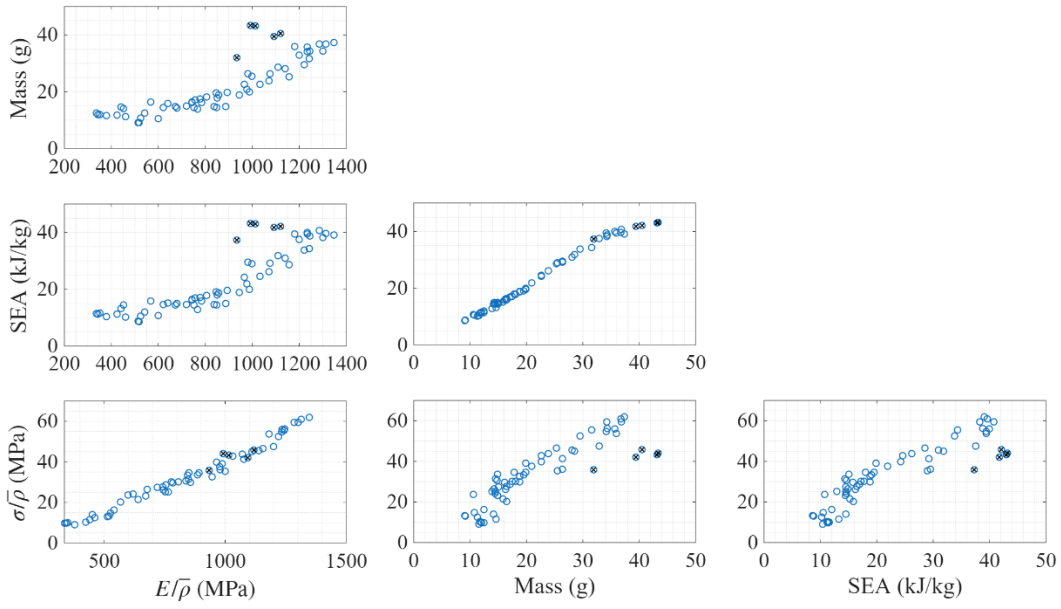


Figure 202 Pareto front of solutions from GA set to maximise SEA, relative specific stiffness and strength while minimising mass. Black markers indicate outliers.

Outliers were observed in the plots of relative specific stiffness and strength against mass. These points correspond to designs with the greatest specific energy absorption along the front, where mass has been increased to satisfy the additional performance objective. This is demonstrated in Figure 203, which indicates the outlying design points occur at greatest relative density and highest SEA. The Pareto solutions provided SEA that outperformed the equivalent octet-truss at relative densities between 26.0-91.5%, offering increases of up to 40.3%. Figures 206-209 below compare the FEA results for example solutions that gave the largest increases in SEA. Corresponding geometries of the examples are shown in Figure 210. The improvements came at the cost of reduced stiffness and strength compared to the optimisations without the SEA objective, however the majority of solutions along the front remained stiffer and stronger than the octet-truss (Figure 204 and Figure 205). The largest increases in stiffness compared to the octet-truss dropped from 186.3% to 123.5%. Maximum increases in strength fell from 126.7% to 75.2%, demonstrating the new trade-off between performance objectives that must be considered by the optimiser. Three of the outliers with the highest energy absorption showed reduced strength compared to the octet-truss (Figure 205). However these designs featured relative densities greater than 87% which is in a range unlikely to be targeted for practical design purposes.

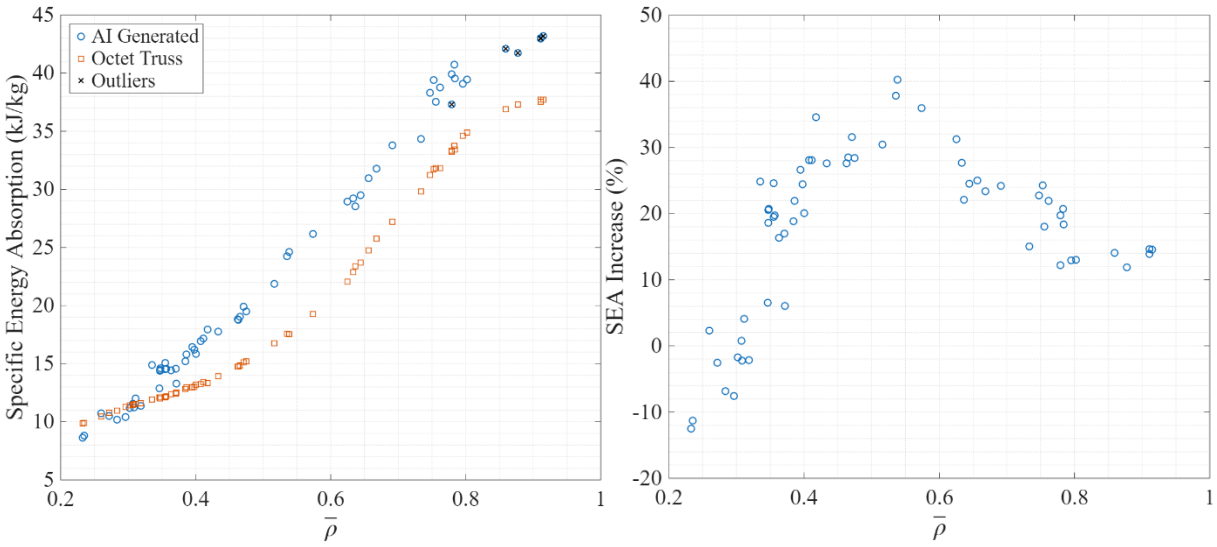


Figure 203 Comparison of specific energy absorption of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in SEA from AI generated lattice (right)

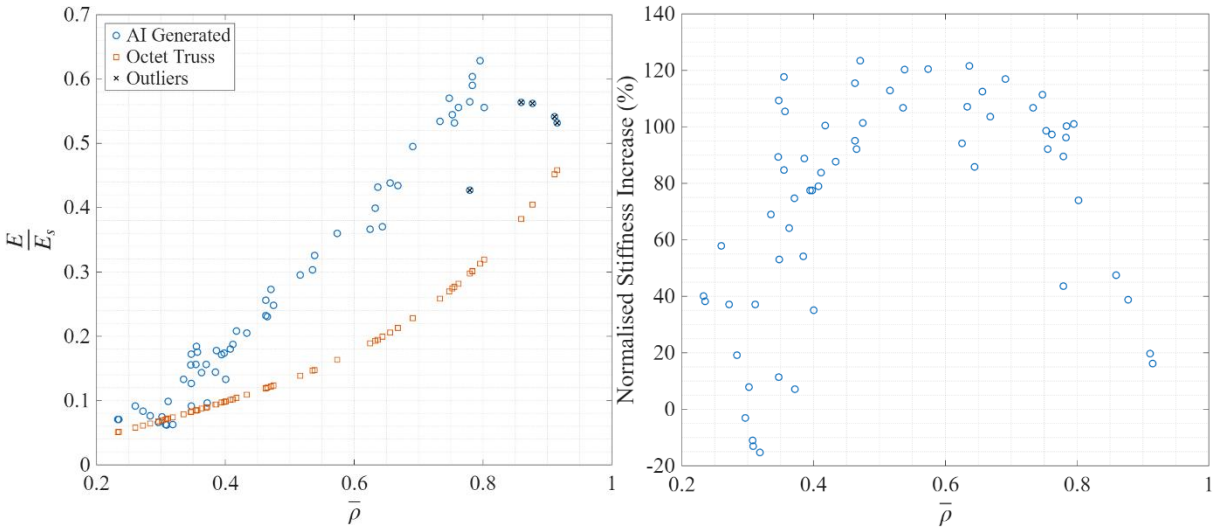


Figure 204 Comparison of normalised stiffness of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in normalised stiffness from AI generated lattice (right)

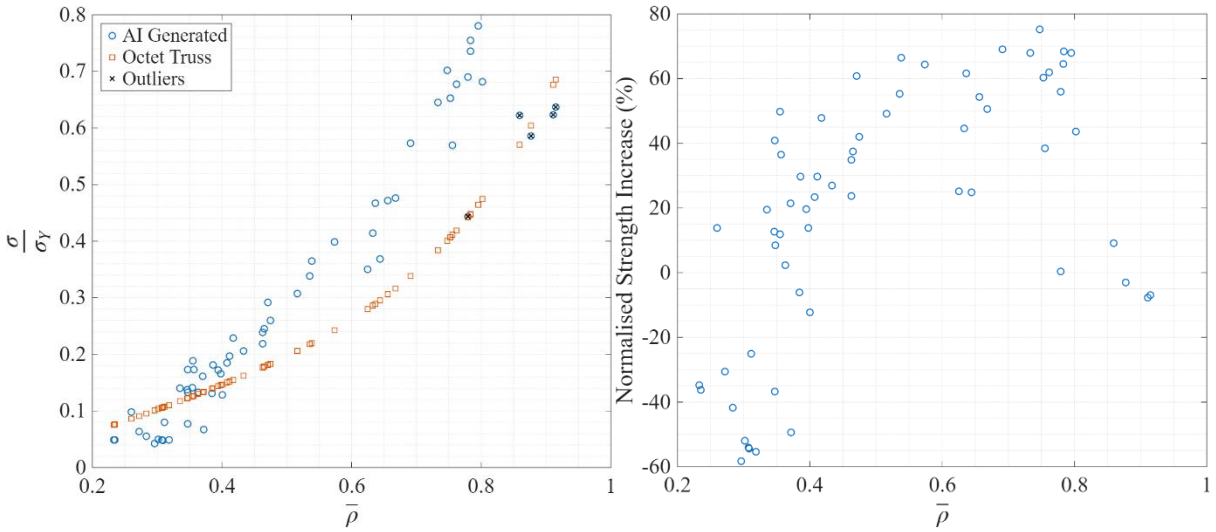


Figure 205 Comparison of normalised strength of Pareto front solutions against equivalent octet-truss at same relative density (left) and percentage increase in normalised strength from AI generated lattice (right)

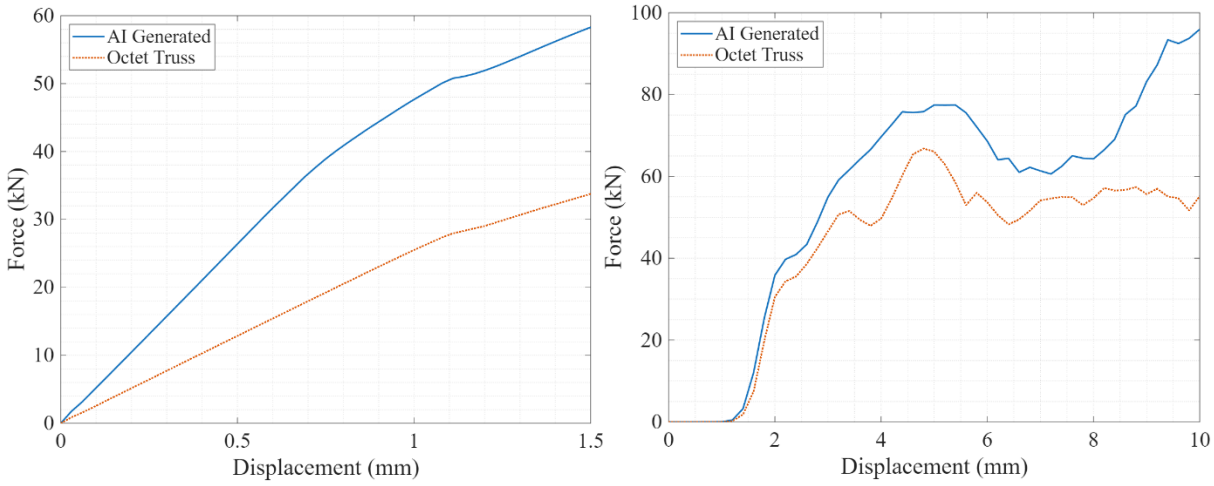


Figure 206 Example Pareto solution 31 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 40.3%, 120.4% and 66.4% respectively at relative density of 53.8%.

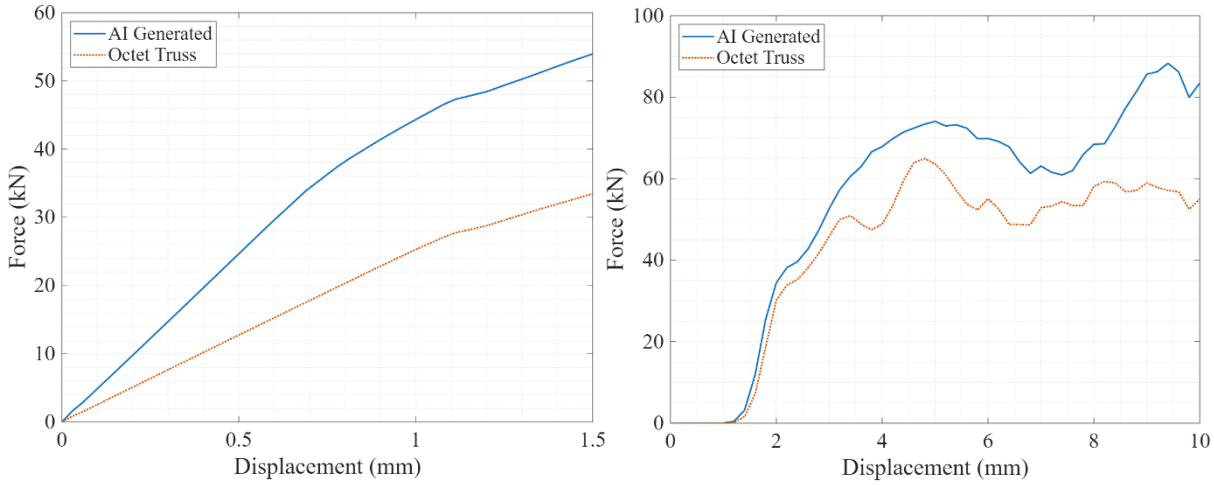


Figure 207 Example Pareto solution 56 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 37.8%, 106.7% and 55.3% respectively at relative density of 53.5%.

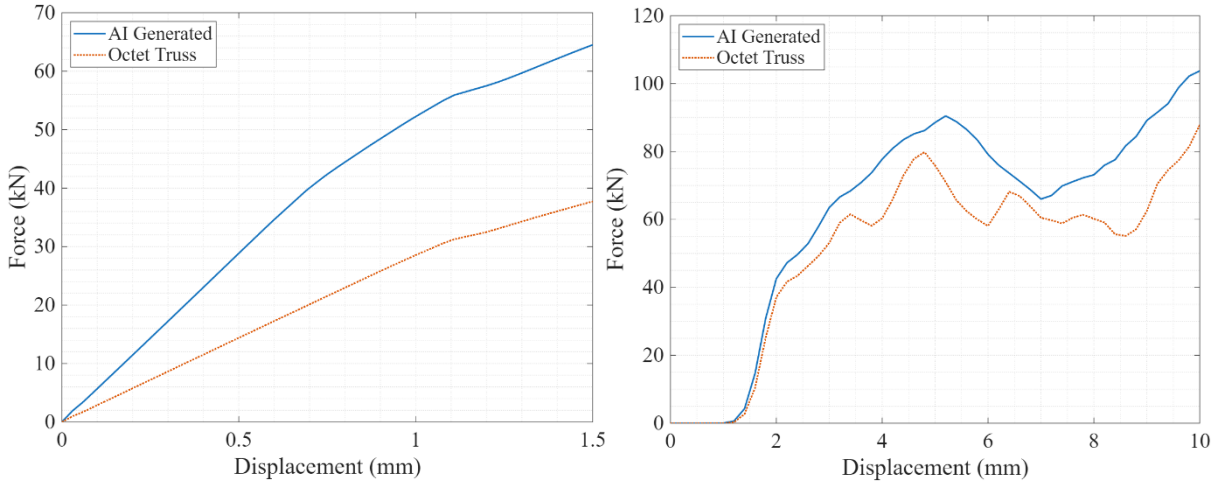


Figure 208 Example Pareto solution 25 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 35.9%, 120.5% and 64.4% respectively at relative density of 57.4%.

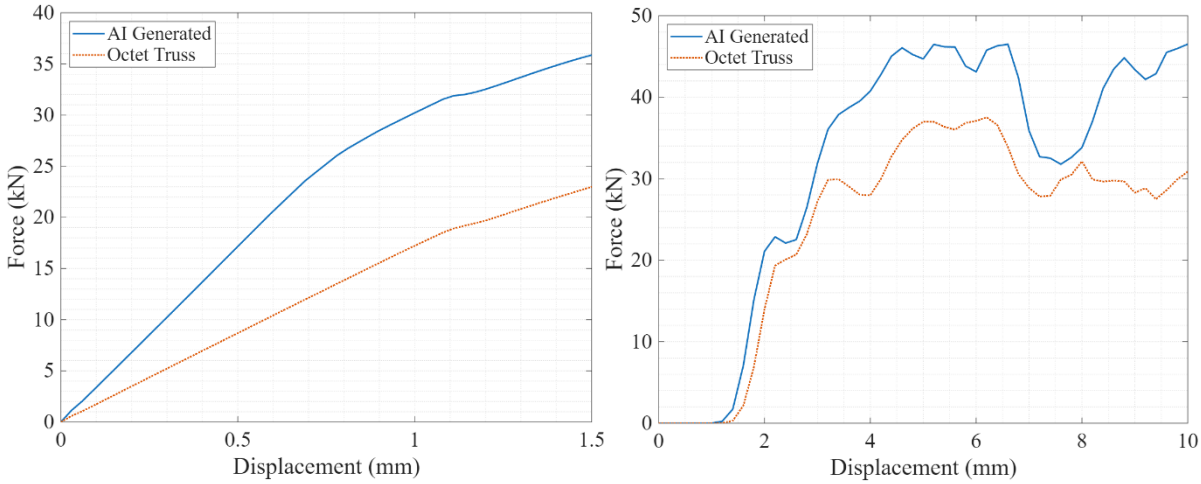


Figure 209 Example Pareto solution 42 FEA comparisons using static model (left) and dynamic model (right). Increases in SEA, relative specific stiffness and strength are 34.6%, 100.4% and 47.8% respectively at relative density of 41.7%.

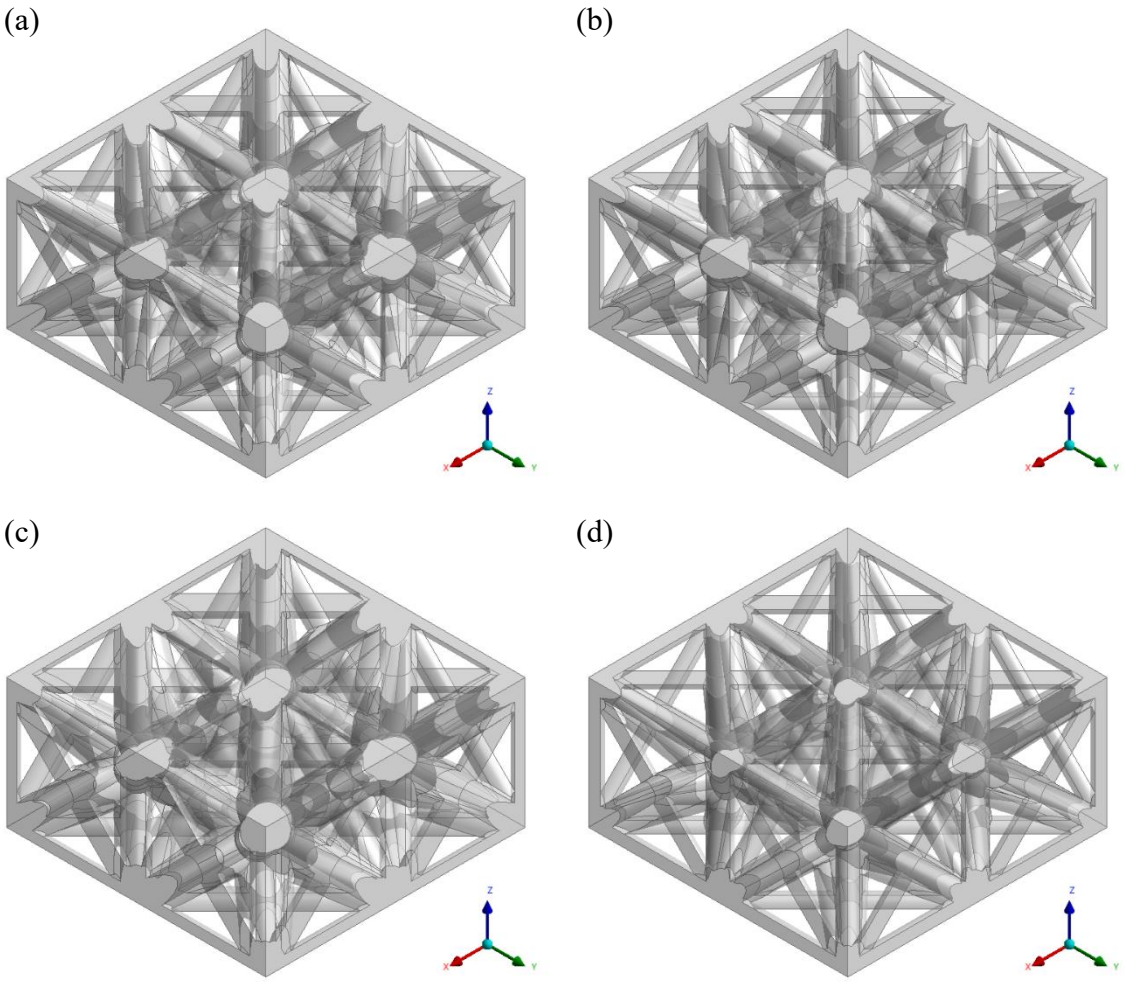


Figure 210 Supercell geometries for Pareto solutions; (a) solution 31; (b) solution 56; (c) solution 25; (d) solution 42

6.5 Conclusion

These works assessed the potential for use of machine learning to tune the performance trade-offs between stiffness, strength and energy absorption to enable a new generation of architected metamaterials with enhanced high velocity impact resistance for improvements in crash safety. Convolutional neural networks were trained to predict the relative specific stiffness and strength under quasi-static loading of a parametric lattice structure in addition to the specific energy absorption under high velocity dynamic impacts from datasets generated by implicit and explicit FEA models. On inputting the locations of nodes and strut radii into the trained networks, they were able to return predictions of the material properties that exceeded 96% regression across all parameters when compared against known values. The predictor networks were combined with a genetic algorithm optimiser to generate designs with maximum mechanical performance. Predictions were returned near instantaneously on making a request. This speed allowed for large populations of results to be calculated rapidly at every generation of the genetic algorithm, allowing for convergence of the solver within less than 2 hours on a single processor. The optimiser produced designs that were up to 35.6% stiffer and 31.5% stronger than the largest values found in the original training dataset of 3226 observations. Setting the algorithm to directly minimise mass in addition to maximising the performance objectives provided the greatest range of optimal solutions, allowing for the most flexibility during practical design scenarios. The user may choose a point along the Pareto front depending on the constraints of the design criteria to obtain an optimal solution according to a specific real-world application. On setting the solver to maximise static mechanical properties the solutions showed increased normalised stiffness by between 32.2-186.3% compared to an equivalent octet-truss lattice at relative densities between 15.3-78.2% and improvements in strength of up to 126.7%. These improvements are estimated to provide lattices that are between 25.8-52.2% lighter than the octet-truss for a given stiffness and up to 45.2% lighter for a given strength requirement. By setting the optimiser to maximise both dynamic and static properties at the same time, the solutions showed up to 40.3% greater specific energy absorption than the equivalent octet-truss under high velocity impacts at 100m/s. This came at the cost of reductions in stiffness and strength due to the inherent performance trade-offs. However designs optimised for dynamic impact resistance remained up to 123.5% stiffer and 75.2% stronger than the octet-truss. The machine learning approach has enabled simultaneous optimisation of both quasi-static and high velocity material responses in the same analysis. Synthesis of metamaterial architectures by use of artificial intelligence based workflows may in future lead to lighter vehicles with improved performance, reduced fuel consumption and carbon emissions and greater impact resistance.

Chapter 7 Layup Optimisation of Carbon Fibre Reinforced Polymer Composites by Machine Learning

There is currently limited governing theory that can be applied to determine the angles of plies in composite laminate materials that will reliably provide the strongest solution for a given structural application. Conventional composite component designs often rely on standardised stacking sequences that may not offer the optimal structural performance for the given loading scenario. A research gap exists for design processes to ensure plies are oriented in directions that give the optimal mechanical properties. The ability to tailor layup stacking sequences according to bespoke design scenarios would maximise the potential of composite materials for lightweighting, leading to increased performance, reduced fuel consumption and improved emissions in road vehicles and aircraft. This section of the project aims to develop optimisation processes for tuning the strength of carbon fibre composites for minimum stress design through the use of machine learning neural networks. A major benefit of machine learning is that it does not require analytical solutions, making it ideal for design of composites. It can be employed to solve any parametric design problem as long as a dataset can be produced for the network to train on. An inverse design approach is used to tailor stress distributions by inputting requests of desired mechanical responses to a trained convolutional neural network. The network then outputs the layups necessary to provide the requests and the model is used to reduce stress concentrations and dictate load paths through the component. An alternative approach is then explored in which a network is used to predict the stress responses for a given input layup. The predictor model is then combined with genetic algorithm and particle swarm evolutionary optimisation algorithms to minimise the maximum stress in the component.

7.1 Inverse Design Approach for Tailored Stress Distribution

7.1.1 Dataset Generation

For the inverse design model a desired stress distribution is supplied as input to a convolutional neural network as a vector containing the Von Mises stress values in every element of an FEA mesh. The model then outputs the layup stacking sequence angles necessary to provide a response as close as possible to a requested set of target values. This allows tailoring of the layup to give a desired stress contour for a certain loading and geometry. The model is built by training the neural network on a large dataset of example stress responses with varied layup sequences. For proof of concept an example is presented of a 100x100mm square composite plate under uniaxial load. The dataset was generated in ANSYS Mechanical APDL, using a script to cycle through layups and automatically extract the Von Mises stress data for each design point. The plate was meshed with SHELL181 elements 10x10mm in size, providing a vector of 100 stress values for each layup. The loads and boundary conditions are displayed in Figure 211 below.

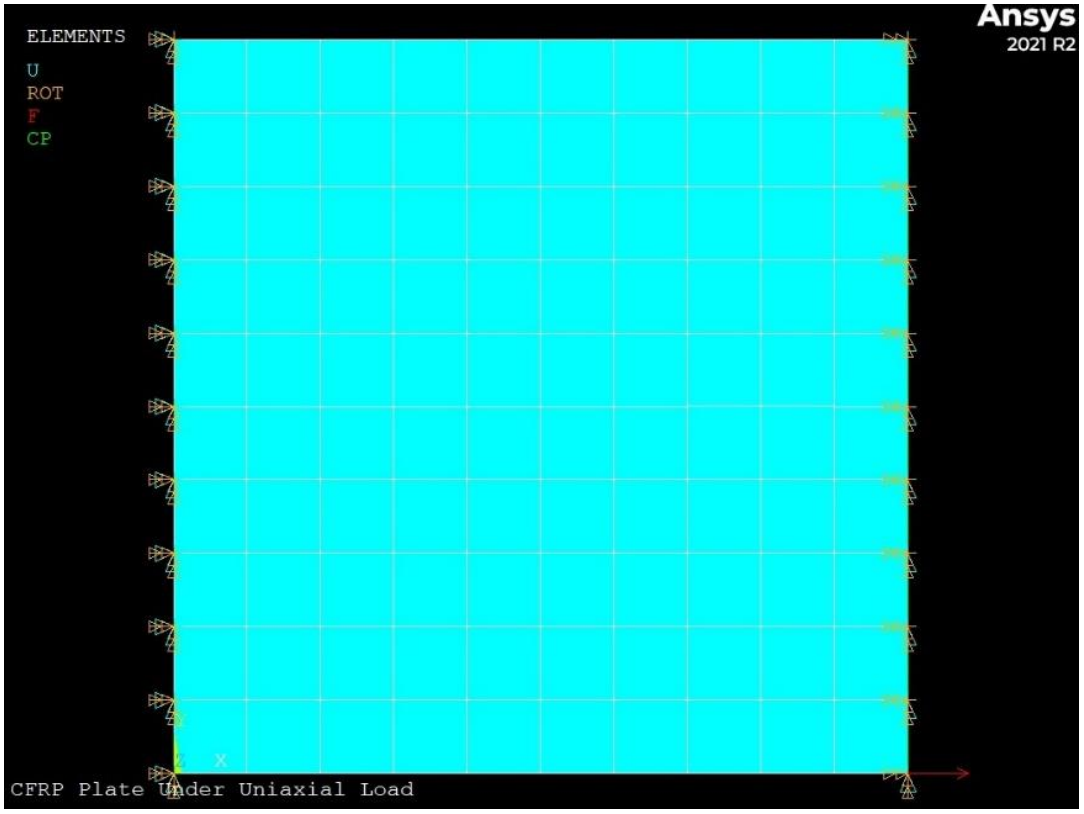


Figure 211 Composite plate FEA model mesh and boundary conditions

All nodes at the left hand edge of the plate, at coordinates where x=0 were fully fixed in all degrees of freedom. The right hand edge at x=100 was constrained to allow displacement in the x-direction only, with all rotational degrees of freedom removed to ensure only uniaxial in-plane deflections occur. The degrees of freedom of all the nodes where x=100 were coupled and a 10kN force was applied to the bottom right hand node in the positive x-direction to simulate in-plane loading applied to the right hand edge of the plate. The layup stacking sequence was defined as a set of 3 angle parameters that were used to create a 6-ply symmetric layup. Each ply had thickness of 0.125mm. The angles were arranged in the order $[\theta_1/\ \theta_2/\ \theta_3]_S$ as shown in Figure 212. Every ply used the same orthotropic material model. Values of the assigned material properties can be found in Table 39.

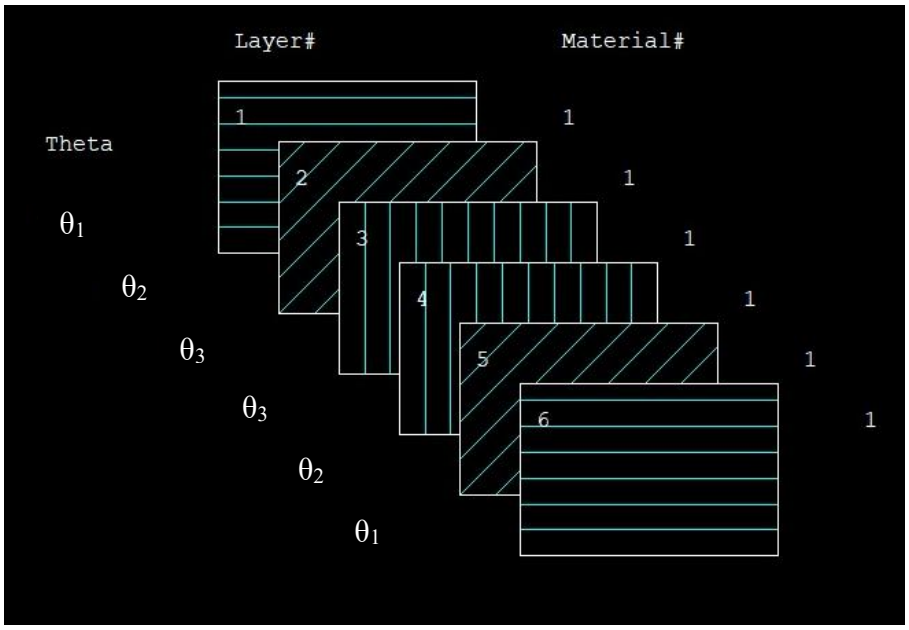


Figure 212 Composite plate FEA model layup stacking sequence

Table 39 FEA model orthotropic material data

Material Property	Direction	Value
Young's Modulus (MPa)	EX	150
	EY	11
	EZ	11
Poisson's Ratio	PRXY	0.25
	PRYZ	0.45
	PRXZ	0.25
Shear Modulus (MPa)	GXY	6
	GYZ	3.7
	GXZ	6

Each angle parameter was allowed to vary between 0° and 180°. A table of design points containing the stacking angles for each run were contained in a .csv file placed in the working directory and read into ANSYS automatically by the script using the TREAD function. The first column contained the design point number which was used to index into the table to update the layup, repeating the analysis with a *DO command loop. On the completion of each run the Von Mises equivalent stress in every element was saved to an array parameter and written to output .dat files in the working directory. In addition the locations of every node, the x and y coordinates of the centroid of every element, and the corresponding node number at the vertices of each element were also output to data files to allow the mesh and stress contours to be plotted in MATLAB for each result.

7.1.2 Convolutional Neural Network Training and Performance Evaluation

The minimum and maximum stress values that appeared anywhere in the dataset were extracted and used to rescale the stress data, normalising all values to between 0 and 1 to aid the statistical mapping process of the neural network. The normalised stress vectors were then formatted for sequence input to a convolutional neural network (CNN) model implemented in MATLAB. The network featured a series of 12 convolution layers with 1-dimensional convolution filters, followed by a global average pooling layer to flatten the extracted feature maps to a single vector. This vector was then fed through a network of 9 fully connected layers to a regression layer with 3 outputs corresponding to the 3 angle parameters. The network architecture was repurposed via transfer learning from the architecture developed in Chapter 3. Details of the architecture including normalisation, pooling and activation measures can be found in section 3.2.

The training was repeated multiple times to tune network hyperparameters and establish the size of the dataset necessary for predictions of sufficient quality through a process of trial and error. Several key test runs have been selected from the iterations to demonstrate the evolution of the predictive performance of the network throughout the process and are summarised in Table 40. For every test the datasets were partitioned with 70% of the data used to train the network, allowing 15% to be used for validation checks as the training progressed and 15% left unseen by the network during training so that it could be used to test predictions against known values after the training was complete. During training, at each validation check the root mean square error (RMSE) and loss for the validation data was compared against values from the training data. The final RMSE and loss from the last validation check at the end of the training are shown in Table 40 for each run, with lower values corresponding to greater network accuracy and agreement between training and validation results indicating

repeatability. After the training was complete, at every design point of the test partition data the stress vector corresponding to the known layup was input to the trained network. The network output a prediction of the stacking sequence needed to produce the stress distribution. The coefficients of regression between predicted and known values of the 3 stacking angles were then calculated across the entire test dataset to evaluate the performance of the network when tested on unfamiliar data. The evolution of this performance as the network was iterated is displayed in Figure 213.

Table 40 Convolutional neural network training iterations

Test	Data		Options		Training Performance			Regression	
A	No. Samples	2000	Training Algorithm	Adam	RMSE	Training	0.16	Parameter	R Value
			Mini Batch Size	100		Validation	0.41		
	Sampling Method	Latin Hypercube	No. Epochs	1000	Loss	Training	0.01		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.10	θ1	0.94236
	Validation (%)	15	Validation Frequency	1	Training Time	50min 49sec		θ2	0.49712
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	14000		θ3	0.53257
B	No. Samples	2000	Training Algorithm	Adam	RMSE	Training	0.33	Parameter	R Value
			Mini Batch Size	10		Validation	0.40		
	Sampling Method	Latin Hypercube	No. Epochs	50	Loss	Training	0.05		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.08	θ1	0.89376
	Validation (%)	15	Validation Frequency	1	Training Time	27min 33sec		θ2	0.59821
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	7000		θ3	0.56531
C	No. Samples	2000	Training Algorithm	Adam	RMSE	Training	0.38	Parameter	R Value
			Mini Batch Size	10		Validation	0.37		
	Sampling Method	Latin Hypercube	No. Epochs	250	Loss	Training	0.07		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.07	θ1	0.92435
	Validation (%)	15	Validation Frequency	1	Training Time	2hr 15min 37sec		θ2	0.56443
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	35000		θ3	0.59508
D	No. Samples	100000	Training Algorithm	Adam	RMSE	Training	0.30	Parameter	R Value
			Mini Batch Size	1000		Validation	0.30		
	Sampling Method	Random	No. Epochs	350	Loss	Training	0.04		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.05	θ1	0.99582
	Validation (%)	15	Validation Frequency	1	Training Time	1hr 35min 14sec		θ2	0.67585
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	24500		θ3	0.68171
E	No. Samples	100000	Training Algorithm	Adam	RMSE	Training	0.31	Parameter	R Value
			Mini Batch Size	100		Validation	0.31		
	Sampling Method	Random	No. Epochs	50	Loss	Training	0.05		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.05	θ1	0.97642
	Validation (%)	15	Validation Frequency	1	Training Time	2hr 3min 8sec		θ2	0.66787
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	35000		θ3	0.66715
F	No. Samples	100000	Training Algorithm	Adam	RMSE	Training	0.30	Parameter	R Value
			Mini Batch Size	1000		Validation	0.31		
	Sampling Method	Random	No. Epochs	2000	Loss	Training	0.04		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.05	θ1	0.99158
	Validation (%)	15	Validation Frequency	1	Training Time	9hr 4min 43sec		θ2	0.66809
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	140000		θ3	0.66005
G	No. Samples	500000	Training Algorithm	Adam	RMSE	Training	0.29	Parameter	R Value
			Mini Batch Size	1000		Validation	0.32		
	Sampling Method	Random	No. Epochs	50	Loss	Training	0.04		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.05	θ1	0.98750
	Validation (%)	15	Validation Frequency	1	Training Time	1hr 8min 56sec		θ2	0.67772
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	17500		θ3	0.64321
H	No. Samples	500000	Training Algorithm	Adam	RMSE	Training	0.30	Parameter	R Value
			Mini Batch Size	1000		Validation	0.30		
	Sampling Method	Random	No. Epochs	350	Loss	Training	0.04		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.04	θ1	0.99730
	Validation (%)	15	Validation Frequency	1	Training Time	8hr 1min 17sec		θ2	0.69705
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	122500		θ3	0.69601
I	No. Samples	500000	Training Algorithm	Adam	RMSE	Training	0.29	Parameter	R Value
			Mini Batch Size	1000		Validation	0.31		
	Sampling Method	Random	No. Epochs	1000	Loss	Training	0.04		
	Train (%)	70	Initial Learning Rate	0.001		Validation	0.05	θ1	0.97374
	Validation (%)	15	Validation Frequency	1	Training Time	23hr 12min 25sec		θ2	0.68360
	Test (%)	15	Shuffle	Every Epoch	No. Iterations	350000		θ3	0.63599

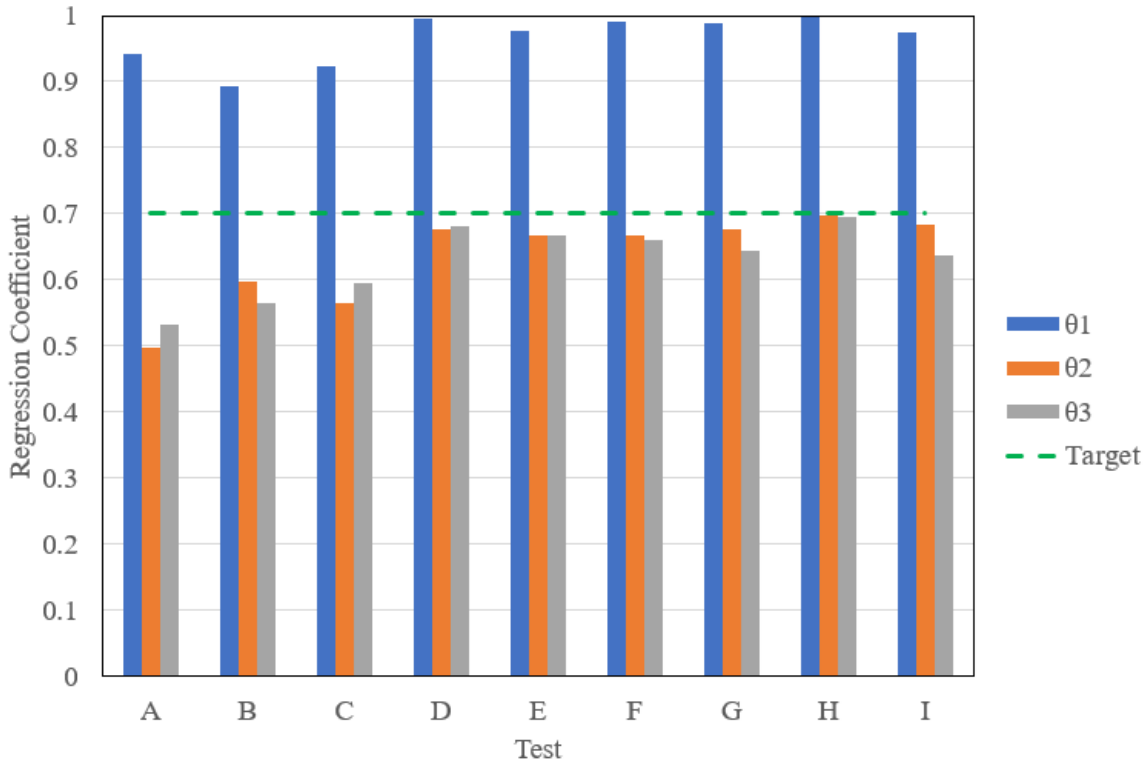


Figure 213 Network training performance evolution

All iterations were trained on identical hardware with an Intel i7-10700K processor at 3.8GHz and 32GB RAM using GPU acceleration on an Nvidia RTX 3060. In each case the initial learning rate was set at 0.001, the data was shuffled after every epoch and the adaptive moment estimation (Adam) optimiser was used to scale the learning rate for each parameter, using momentum to smooth the gradient descent. Tests A-C were trained on a limited dataset of 2000 design points sampled using a Latin hypercube design of experiments test plan to select the layup sequences. Initially the dataset was passed through the network in mini batches of 100 samples and showed strong regression for the first angle parameter after training for 1000 epochs, but poor predictions of angles 2 and 3. This was improved by passing the data more slowly in batches of 10, with regression approaching 60% and stronger agreement between training and validation values of RMSE and loss observed at the cost of increasing the training time. Tests D-F showed more consistent agreement between training and validation performance, indicating improvement in the generalisation of the dataset as a result of increasing the number of design points to 100,000 random samples. In addition the regression improved to no less than 66% across all parameters. The best trade-off between performance and training time was identified at a mini batch size of 1000. Increasing the size of the dataset further to 500,000 samples showed further improvements in the performance of networks G-I, however by a considerably smaller amount which may indicate the study approaching the limit of what is practically achievable on the current hardware for this

specific physics problem. Regression values close to 1 indicate near perfect predictions, where use of the layup sequence is expected to produce stress contours that are an exact match of the requested values in every element of the FEA mesh. In practice this is likely to be technically unachievable across all angle parameters. However it is unlikely to be a necessary requirement for enabling practical design of composite components. Instead the network was iterated until it was able to identify the underlying trends in the stress distributions sufficiently to allow for tailoring of the stress contours and agreement in maximum stress predictions. This was realised at approximately 70% regression across all parameters. At this level good agreement was observed in overall shape and magnitude between requested stress contours and those obtained via FEA from the predicted layup. Furthermore changes in the requests such as movement of contour boundaries or scaling of values to reduce the stresses were reflected in the results from the layups received back from the neural network, with strong responsiveness that allowed tailoring of the stress distribution for the given loading, as evidenced in section 7.1.3. The best performing network from test H was taken forward to demonstrate examples of tailoring the layup for desired stress responses. It was trained from a dataset of 500,000 samples over 350 epochs in mini batches of 1000, taking 8hr 1min 17sec to train. Figure 214 shows the training history of RMSE and loss. Agreement remained consistent between training and validation values throughout the training process with no indication of overfitting. Increasing the number of epochs to 1000 lead to a small drop off in regression as shown in test I, indicating increasing the training time further is unlikely to improve results. The regression between known and predicted values of the 3 angle parameters is plotted in Figure 215. Each plot features 75,000 datapoints from the test dataset partition. The regression results indicate overall trends of correlation with some uncommon outliers near the edges of the design space at angles of 0° or 180°. Some irregularities in the predictions here is likely to be expected as it is where the design space overlaps itself and the network has not been provided with the knowledge that a ply angled at 0° is equivalent to one at 180°.

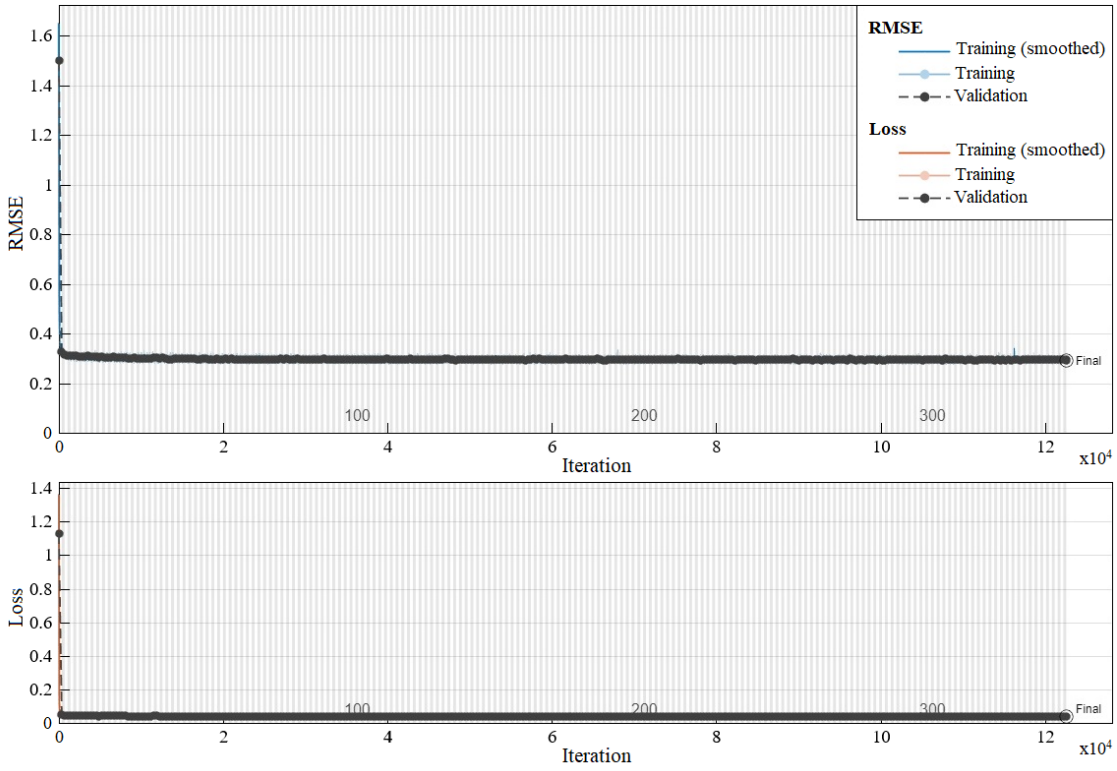


Figure 214 Network H training history

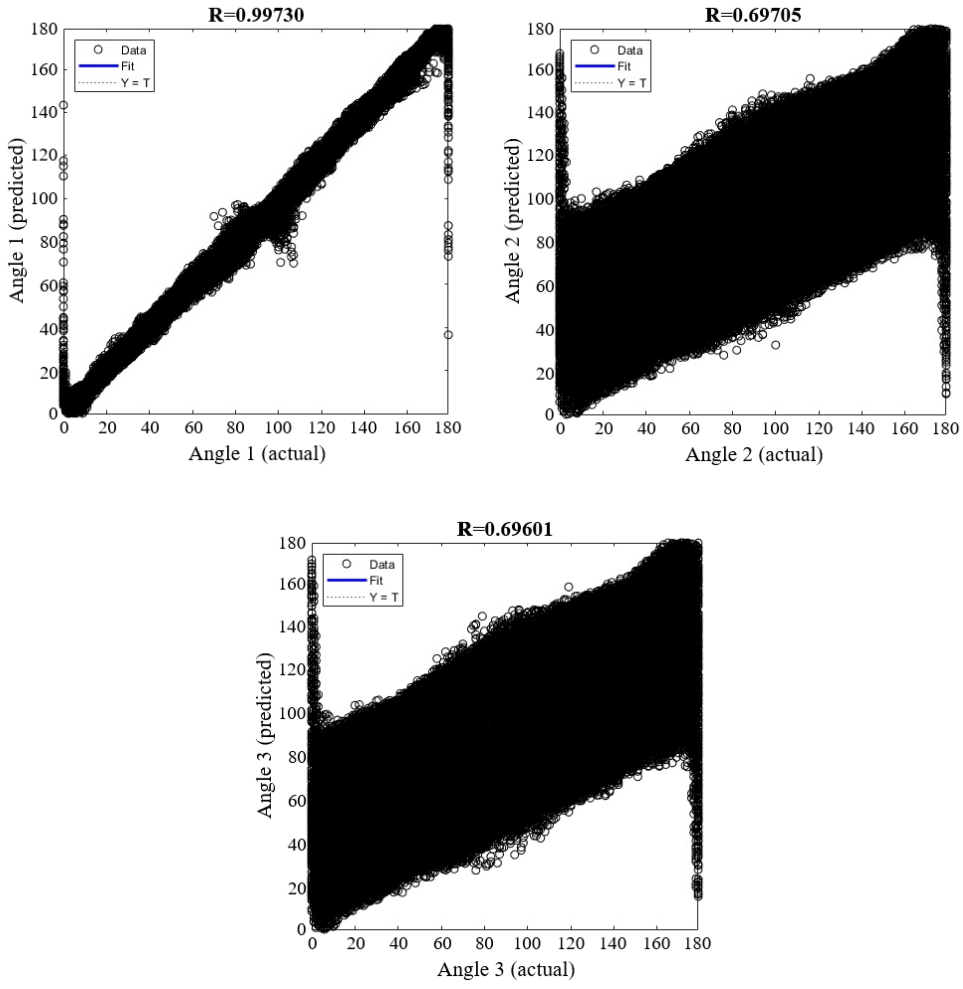


Figure 215 Network H regression plots between predicted and actual known layup angles

7.1.3 Tailoring Composite Layup for Desired Stress Responses

To evaluate the potential of the trained neural network for tailoring stress distributions in design of composite components, initial layups with high stress responses were selected from the training dataset for use as baseline examples. These represent the design scenario where a highly stressed composite is to be strengthened by only altering the layup stacking sequence without adding any extra plies or additional mass, enabling enhanced lightweighting performance. The baseline responses are shown in Figure 216 below.

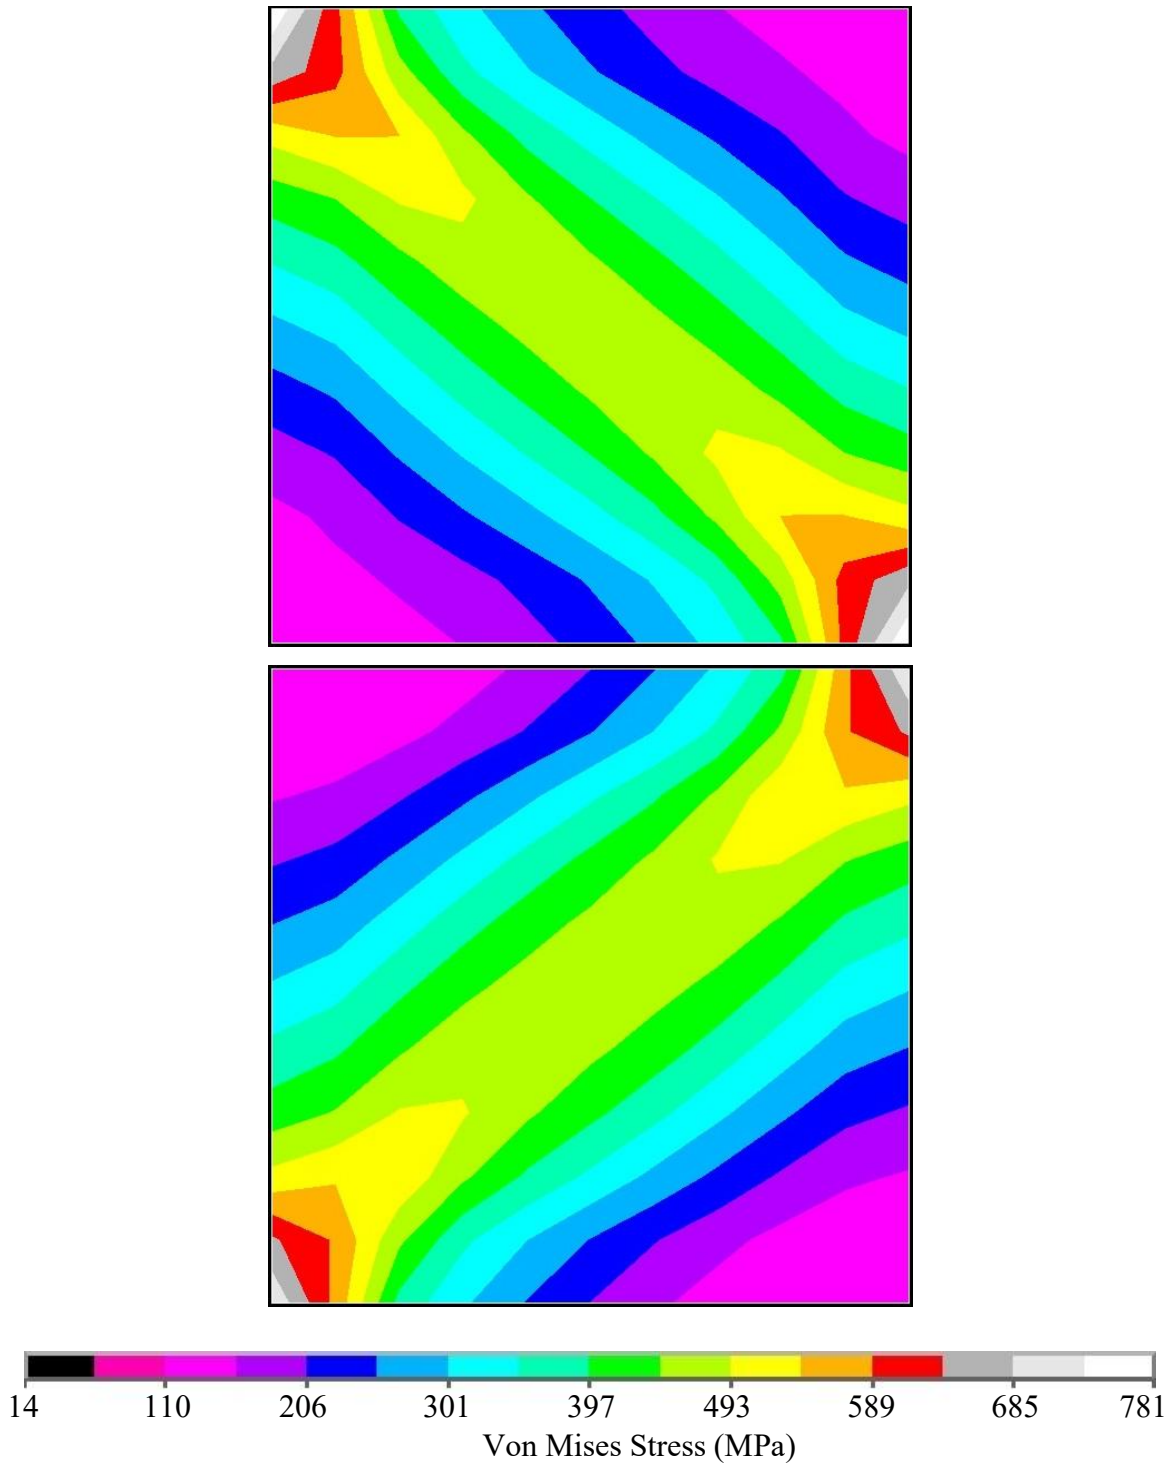


Figure 216 Baseline composite stress responses; example 1 (top) and example 2 (bottom)

The initial Von Mises stress distributions were edited in MATLAB to create a set of desirable responses to be requested of the neural network. The boundaries of each contour were redefined and the values scaled to reduce the maximum stress concentrations, while maintaining the major load paths through the material. The stress vectors were input into the trained CNN which provided a set of stacking sequences to satisfy the request, effectively attempting to match the response surface contours as closely as possible. The predicted layup sequences were then run back through the same FEA model used to generate the original dataset and the resulting stress contours compared. The new stress distributions were then re-edited and the process iterated to continue reduction of the maximum and average stresses across the plate. The requests at every iteration are summarised and compared against the actual responses provided by the CNN for each example in Table 41. The corresponding contour plots are displayed side-by side for comparison in Figure 217 and Figure 218.

Table 41 Summary of requested stress responses, layup sequences returned by the CNN and resulting actual stress responses measured via FEA

	Iteration No.	Requested Stress Response (MPa)			Predicted Layup Sequence (°)			Actual Stress Response (MPa)		
		Maximum	Minimum	Mean	Angle 1	Angle 2	Angle 3	Maximum	Minimum	Mean
Example 1	Baseline	N/A	N/A	N/A	145	71	121	778.50	125.88	352.61
	1	640.48	153.30	346.08	162	102	103	582.77	139.03	355.89
	2	491.90	223.56	319.32	176	123	124	474.23	269.06	330.72
	3	195.00	84.28	135.34	87	108	108	212.68	106.42	165.70
	4	145.76	21.97	56.55	135	155	157	192.35	24.61	76.49
Example 2	Baseline	N/A	N/A	N/A	37	61	77	728.13	110.33	350.98
	1	594.44	138.21	318.90	17	69	71	588.10	156.64	354.31
	2	411.53	50.41	195.78	25	36	37	512.93	61.62	219.90
	3	377.16	27.62	159.41	32	33	34	396.52	29.96	169.00
	4	181.75	21.01	86.06	43	40	41	420.26	35.22	145.13
	5	94.96	21.70	50.22	54	43	44	221.30	55.04	103.26

Example 1

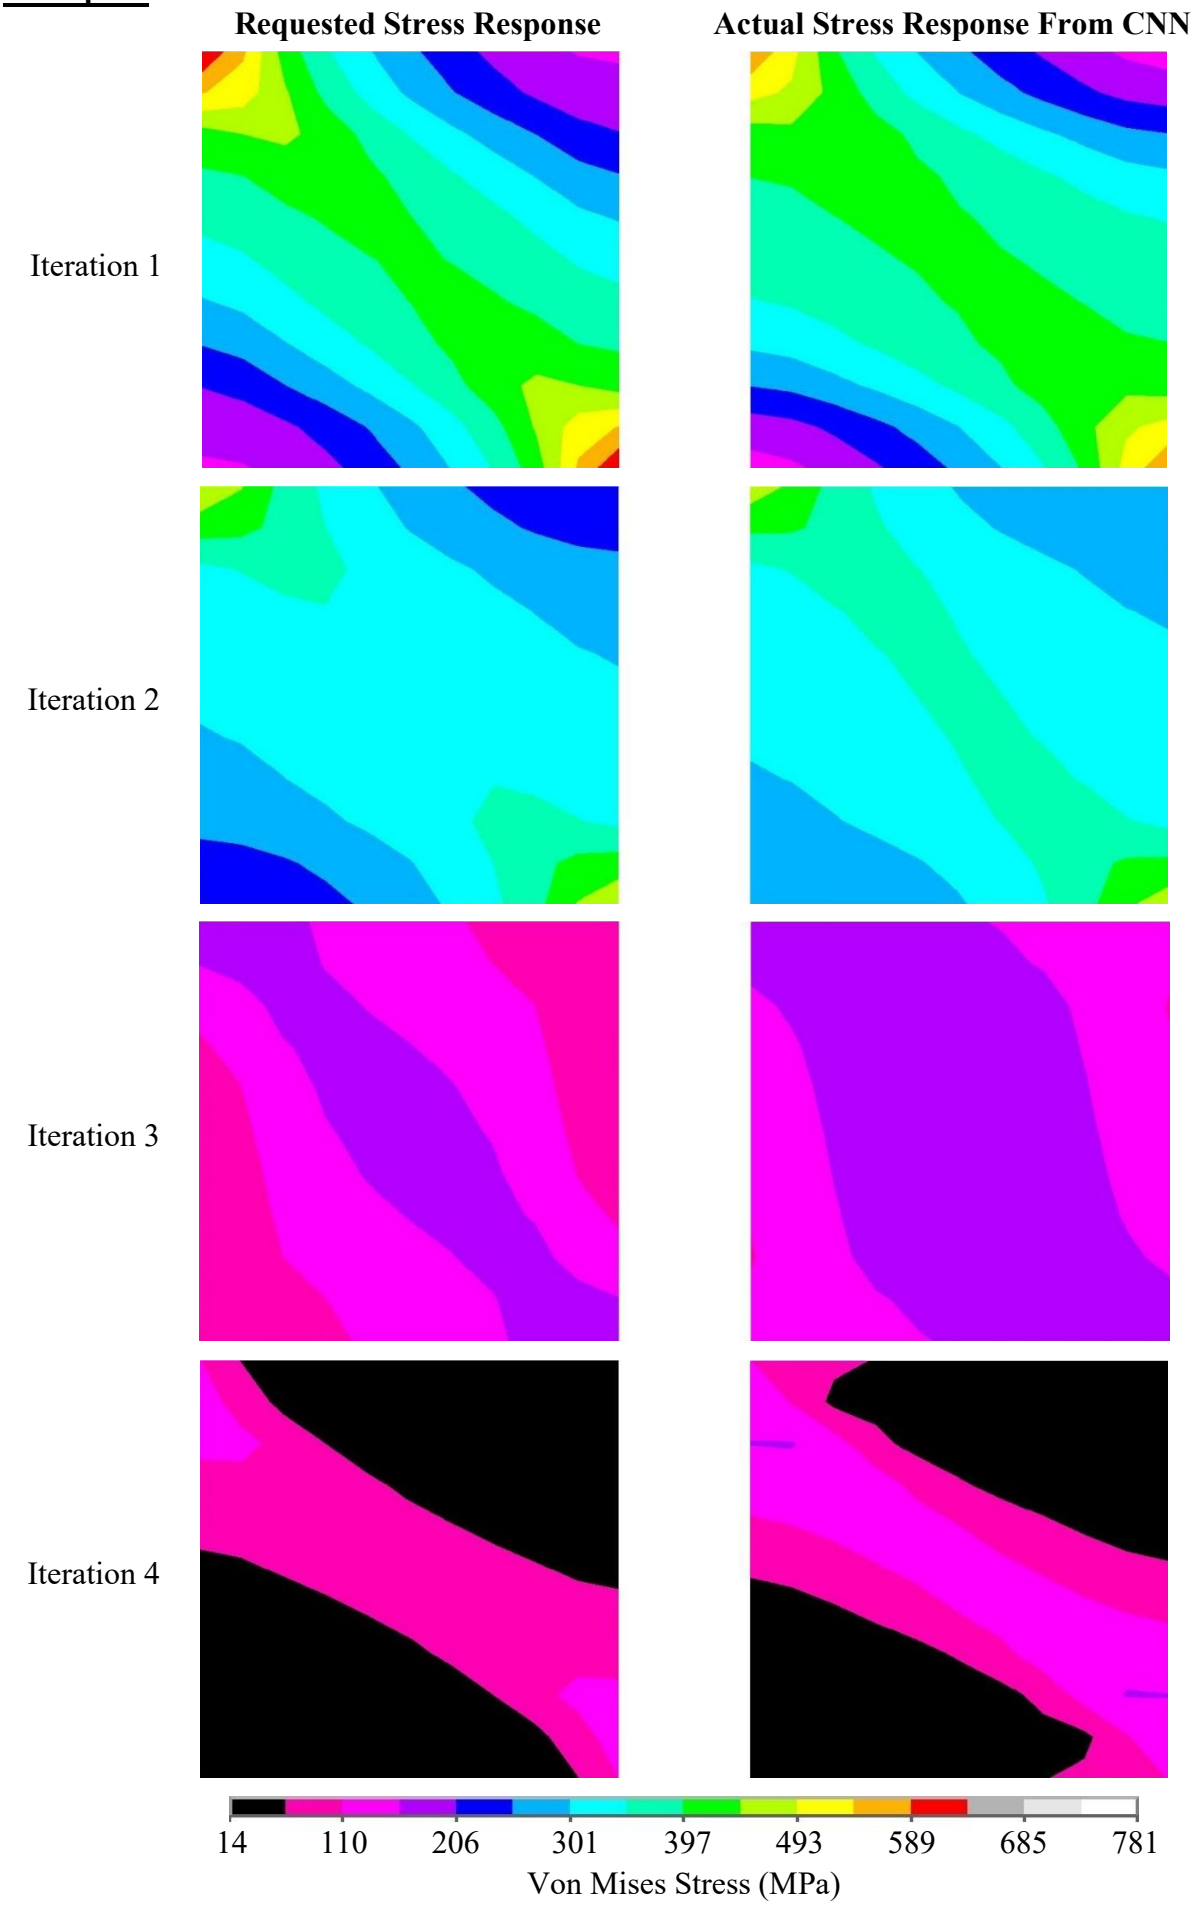
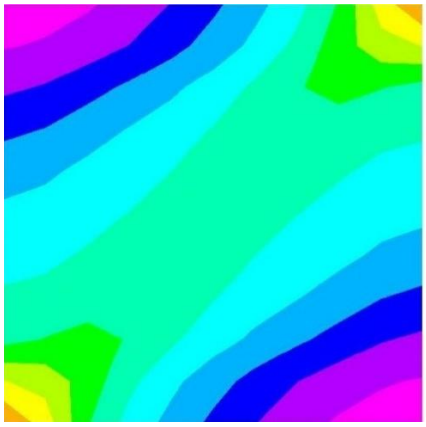
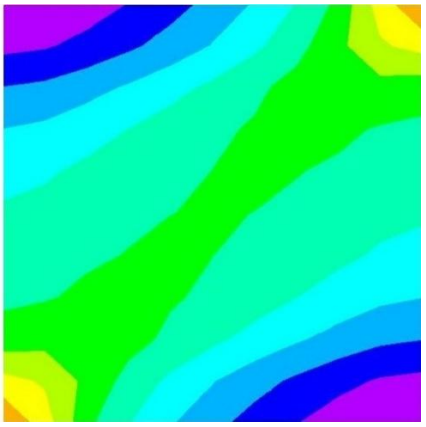
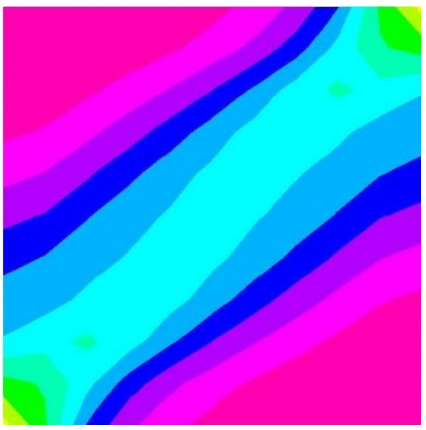
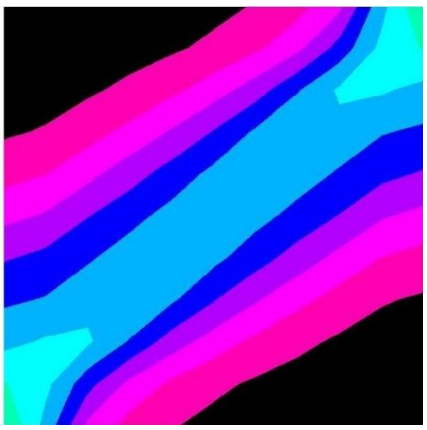
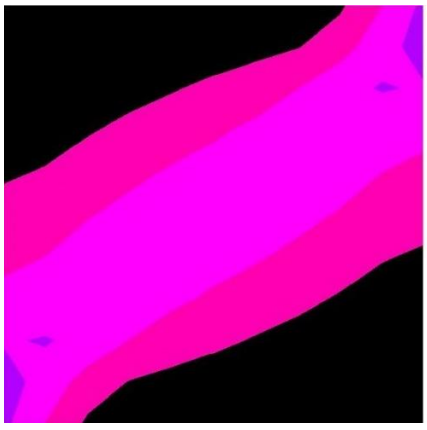
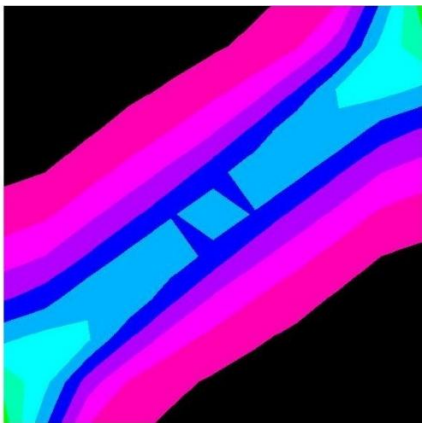


Figure 217 Comparison of requested and actual stress responses - Example 1

Example 2

	Requested Stress Response	Actual Stress Response From CNN
Iteration 1		
Iteration 2		
Iteration 3		
Iteration 4		

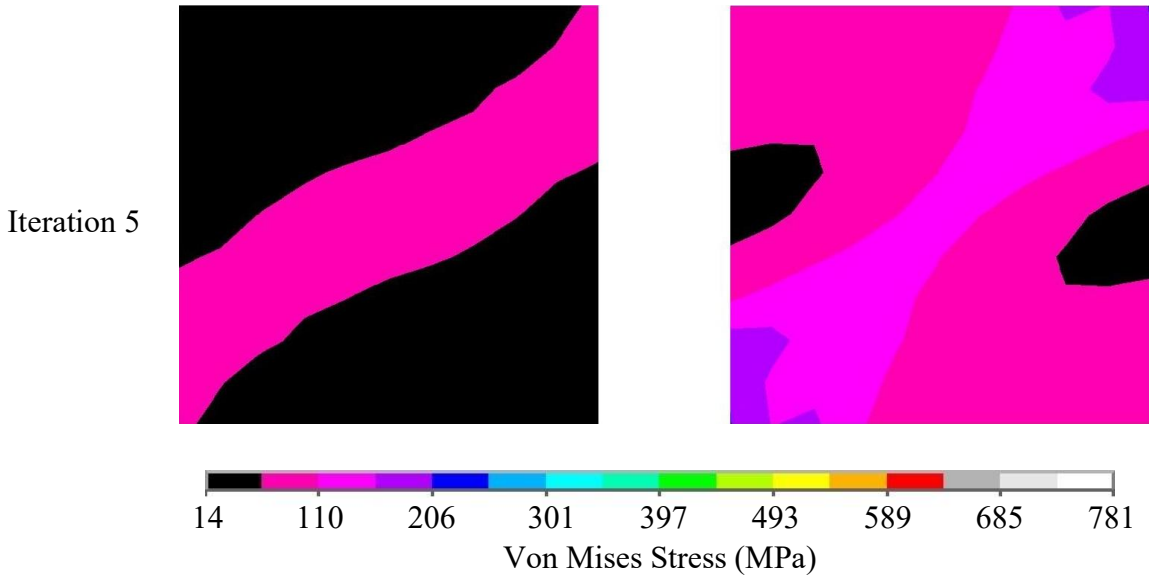


Figure 218 Comparison of requested and actual stress responses - Example 2

The trained CNN model showed the ability to recognise the major features of the requested stress contours, providing layup sequences with stress responses that match the overall shape and magnitude with good agreement. A strong level of responsiveness to new inputs was observed as the requests were updated and the designs reiterated, giving the user the ability to dictate the load paths and regions of high stress concentrations. The responses are unlikely to be a perfect match, the main limitation of the inverse design approach is that it may be difficult to choose a plausible request that is close to something physically achievable. However the network is likely to provide the nearest result to the input as long as the request is reasonable. The model successfully reduced the maximum stress by 75.3% for example 1 and 69.6% for example 2. The mean stress across the plate was reduced by 78.3% and 70.6% respectively. The study provides a proof of concept that may be expanded in future works to include more complex geometry and loading conditions for practical design of 3-dimensional components. Components featuring large numbers of finite elements are likely to require neural networks of greater depth and complexity, especially if the number of layup angle parameters is increased. The dimensionality of the convolution filters can be increased to use 2D convolution, formatting the stress input as a matrix of values. High performance computing clusters are likely to be needed for dataset generation and network training. Alternatively the level of complexity of the problem can be controlled by sampling a coarser grid of stress values. For example taking the stress in every other element may reduce the computational requirements while still providing the necessary amount of information to tune the overall response. The stress can also be probed selectively in critical locations. The current network may be applied to a 3D component by choosing 100 key points throughout the geometry following an initial finite element analysis to identify major stress concentrations. Furthermore for minimum stress design the network can be trained on a

smaller subset of the possible design space, training only on samples where the maximum stress is below a certain value. Limiting the size of the design space and the range of possible stress values is likely to improve regression further with negligible drawbacks as only the low stress region is of interest. The inverse CNN model offers the user design flexibility, giving control for fine-tuning the response surface and updating the request as necessary. However instead of manually tuning the stress distributions, optimisation via evolutionary approaches such as genetic algorithm and particle swarm optimisation may further improve the results when seeking the global minimal stress solution. These approaches are explored in the following sections.

7.2 Stress Optimisation Via Predictor Neural Networks

Instead of inputting the desired stress vector to the neural network and outputting the required layup, an alternative approach is to work in the opposite direction. The network is trained to predict the stress at every location within the composite for a given stacking sequence. Once trained it is able to predict the stresses considerably faster than FEA runs. Stress values are returned near instantaneously which allows the regression model to be used as the objective function in an evolutionary optimisation algorithm to minimise the stresses within the component. The proposed predictor network model is illustrated in Figure 219. The stacking sequence angles are now supplied as input to the network along with the x and y coordinates of the location where the stress is to be probed. These are mapped to the Von Mises stress value at the solution location. By requesting multiple locations the response is built up for the full FEA mesh and compared against known values during training. This approach considerably reduces the complexity of the network compared to the inverse design method, allowing for a greater number of inputs. The stacking sequences were increased to 8 angle parameters corresponding to an 8 ply layup. As there are now 10 inputs mapping to a single scalar output the problem can be solved using a shallow artificial neural network (ANN). This is faster and more practical to implement compared to deep learning as there is only a single hidden layer so less time is spent tuning network architecture. Details are provided below of the shallow ANN model. Then the concept is implemented using a deep neural network (DNN) and convolutional neural network (CNN) to compare the difference in the results.

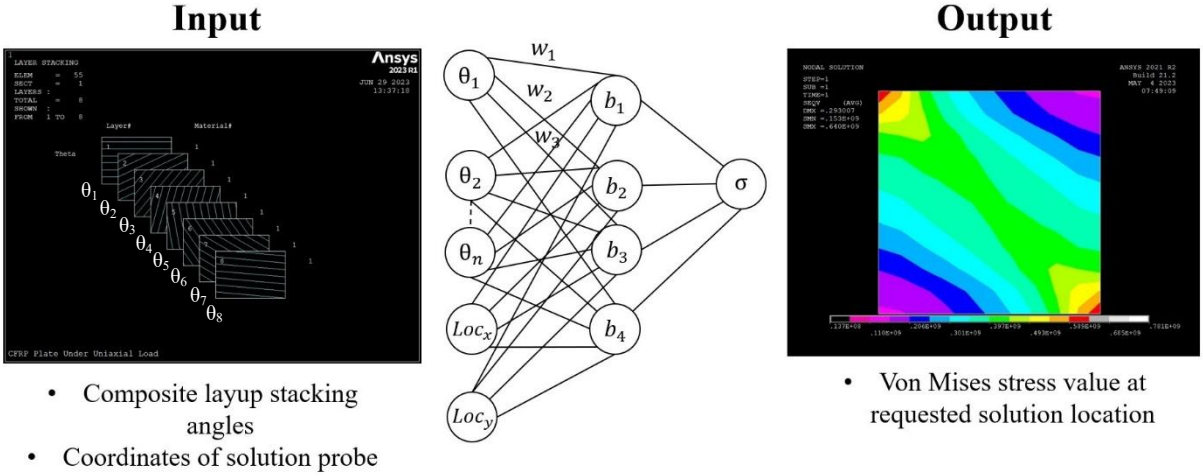


Figure 219 Schematic of predictor neural network model. Network architecture is figurative and simplified for illustrative purposes. Actual architectures used featured a greater number of neurons and hidden layers. Labels w and b represent weights and biases of network channels.

7.2.1 Shallow Artificial Neural Network Predictor Model

Dataset & Network Architecture

The ANSYS model from section 7.1.1 was used to generate a dataset from 2000 FEA runs with layups determined using a Latin hypercube design of experiments test plan. The angle parameters were able to vary between 0° and 180°. The pairwise plot of the design space featured in the dataset is shown in Figure 220. The ANSYS script was modified to extend the layup to 8 angles, using the same geometry, loading and mesh as previously. For each FEA simulation the Von Mises stress values were extracted from every element along with the coordinates of the centroid of the element. The stacking angles and element locations were formatted for input into the first layer of the ANN as an m x n matrix, where m is the number of parameters and n the number of observations. Instead of a vector of stress values, each output observation was now a single scalar value representing the stress in a single element of the FEA mesh. As each mesh contained 100 elements, the 2000 FEA runs generated a total of 200,000 observations.

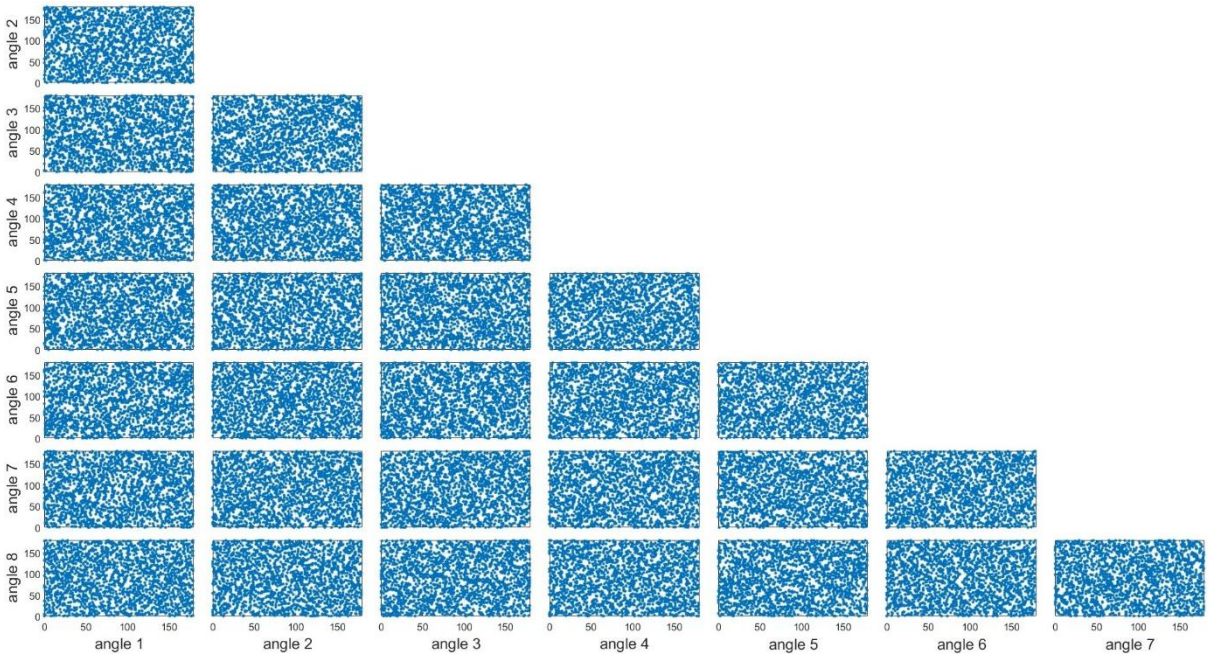


Figure 220 Training dataset design space - 2000 sample Latin hypercube test plan

The dataset was partitioned with 80% used for training and 20% used for testing. In this case a validation partition was not required as Bayesian regularisation was used as the training algorithm, which has its own form of validation built into the algorithm. It is designed to minimise the sum of squared errors between the network predictions and known values, as well as either the weighted sum of squared weights or weighted sum of absolute weight values. This helps prevent excessive numbers of weight parameters to aid generalisation and

avoid overfitting. As training progresses the weights typically become large to provide accurate predictions. Validation checks against known values from a separate data partition are used to prevent overfitting by stopping the training early if the validation error begins to increase for a certain number of epochs. Bayesian error depends not only on the performance of the network predictions, but also on how large the weights are. Conventional validation checks therefore may stop the training before the network is able to explore larger weights and prevent convergence on the global minimum. It is therefore recommended to disable validation stops and remove the validation partition. During Bayesian regularisation a backpropagation algorithm is used to calculate the Jacobian jX of the performance parameter with respect to the weight and bias variables, X [159]. In this case the performance was the mean squared error (MSE) between the predicted and actual stress values. The Levenberg-Marquardt algorithm is used to adjust the weights and biases of the channels between every neuron in the network according to:

$$jj = jX * jX \tag{7.1}$$

$$je = jX * E \tag{7.2}$$

$$dX = -(jj + I*mu) / je \tag{7.3}$$

where E is all errors and I is the identity matrix.

An adaptive variable, mu is increased incrementally until the updated weight and bias values result in a reduction in MSE. The training continues until a maximum number of epochs or a maximum time limit is reached, the error falls below a target value for a certain number of iterations, the performance gradient falls below a certain value, or the mu variable exceeds a preset limit. The process was implemented in MATLAB using the shallow artificial neural network architecture shown in Figure 221**Error! Reference source not found.**(a), featuring a single hidden layer with sigmoid activation connected to a linear output. The maximum number of epochs was set to 1000, the minimum performance gradient was $1x10^{-7}$ and the maximum value of mu $1x10^{10}$. The training was repeated to determine the number of neurons in the hidden layer necessary to provide stress predictions as close as possible to a target of 99% regression within a practical timeframe. The effects of the number of neurons on the mean squared error and regression are displayed in Figure 221(b), with the corresponding trade-off in increased training time highlighted by Figure 221(c). Error and regression values were calculated from the test partition of 40,000 observations that remained unseen by the network during training. The final chosen architecture featured 200 hidden layer neurons and was able to train in a single day over 7hours 51minutes 41 seconds on the RTX 3060 GPU.

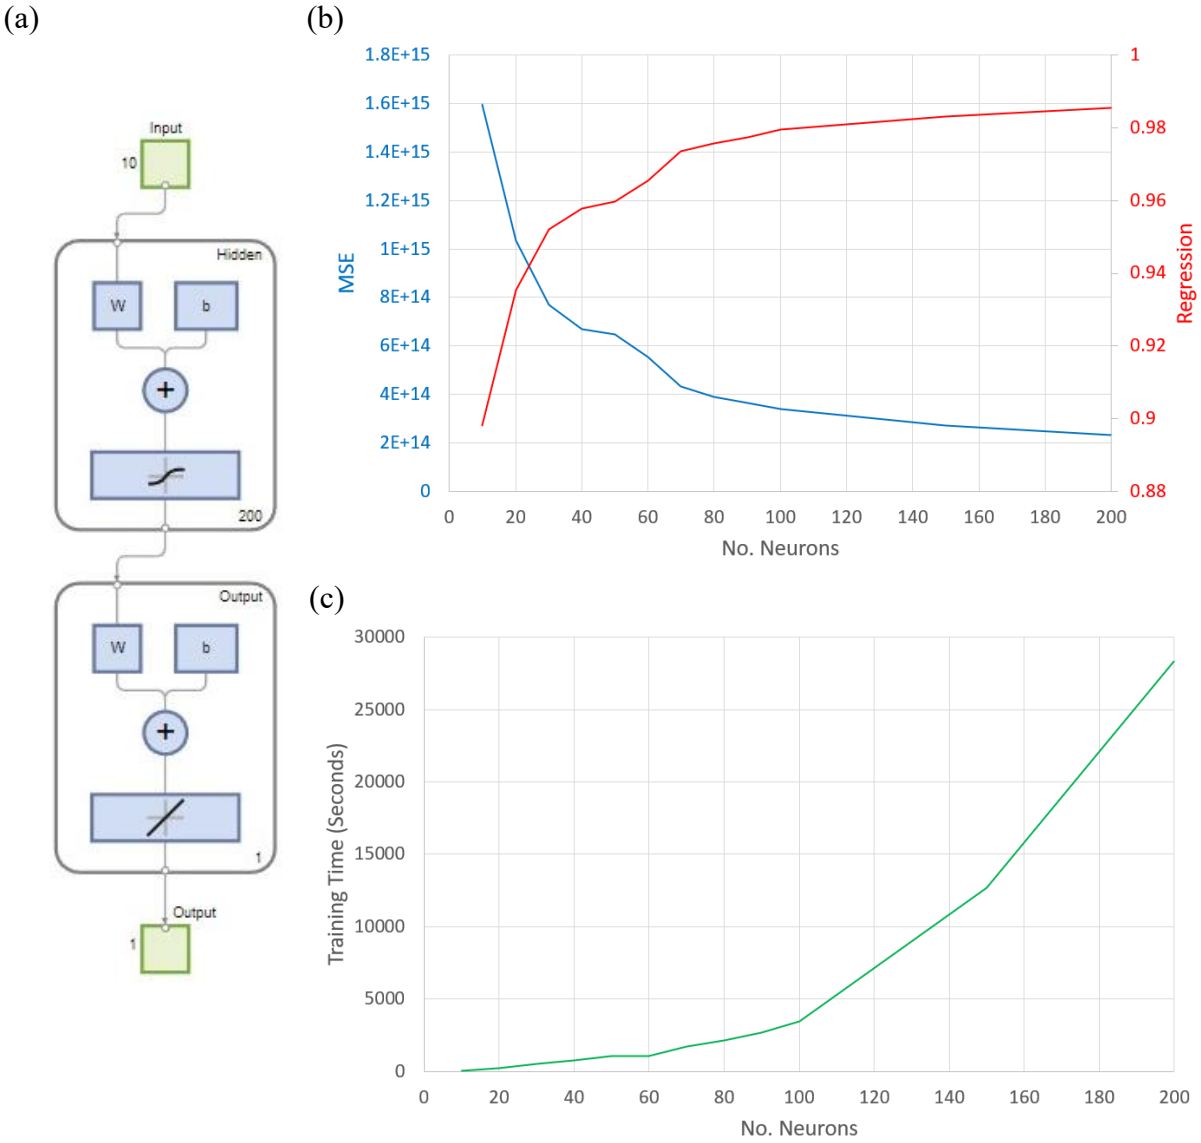


Figure 221 (a) Shallow artificial neural network architecture; (b) Test mean squared error and regression vs number of hidden layer neurons; (c) Training time vs number of hidden layer neurons

Network Performance

Details of the training progress and performance for the best performing ANN are summarised in Table 42. The training was stopped automatically as a result of reaching the maximum target number of epochs. Each epoch indicates the entire training dataset being passed through the network once. The state of the training parameters at every epoch is plotted in Figure 222. The number of parameters, gamK refers to the neural network's weight and bias vector. The sum of the squared parameters, ssX is the sum of the squared weight and bias vector. Both gamK and ssX increase as the performance of the network improves before levelling off after approximately 500 epochs and 800 epochs respectively, indicating less change in updating of the weights and biases towards the end of the training as the regression

model converges. This suggests additional training time is unlikely to significantly improve the results further. Figure 223 shows the mean squared error between the predicted and known stress values for the train and test data partitions throughout the training history. The test mean squared error closely follows the training error, indicating the network shows similar performance on new unseen data and generalises well with no evidence of overfitting observed.

Table 42 ANN training progress and performance summary

Network Architecture		**Training Progress**				**Training Performance**			
Type	ANN	**Unit**	**Initial Value**	**Stopped Value**	**Target Value**		**Observations**	**MSE**	**R**
		Epoch	0	1000	1000				
		Elapsed Time	-	07:51:44	-				
		Performance	1.81E+18	2.24E+14	0	**Training**	160000	2.243E+14	0.986
No. Hidden Layer Neurons	200	Gradient	3.56E+18	1.32E+14	1.00E-07				
Activation	Sigmoid	Mu	0.005	5.00E+01	1.00E+10				
Training Algorithm	Bayesian Regularisation	Effective # Param	2.40E+03	2.37E+03	0	**Test**	40000	2.311E+14	0.986
Performance Metric	Mean Squared Error	Sum Squared Param	1.58E+03	1.55E+03	0				

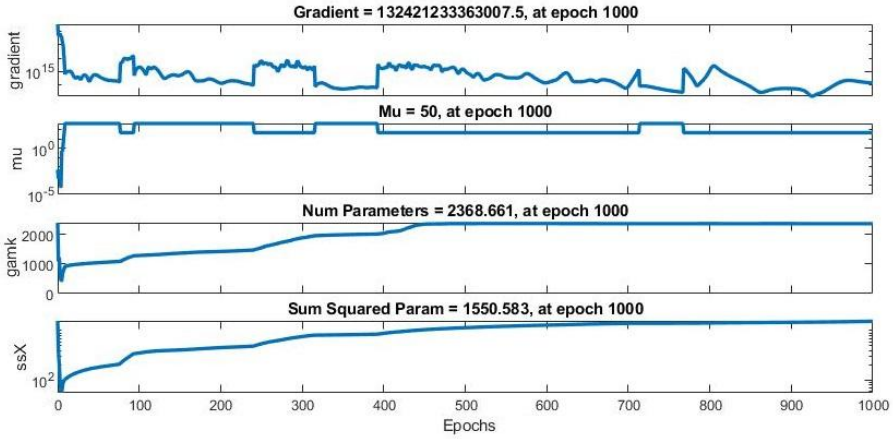


Figure 222 ANN training state

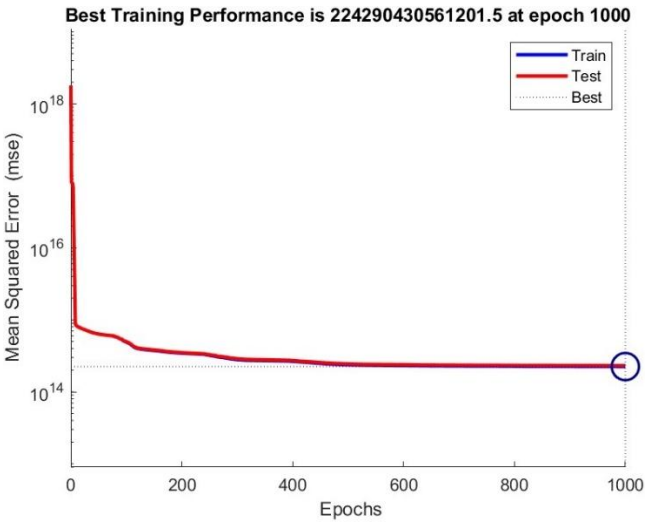


Figure 223 ANN training performance

Figure 224 plots the regression between the known target stress values and the predictions output by the trained ANN for the test dataset of 40,000 observations. Strong positive correlation is observed with few extreme outliers, which typically occur towards the higher end of the stress range. The coefficient of regression between the target and output stress values is 0.98556. Figure 225 shows the error between the known and predicted stress values for the entire dataset of 200,000 observations. The difference between the target and output value is negative more often than positive, indicating the network is more likely to overestimate than underestimate the stresses. The number of instances suggest an 80.9% chance that the error will be below 13.9MPa and a 96.8% chance the error will be below 31.8MPa when predicting stresses anywhere in the plate.

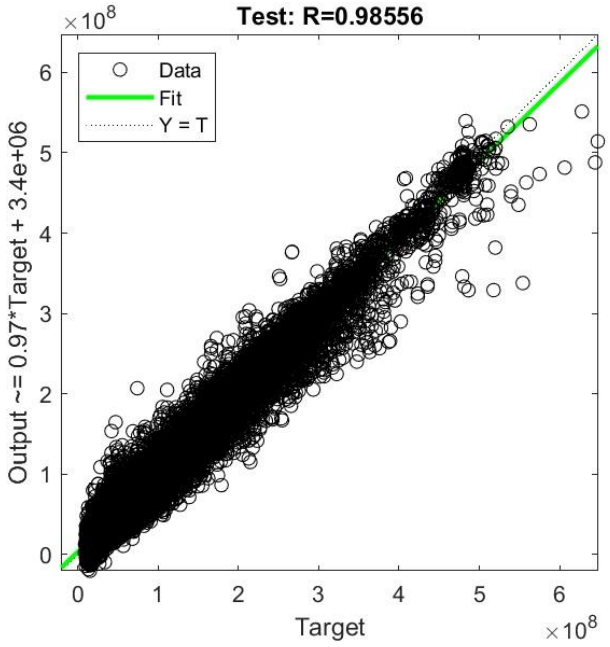


Figure 224 ANN test regression between known target stress values and output predictions

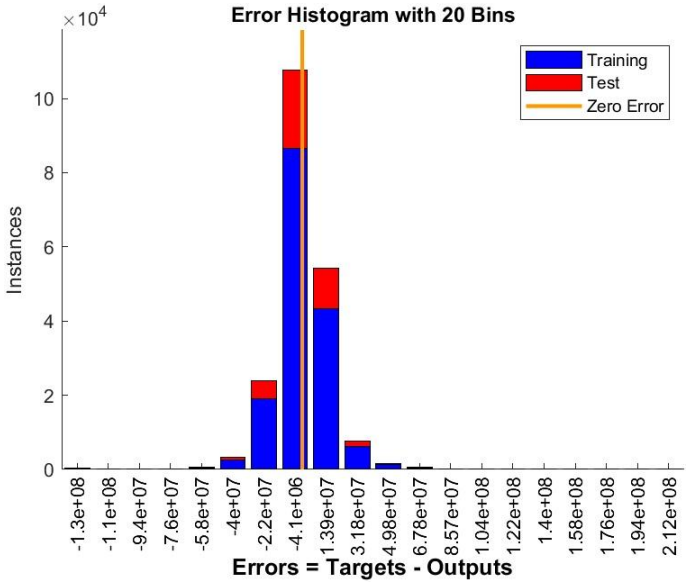


Figure 225 Error histogram for entire dataset of 200,000 observations

To further evaluate the location specific error at every element in the FEA mesh, an additional test dataset of 500 possible layup combinations was generated using a Latin hypercube test plan. Each layup sequence was supplied as input to the trained ANN and predictions were requested at every location corresponding to the 100 element mesh. The predicted stress values are compared against true values obtained by FEA in Figure 226, which shows the magnitude of the absolute error for every element of the mesh individually. The box and whisker for each element corresponds to 500 datapoints, where each datapoint represents a different layup. In all cases the median absolute error remains below 30MPa regardless of the location within the mesh. The upper quartile range limit exceeds 50MPa in only 4 elements. However the range of data signifies large errors may occur using the shallow network architecture. Outliers are observed suggesting in extreme cases the error may exceed 100MPa. Figure 227 shows the mean absolute error averaged across the plate for each of the 500 test layups. When averaged across the mesh the network provided predictions with a mean error of 28.7MPa, however the worst cases exceed 120MPa. The shallow network gave sufficient regression to indicate it is able to extract the critical features and identify the underlying patterns in the dataset. This may lead to a practical method for stress optimisation of composite components when combined with optimisers to minimise the maximum stress. However for the most accurate full stress analysis where a comparable alternative to FEA is required, deeper network models are likely to be necessary.

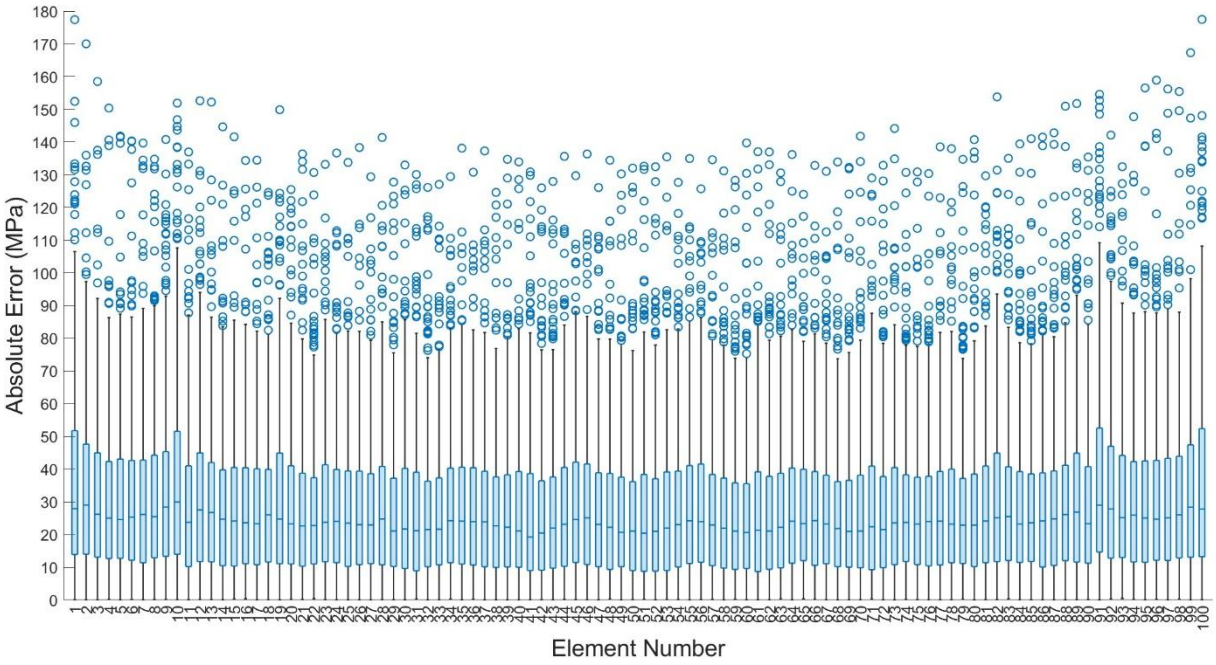


Figure 226 Absolute error between known stress values and values predicted by the ANN for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan.

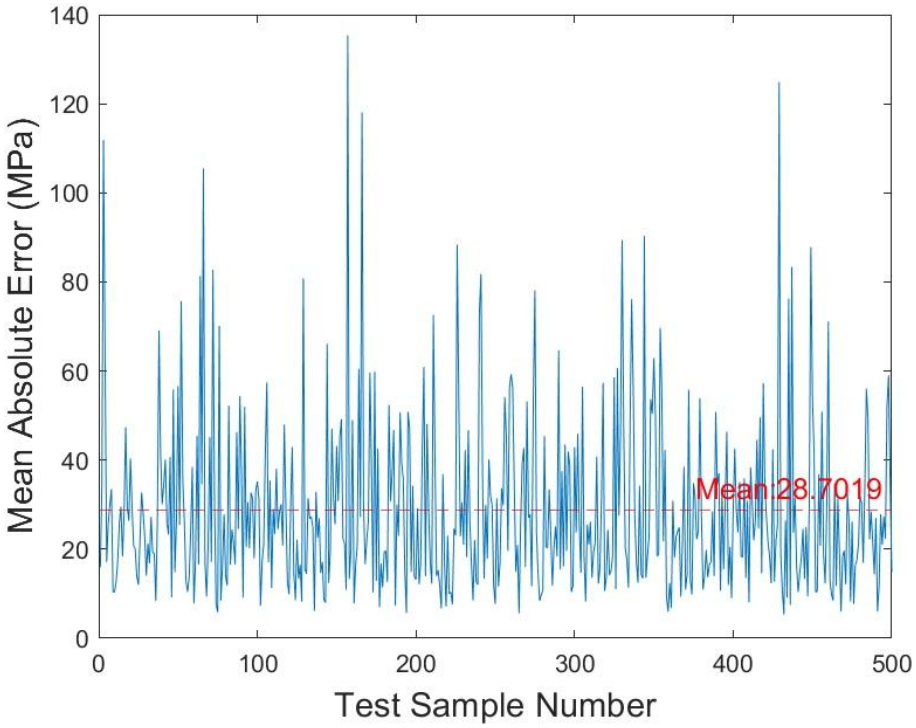


Figure 227 Mean absolute error averaged across the composite plate for 500 test layups – ANN

7.2.2 Deep Learning Network Predictor Models

Dataset & Network Architecture

For the deep learning models the dataset of 2000 samples from the Latin hypercube test plan displayed in Figure 220 was repartitioned to allow 70% for training, 15% for validation and 15% for testing. The angle and stress parameters were normalised, rescaling between 0 and 1. The minimum and maximum stress values in the dataset were recorded and used to scale the network predictions back again after training. Deep network architectures were developed using a DNN of entirely fully connected hidden layers and a CNN which featured a series of convolution layers before the fully connected layers. In both cases architectures were iterated using a process of trial and error, varying the numbers of neurons and convolution filters in each layer. Details of the iterations can be found in Appendix P. Multiple networks were trained for 50 epochs with a mini batch size of 1000, using the Adam optimiser to minimise the root mean square error performance metric. The stress predictions were evaluated using the test dataset and compared against values obtained via FEA. The architectures that provided the strongest regression coefficient between the predicted and known values were selected and trained further for 1000 epochs. The best performing DNN architecture is detailed in Table 43. For this case the data was formatted for feature input, requiring a matrix of normalised angles where each row is a different observation and each column a separate

parameter. The feature input was fed to a combination of fully connected layers, starting at 50 neurons and increasing by 50 each time until 200 before reducing in likewise manner. Each fully connected layer was immediately followed by batch-wise normalisation and then a ReLU activation, except for the final layer which featured a single neuron connected directly to a regression output to provide the stress predictions. The resulting network consisted of 24 layers and 102.9k total learnable parameters that were tuned during training.

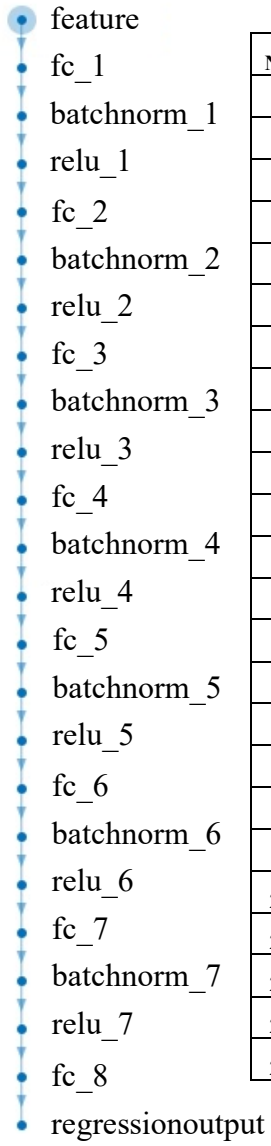


Table 43 Predictor network DNN architecture

No.	Name	Type	Activations	Learnable Parameters
1	feature	Feature Input	10 x 1	-
2	fc 1	50 Neuron Fully Connected	50 x 1	Weights 50 x 10, Bias 50 x 1
3	batchnorm 1	Batch Normalisation	50 x 1	Offset 50 x 1, Scale 50 x 1
4	relu 1	ReLU Activation	50 x 1	-
5	fc 2	100 Neuron Fully Connected	100 x 1	Weights 100 x 50, Bias 100 x 1
6	batchnorm 2	Batch Normalisation	100 x 1	Offset 100 x 1, Scale 100 x 1
7	relu 2	ReLU Activation	100 x 1	-
8	fc 3	150 Neuron Fully Connected	150 x 1	Weights 150 x 100, Bias 150 x 1
9	batchnorm 3	Batch Normalisation	150 x 1	Offset 150 x 1, Scale 150 x 1
10	relu 3	ReLU Activation	150 x 1	-
11	fc 4	200 Neuron Fully Connected	200 x 1	Weights 200 x 150, Bias 200 x 1
12	batchnorm 4	Batch Normalisation	200 x 1	Offset 200 x 1, Scale 200 x 1
13	relu 4	ReLU Activation	200 x 1	-
14	fc 5	150 Neuron Fully Connected	150 x 1	Weights 150 x 200, Bias 150 x 1
15	batchnorm 5	Batch Normalisation	150 x 1	Offset 150 x 1, Scale 150 x 1
16	relu 5	ReLU Activation	150 x 1	-
17	fc 6	100 Neuron Fully Connected	100 x 1	Weights 100 x 150, Bias 100 x 1
18	batchnorm 6	Batch Normalisation	100 x 1	Offset 100 x 1, Scale 100 x 1
19	relu 6	ReLU Activation	100 x 1	-
20	fc 7	50 Neuron Fully Connected	50 x 1	Weights 50 x 100, Bias 50 x 1
21	batchnorm 7	Batch Normalisation	50 x 1	Offset 50 x 1, Scale 50 x 1
22	relu 7	ReLU Activation	50 x 1	-
23	fc 8	1 Neuron Fully Connected	1 x 1	Weights 1 x 50, Bias 1 x 1
24	regressionoutput	Regression Output	1 x 1	-

To train the CNN the dataset was reformatted for sequence input to allow for 1D convolution layers. The input was a cell array of sequence vectors containing the set of parameter values for each layup and location. The CNN architecture featured 43 layers with 276.7k learnables. The arrangement of layers is displayed in Figure 228 and summarised in Table 44. The sequence input was passed through a set of 5 convolutional layers. Each layer featured filters of size 3 and stride 1 with no padding and was immediately followed by a ReLU activation and layer-wise normalisation. After the convolution layers the extracted feature maps were flattened into a single vector using global average pooling and sent through a fully connected network of 9 dense hidden layers. The first had 512 neurons and then the number of neurons was halved with each successive layer until the final layer, which had a single activation followed by a regression output.

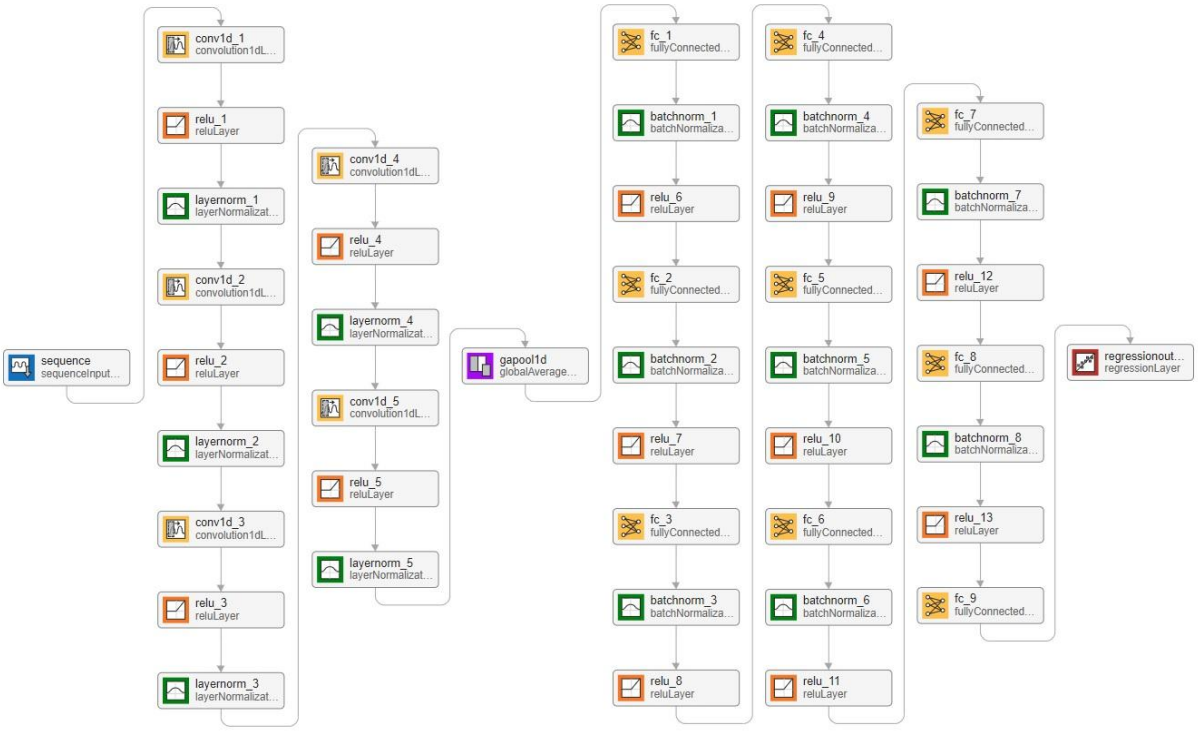


Figure 228 Predictor network CNN schematic layout

Table 44 Predictor network CNN architecture

No.	Name	Type	Activations	Learnable Parameters
1	sequence	Sequence Input	1 x 1 x 1	-
2	conv1d 1	8 Filter 1-D Convolution	8 x 1 x 1	Weights 3 x 1 x 8, Bias 1 x 8
3	relu 1	ReLU Activation	8 x 1 x 1	-
4	layernorm 1	Layer Normalisation	8 x 1 x 1	Offset 8 x 1, Scale 8 x 1
5	conv1d 2	16 Filter 1-D Convolution	16 x 1 x 1	Weights 3 x 8 x 16, Bias 1 x 16
6	relu 2	ReLU Activation	16 x 1 x 1	-
7	layernorm 2	Layer Normalisation	16 x 1 x 1	Offset 16 x 1, Scale 16 x 1
8	conv1d 3	32 Filter 1-D Convolution	32 x 1 x 1	Weights 3 x 16 x 32, Bias 1 x 32
9	relu 3	ReLU Activation	32 x 1 x 1	-
10	layernorm 3	Layer Normalisation	32 x 1 x 1	Offset 32 x 1, Scale 32 x 1
11	conv1d 4	64 Filter 1-D Convolution	64 x 1 x 1	Weights 3 x 32 x 64, Bias 1 x 64
12	relu 4	ReLU Activation	64 x 1 x 1	-
13	layernorm 4	Layer Normalisation	64 x 1 x 1	Offset 64 x 1, Scale 64 x 1
14	conv1d 5	128 Filter 1-D Convolution	128 x 1 x 1	Weights 3 x 64 x 128, Bias 1 x 128
15	relu 5	ReLU Activation	128 x 1 x 1	-
16	layernorm 5	Layer Normalisation	128 x 1 x 1	Offset 128 x 1, Scale 128 x 1
17	gapool1d	1-D Global Average Pooling	128 x 1 x 1	-
18	fc 1	512 Neuron Fully Connected	512 x 1	Weights 512 x 128, Bias 512 x 1
19	batchnorm 1	Batch Normalisation	512 x 1	Offset 512 x 1, Scale 512 x 1
20	relu 6	ReLU Activation	512 x 1	-
21	fc 2	256 Neuron Fully Connected	256 x 1	Weights 256 x 512, Bias 256 x 1
22	batchnorm 2	Batch Normalisation	256 x 1	Offset 256 x 1, Scale 256 x 1
23	relu 7	ReLU Activation	256 x 1	-
24	fc 3	128 Neuron Fully Connected	128 x 1	Weights 128 x 256, Bias 128 x 1
25	batchnorm 3	Batch Normalisation	128 x 1	Offset 128 x 1, Scale 128 x 1
26	relu 8	ReLU Activation	128 x 1	-
27	fc 4	64 Neuron Fully Connected	64 x 1	Weights 64 x 128, Bias 64 x 1
28	batchnorm 4	Batch Normalisation	64 x 1	Offset 64 x 1, Scale 64 x 1
29	relu 9	ReLU Activation	64 x 1	-
30	fc 5	32 Neuron Fully Connected	32 x 1	Weights 32 x 64, Bias 32 x 1
31	batchnorm 5	Batch Normalisation	32 x 1	Offset 32 x 1, Scale 32 x 1
32	relu 10	ReLU Activation	32 x 1	-
33	fc 6	16 Neuron Fully Connected	16 x 1	Weights 16 x 32, Bias 16 x 1
34	batchnorm 6	Batch Normalisation	16 x 1	Offset 16 x 1, Scale 16 x 1
35	relu 11	ReLU Activation	16 x 1	-
36	fc 7	8 Neuron Fully Connected	8 x 1	Weights 8 x 16, Bias 8 x 1
37	batchnorm 7	Batch Normalisation	8 x 1	Offset 8 x 1, Scale 8 x 1
38	relu 12	ReLU Activation	8 x 1	-
39	fc 8	4 Neuron Fully Connected	4 x 1	Weights 4 x 8, Bias 4 x 1
40	batchnorm 8	Batch Normalisation	4 x 1	Offset 4 x 1, Scale 4 x 1
41	relu 13	ReLU Activation	4 x 1	-
42	fc 9	1 Neuron Fully Connected	1 x 1	Weights 1 x 4, Bias 1 x 1
43	regressionoutput	Regression Output	1 x 1	-

Network Performance

The root mean square error, loss and regression metrics for the chosen network architectures are summarised in Table 45. The results indicate a performance advantage exists when convolutional layers are applied. The best performing DNN offered 98.9% regression after training for 1000 epochs in 1hr 41min 44sec. The CNN produced comparable RMSE and regression statistics after only 5% of the training epochs, giving 98.9% regression in only 13min 6sec. On running for the full 1000 epochs over 4hr 25min 29sec the regression improved beyond 99% though less agreement is observed between the final training and validation RSME and loss. The training histories displayed on page 326 show no clear evidence of overfitting, however the final values suggest the CNN may be at risk of overfitting if the training were to be extended further. On testing the stress predictions against known values for the partition of 30,000 observations that remained unseen by the network during training, strong positive correlation is observed. Figure 231 indicates the network is sufficiently well generalised to provide high levels of regression when presented with new data. To further evaluate the accuracy of the network the error in the stress values obtained at the locations of every element were calculated using the same 500 layups used to test the ANN. In each case the mesh contained 100 elements, leading to 50,000 additional stress predictions in total. The error between the values obtained via FEA and the predictions from the best performing DNN and CNN are shown in Figure 232 and Figure 233 respectively.

Table 45 Summary of training performance for chosen stress predictor networks

Test Number	Architecture		Training Options		Training Performance			Regression	
DNN-10	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.01788	Parameter	R Value
			Mini Batch Size	1000		Validation	0.01913		
			No. Epochs	1000	Loss	Training	0.00016		
	No. Hidden Layer Neurons	[50, 100, 150, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	0.00018	Von Mises Stress (Pa)	0.98948
	Activation	ReLU	Validation Frequency	140	Training Time	1hr 41min 44sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	7000			
CNN-5	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0188	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0191		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00018		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00018	Von Mises Stress (Pa)	0.98907
	Activation	ReLU	Validation Frequency	140	Training Time	13min 6sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	7000			
CNN-12	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0192	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0141		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.00019		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00010	Von Mises Stress (Pa)	0.99390
	Activation	ReLU	Validation Frequency	140	Training Time	4hr 25min 29sec			
	Normalisation	Yes	Shuffle	Every Epoch	No. Iterations	140000			

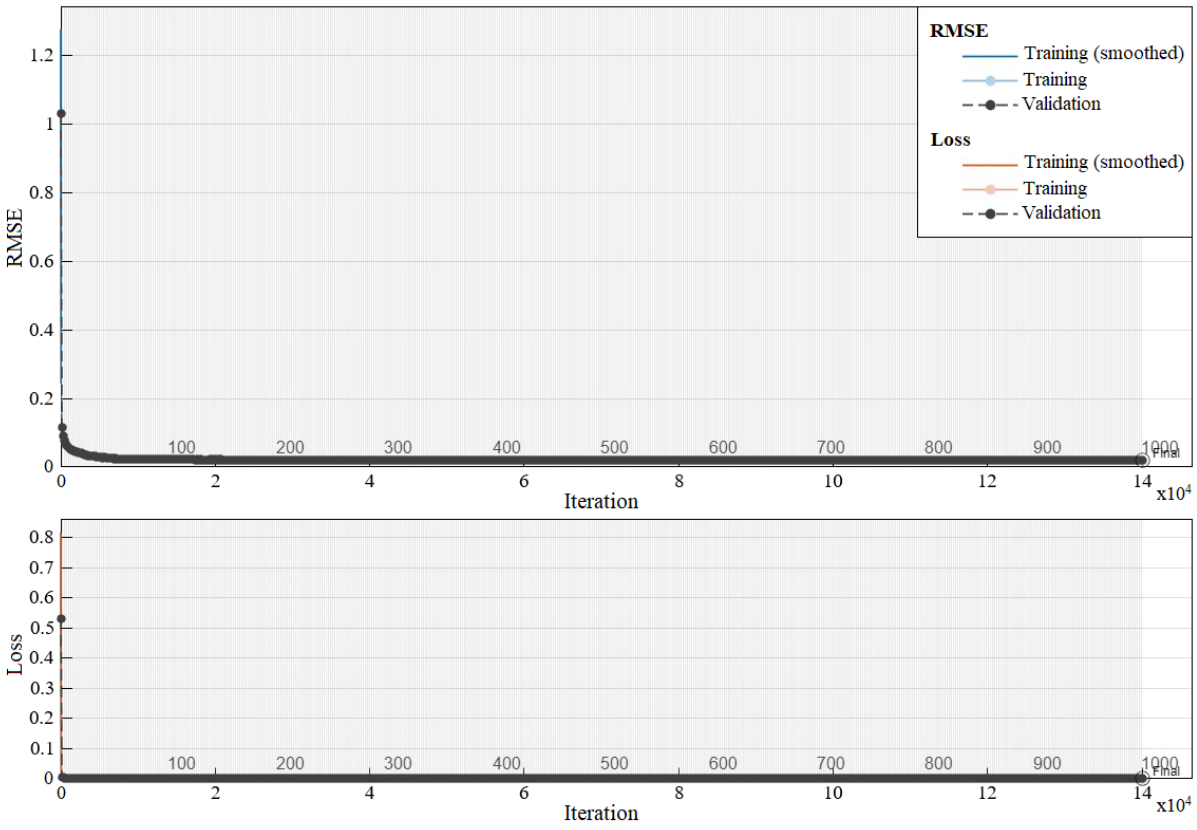


Figure 229 Training history of network DNN-10

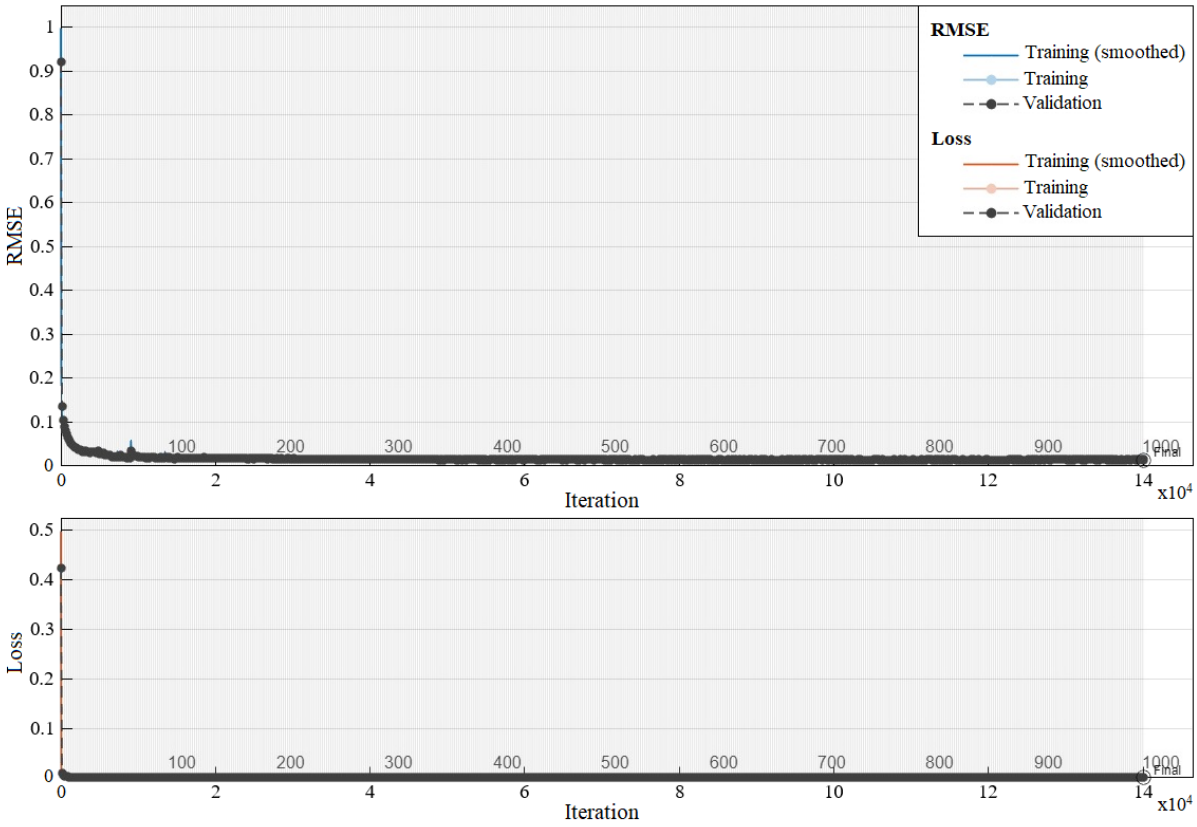
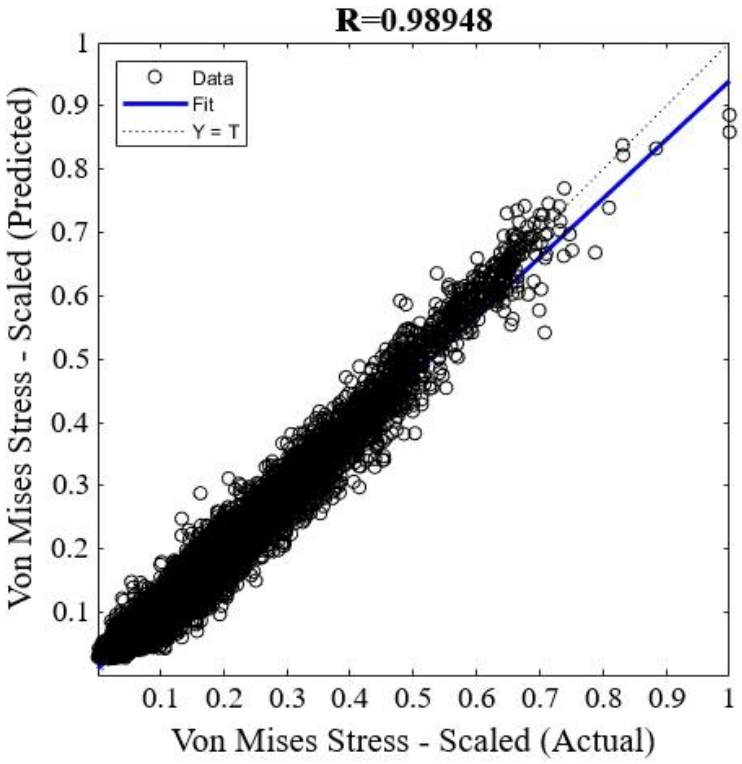


Figure 230 Training history of network CNN-12

(a)



(b)

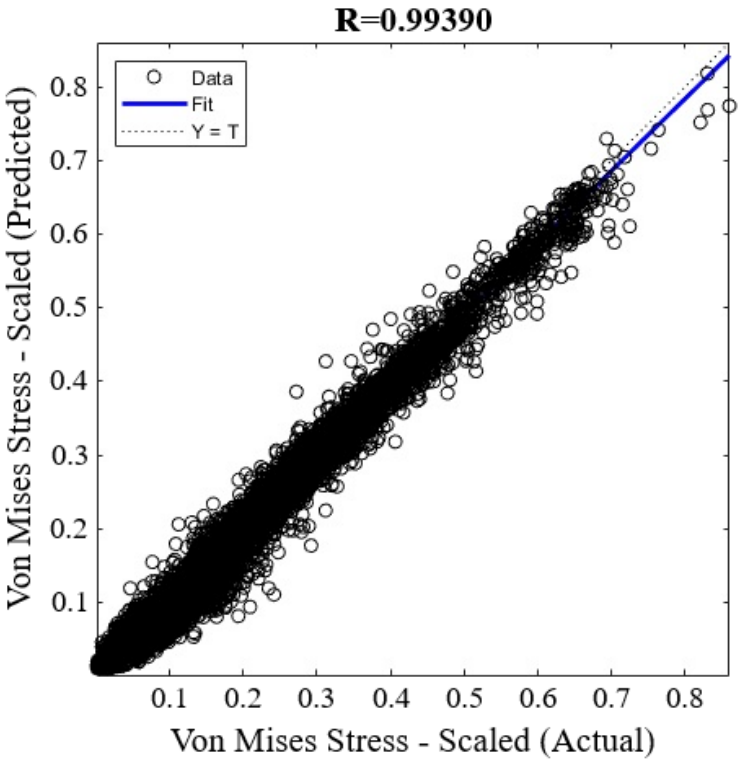


Figure 231 Regression plots between predicted and actual normalised stress values for networks (a) DNN-10 and (b) CNN-12. Each plot features 30,000 datapoints generated from the test partition of the dataset.

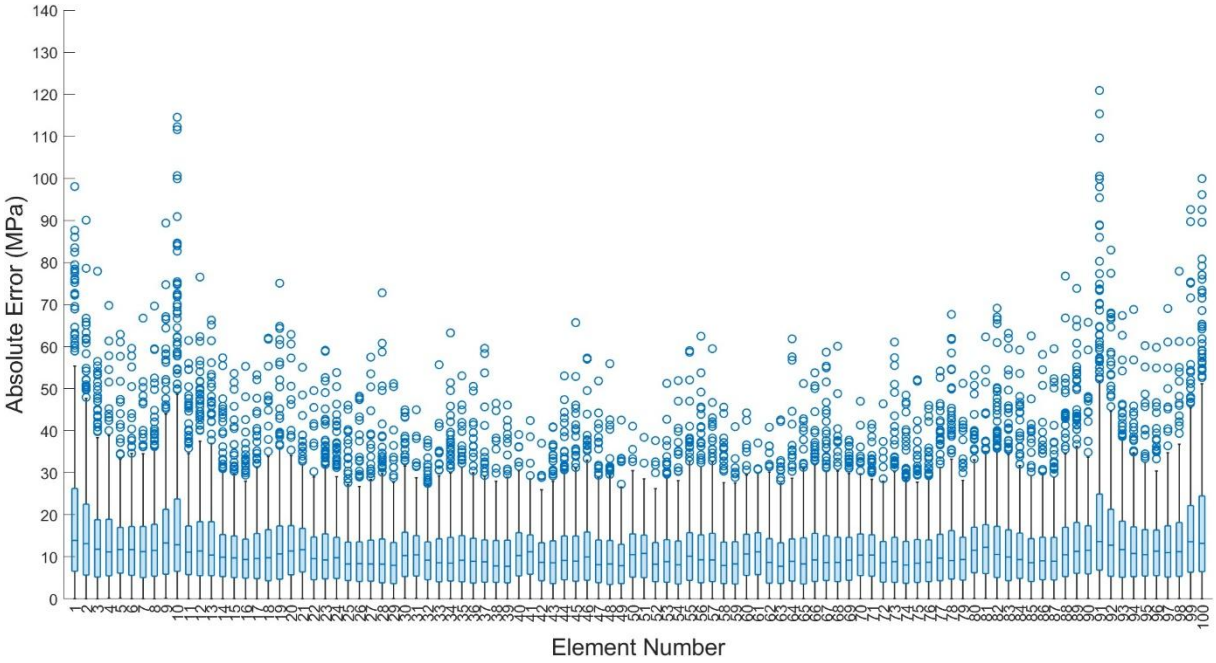


Figure 232 Absolute error between known stress values and values predicted by DNN-10 for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan.

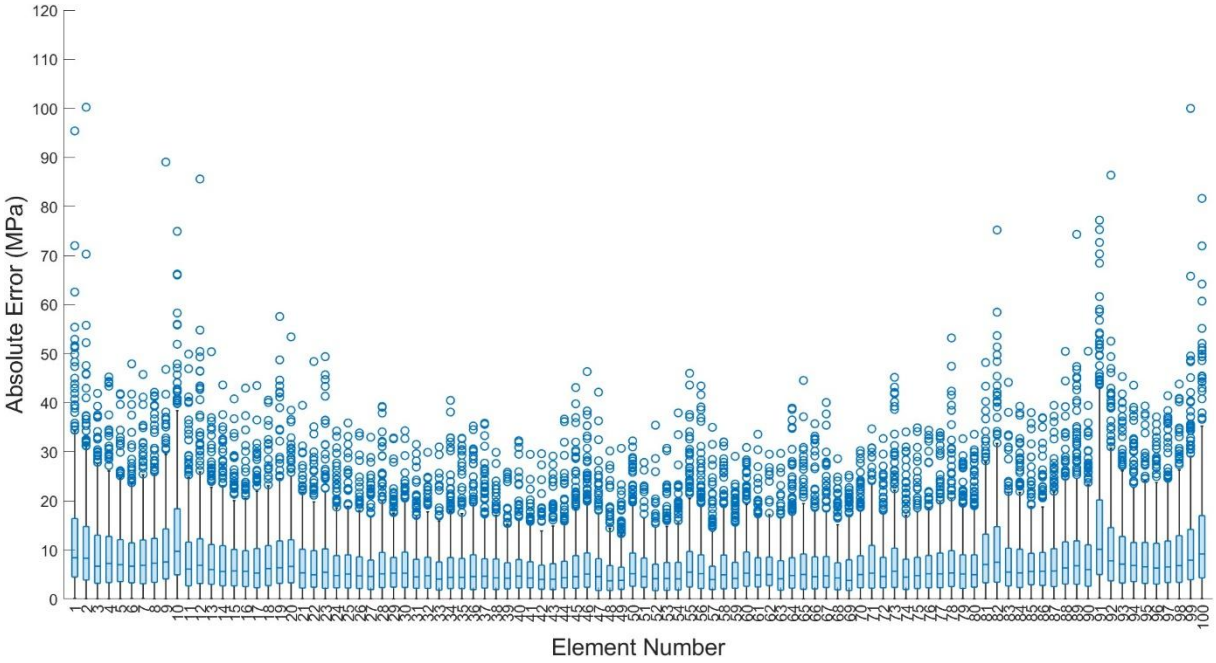
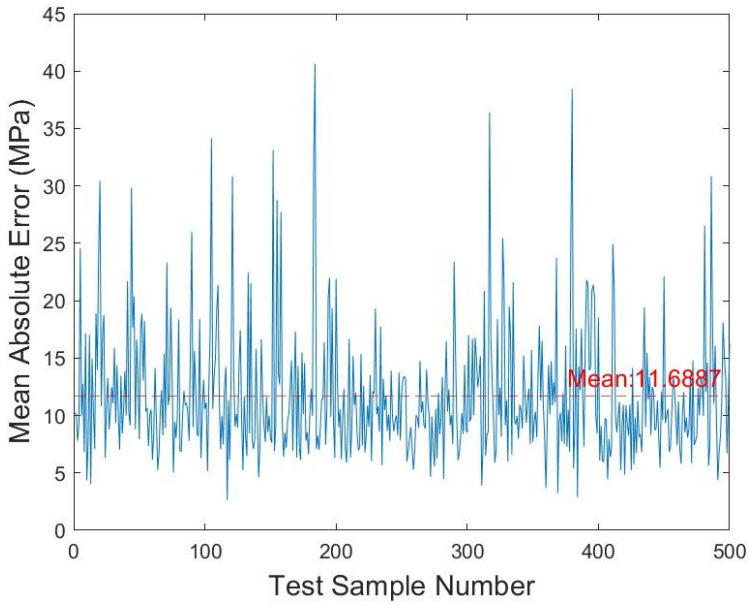


Figure 233 Absolute error between known stress values and values predicted by CNN-12 for every element in the FEA mesh. Data collected from 500 test layups defined from a Latin hypercube test plan.

When compared to the equivalent ANN result in Figure 226 significant improvements are observed in the both the magnitude of the error and the spread of the data as a result of the additional network depth. The CNN trained for 1000 epochs provided the best predictions. The median error remained below 8MPa across all elements and the upper quartile limit was less than 15MPa in all but 4 elements. Figure 234 shows that the error averaged across the plate for the 500 test layups has mean of 11.69MPa for the DNN. The CNN architecture predicted the stresses in the composite to a mean of 7.29MPa, within 1% of the maximum value that occurred in the original training dataset. In addition, the mean error remained below 30MPa for all 500 layups.

(a)



(b)

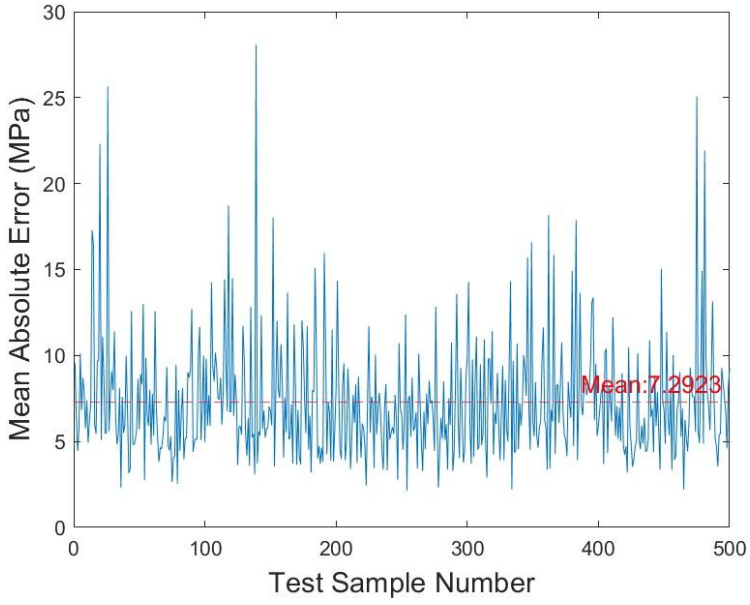


Figure 234 Mean absolute error averaged across the composite plate for 500 test layups (a) DNN-10 and (b) CNN-12

A possible source of error in the model may come from anomalies leading to artificially high stress values in the FEA mesh close to the boundary conditions and the locations where the loading is applied. The error analysis was repeated discarding the elements at the edges of the mesh. The centroid locations of the remaining 64 elements where the stress was probed are highlighted in Figure 235.

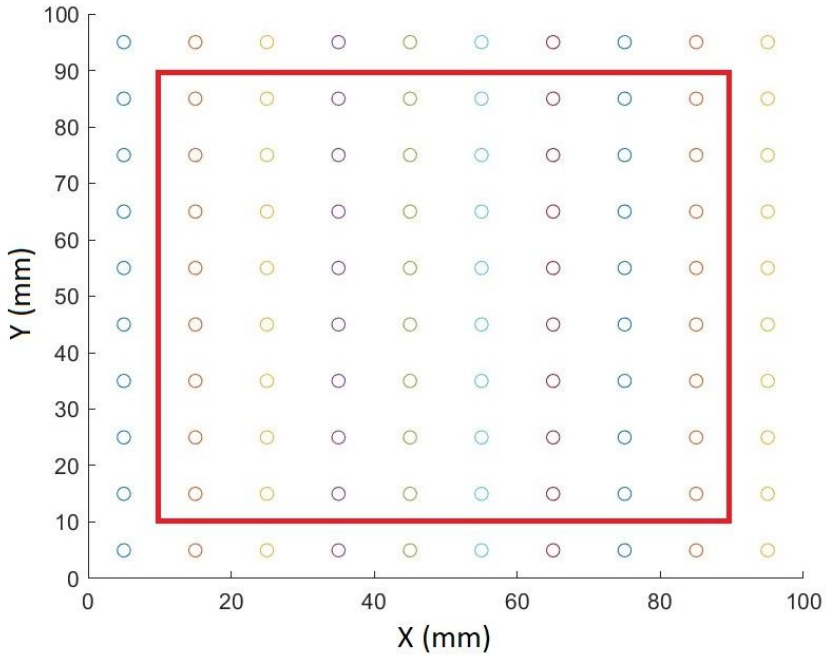


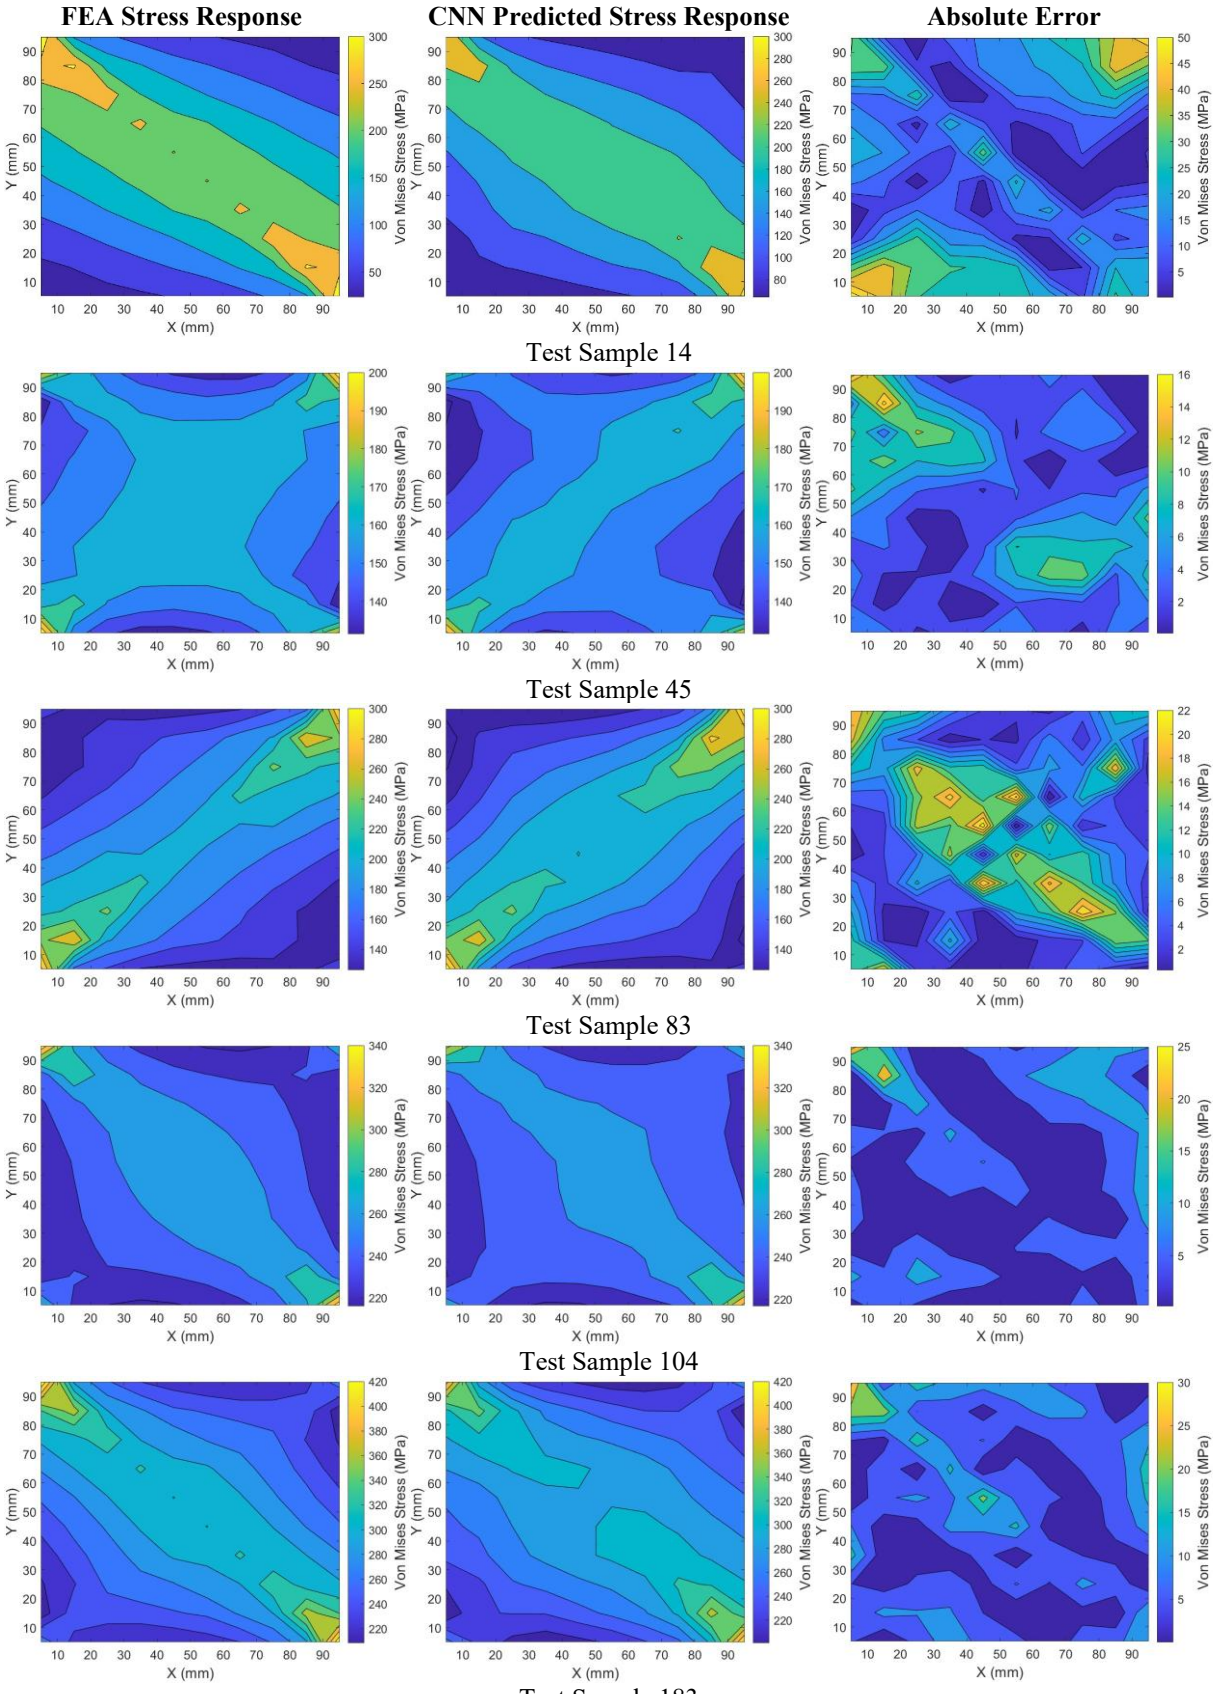
Figure 235 Reduced mesh element centroid locations

Table 46 summarises the mean error for all network architectures with and without the edges excluded. In all cases the error is improved when the edges are removed, suggesting the predictions may be more accurate away from the boundaries. The absolute error is plotted alongside comparisons of the FEA and CNN predicted stress contours for a selection of the test layups in Figure 236. Additional examples are provided in Appendix Q. The neural network is able to capture the shape and layout of the stress distributions with overall agreement to the FEA model. Regions of the highest stress concentrations are approximated at locations that are near to the correct solution and load paths are well represented. The network is unlikely to give a perfect recreation of the stress contours. Some degree of error must be expected and taken into account. The target of the model may not be to provide faultless predictions, but to provide predictions with errors that consistently fall below acceptable limits so that the overall goal of the model can be achieved. The convolutional predictor network shows accuracy that is likely to be sufficient for many practical design purposes. The network returns the stress predictions near instantaneously. The speed of solution allows the regression model to be applied as the objective function in evolutionary algorithms for optimisation of layups of composite components for stress reduction.

Table 46 Composite stress predictor network models mean error comparison summary

	Mean Error Original Mesh (MPa)			
	ANN 7hr 51min 44sec	**DNN 1hr 41min 44sec**	**CNN 13min 6sec**	**CNN 4hr 25min 29sec**
Mean	28.70	11.69	10.69	7.29
Standard Deviation	19.49	5.54	4.87	3.46
Max	135.38	40.64	32.22	28.06
Min	5.32	2.63	2.28	2.15
Range	130.07	38.01	29.94	25.91

	Mean Error Edges Excluded (MPa)			
	ANN 7hr 51min 44sec	**DNN 1hr 41min 44sec**	**CNN 13min 6sec**	**CNN 4hr 25min 29sec**
Mean	27.97	10.74	9.99	6.53
Standard Deviation	20.27	6.12	5.41	3.72
Max	134.25	40.54	35.29	30.86
Min	3.45	1.94	1.57	1.26
Range	130.80	38.59	33.72	29.60



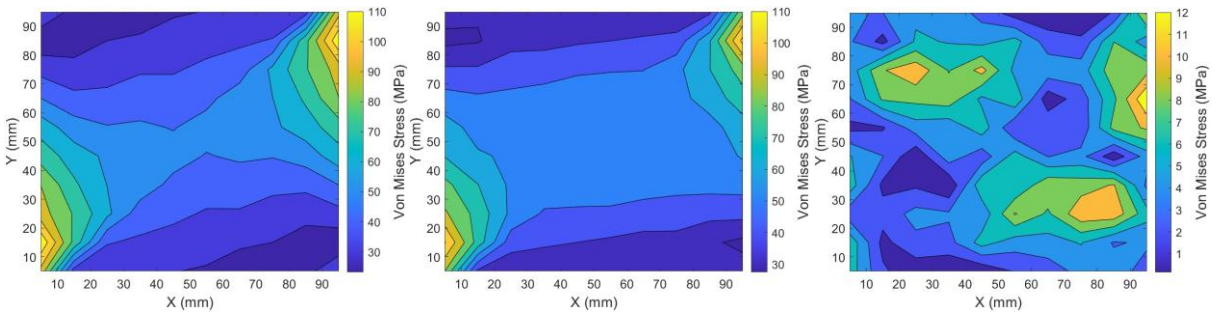
Test Sample 14

Test Sample 45

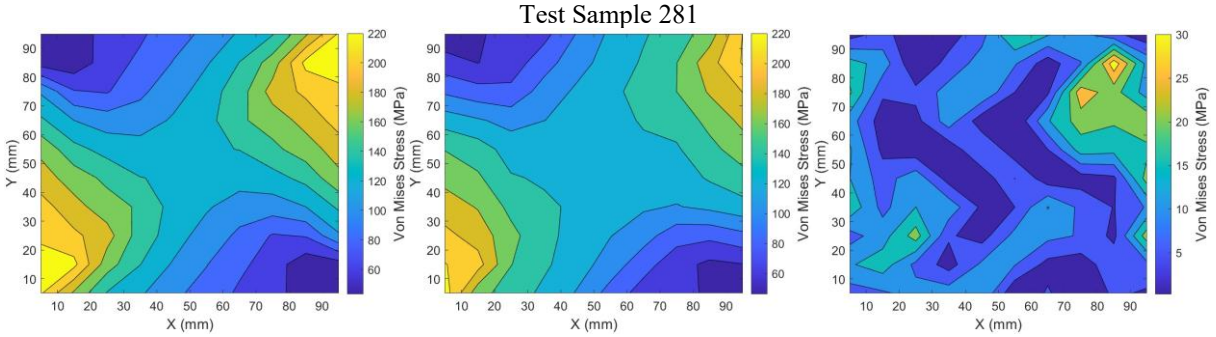
Test Sample 83

Test Sample 104

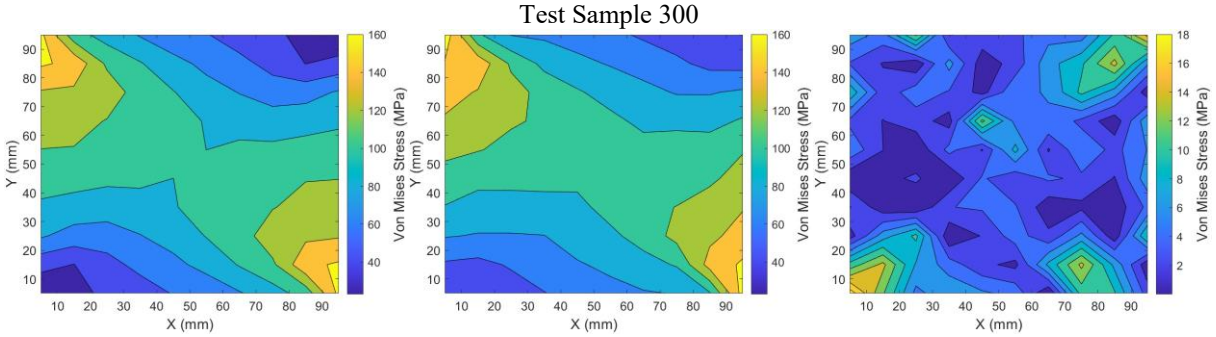
Test Sample 183



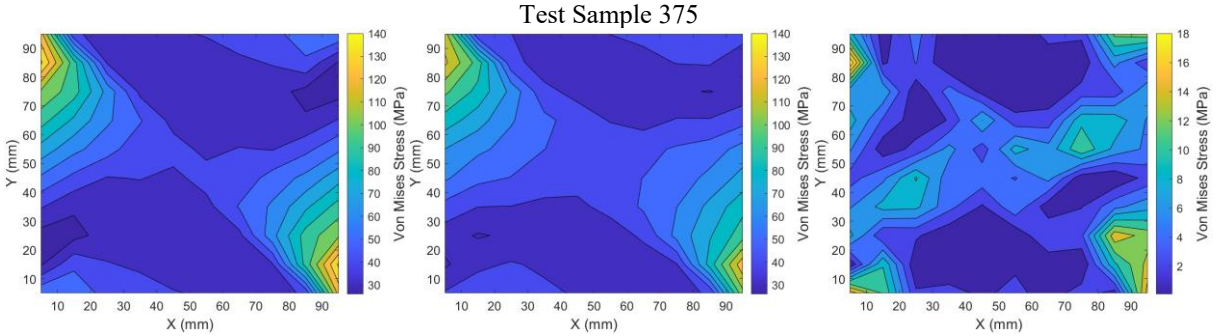
Test Sample 281



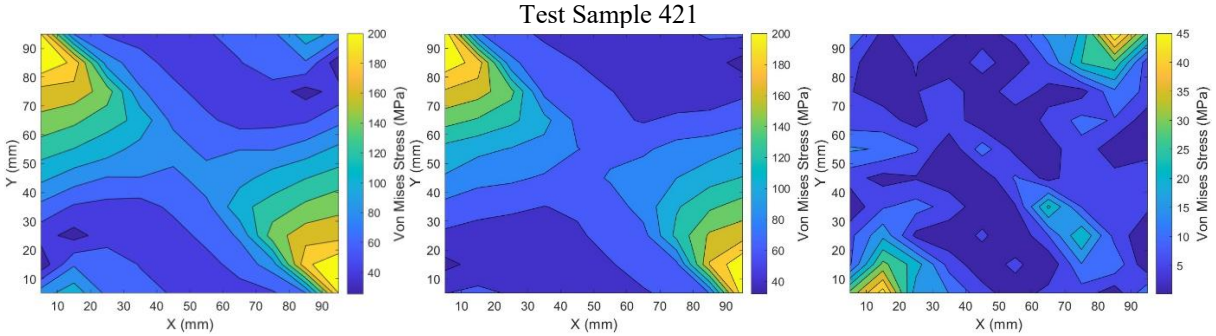
Test Sample 300



Test Sample 375



Test Sample 421



Test Sample 494

Figure 236 Comparison between FEA and CNN predicted composite stress contours. Test sample numbers correspond to those in Figure 234.

7.2.3 Maximum Stress Optimisation

In order to generate a model able to output the layup stacking sequence of a composite for minimum stress design, the predictor network was modified to output the maximum stress for a given layup. The network was then combined with evolutionary optimisation algorithms to minimise the maximum stress that occurs anywhere in the plate. The dataset of 2000 FEA simulations was reformatted. For each design point the maximum Von Mises stress was extracted. The location coordinate parameters were removed, leaving only the stacking angles and reducing the network to 8 inputs mapped to a single output. As now only the maximum stress was used rather than the values in each of the 100 elements, 2000 FEA runs lead to 2000 observations. An additional dataset was generated from 200,000 FEA simulations to match the size of the previous dataset. The training performance is compared in Table 47, using the best performing CNN architecture from section 7.2.2. The training history and regression obtained from the CNN trained on the dataset of 200,000 samples is displayed on page 335. Equivalent results for the ANN and DNN architectures are included in Appendix R.

Table 47 Maximum stress predictor network training performance

Dataset		Architecture		Training Options		Training Performance			Regression	
No. FEA Samples	2000	**Type**	CNN	**Training Algorithm**	Adam	**RMSE**	**Training**	0.0233		
Total Observations	2000	**Convolution Filter Size**	3	**Mini Batch Size**	100		**Validation**	0.0564	**Parameter**	**R Value**
Sampling Method	Latin Hypercube	**No. Filters**	[8, 16, 32, 64, 128]	**No. Epochs**	1000	**Loss**	**Training**	0.00027		
Train (%)	70	**No. Dense Layer Neurons**	[512, 256, 128, 64, 32, 16, 8, 4, 1]	**Initial Learning Rate**	0.001		**Validation**	0.00159		
Validation (%)	15	**Activation**	ReLU	**Validation Frequency**	14	**Training Time**	23min 39sec		**Von Mises Stress (Pa)**	0.9155
Test (%)	15	**Normalisation**	Yes	**Shuffle**	Every Epoch	**No. Iterations**	14000			
No. FEA Samples	200000	**Type**	CNN	**Training Algorithm**	Adam	**RMSE**	**Training**	0.0146		
Total Observations	200000	**Convolution Filter Size**	3	**Mini Batch Size**	1000		**Validation**	0.0137	**Parameter**	**R Value**
Sampling Method	Latin Hypercube	**No. Filters**	[8, 16, 32, 64, 128]	**No. Epochs**	1000	**Loss**	**Training**	0.00011		
Train (%)	70	**No. Dense Layer Neurons**	[512, 256, 128, 64, 32, 16, 8, 4, 1]	**Initial Learning Rate**	0.001		**Validation**	0.00009		
Validation (%)	15	**Activation**	ReLU	**Validation Frequency**	140	**Training Time**	4hr 23min 17sec		**Von Mises Stress (Pa)**	0.9956
Test (%)	15	**Normalisation**	Yes	**Shuffle**	Every Epoch	**No. Iterations**	140000			

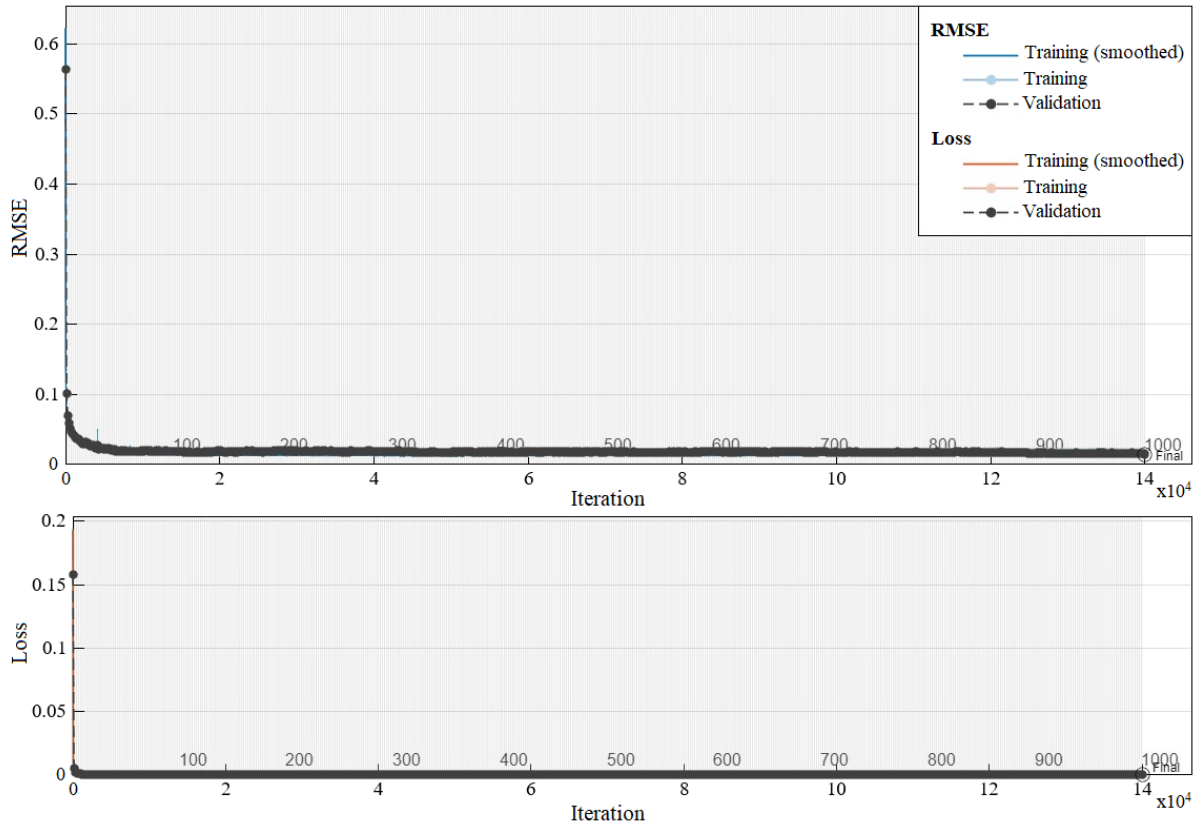


Figure 237 Training history of best performing maximum stress predictor CNN, trained on a dataset of 200,000 observations

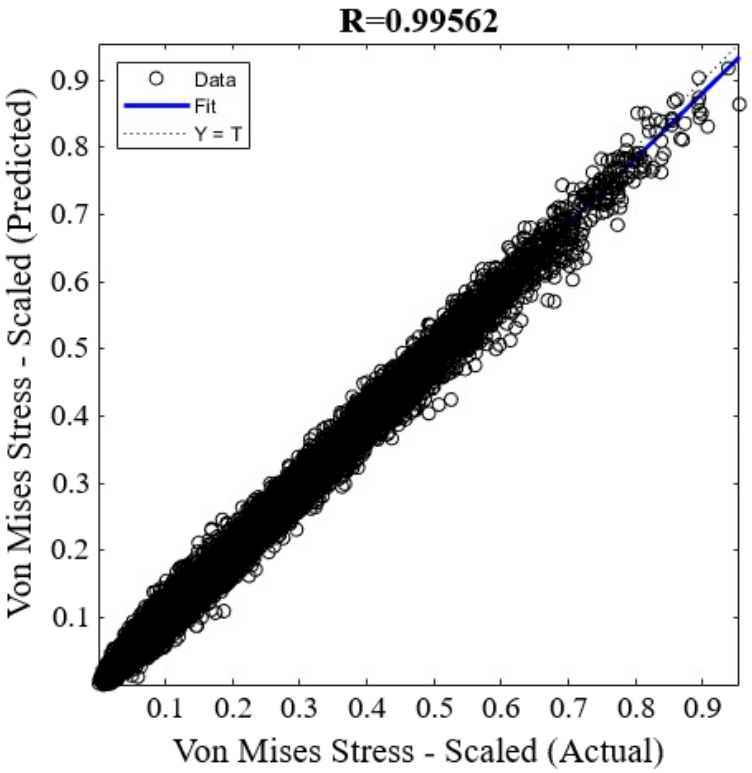


Figure 238 Regression between normalised predicted and actual known values for the best performing maximum stress predictor CNN

The developed CNN architecture was able to provide similarly strong regression when it was reapplied to predict the maximum stress anywhere in the composite. However increasing the dataset size to 200,000 observations was necessary to exceed 99% regression. The RMSE and loss was improved, showing more agreement between training and validation values compared to the full predictor model after training for the same number of epochs.

The absolute error between predicted and known maximum stress values was evaluated using 500 additional test layups and is displayed in Table 48. The layups were determined using a Latin hypercube design of experiments test plan to cover the possible design space. The best CNN that was trained over 4hr 23min 17sec gave predictions to a mean error of 8.52MPa. This network was selected for combination with the evolutionary optimisers. The absolute error of the network predictions is shown for each of the test layups in Figure 239.

Table 48 Comparison summary of absolute error between predicted and known maximum stress values for 500 test layups

	Max Stress Error (MPa)				
	2000 Sample Training Dataset			**200000 Sample Training Dataset**	
	ANN 4min 16sec	**DNN 9min 2 sec**	**CNN 23min 39sec**	**DNN 1hr 41min 13sec**	**CNN 4hr 23min 17sec**
Mean	36.90	32.25	30.08	11.49	8.52
Standard Deviation	30.10	29.42	26.24	10.32	7.10
Max	152.67	199.53	172.09	76.42	53.00
Min	0.19	0.08	0.10	0.01	3.62E-05
Range	152.49	199.45	171.99	76.41	53.00

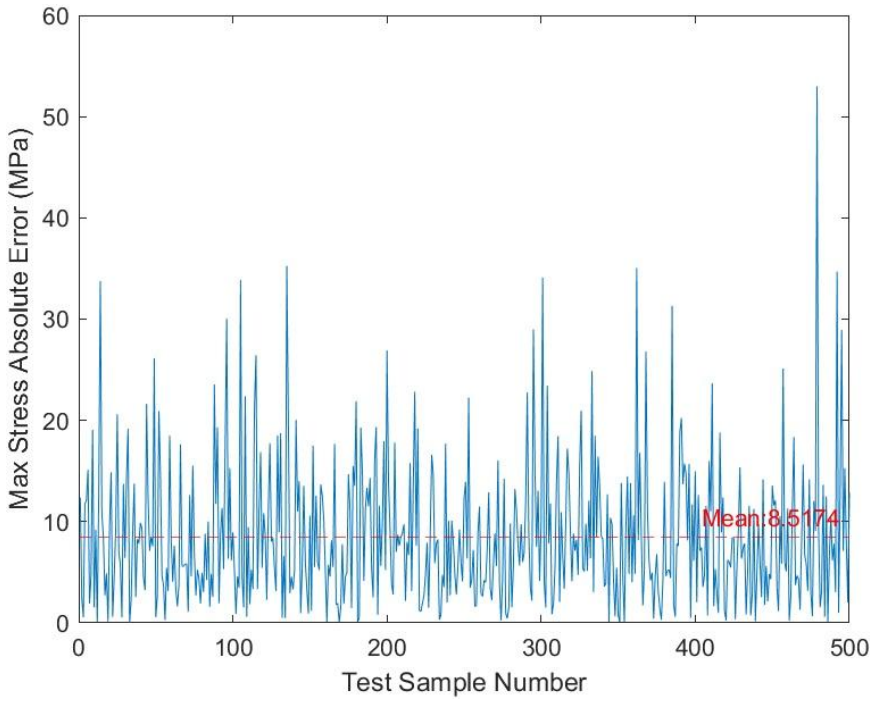


Figure 239 Absolute error between predicted and known maximum stress values for best performing CNN

The maximum stress predictor network was combined with the optimisers using a script implemented in MATLAB's optimisation toolbox, setup to minimise a non-linear objective function according to 8 input variables that correspond to the stacking angles of the layup. The script then returned the optimised objective value and the set of solution parameters necessary to achieve it to variables in the workspace. The optimisers were set to use non integer constraints for the angle variables with lower and upper bounds of 0° and 180° respectively. The objective function was assigned from a custom script placed in the working directory. The script loaded the pretrained neural network and scaled the stacking angles from each member of the population at every stage of the optimiser to normalise between 0 and 1 according to the maximum and minimum values in the original training dataset. It then converted to sequence data and ran the prediction through the neural network to output the normalised maximum stress before rescaling back to values in pascals. Optimisation tests were repeated with three different algorithms and the solutions compared. Genetic algorithm (GA), particle swarm optimisation (PSO) and simulated annealing (SA) optimisers were used. Each was set to continue until the average change in the best fitness value remained less than $1x10^{-6}$ over a given maximum number of stall iterations. For the genetic algorithm this limit was 50 generations and the maximum possible was capped at 800. The particle swarm optimisation was stopped at 20 stall iterations, with a maximum of 1600. Simulated annealing had a stall limit of 4000 iterations and the objective function could be evaluated a maximum of 24000 times. In all cases no time limits were set and there were no lower limits on the fitness values. Settings used for each optimiser are summarised in Table 49.

Table 49 Optimisation algorithm settings

GA	**Solver**	Genetic Algorithm	**Population Size**	200
	Creation Function	Uniform Distribution	**Max Generations**	800
	Crossover Function	Scattered	**Max Stall Generations**	50
	Crossover Fraction	0.8	**Stall Time Limit**	Inf
	Fitness Scaling Function	Rank	**Time Limit**	Inf
	Mutation Function	Adapt Feasible	**Function Tolerance**	1.00E-06
	Selection Function	Stochastic Uniform	**Constraint Tolerance**	0.001
	Number of Elite Members	10	**Fitness Limit**	Negative Inf
PSO	**Solver**	Particle Swarm	**Max Stall Iterations**	20
	Creation Function	Uniform Distribution	**Stall Time Limit**	Inf
	Inertia Range	0.1 to 1.1	**Time Limit**	Inf
	Min Neighbour Fraction	0.25	**Swarm Size**	80
	Self Adjustment Weight	1.49	**Initial Swarm Span**	2000
	Social Adjustment Weight	1.49	**Function Tolerance**	1.00E-06
	Max Iterations	1600	**Objective Limit**	Negative Inf
SA	**Solver**	Simulated Annealing	**Max Function Evaluations**	24000
	Initial Starting Point	[0 0 0 0 0 0 0 0]	**Max Iterations**	Infinity
	Acceptance Function	sa	**Stall Iteration Limit**	4000
	Temperature Function	exp	**Time Limit**	Infinity
	Initial Temperature	100	**Function Tolerance**	1.00E-06
	Reannealing Interval	100	**Objective Limit**	Negative Inf

The genetic algorithm created a random uniformly distributed initial population of size 200. The fitness values were calculated for each member of the population and scaled according to rank. A stochastic uniform selection function was used to choose parents for following generations. At each generation 10 elite members were selected for guaranteed survival from the previous generation. 80% of the remaining members were crossover children produced by combining parents according to the scattered crossover function. In this case a randomised binary vector is generated. Where the vector has the value 1 a gene is selected from the first parent, when it is 0 the gene is instead selected from the second parent. The remaining 20% was created via mutation by introducing random changes in the genes of individual parents. This enables genetic diversity and allows the limits of the design space to be explored. The particle swarm optimiser was initialised with a random uniform swarm of size 80 and span 2000, while the simulated annealing process began from an initial starting point with all angle parameters set to zero.

Each optimisation process was repeated 7 times. The stacking angle values output by the optimisers were rounded to the nearest degree. FEA simulations were then carried out using the solution layups in the same ANSYS model that was used to generate the training dataset. The maximum stress in the composite plate for each test run is compared against the smallest value that was found in the original training dataset in Figure 240. The best layup solutions returned by each algorithm are displayed in Table 50. The stacking angles returned by the optimisers in every run can be found in Appendix S. Contour plots of the stress distributions obtained from the best performing layups are displayed alongside the corresponding optimisation history on page 340.

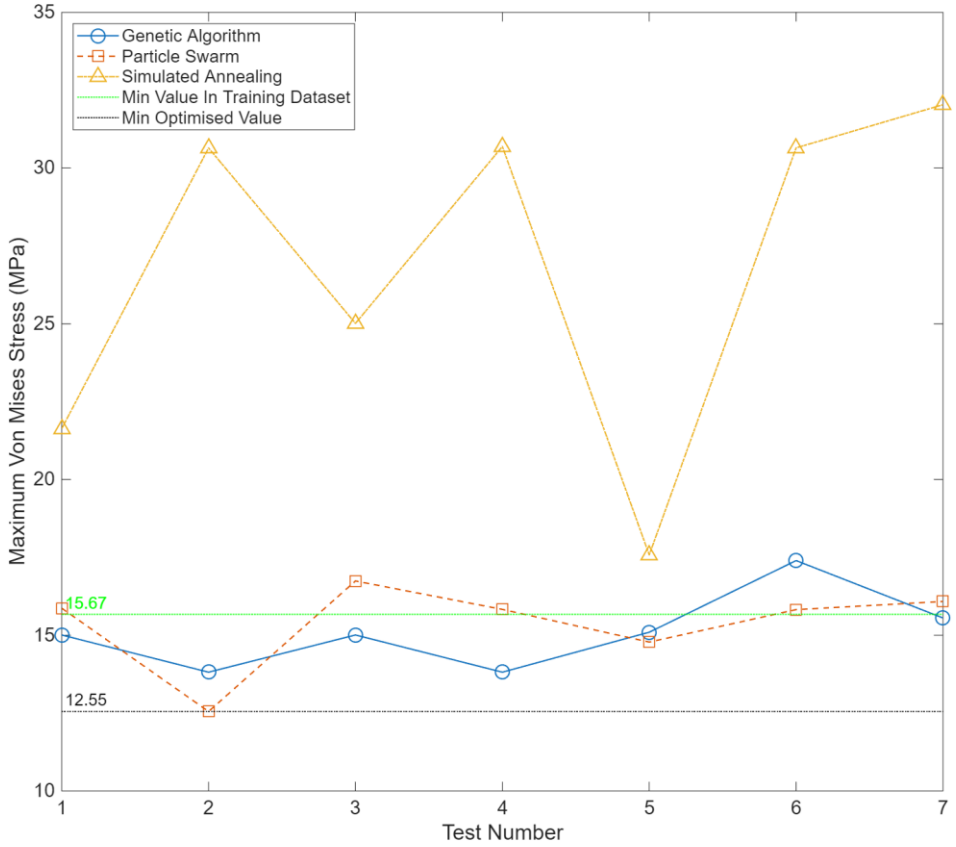


Figure 240 Maximum Von Mises stress in the composite plate for the optimised layup sequences

Table 50 Best performing solution layup stacking sequences

Optimiser	Solution Stacking Sequence (°)		FEA Max Equivalent Stress (MPa)
GA Test 2/4	**Angle 1**	180	13.81
	Angle 2	180	
	Angle 3	180	
	Angle 4	98	
	Angle 5	180	
	Angle 6	3	
	Angle 7	103	
	Angle 8	98	
PSO Test 2	**Angle 1**	83	12.55
	Angle 2	180	
	Angle 3	6	
	Angle 4	180	
	Angle 5	180	
	Angle 6	93	
	Angle 7	0	
	Angle 8	99	
SA Test 5	**Angle 1**	7	17.58
	Angle 2	66	
	Angle 3	2	
	Angle 4	0	
	Angle 5	0	
	Angle 6	3	
	Angle 7	71	
	Angle 8	103	

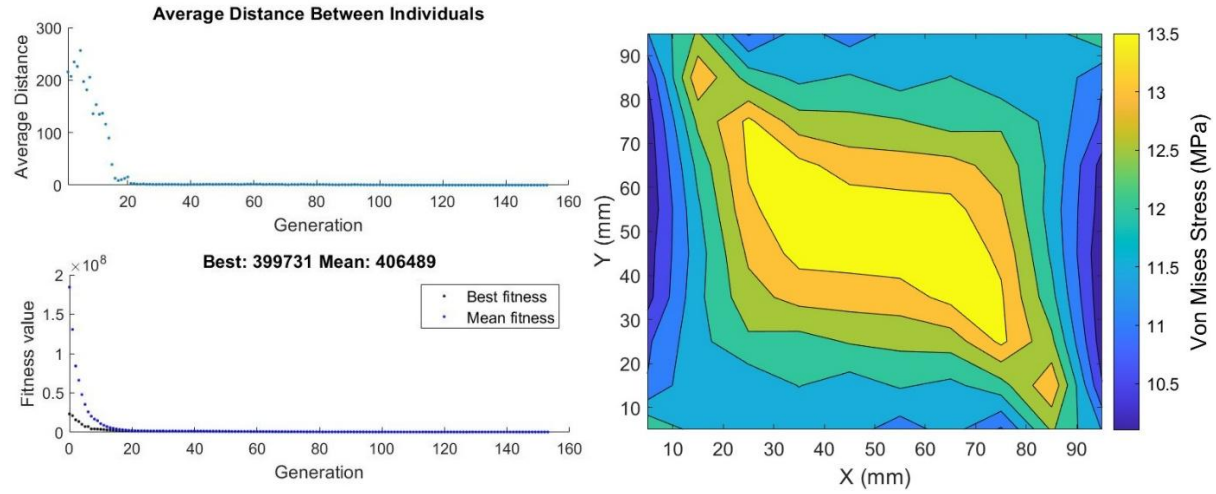


Figure 241 Genetic Algorithm Test 2 optimisation history (left) and solution layup FEA stress response (right)

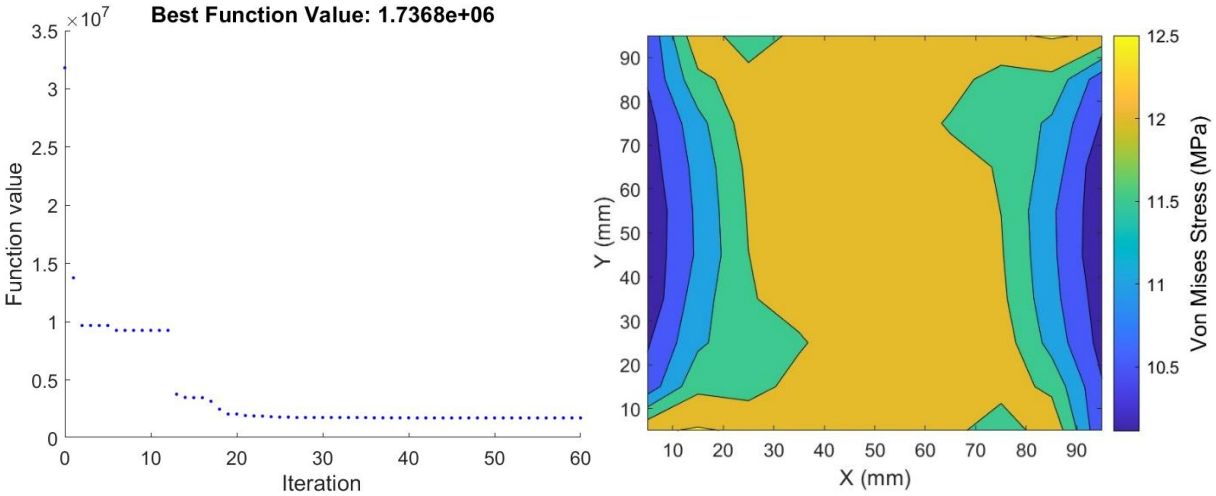


Figure 242 Particle Swarm Test 2 optimisation history (left) and solution layup FEA stress response (right)

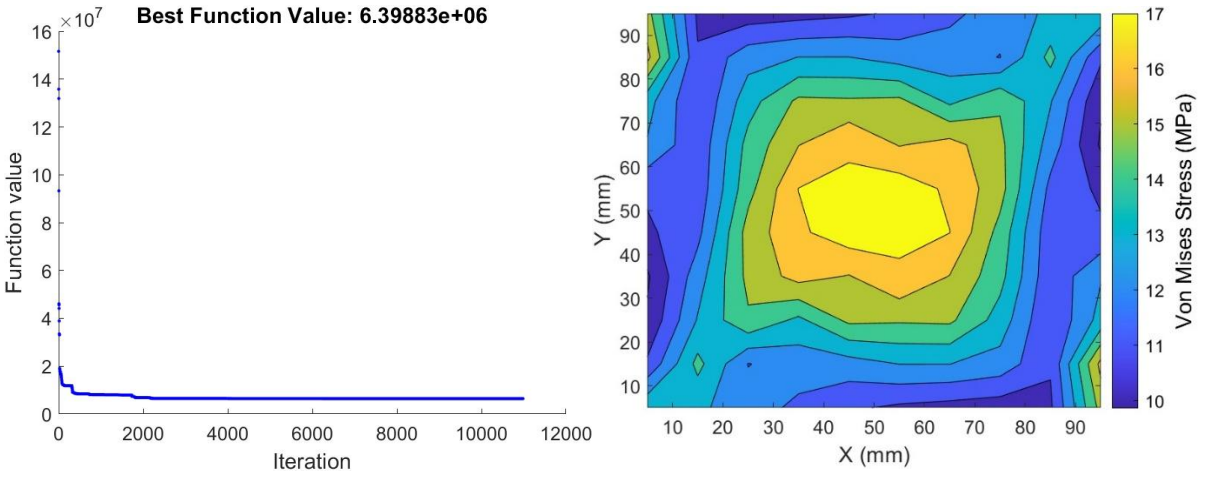


Figure 243 Simulated Annealing Test 5 optimisation history (left) and solution layup FEA stress response (right)

In all cases the optimisation stopping criteria was triggered prior to the maximum number of iterations, indicating negligible further improvement would be expected if the algorithm was run for a greater amount of time. All history plots show a sharp initial reduction in fitness value that levels off as the optimisation progresses. The average distance between individual members of the GA population fluctuates initially and then begins to drop rapidly after 10 generations, indicating the solver honing in on an increasingly reduced region of the design space as it converges on a minima. Variations in the returned layups suggest it is likely that running the optimisers multiple times will result in different solutions that will depend on the initial population. It is therefore recommended that the analysis should be repeated using multiple solvers. A script may be implemented to automate this process and return the best performing layup. However in all cases the algorithms successfully returned low stress layups that gave maximum stresses close to the minimum value in the original dataset that the CNN was trained on. The greatest value of maximum stress found in the training dataset was 855.58MPa. The smallest maximum stress was 15.67MPa. The genetic algorithm optimiser was most consistent at providing layups with maximum stress below this value. It also returned identical solutions on multiple occasions. Layups from tests 1 and 3 were the same, as were those from tests 2 and 4. Both genetic algorithm and particle swarm outperformed simulated annealing. Simulated annealing was effective at returning layups that lead to maximum stresses that were less than 4% of the largest value and 3 of the 7 tests returned similar layups that differed by as little as 1 degree between angles. The particle swarm optimiser converged on different solutions each time however gave the best overall result, 3.12MPa below the minimum value found anywhere in the original training dataset of 200,000 possible layups. The speed of the CNN allows for hundreds of generations worth of stress predictions to occur within a practical timeframe. The best GA solution took 1 hour 10 minutes for the optimiser to converge, the PSO took 15 minutes and SA 45 minutes. The model framework can be readily applied to design of full composite components. As only the maximum stress that occurs in the component is necessary to carry out the optimisation, the existing network architectures are likely to be sufficient and no additional complexity is required for layup sequences of similar length. The networks can simply be retrained on a new FEA dataset. The main limiting factor for widespread uptake is the practicality of generating large enough quantities of data, which may require parallel processing and high performance computing depending on the size of the model.

7.3 Artificial Intelligence Based Parametric Stress Analysis

The predictor model developed in section 7.2 can be applied to other design problems, however a new dataset is required for each problem. A more desirable goal is a single pretrained artificial intelligence programme that can be used for any combination of geometry, loading and boundary conditions. Such a model would serve as a faster alternative to FEA for parametric stress analysis, applicable across all engineering design practices. It would also provide the additional benefit of automatically optimising layups to maximise the strength of composite components, suggesting the optimal layup and returning the stress distribution alongside for the given loading. The aim of this section of the project is an initial proof of concept for such a model based on simplified 2-dimensional geometry and in-plane loading.

The predictor network framework was extended to allow the geometry and boundary conditions to vary, reapplying the CNN architecture outlined in Figure 228. The geometry and loading were supplied as additional input parameters to the network. Table 51 summarises the input and output parameters the CNN was trained on. The geometry was kept a simplified 4-sided polygon. The locations of the 4 vertices were represented as integer values between 1 and 121 corresponding to points mapped onto a grid in 2D space. Figure 244 displays an example geometry with the possible vertex locations and line numbers labelled. Each vertex was constrained to a 100x100mm grid as shown. The polygon was able to form any shape within these grids, the smallest possible geometry being a 100x100mm square and the largest 300x300mm. The boundary conditions were represented by the line numbers that a displacement load and a fully fixed support were applied to, with line 1 joining vertices 1 & 2, line 2 joining vertices 2 & 3, line 3 joining vertices 3 & 4, and line 4 joining vertices 4 & 1. It should be noted that a constraint must be applied when selecting the dataset test plan to ensure the load and support are not applied at the same location. For simplicity in these initial tests the support was always applied to the line opposite the load and only a single parameter was used to define the loading location. Separate parameters were used to define the components of the displacement load in the x and y directions, which were allowed to vary between 10mm and 100mm. The number of stacking angle parameters was reduced to 3 for a 6-ply symmetric layup. For full stress predictions the x and y coordinates of the element centroid location where the solution was to be probed was also included. In total this lead to 12 input parameters mapped to a single scalar output, the Von Mises equivalent stress in the element. The equivalent model for maximum stress prediction and optimisation featured 10 inputs, without the probe locations.

Table 51 Parametric stress predictor model parameters

	Parameter Name	Units	Minimum	Maximum
Inputs	elem_x_cent_loc	metres	0.00315	0.29639
	elem_y_cent_loc	metres	0.00352	0.29673
	vertex_loc_1	integer	1	121
	vertex_loc_2	integer	1	121
	vertex_loc_3	integer	1	121
	vertex_loc_4	integer	1	121
	load_line_no	integer	1	4
	displacement_x	metres	0.01	0.1
	displacement_y	metres	0.01	0.1
	angle_1	degrees	0	180
	angle_2	degrees	0	180
	angle_3	degrees	0	180
Output	elem_eqv_stress	pascals	2.77E+03	4.13E+08

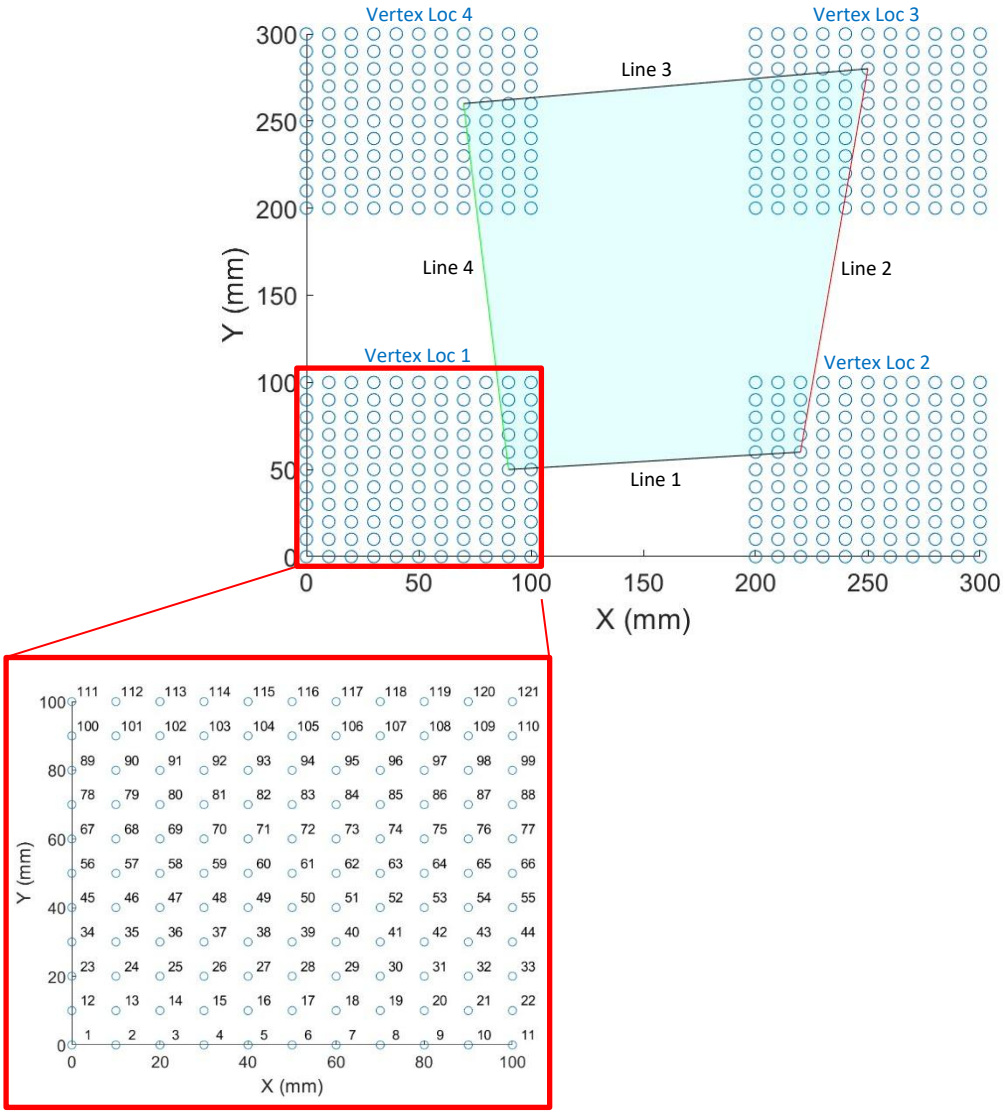


Figure 244 Example geometry and boundary conditions for the parametric predictor model. Red line indicates load, green line is fully fixed support

The design of experiments test plan displayed in Figure 245 was used to generate a training dataset, featuring 2000 FEA simulations at design points chosen using Latin hypercube sampling with the parameter ranges outlined in Table 51. A MATLAB script was used to convert the vertex location integer values to x and y coordinates in 2D space. The geometries and boundary conditions were then defined using an ANSYS Mechanical APDL script which meshed each geometry with elements of size 10mm. Generation of the dataset of stress responses was automated by the script, which also returned the number of nodes and elements in each mesh. The 2000 FEA runs took 3hrs 36min 38sec, generating a total of 1,120,878 observations corresponding to the individual stress readings in every element of every mesh.

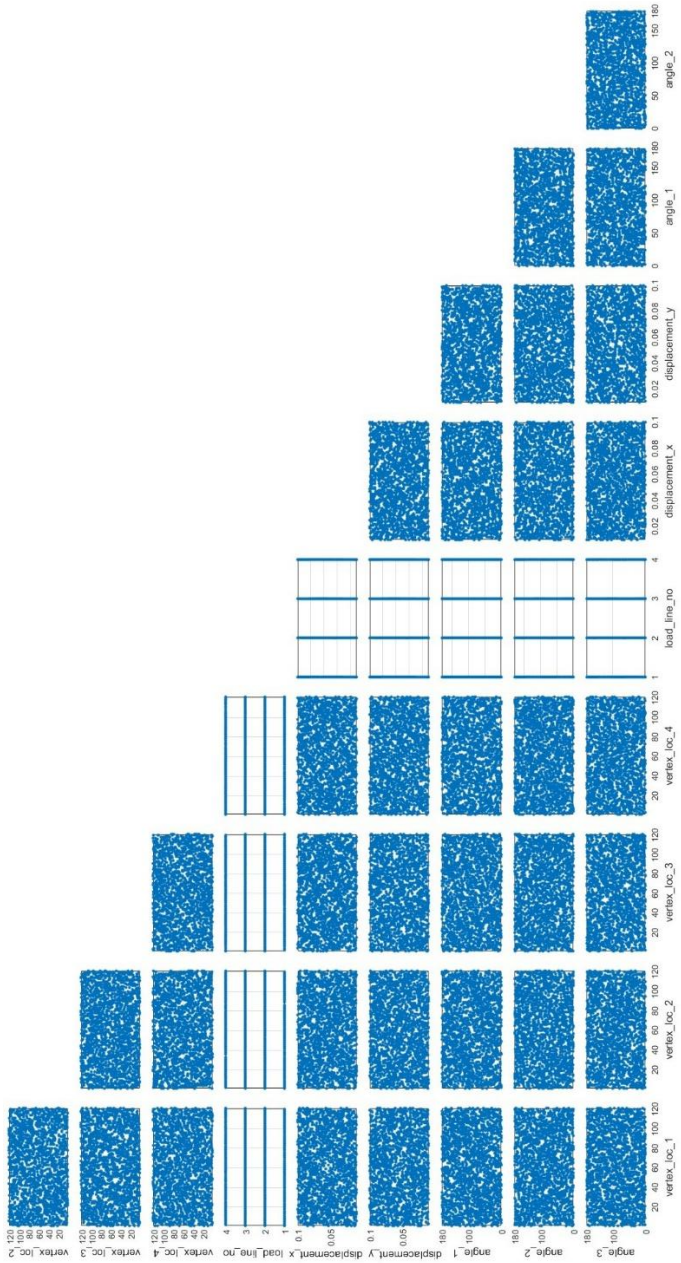


Figure 245 Parametric predictor model training dataset design of experiments test plan. 2000 datapoints generated via Latin hypercube sampling.

Of the observations 70% were used to train the network and 15% for validation checks during the training. The observations were passed through the CNN in mini batches of 1000. Figure 246 displays the loss and root mean square error between the predicted and known stress values at every training iteration and validation check for the best performing network. Each vertical stripe indicates one epoch, where the entire dataset has been passed through the network once. Training for 1000 epochs took 15hrs 16sec and the regression fell off to 62.3%, likely due to overfitting. The best performance was found by limiting the training to 100 epochs over 1hr 21min 38sec. For this case a rapid drop in RMSE and loss were observed over the first epoch. After this small spikes in RMSE can be seen between epochs 5 and 7. These fluctuations smooth out into gradual improvements for the remaining epochs and strong agreement is maintained between training and validation values throughout.

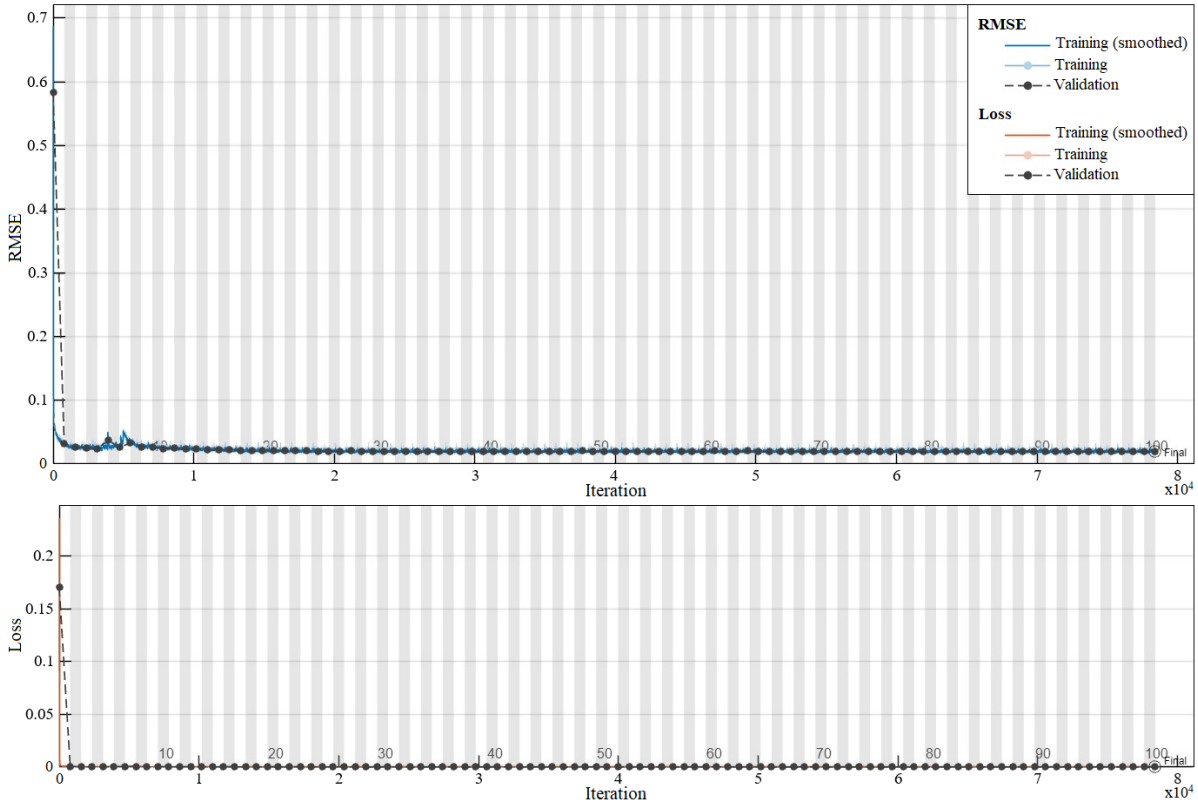


Figure 246 Training history of best performing CNN model for parametric stress prediction

The partition of 15% that remained unseen by the network during training was used to test the stress predictions against known values obtained by FEA. The regression between predicted and actual values is shown in Figure 247. A greater spread in datapoints is observed with a larger number of outliers compared to the previous predictor model, likely due to the increased complexity and greater degree of variation coming from allowing the geometry and loading to vary. The greatest discrepancies occur when attempting to predict the highest actual stress values. The network is likely to be vulnerable to anomalies in the FEA model

used to generate the training dataset. Additional steps to filter out anomalous values in the data before training may be recommended. Predictions may also suffer when approaching the limits of the design space. However a strong regression coefficient was maintained overall. The trained model provided 91.8% regression across 168,131 test observations. The test partition was assigned randomly. This value therefore corresponds to stress predictions in randomised locations for random combinations of geometry, loading and layups across the initial dataset.

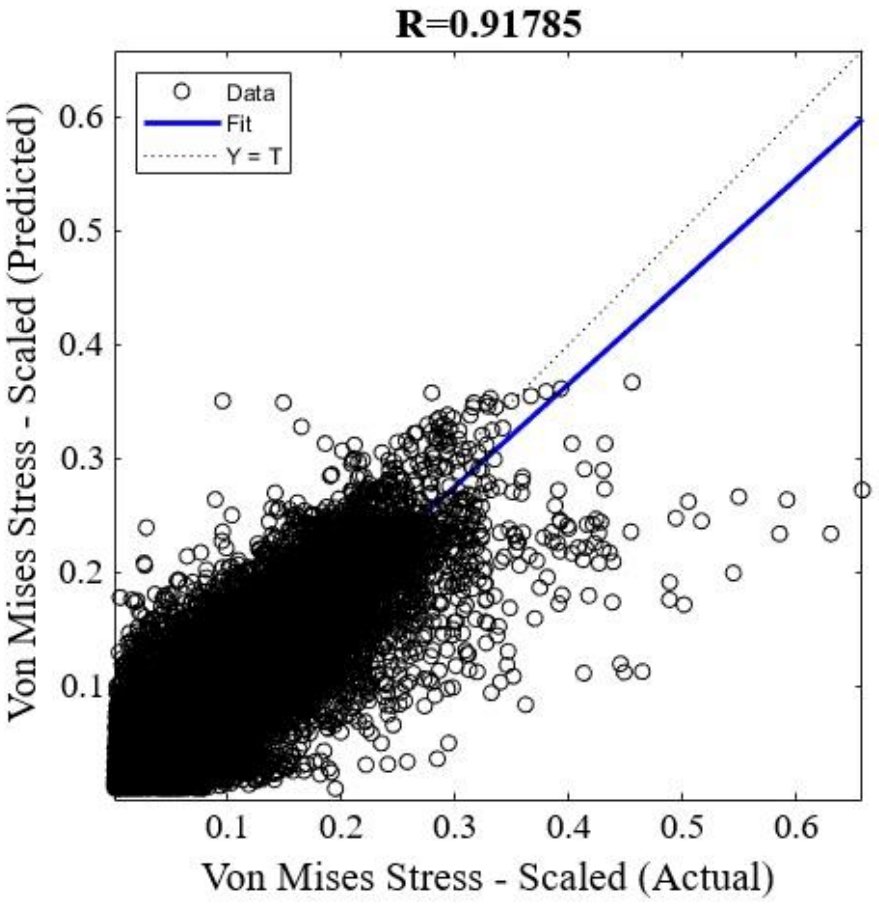


Figure 247 Regression between predicted and actual known normalised Von Mises stress values for parametric predictor CNN. Plot features 168,131 datapoints corresponding to separate test observations.

Further analysis was carried out to test full stress predictions, considering the entire mesh for a specific selection of geometry and loading scenarios. A selection of 500 test cases were chosen to represent the design space of possible combinations using Latin hypercube sampling. The chosen parameter combinations are displayed in Appendix T. At each design point the full stress distribution across the composite plate was obtained using the same FEA model the training dataset was generated from. The coordinates of every element were extracted and the locations used to probe stress predictions with the trained CNN, building up

the equivalent full distributions across the plate for comparison. For every element the difference was calculated between the known and predicted Von Mises stress. The mean absolute error averaged across the plate was then determined. Values for each of the 500 test cases are shown in Figure 248. The CNN gave predictions to an average mean absolute error of 7.28MPa across the test runs, within 1.8% of the maximum stress value found in the original dataset of 413.29MPa. The mean error remained below 20 MPa with a probability of 96.6%, below 30MPa with a probability of 98.6%, and remained below 40MPa in all but 2 cases. Reductions in accuracy of predictions compared to FEA may be outweighed by considerable improvements in solve time. To test the speed of the solver a dataset of full stress predictions at 200,000 different design points was generated by FEA from a Latin hypercube test plan. The same analysis was then repeated using the CNN to provide the stress predictions. The FEA simulations were automated by the APDL script in 2 batch runs of 100,000, taking 31hrs 25mins 20sec and 34hrs 26mins 22sec with a total run time of 65hrs 51mins 42sec. The equivalent CNN predictor model carried out the analysis in only 47min 40sec on identical computer hardware. Evaluation of the machine learning model suggests performance that is likely to be sufficient to enable rapid design iterations of composite components. Final results may then be validated using FEA and real-world experimental testing.

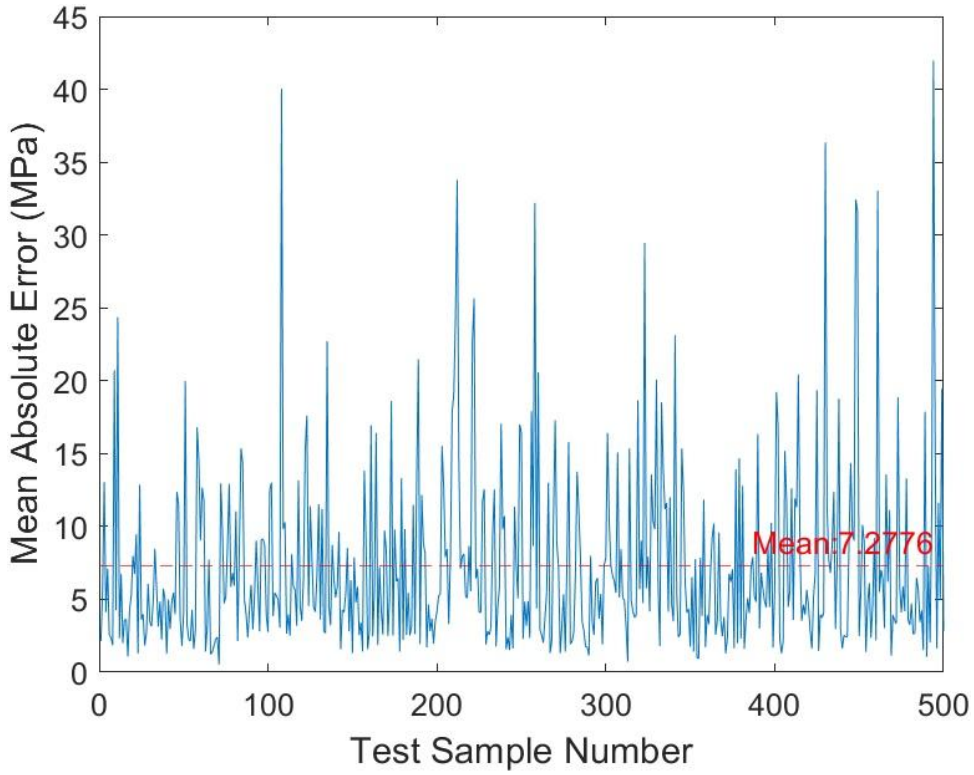


Figure 248 Mean absolute error between CNN predictions and FEA values of Von Mises stress averaged across each composite plate for 500 test cases of selected geometry, loading and layup combinations

7.4 Composite Layup Optimisation for Specific Geometry and Boundary Conditions

The parametric predictor model was adapted for prediction of the maximum stress and combined with genetic algorithm and particle swarm optimisers to output the layup that minimises the maximum stress for given input geometry and boundary conditions. The network parameters outlined in Table 51 were altered by removing the element locations and the element equivalent stress output was replaced by the maximum stress anywhere in the composite. The dataset size was increased to 200,000 FEA simulations sampled using Latin hypercube to provide 200,000 discrete observations for the network to train on. The best CNN from section 7.3 was retrained using the same architecture, hyperparameters and data partition ratio to predict the maximum stress. After 100 epochs the CNN provided approximately 90.6% regression between predicted and known values for 30,000 test observations. Increasing the training to 500 and 1000 epochs resulted in 90.1% and 90.2% regression respectively and these networks were discarded. The 100 epoch network took 13min 43sec to train and was selected for use in the optimisation objective functions. The training history and regression for the chosen network are displayed in Figure 249 and Figure 250.

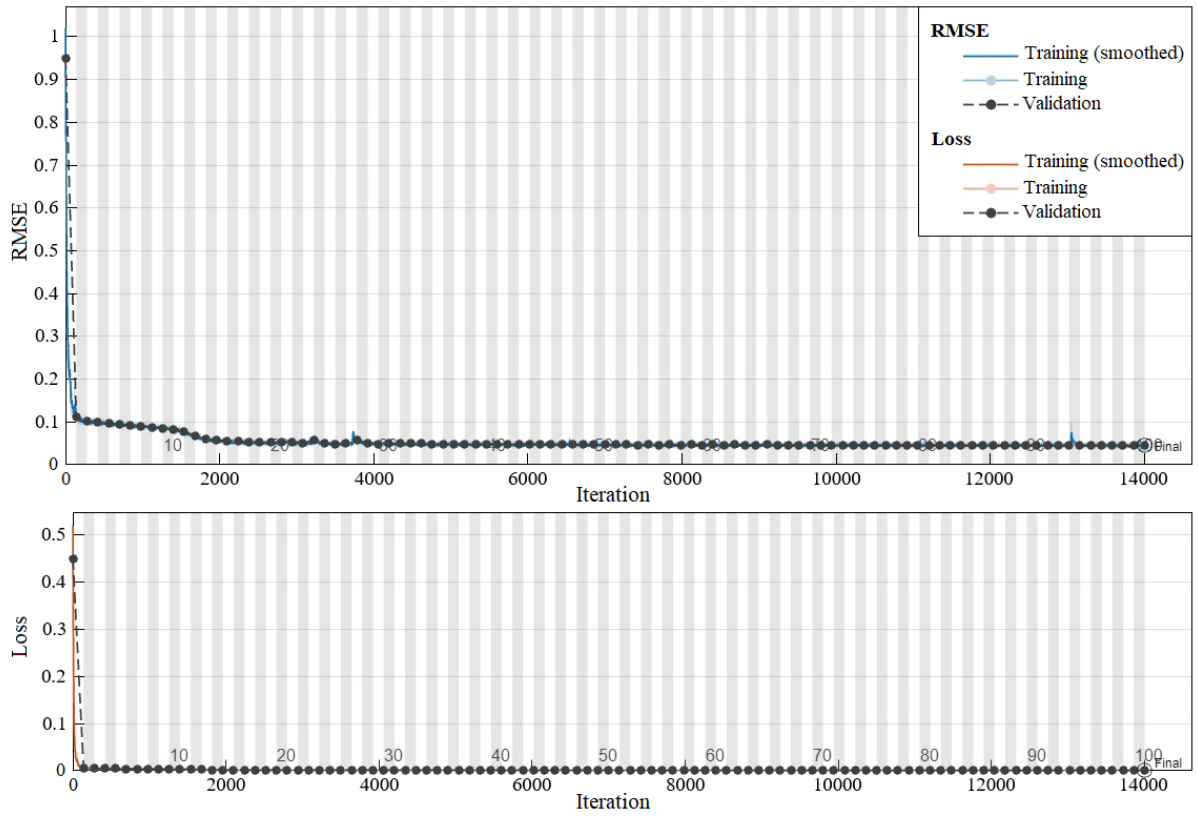


Figure 249 Maximum stress parametric predictor CNN training history

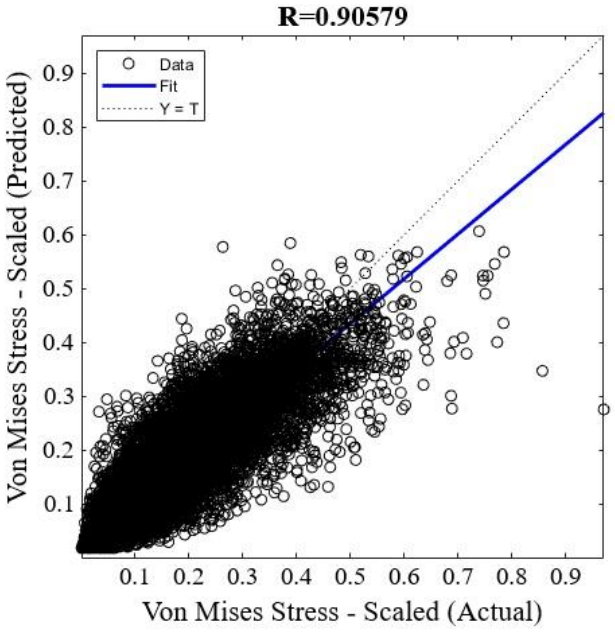


Figure 250 Maximum stress parametric predictor CNN regression between predicted and actual normalised maximum Von Mises Stress

The error histogram in Figure 251 shows the difference between the predicted and actual maximum stress values for the 30,000 test cases, which were partitioned at random from the original dataset and remained unseen during training. The error distribution indicates the network is more likely to overestimate than underestimate the maximum stress. The mean error across all 30,000 observations was 28.32MPa and the error remained below 50MPa for 80.3% of the test observations.

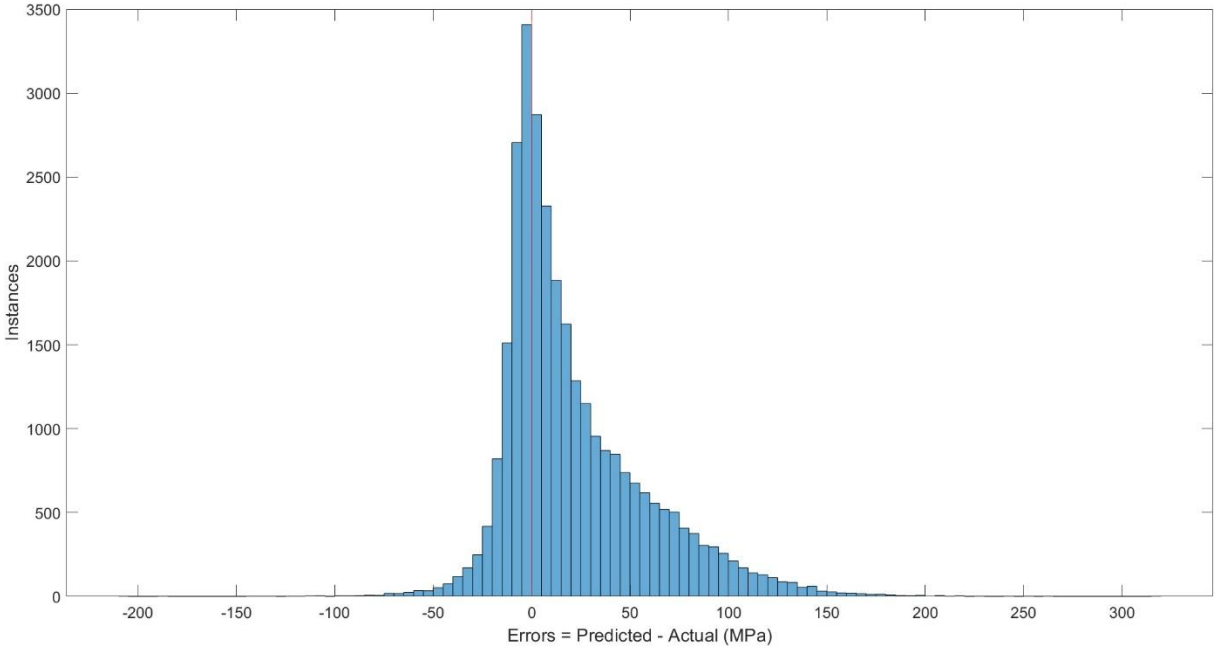


Figure 251 Maximum stress predictor CNN error histogram for test dataset of 30,000 observations

The genetic algorithm and particle swarm optimisation scripts were altered to output the stacking sequence angles to minimise the maximum stress in the composite for a specific set of geometry and loading parameters. The vertex coordinates of the geometry, the boundary condition location and components of the loading in the x and y directions were defined at the start of the script prior to the optimisation. These values were then locked and imported into the objective function automatically, while the layup angles were able to vary for each member of the optimiser population. A script was written to set the objective function, calling the pretrained CNN to predict the maximum stresses at each iteration of the optimisation algorithm. To evaluate the ability of the model to output minimal stress layups, optimisations were repeated for a selection of 50 test combinations of geometry and loading conditions. These test design points were chosen using Latin hypercube sampling to represent the possible ranges of combinations. The pairwise plots of the chosen design points are shown in Figure 252. For every design point genetic algorithm and particle swarm optimisers were tested and the maximum stress from the output layups determined by FEA. An exhaustive search was then carried out for comparison, evaluating the maximum stress from each of the 1000 layups displayed in Figure 253 using 50,000 FEA runs in total.

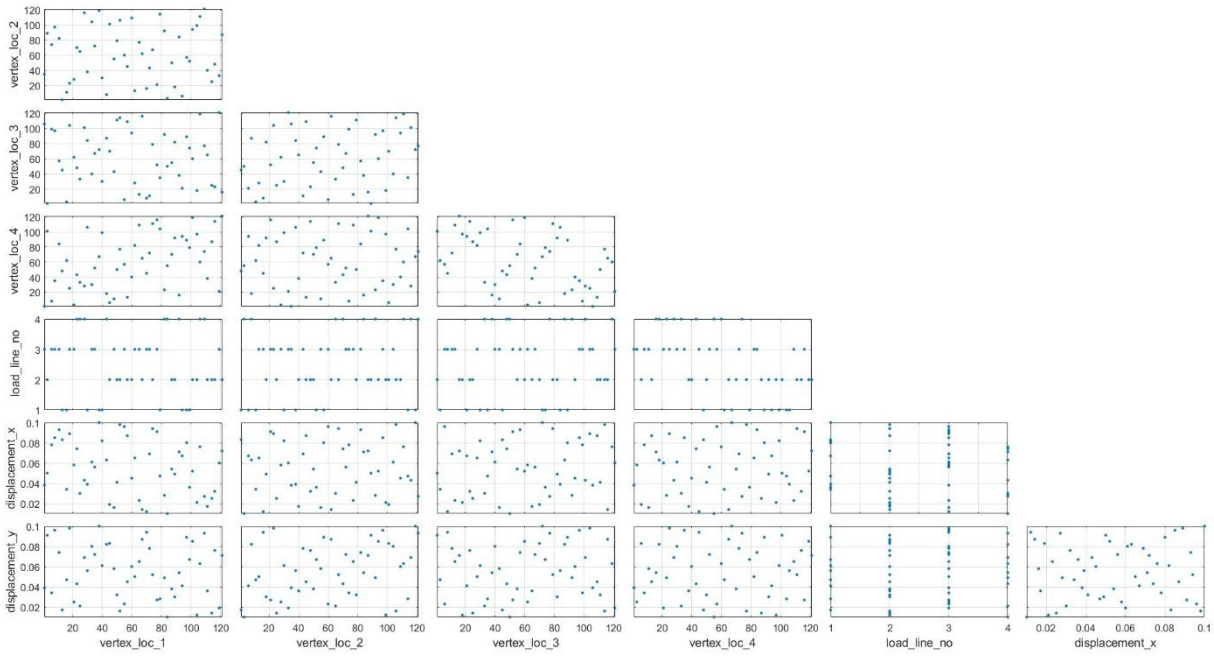


Figure 252 Pairwise plots of 50 test geometry and loading conditions chosen by Latin hypercube sampling

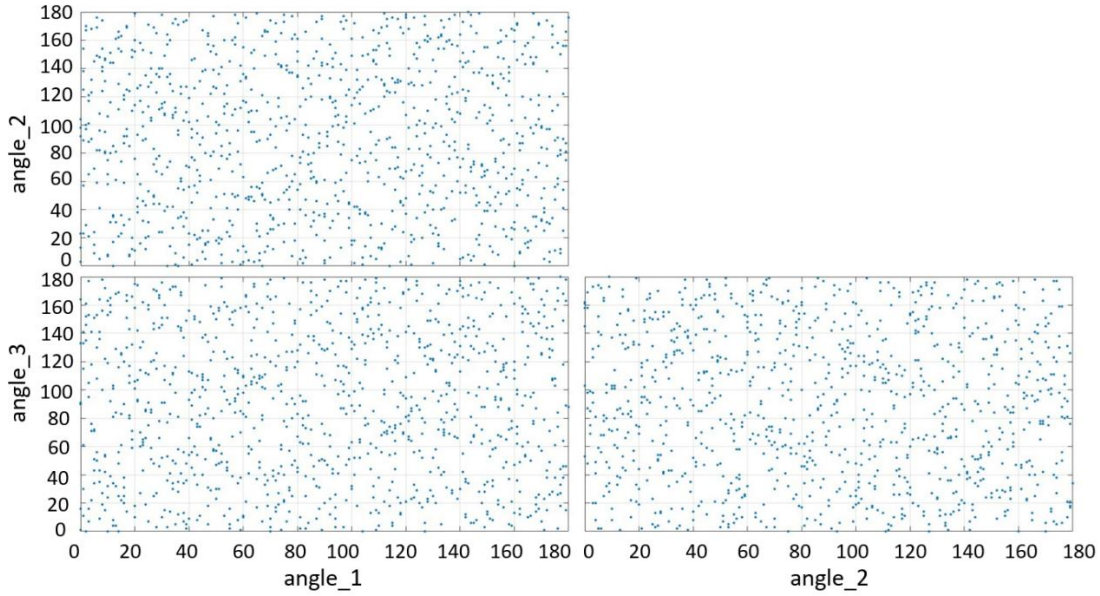


Figure 253 Exhaustive search test plan. 1000 layups chosen by Latin hypercube sampling.

The maximum stresses for the best performing layups returned by the optimisers are compared against the largest and smallest values found in the exhaustive search for each of the 50 test cases in Figure 254. Both optimisers successfully consistently provided low stress layups close to the smallest values found in the exhaustive search. The mean difference between the optimiser and exhaustive search smallest maximum stress values was 2.90MPa for the genetic algorithm and 2.94MPa for the particle swarm. Genetic algorithm provided layups that outperformed the exhaustive search in 10% of the tests, while particle swarm was better than the exhaustive search 12% of the time. The AI assisted optimisers both gave layups that provided maximum stresses that were on average less than 2% of the greatest values found in the exhaustive search. The mean values across the study were 1.96% for GA and 1.94% for PSO. A summary table of the largest and smallest maximum stress values in the exhaustive search, and the smallest values obtained by GA and PSO is provided at each test point in Appendix U.

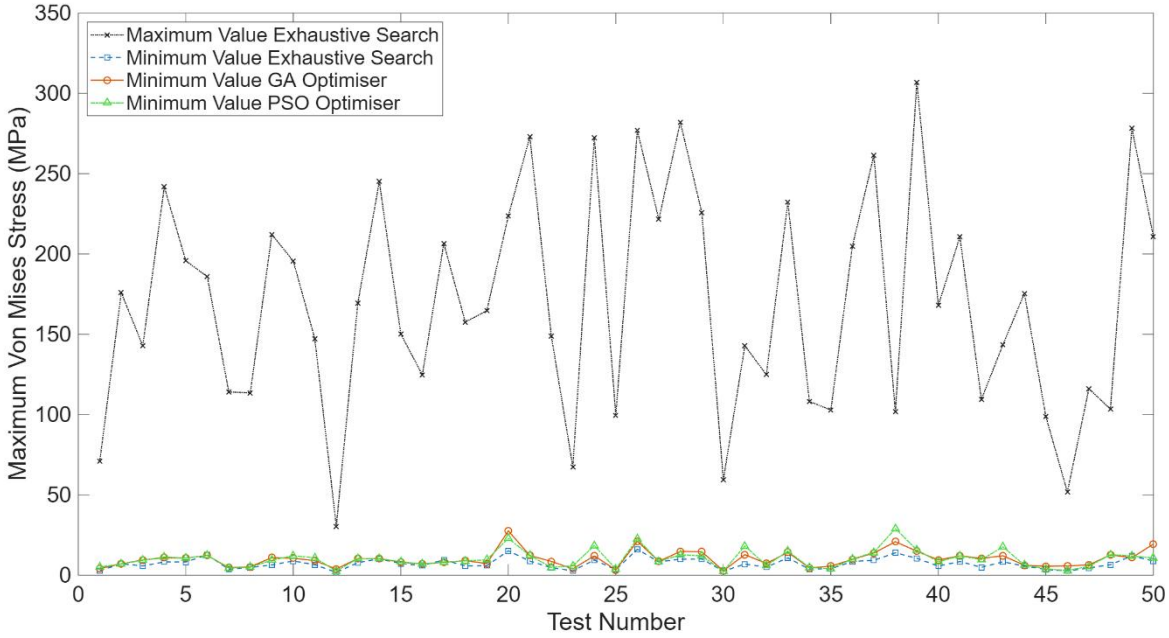


Figure 254 Comparison of maximum Von Mises stress for optimised test layups against smallest and largest values found by exhaustive search

Each optimisation was repeated 5 times to determine the consistency of the predictions. Figure 255 shows the range between the largest and smallest maximum stresses obtained using the layups output by the optimisers. The maximum stress values in every run can be found in Appendix V. Both algorithms are likely to converge on different solutions when the optimisations are repeated, however for 80% of all the test runs in the study the prediction range was below 10MPa. Occasionally running the optimiser multiple times can show more extreme variations in the results. The greatest variation was 27.73MPa using particle swarm optimisation. Both algorithms consistently provided layups that offer low stress solutions and each may outperform the other on separate occasions. For best results it is recommended to use both solvers and run them multiple times from an automated script.

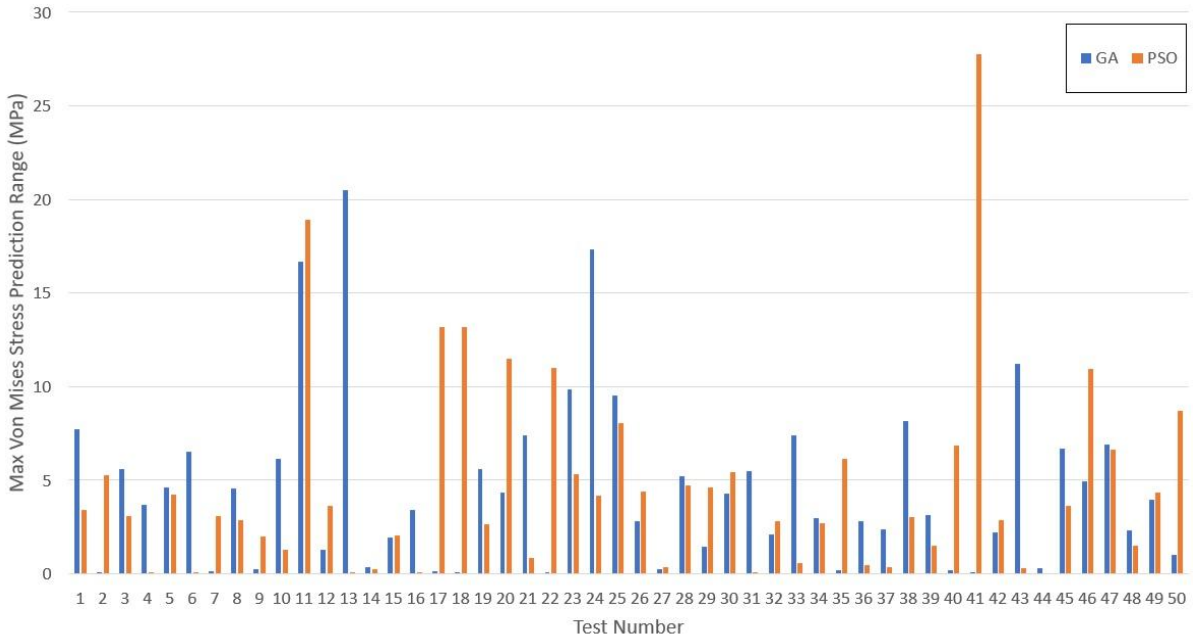


Figure 255 Optimised composite layup prediction consistency for genetic algorithm and particle swarm optimisation

The study demonstrates evidence that machine learning models may be developed that are able to take the parameters traditionally used to define an FEA stress analysis and return an optimised layup for minimum stress design of composite components, alongside performing a full stress analysis in a fraction of the time that a conventional FEA simulation would take to run. Allowing the geometry and boundary conditions to vary significantly increased the complexity of the solver compared to the previous model where the parameters were fixed, leading to increased error in the stress predictions. However the accuracy was sufficient for successful optimisation of composite layup sequences for minimum overall stress, offering results comparable to exhaustive search that do not require extensive FEA simulations if the model has been pretrained in advance to work on any combination of parameters. The CNN based optimiser may still be able to successfully minimise the maximum stress if the accuracy of the predictor is limited, as long as the training has extracted enough features to identify the underlying patterns in the dataset and offer strong regression performance.

7.5 Conclusion

Stress optimisation of composite components has previously been limited by lack of governing theory. The proposed machine learning framework circumvents this requirement, enabling rapid stress optimisations without extensive trial and error or exhaustive search processes. A regression model can be built from empirical or numerical simulation data without the need for analytical solutions. The solver is able to return the layup to maximise the strength of the composite with high consistency. This may be used for enhanced lightweighting in real-world product development, enhancing the speed of the design cycle while offering improved design flexibility. Geometries and loading scenario can be updated to suit evolving requirements such as packaging and aerodynamics and the model will immediately return the optimal layup and accompanying stress distribution without requiring new optimisers to be rebuilt from scratch. The proposed neural network models may be applied to aid design challenges associated with unique failure mechanisms observed within composite materials. In future works the stress criterion may be substituted for alternative failure metrics such as "Tsai-Wu and Tsai-Hill" and the networks retrained while the dimensionality of the inputs and outputs remains the same. Machine learning models developed for optimisation of carbon fibre composite layup for minimum stress design offer significant potential for lightweighting in passenger vehicles, allowing strengthening of components purely through tuning of stacking angles without adding additional plies and extra mass. Layups can be optimised according to a given practical loading scenario, ensuring the minimum amount of weight for a given application and leading to reduced fuel consumption and carbon emissions. Successful proof of concept has been obtained for a fully parametric predictor model able to perform stress analyses comparable to FEA significantly faster, with potential for application across all engineering industries.

Chapter 8 Conclusion and Future Works

8.1 Conclusion

The aim of this thesis was to explore the potential for applications of machine learning to allow for optimisation and tailoring of materials for improved mechanical performance, developing foundational frameworks that can be applied to design future generations of architected metamaterials while enabling new levels of multifunctionality. A comprehensive review of the current state of the art of metamaterials has been presented. From this analysis multiple key workflows have been developed to allow for multi-objective tailoring of material properties through the use of machine learning. An inverse design approach has allowed for property relations to be dictated by the engineer. The AI model then returns the geometry necessary to meet the request, reversing traditional trial and error based development workflows and offering on demand design iterations. The inverse model offers greatest control for fine-tuning of material responses to bespoke real-world applications. The alternative method trains neural networks to predict material performance for a given geometry, which enables optimisation via evolutionary algorithms to simulate the processes of variation and selection found in nature. The predictor direction model allowed for increased levels of complexity, the largest amount of geometry parameters and the greatest number of objectives to be analysed simultaneously. Machine learning offers the ability to perform optimisations through statistical mapping of neural networks without the need for governing theory, allowing for relation of diverse physical properties that cannot be linked analytically. Therefore potential for applications is vast and ranges across all engineering disciplines. This research has studied the effectiveness for optimisation of quasi-static and dynamic mechanical properties, sound absorption and impact resistance. Both developed approaches may be readily applied to problems of similar dimensionality through the means of transfer learning.

The inverse method was able to augment absorption responses of micro-perforated panel sound absorbers to maximise absorption amplitude and range at subwavelength thicknesses. The solver showed the ability to rapidly compensate for discrepancies between theoretical and experimental results that may also be applicable to other problems such as predicting the stress-strain response of a lattice unit cell. If errors are observed under real-world materials testing, the design can be quickly reiterated by the CNN to compensate. The inverse CNN approach enhances speed of predictions and design flexibility. Manufacturing constraints and design space specification can be taken into account during the optimisation by only including

parameters that fall within the desired ranges in the training dataset. Optimisation objectives can be set by applying limits to the design space and retraining the model. For example the maximum thickness requirement can be set by training the model on a dataset where only values of the back cavity depth below a certain threshold are included. Constraints can be placed on the parameters to only train on MPPs above a certain perforation diameter and below a given number of perforations to improve manufacturability. The design objectives can be refined and updated on the fly as the project progresses by retraining the network without having to redesign its architecture. If initial results are unsatisfactory, the model can be improved by adding more training data, increasing the training time or broadening the depth of the network. However additional considerations must be taken when planning to optimise a metamaterial via deep learning. The process of developing the model is currently time consuming. Extensive trial and error is required if working from scratch for a bespoke problem and existing previous models from similar studies are not available for reference. The quality of results is largely dependent on the quality and size of the training dataset. Large dataset sizes can be necessary, which if FEA or practical experimentation is required may take a long time to generate, requiring simplification of the problem by reducing the dimensionality and number of design variables.

The preliminary studies indicated the suitability of the inverse CNN optimisation model for mapping metamaterial geometry parameters to a 2-dimensional material response. In following works this was applied to optimising non-linear stress-strain characteristics of architected cellular materials to provide maximum strength and stiffness at lowest mass. A desired stress response may be tailored for a given application. The stress plateau could be augmented to provide progressive stable deformation under compression, allowing for an increase in energy absorption for impact resistance applications. Networks can be adapted to allow for multi-objective optimisations by training on multiple performance criteria, leading to multifunctional components. This was demonstrated using both inverse and predictor approaches, which were able to simultaneously tune sound absorption and mechanical properties of a quarter-wavelength resonator labyrinth unit cell by combining datasets from analytical theory and numerical simulations. The models were also able to combine static and dynamic property optimisations in the same analysis, synthesising lattice structures that provided greater specific stiffness, strength and energy absorption than pre-existing cellular architectures formed from geometric primitives. The workflows may be used to optimise layup of composite facesheets in addition to core materials for reducing stress in composite sandwich structures. An early proof of concept was obtained that machine learning models may in future be used to perform parametric stress analysis faster than conventional FEA.

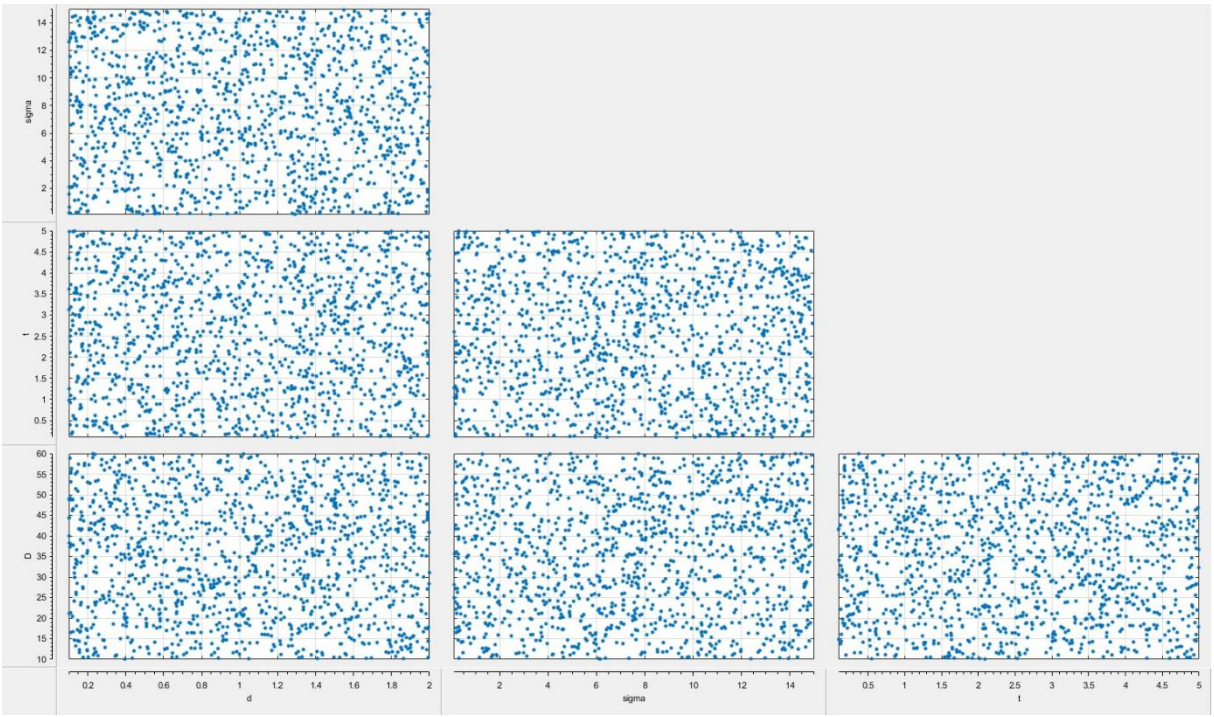
8.2 Future Works

The research undertaken as part of this project has demonstrated the potential advantages of using machine learning to design lightweight multifunctional metamaterials with mechanical properties that can be more closely tailored to a specific application than conventional alternatives. Further research may enable large-scale uptake of such approaches as standard engineering design practice in the future. The processes can be revised to allow optimisation of parametric component designs without time consuming dataset generation stages, the end goal being to produce models that will return an immediate optimal lattice structure when provided with macro level geometry, loading and boundary conditions. In future studies the CNN optimiser framework can be expanded for additional complexity, allowing more complex shapes, additional types of loads and support conditions, longer composite stacking sequences and different material properties. The inverse CNN architectures developed across this study offer potential for at least 100 input parameters mapped to 3 outputs in the same network. Future works may develop this foundation for application to 3-dimensional components. A particular strength of the CNN optimiser is the ability to map layup parameters directly to real-world loading scenarios, and this should be tested by component case studies. Additional studies are required to assess the levels of regression and dataset size requirements necessary to extract the levels of complexity needed for successful stress optimisation in more complex design problems. A potential approach for increasing the number of parameters of the model is dividing the design space and training combinations of multiple neural networks for each partition. In the future AI networks may be trained by other AI models to automate the process, covering any possible combination of loading and geometry. Machine learning predictor models such as those developed in this study can be used to rapidly generate large quantities of stress analysis data. This may lead to a snowball effect, overcoming the considerable time requirements necessary for dataset generation that is currently the key limiting factor for widespread adoption of the technology. The datasets can then in turn be used to train future optimiser models. In addition to optimisation of material architecture, similar models could be used to optimise macro level geometry of components. Solvers could be developed based on alternative simulation techniques, training on results from thermal, electrical analysis or computational fluid dynamics. Additional objectives such as composite failure criteria, thermal regulation and fatigue resistance can be included and optimised simultaneously. In place of stress distributions models may optimise pressure distributions, minimising drag and maximising lift or downforce. One model may be used to tune aerodynamics of a component, locking the macro-scale geometry while another network tunes the micro-scale unit cell of the material to tune mechanical response without disrupting

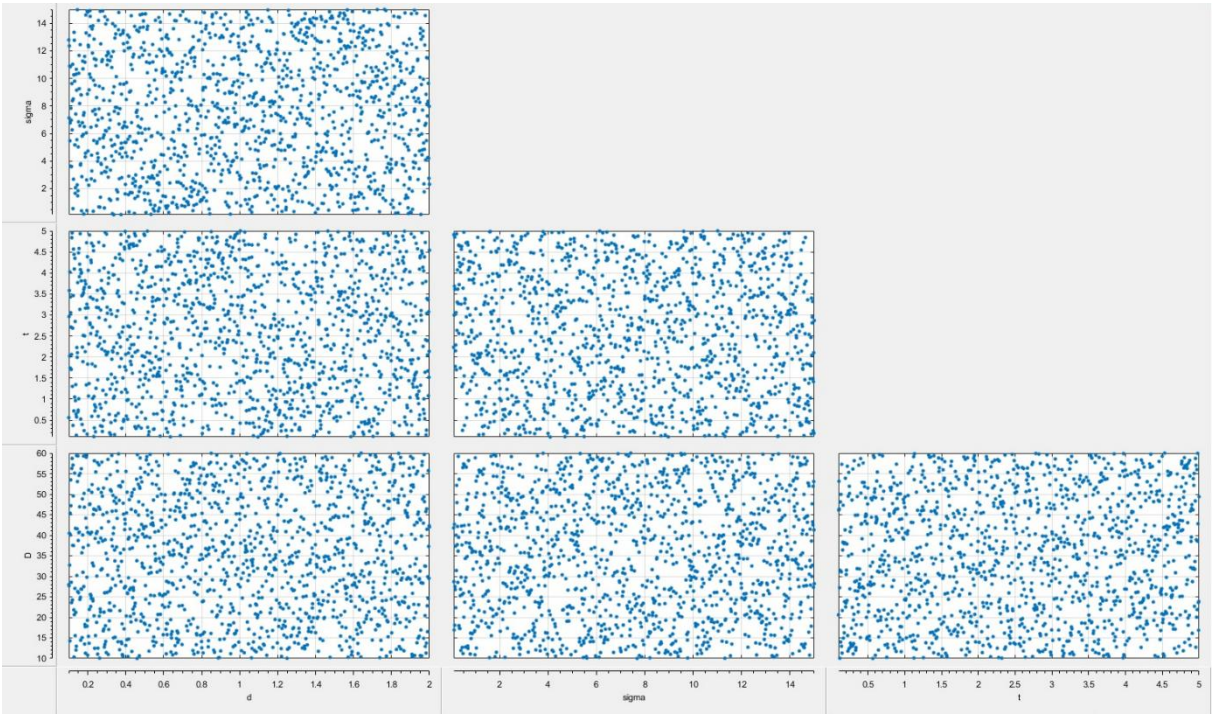
aerodynamics. Both composite layup and core material optimisations can be included in the same analysis, modelling the interactions between the composite face sheets and lattice core struts as delamination occurs. Optimisations can be performed across multiple levels within hierarchical and graded density lattice structures. Networks may be used to assign regions of stiff or compliant material within a design space for multi-material crash structures. Further studies involving explicit dynamics simulations may be validated through experimental impact testing and expanded to include consideration of strain rate effects and examination of collapse mechanisms. AI models can be used to tune active absorber systems with variable geometry or input oscillations that respond to external frequencies detected by sensors to damp sound or vibration in real time. Additional studies are required to assess the levels of network regression necessary to perform successful optimisations, particularly when used in combination with evolutionary optimisers. The current works offer strong predictive performance, however similar end results may be possible in a reduced timeframe with less accurate predictor networks. Further investigation may enable streamlining of the process workflows for real-world applications. In the future machine learning based stress analysis and design optimisation is likely to become increasingly relevant with persistent improvements in computing power, parallel processing and high performance computing, enabling continuous revolutions in design of architected materials.

Appendix A MPP CNN Model Design of Experiments Test Plans

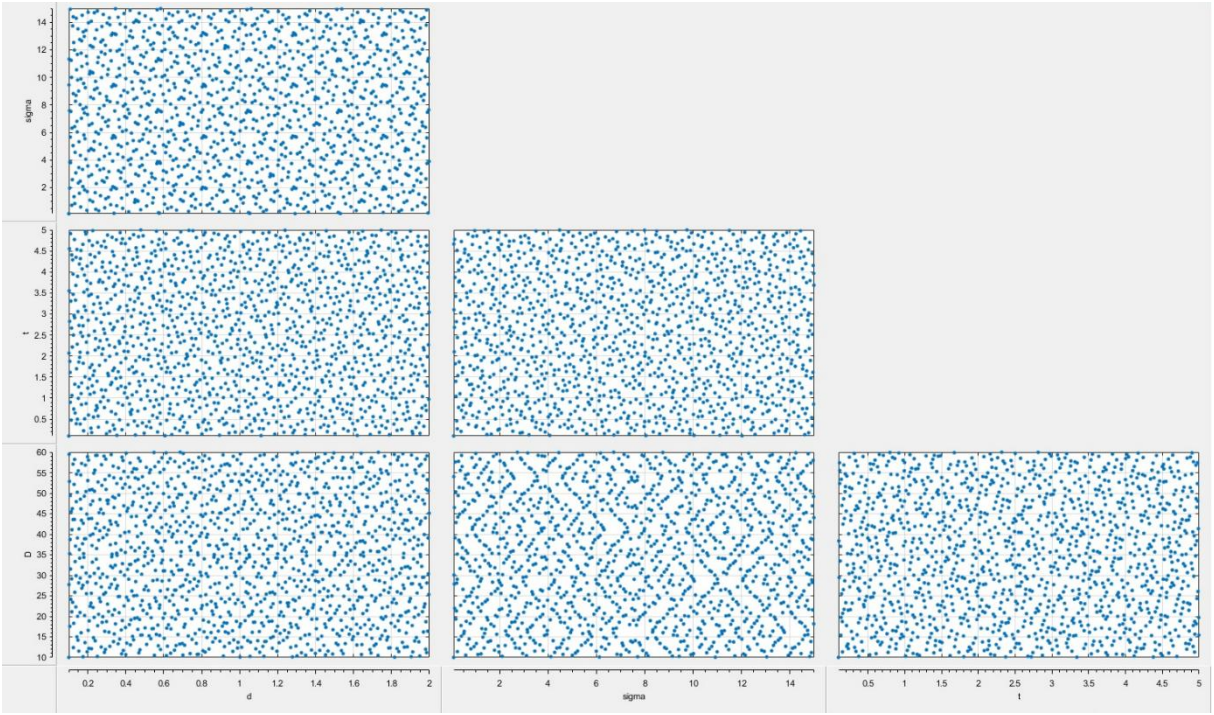
Uniform Random Distribution Pairwise Projection – 1300 Samples



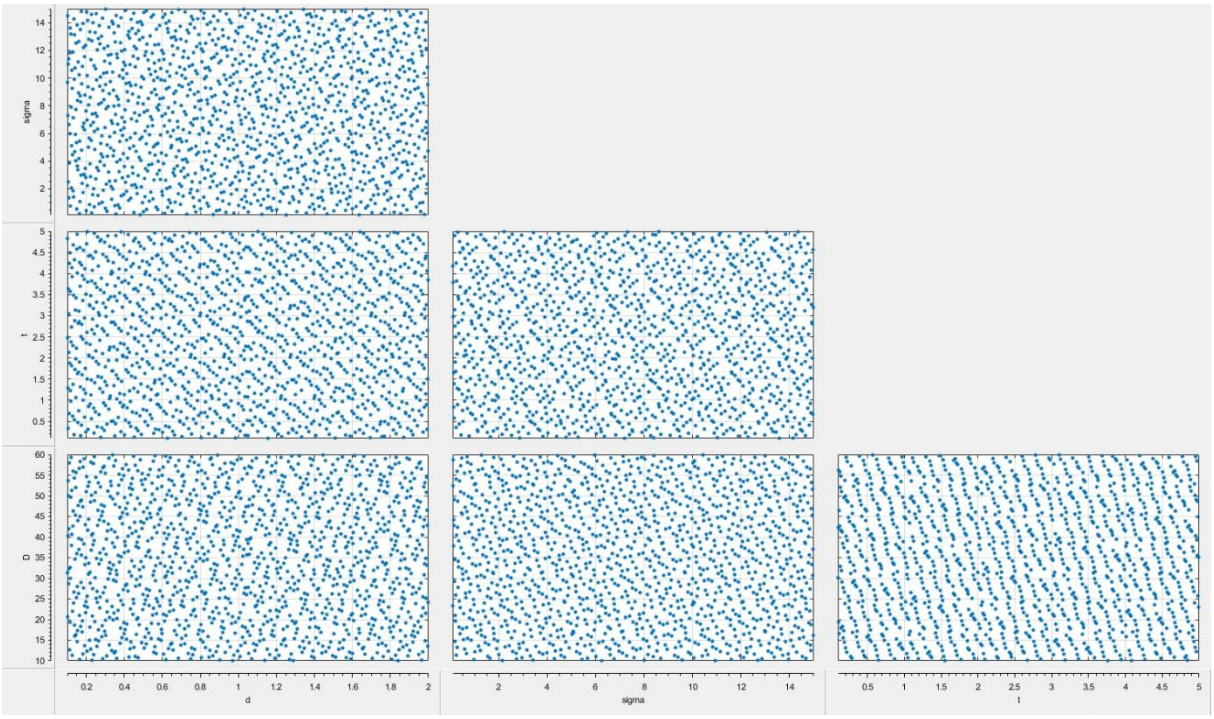
Latin Hypercube Pairwise Projection – 1300 Samples



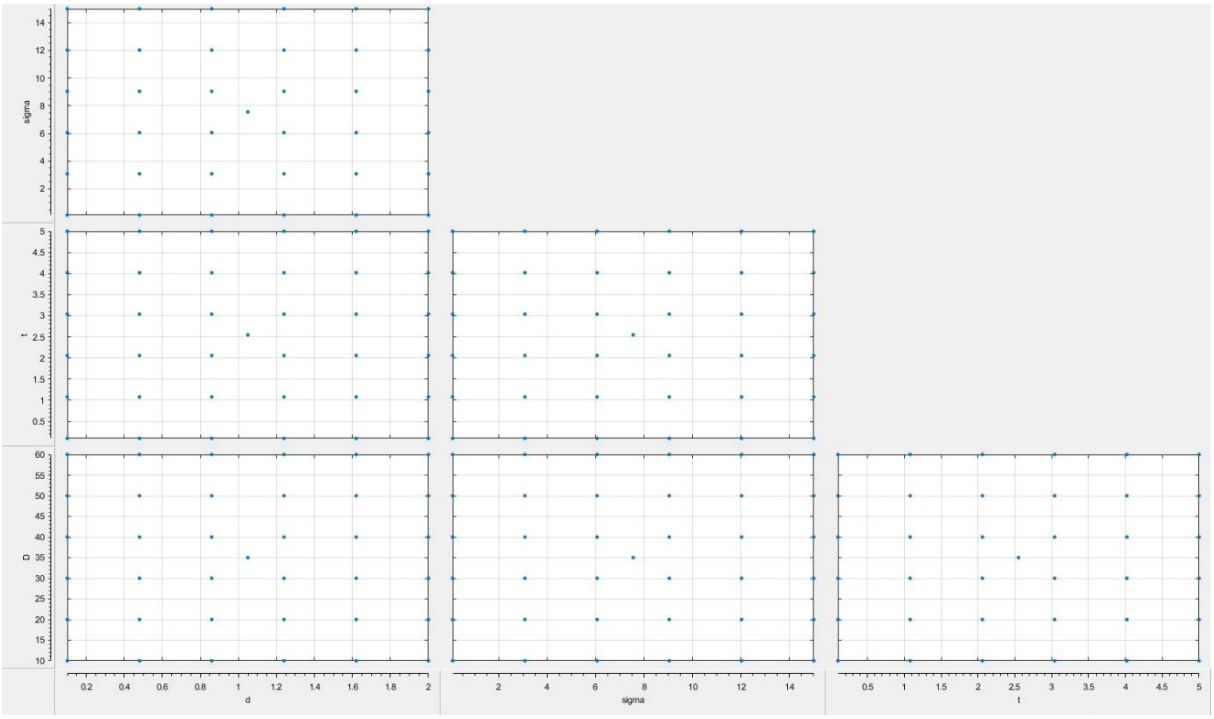
Sobol Sequence Pairwise Projection – 1300 Samples



Halton Sequence Pairwise Projection – 1300 Samples

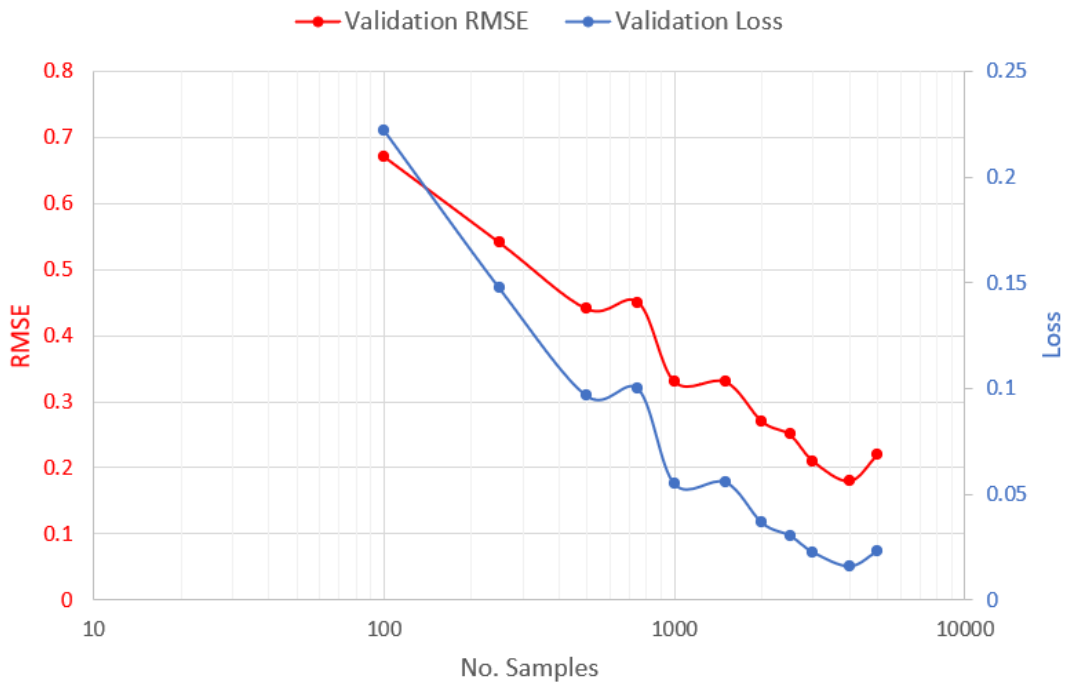


Full Factorial With 6 Levels Pairwise Projection – 1300 Samples

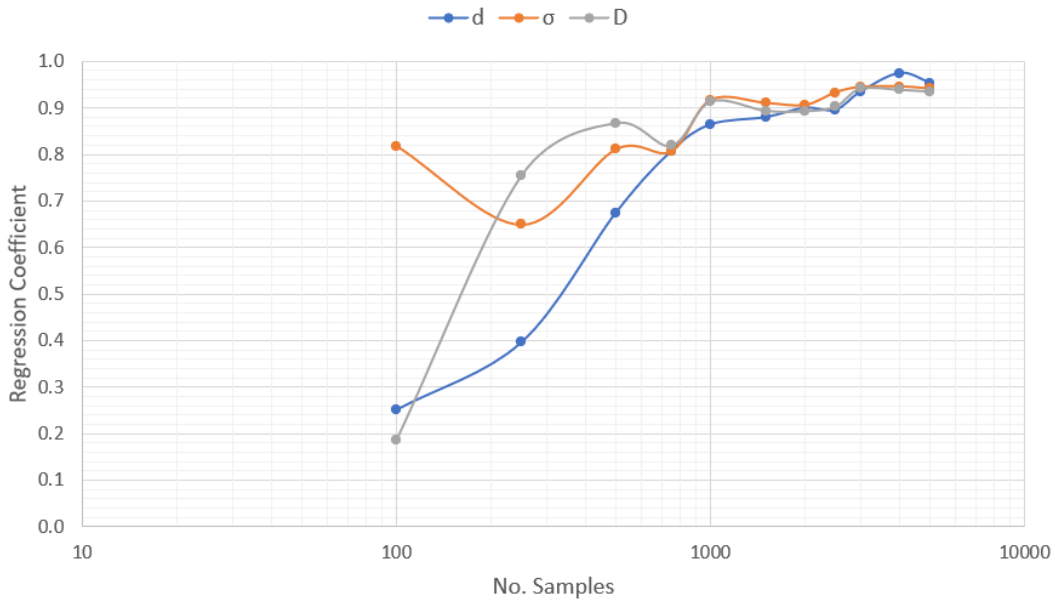


Appendix B MPP CNN Model Design of Experiments Study Validation Error, Loss and Regression Performance Plots

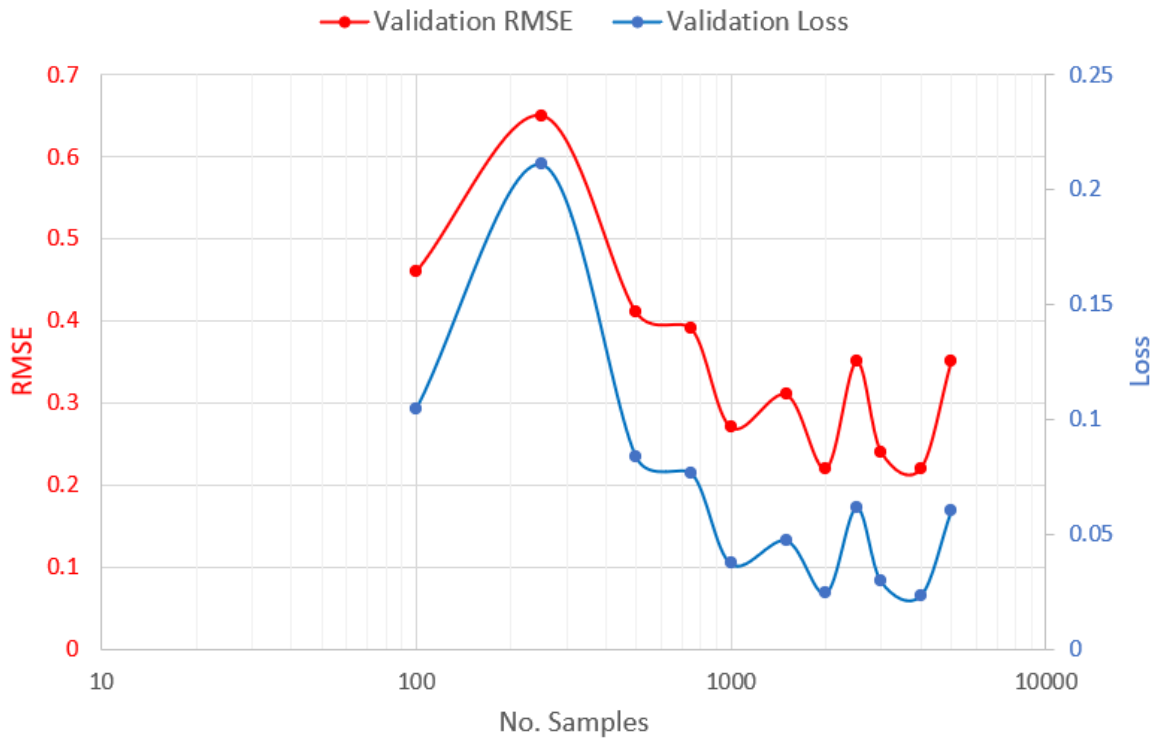
Uniform Random Distribution - Validation Root Mean-Squared Error and Loss



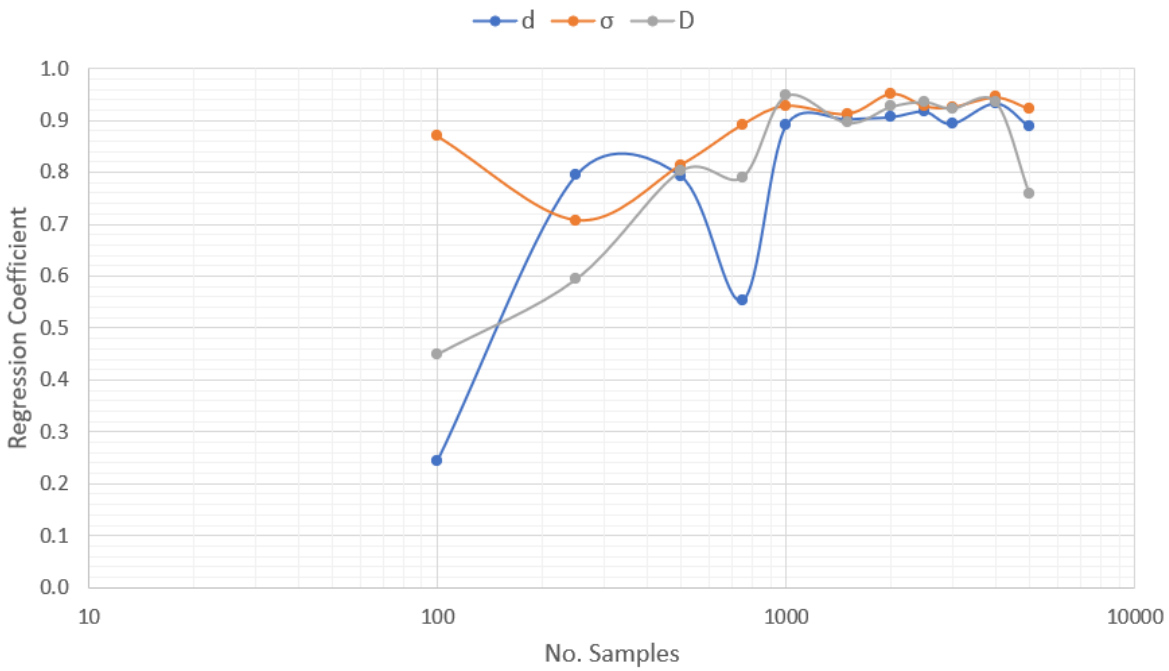
Uniform Random Distribution – Regression Coefficient



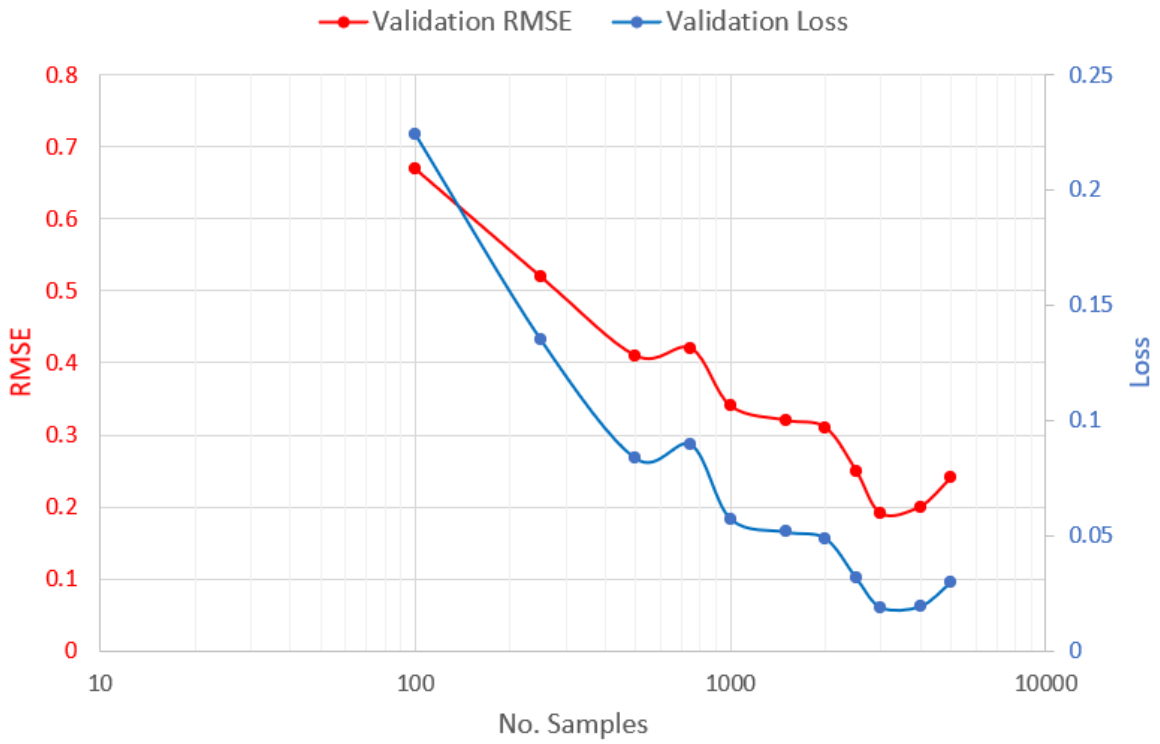
Latin Hypercube – Validation Root Mean-Squared Error And Loss



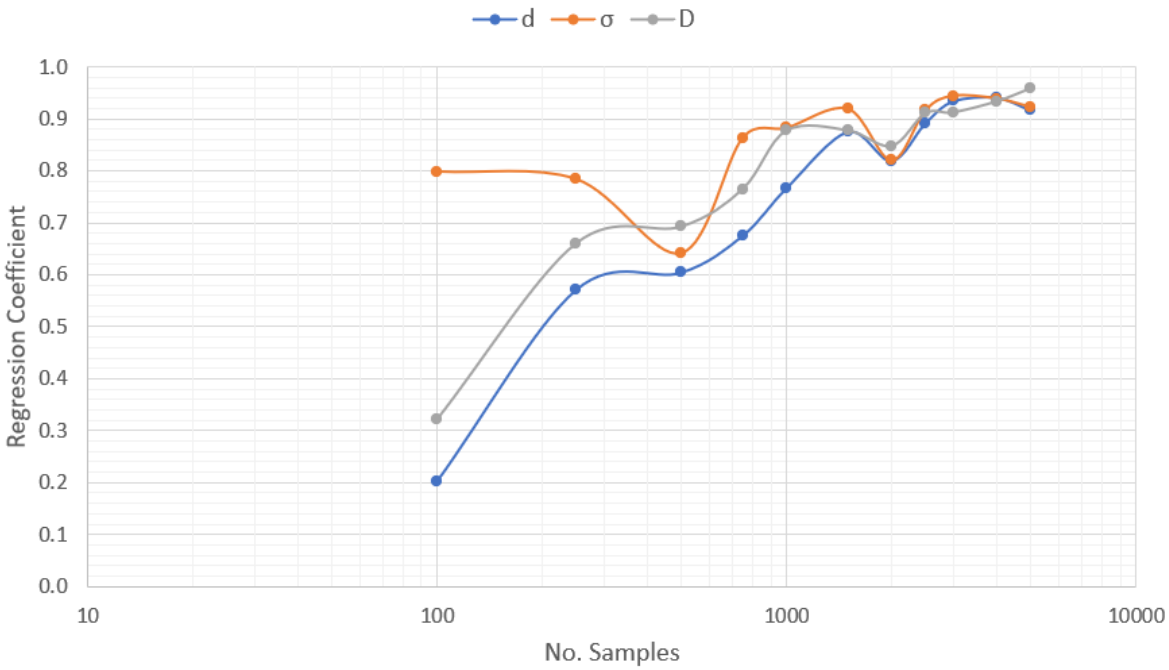
Latin Hypercube – Regression Coefficient



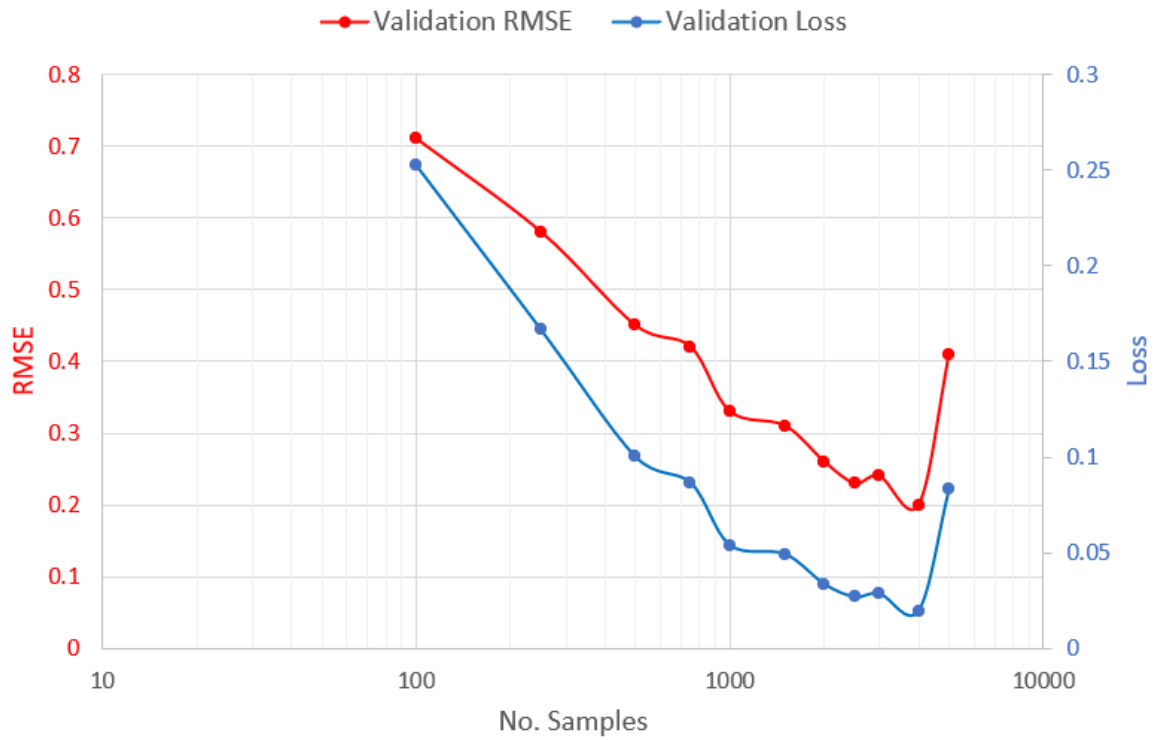
Sobol Sequence – Validation Root Mean-Squared Error And Loss



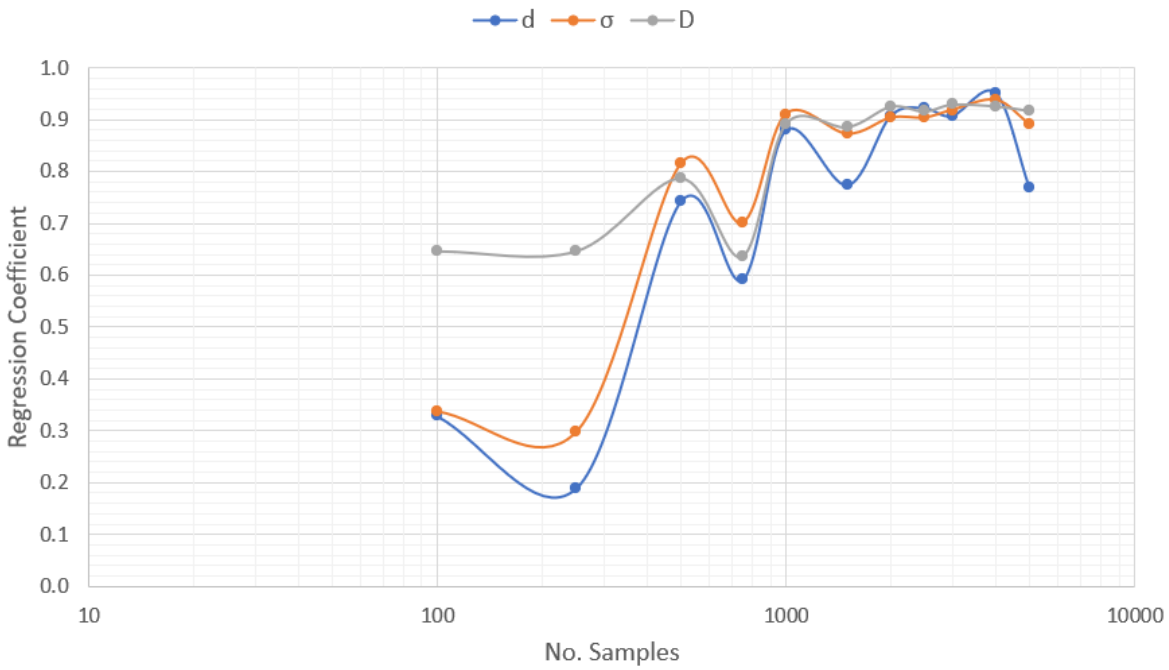
Sobol Sequence – Regression Coefficient



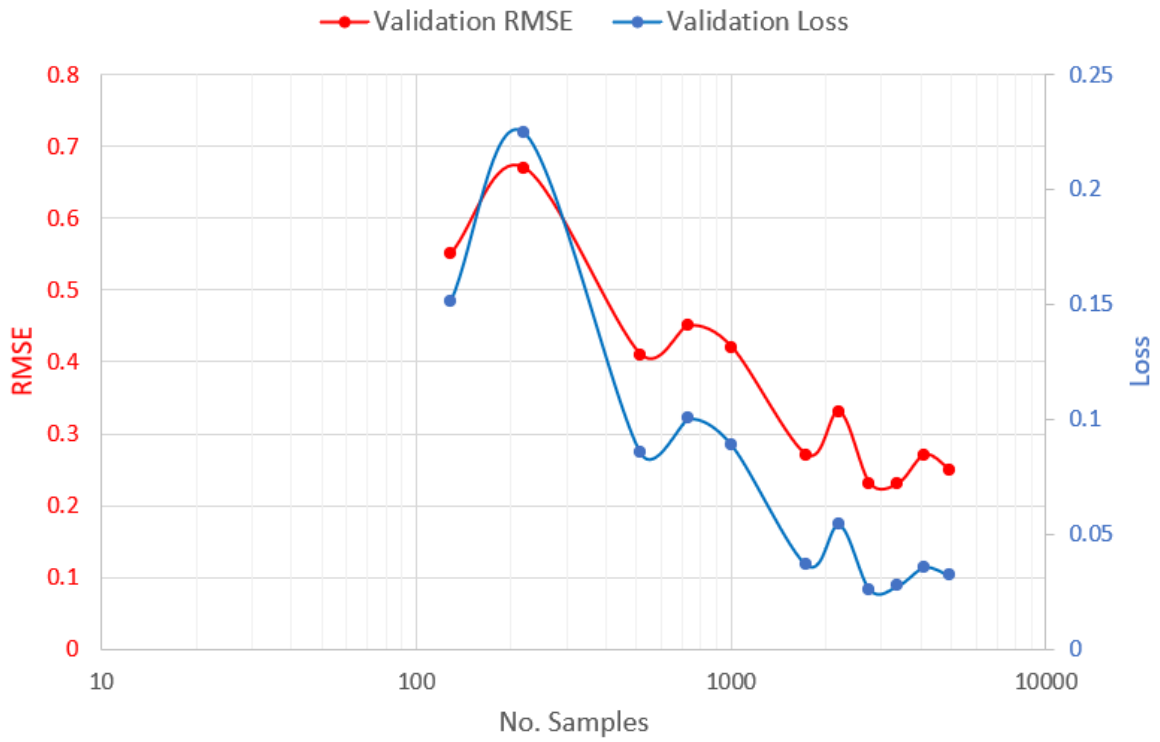
Halton Sequence – Validation Root Mean-Squared Error And Loss



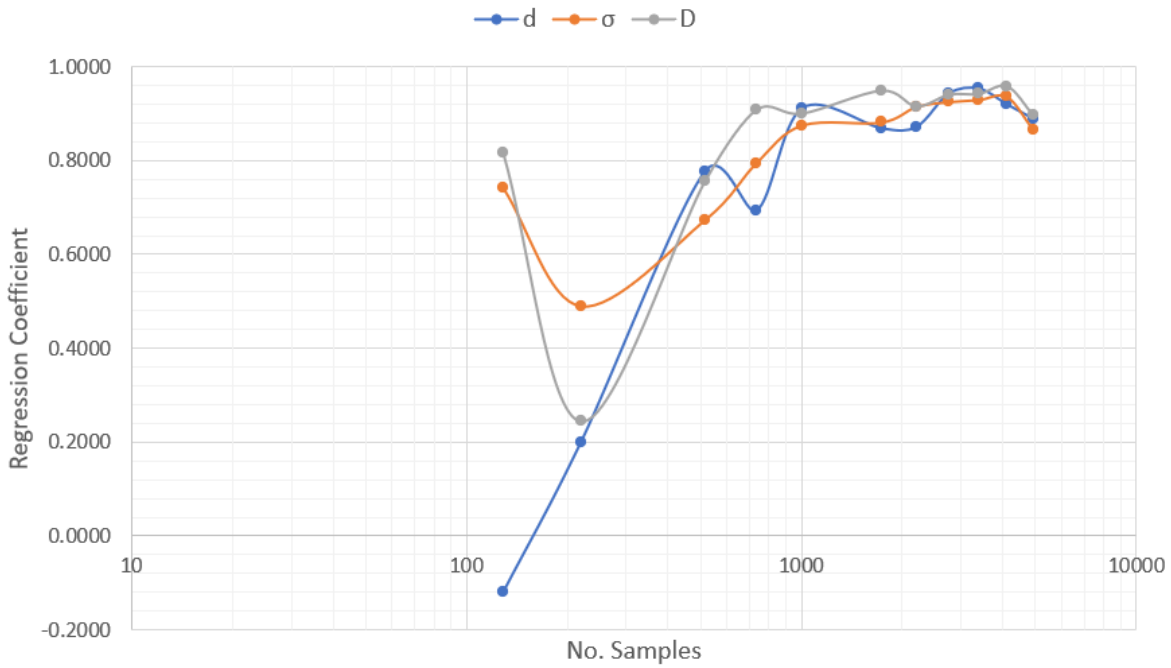
Halton Sequence – Regression Coefficient



Full Factorial With 6 Levels – Validation Root Mean-Squared Error And Loss



Full Factorial With 6 Levels – Regression Coefficient



Appendix C ASTM D638 Tensile Coupon Measurements

Material	Sample No.	Layer Height (mm)	Infill Line Orientation wrt. Loading (Degrees)	Average Gauge Width (mm)	Average Gauge Thickness (mm)	Cross Sectional Area (mm^2)
Polymaker CoPA	1	0.15	±45	12.97	3.22	41.7634
	2	0.15	±45	12.95	3.20	41.4400
	3	0.15	±45	12.97	3.20	41.5040
	4	0.15	±45	12.95	3.20	41.4400
	5	0.15	±45	12.96	3.20	41.4720
Fiberthree F3 PA-CF Pro	1	0.15	±45	13.03	3.20	41.6960
	2	0.15	±45	13.00	3.24	42.1200
	3	0.15	±45	13.00	3.28	42.6400
	4	0.15	±45	13.00	3.20	41.6000
	5	0.15	±45	13.04	3.21	41.8584
Esun ePA	1	0.15	±45	12.94	3.18	41.1492
	2	0.15	±45	12.99	3.15	40.9185
	3	0.15	±45	12.97	3.18	41.2446
	4	0.15	±45	13.01	3.20	41.6320
	5	0.15	±45	13.01	3.20	41.6320
Esun ePA-cf	1	0.15	±45	12.98	3.22	41.7956
	2	0.15	±45	12.96	3.20	41.4720
	3	0.15	±45	12.95	3.20	41.4400
	4	0.15	±45	12.94	3.18	41.1492
	5	0.15	±45	13.00	3.20	41.6000
ColorFabb XT-CF20	1	0.15	±45	13.02	3.28	42.7056
	2	0.15	±45	13.02	3.28	42.7056
	3	0.15	±45	13.06	3.26	42.5756
	4	0.15	±45	13.08	3.24	42.3792
	5	0.15	±45	13.04	3.26	42.5104

Appendix D ASTM D695 Compressive Coupon Measurements

Compression ASTM D695 - Cylindrical

Material	Sample No.	Layer Height (mm)	Build Orientation wrt. Loading (Degrees)	Average Gauge Diameter (mm)	Average Gauge Length (mm)	Cross Sectional Area (mm^2)
Polymaker CoPA	1	0.15	90	12.61	25.38	124.8878
	2	0.15	90	12.59	25.42	124.4920
	3	0.15	90	12.63	25.41	125.2843
	4	0.15	90	12.67	25.36	126.0791
	5	0.15	90	12.60	25.40	124.6898
Fiberthree F3 PA-CF Pro	1	0.15	90	12.62	25.40	125.0860
	2	0.15	90	12.59	25.40	124.4920
	3	0.15	90	12.65	25.41	125.6814
	4	0.15	90	12.69	25.40	126.4775
	5	0.15	90	12.62	25.42	125.0860
Esun ePA	1	0.15	90	12.55	25.21	123.7022
	2	0.15	90	12.54	25.21	123.5051
	3	0.15	90	12.53	25.25	123.3082
	4	0.15	90	12.56	25.26	123.8994
	5	0.15	90	12.58	25.24	124.2943
Esun ePA-cf	1	0.15	90	12.59	25.31	124.4920
	2	0.15	90	12.61	25.39	124.8878
	3	0.15	90	12.64	25.39	125.4828
	4	0.15	90	12.63	25.39	125.2843
	5	0.15	90	12.66	25.41	125.8802
ColorFabb XT-CF20	1	0.15	90	12.70	25.42	126.6769
	2	0.15	90	12.69	25.41	126.4775
	3	0.15	90	12.70	25.38	126.6769
	4	0.15	90	12.69	25.40	126.4775
	5	0.15	90	12.69	25.44	126.4775

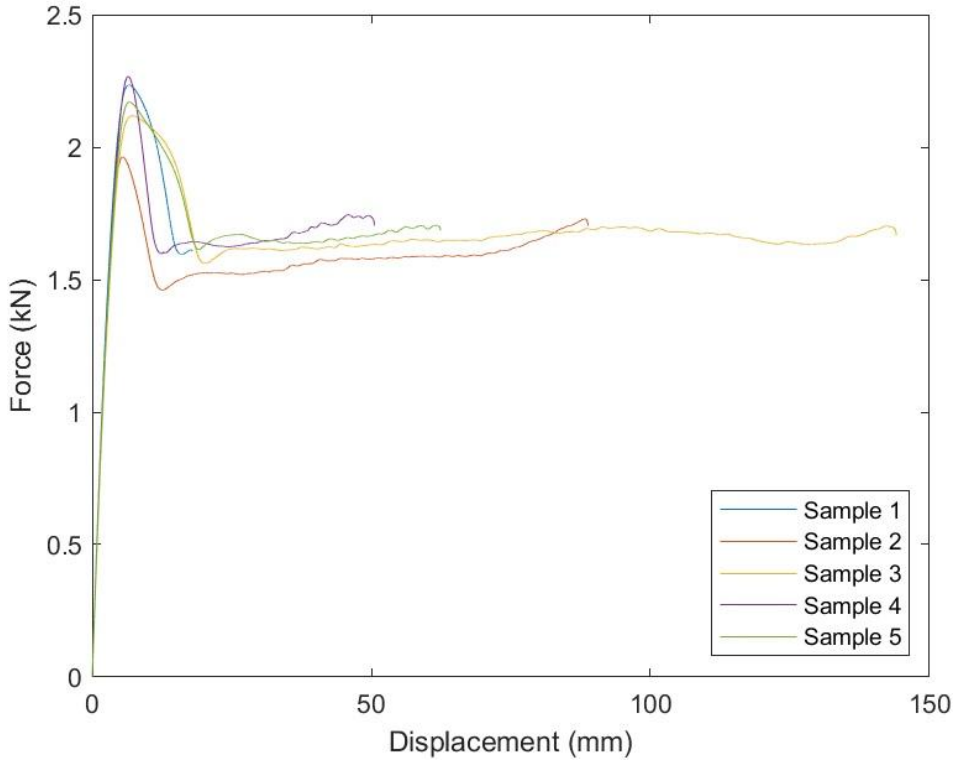
Compression ASTM D695 - Cuboid

Material	Build No.	Sample No.	Layer Height (mm)	Build Orientation wrt. Loading (Degrees)	Average Gauge Width (mm)	Average Gauge Depth (mm)	Average Gauge Length (mm)	Cross Sectional Area (mm^2)
Polymaker CoPA	1	1	0.15	0	12.70	12.56	25.22	159.5120
		2	0.15	0	12.82	12.47	25.36	159.8654
		3	0.15	0	12.76	12.54	25.42	160.0104
		4	0.15	0	12.84	12.64	25.47	162.2976
		5	0.15	0	12.85	12.60	25.43	161.9100
Polymaker CoPA	2	1	0.15	45	12.71	12.68	25.39	161.1628
		2	0.15	45	12.69	12.70	25.39	161.1630
		3	0.15	45	12.71	12.70	25.42	161.4170
		4	0.15	45	12.69	12.66	25.43	160.6554
		5	0.15	45	12.72	12.67	25.43	161.1624
Polymaker CoPA	3	1	0.15	90	12.68	12.75	25.35	161.6700
		2	0.15	90	12.76	12.71	25.34	162.1796
		3	0.15	90	12.70	12.70	25.37	161.2900
		4	0.15	90	12.70	12.71	25.35	161.4170
		5	0.15	90	12.70	12.70	25.38	161.2900

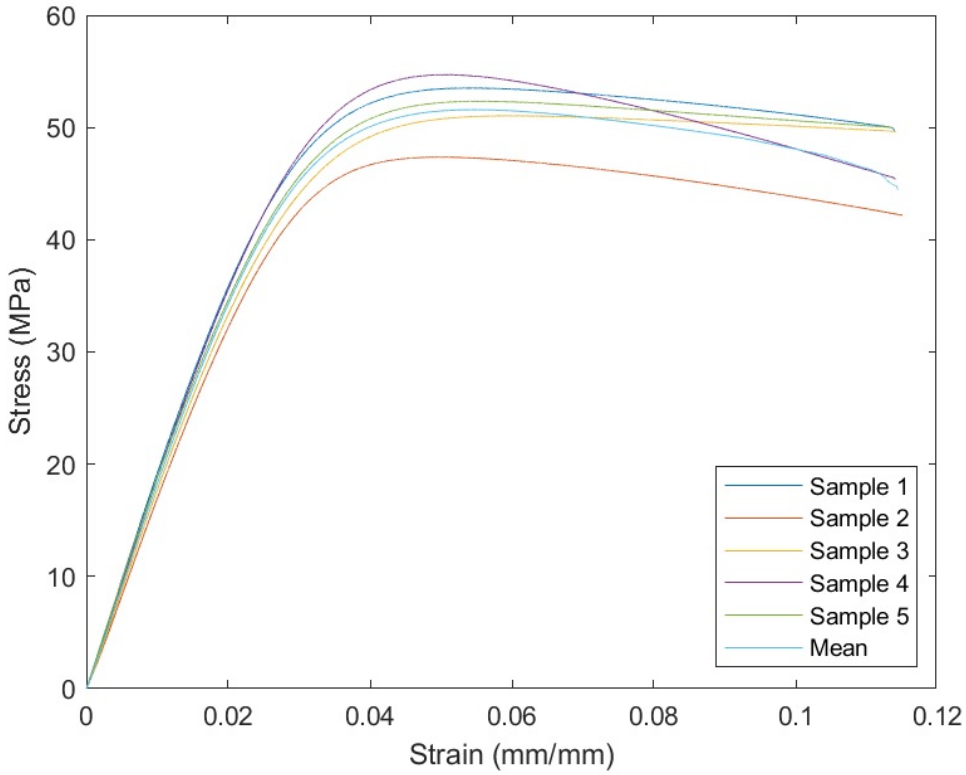
Appendix E ASTM D638 Tensile Test Results

Polymaker CoPA

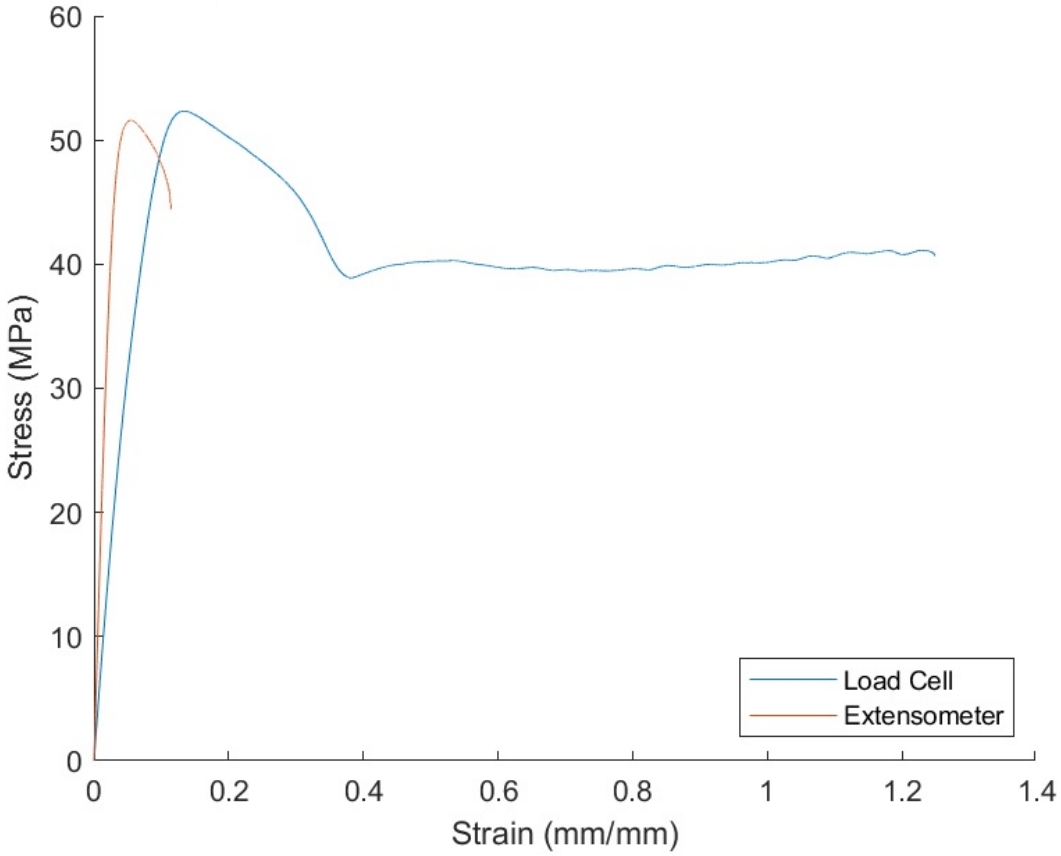
All Results – Load Cell



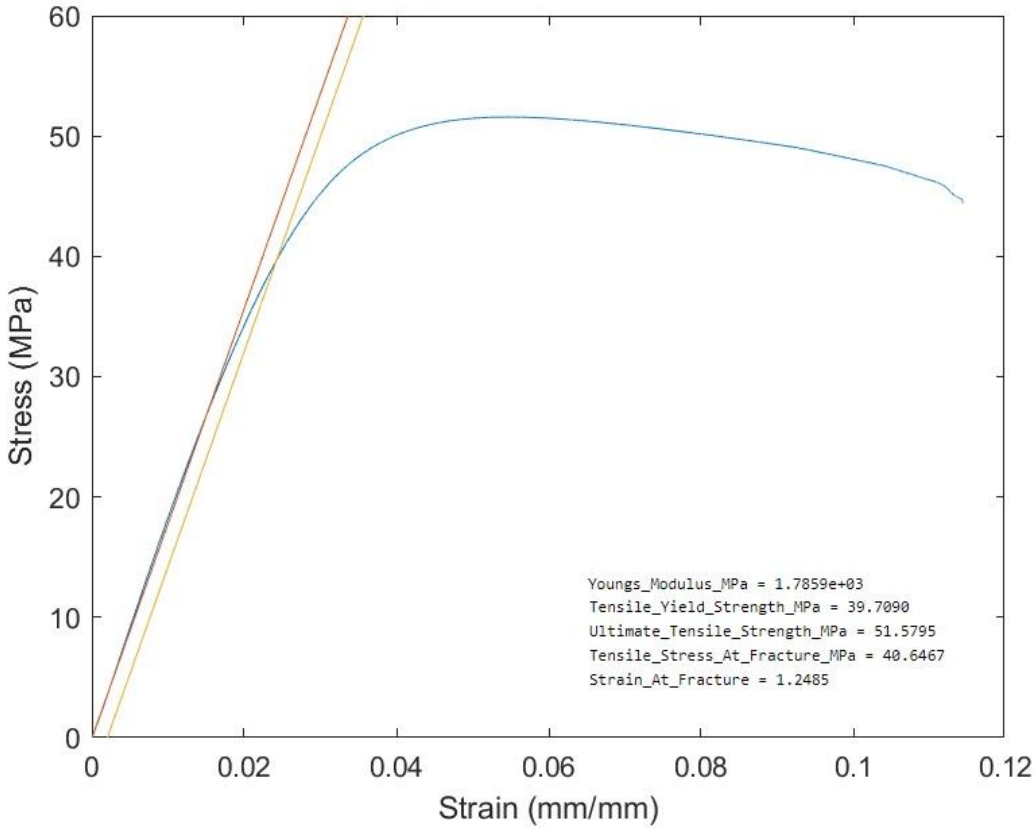
All Results – Extensometer



Mean Response – Load Cell Vs Extensometer

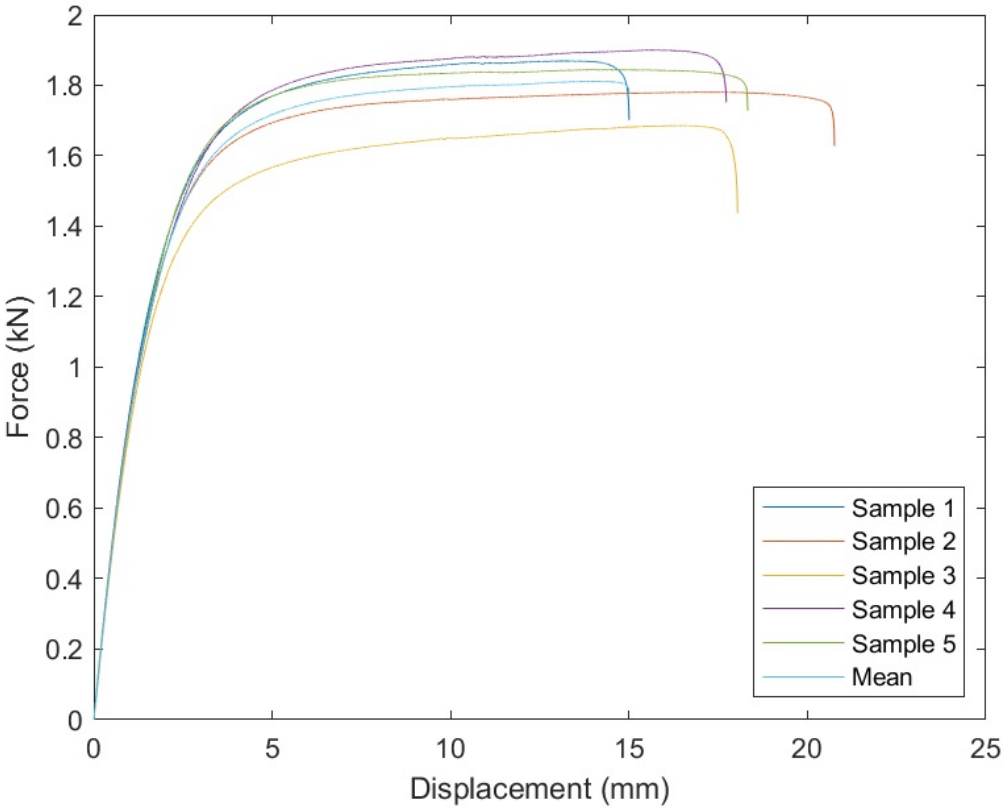


0.2% Offset Proof Stress Calculation - Extensometer

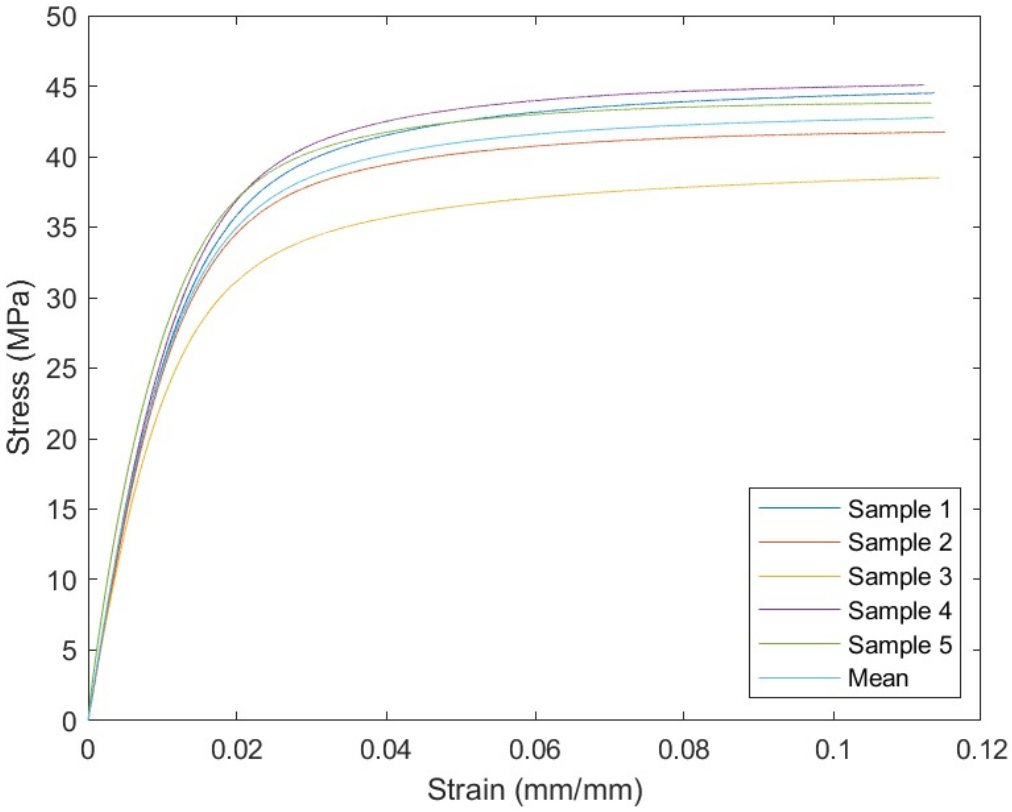


Fiberthree F3 PA-CF Pro

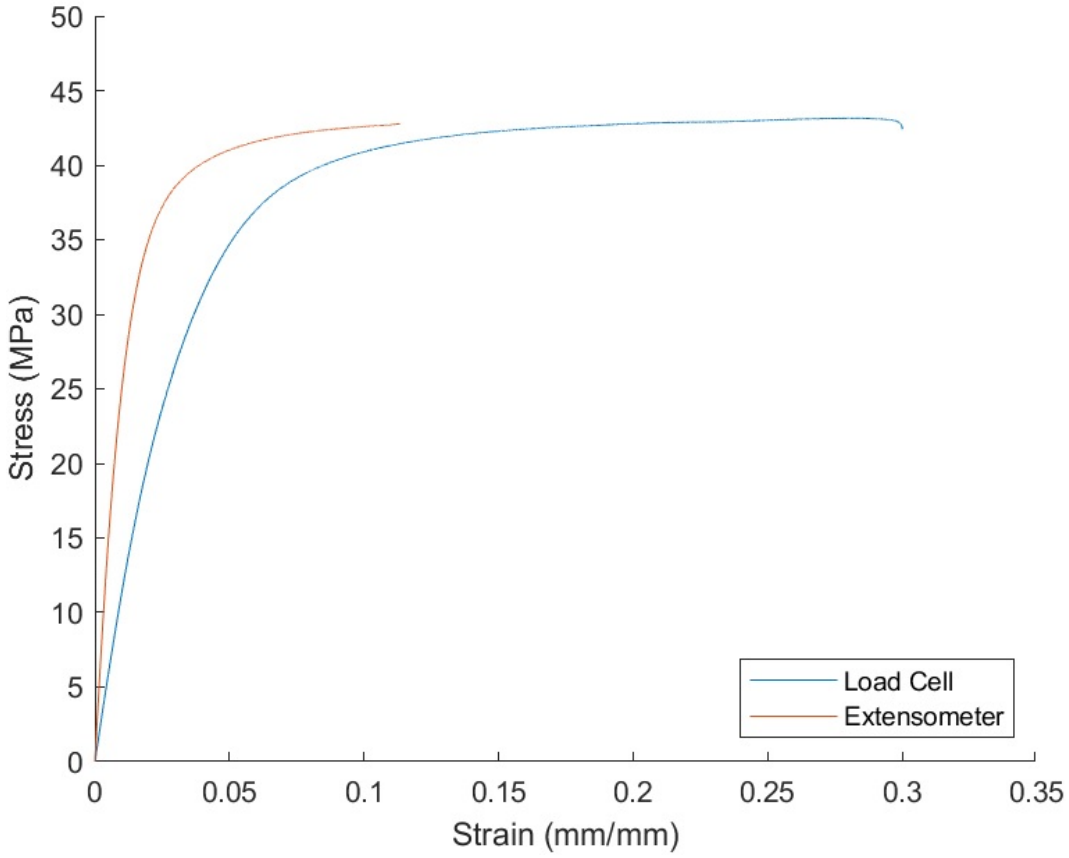
All Results – Load Cell



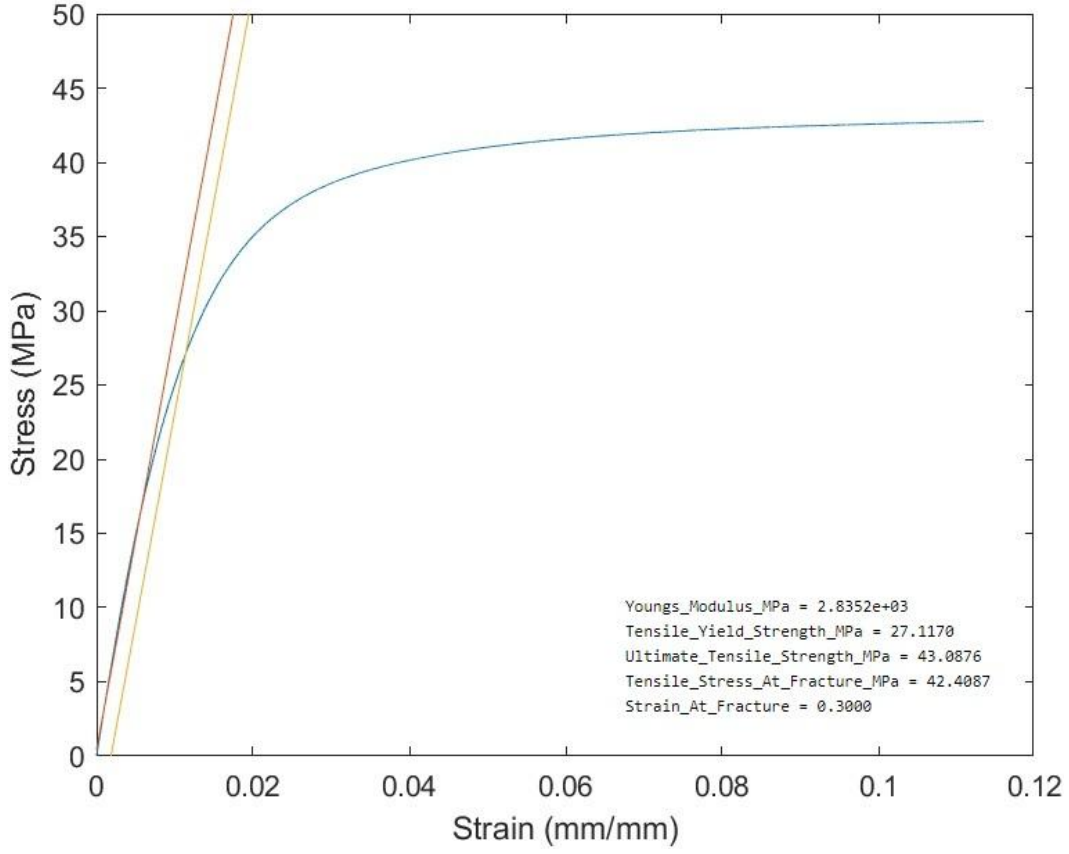
All Results – Extensometer



Mean Response – Load Cell Vs Extensometer

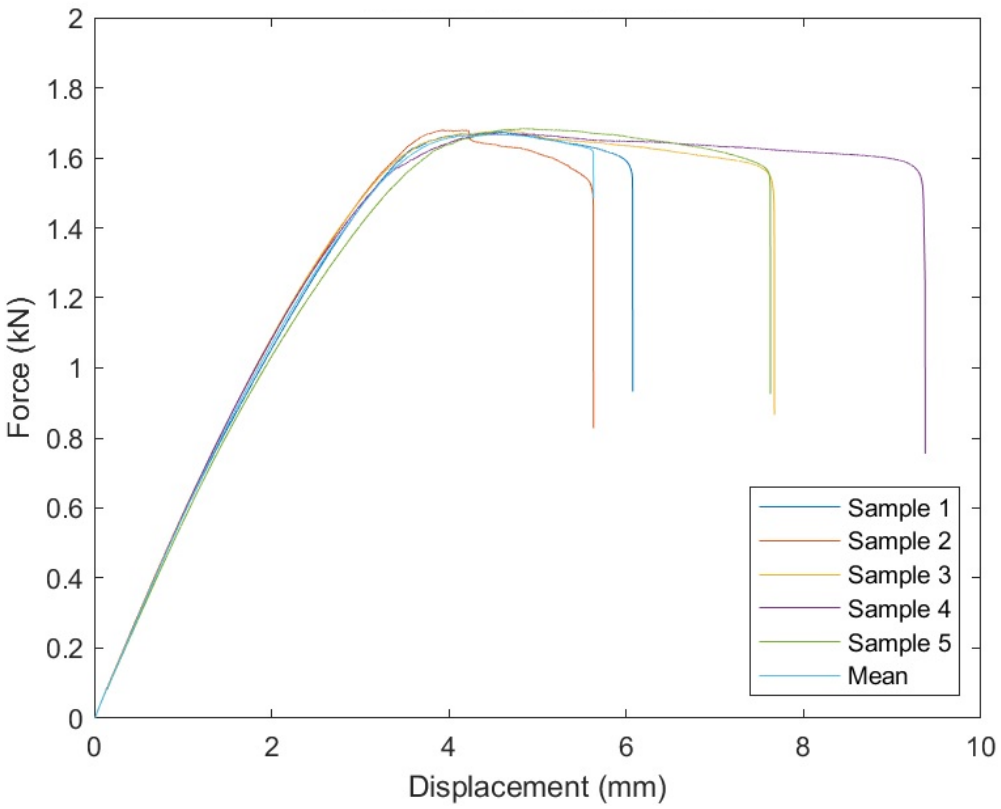


0.2% Offset Proof Stress Calculation - Extensometer

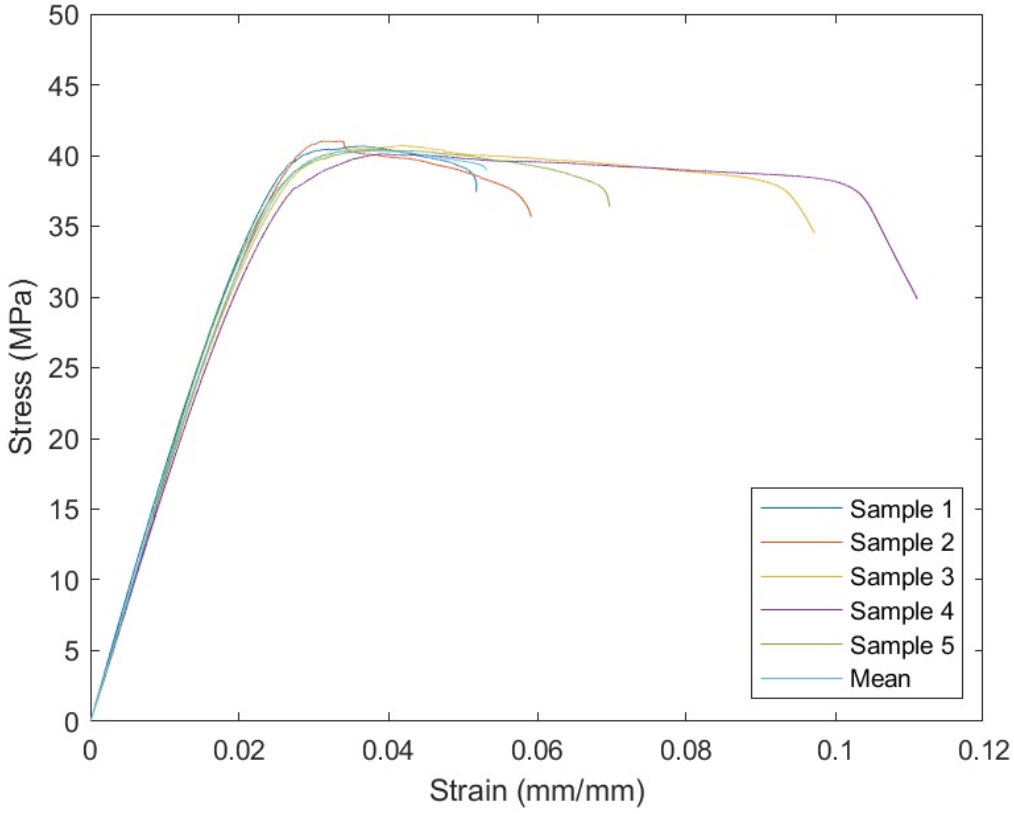


Esun ePA

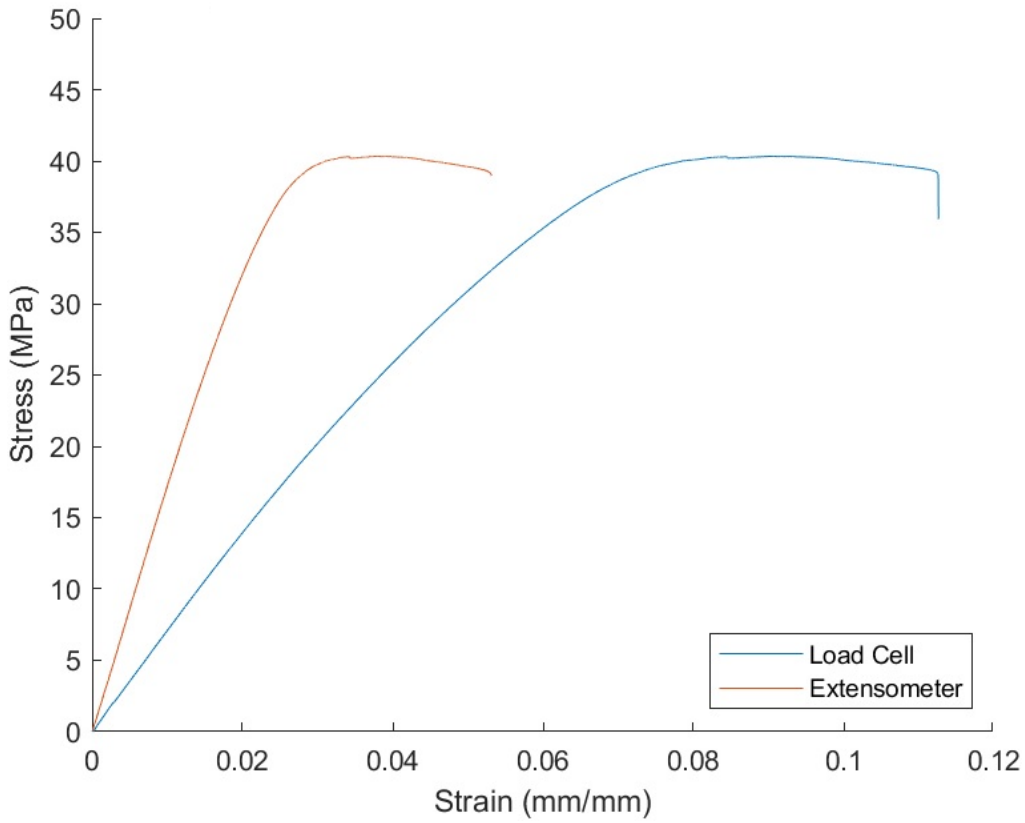
All Results – Load Cell



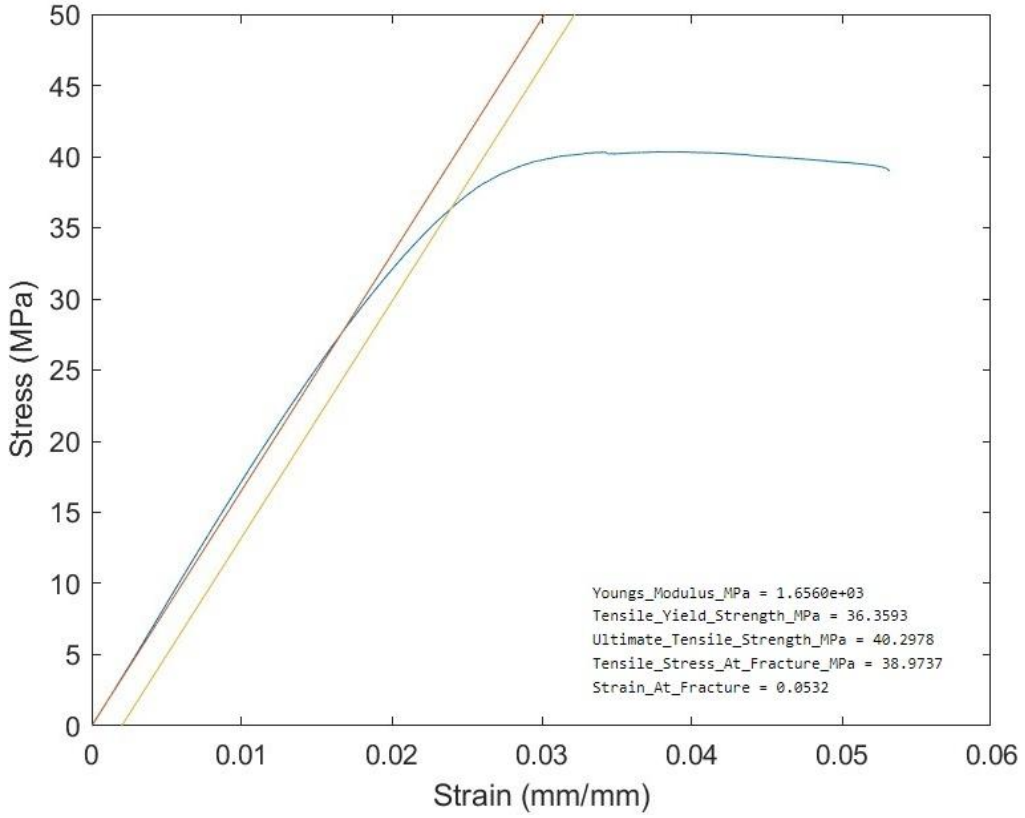
All Results – Extensometer



Mean Response – Load Cell Vs Extensometer

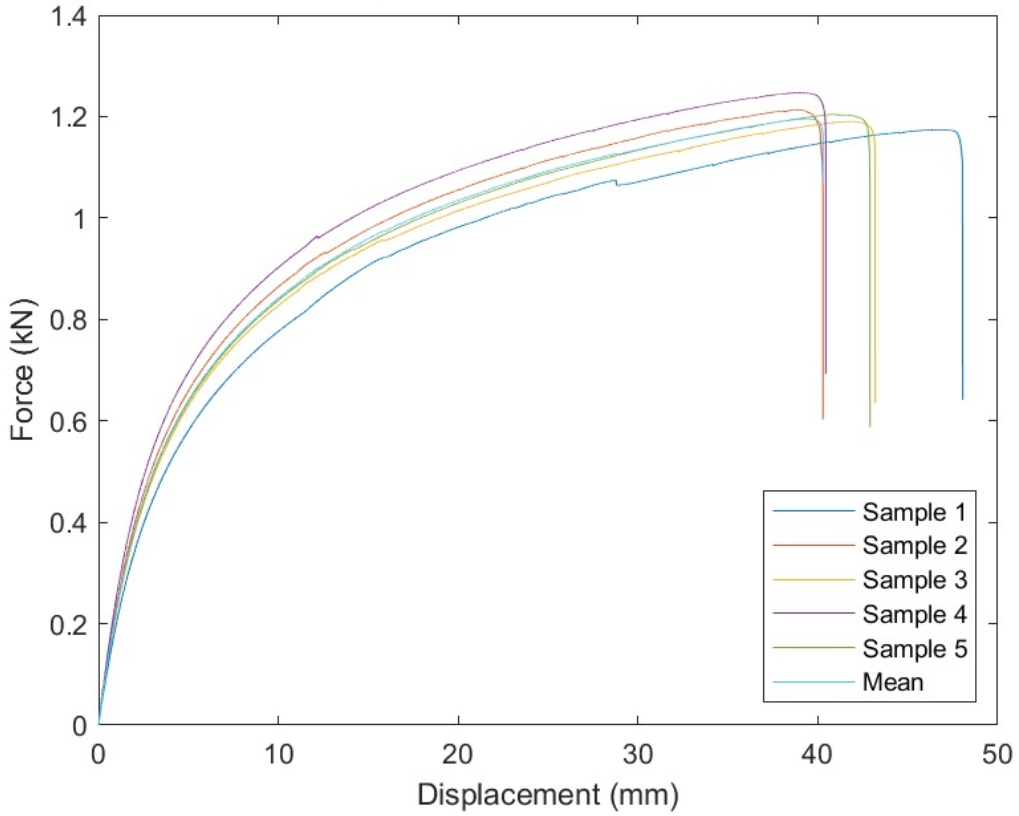


0.2% Offset Proof Stress Calculation – Extensometer

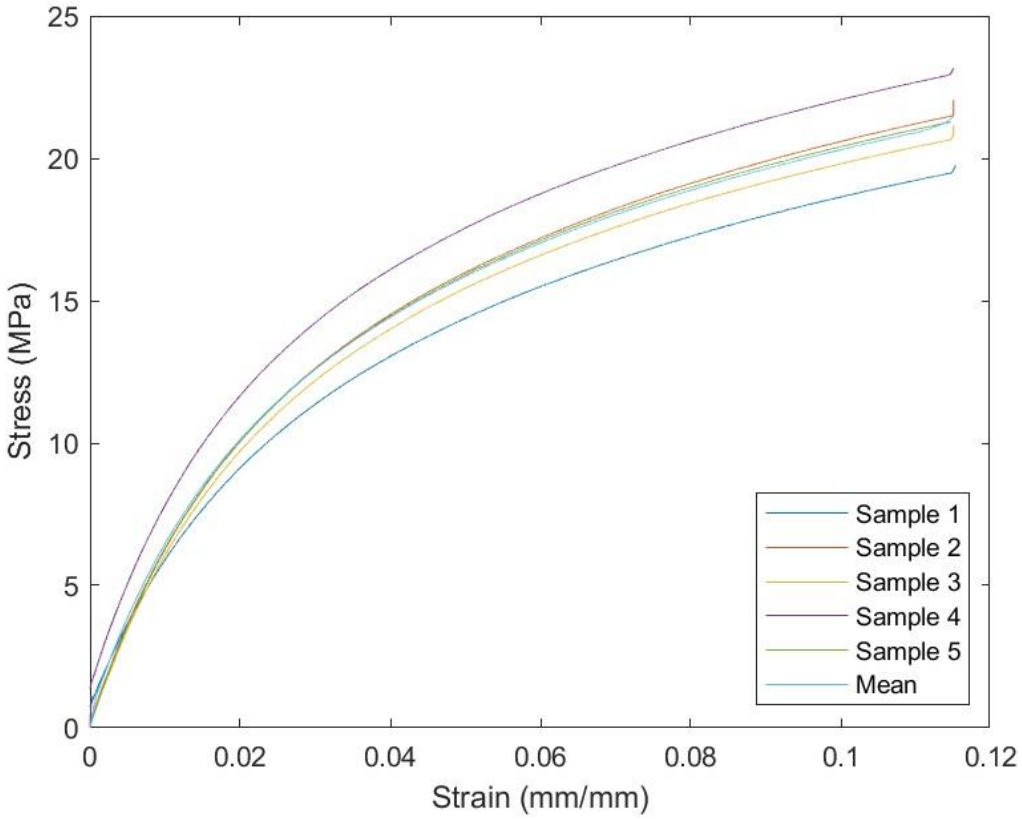


Esun ePA-CF

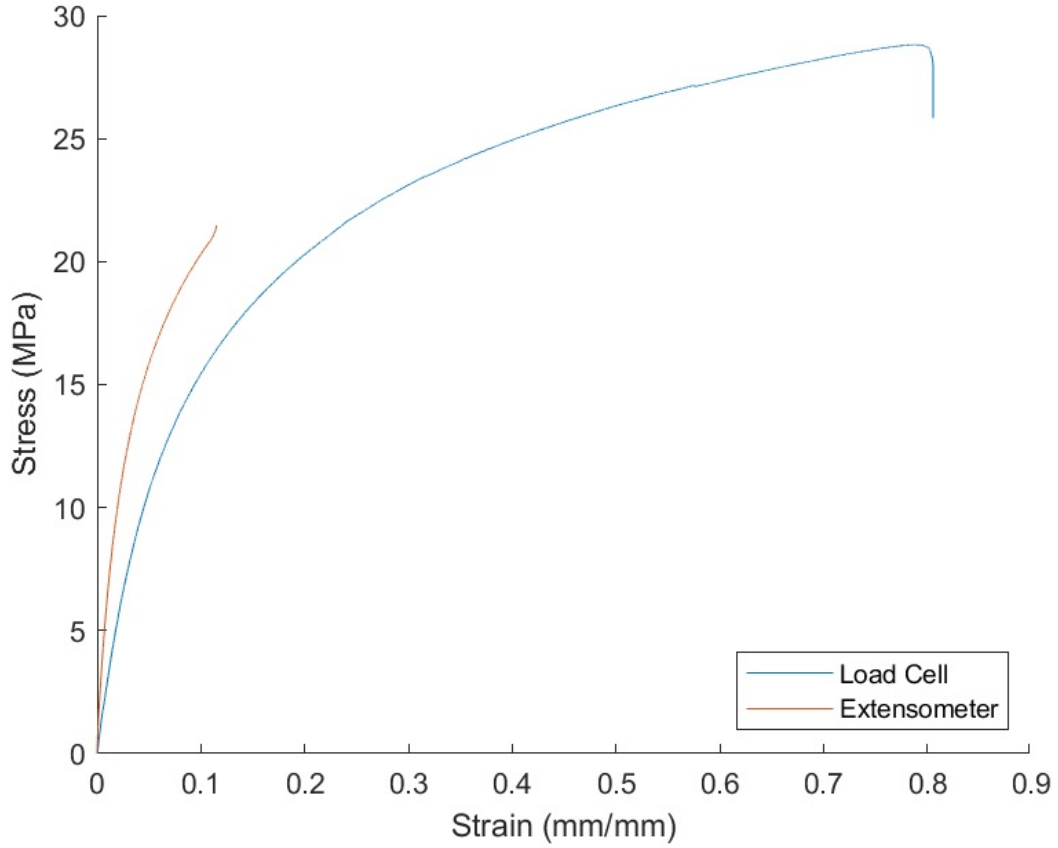
All Results – Load Cell



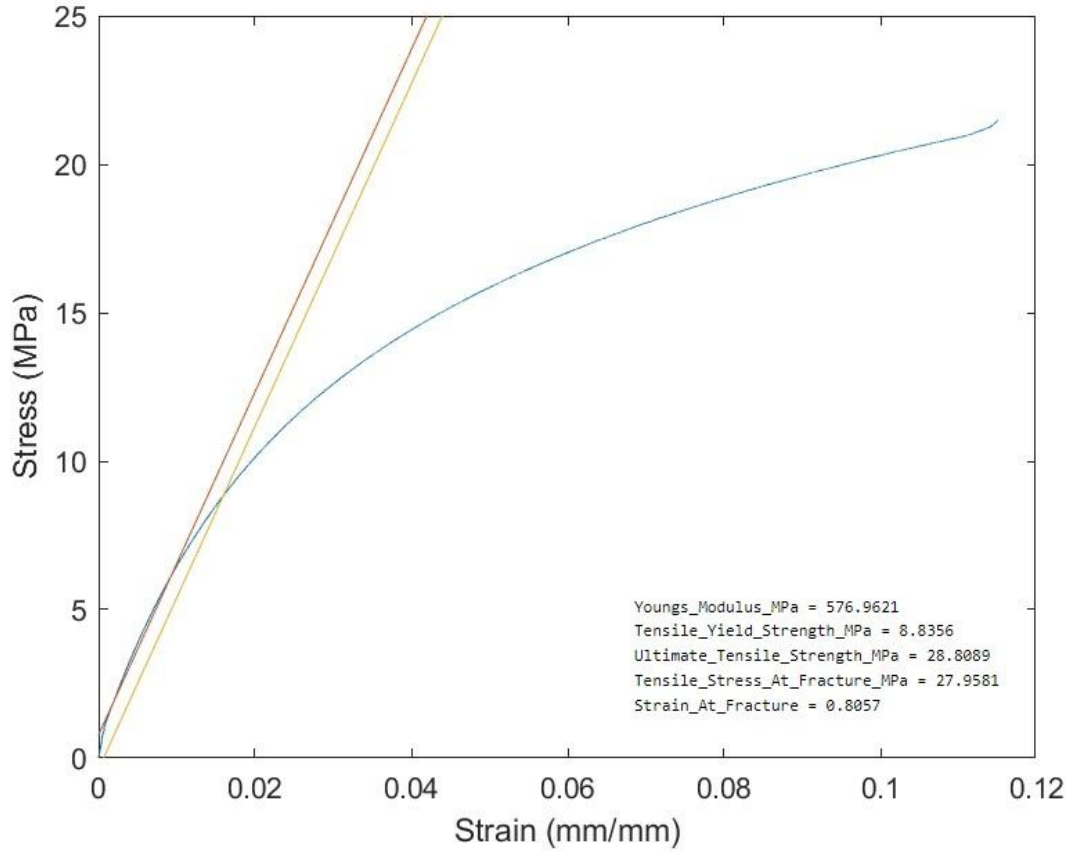
All Results – Extensometer



Mean Response – Load Cell Vs Extensometer

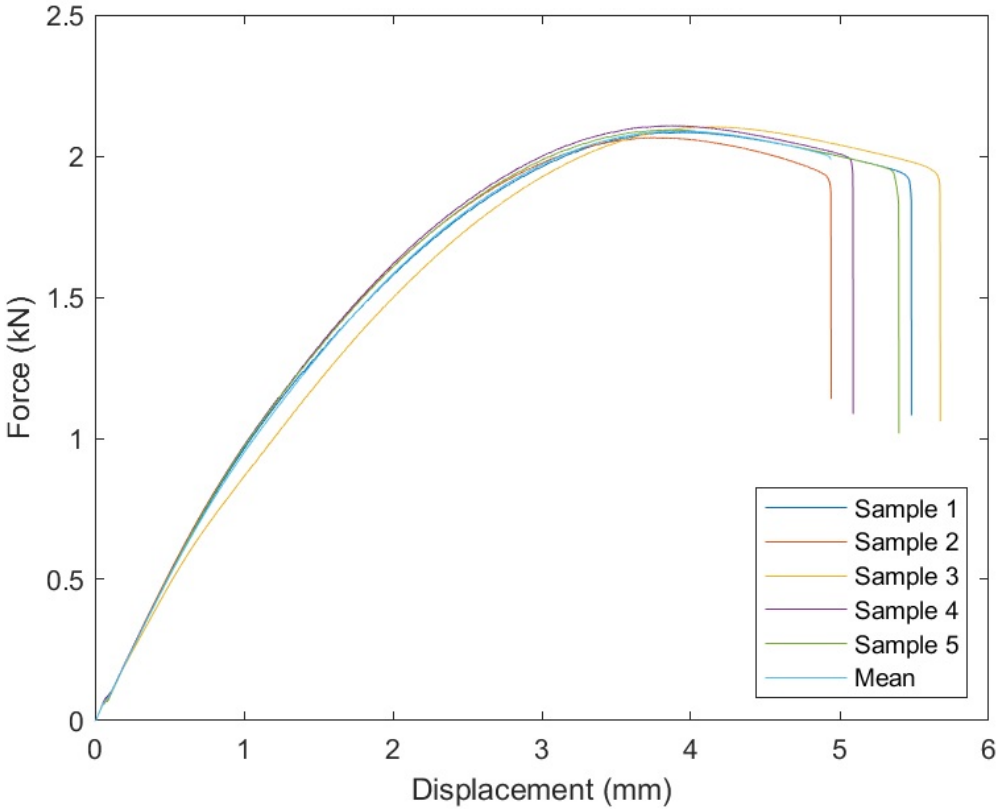


0.2% Offset Proof Stress Calculation – Extensometer

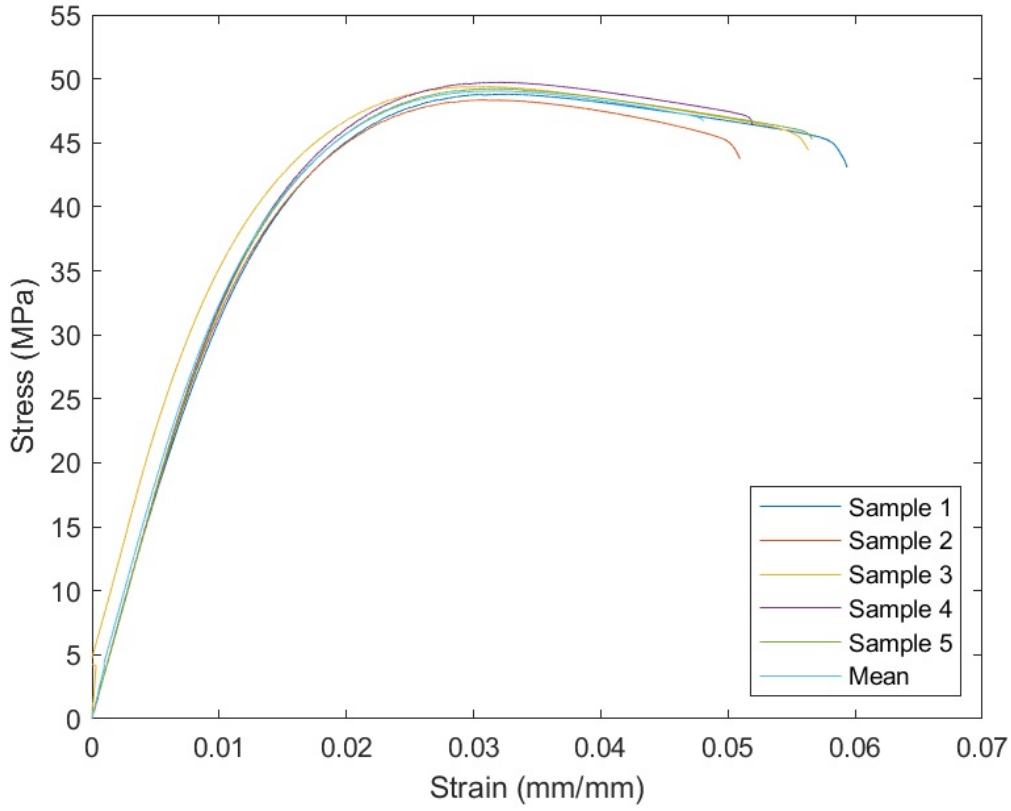


ColorFabb XT-CF20

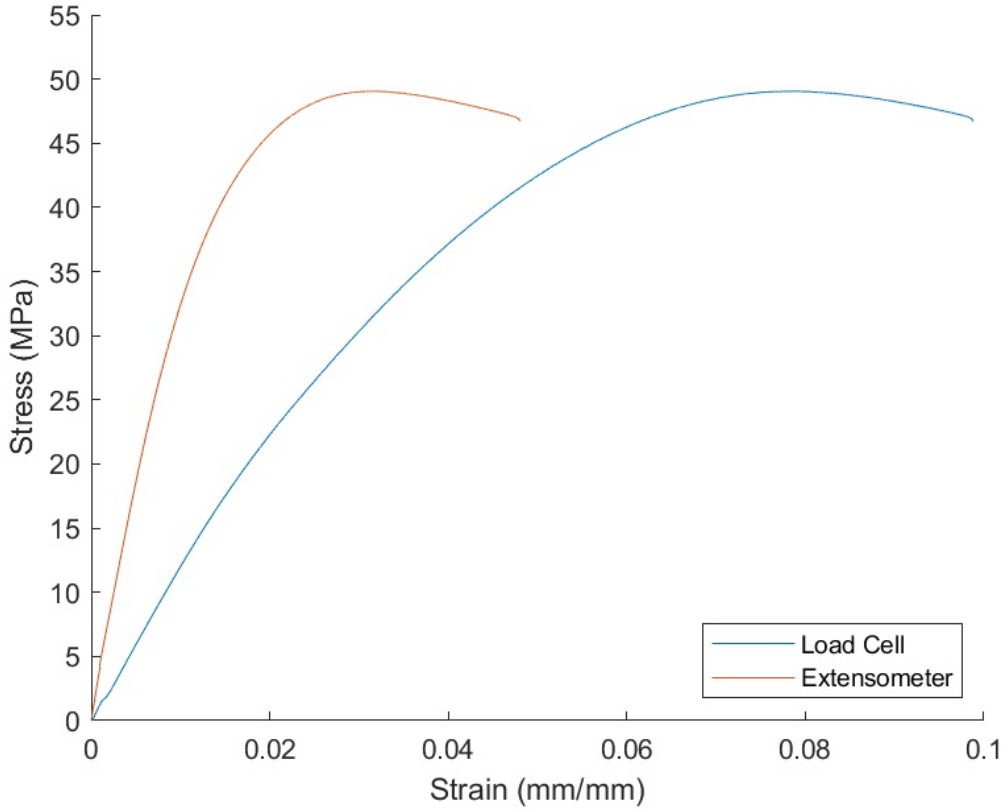
All Results – Load Cell



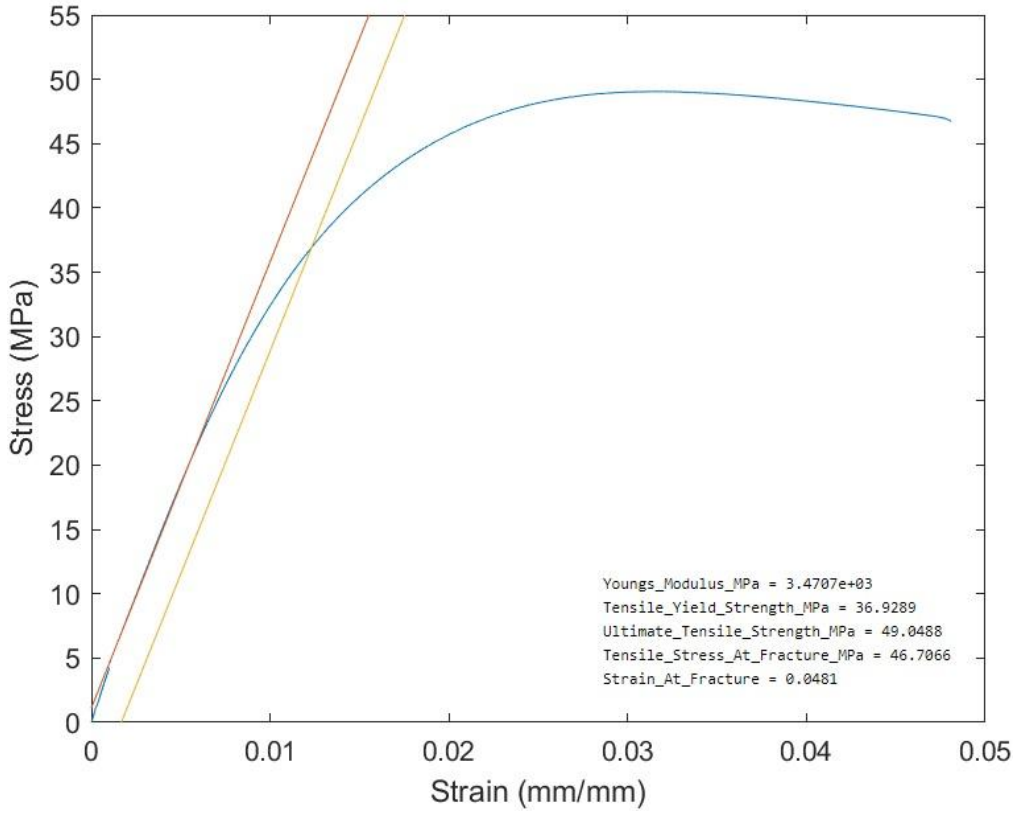
All Results – Extensometer



Mean Response – Load Cell Vs Extensometer



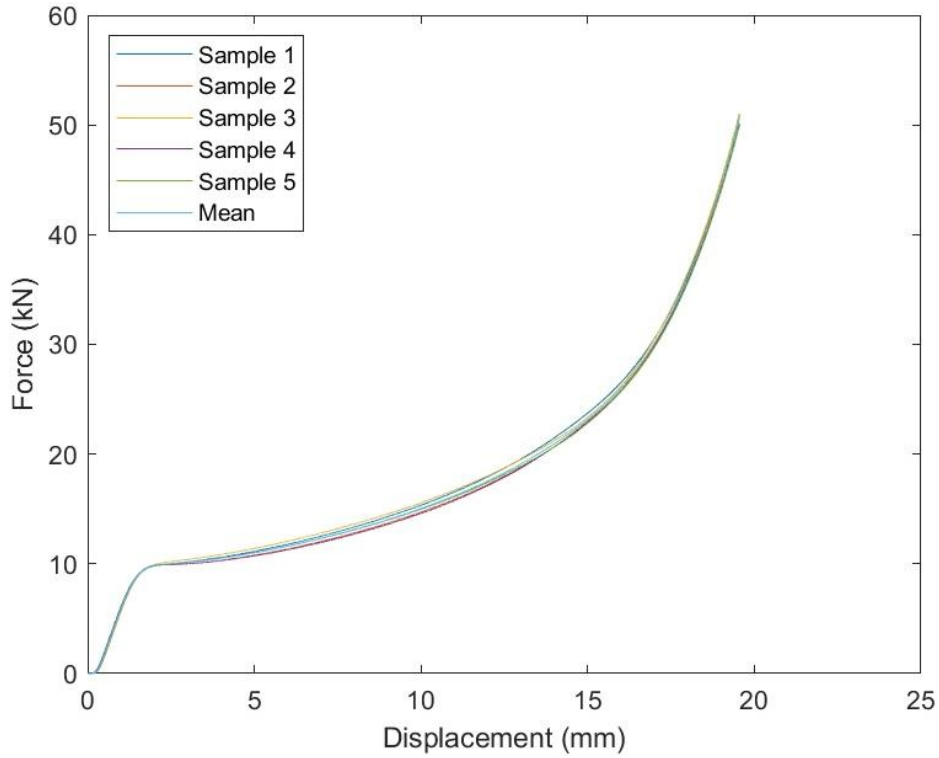
0.2% Offset Proof Stress Calculation - Extensometer



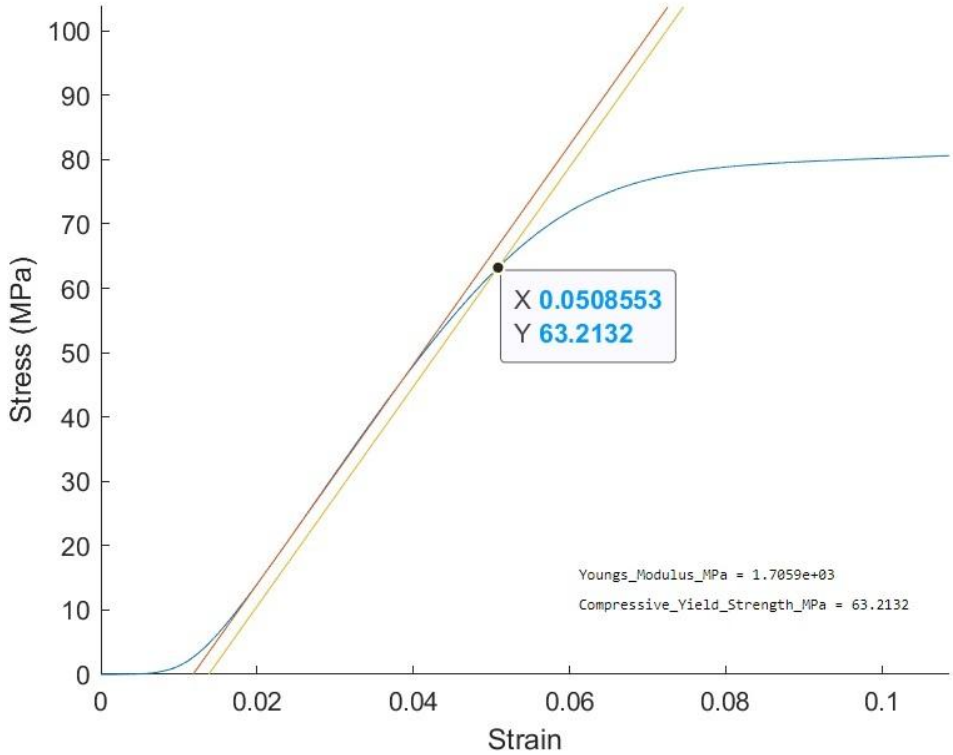
Appendix F ASTM D695 Compressive Test Results

Polymaker CoPA - Cylindrical

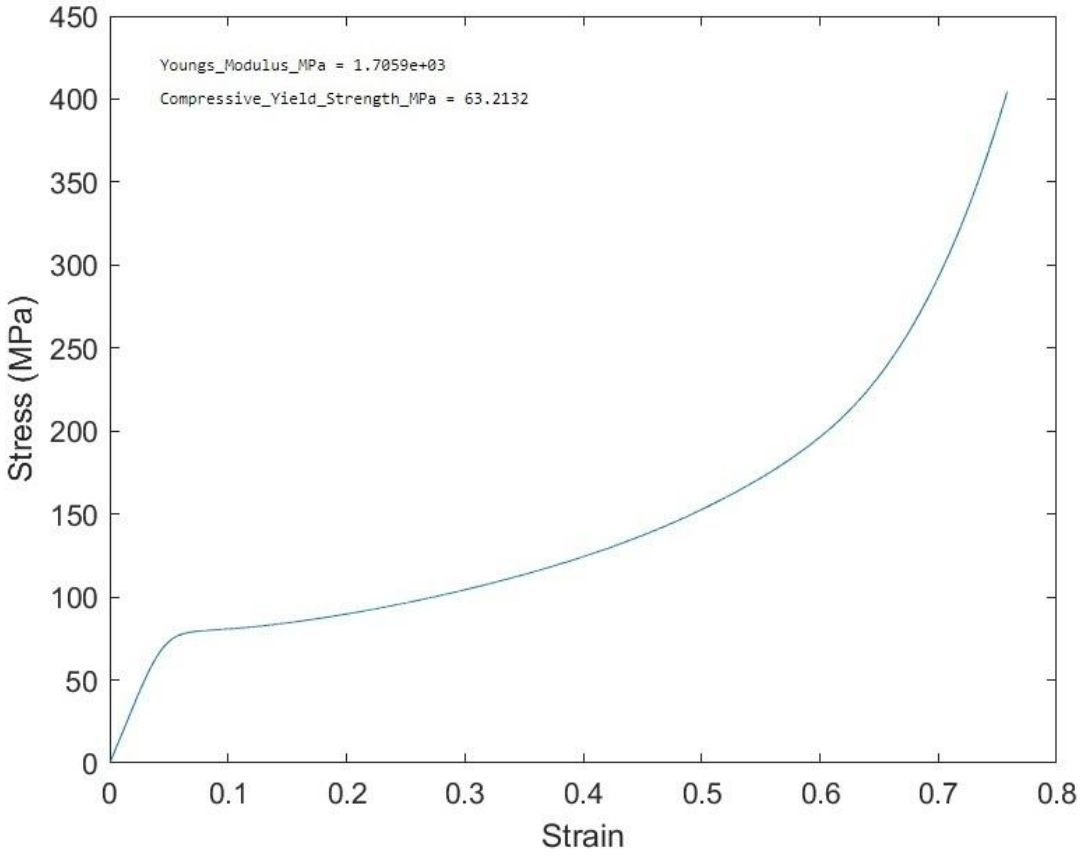
All Force-Displacement Results



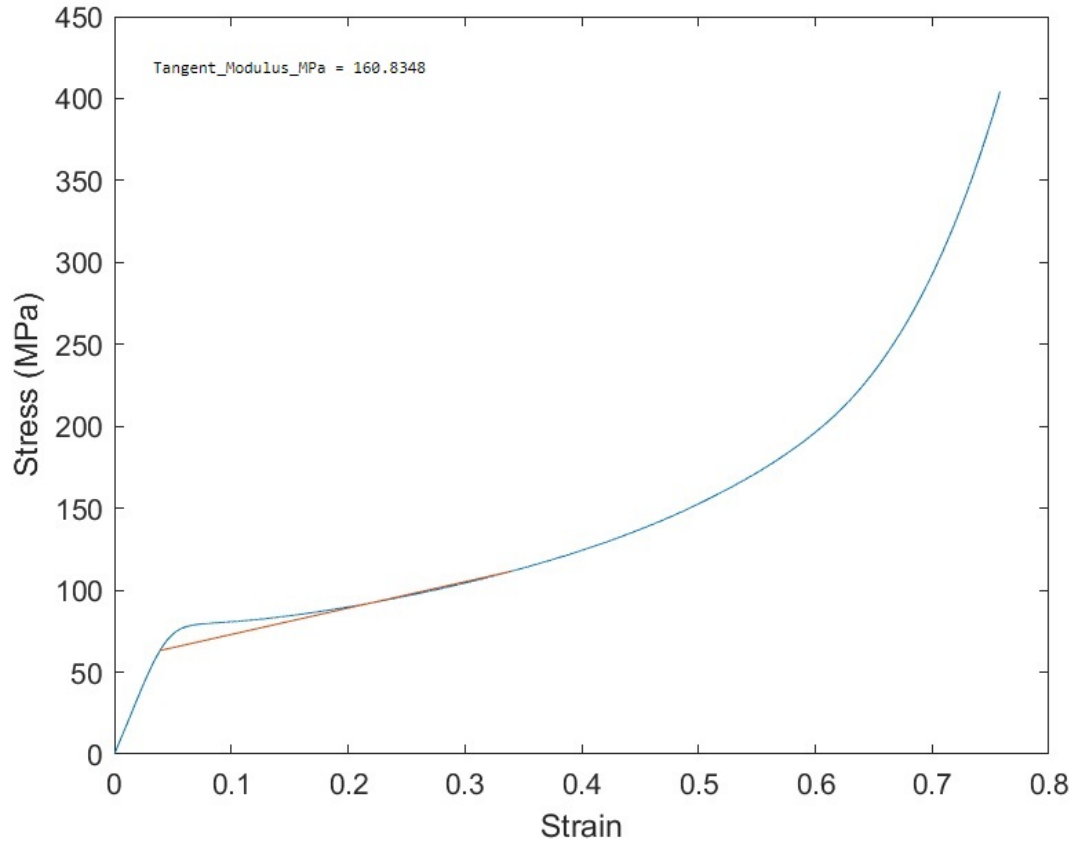
0.2% Offset Proof Stress Calculation



Toe Compensated Mean Stress-Strain Response

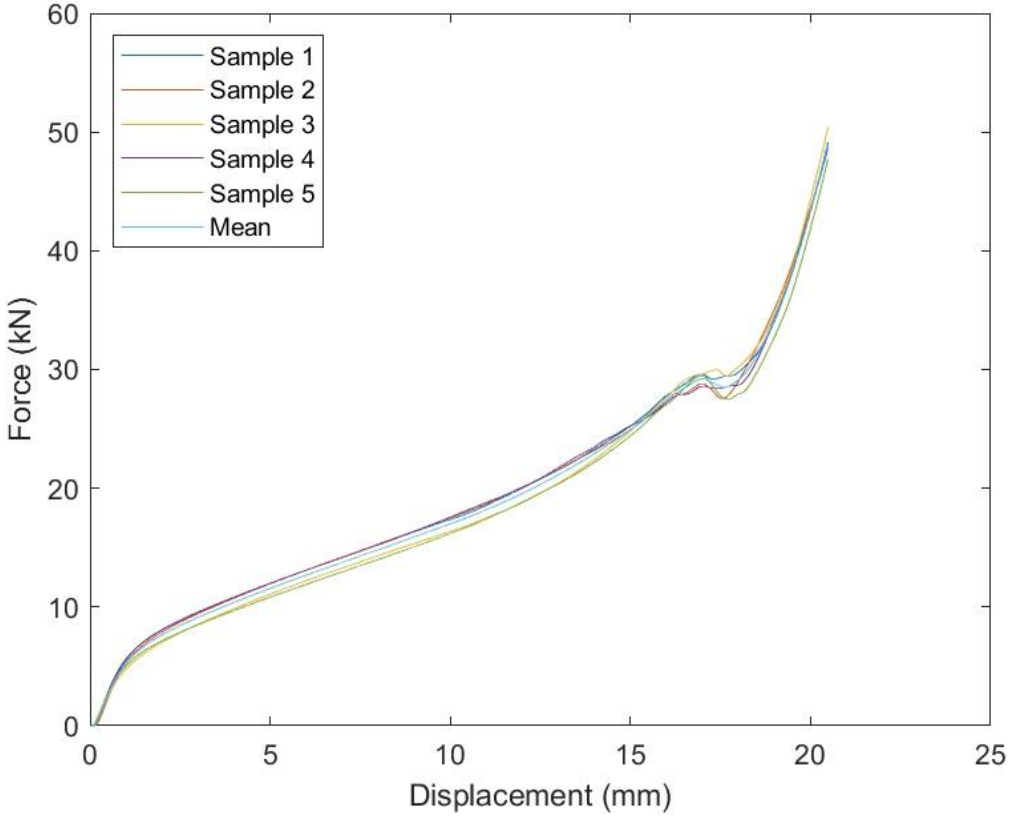


Tangent Modulus Estimation

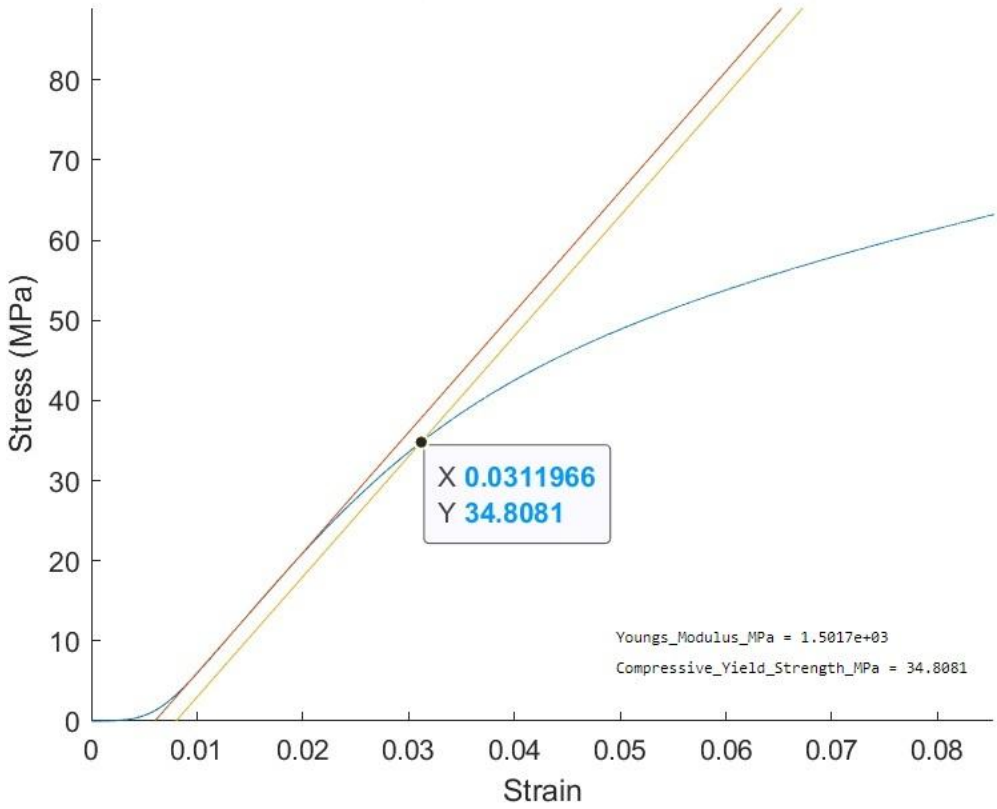


Fiberthree F3 PA-CF Pro - Cylindrical

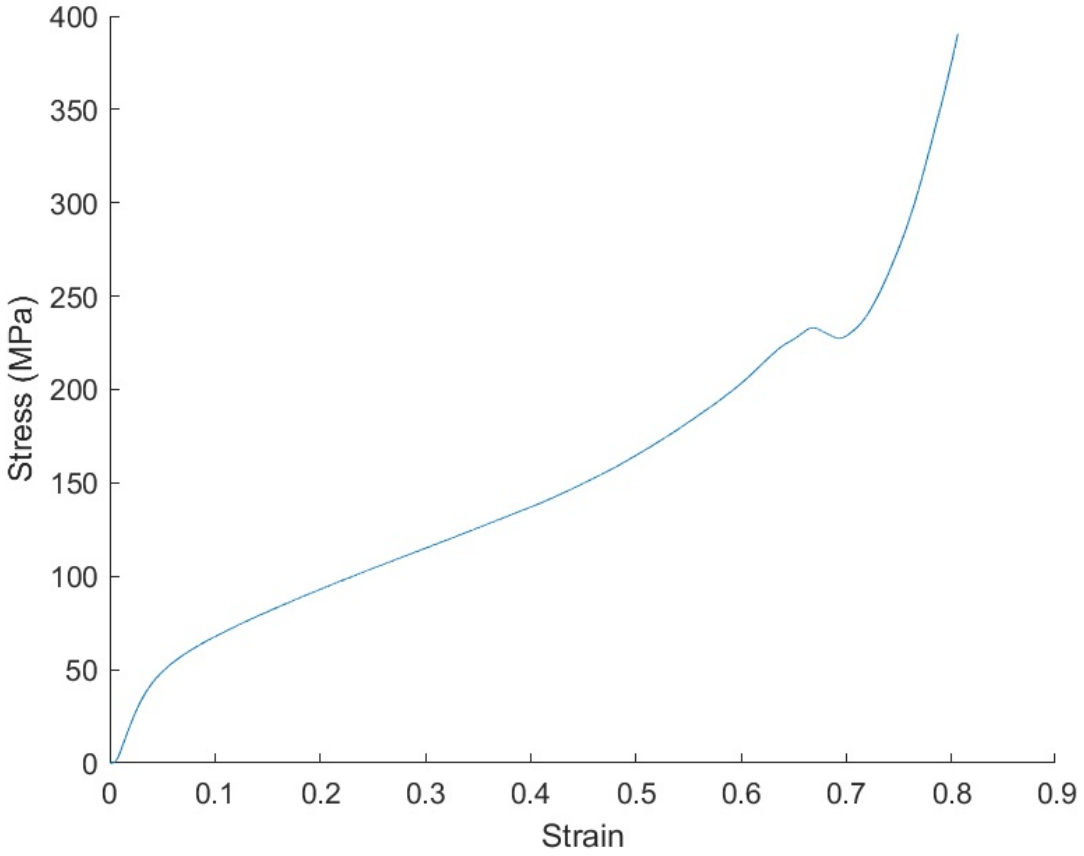
All Force-Displacement Results



0.2% Offset Proof Stress Calculation

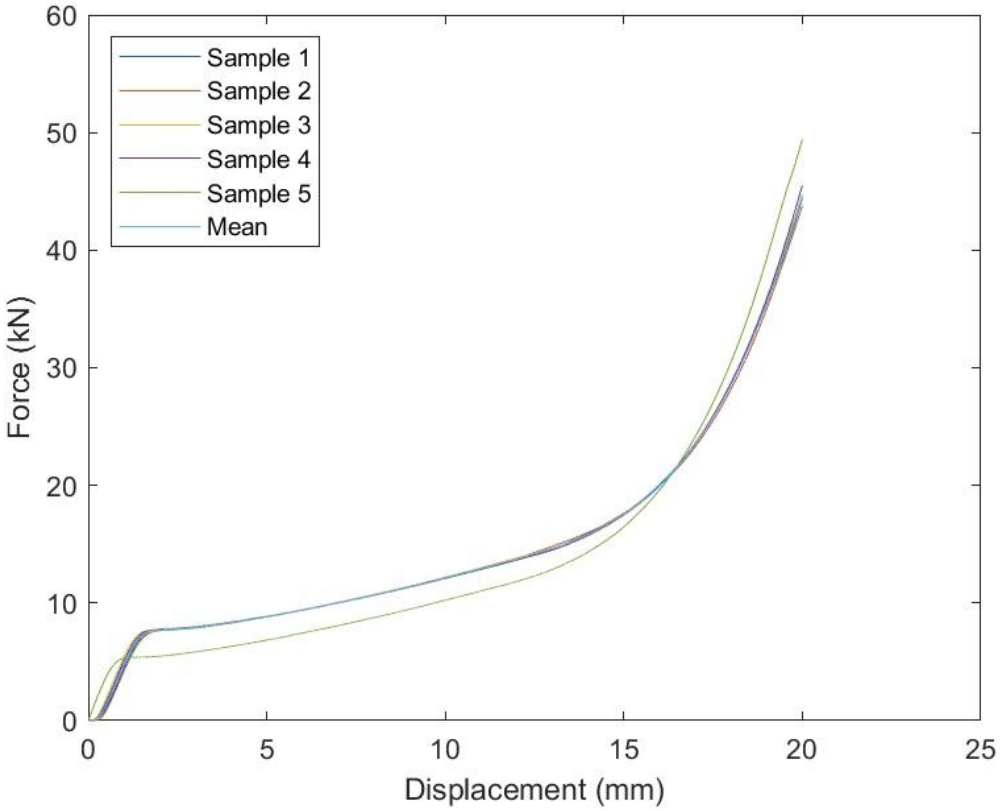


Mean Stress-Strain Response

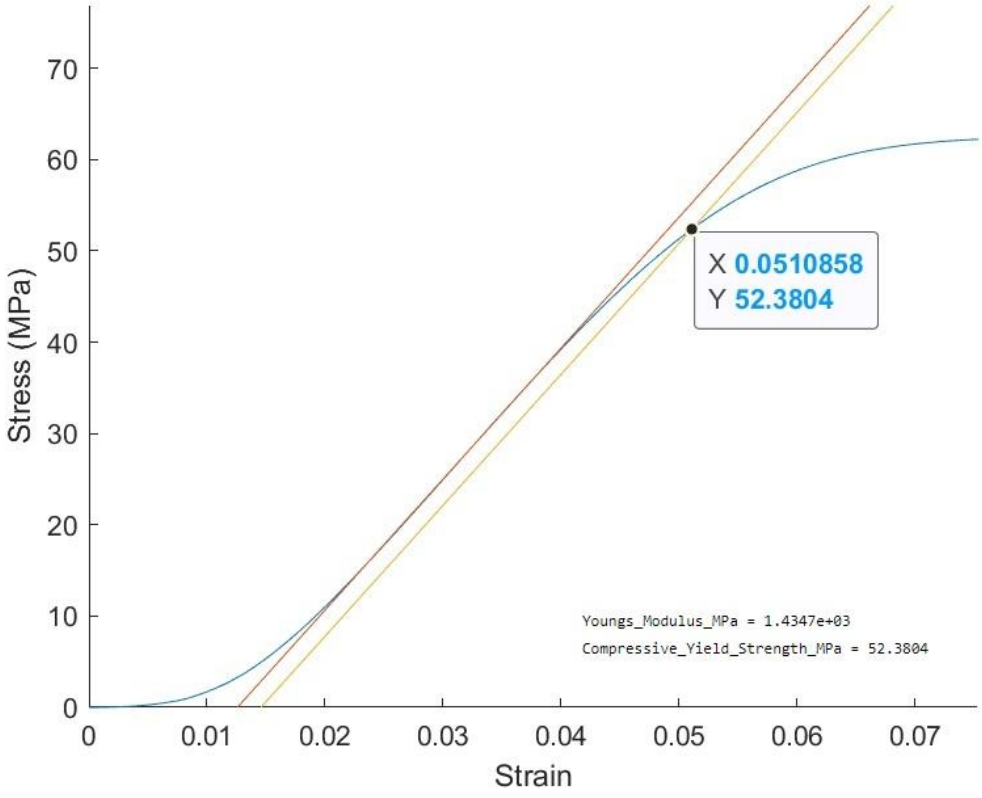


Esun ePA - Cylindrical

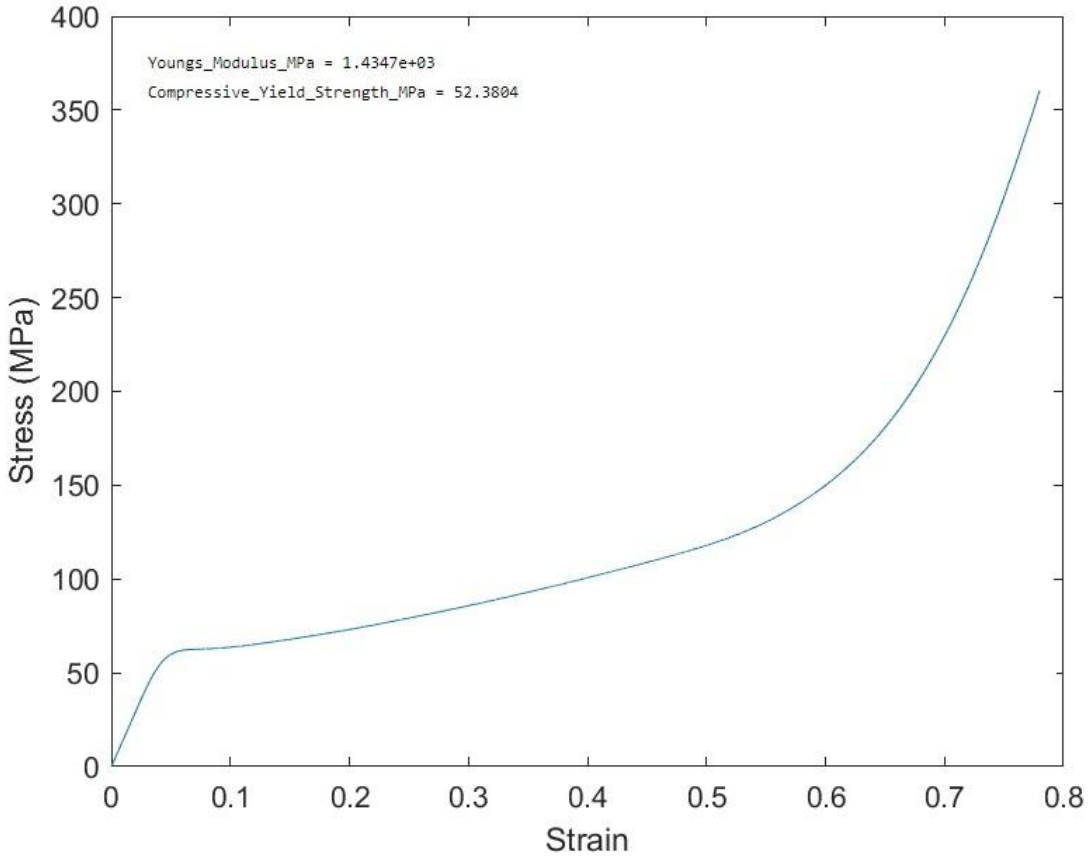
All Force-Displacement Results



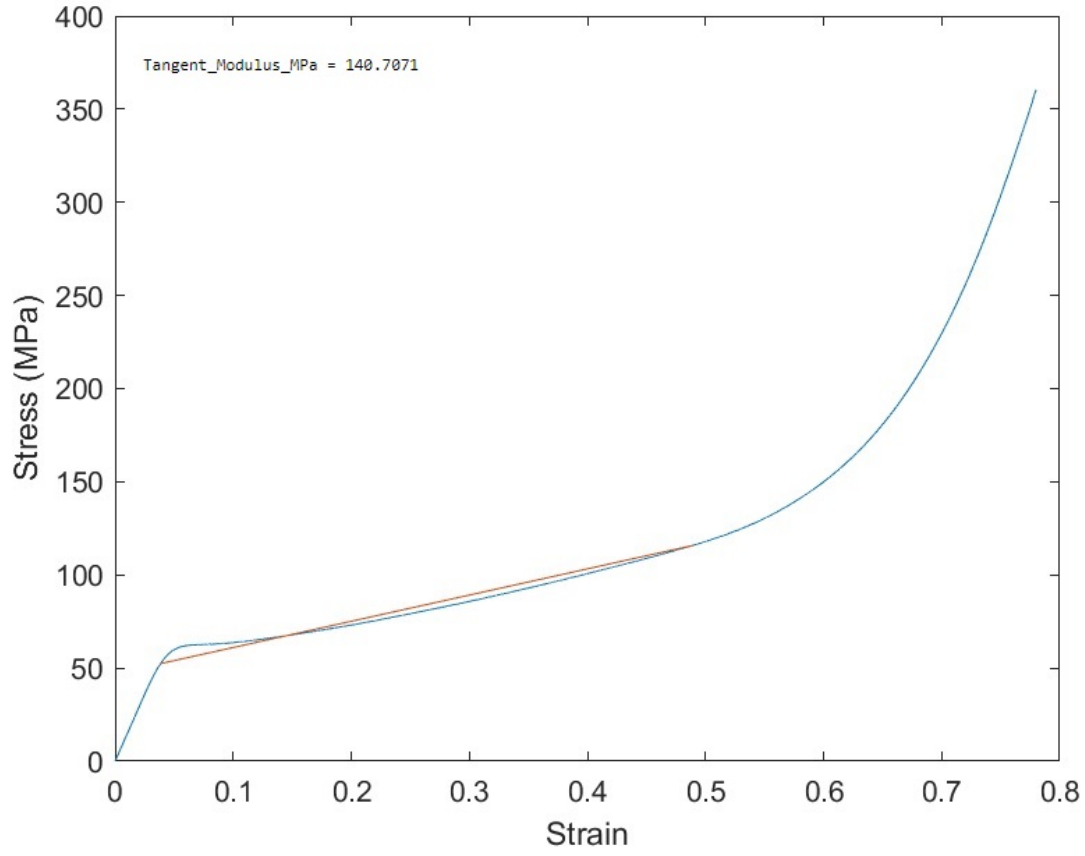
0.2% Offset Proof Stress Calculation



Toe Compensated Mean Stress-Strain Response

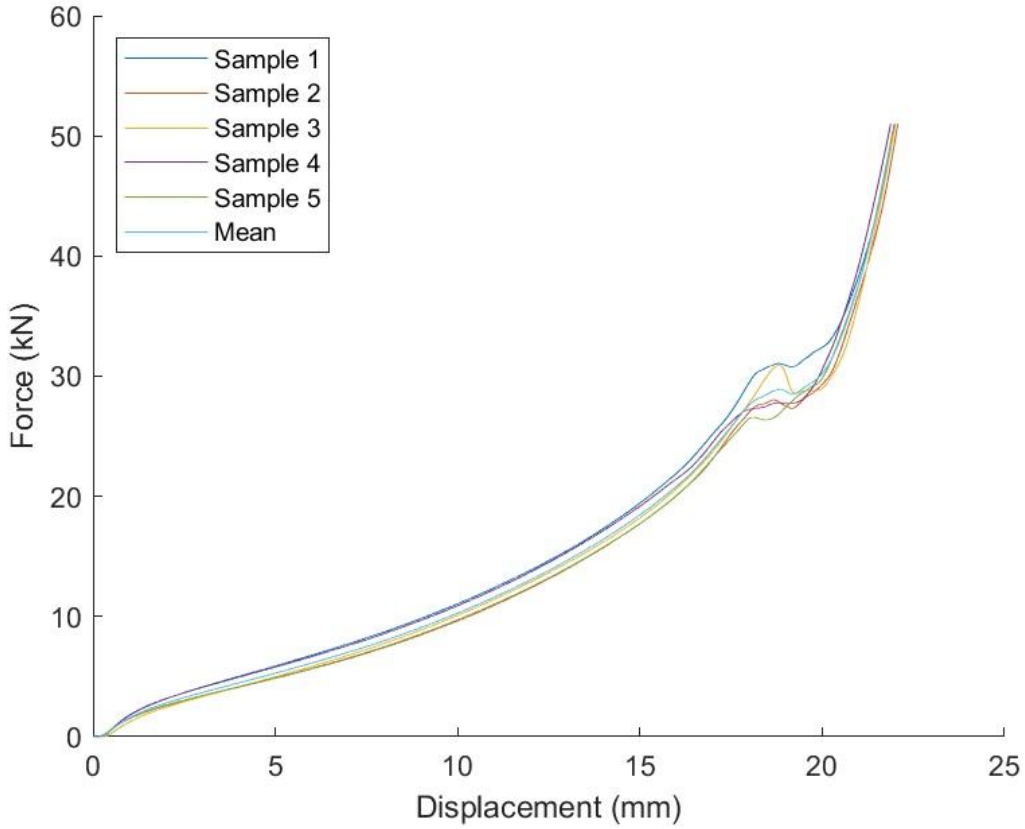


Tangent Modulus Estimation

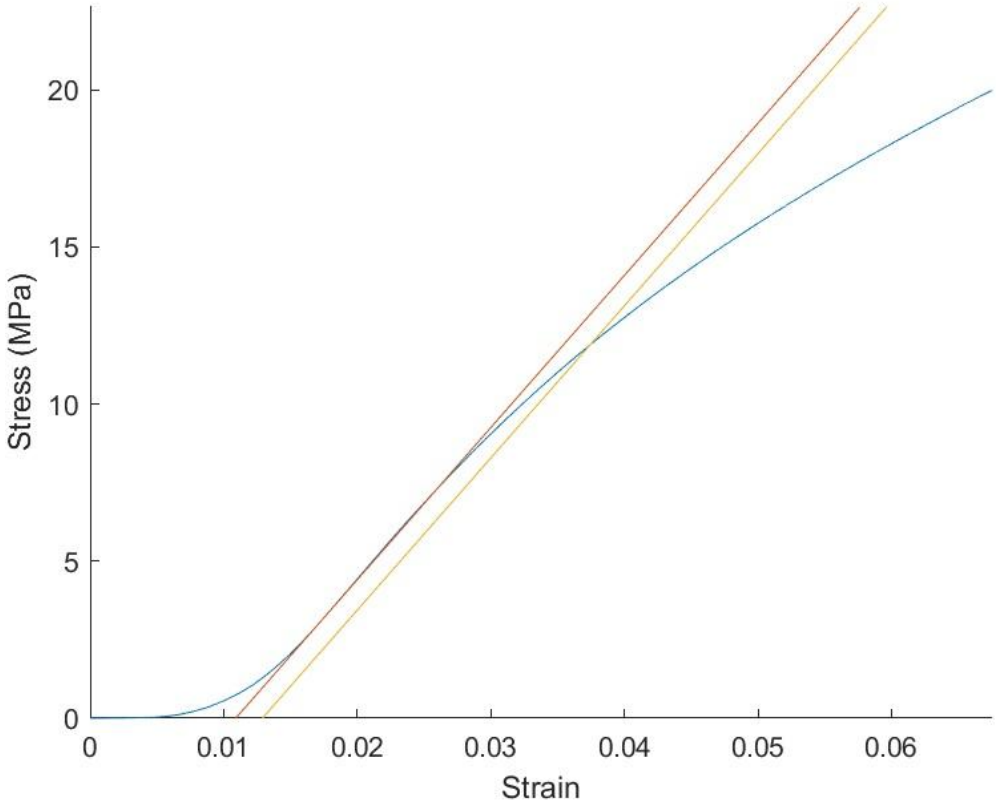


Esun ePA-CF - Cylindrical

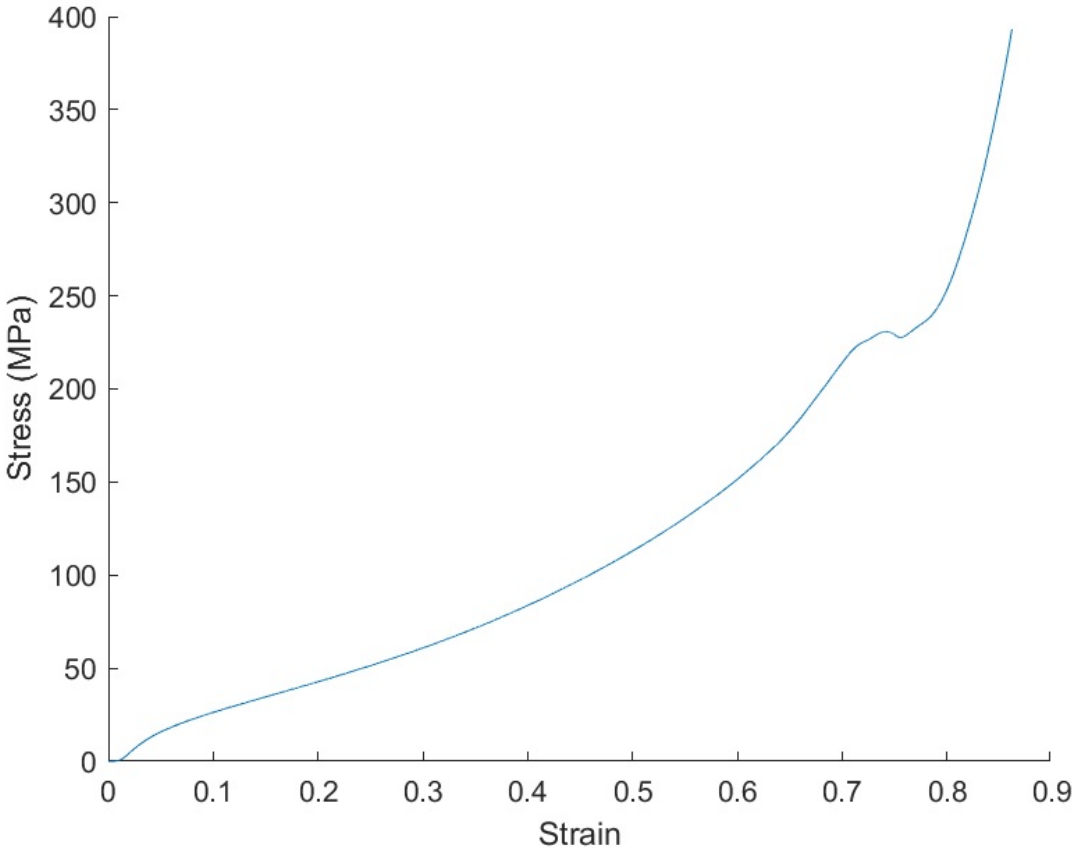
All Force-Displacement Results



0.2% Offset Proof Stress Calculation

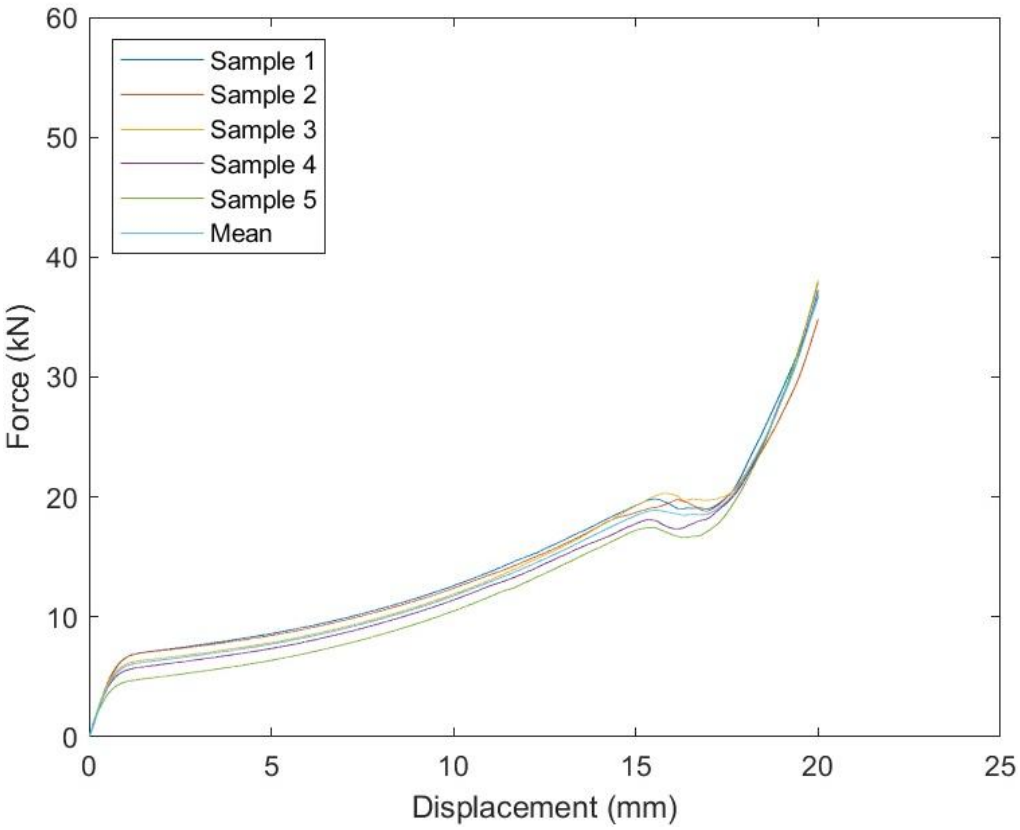


Mean Stress-Strain Response

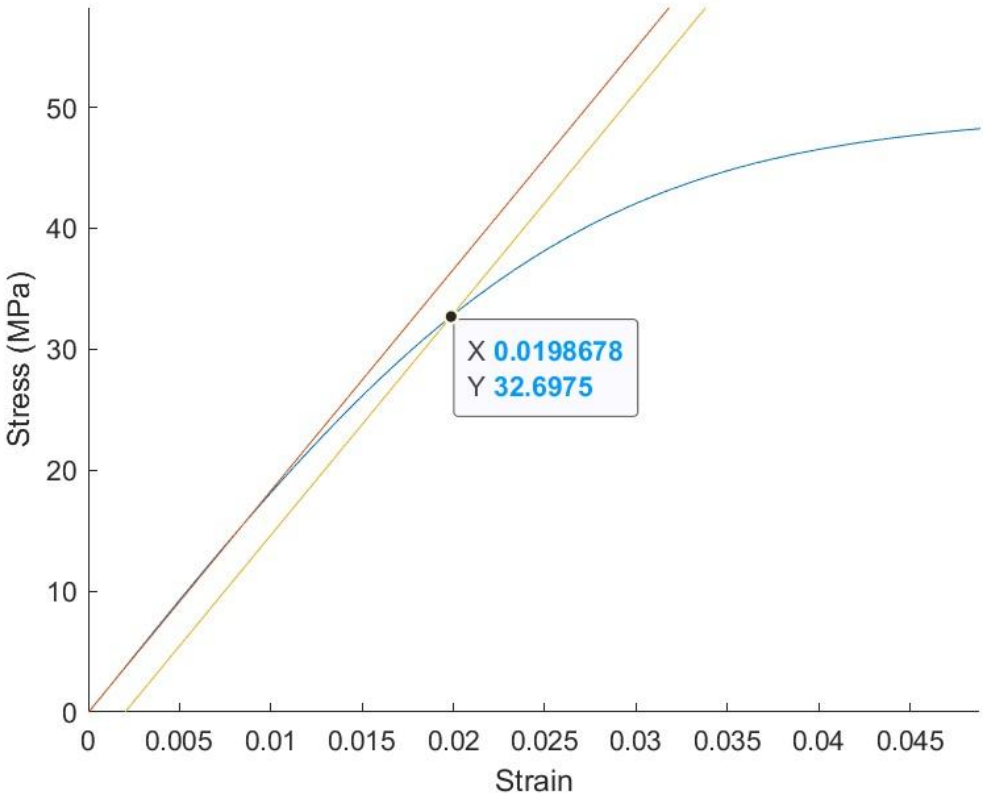


ColorFabb XT-CF20 - Cylindrical

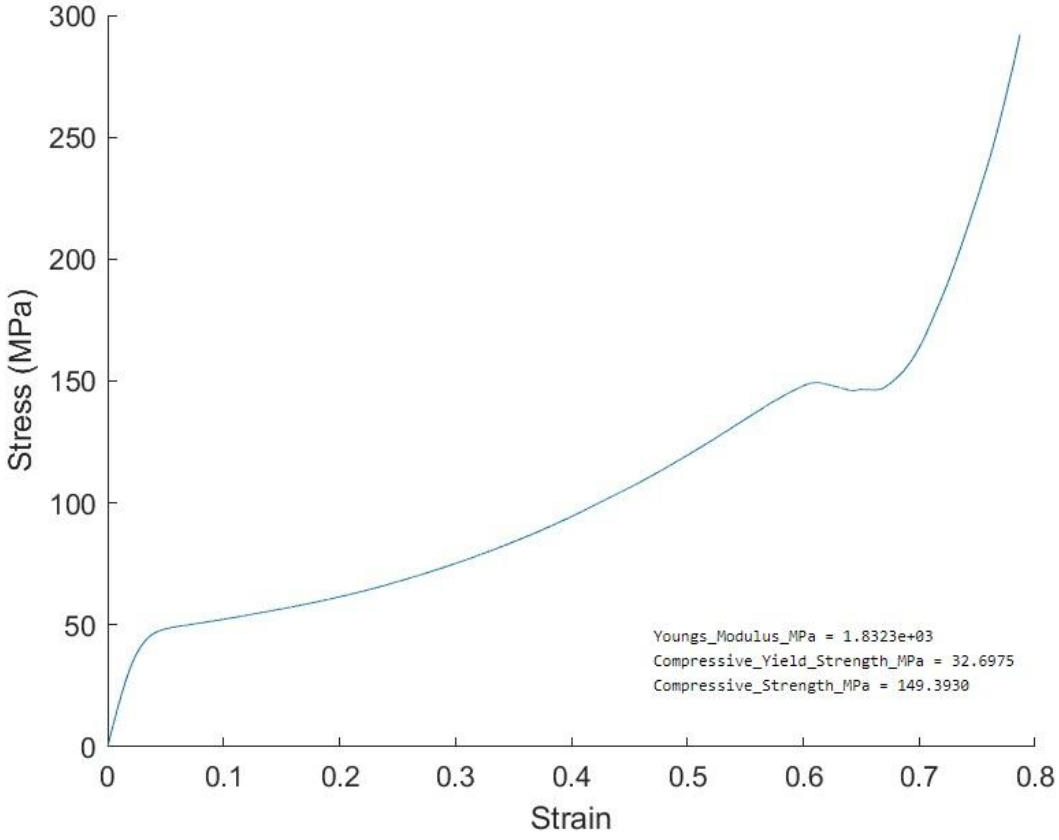
All Force-Displacement Results



0.2% Offset Proof Stress Calculation

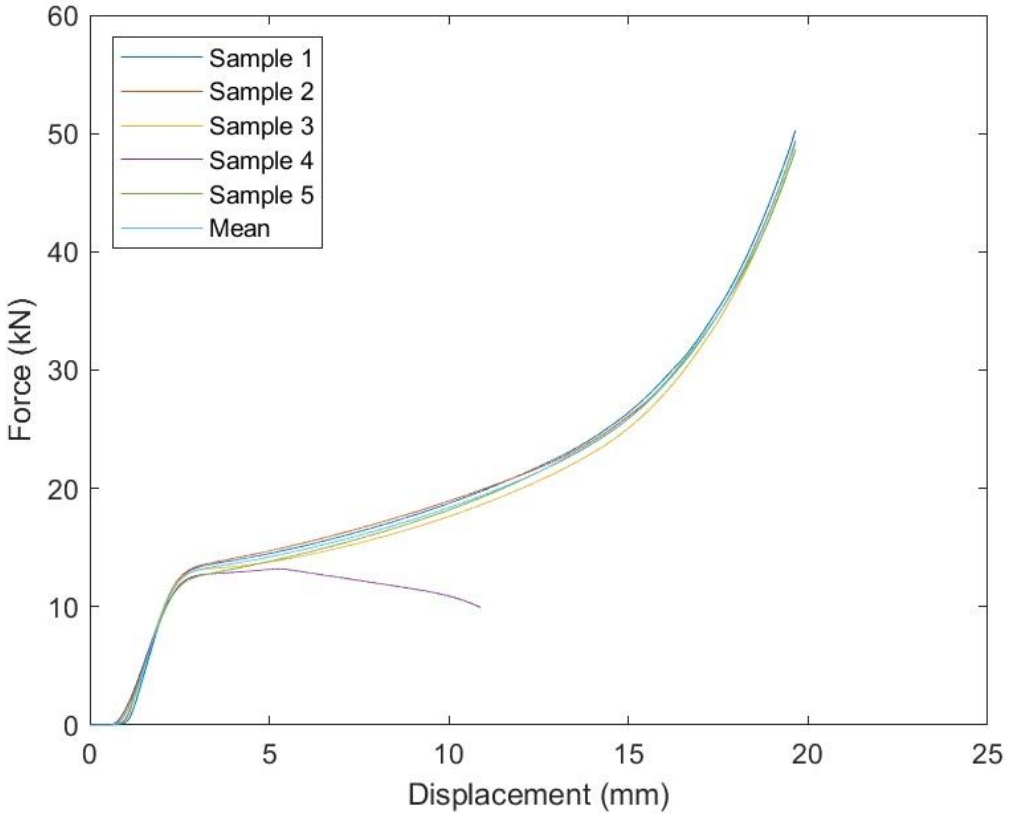


Toe Compensated Mean Stress-Strain Response

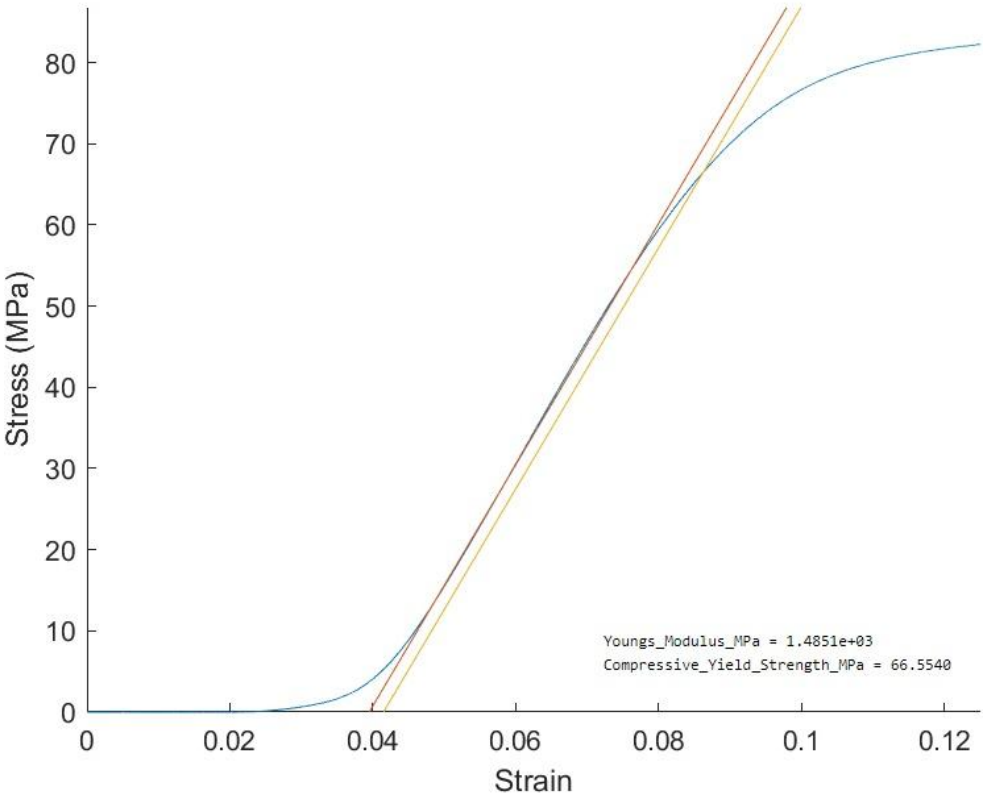


Polymaker CoPA - Cuboid 0 Degrees

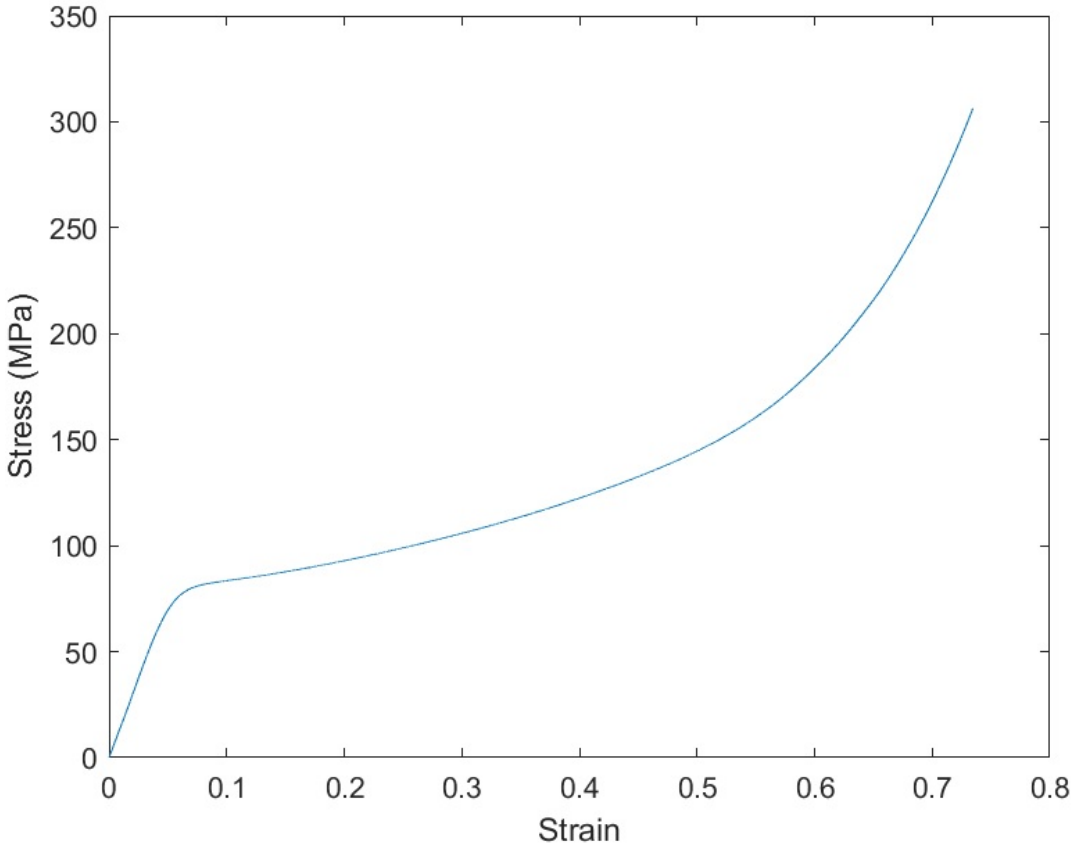
All Force-Displacement Results



0.2% Offset Proof Stress Calculation

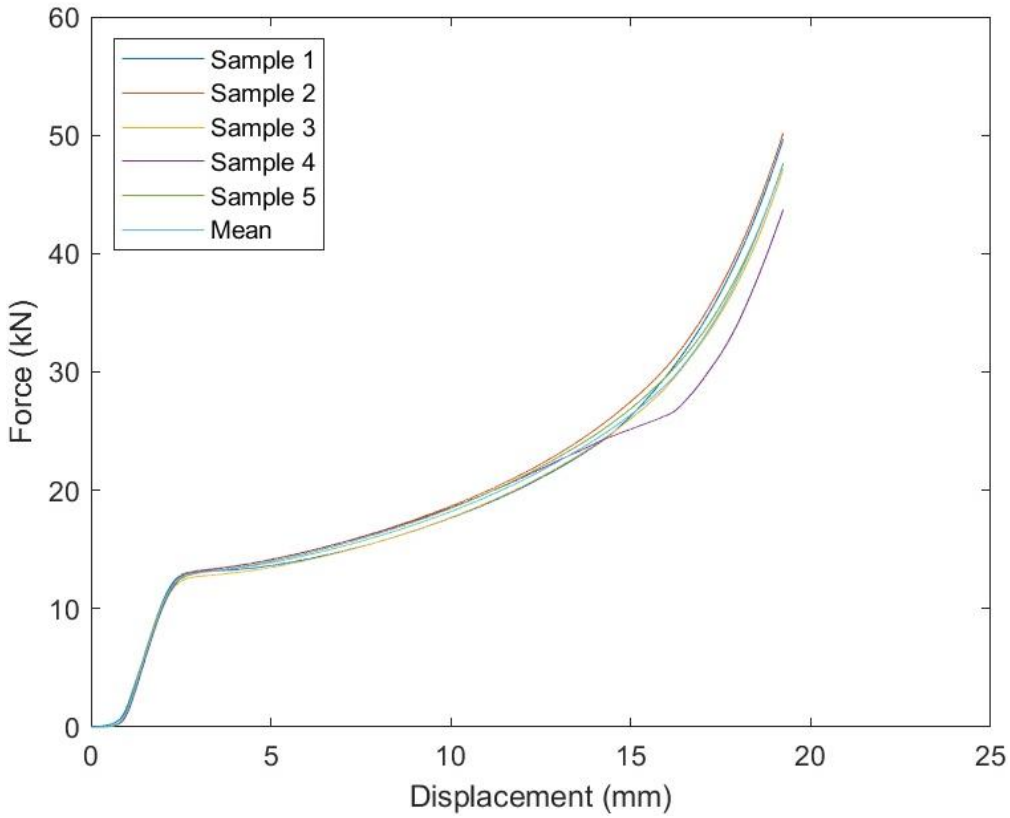


Toe Compensated Mean Stress-Strain Response

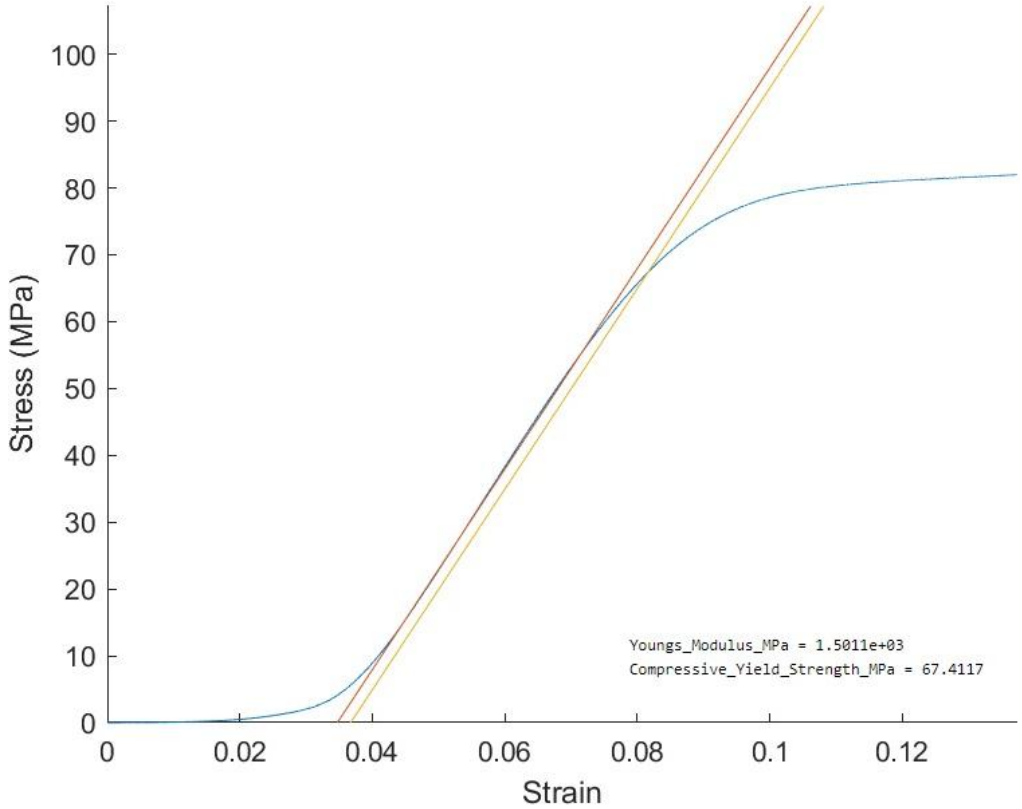


Polymaker CoPA - Cuboid 45 Degrees

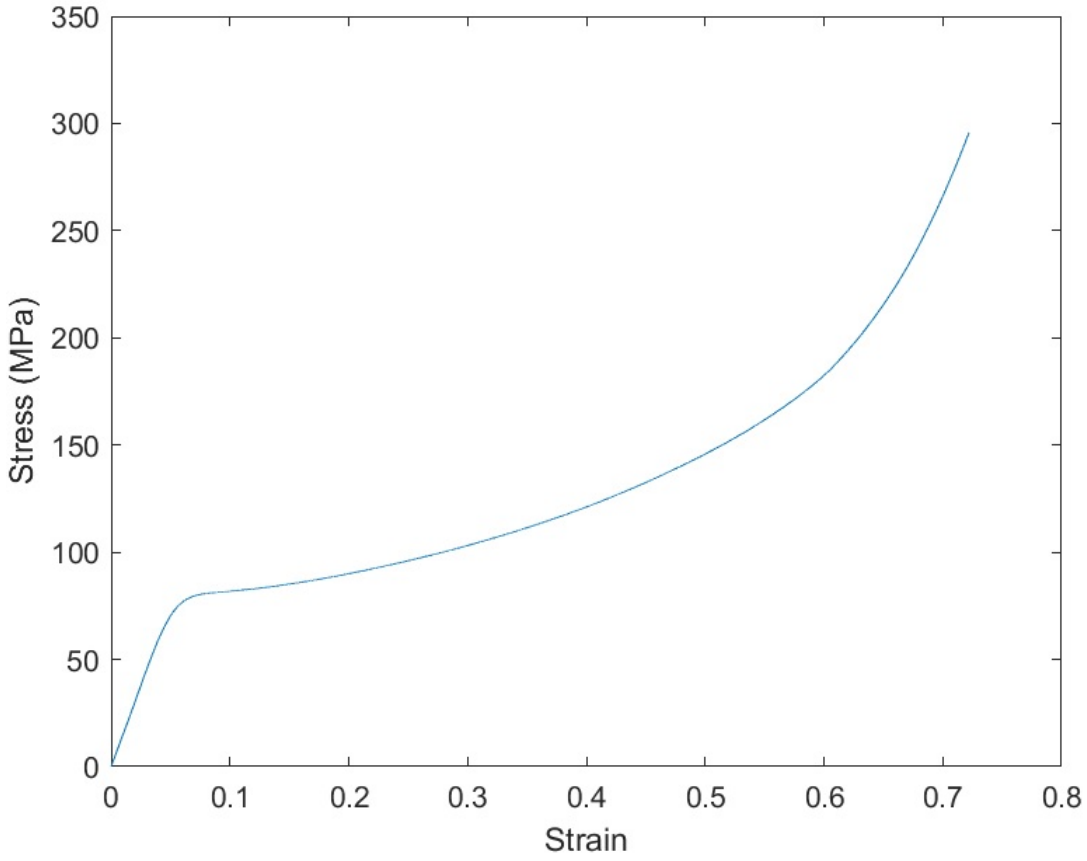
All Force-Displacement Results



0.2% Offset Proof Stress Calculation

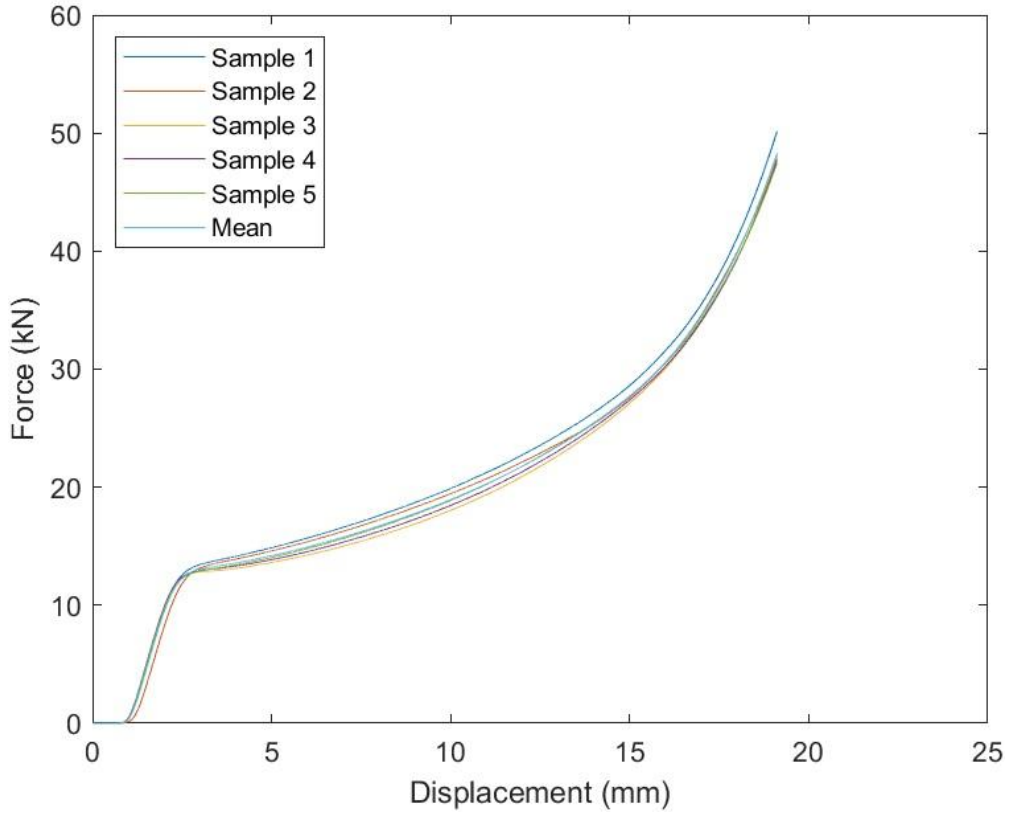


Toe Compensated Mean Stress-Strain Response

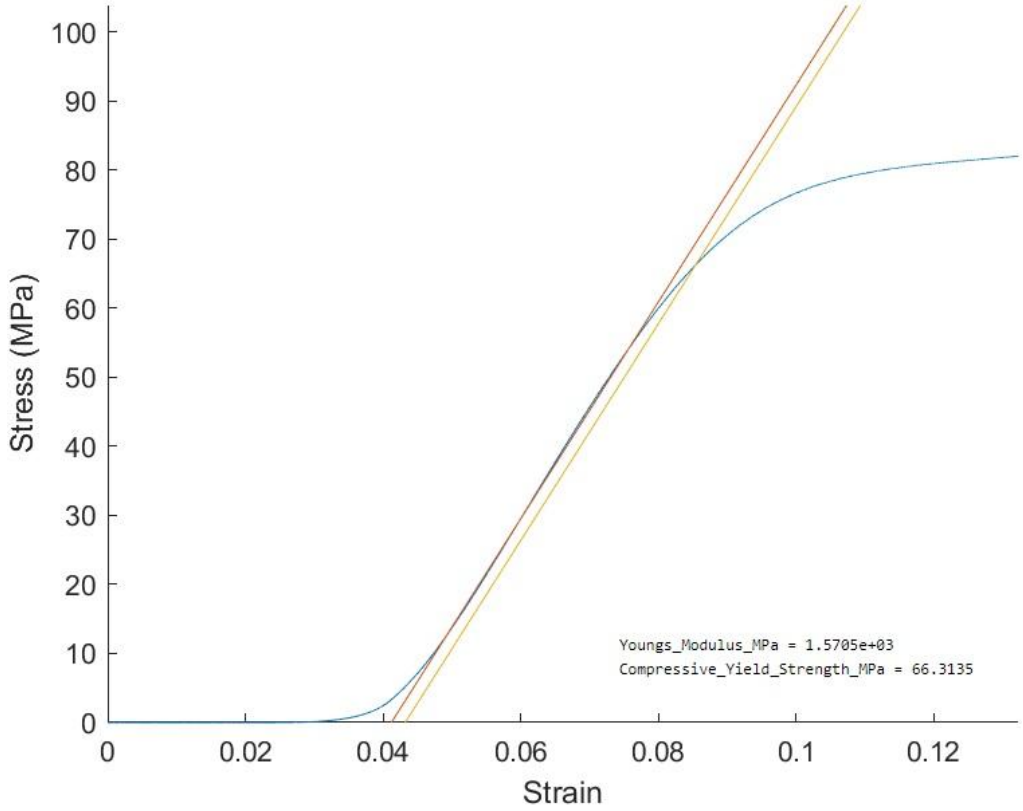


Polymaker CoPA - Cuboid 90 Degrees

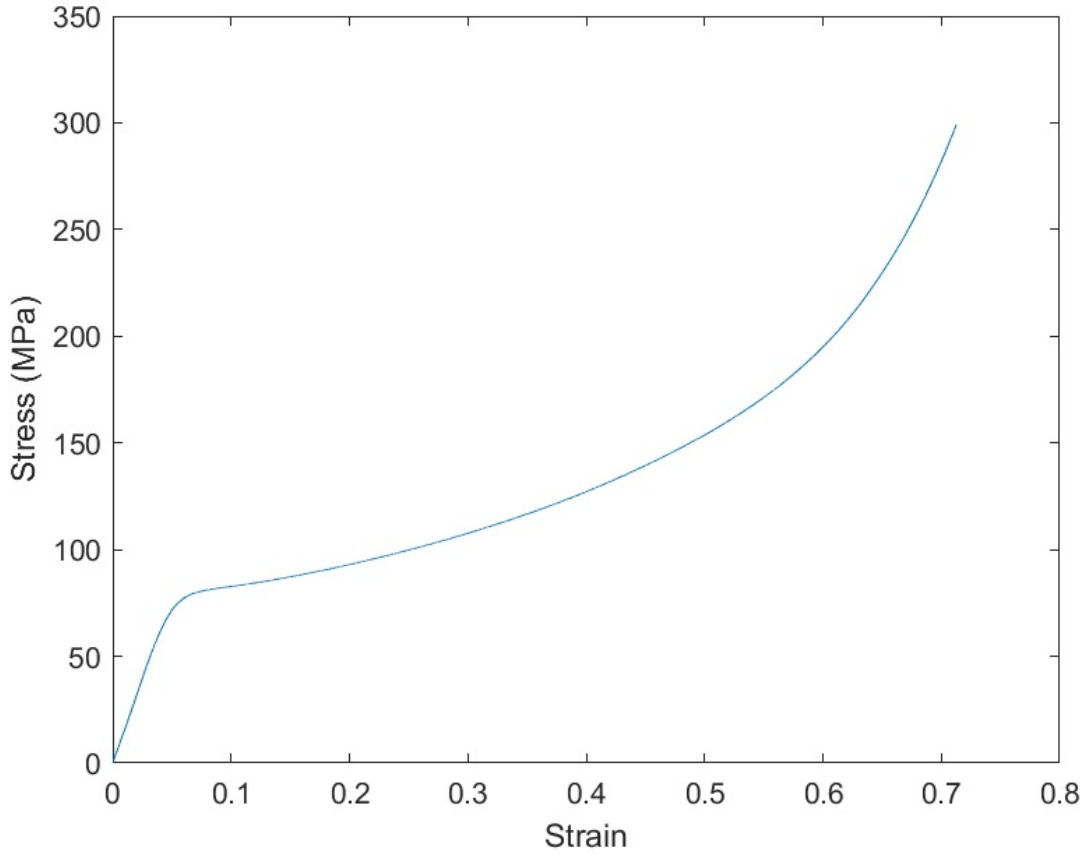
All Force-Displacement Results



0.2% Offset Proof Stress Calculation



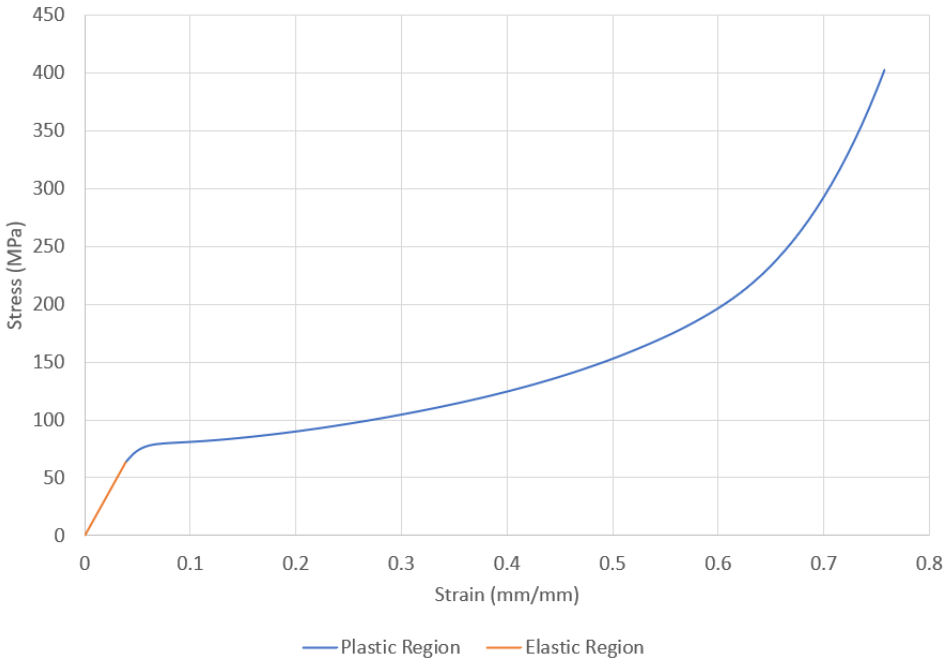
Toe Compensated Mean Stress-Strain Response



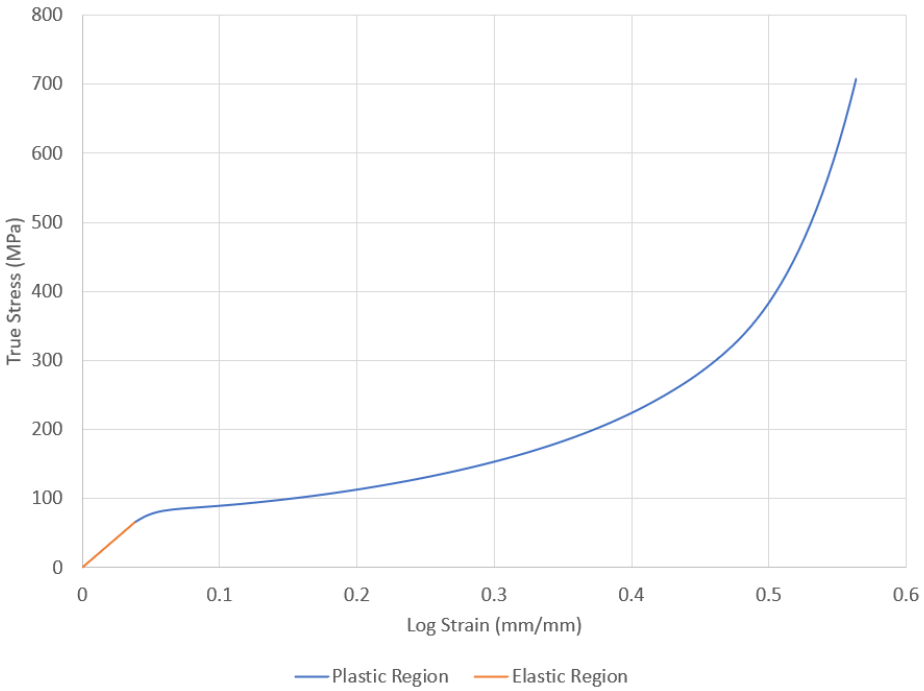
Appendix G Engineering and True Stress-Strain Compressive Material Data Used in the ANSYS Non-linear Material Model

Polymaker CoPA – 90° Circular

Engineering Stress-Strain Response – 100 Datapoints

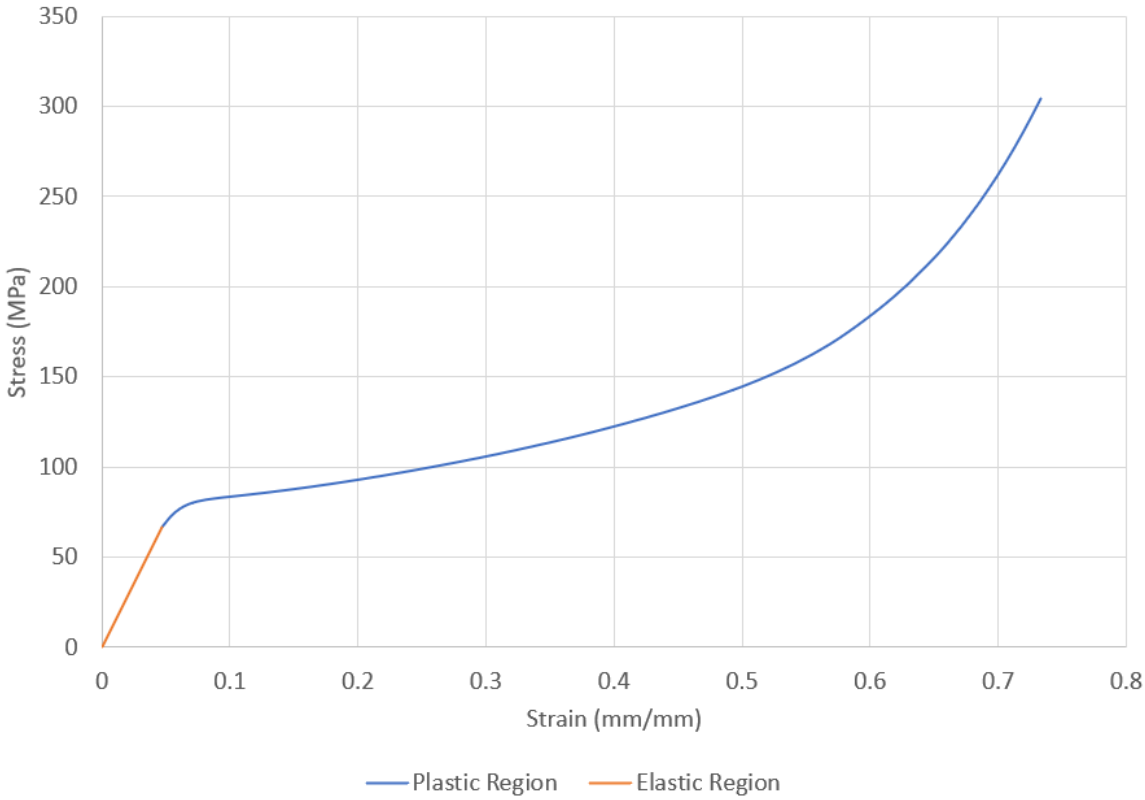


True Stress-Strain Response – 100 Datapoints

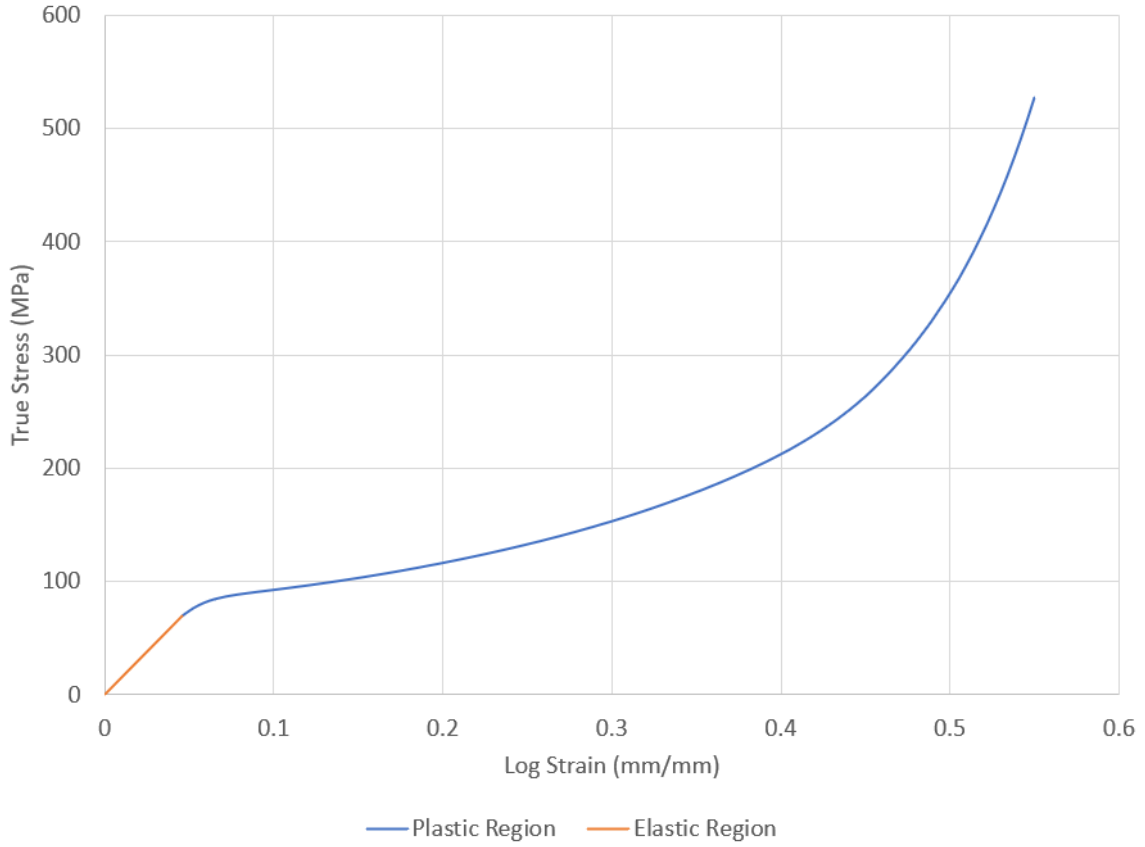


Polymaker CoPA – 0° Square

Engineering Stress-Strain Response – 100 Datapoints

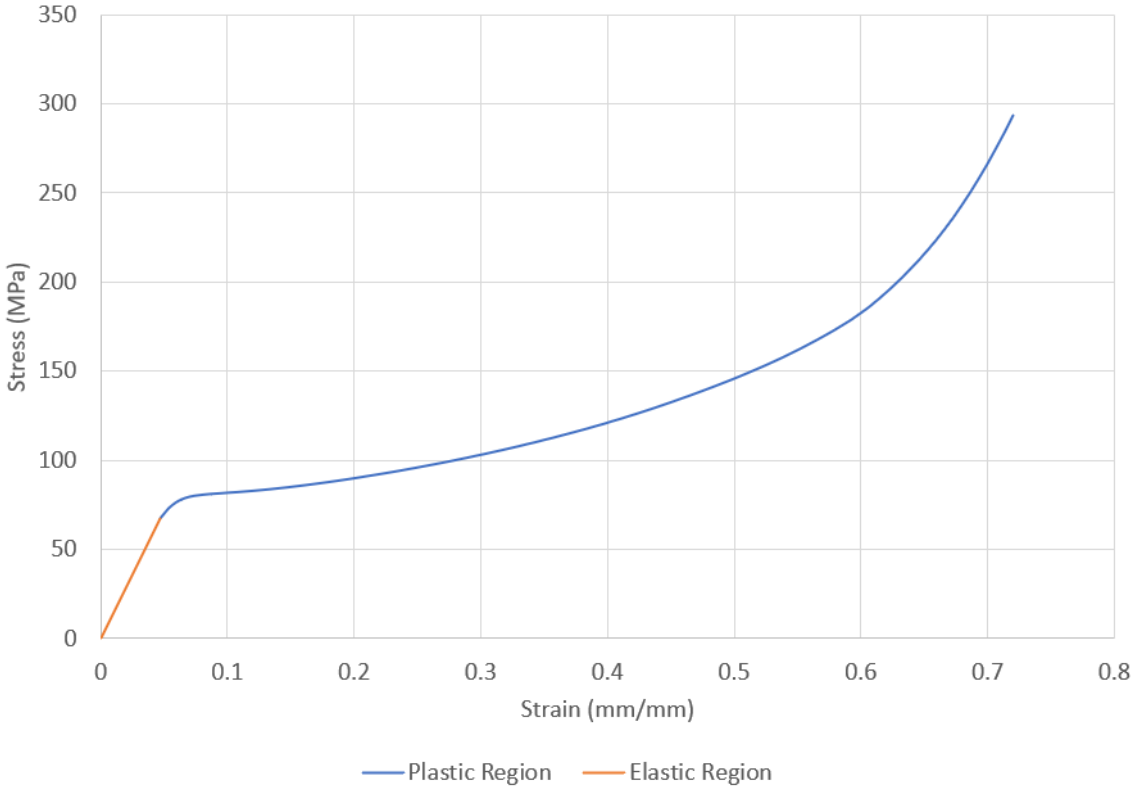


True Stress-Strain Response – 100 Datapoints

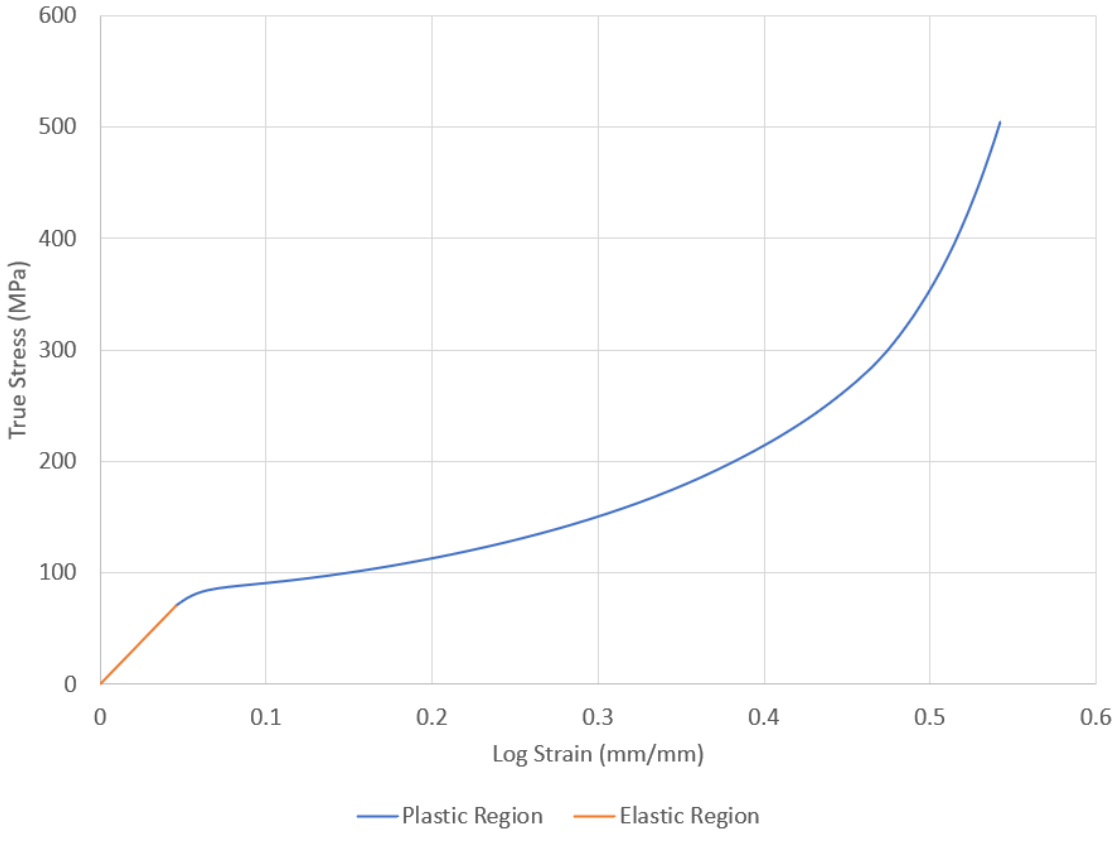


Polymaker CoPA – 45° Square

Engineering Stress-Strain Response – 100 Datapoints

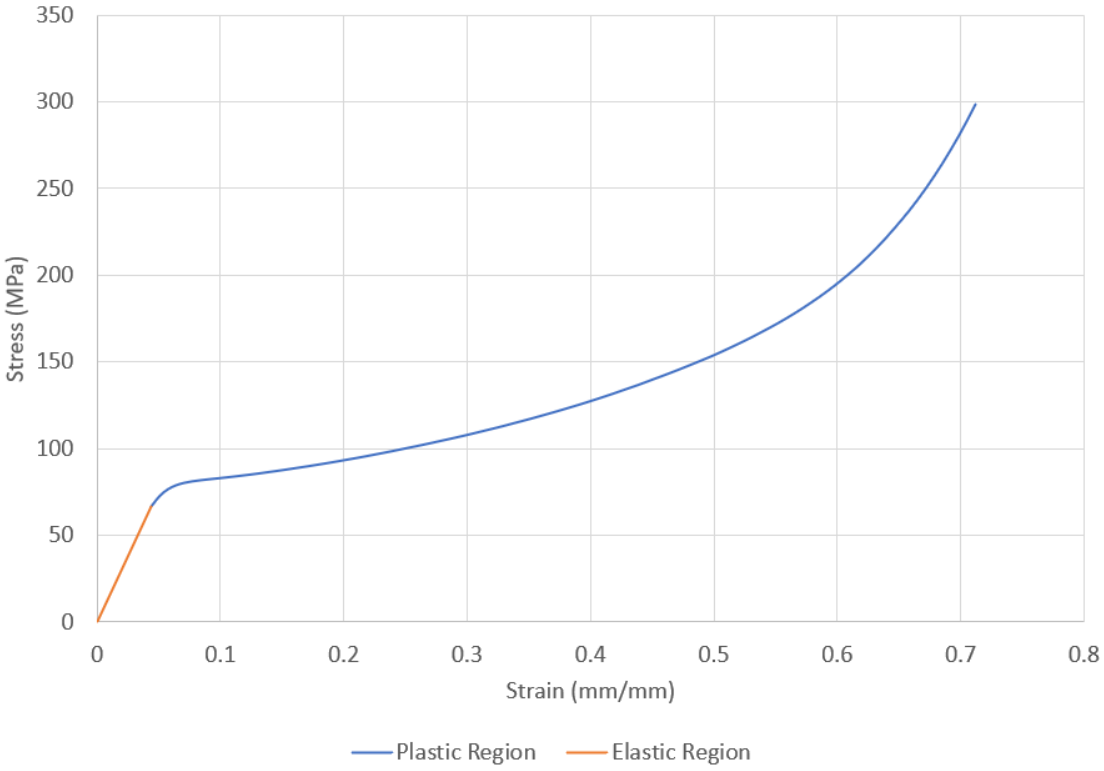


True Stress-Strain Response – 100 Datapoints

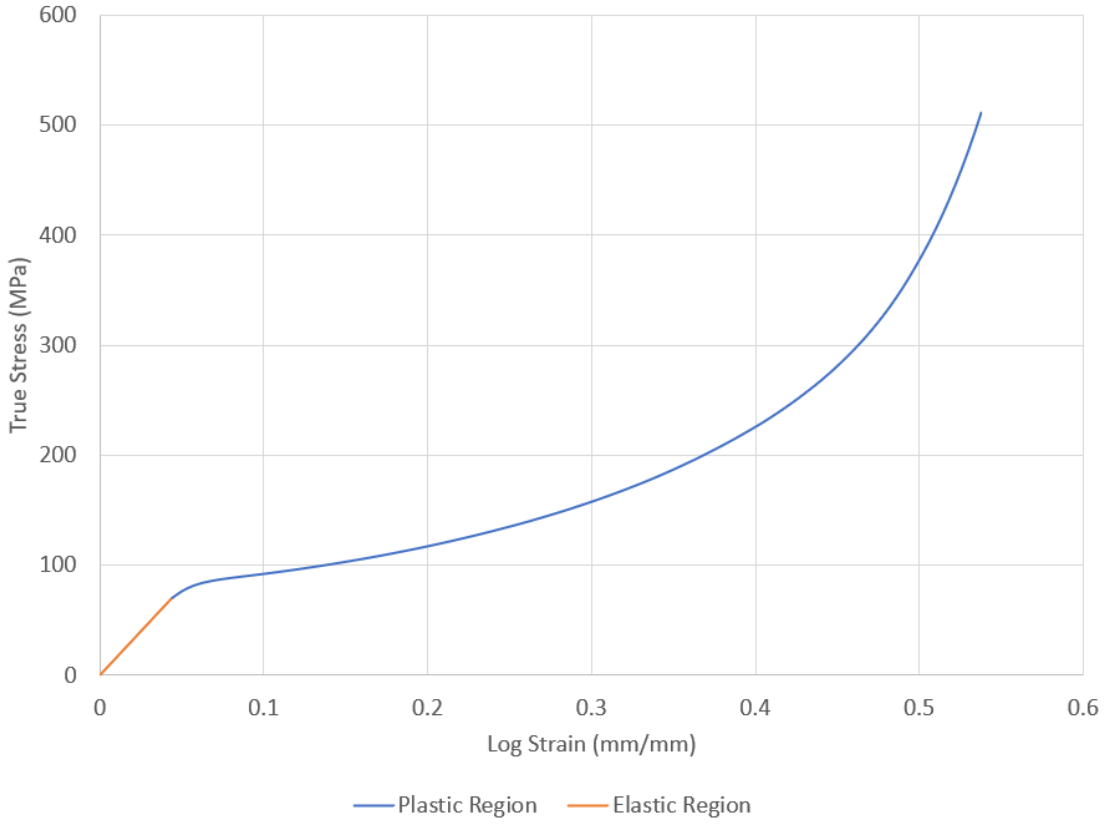


Polymaker CoPA – 90° Square

Engineering Stress-Strain Response – 100 Datapoints



True Stress-Strain Response – 100 Datapoints

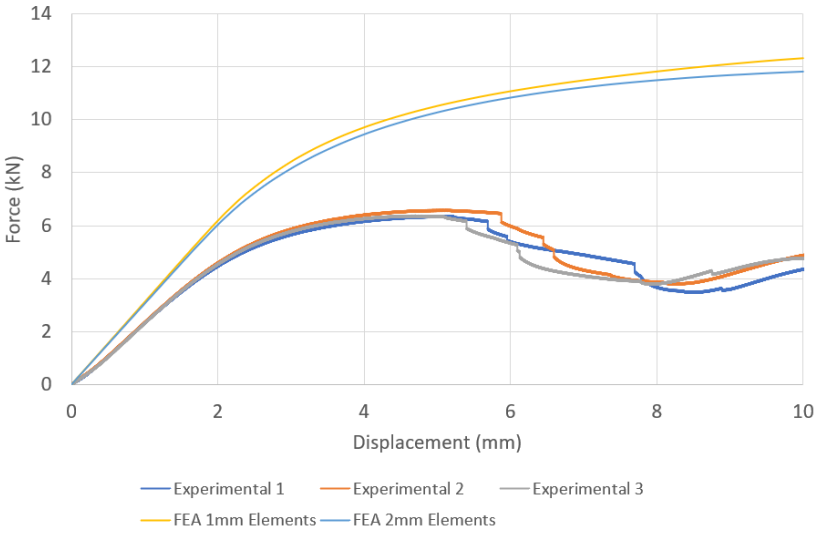


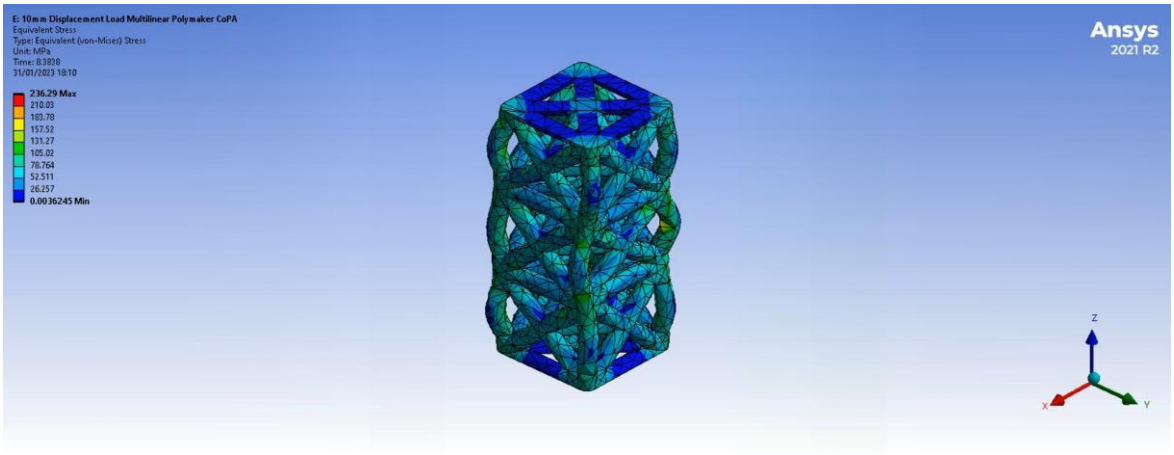
Appendix H Solid Body FEA Model Element Size Comparison

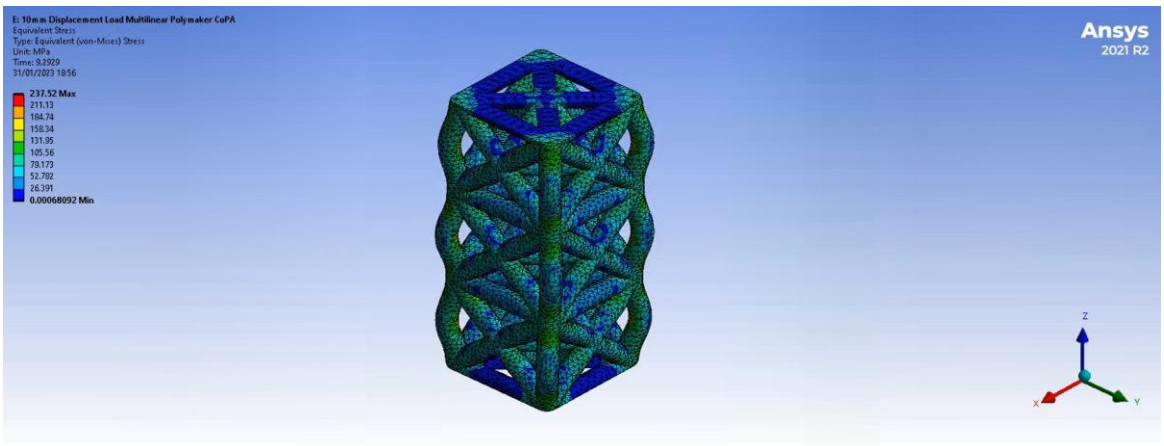
Material model: Polymaker CoPA non-linear - 90° circular test coupons

20x20x60mm lattice compression block

All strut radii set to 2mm

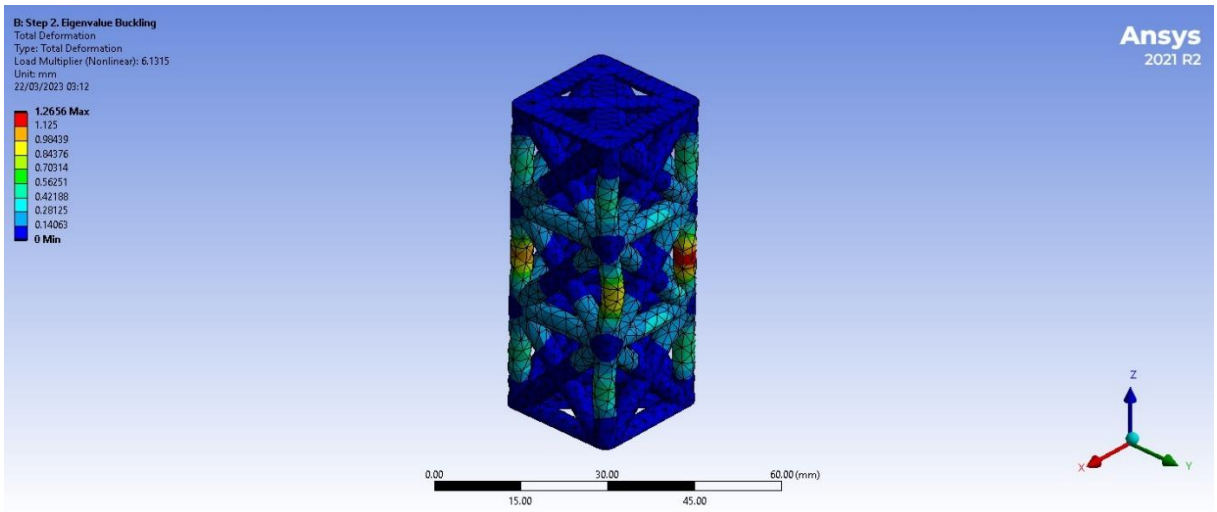




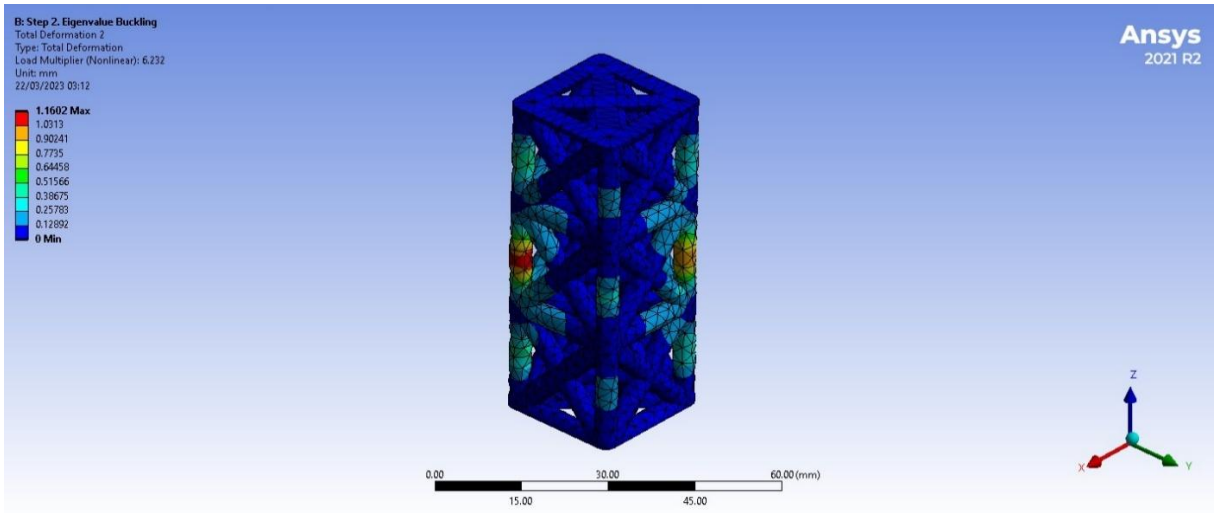


Appendix I Non-linear Eigenvalue Buckling Analysis Mode Shapes for 2mm Strut Radius Lattice Compression Block

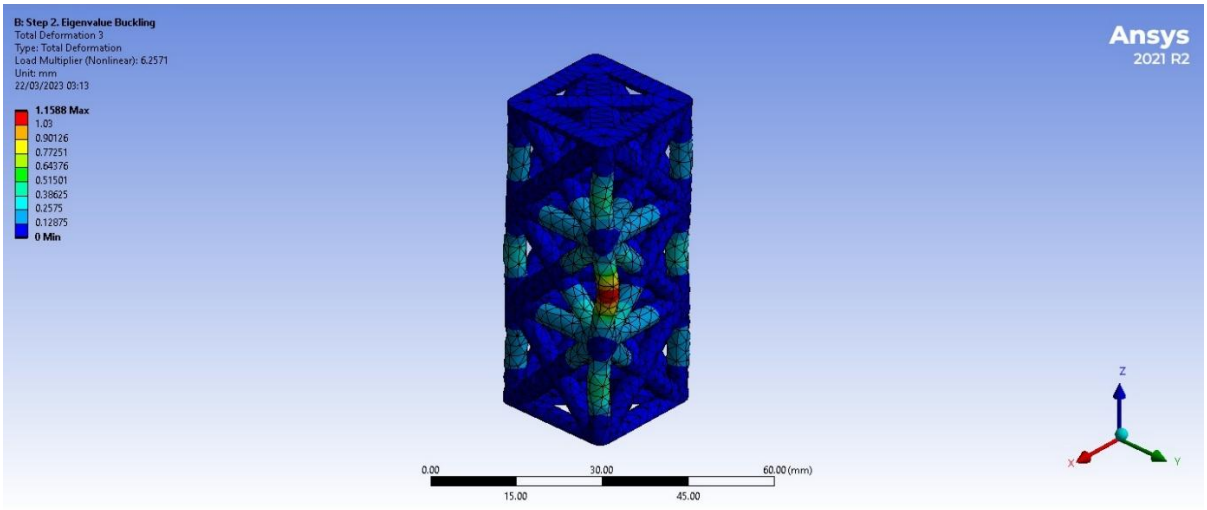
Mode 1 – Load Multiplier 6.1315



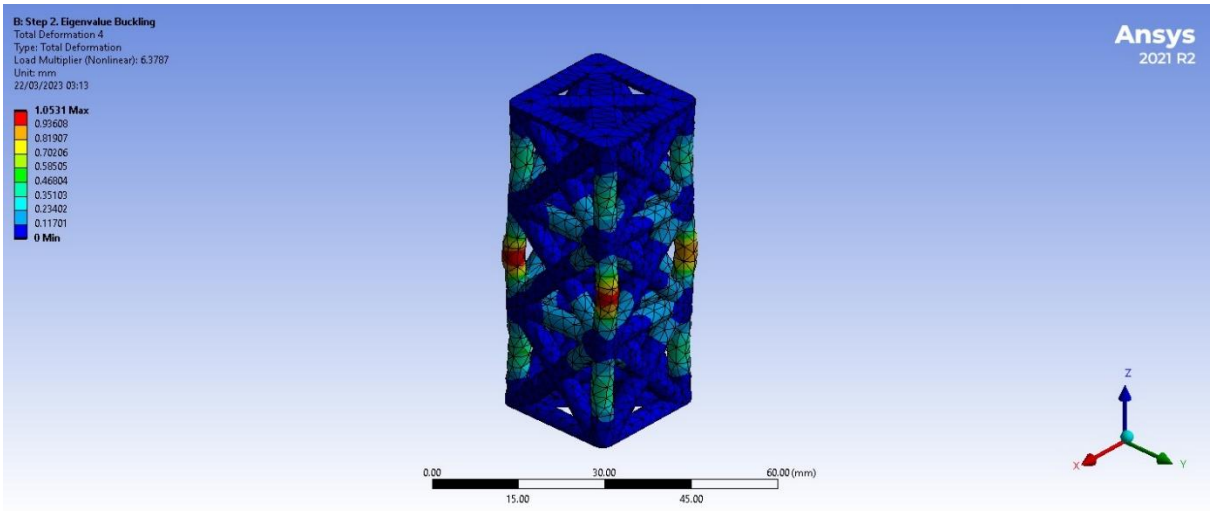
Mode 2 – Load Multiplier 6.2320



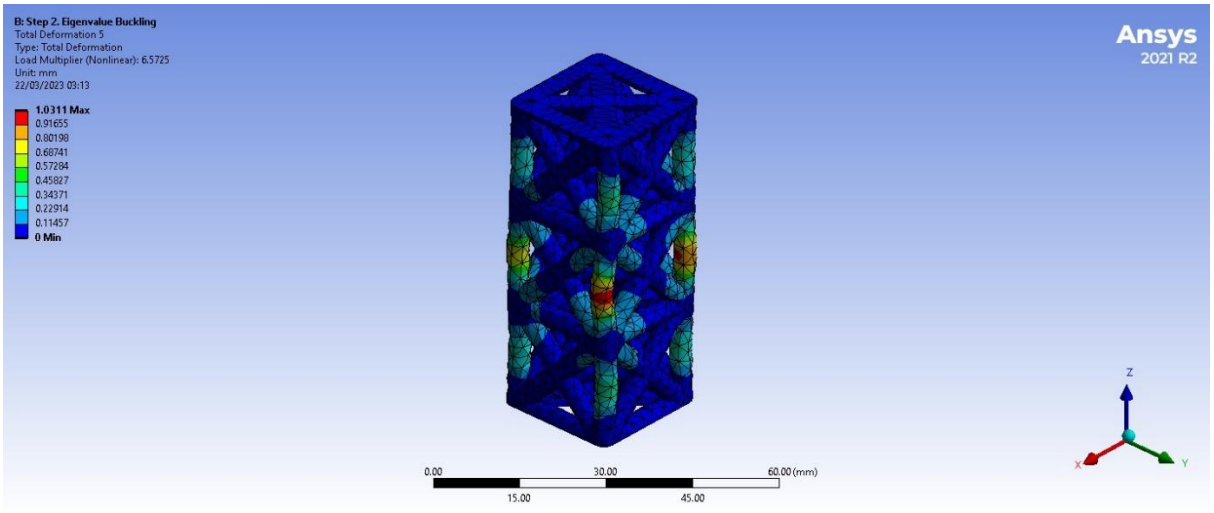
Mode 3 – Load Multiplier 6.2571



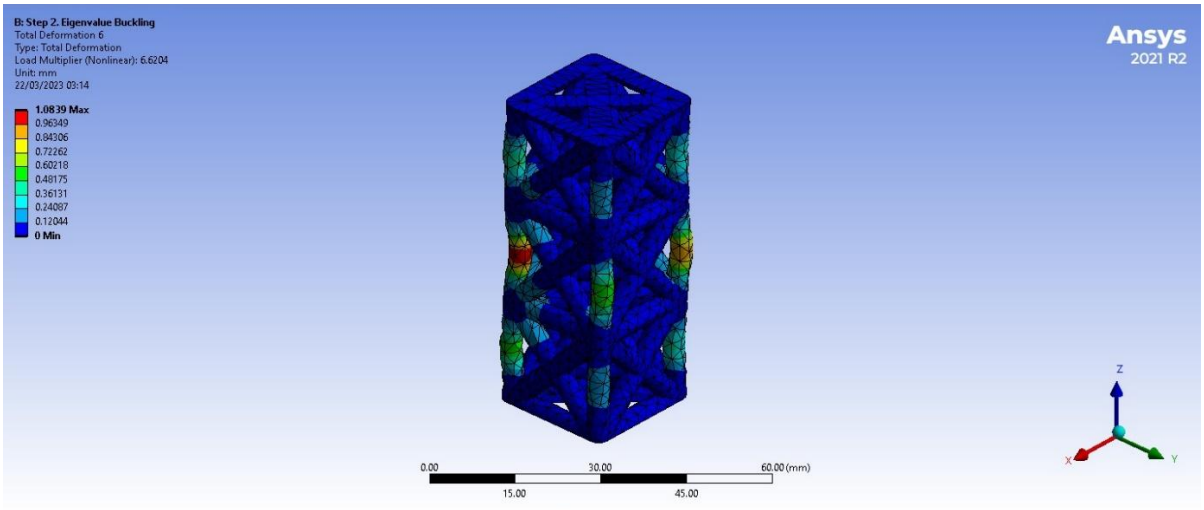
Mode 4 – Load Multiplier 6.3787



Mode 5 – Load Multiplier 6.5725

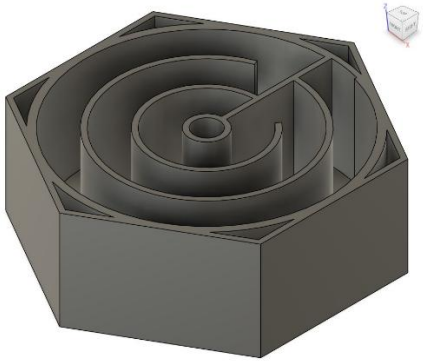


Mode 6 – Load Multiplier 6.6204

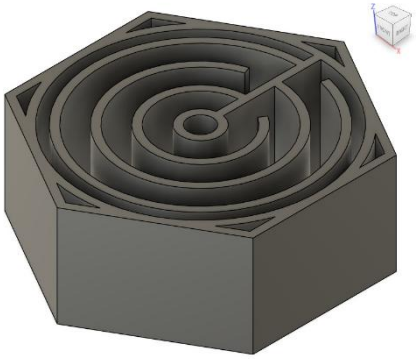


Appendix J Solution Unit Cell Geometries of Inverse Optimised Metamaterials

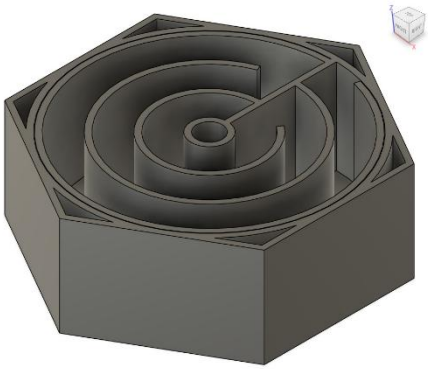
Example A



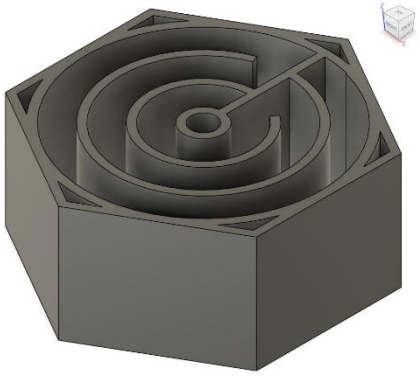
Example B



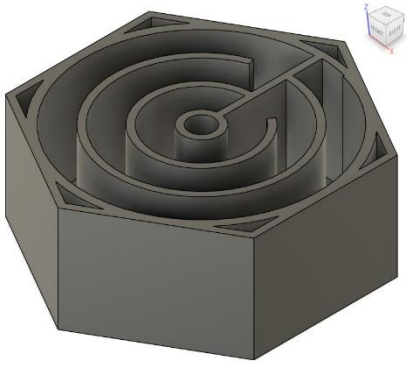
Example C



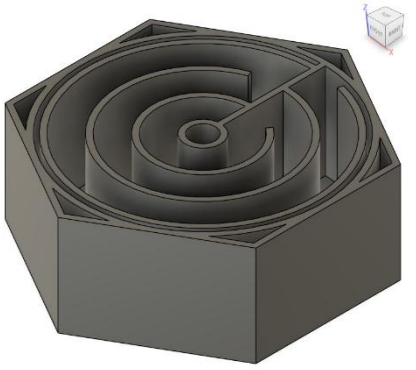
Example D



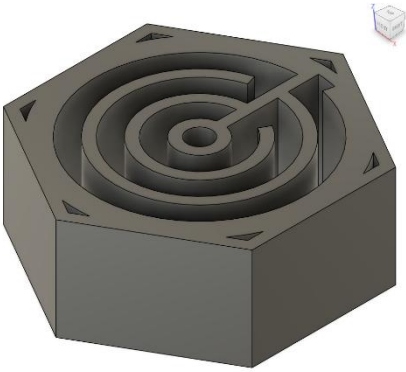
Example E



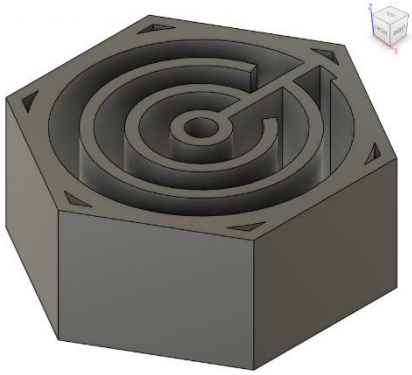
Example F



Example G

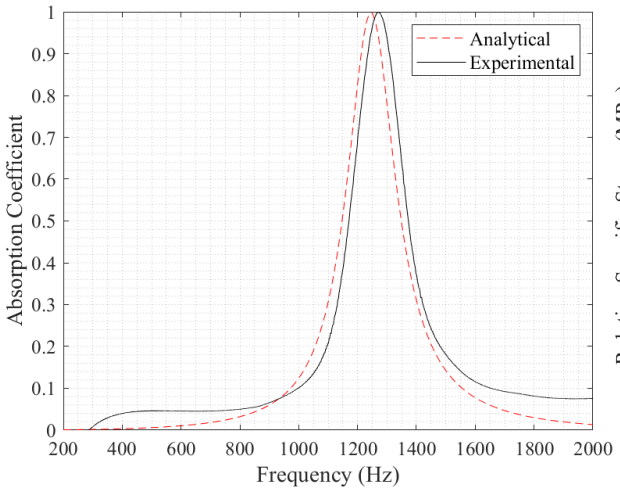


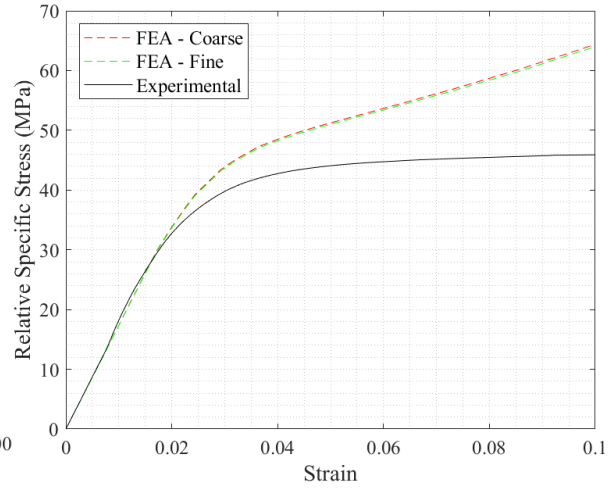
Example H



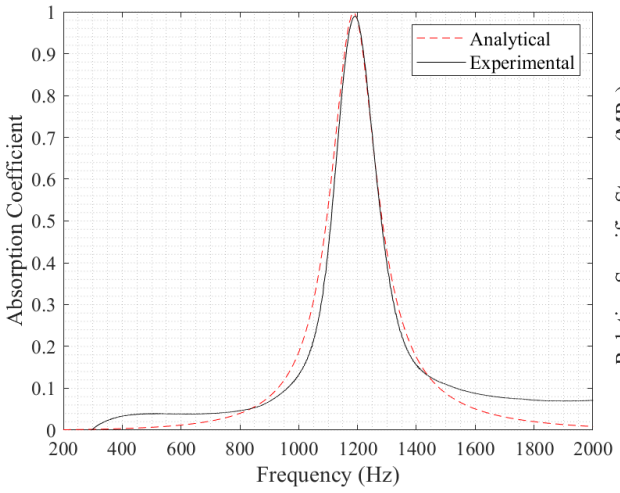
Appendix K Material Properties of Genetic Algorithm Optimised Pareto Front Solutions

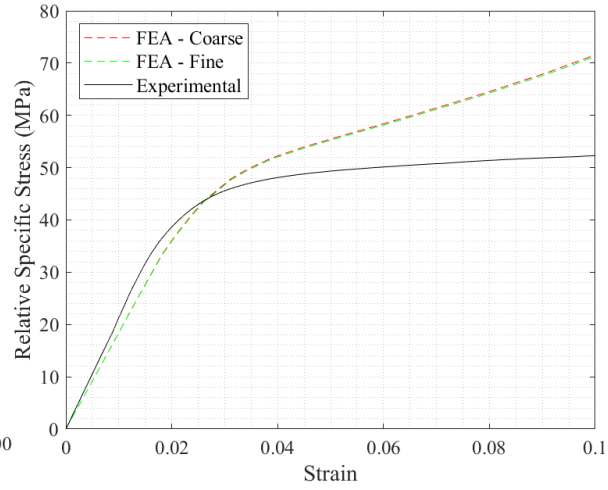
Solution 1



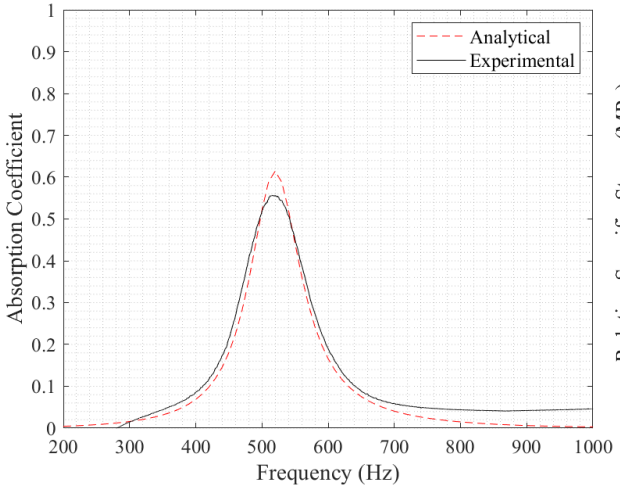


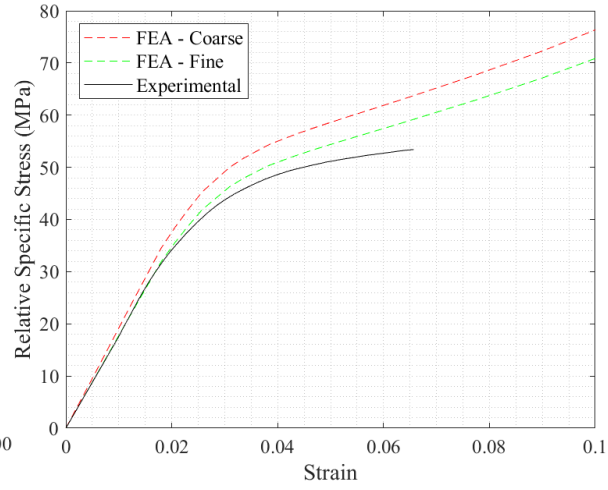
Solution 2



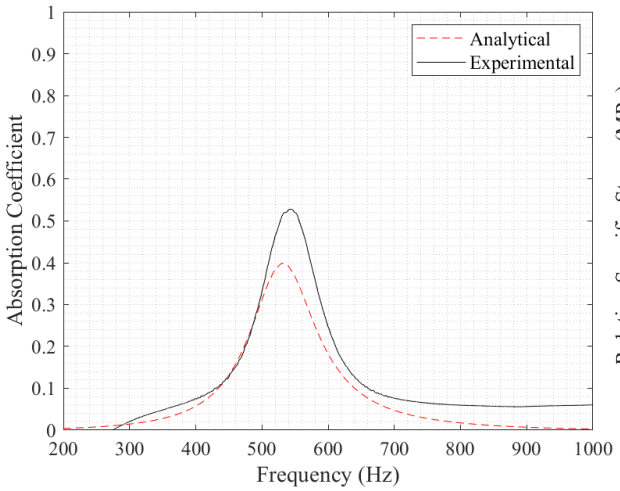


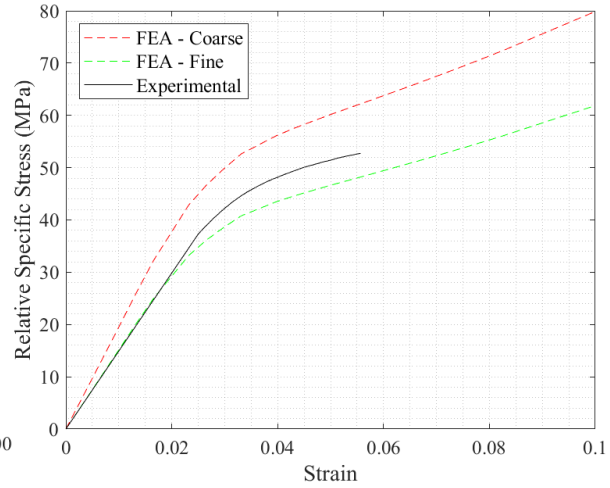
Solution 3



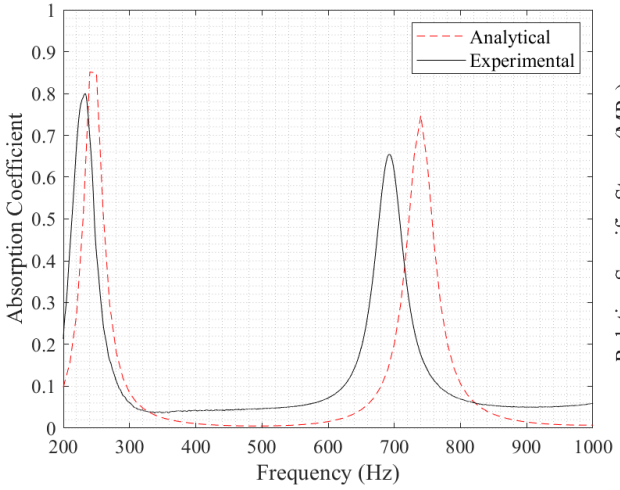


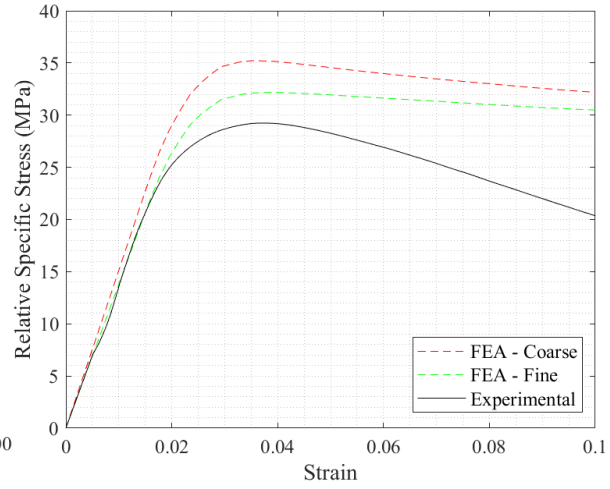
Solution 4



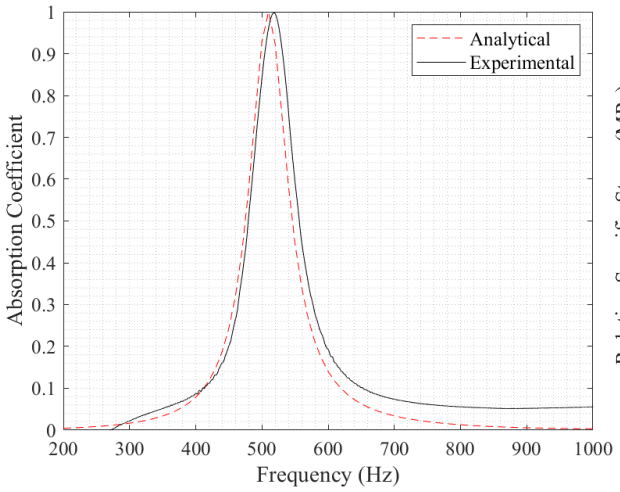


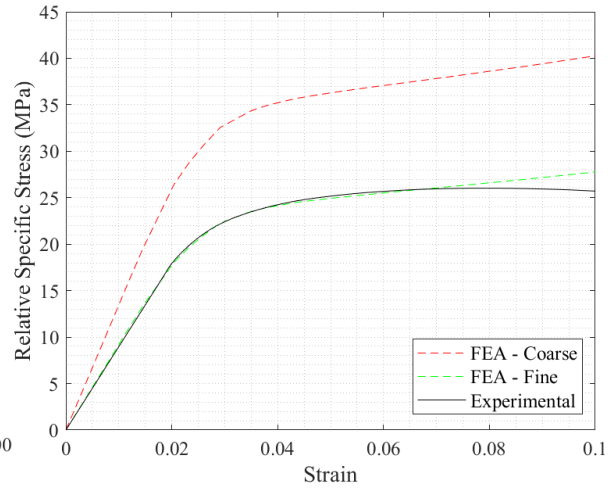
Solution 5



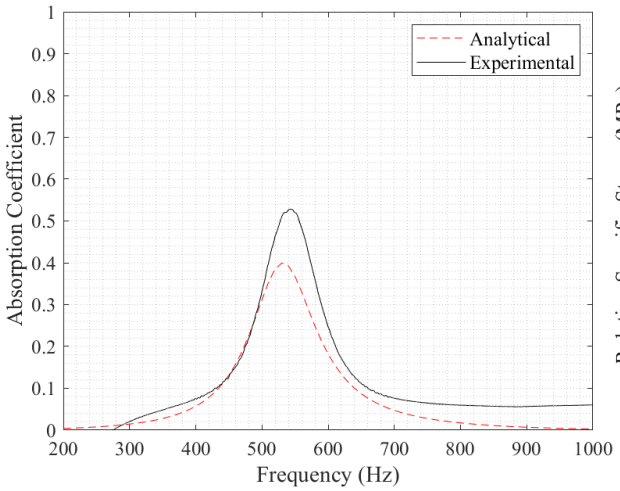


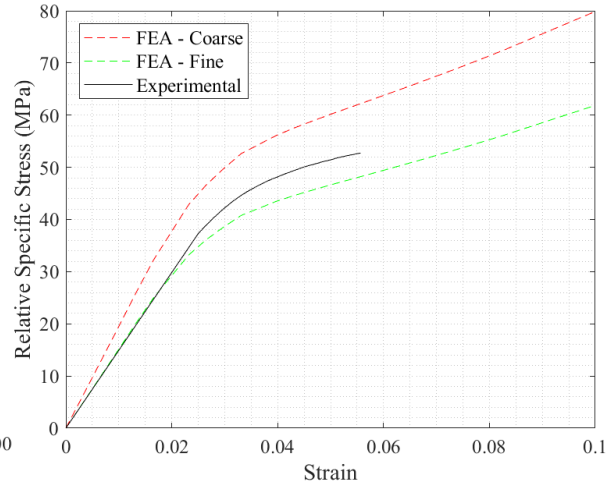
Solution 6



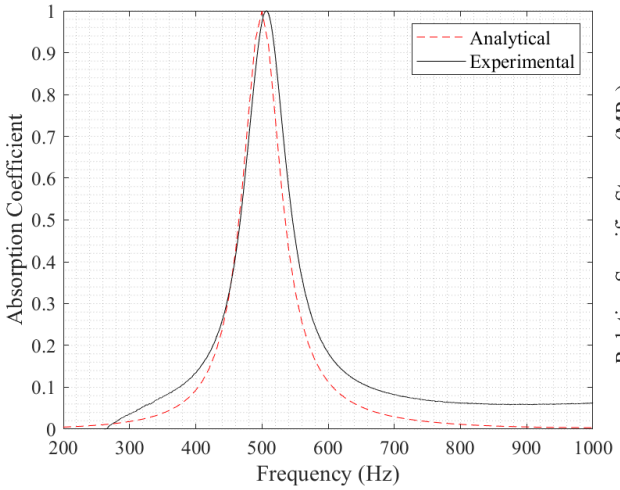


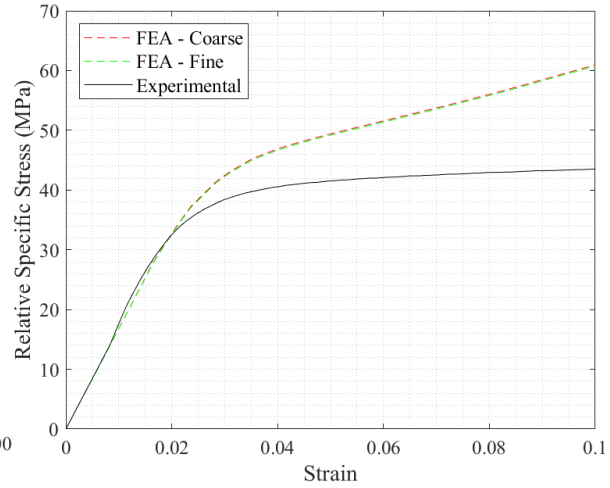
Solution 7



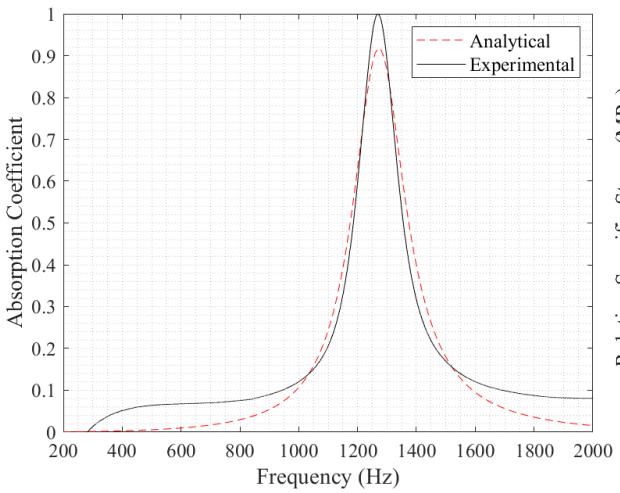


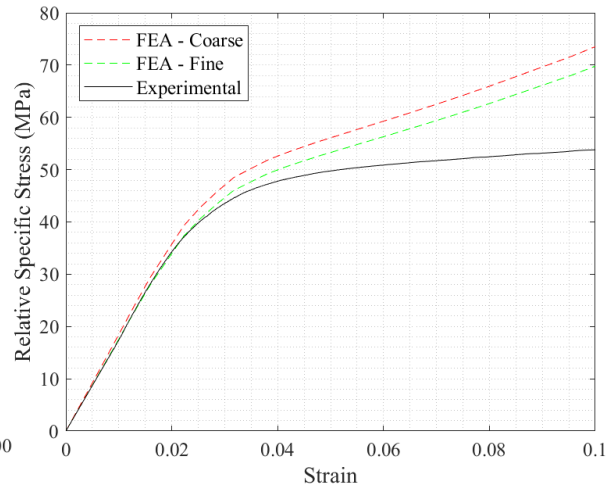
Solution 8



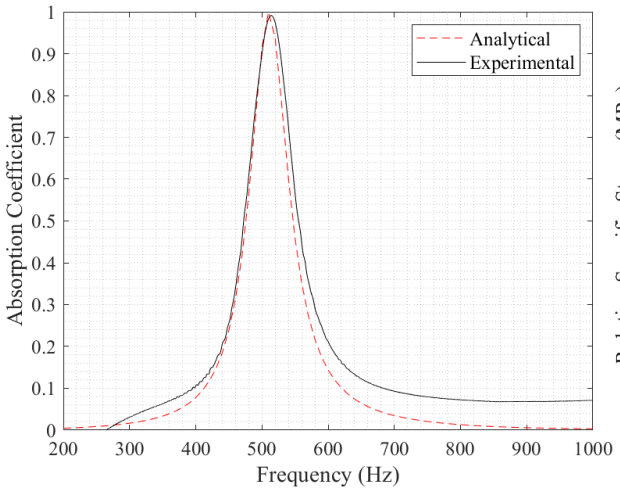


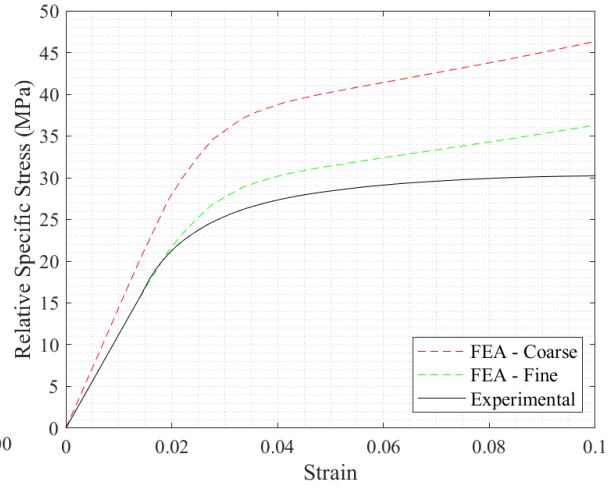
Solution 9



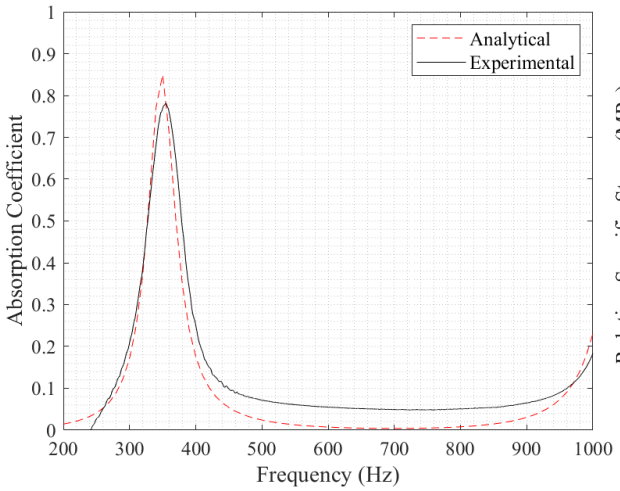


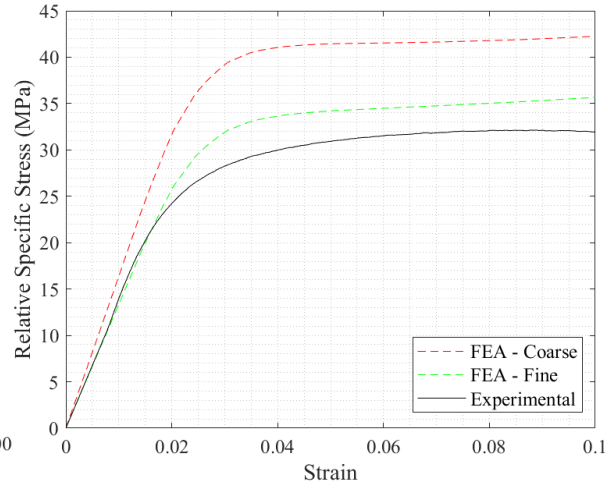
Solution 10



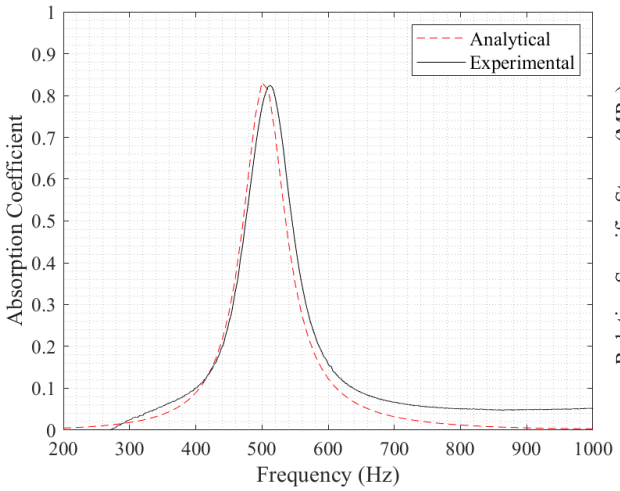


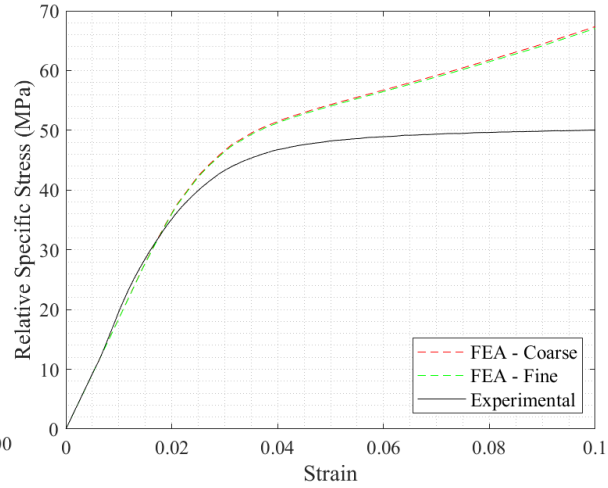
Solution 11



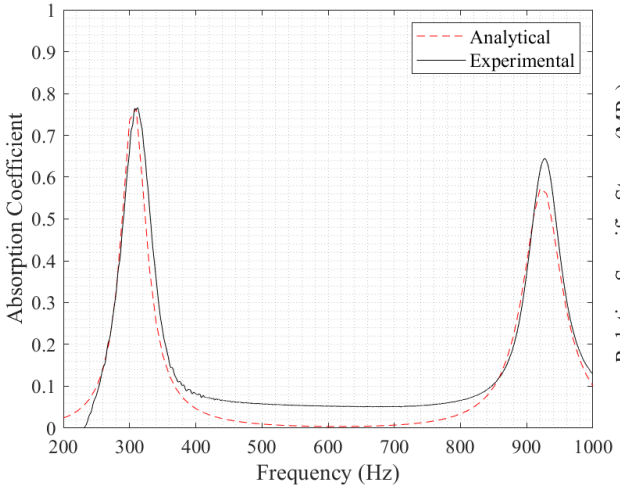


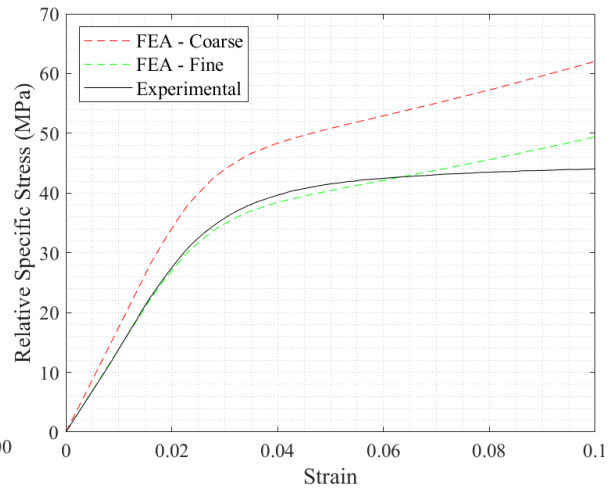
Solution 12



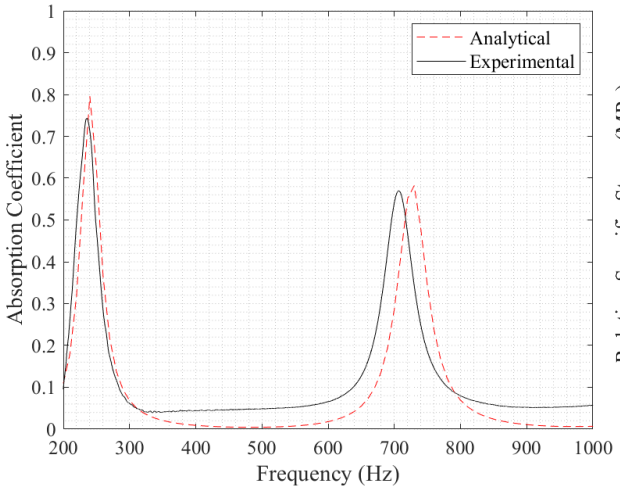


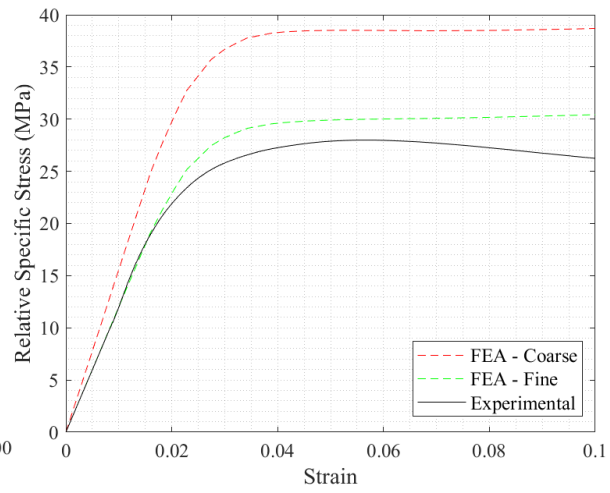
Solution 13



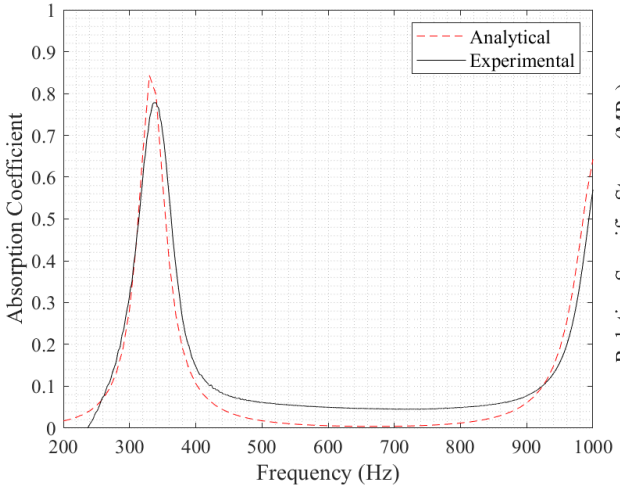


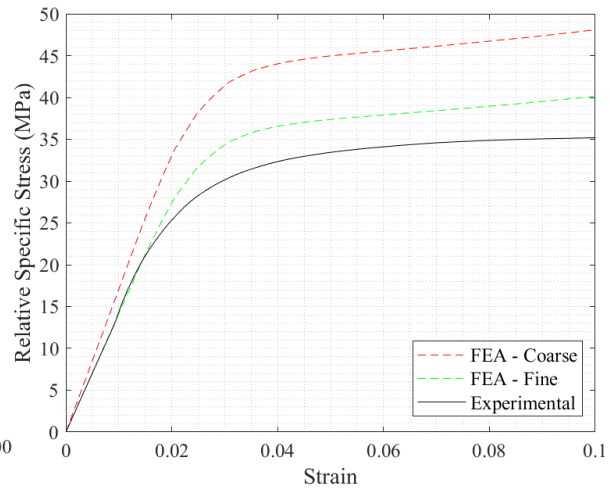
Solution 14



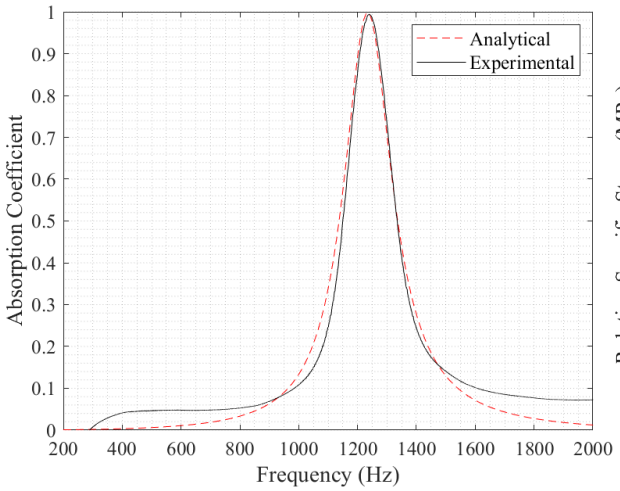


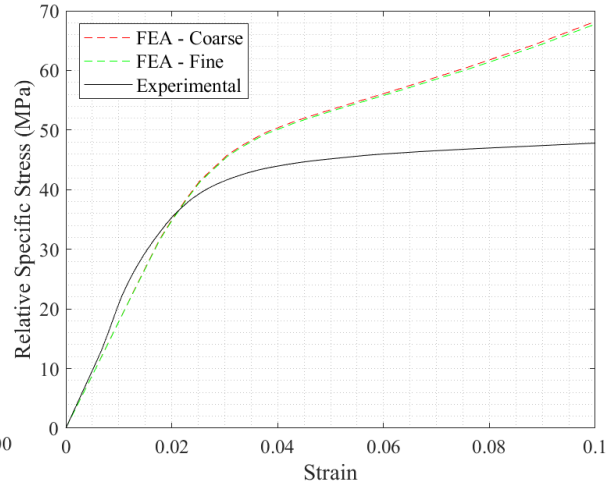
Solution 15



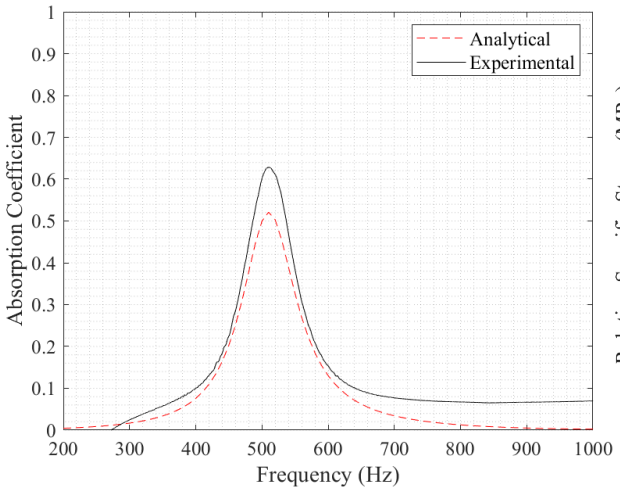


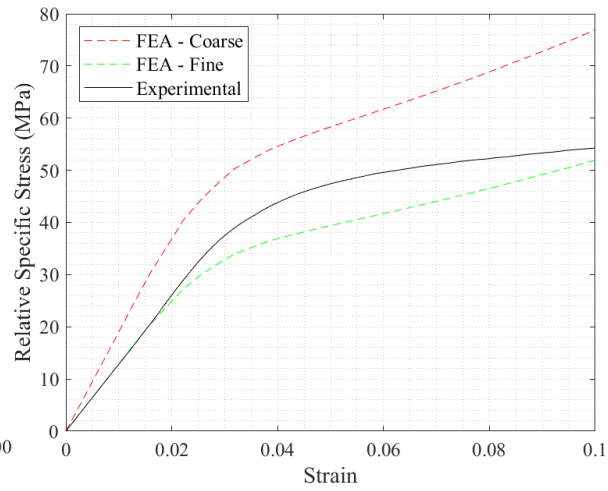
Solution 16



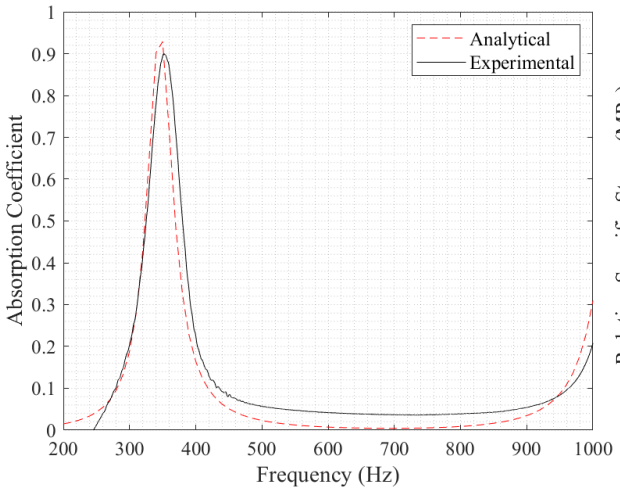


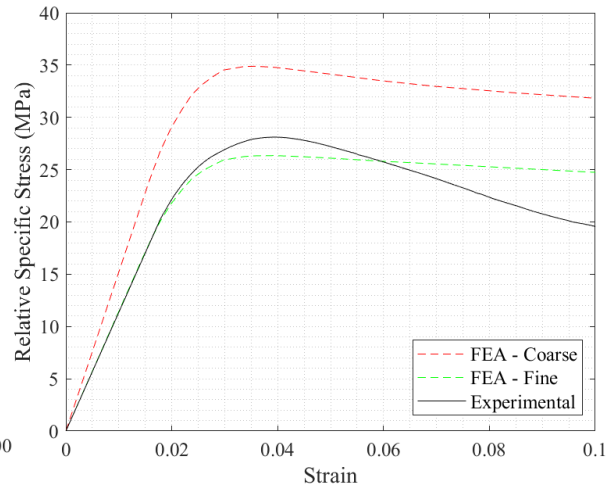
Solution 17





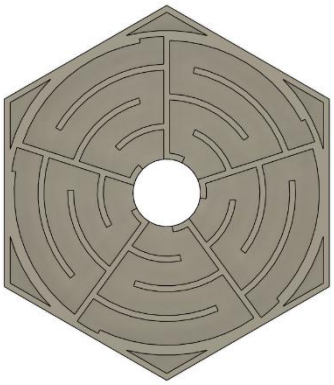
Solution 18





Appendix L Solution Geometries of Low Thickness Inverse Optimised Metamaterials

Example I



Example J



Example K



Example L



Example M



Example N



Example O



Appendix M Pareto Front Solution Parameters Optimised by GA to Maximise Relative Specific Stiffness and Strength

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Supercell Mass (g)	Relative Density
1	8.98	10.98	9.97	10.99	10.85	10.13	10.99	10.47	9.89	10.71	12.05	11.65	2.75	2.75	3.00	2.18	1291.80	60.64	35.94	0.77
2	8.91	10.66	10.64	10.95	10.63	10.29	10.55	10.18	9.95	10.64	11.93	12.21	2.74	2.77	3.00	2.19	1298.58	59.69	36.09	0.78
3	9.60	11.70	9.88	11.24	11.27	10.39	11.29	10.82	9.71	10.68	13.23	10.61	2.75	2.53	3.00	2.18	1253.22	62.80	34.22	0.74
4	8.86	10.68	10.71	11.07	10.62	10.31	10.32	10.29	10.60	10.68	11.93	12.39	2.75	2.77	3.00	2.22	1302.26	59.07	36.20	0.78
5	8.85	10.88	10.95	11.12	10.51	10.34	10.31	10.43	10.71	10.88	11.86	12.46	2.76	2.80	3.00	2.26	1303.92	58.70	36.36	0.78
6	9.49	11.70	9.95	11.31	11.32	10.38	11.30	10.81	9.70	10.67	13.23	10.63	2.75	2.52	3.00	2.18	1251.82	62.82	34.58	0.75
7	8.85	10.65	10.74	10.98	10.59	10.28	10.52	10.26	10.20	10.66	11.91	12.30	2.75	2.77	3.00	2.21	1300.40	59.37	36.15	0.78
8	8.91	10.61	10.89	11.08	10.44	10.14	10.27	10.41	10.94	11.18	11.89	12.96	2.76	2.80	3.00	2.35	1305.75	57.75	36.62	0.79
9	9.15	11.46	9.98	11.08	11.09	10.43	11.08	10.59	9.71	10.69	12.68	11.40	2.76	2.66	3.00	2.17	1280.71	61.61	35.49	0.76
10	8.77	10.83	10.95	11.13	10.52	10.40	10.31	10.43	10.65	10.79	11.86	12.41	2.76	2.80	3.00	2.22	1303.41	58.85	36.35	0.78
11	9.20	11.49	9.83	11.11	11.19	10.49	11.20	10.71	9.64	10.69	12.87	11.29	2.75	2.65	3.00	2.17	1276.26	61.98	35.45	0.76
12	8.98	10.79	10.10	11.00	10.85	10.15	10.96	10.46	9.95	10.68	11.99	11.75	2.75	2.76	3.00	2.19	1292.86	60.42	35.99	0.78
13	8.87	10.66	10.61	11.09	10.56	10.29	10.36	10.43	10.62	10.71	11.89	12.36	2.75	2.78	3.00	2.19	1302.00	59.14	36.23	0.78
14	8.93	10.65	10.65	11.06	10.89	10.42	10.71	10.31	10.06	10.70	11.93	12.15	2.76	2.77	3.00	2.19	1297.54	59.77	36.18	0.78
15	9.20	11.49	9.84	11.08	11.07	10.45	11.08	10.59	9.79	10.70	12.49	11.46	2.75	2.66	3.00	2.17	1282.10	61.53	35.42	0.76
16	9.11	11.29	9.89	11.11	11.00	10.35	11.02	10.49	9.83	10.69	12.42	11.68	2.75	2.69	3.00	2.18	1286.78	61.17	35.67	0.77
17	9.19	11.48	9.84	11.09	11.12	10.43	11.10	10.66	9.70	10.69	12.69	11.34	2.75	2.66	3.00	2.17	1279.59	61.70	35.45	0.76
18	8.97	11.01	10.15	11.05	10.76	10.13	10.87	10.34	9.97	10.73	12.01	11.91	2.74	2.75	3.00	2.19	1293.55	60.35	35.97	0.78
19	9.09	11.08	10.15	11.03	10.77	10.20	10.94	10.39	9.90	10.70	12.02	11.75	2.75	2.75	3.00	2.18	1291.98	60.57	35.88	0.77
20	8.89	10.66	10.90	11.00	10.16	10.14	10.33	10.44	10.98	10.94	11.63	12.73	2.76	2.79	3.00	2.33	1305.71	57.99	36.47	0.79
21	9.18	11.40	9.97	11.09	11.01	10.40	11.09	10.57	9.72	10.69	12.39	11.54	2.75	2.70	3.00	2.17	1285.35	61.27	35.66	0.77
22	9.34	11.49	9.89	11.16	11.22	10.36	11.18	10.68	9.74	10.71	12.89	11.23	2.75	2.63	3.00	2.17	1274.57	62.06	35.25	0.76
23	8.94	10.67	10.71	11.07	10.47	10.28	10.58	10.28	10.10	10.69	11.93	12.19	2.75	2.77	3.00	2.20	1299.22	59.51	36.12	0.78
24	8.83	10.62	10.99	11.09	10.46	10.17	10.28	10.40	10.88	10.82	11.88	12.49	2.76	2.80	3.00	2.34	1305.54	58.26	36.44	0.79
25	9.51	11.56	9.90	11.17	11.29	10.37	11.30	10.78	9.70	10.67	13.23	10.68	2.75	2.55	3.00	2.17	1257.31	62.75	34.77	0.75
26	9.37	11.48	9.89	11.19	11.18	10.43	11.18	10.68	9.75	10.69	13.11	10.86	2.75	2.61	3.00	2.15	1267.86	62.45	35.12	0.76
27	8.87	10.67	10.69	11.01	10.60	10.32	10.48	10.21	10.26	10.85	11.90	12.33	2.75	2.78	3.00	2.22	1301.21	59.21	36.25	0.78
28	8.86	10.71	10.74	11.10	10.47	10.28	10.50	10.25	10.11	10.69	11.93	12.34	2.75	2.78	3.00	2.23	1300.90	59.32	36.24	0.78
29	8.92	10.66	10.74	11.10	10.47	10.28	10.50	10.25	10.11	10.69	11.93	12.32	2.75	2.78	3.00	2.16	1300.10	59.46	36.21	0.78
30	8.83	10.62	11.01	11.08	10.46	10.17	10.28	10.41	10.85	10.82	11.88	12.47	2.76	2.80	3.00	2.29	1305.34	58.39	36.44	0.79
31	8.84	10.66	10.94	11.10	10.48	10.23	10.28	10.40	10.74	10.81	11.89	12.45	2.76	2.80	3.00	2.29	1304.92	58.52	36.40	0.78
32	9.13	11.14	10.09	11.07	11.06	10.31	10.98	10.51	9.85	10.69	12.35	11.46	2.75	2.72	3.00	2.18	1287.27	61.03	35.67	0.77
33	9.54	11.61	9.89	11.31	11.29	10.45	11.29	10.81	9.71	10.69	13.23	10.69	2.75	2.58	3.00	2.15	1259.67	62.74	34.93	0.75
34	9.57	11.55	9.87	11.34	11.18	10.47	11.29	10.66	9.72	10.69	13.07	11.17	2.76	2.63	3.00	2.17	1271.15	62.27	35.30	0.76
35	8.87	10.65	10.81	11.12	10.51	10.27	10.30	10.41	10.67	10.78	11.89	12.40	2.75	2.80	3.00	2.20	1303.38	58.90	36.33	0.78
36	9.47	11.51	9.94	11.26	11.26	10.36	11.29	10.78	9.72	10.65	13.16	10.70	2.75	2.58	3.00	2.17	1261.03	62.68	34.84	0.75
37	8.91	10.64	10.86	11.07	10.60	10.35	10.31	10.25	10.26	10.71	11.89	12.46	2.75	2.79	3.00	2.25	1303.22	58.97	36.29	0.78
38	8.92	10.68	10.40	10.97	10.64	10.41	10.55	10.21	9.88	10.70	11.94	12.12	2.75	2.77	3.00	2.18	1297.24	59.89	36.09	0.78

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Supercell Mass (g)	Relative Density
39	8.86	10.81	10.75	11.10	10.53	10.32	10.35	10.43	10.66	10.80	11.89	12.43	2.76	2.79	3.00	2.20	1302.49	59.02	36.28	0.78
40	8.88	10.66	10.69	11.09	10.61	10.36	10.70	10.30	10.09	10.70	11.93	12.31	2.75	2.78	3.00	2.20	1299.03	59.55	36.21	0.78
41	8.98	11.15	10.05	11.03	10.85	10.26	11.00	10.49	9.86	10.71	12.12	11.70	2.75	2.74	3.00	2.18	1290.79	60.74	35.89	0.77
42	8.92	10.79	10.64	11.05	10.71	10.30	10.75	10.28	9.86	10.69	11.99	12.14	2.74	2.76	3.00	2.19	1296.65	59.96	36.11	0.78
43	9.55	11.62	9.92	11.18	11.23	10.39	11.24	10.82	9.79	10.68	13.22	10.79	2.75	2.60	3.00	2.16	1263.98	62.56	35.02	0.75
44	9.39	11.50	9.91	11.19	11.20	10.41	11.24	10.70	9.80	10.69	13.14	10.84	2.75	2.61	3.00	2.15	1266.80	62.49	35.11	0.76
45	9.01	11.05	9.97	11.07	10.86	10.19	11.01	10.48	9.80	10.69	12.27	11.36	2.75	2.74	3.00	2.17	1288.58	60.94	35.87	0.77
46	9.14	11.33	9.86	11.11	11.11	10.42	11.04	10.54	9.78	10.69	12.43	11.38	2.75	2.68	3.00	2.18	1283.37	61.41	35.51	0.77
47	9.33	11.46	9.91	11.11	11.23	10.39	11.24	10.64	9.74	10.68	13.17	10.86	2.75	2.62	3.00	2.14	1268.14	62.40	35.27	0.76
48	9.13	11.22	9.93	11.16	10.90	10.18	11.04	10.51	9.87	10.68	11.99	11.42	2.75	2.76	3.00	2.18	1289.34	60.83	35.85	0.77
49	9.35	11.52	9.86	11.16	11.14	10.33	11.25	10.76	9.73	10.61	13.09	10.77	2.75	2.60	3.00	2.13	1265.58	62.55	35.03	0.76
50	9.17	11.32	9.83	11.11	11.05	10.35	11.02	10.49	9.83	10.69	12.42	11.75	2.75	2.69	3.00	2.22	1287.00	61.11	35.68	0.77
51	9.22	11.25	9.88	11.16	11.12	10.33	11.18	10.63	9.70	10.71	12.77	11.28	2.77	2.65	3.00	2.17	1278.29	61.79	35.47	0.76
52	9.34	11.53	9.84	11.21	11.20	10.51	11.29	10.65	9.72	10.70	13.00	11.18	2.76	2.63	3.00	2.17	1272.68	62.23	35.35	0.76
53	8.96	10.89	10.18	11.05	10.75	10.15	10.77	10.29	9.87	10.70	12.00	12.03	2.74	2.75	3.00	2.19	1294.48	60.24	36.03	0.78
54	8.83	10.79	10.91	11.13	10.52	10.34	10.30	10.43	10.67	10.83	11.86	12.42	2.75	2.80	3.00	2.25	1303.74	58.78	36.33	0.78
55	9.18	11.42	9.91	11.13	11.15	10.29	11.08	10.52	9.82	10.68	12.69	11.32	2.75	2.63	3.00	2.15	1277.45	61.83	35.30	0.76
56	9.06	11.13	9.92	11.08	10.91	10.23	11.11	10.56	9.83	10.69	12.26	11.45	2.75	2.74	3.00	2.17	1288.19	61.00	35.85	0.77
57	9.28	11.52	9.86	11.17	11.07	10.40	11.23	10.70	9.73	10.64	13.05	11.07	2.75	2.62	3.00	2.16	1270.72	62.30	35.20	0.76
58	8.86	10.67	10.63	11.06	10.79	10.32	11.02	10.23	9.86	10.67	12.17	11.96	2.77	2.77	3.00	2.22	1295.44	60.08	36.33	0.78
59	9.15	11.47	10.08	11.07	11.03	10.39	10.99	10.55	9.84	10.69	12.37	11.46	2.75	2.67	3.00	2.17	1284.28	61.34	35.41	0.76
60	9.06	10.96	10.16	11.06	10.90	10.31	10.97	10.30	9.85	10.69	12.11	11.49	2.75	2.73	3.00	2.19	1289.78	60.80	35.77	0.77
61	8.83	10.72	10.98	11.12	10.48	10.21	10.29	10.42	10.87	10.82	11.87	12.46	2.76	2.80	3.00	2.31	1305.06	58.43	36.42	0.78
62	9.27	11.52	9.84	11.17	11.13	10.43	11.23	10.65	9.72	10.69	12.86	11.31	2.76	2.66	3.00	2.17	1277.28	61.90	35.47	0.76
63	9.42	11.68	9.99	11.22	11.15	10.39	11.26	10.63	9.75	10.67	13.19	11.03	2.75	2.62	3.00	2.17	1268.70	62.35	35.25	0.76
64	9.46	11.58	9.87	11.18	11.18	10.34	11.27	10.76	9.73	10.65	13.15	10.72	2.75	2.60	3.00	2.13	1263.87	62.62	35.02	0.76
65	9.18	11.47	9.90	11.09	11.12	10.43	11.09	10.66	9.71	10.69	12.69	11.38	2.76	2.66	3.00	2.17	1280.13	61.66	35.46	0.76
66	9.27	11.49	9.85	11.13	11.18	10.49	11.20	10.69	9.73	10.69	12.99	11.19	2.78	2.63	3.00	2.16	1273.56	62.15	35.40	0.76
67	8.83	10.62	10.99	11.09	10.46	10.17	10.28	10.40	10.88	10.85	11.88	12.55	2.76	2.80	3.00	2.36	1305.65	58.17	36.48	0.79
68	9.24	11.49	9.88	11.14	11.18	10.48	11.20	10.69	9.72	10.71	12.99	11.19	2.75	2.63	3.00	2.16	1273.37	62.16	35.31	0.76
69	9.04	10.88	10.65	10.98	10.72	10.50	10.88	10.31	9.92	10.68	11.93	12.21	2.75	2.76	3.00	2.19	1295.81	60.02	36.04	0.78
70	8.86	10.66	10.94	11.11	10.48	10.25	10.28	10.40	10.74	10.81	11.89	12.45	2.76	2.80	3.00	2.26	1304.67	58.61	36.39	0.78
																Max	1305.75	62.82	36.62	0.79
																Min	1251.82	57.75	34.22	0.74

Appendix N Pareto Front Solution Parameters Optimised by GA to Maximise Relative Specific Stiffness and Strength While Minimising Mass

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Supercell Mass (g)	Relative Density
1	6.08	6.42	7.96	8.19	7.85	8.57	6.69	7.53	6.67	7.83	5.56	7.41	1.52	1.23	0.83	1.45	584.48	24.38	10.33	0.25
2	11.52	8.40	9.38	10.68	10.22	9.78	9.22	10.68	11.01	11.50	11.00	9.86	2.69	2.22	3.00	2.12	1353.13	61.38	33.26	0.72
3	8.38	6.10	14.33	9.28	8.55	13.11	8.32	14.58	9.86	9.02	13.29	12.05	1.06	0.90	0.56	1.09	476.55	10.90	6.30	0.16
4	11.15	9.02	9.40	10.53	10.67	9.02	9.34	10.12	10.98	10.27	10.33	10.97	2.77	2.80	3.00	2.09	1354.12	61.30	36.34	0.78
5	11.15	9.02	9.40	10.54	10.67	9.02	9.34	10.12	10.99	10.27	10.33	10.97	2.77	2.80	3.00	2.09	1354.12	61.30	36.34	0.78
6	8.06	7.01	8.99	8.12	7.84	9.54	7.79	8.48	7.94	8.51	6.96	8.28	1.55	1.04	0.90	1.26	691.37	29.11	9.02	0.22
7	8.33	5.97	14.83	8.40	8.57	13.18	8.76	14.74	9.88	8.87	13.25	11.78	1.28	0.94	0.59	1.07	509.75	14.81	7.32	0.18
8	8.46	8.46	9.06	9.67	9.31	9.38	9.15	9.95	9.27	10.01	9.57	10.05	2.50	2.13	2.90	1.97	1339.98	60.40	30.94	0.68
9	8.37	7.39	8.97	8.42	8.64	9.13	8.26	8.92	8.13	9.09	8.29	9.05	1.71	1.48	2.51	1.58	1242.47	53.50	21.51	0.51
10	6.06	6.80	8.37	8.17	7.89	7.91	7.99	8.84	7.44	8.36	6.65	7.74	1.61	1.19	0.82	1.39	628.12	23.27	10.05	0.24
11	8.27	6.07	14.48	9.29	8.60	13.06	8.60	14.60	9.85	8.92	13.27	11.15	1.14	0.91	0.56	1.08	497.03	12.58	6.54	0.16
12	8.43	6.04	11.23	9.06	7.87	10.41	8.13	10.36	9.88	8.89	11.38	10.31	1.98	0.96	0.76	1.31	667.01	26.74	10.69	0.25
13	8.36	6.03	14.49	8.91	8.52	13.15	8.45	14.71	9.87	8.84	13.26	12.17	1.16	0.92	0.58	1.09	490.29	12.51	6.88	0.17
14	5.89	6.37	7.90	6.18	6.97	7.22	6.72	6.40	6.30	8.18	5.32	7.08	1.36	1.37	0.83	1.43	562.01	19.07	10.33	0.25
15	6.41	7.03	8.45	8.34	7.90	9.29	8.22	8.95	7.82	8.51	6.60	7.74	1.67	1.09	0.82	1.37	632.97	25.64	10.06	0.24
16	7.73	6.89	8.79	8.32	8.06	8.53	8.28	8.89	8.17	8.91	9.19	8.87	1.68	1.15	2.00	1.27	1053.69	43.54	16.73	0.40
17	7.18	6.86	8.38	8.21	7.70	8.57	7.70	8.14	7.97	8.46	6.67	8.48	1.63	1.31	2.37	1.53	1151.76	48.78	19.68	0.47
18	7.91	7.87	9.20	8.55	8.03	8.58	8.20	9.15	8.33	8.91	9.18	9.04	1.77	1.37	2.62	1.60	1302.18	56.44	22.10	0.52
19	6.75	6.97	8.55	8.30	7.73	8.81	7.97	8.45	7.45	8.57	7.25	8.20	1.90	1.31	1.52	1.51	846.89	35.47	15.41	0.36
20	7.49	7.01	8.47	8.16	7.73	8.71	7.79	8.27	8.06	8.58	6.89	7.94	1.51	1.18	2.18	1.48	1099.25	44.85	17.45	0.42
21	10.33	8.23	9.10	8.59	10.16	8.94	8.50	8.98	9.77	9.67	9.18	10.41	2.47	2.23	2.95	1.67	1346.79	61.12	31.72	0.70
22	8.28	6.09	14.49	9.20	8.61	13.06	8.67	14.45	9.87	9.33	13.28	11.13	1.23	0.91	0.56	1.08	513.03	13.54	6.78	0.17
23	8.15	6.93	8.44	8.33	7.88	8.65	8.01	8.24	8.10	8.46	7.96	8.34	1.64	1.34	2.42	1.53	1230.28	51.91	19.58	0.47
24	7.09	6.86	8.35	8.21	7.49	8.53	7.68	8.14	7.92	8.43	6.61	8.49	1.62	1.31	2.53	1.54	1202.64	51.40	20.82	0.50
25	7.62	7.05	8.82	8.31	8.15	8.85	8.11	8.61	8.02	8.60	7.24	8.39	1.65	1.42	2.47	1.55	1170.90	49.92	21.37	0.51
26	8.26	6.08	14.38	9.29	8.60	13.05	8.56	14.53	9.85	9.03	13.28	11.11	0.93	0.91	0.56	1.08	450.29	9.42	6.16	0.16
27	7.80	6.89	8.79	7.63	7.59	8.61	7.66	7.93	7.74	8.60	6.90	8.45	1.63	1.32	2.08	1.52	1059.37	44.07	17.39	0.41
28	8.02	7.13	8.39	8.16	7.62	9.05	7.81	8.05	8.03	8.90	7.66	8.34	1.61	1.35	1.91	1.56	1018.14	41.61	16.19	0.39
29	7.24	6.76	9.28	7.54	8.13	8.91	8.03	7.41	7.71	8.58	7.75	8.02	1.48	1.16	1.71	1.43	911.14	35.38	14.27	0.35
30	7.42	7.26	8.70	8.39	8.77	8.61	8.67	8.86	7.68	8.56	7.79	8.40	1.62	1.37	1.84	1.57	973.75	39.23	16.34	0.39
31	7.15	6.90	8.25	8.18	7.75	8.71	7.72	7.93	7.60	8.45	6.67	8.17	1.60	1.28	2.05	1.45	1030.64	42.13	17.12	0.41
32	7.27	7.05	8.81	8.27	8.11	8.89	7.92	8.43	8.07	8.52	7.15	8.61	1.75	1.31	2.56	1.42	1214.00	52.90	21.85	0.52
33	6.49	6.86	8.22	7.83	7.43	8.72	7.71	8.05	7.30	8.36	6.07	8.05	1.54	1.30	1.96	1.52	943.37	37.87	17.13	0.41
34	7.13	6.99	8.17	8.21	7.86	8.60	7.14	7.69	7.99	8.40	6.17	7.49	1.52	1.23	1.79	1.46	929.15	36.61	15.02	0.36
35	7.24	7.00	8.39	8.20	7.70	8.69	7.71	8.20	8.00	8.55	6.71	8.17	1.61	1.27	2.31	1.51	1133.60	47.57	19.02	0.45
36	7.74	7.49	8.85	8.56	8.12	8.64	8.08	9.04	8.36	8.79	7.54	9.22	1.74	1.34	2.58	1.56	1268.90	55.19	21.63	0.51

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Supercell Mass (g)	Relative Density
37	7.47	7.80	8.53	8.32	8.06	8.73	8.00	9.10	8.32	8.70	7.10	8.93	1.60	1.29	2.20	1.59	1113.96	46.52	18.34	0.44
38	8.12	7.02	8.99	7.44	7.28	9.29	7.83	7.63	8.72	8.70	7.04	7.98	1.96	1.60	0.77	1.52	668.49	26.38	13.01	0.30
39	7.58	6.83	10.23	8.45	9.16	8.76	8.04	8.27	9.05	8.65	7.85	8.68	2.08	1.29	2.67	1.51	1263.72	57.07	23.87	0.55
40	7.81	6.39	8.66	8.32	7.85	8.73	7.92	8.96	8.89	7.87	10.43	9.14	1.96	1.15	0.79	1.43	676.22	25.90	11.12	0.26
41	8.42	6.02	14.33	9.31	8.54	13.16	8.31	14.62	9.86	8.73	13.27	12.32	0.93	0.90	0.56	1.09	450.49	9.40	6.01	0.15
42	6.59	6.97	8.59	8.34	7.57	9.25	7.49	7.05	7.54	8.46	6.23	7.41	1.42	1.35	0.83	1.37	622.90	21.90	9.83	0.24
43	9.57	8.69	9.35	10.41	9.36	9.33	9.16	9.84	10.35	9.98	10.42	10.06	2.02	2.30	2.84	1.94	1335.82	59.52	29.00	0.67
44	7.61	7.22	8.69	8.25	7.75	8.96	7.79	8.15	8.32	8.57	6.97	8.47	1.63	1.22	2.34	1.54	1183.43	50.30	18.64	0.44
45	11.23	8.57	9.37	10.62	10.19	9.64	9.22	10.22	11.00	10.42	11.00	9.87	2.48	2.27	2.99	1.97	1367.95	61.72	32.25	0.71
46	5.90	6.74	7.98	6.35	7.37	8.06	6.98	6.45	6.56	8.26	5.37	7.34	1.37	1.37	0.83	1.44	567.54	19.36	10.34	0.25
47	7.86	7.13	8.40	8.16	7.72	8.75	7.81	8.16	8.01	8.83	7.47	8.28	1.67	1.26	1.96	1.51	1052.10	43.35	16.21	0.39
48	6.15	7.03	8.00	7.70	8.15	7.91	7.31	6.78	6.71	8.18	5.93	9.09	1.74	1.36	1.45	1.49	746.56	30.22	15.31	0.36
49	8.00	7.42	8.83	8.35	8.09	9.13	8.48	9.10	8.06	8.55	7.46	8.49	1.83	1.26	1.00	1.49	739.82	30.11	11.49	0.27
50	8.78	8.71	9.32	10.53	8.85	8.92	9.31	8.94	7.99	10.20	7.79	8.04	1.97	1.40	2.94	1.97	1423.68	62.03	25.78	0.60
51	8.96	7.56	9.68	9.28	9.49	9.77	9.12	10.10	8.74	10.06	9.44	10.62	2.34	1.35	1.15	1.93	785.82	34.26	15.76	0.35
52	11.26	8.58	9.40	10.60	10.50	9.08	9.31	10.15	11.00	11.21	10.34	10.46	2.76	2.48	3.00	2.10	1355.38	61.33	34.75	0.75
53	7.13	6.98	8.51	8.28	7.85	8.83	7.29	7.52	7.84	8.41	6.20	7.43	1.52	1.27	1.69	1.37	873.30	34.17	14.80	0.36
54	6.87	7.06	8.40	7.13	7.76	8.64	7.88	8.23	7.11	8.38	7.13	8.07	1.74	1.31	1.68	1.50	886.99	36.17	15.56	0.37
55	7.40	7.09	8.60	8.52	8.19	8.45	8.15	8.93	8.63	8.84	7.53	9.26	1.74	1.36	2.15	1.59	1080.32	46.04	18.91	0.45
56	9.97	8.28	8.84	10.35	9.48	8.83	8.99	9.54	10.19	9.73	10.15	10.17	2.22	1.83	2.89	1.97	1379.45	61.36	28.03	0.63
57	8.55	5.88	14.40	8.26	8.44	13.18	8.42	14.83	9.85	8.45	13.27	12.13	1.04	0.90	0.57	1.10	468.35	10.71	6.43	0.16
58	7.13	7.02	8.47	7.52	7.51	8.91	7.80	6.71	8.18	8.59	7.00	7.88	1.71	1.29	0.83	1.40	676.67	25.93	10.36	0.25
59	7.68	7.57	8.76	10.03	8.80	8.89	8.73	9.41	10.41	10.21	8.05	10.54	1.80	1.94	2.96	1.93	1324.82	58.87	28.51	0.67
60	8.06	7.01	8.99	8.12	7.84	9.54	7.79	8.48	7.94	8.51	6.96	8.28	1.55	1.04	0.90	1.26	691.41	28.66	9.02	0.22
61	11.51	8.26	9.31	10.26	10.26	9.39	9.22	10.30	10.77	11.08	10.75	10.39	2.70	2.24	3.00	1.98	1358.25	61.41	33.33	0.72
62	5.91	6.86	8.08	7.61	7.38	8.19	7.63	6.50	6.88	8.30	5.58	7.36	1.37	1.35	0.83	1.44	575.71	19.66	10.28	0.25
63	7.18	6.86	8.38	8.21	7.70	8.57	7.70	8.15	7.97	8.46	6.67	8.48	1.63	1.31	2.37	1.53	1151.76	48.78	19.68	0.47
64	8.28	6.14	14.51	9.10	8.62	13.06	8.67	14.43	9.88	9.37	13.31	11.10	1.32	0.93	0.56	1.08	527.79	14.40	7.12	0.17
65	8.28	6.08	14.48	9.25	8.60	13.06	8.61	14.46	9.86	9.12	13.28	11.11	0.97	0.91	0.56	1.08	467.84	9.92	6.09	0.15
66	11.15	9.02	9.40	10.54	10.67	9.02	9.34	10.12	10.99	10.27	10.33	10.97	2.77	2.80	3.00	2.09	1354.14	61.30	36.34	0.78
																Max	1423.68	62.03	36.34	0.78
																Min	450.29	9.40	6.01	0.15

Appendix O Pareto Front Solution Parameters Optimised by GA to Maximise SEA, Relative Specific Stiffness and Strength While Minimising Mass

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Specific Energy Absorption (kJ/kg)	Supercell Mass (g)	Relative Density
1	6.55	8.99	7.13	10.22	7.20	4.94	7.45	6.47	6.24	7.63	5.76	6.00	0.87	1.09	0.98	1.54	519.26	13.31	8.64	9.09	0.23
2	8.56	9.76	7.79	11.83	9.27	11.18	12.09	10.87	10.35	11.78	12.01	9.44	2.61	2.30	3.00	2.05	1301.38	59.35	38.30	34.23	0.75
3	4.21	1.19	2.00	13.50	9.08	12.08	7.99	11.46	20.00	19.42	13.55	3.81	3.00	2.99	3.00	0.98	991.96	44.02	43.20	43.37	0.92
4	6.47	8.83	6.17	10.28	7.31	7.26	7.69	7.92	6.80	7.62	7.84	6.21	1.61	1.01	1.65	1.77	847.53	31.44	14.36	14.53	0.35
5	6.64	8.40	4.73	9.37	8.01	7.47	6.01	4.73	7.50	7.49	7.24	7.29	1.38	1.13	1.72	1.95	744.00	26.30	16.21	16.29	0.40
6	6.17	8.46	4.69	9.94	7.89	6.60	6.36	5.72	5.36	8.87	7.26	6.38	1.35	0.96	1.65	2.04	752.03	25.10	14.52	14.49	0.35
7	6.68	8.71	5.91	10.39	7.61	7.41	7.68	7.78	7.05	7.99	7.56	6.81	1.51	1.15	1.65	1.89	837.06	30.66	14.59	14.79	0.36
8	9.30	10.66	9.41	11.77	11.66	10.16	10.71	10.18	10.34	11.34	10.85	10.82	2.90	2.91	3.00	2.99	1347.51	62.00	39.07	37.38	0.80
9	3.29	1.21	2.06	11.75	9.10	12.56	8.48	12.90	19.81	19.50	13.14	5.81	3.00	2.96	3.00	2.38	1013.82	43.25	43.02	43.17	0.91
10	5.21	7.98	0.05	9.70	5.84	1.99	1.22	3.12	4.62	4.20	3.01	3.59	0.87	0.89	1.14	2.37	342.40	9.90	11.23	12.04	0.31
11	6.34	9.17	5.82	10.11	7.29	5.19	7.49	6.63	6.08	6.58	5.76	6.06	1.05	1.23	1.02	1.84	525.40	14.67	10.49	10.79	0.27
12	6.01	8.17	3.32	9.69	6.96	4.79	4.42	6.40	6.57	7.55	6.83	6.78	1.42	1.03	1.60	2.03	640.45	21.49	15.22	15.81	0.38
13	6.72	6.51	4.99	11.54	6.99	5.83	6.66	6.32	14.16	14.85	7.40	6.22	1.39	1.90	2.85	2.11	934.54	35.97	37.32	31.96	0.78
14	6.36	8.71	6.04	10.37	7.03	6.83	7.80	6.71	7.31	8.39	8.20	6.60	1.16	1.16	1.59	1.72	765.81	25.15	12.87	13.89	0.35
15	8.38	9.84	8.00	11.83	9.54	11.18	12.42	11.17	10.43	12.90	12.02	9.45	2.90	2.54	3.00	2.34	1285.41	59.39	40.74	36.81	0.78
16	5.83	2.95	4.15	10.24	8.17	9.73	8.14	9.46	7.22	6.61	6.30	5.25	1.42	1.52	1.29	1.77	566.91	20.27	15.83	16.47	0.40
17	6.56	9.06	6.06	11.18	7.90	7.24	7.01	7.99	7.47	8.95	6.87	6.53	1.77	1.23	1.23	2.03	678.36	26.47	14.90	14.24	0.34
18	6.85	8.46	7.19	10.26	8.86	8.03	8.09	8.01	6.74	8.09	6.97	6.98	1.70	1.46	2.09	2.06	987.40	39.14	19.91	19.88	0.47
19	6.63	8.70	6.85	10.93	7.33	7.76	7.37	6.24	7.37	8.34	7.64	6.35	1.60	1.42	2.02	2.09	892.46	34.66	19.50	19.84	0.47
20	6.85	8.78	6.36	10.03	7.27	7.90	7.49	8.41	6.97	8.07	7.73	6.87	1.62	1.10	1.69	2.26	885.84	33.62	15.07	14.86	0.36
21	6.69	8.42	5.98	10.16	7.52	7.79	7.54	4.18	7.34	8.18	7.62	5.64	1.77	1.08	1.73	2.23	777.21	30.23	17.16	17.48	0.41
22	6.62	8.73	6.10	9.98	7.45	7.98	7.41	4.18	7.09	8.18	7.68	5.46	1.59	1.08	1.67	2.04	742.44	27.57	16.45	16.48	0.39
23	4.44	2.46	5.34	11.11	9.23	5.15	9.45	12.74	15.97	14.62	7.47	8.60	2.40	2.78	2.97	2.48	1093.04	42.19	41.72	39.42	0.88
24	7.46	7.26	5.93	10.61	7.45	7.72	6.62	5.14	7.14	8.24	7.20	7.84	1.63	1.47	1.93	1.97	845.21	33.30	19.05	19.52	0.47
25	7.13	8.62	7.07	10.90	8.98	7.91	6.39	7.48	8.13	7.83	7.99	8.25	1.57	1.72	2.46	2.04	1070.42	43.87	26.17	23.92	0.57
26	6.67	8.77	6.16	10.01	7.30	7.97	7.43	4.92	6.98	8.17	7.70	6.55	1.60	1.09	1.69	2.18	784.70	29.72	15.79	16.14	0.39
27	7.01	8.01	5.78	11.75	8.24	6.25	8.78	9.84	8.82	10.30	7.30	6.91	1.85	1.91	2.68	2.44	1138.59	45.47	30.96	28.07	0.66
28	6.39	8.77	5.22	10.47	6.95	6.66	7.65	8.46	6.51	6.56	6.46	7.32	1.29	1.30	2.03	1.88	856.61	29.92	18.79	18.81	0.46
29	6.63	8.64	5.12	9.61	7.65	7.49	6.91	4.07	7.50	7.97	7.30	6.19	1.58	1.08	1.89	1.98	806.29	30.05	17.76	18.09	0.43
30	3.29	1.21	2.06	11.75	9.10	12.56	8.48	12.90	19.81	19.50	13.14	5.81	3.00	2.96	3.00	2.38	1013.82	43.25	42.99	43.17	0.91
31	8.28	9.63	7.58	10.93	9.13	8.21	7.85	6.98	9.64	7.92	6.45	7.53	1.68	1.46	2.34	2.18	1031.74	42.88	24.61	22.67	0.54
32	6.64	8.40	4.73	9.37	8.01	7.47	6.01	4.73	7.50	7.49	6.74	7.29	1.63	1.13	1.72	1.95	755.12	28.68	16.96	17.07	0.41
33	6.65	7.32	4.95	10.48	7.49	5.86	6.08	5.27	9.47	12.00	7.29	6.43	1.38	1.66	2.60	1.57	980.20	36.20	29.50	26.39	0.64
34	5.21	7.98	0.05	9.70	6.34	1.99	1.22	3.12	4.62	4.20	3.01	3.59	0.87	0.89	1.14	2.37	350.48	9.98	11.53	11.98	0.31
35	6.53	8.99	7.41	10.74	7.14	5.17	7.49	6.46	6.19	7.45	5.71	6.49	0.87	1.06	0.98	1.79	513.50	13.04	8.80	9.17	0.24

No.	Node 15 X	Node 15 Y	Node 15 Z	Node 16 X	Node 16 Y	Node 16 Z	Node 17 X	Node 17 Y	Node 17 Z	Node 18 X	Node 18 Y	Node 18 Z	R1	R2	R3	R4	Relative Specific Stiffness (MPa)	Relative Specific Strength (MPa)	Specific Energy Absorption (kJ/kg)	Supercell Mass (g)	Relative Density
36	7.86	9.75	6.52	11.84	7.81	8.52	9.38	9.55	9.16	7.80	9.37	7.47	1.07	1.94	2.64	1.92	1156.82	46.42	28.56	25.32	0.64
37	5.83	8.94	3.72	10.50	6.32	4.31	4.53	6.45	6.11	5.86	5.80	6.21	1.07	0.98	1.10	2.01	461.22	12.42	10.19	11.28	0.28
38	6.48	8.58	6.18	10.66	7.20	5.64	6.96	5.79	6.30	7.58	7.43	6.05	1.39	1.32	1.45	2.08	672.22	23.09	14.42	14.89	0.36
39	5.73	8.49	4.84	10.65	6.89	4.39	2.59	6.33	5.59	5.46	4.21	4.57	0.87	1.01	1.14	1.89	423.39	10.37	11.20	11.78	0.30
40	5.25	7.98	1.14	9.73	6.38	5.59	1.30	6.26	5.16	5.02	6.06	4.27	0.89	0.90	1.16	1.88	378.62	8.99	10.43	11.58	0.30
41	7.10	8.74	3.59	11.77	9.24	7.54	6.50	5.97	6.15	8.99	4.51	6.47	1.34	1.22	2.69	2.10	999.21	35.46	28.95	25.50	0.63
42	6.71	7.88	7.36	10.92	7.43	6.91	8.45	7.05	8.58	9.34	8.98	7.87	1.85	1.29	1.74	1.81	851.23	34.59	17.96	17.85	0.42
43	6.79	7.68	7.60	11.63	9.75	8.06	7.26	6.08	9.84	11.91	10.29	6.84	1.83	1.69	2.96	2.08	1221.30	52.39	33.81	29.52	0.69
44	6.13	8.95	0.70	11.07	7.26	2.22	6.07	6.70	6.95	5.16	3.76	4.41	1.36	0.90	1.26	2.27	450.80	14.10	14.49	14.18	0.35
45	6.69	7.79	5.24	11.32	8.18	6.22	7.91	7.80	9.24	11.01	7.30	6.71	1.55	1.77	2.61	2.42	1076.26	41.37	29.20	26.34	0.63
46	6.21	7.58	5.49	11.41	7.24	11.09	8.81	8.32	9.48	11.62	10.95	8.77	2.37	2.58	2.95	2.11	1181.37	53.71	39.45	35.97	0.80
47	6.13	8.46	4.87	10.31	7.28	7.13	3.51	5.22	5.95	7.27	6.77	5.65	1.16	1.02	1.24	1.84	541.26	16.30	12.01	12.48	0.31
48	6.16	8.75	7.67	11.51	8.35	11.44	10.78	10.49	10.71	11.67	12.00	7.83	2.60	2.32	3.00	2.10	1234.96	55.96	39.92	35.68	0.78
49	5.22	7.39	0.20	9.81	6.60	2.29	8.25	5.60	6.52	6.65	3.75	5.15	0.94	1.32	1.33	1.94	440.89	11.47	13.29	14.60	0.37
50	6.09	9.00	4.86	10.31	8.16	9.05	7.05	9.15	7.45	7.23	6.73	6.51	1.52	1.35	1.23	1.84	620.20	24.21	14.62	14.43	0.35
51	8.53	9.57	6.10	11.64	7.65	9.55	5.95	10.82	10.23	8.23	11.92	9.02	2.50	2.16	2.98	2.06	1233.53	54.83	39.43	34.15	0.75
52	8.10	7.32	6.56	12.42	8.54	5.04	9.14	10.78	11.69	9.70	11.31	6.83	2.03	2.14	2.93	2.07	1199.91	47.59	37.56	32.87	0.76
53	6.83	7.55	7.25	11.25	9.01	8.30	9.32	10.81	10.05	11.19	10.34	9.46	2.43	2.49	2.92	2.11	1244.63	56.23	38.77	34.33	0.76
54	6.57	8.84	6.26	10.16	7.15	5.69	7.43	6.48	7.02	7.40	6.74	6.59	1.27	1.14	1.06	1.94	599.52	23.79	10.72	10.55	0.26
55	6.27	7.05	7.69	11.18	8.52	6.63	7.70	8.33	8.43	9.11	7.56	7.41	1.26	1.55	2.22	2.31	976.49	37.66	21.89	20.90	0.52
56	6.96	7.81	6.41	10.68	8.43	8.49	7.89	5.71	10.95	9.00	10.50	8.33	1.73	1.18	2.29	2.00	965.27	39.91	24.23	22.65	0.54
57	6.69	8.91	5.89	10.15	7.46	6.09	8.42	6.26	7.04	7.55	6.50	5.13	1.31	1.15	2.18	2.15	945.63	32.60	18.81	18.82	0.46
58	5.21	8.98	1.05	10.70	5.84	2.99	2.22	3.12	4.62	5.20	3.01	3.59	0.87	0.89	1.14	2.37	336.05	9.76	11.35	12.43	0.32
59	7.12	8.51	5.79	9.60	8.79	5.68	7.49	9.23	8.47	8.76	10.69	9.23	1.26	1.42	1.45	2.47	719.00	27.54	14.55	15.01	0.37
60	5.49	2.93	1.57	12.17	9.27	12.01	8.92	11.39	16.44	14.20	11.65	9.15	2.93	2.70	2.99	2.57	1119.40	45.77	42.12	40.49	0.86
61	8.01	8.01	5.78	12.75	8.24	6.25	8.78	9.84	8.82	10.30	6.30	7.91	1.85	1.91	2.68	2.44	1109.31	45.10	31.79	28.57	0.67
62	8.17	10.45	9.12	11.89	10.49	10.72	11.62	10.61	10.35	11.93	12.01	10.29	2.87	2.69	3.00	1.69	1313.82	60.86	39.56	36.72	0.78
63	9.44	8.25	8.69	12.17	9.68	11.04	7.92	11.48	11.81	10.21	12.68	9.44	1.93	2.11	2.97	2.02	1243.20	55.59	34.33	31.60	0.73
																Max	1347.51	62.00	43.20	43.37	0.92
																Min	336.05	8.99	8.64	9.09	0.23

Appendix P Composite Stress Predictor Model Deep Learning Network Architecture Iterations

Dataset	
No. FEA Samples	2000
Total Observations	200000
Sampling Method	Latin Hypercube
Train (%)	70
Validation (%)	15
Test (%)	15

Iterations that improved coefficient of regression between predicted and known Von Mises stress values are highlighted in green.

Final chosen networks are highlighted in yellow.

All models were trained on an identical dataset from 2000 FEA simulations.

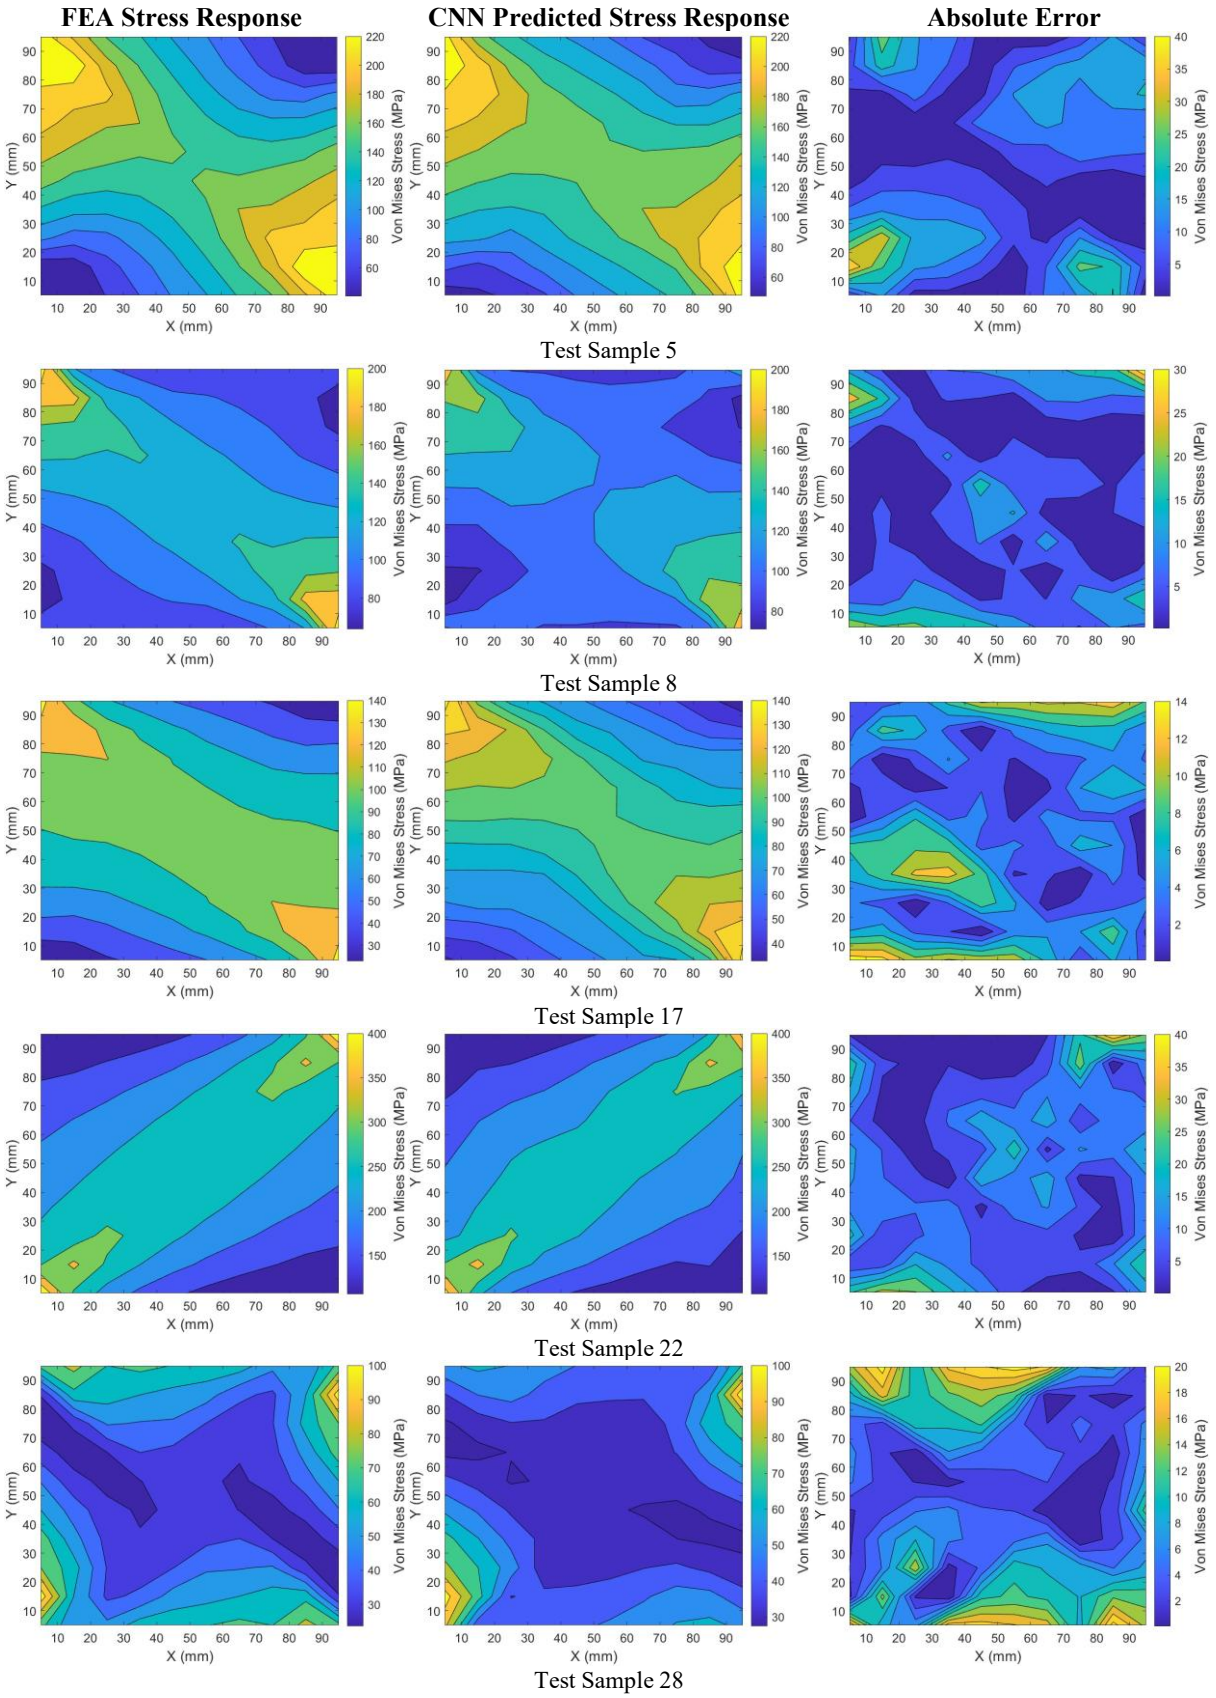
Deep Neural Network (DNN)

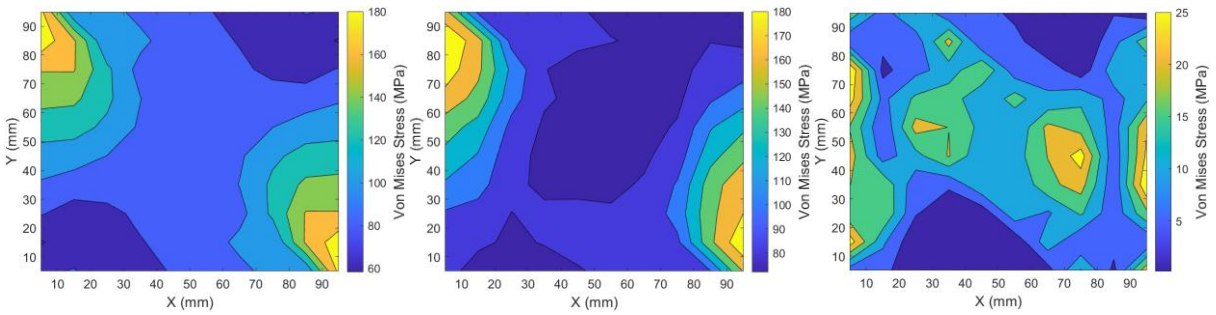
Test Number	Architecture		Options		Training Performance			Regression	
1	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.051	Parameter	R Value
			Mini Batch Size	1000		Validation	0.050		
			No. Epochs	50	Loss	Training	0.001		
	No. Hidden Layer Neurons	[9, 8, 7, 6, 5, 4, 3, 2, 1]	Initial Learning Rate	0.001		Validation	0.001	Von Mises Stress (Pa)	0.9212
	Activation	ReLU	Validation Frequency	140	Training Time	5min 46sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
2	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.041	Parameter	R Value
			Mini Batch Size	1000		Validation	0.043		
			No. Epochs	500	Loss	Training	9E-04		
	No. Hidden Layer Neurons	[9, 8, 7, 6, 5, 4, 3, 2, 1]	Initial Learning Rate	0.001		Validation	9E-04	Von Mises Stress (Pa)	0.9425
	Activation	ReLU	Validation Frequency	140	Training Time	57min 46sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	70000			
3	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.031	Parameter	R Value
			Mini Batch Size	1000		Validation	0.025		
			No. Epochs	50	Loss	Training	5E-04		
	No. Hidden Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 2, 1]	Initial Learning Rate	0.001		Validation	3E-04	Von Mises Stress (Pa)	0.9807
	Activation	ReLU	Validation Frequency	140	Training Time	5min 31sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
4	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.019	Parameter	R Value
			Mini Batch Size	1000		Validation	0.020		
			No. Epochs	1000	Loss	Training	2E-04		
	No. Hidden Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 2, 1]	Initial Learning Rate	0.001		Validation	2E-04	Von Mises Stress (Pa)	0.9886
	Activation	ReLU	Validation Frequency	140	Training Time	1hr 50min 46sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	140000			
5	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.032	Parameter	R Value
			Mini Batch Size	1000		Validation	0.026		
			No. Epochs	50	Loss	Training	5E-04		
	No. Hidden Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 2, 1]	Initial Learning Rate	0.001		Validation	3E-04	Von Mises Stress (Pa)	0.9786
	Activation	ReLU	Validation Frequency	140	Training Time	6min 10sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
6	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.025	Parameter	R Value
			Mini Batch Size	1000		Validation	0.026		
			No. Epochs	50	Loss	Training	3E-04		
	No. Hidden Layer Neurons	[128, 64, 32, 16, 8, 4, 2, 1]	Initial Learning Rate	0.001		Validation	3E-04	Von Mises Stress (Pa)	0.9788
	Activation	ReLU	Validation Frequency	140	Training Time	5min 4sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
7	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.023	Parameter	R Value
			Mini Batch Size	1000		Validation	0.022		
			No. Epochs	50	Loss	Training	3E-04		
	No. Hidden Layer Neurons	[50, 100, 150, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	2E-04	Von Mises Stress (Pa)	0.9843
	Activation	ReLU	Validation Frequency	140	Training Time	5min 2sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
8	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.023	Parameter	R Value
			Mini Batch Size	1000		Validation	0.022		
			No. Epochs	50	Loss	Training	3E-04		
	No. Hidden Layer Neurons	[50, 100, 150, 200, 250, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	3E-04	Von Mises Stress (Pa)	0.9839
	Activation	ReLU	Validation Frequency	140	Training Time	6min 0sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
9	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.026	Parameter	R Value
			Mini Batch Size	1000		Validation	0.026		
			No. Epochs	50	Loss	Training	3E-04		
	No. Hidden Layer Neurons	[20, 40, 80, 160, 320, 160, 80, 40, 20, 1]	Initial Learning Rate	0.001		Validation	3E-04	Von Mises Stress (Pa)	0.9795
	Activation	ReLU	Validation Frequency	140	Training Time	5min 58sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
10	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.018	Parameter	R Value
			Mini Batch Size	1000		Validation	0.019		
			No. Epochs	1000	Loss	Training	2E-04		
	No. Hidden Layer Neurons	[50, 100, 150, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	2E-04	Von Mises Stress (Pa)	0.9895
	Activation	ReLU	Validation Frequency	140	Training Time	1hr 41min 44sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			

Convolutional Neural Network (CNN)

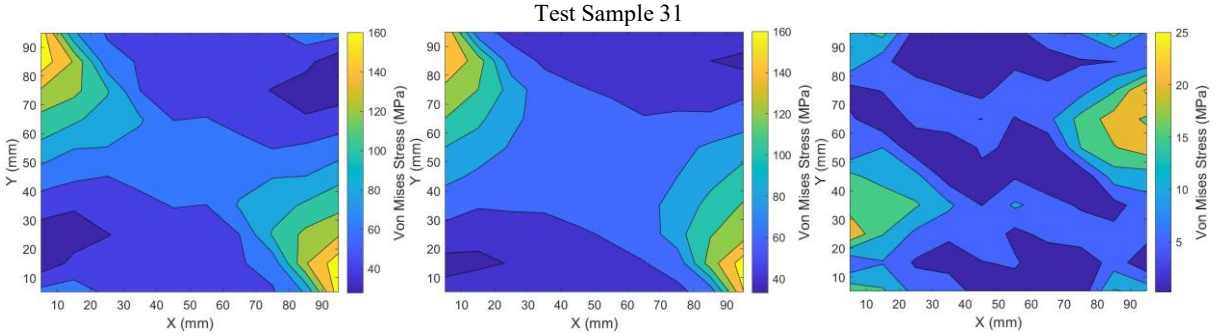
Test Number	Architecture		Options		Training Performance			Regression	
1	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0228	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0230		
	No. Filters	[8, 16, 24, 32]	No. Epochs	50	Loss	Training	0.00026		
	No. Dense Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00026	Von Mises Stress (Pa)	0.9838
	Activation	ReLU	Validation Frequency	140	Training Time	11min 16sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
2	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0418	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0444		
	No. Filters	[2, 4, 6, 8]	No. Epochs	50	Loss	Training	0.00087		
	No. Dense Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00099	Von Mises Stress (Pa)	0.9373
	Activation	ReLU	Validation Frequency	140	Training Time	11min 2sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
3	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0186	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0193		
	No. Filters	[8, 16, 32, 64]	No. Epochs	50	Loss	Training	0.00017		
	No. Dense Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00019	Von Mises Stress (Pa)	0.9889
	Activation	ReLU	Validation Frequency	140	Training Time	10min 49sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
4	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0265	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0239		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00035		
	No. Dense Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00029	Von Mises Stress (Pa)	0.9825
	Activation	ReLU	Validation Frequency	140	Training Time	11min 57sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
5	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0188	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0191		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00018		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00018	Von Mises Stress (Pa)	0.9891
	Activation	ReLU	Validation Frequency	140	Training Time	13min 6sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
6	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0273	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0230		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00037		
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00026	Von Mises Stress (Pa)	0.9834
	Activation	ReLU	Validation Frequency	140	Training Time	14min 3sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
7	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0383	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0363		
	No. Filters	[8, 16, 32, 64]	No. Epochs	50	Loss	Training	0.00073		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00066	Von Mises Stress (Pa)	0.9591
	Activation	ReLU	Validation Frequency	140	Training Time	12min 13sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
8	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0264	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0205		
	No. Filters	[8, 16, 32, 64]	No. Epochs	50	Loss	Training	0.00035		
	No. Dense Layer Neurons	[256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00021	Von Mises Stress (Pa)	0.9875
	Activation	ReLU	Validation Frequency	140	Training Time	11min 8sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
9	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0302	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0281		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00046		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00039	Von Mises Stress (Pa)	0.9751
	Activation	ReLU	Validation Frequency	140	Training Time	13min 25sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
10	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0320	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0331		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	50	Loss	Training	0.00051		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00055	Von Mises Stress (Pa)	0.9658
	Activation	ReLU	Validation Frequency	140	Training Time	14min 4sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
11	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0313	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0288		
	No. Filters	[8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88,	No. Epochs	50	Loss	Training	0.00049		
	No. Dense Layer Neurons	[1024, 512, 256, 128, 64, 32, 16, 8, 1]	Initial Learning Rate	0.001		Validation	0.00042	Von Mises Stress (Pa)	0.9737
	Activation	ReLU	Validation Frequency	140	Training Time	25min 27sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	7000			
12	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0192	Parameter	R Value
	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0141		
	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.00019		
	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00010	Von Mises Stress (Pa)	0.9939
	Activation	ReLU	Validation Frequency	140	Training Time	4hr 25min 29sec			
	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	140000			

Appendix Q Comparison Between FEA and CNN Predicted Composite Stress Contours

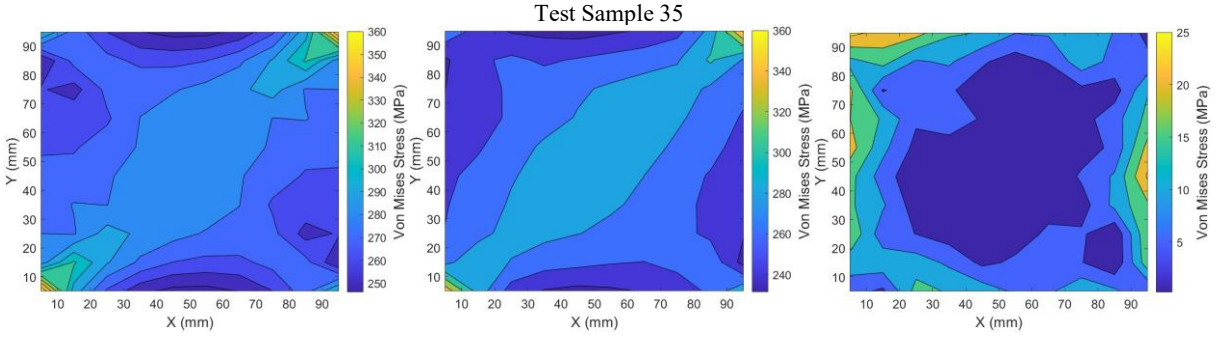




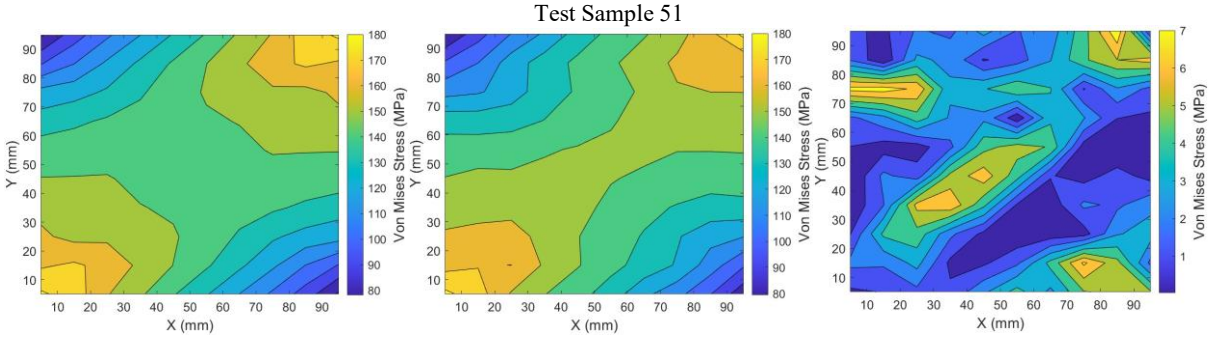
Test Sample 31



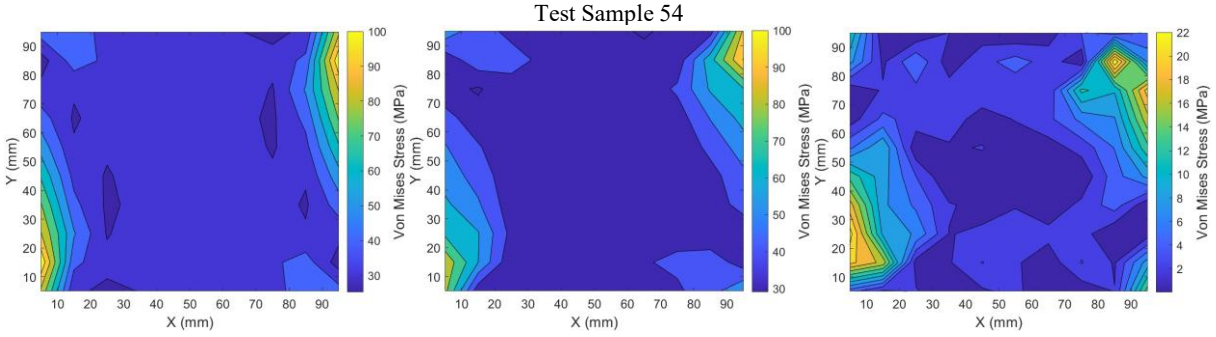
Test Sample 35



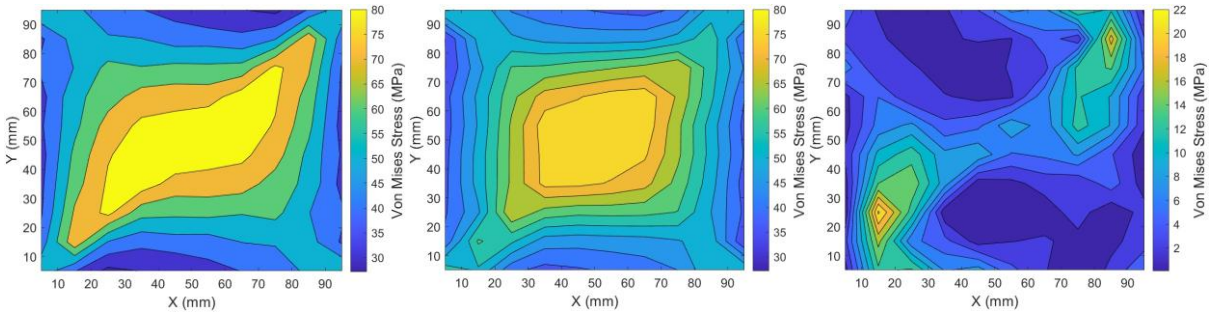
Test Sample 51



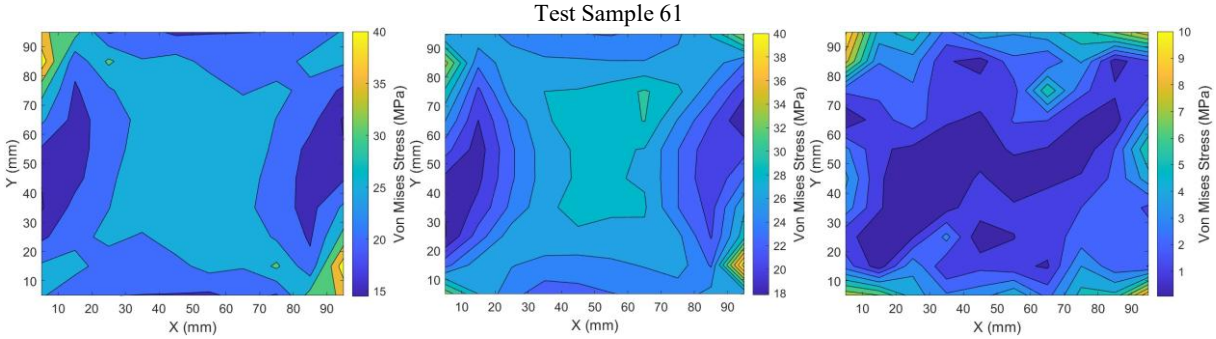
Test Sample 54



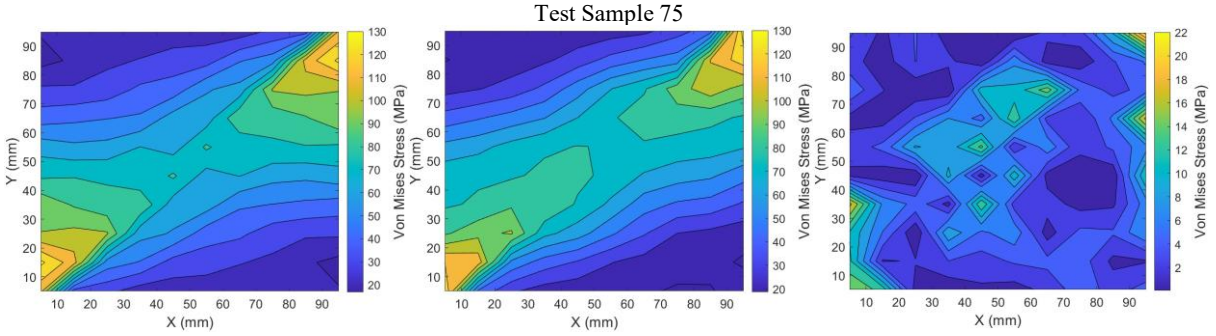
Test Sample 59



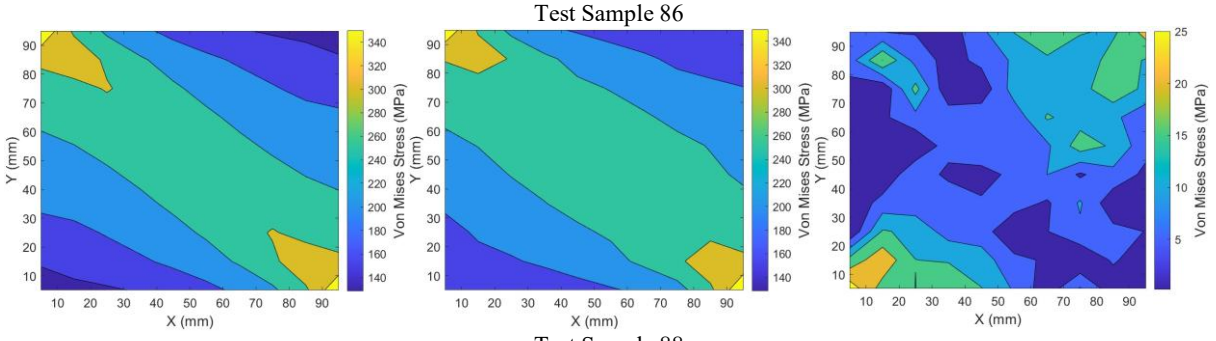
Test Sample 61



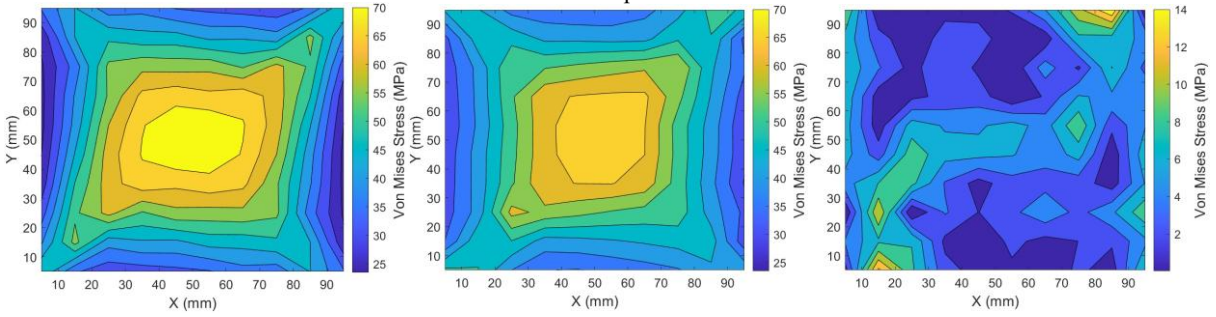
Test Sample 75



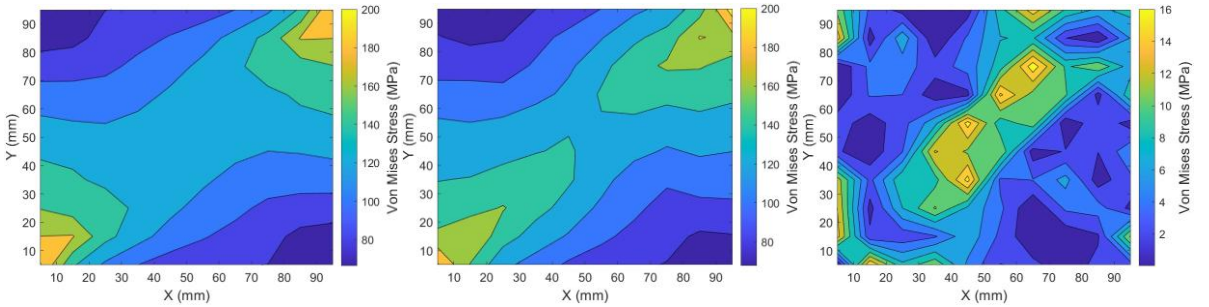
Test Sample 86



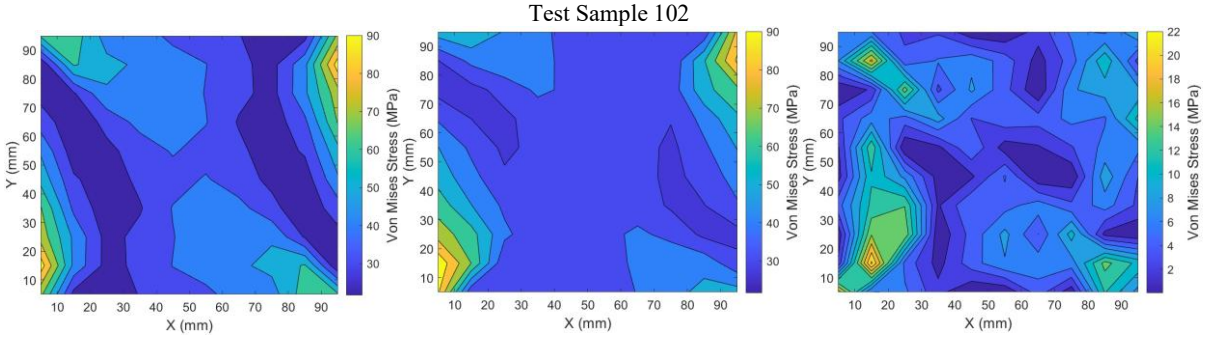
Test Sample 88



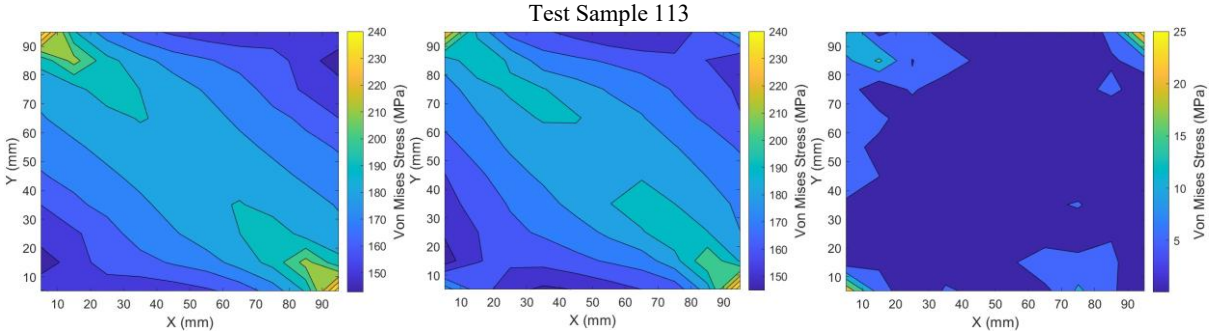
Test Sample 91



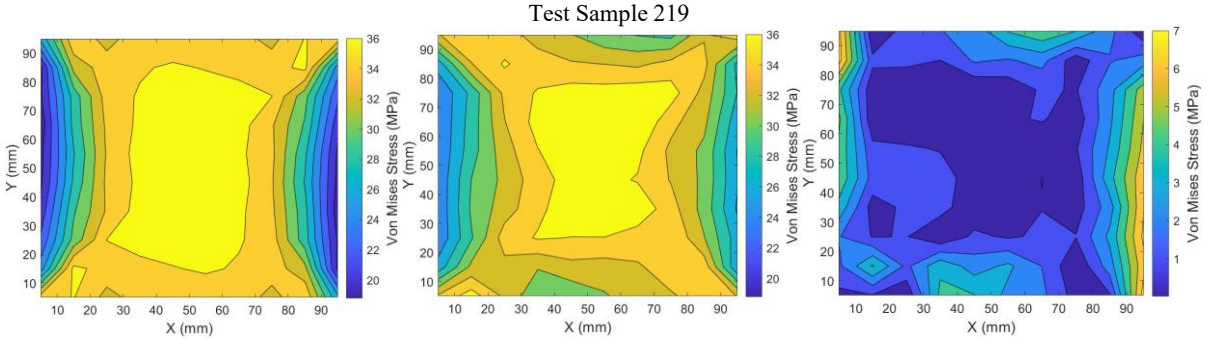
Test Sample 102



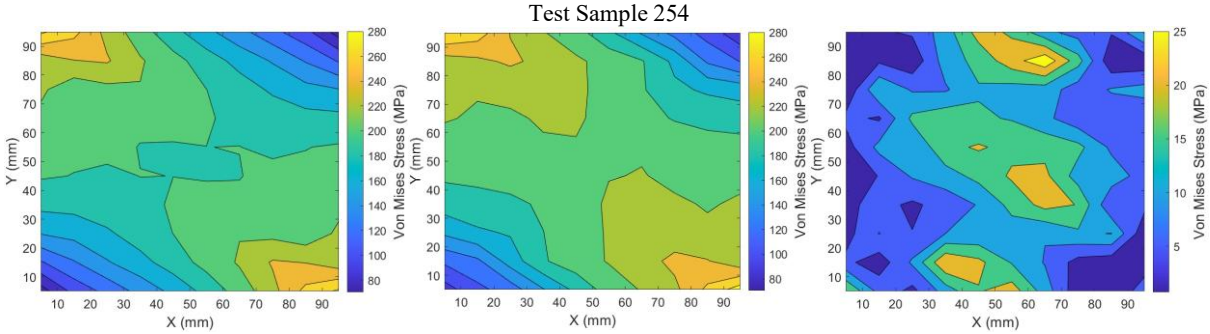
Test Sample 113



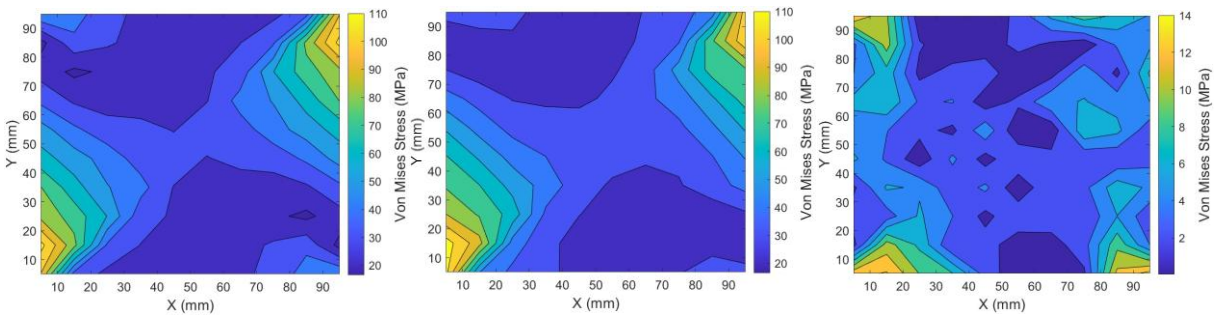
Test Sample 219



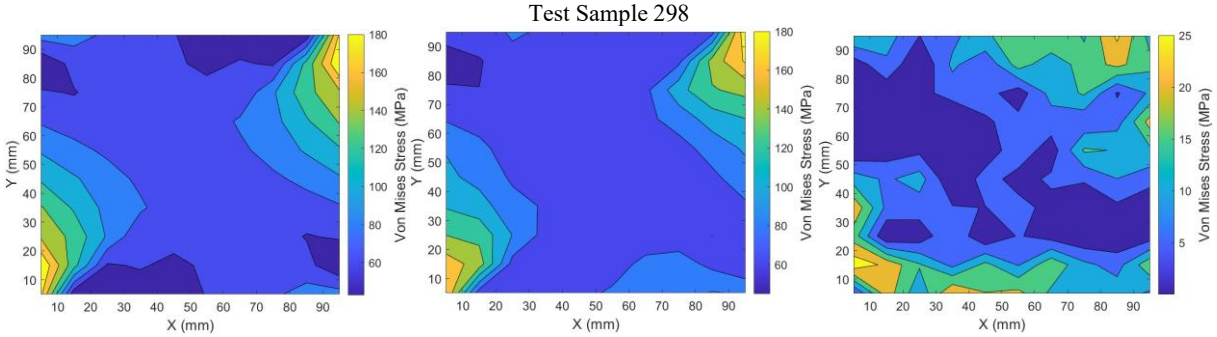
Test Sample 254



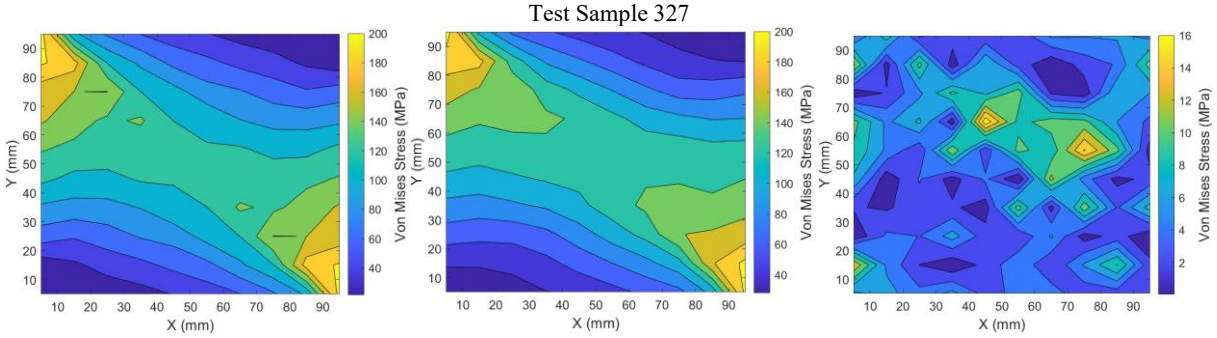
Test Sample 288



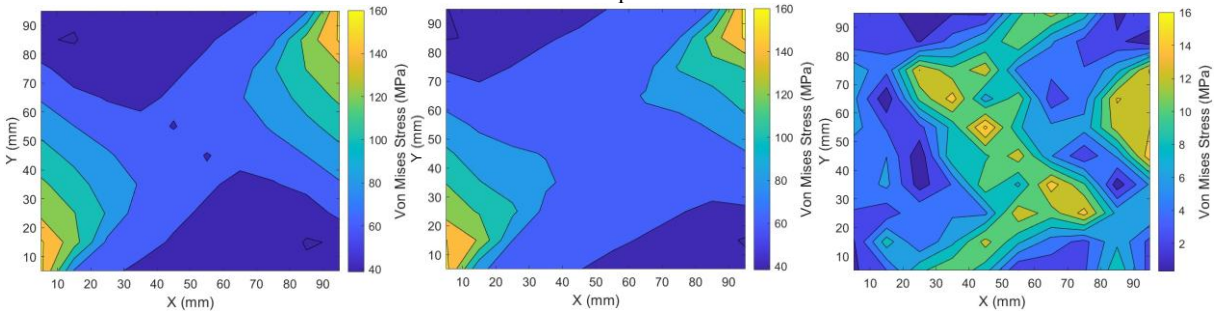
Test Sample 298



Test Sample 327



Test Sample 353



Test Sample 441

Appendix R Composite Maximum Stress Predictor Network Performance Evaluations

Training Performance Summary

Shallow ANN

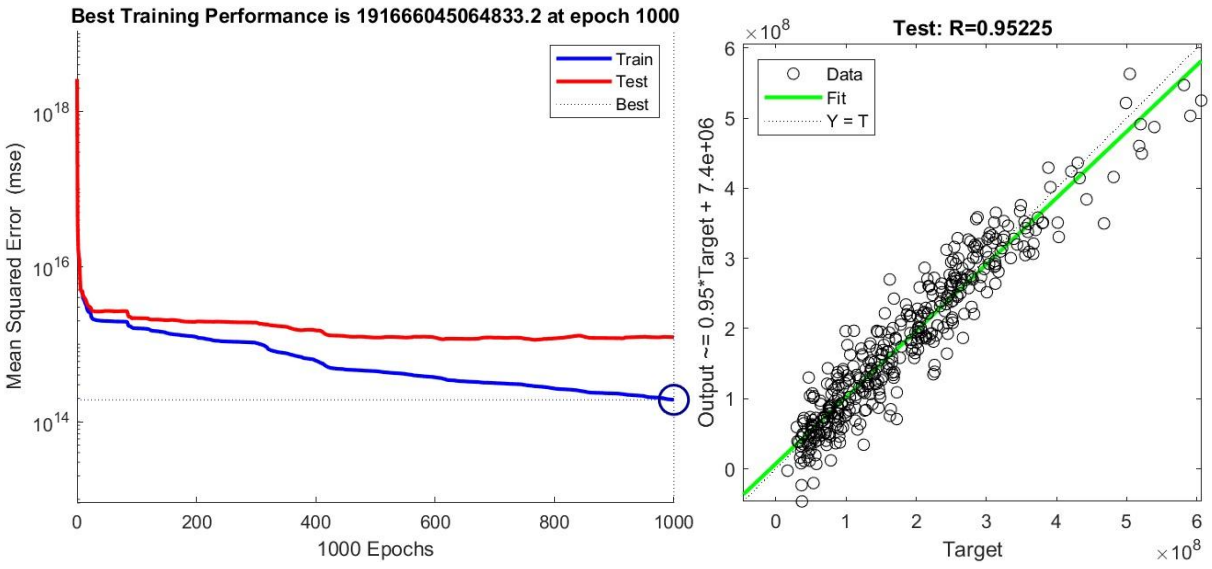
Test Number	Data		Architecture		Training Progress				Training Performance			
ANN-11	No. FEA Samples	2000	Type	ANN	Unit	Initial Value	Stopped Value	Target Value		Observations	MSE	R
	No. FEA Mesh Elements	100			Epoch	0	1000	1000				
	Total Observations	2000			Elapsed Time	-	00:04:16	-				
	Sampling Method	Latin Hypercube			Performance	2.58E+18	1.92E+14	0	Training	1600	1.9167E+14	0.9926
	Train (%)	80	No. Hidden Layer Neurons	150	Gradient	6.63E+18	2.77E+14	1.00E-07				
			Activation	Sigmoid	Mu	0.005	5.00E+00	1.00E+10				
	Test (%)	20	Training Algorithm	Bayesian Regularization	Effective # Param	1.50E+03	7.63E+02	0	Test	400	1.232E+15	0.9522
			Performance Metric	Mean Squared Error	Sum Squared Param	1.43E+03	3.57E+02	0				

Deep DNN & CNN

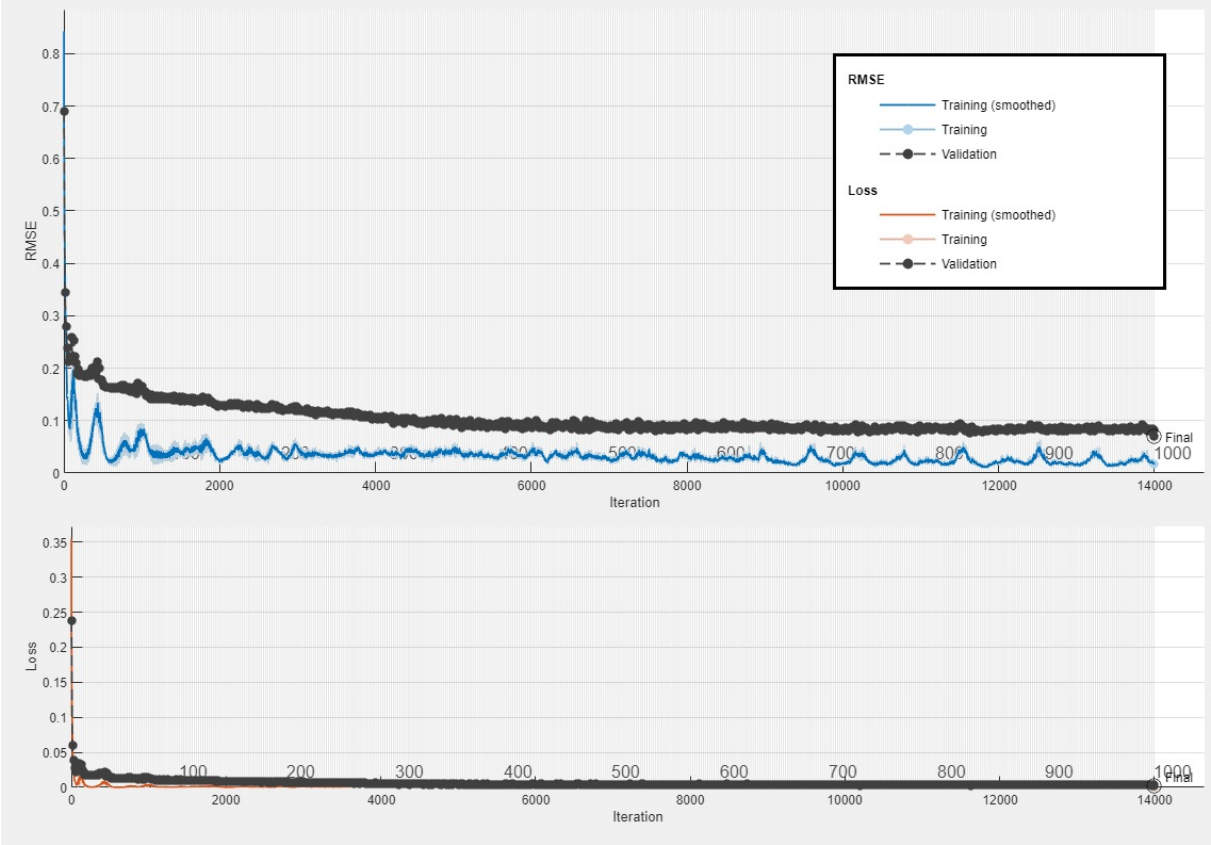
Test Number	Data		Architecture		Training Options		Training Performance			Regression	
DNN-4	No. FEA Samples	2000	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.01747	Parameter	R Value
	Total Observations	2000			Mini Batch Size	100		Validation	0.06912		
	Sampling Method	Latin Hypercube			No. Epochs	1000	Loss	Training	0.00015		
	Train (%)	70	No. Hidden Layer Neurons	[50, 100, 150, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	0.00239	Von Mises Stress (Pa)	0.9179
	Validation (%)	15	Activation	ReLU	Validation Frequenc	14	Training Time	9min 2sec			
	Test (%)	15	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	14000			
DNN-7	No. FEA Samples	200000	Type	DNN	Training Algorithm	Adam	RMSE	Training	0.02493	Parameter	R Value
	Total Observations	200000			Mini Batch Size	1000		Validation	0.01969		
	Sampling Method	Latin Hypercube			No. Epochs	1000	Loss	Training	0.00031		
	Train (%)	70	No. Hidden Layer Neurons	[50, 100, 150, 200, 150, 100, 50, 1]	Initial Learning Rate	0.001		Validation	0.00019	Von Mises Stress (Pa)	0.9907
	Validation (%)	15	Activation	ReLU	Validation Frequenc	140	Training Time	1hr 41min 13sec			
	Test (%)	15	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	140000			
CNN-4	No. FEA Samples	2000	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0233	Parameter	R Value
	Total Observations	2000	Convolution Filter Size	3	Mini Batch Size	100		Validation	0.0564		
	Sampling Method	Latin Hypercube	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.00027		
	Train (%)	70	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00159	Von Mises Stress (Pa)	0.9155
	Validation (%)	15	Activation	ReLU	Validation Frequenc	14	Training Time	23min 39sec			
	Test (%)	15	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	14000			
CNN-7	No. FEA Samples	200000	Type	CNN	Training Algorithm	Adam	RMSE	Training	0.0146	Parameter	R Value
	Total Observations	200000	Convolution Filter Size	3	Mini Batch Size	1000		Validation	0.0137		
	Sampling Method	Latin Hypercube	No. Filters	[8, 16, 32, 64, 128]	No. Epochs	1000	Loss	Training	0.00011		
	Train (%)	70	No. Dense Layer Neurons	[512, 256, 128, 64, 32, 16, 8, 4, 1]	Initial Learning Rate	0.001		Validation	0.00009	Von Mises Stress (Pa)	0.9956
	Validation (%)	15	Activation	ReLU	Validation Frequenc	140	Training Time	4hr 23min 17sec			
	Test (%)	15	Normalization	Yes	Shuffle	Every Epoch	No. Iterations	140000			

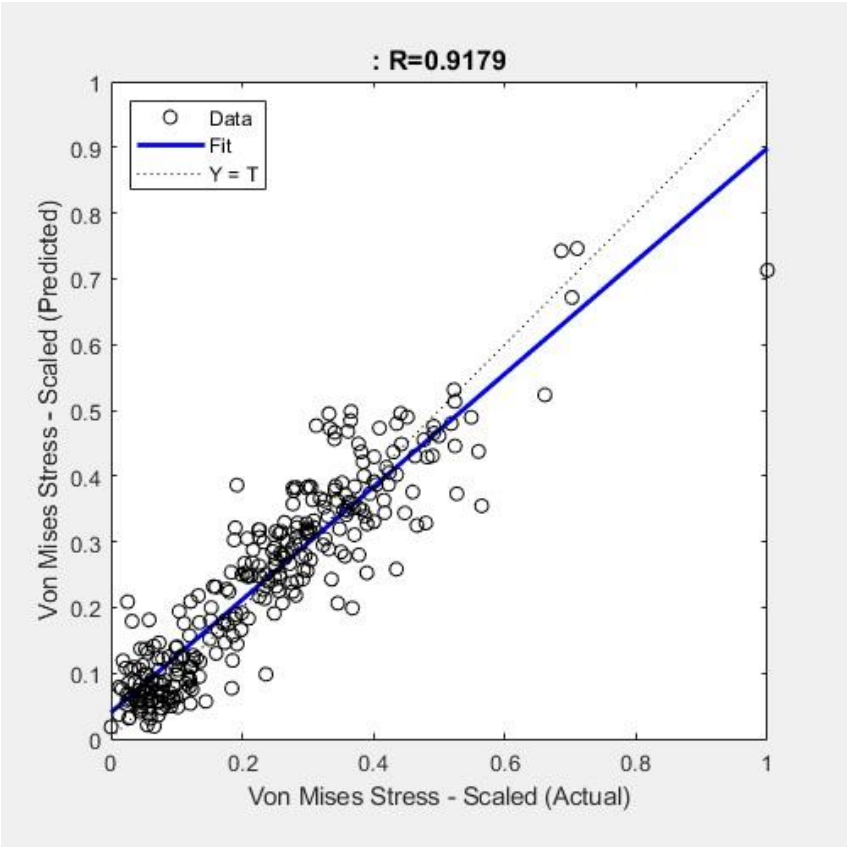
Training History & Regression

ANN-11

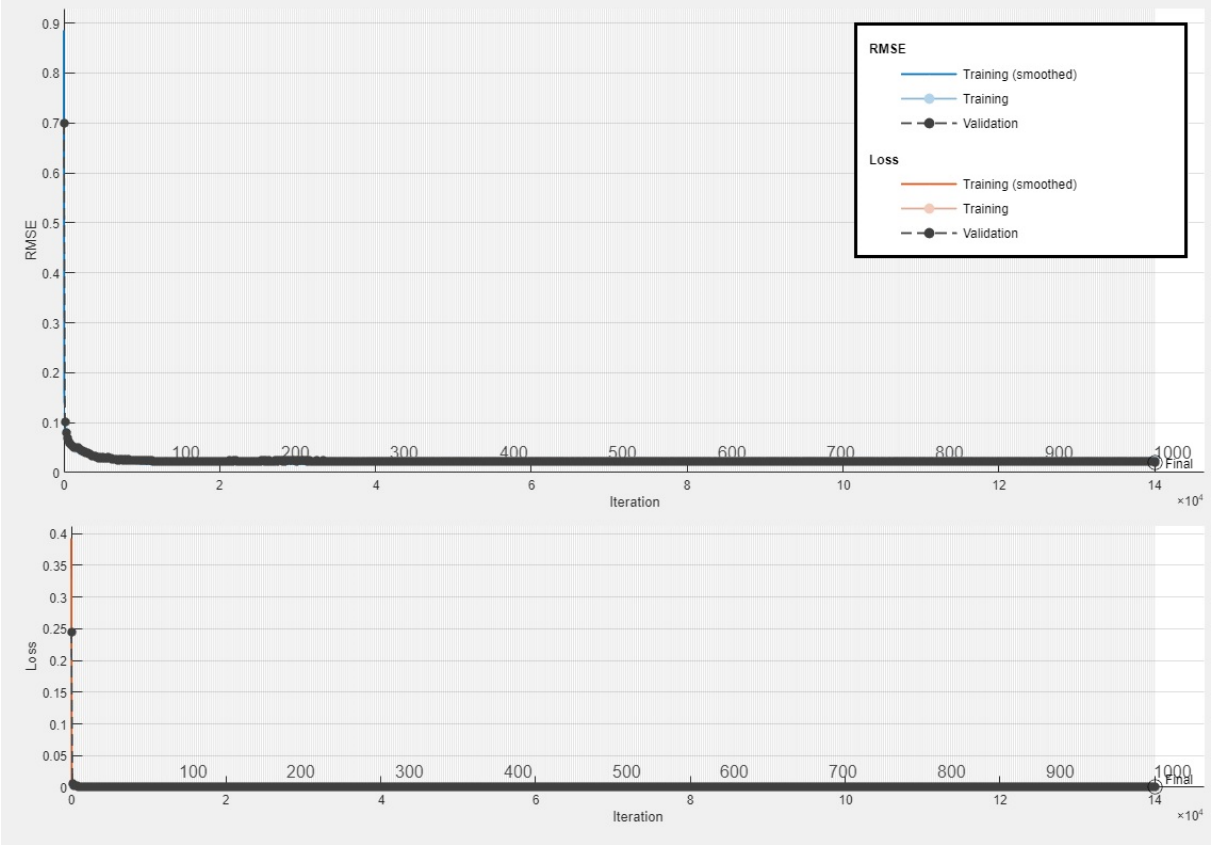


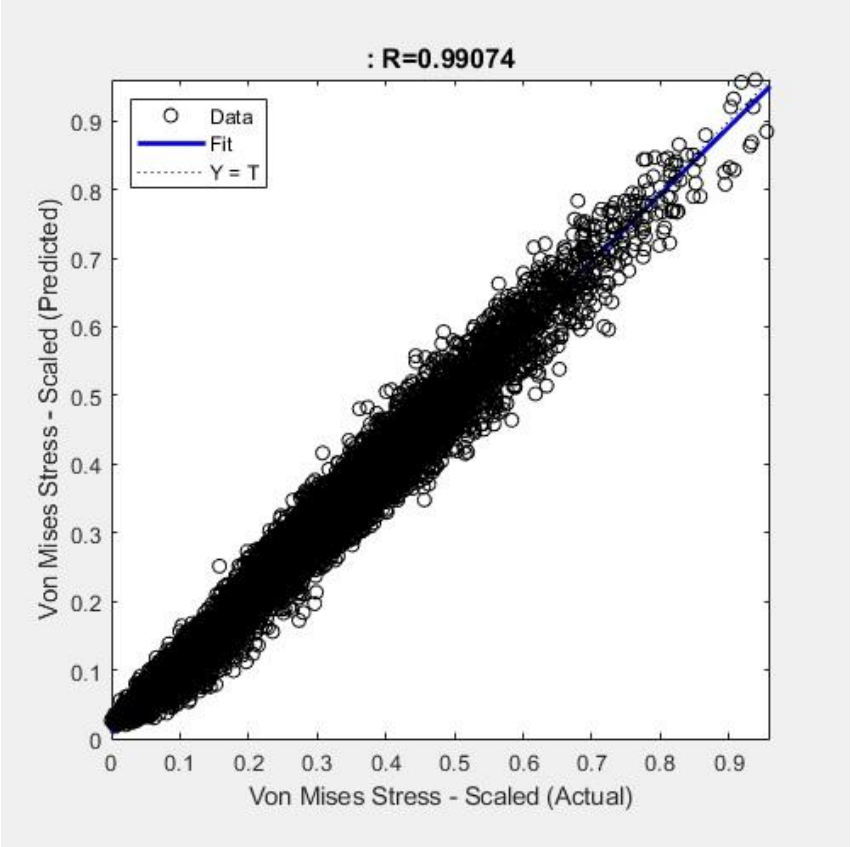
DNN-4


RMSE
Training (smoothed)
Training
Validation
Loss
Training (smoothed)
Training
Validation
RMSE
Iteration
Final
1000
Loss
Iteration
Final

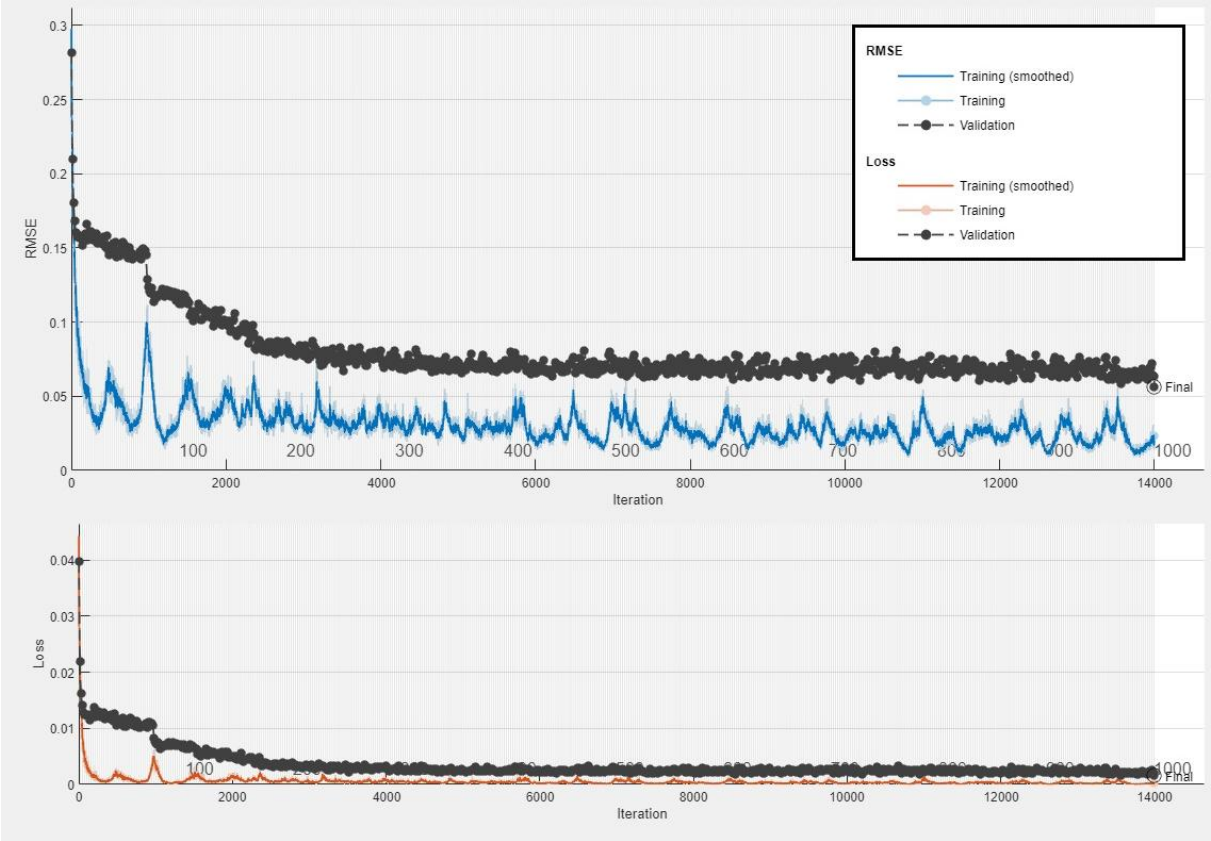

: R=0.9179
Data
Fit
Y = T
Von Mises Stress - Scaled (Predicted)
Von Mises Stress - Scaled (Actual)

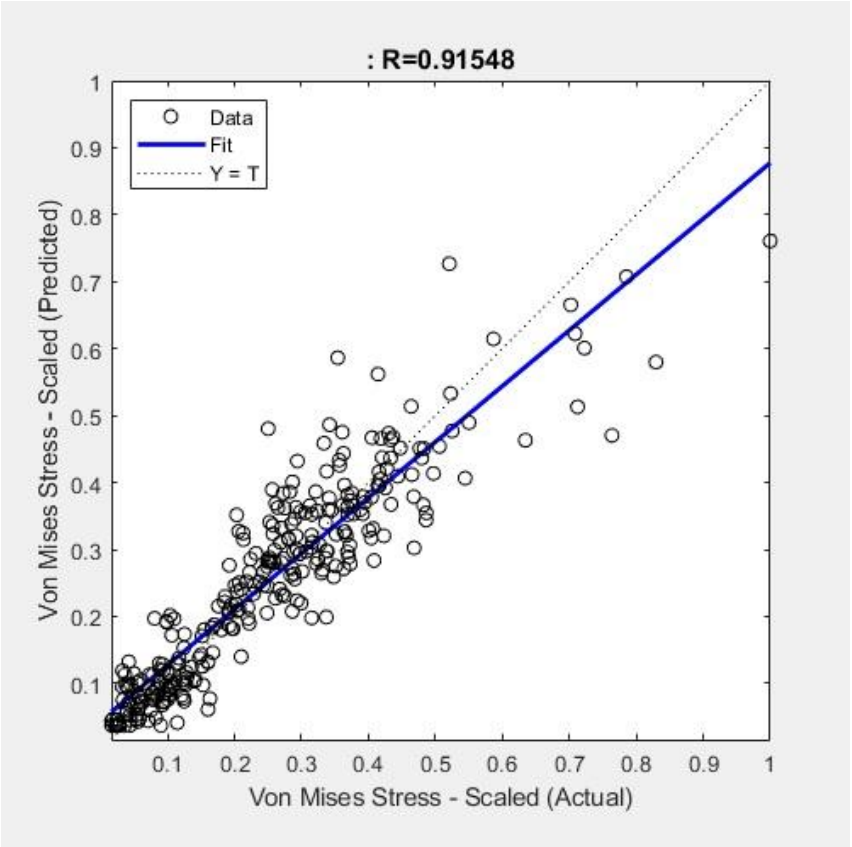
DNN-7


RMSE
Training (smoothed)
Training
Validation
Loss
Training (smoothed)
Training
Validation
RMSE
Iteration
Final
Loss
Iteration
Final

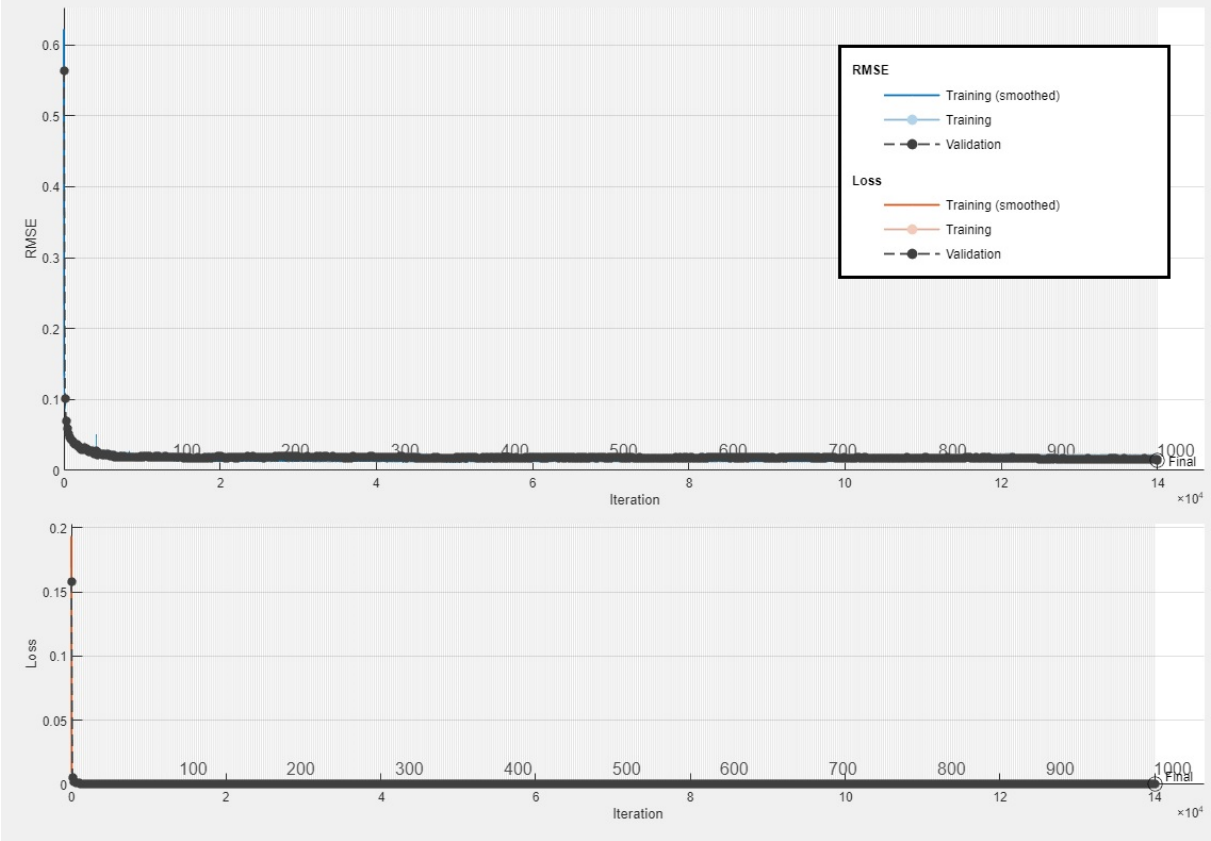

: R=0.99074
Data
Fit
Y = T
Von Mises Stress - Scaled (Predicted)
Von Mises Stress - Scaled (Actual)

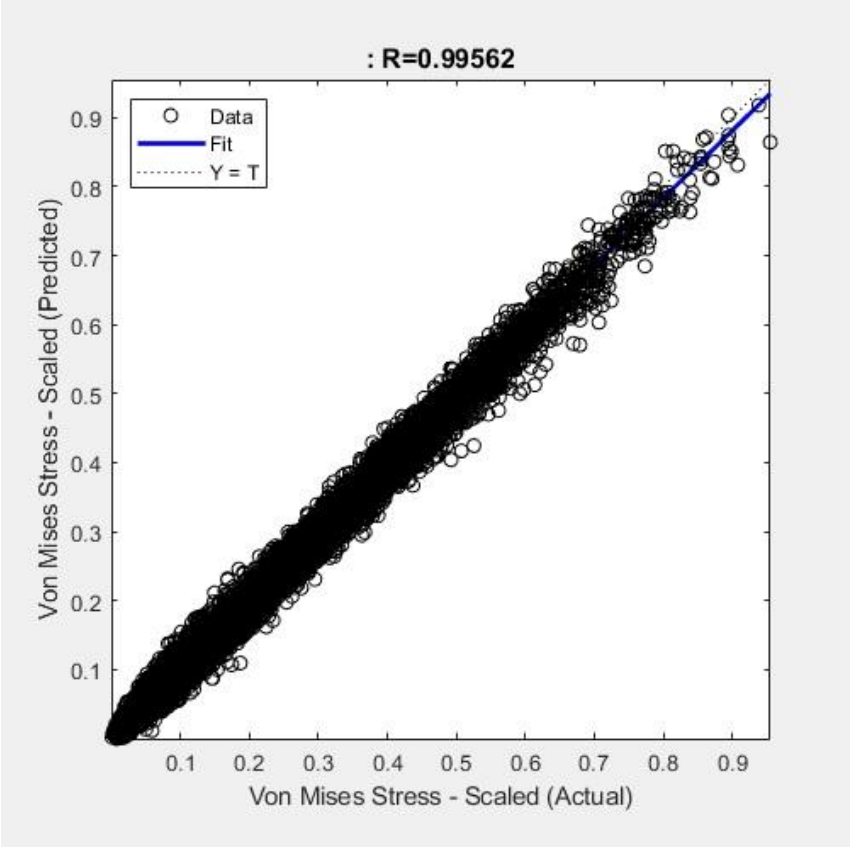
CNN-4


RMSE
Training (smoothed)
Training
Validation
Loss
Training (smoothed)
Training
Validation
RMSE
Iteration
Final
Loss
Iteration
Final


: R=0.91548
Data
Fit
Y = T
Von Mises Stress - Scaled (Predicted)
Von Mises Stress - Scaled (Actual)

CNN-7

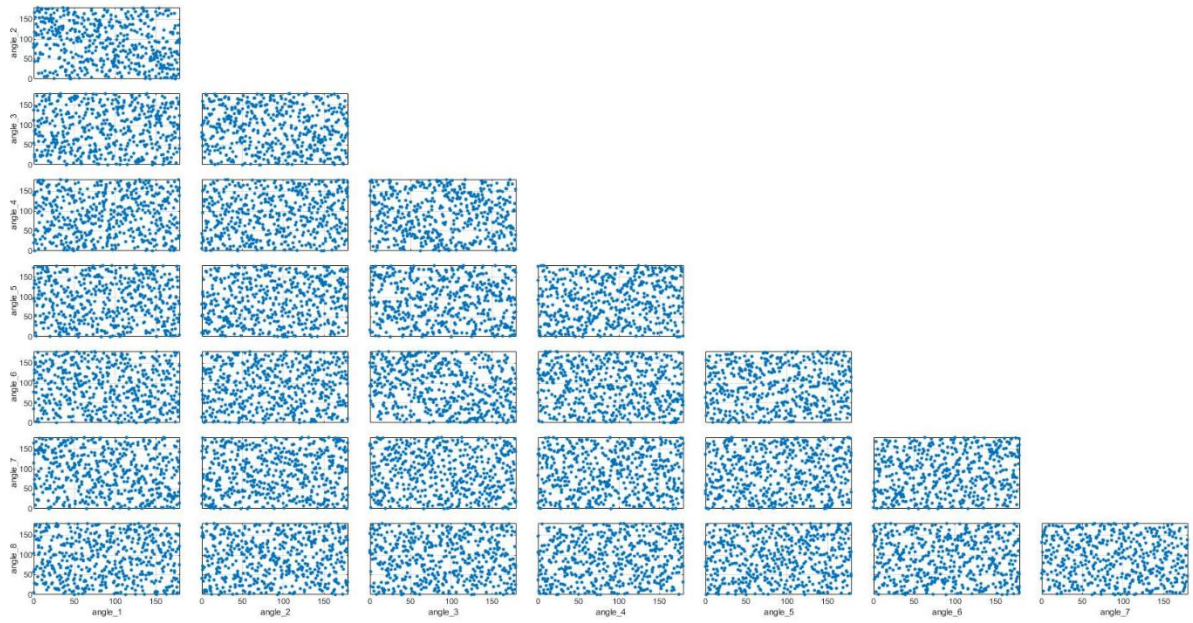

RMSE
Training (smoothed)
Training
Validation
Loss
Training (smoothed)
Training
Validation
RMSE
Iteration
Final
Loss
Iteration
Final


: R=0.99562
Data
Fit
Y = T
Von Mises Stress - Scaled (Predicted)
Von Mises Stress - Scaled (Actual)

Absolute Error

Test Layup Pairwise Design Space – 500 Design Points, Latin hypercube sampling

Absolute error magnitude between predicted (neural network) and actual (FEA) maximum stress calculated at each design point

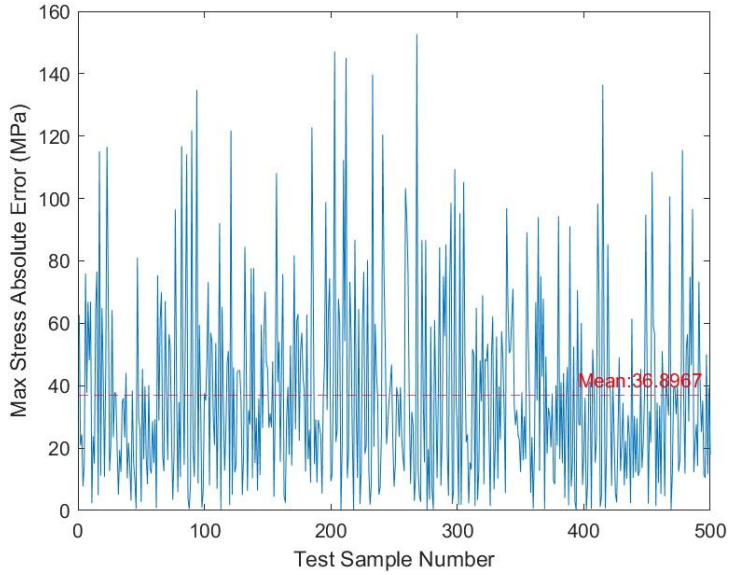


Error Comparison Summary

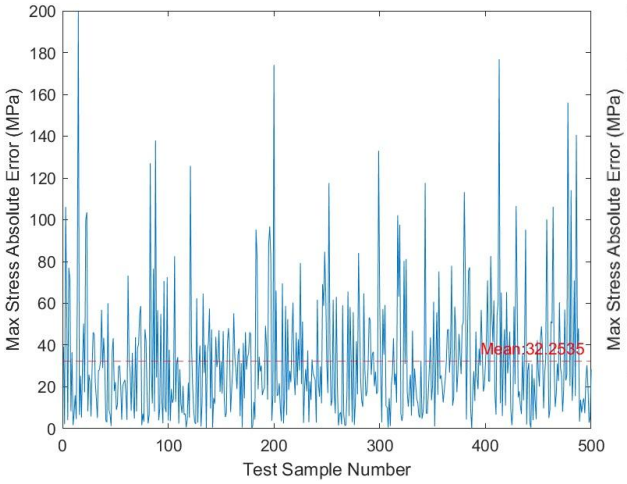
	Max Stress Error (MPa)				
	2000 Sample Training Dataset			200000 Sample Training Dataset	
	ANN - 4min 16sec	DNN - 9min 2sec	CNN - 23min 39sec	DNN - 1hr 41min 13sec	CNN - 4hr 23min 17sec
Mean	36.90	32.25	30.08	11.49	8.52
Standard Deviation	30.10	29.42	26.24	10.32	7.10
Max	152.67	199.53	172.09	76.42	53.00
Min	0.19	0.08	0.10	0.01	3.62E-05
Range	152.49	199.45	171.99	76.41	53.00

Absolute Error Plots

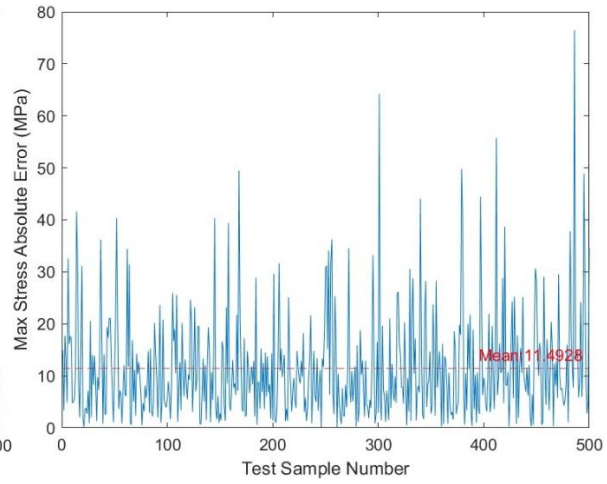
ANN-11, 2000 Sample Training Dataset



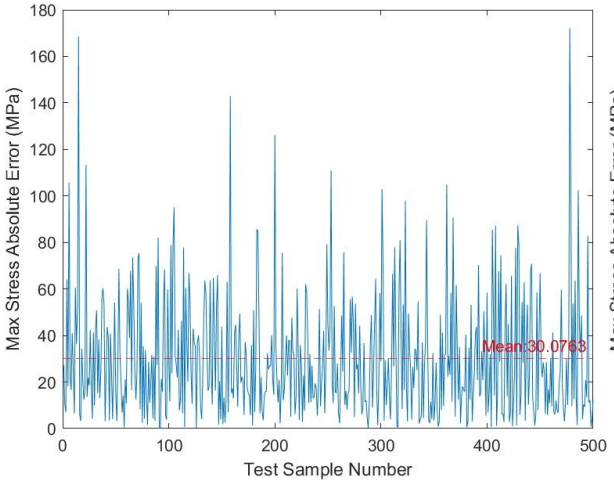
DNN-4, 2000 Sample Training Dataset



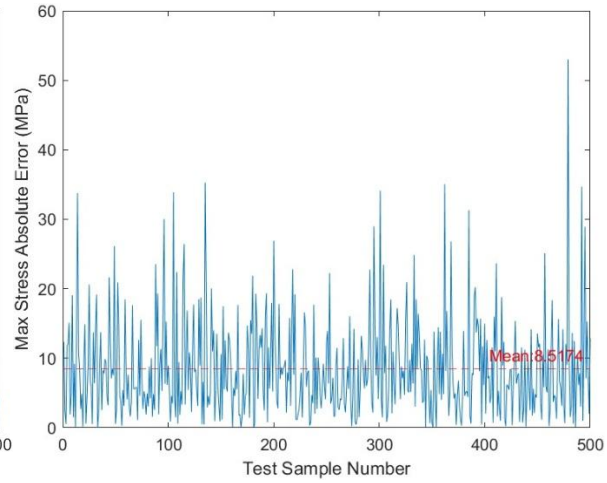
DNN-7, 200000 Sample Training Dataset



CNN-4, 2000 Sample Training Dataset



CNN-7, 200000 Sample Training Dataset



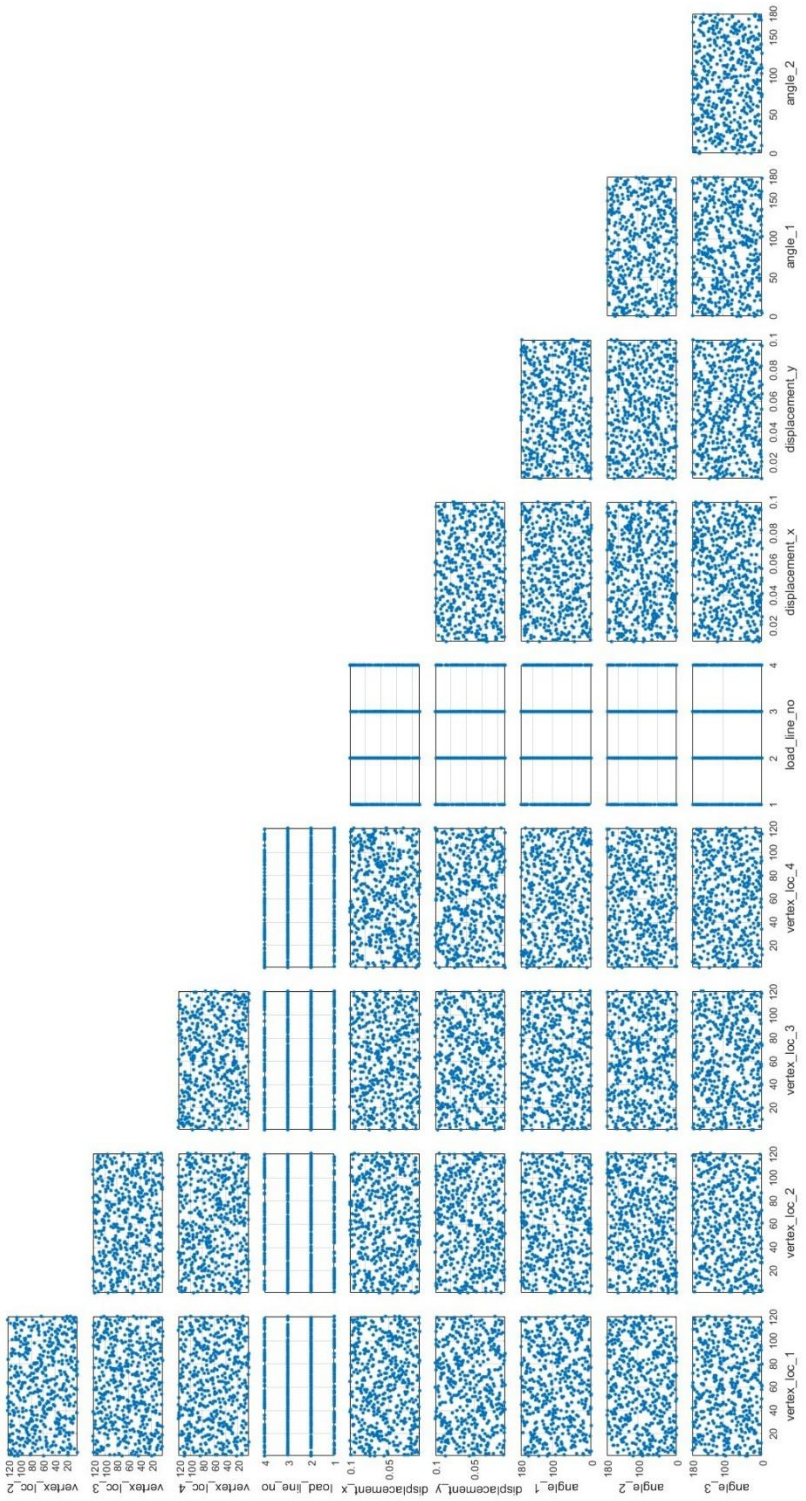
Appendix S Composite Maximum Stress Optimisation Test Solution Layup Stacking Angles

	Genetic Algorithm			Particle Swarm Optimisation			Simulated Annealing		
Test No.	Solution Stacking Sequence (Degrees)		FEA Max Eqv Stress (MPa)	Solution Stacking Sequence (Degrees)		FEA Max Eqv Stress (MPa)	Solution Stacking Sequence (Degrees)		FEA Max Eqv Stress (MPa)
1	Angle 1	175	15.01	Angle 1	74	15.87	Angle 1	0	21.63
	Angle 2	15		Angle 2	180		Angle 2	0	
	Angle 3	179		Angle 3	10		Angle 3	0	
	Angle 4	179		Angle 4	168		Angle 4	2	
	Angle 5	7		Angle 5	82		Angle 5	89	
	Angle 6	92		Angle 6	9		Angle 6	0	
	Angle 7	5		Angle 7	1		Angle 7	0	
	Angle 8	100		Angle 8	101		Angle 8	68	
2	Angle 1	180	13.81	Angle 1	83	12.55	Angle 1	2	30.64
	Angle 2	180		Angle 2	180		Angle 2	0	
	Angle 3	180		Angle 3	6		Angle 3	2	
	Angle 4	98		Angle 4	180		Angle 4	0	
	Angle 5	180		Angle 5	180		Angle 5	1	
	Angle 6	3		Angle 6	93		Angle 6	0	
	Angle 7	103		Angle 7	0		Angle 7	0	
	Angle 8	98		Angle 8	99		Angle 8	62	
3	Angle 1	175	15.01	Angle 1	77	16.74	Angle 1	0	25.01
	Angle 2	15		Angle 2	180		Angle 2	0	
	Angle 3	179		Angle 3	9		Angle 3	2	
	Angle 4	179		Angle 4	180		Angle 4	3	
	Angle 5	7		Angle 5	1		Angle 5	0	
	Angle 6	92		Angle 6	0		Angle 6	0	
	Angle 7	5		Angle 7	72		Angle 7	62	
	Angle 8	100		Angle 8	99		Angle 8	65	
4	Angle 1	180	13.81	Angle 1	175	15.83	Angle 1	2	30.69
	Angle 2	180		Angle 2	10		Angle 2	0	
	Angle 3	180		Angle 3	180		Angle 3	0	
	Angle 4	98		Angle 4	12		Angle 4	0	
	Angle 5	180		Angle 5	180		Angle 5	2	
	Angle 6	3		Angle 6	7		Angle 6	0	
	Angle 7	103		Angle 7	91		Angle 7	0	
	Angle 8	98		Angle 8	98		Angle 8	62	
5	Angle 1	179	15.10	Angle 1	178	14.78	Angle 1	7	17.58
	Angle 2	178		Angle 2	180		Angle 2	66	
	Angle 3	85		Angle 3	179		Angle 3	2	
	Angle 4	0		Angle 4	180		Angle 4	0	
	Angle 5	2		Angle 5	95		Angle 5	0	
	Angle 6	4		Angle 6	98		Angle 6	3	
	Angle 7	72		Angle 7	8		Angle 7	71	
	Angle 8	99		Angle 8	90		Angle 8	103	
6	Angle 1	180	17.40	Angle 1	180	15.82	Angle 1	1	30.65
	Angle 2	12		Angle 2	6		Angle 2	1	
	Angle 3	85		Angle 3	15		Angle 3	1	
	Angle 4	10		Angle 4	170		Angle 4	0	
	Angle 5	167		Angle 5	0		Angle 5	2	
	Angle 6	10		Angle 6	0		Angle 6	0	
	Angle 7	78		Angle 7	82		Angle 7	0	
	Angle 8	99		Angle 8	99		Angle 8	62	
7	Angle 1	180	15.55	Angle 1	174	16.08	Angle 1	19	32.03
	Angle 2	13		Angle 2	12		Angle 2	0	
	Angle 3	85		Angle 3	180		Angle 3	0	
	Angle 4	3		Angle 4	7		Angle 4	2	
	Angle 5	11		Angle 5	81		Angle 5	2	
	Angle 6	168		Angle 6	2		Angle 6	0	
	Angle 7	92		Angle 7	13		Angle 7	0	
	Angle 8	97		Angle 8	101		Angle 8	62	

Appendix T Composite Parametric Stress Predictor Model Error Evaluation Test Dataset

Test dataset features 500 combinations of geometry, loading and layups selected using a Latin hypercube design of experiments test plan.

At each design point the full distribution of stress values in every element within the FEA mesh was compared against equivalent values predicted by the trained CNN.



Appendix U Comparison of Maximum Von Mises Stress Values for Optimised Layups Against Values Obtained Via Exhaustive Search

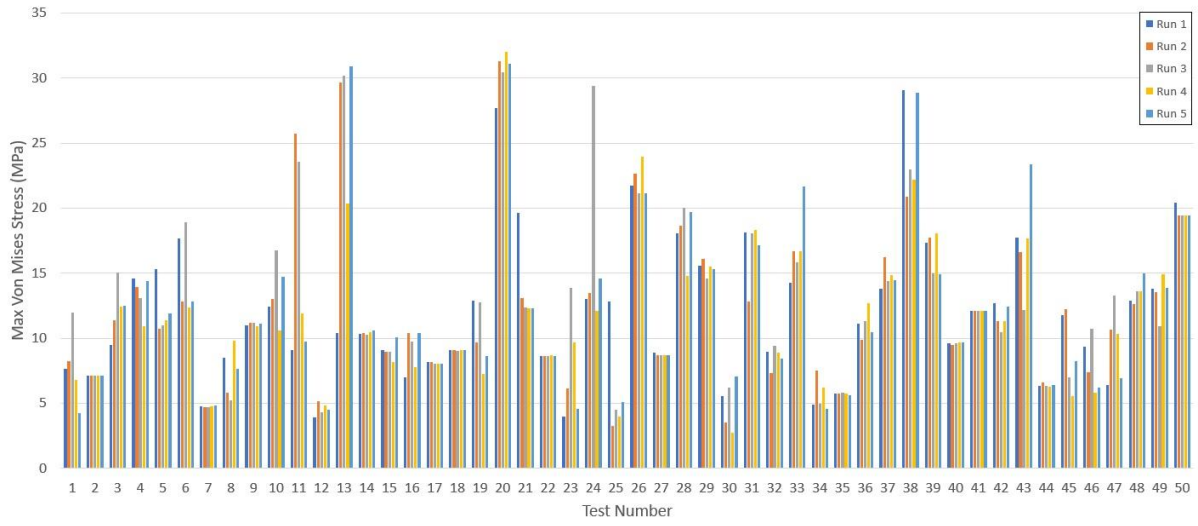
Test No.	Largest Max Stress Exhaustive Search (MPa)	Smallest Max Stress Exhaustive Search (MPa)	Smallest Max Stress GA (MPa)	Smallest Max Stress PSO (MPa)
1	70.89	2.85	4.23	5.08
2	175.88	7.11	7.10	7.04
3	142.81	5.80	9.48	9.34
4	241.75	8.32	10.93	11.33
5	195.79	8.26	10.70	10.70
6	185.88	12.65	12.37	12.37
7	114.16	3.74	4.71	4.44
8	113.47	4.81	5.24	5.02
9	212.15	6.45	10.94	9.39
10	195.57	8.92	10.61	12.17
11	147.28	6.54	9.06	10.77
12	30.43	1.80	3.91	2.64
13	169.26	7.81	10.38	10.33
14	245.10	10.25	10.28	10.38
15	150.15	7.23	8.14	8.38
16	124.77	6.06	6.97	6.99
17	206.48	9.53	8.05	8.22
18	157.61	5.81	9.05	8.92
19	164.67	5.98	7.27	9.63
20	223.66	15.18	27.66	22.96
21	273.14	8.80	12.27	12.27
22	148.94	4.69	8.61	4.95
23	67.42	2.80	4.00	5.81
24	272.50	9.45	12.08	18.48
25	99.58	3.51	3.26	3.91
26	276.92	16.15	21.15	22.65
27	221.58	8.30	8.67	8.60
28	282.00	9.98	14.80	12.74
29	225.61	10.10	14.61	12.01
30	59.36	2.53	2.76	2.52
31	142.88	6.95	12.83	18.20
32	124.85	5.08	7.32	6.09
33	232.24	10.77	14.27	14.99
34	108.18	3.79	4.56	4.82
35	102.86	4.12	5.64	4.10
36	204.74	8.56	9.90	10.16
37	261.49	9.28	13.82	14.20
38	101.72	13.97	20.88	29.05
39	306.79	10.55	14.95	15.80
40	168.00	5.66	9.50	8.32
41	210.70	8.55	12.09	12.12

Test No.	Largest Max Stress Exhaustive Search (MPa)	Smallest Max Stress Exhaustive Search (MPa)	Smallest Max Stress GA (MPa)	Smallest Max Stress PSO (MPa)
42	109.46	4.91	10.47	9.69
43	143.44	8.36	12.13	17.68
44	175.44	5.72	6.25	6.35
45	98.68	3.64	5.56	3.93
46	51.85	2.66	5.82	2.97
47	116.14	4.50	6.41	5.52
48	103.41	6.59	12.64	12.82
49	278.36	12.03	10.93	11.97
50	210.63	8.82	19.41	10.71

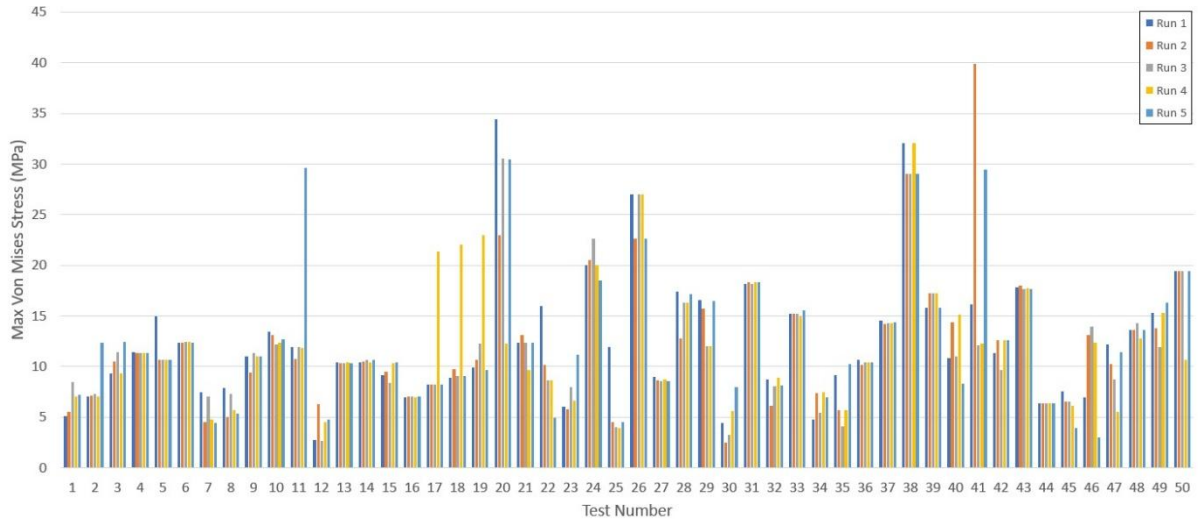
Appendix V Composite Layup Stress Optimisation Consistency

Each optimiser is tested using 50 combinations of geometry and loading conditions. For each case the optimisation is run 5 times and the maximum Von Mises stress from the resulting layups are determined by FEA.

Genetic Algorithm



Particle Swarm Optimisation



List of References

[1] Russell, S. and Norvig, P., 2016. *Artificial Intelligence: A Modern Approach.* Harlow: Harlow: Pearson Education, Limited.

[2] *Energy Use For Transportation* [Online]. Available from: https://www.eia.gov/energyexplained/use-of-energy/transportation.php [Accessed 4th October 2021].

[3] *Lightweight Materials For Cars And Trucks* [Online]. Available from: https://www.energy.gov/eere/vehicles/lightweight-materials-cars-and-trucks [Accessed 4th October 2021].

[4] Lipshaw, J., *What Is Lightweighting And How Does It Improve Fuel Economy In Vehicles?* [Online]. Available from: https://blog.ucsusa.org/science-blogger/lightweighting-and-fuel-economy-in-vehicles/ [Accessed 4th October 2021].

[5] Ward, J., Brooker, A.D. and Wang, L., 2013. Lightweighting Impacts on Fuel Economy, Cost, and Component Losses. SAE International. Available from: https://doi.org/10.4271/2013-01-0381 [Accessed 4th October 2021].

[6] Ashby, M.F., 2017. *Materials selection in mechanical design.* Fifth edition. ed. Amsterdam: Amsterdam : Butterworth-Heinemann.

[7] Ashby, M., 2013. Designing architectured materials. *Scripta Materialia,* 68(1), pp. 4-7.

[8] Fleck, N.A., Deshpande, V.S. and Ashby, M.F., 2010. Micro-architectured materials: past, present and future. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences,* 466(2121), pp. 2495-2516.

[9] Ashby, M.F. and Bréchet, Y.J.M., 2003. Designing hybrid materials. *Acta Materialia,* 51(19), pp. 5801-5821.

[10] Gibson, L.J. and Ashby, M.F., 1997. *Cellular Solids: Structure and Properties.* 2 ed. Cambridge: Cambridge University Press.

[11] Hambli, R. and Hattab, N., 2013. Application of Neural Network and Finite Element Method for Multiscale Prediction of Bone Fatigue Crack Growth in Cancellous Bone.

[12] *Evolution of Birds* [Online]. Available from: https://www2.gwu.edu/~darwin/BiSc151/Birds/Birds.HTML [Accessed 10th August 2022].

[13] Zheng, X., Lee, H., Weisgraber, T.H., Shusteff, M., DeOtte, J., Duoss, E.B., Kuntz, J.D., Biener, M.M., Ge, Q., Jackson, J.A., Kucheyev, S.O., Fang, N.X. and Spadaccini, C.M., 2014. Ultralight, Ultrastiff Mechanical Metamaterials. *Science (American Association for the Advancement of Science),* 344(6190), pp. 1373-1377.

[14] Gibson, L.J., 2003. Cellular Solids. *MRS Bulletin,* 28(4), pp. 270-274.

[15] Deshpande, V.S., Ashby, M.F. and Fleck, N.A., 2001. Foam topology: bending versus stretching dominated architectures. *Acta Materialia,* 49(6), pp. 1035-1040.

[16] Maxwell, J.C., 1864. On the calculation of the equilibrium and stiffness of frames. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science,* 27(182), pp. 294-299.

[17] Calladine, C.R., 1978. Buckminster Fuller's “Tensegrity” structures and Clerk Maxwell's rules for the construction of stiff frames. *International Journal of Solids and Structures,* 14(2), pp. 161-172.

[18] Greer, J.R. and Deshpande, V.S., 2019. Three-dimensional architected materials and structures: Design, fabrication, and mechanical behavior. *MRS Bulletin,* 44(10), pp. 750-757.

[19] Wadley, H.N.G., 2006. Multifunctional periodic cellular metals. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences,* 364(1838), pp. 31-68.

[20] Wicks, N. and Hutchinson, J.W., 2001. Optimal truss plates. *International Journal of Solids and Structures,* 38(30), pp. 5165-5183.

[21] Wadley, H.N.G., Fleck, N.A. and Evans, A.G., 2003. Fabrication and structural performance of periodic cellular metal sandwich structures. *Composites Science and Technology,* 63(16), pp. 2331-2343.

[22] Deshpande, V.S. and Fleck, N.A., 2001. Collapse of truss core sandwich beams in 3-point bending. *International Journal of Solids and Structures,* 38(36), pp. 6275-6305.

[23] Chiras, S., Mumm, D.R., Evans, A.G., Wicks, N., Hutchinson, J.W., Dharmasena, K., Wadley, H.N.G. and Fichter, S., 2002. The structural performance of near-optimized truss core panels. *International Journal of Solids and Structures,* 39(15), pp. 4093-4115.

[24] Wicks, N. and Hutchinson, J.W., 2004. Performance of sandwich plates with truss cores. *Mechanics of Materials,* 36(8), pp. 739-751.

[25] Kooistra, G.W., Deshpande, V.S. and Wadley, H.N.G., 2004. Compressive behavior of age hardenable tetrahedral lattice truss structures made from aluminium. *Acta Materialia,* 52(14), pp. 4229-4237.

[26] Wallach, J.C. and Gibson, L.J., 2001. Mechanical behavior of a three-dimensional truss material. *International Journal of Solids and Structures,* 38(40), pp. 7181-7196.

[27] Zok, F.W., Waltner, S.A., Wei, Z., Rathbun, H.J., McMeeking, R.M. and Evans, A.G., 2004. A protocol for characterizing the structural performance of metallic sandwich panels: application to pyramidal truss cores. *International Journal of Solids and Structures,* 41(22), pp. 6249-6271.

[28] Hyun, S., Karlsson, A.M., Torquato, S. and Evans, A.G., 2003. Simulated properties of Kagomé and tetragonal truss core panels. *International Journal of Solids and Structures,* 40(25), pp. 6989-6998.

[29] Hyun, S. and Torquato, S., 2002. Optimal and Manufacturable Two-Dimensional, Kagome-Like Cellular Solids. *Journal of Materials Research - J MATER RES,* 17, pp. 137-144.

[30] Wang, J., Evans, A.G., Dharmasena, K. and Wadley, H.N.G., 2003. On the Performance of Truss Panels With Kagome Cores. *International Journal of Solids and Structures,* 40, pp. 6981-6988.

[31] Deshpande, V.S., Fleck, N.A. and Ashby, M.F., 2001. Effective properties of the octet-truss lattice material. *Journal of the Mechanics and Physics of Solids,* 49(8), pp. 1747-1769.

[32] 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), pp. 127-140.

[33] Hashin, Z., 1965. On elastic behaviour of fibre reinforced materials of arbitrary transverse phase geometry. *Journal of the Mechanics and Physics of Solids,* 13(3), pp. 119-134.

[34] Zhai, G. and Zhang, J., 2021. Scalable fiber composite lattice structures via continuous spatial weaving. *Composite Structures,* 262, p. 113651.

[35] Romijn, N.E.R. and Fleck, N.A., 2007. The fracture toughness of planar lattices: Imperfection sensitivity. *Journal of the Mechanics and Physics of Solids,* 55(12), pp. 2538-2564.

[36] Fleck, N.A. and Qiu, X., 2007. The damage tolerance of elastic–brittle, two-dimensional isotropic lattices. *Journal of the Mechanics and Physics of Solids,* 55(3), pp. 562-588.

[37] Nguyen, J., Park, S.-i., Rosen, D., Folgar, L. and Williams, J., 2012. Conformal lattice structure design and fabrication. *23rd Annual International Solid Freeform Fabrication Symposium - An Additive Manufacturing Conference, SFF 2012*, pp. 138-161.

[38] Wang, H., Chen, Y. and Rosen, D., 2005. *A Hybrid Geometric Modeling Method for Large Scale Conformal Cellular Structures.*

[39] Gandy, P.J.F., Bardhan, S., Mackay, A.L. and Klinowski, J., 2001. Nodal surface approximations to the P,G,D and I-WP triply periodic minimal surfaces. *Chemical Physics Letters,* 336(3), pp. 187-195.

[40] Wang, Y., 2007. Periodic surface modeling for computer aided nano design. *Computer-Aided Design,* 39(3), pp. 179-189.

[41] Al-Ketan, O., Rowshan, R. and Abu Al-Rub, R.K., 2018. Topology-mechanical property relationship of 3D printed strut, skeletal, and sheet based periodic metallic cellular materials. *Additive Manufacturing,* 19, pp. 167-183.

[42] Abueidda, D.W., Bakir, M., Abu Al-Rub, R.K., Bergström, J.S., Sobh, N.A. and Jasiuk, I., 2017. Mechanical properties of 3D printed polymeric cellular materials with triply periodic minimal surface architectures. *Materials & Design,* 122, pp. 255-267.

[43] Novak, N., Al-Ketan, O., Krstulović-Opara, L., Rowshan, R., Abu Al-Rub, R.K., Vesenjak, M. and Ren, Z., 2021. Quasi-static and dynamic compressive behaviour of sheet TPMS cellular structures. *Composite Structures,* 266, p. 113801.

[44] Yang, L., Yan, C., Han, C., Chen, P., Yang, S. and Shi, Y., 2018. Mechanical response of a triply periodic minimal surface cellular structures manufactured by selective laser melting. *International Journal of Mechanical Sciences,* 148, pp. 149-157.

[45] Fan, X., Tang, Q., Feng, Q., Ma, S., Song, J., Jin, M., Guo, F. and Jin, P., 2021. Design, mechanical properties and energy absorption capability of graded-thickness triply periodic minimal surface structures fabricated by selective laser melting. *International Journal of Mechanical Sciences,* 204, p. 106586.

[46] Roberts, A.P. and Garboczi, E.J., 2002. Elastic properties of model random three-dimensional open-cell solids. *Journal of the Mechanics and Physics of Solids,* 50(1), pp. 33-55.

[47] Van Der Burg, M.W.D., Shulmeister, V., Van Der Geissen, E. and Marissen, R., 1997. On the Linear Elastic Properties of Regular and Random Open-Cell Foam Models. *Journal of Cellular Plastics,* 33(1), pp. 31-54.

[48] Sigmund, O., 1994. *Design of Material Structures Using Topology Optimization*.

[49] Tang, Y., Dong, G., Zhou, Q. and Zhao, Y.F., 2018. Lattice Structure Design and Optimization With Additive Manufacturing Constraints. *IEEE Transactions on Automation Science and Engineering,* 15(4), pp. 1546-1562.

[50] 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), pp. 197-224.

[51] Sigmund, O., 1994. Materials with prescribed constitutive parameters: An inverse homogenization problem. *International Journal of Solids and Structures,* 31(17), pp. 2313-2329.

[52] Messner, M.C., 2016. Optimal lattice-structured materials. *Journal of the Mechanics and Physics of Solids,* 96, pp. 162-183.

[53] Bendsoe, M.P., 2004. *Topology Optimization : Theory, Methods, and Applications.* 2nd 2004. ed. Berlin, Heidelberg: Berlin, Heidelberg : Springer Berlin Heidelberg : Imprint: Springer.

[54] Osher, S. and Sethian, J.A., 1988. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics,* 79(1), pp. 12-49.

[55] Svanberg, K., 1987. The method of moving asymptotes—a new method for structural optimization. *International Journal for Numerical Methods in Engineering,* 24(2), pp. 359-373.

[56] Xie, Y.M. and Steven, G.P., 1993. A simple evolutionary procedure for structural optimization. *Computers & Structures,* 49(5), pp. 885-896.

[57] Querin, O.M., Steven, G.P. and Xie, Y.M., 1998. Evolutionary structural optimisation (ESO) using a bidirectional algorithm. *Engineering Computations,* 15(8), pp. 1031-1048.

[58] Mahdavi, S. and Hanna, S., 2004. *An evolutionary approach to microstructure optimisation of stereolithographic models.*

[59] Graf, G.C., Chu, J., Engelbrecht, S. and Rosen, D.W., 2009. Synthesis Methods for Lightweight Lattice Structures. Volume 2: 29th Computers and Information in Engineering Conference, Parts A and B, pp. 579-589. Available from: https://doi.org/10.1115/DETC2009-86993 [Accessed 8/24/2022].

[60] Kim, P., 2017. *MATLAB Deep Learning : With Machine Learning, Neural Networks and Artificial Intelligence.* 1st 2017. ed. Berkeley, CA: Berkeley, CA : Apress : Imprint: Apress.

[61] Kawato, M., 2009. Reinforcement Models. In: L.R. Squire, ed. *Encyclopedia of Neuroscience.* Oxford: Academic Press, pp. 89-97.

[62] Habibi Aghdam, H., 2017. *Guide to Convolutional Neural Networks : A Practical Application to Traffic-Sign Detection and Classification.* 1st 2017. ed. Cham: Cham : Springer International Publishing : Imprint: Springer.

[63] Duchi, J., Hazan, E. and Singer, Y., 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research,* 12, pp. 2121-2159.

[64] Kingma, D. and Ba, J., 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*.

[65] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. and Salakhutdinov, R., 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research,* 15, pp. 1929-1958.

[66] Yang, C., Kim, Y., Ryu, S. and Gu, G.X., 2020. Prediction of composite microstructure stress-strain curves using convolutional neural networks. *Materials & Design,* 189, p. 108509.

[67] Liu, Z., Wu, C.T. and Koishi, M., 2019. A deep material network for multiscale topology learning and accelerated nonlinear modeling of heterogeneous materials. *Computer Methods in Applied Mechanics and Engineering,* 345, pp. 1138-1168.

[68] Cahn, J.W. and Hilliard, J.E., 1958. Free Energy of a Nonuniform System. I. Interfacial Free Energy. *The Journal of Chemical Physics,* 28(2), pp. 258-267.

[69] Gu, G.X., Chen, C.-T., Richmond, D.J. and Buehler, M.J., 2018. Bioinspired hierarchical composite design using machine learning: simulation, additive manufacturing, and experiment. *Materials Horizons,* 5(5), pp. 939-945.

[70] Frazer, J., 2002. Chapter 9 - Creative Design and the Generative Evolutionary Paradigm. In: P.J. Bentley and D.W. Corne, eds. *Creative Evolutionary Systems.* San Francisco: Morgan Kaufmann, pp. 253-274.

[71] Yan, W.Q., 2021. *Computational methods for deep learning : theoretic, practice and applications.* 1st 2021. ed. Cham, Switzerland: Cham, Switzerland : Springer.

[72] Krish, S., 2011. A practical generative design method. *Computer-Aided Design,* 43(1), pp. 88-100.

[73] Keane, P., 2018. *Generative Design Comes To Fusion 360* [Online]. Available from: https://www.engineering.com/story/generative-design-comes-to-fusion-360 [Accessed 3rd September 2022].

[74] Wong, K., 2018. *Generating Collaboration at GM and Adidas with Generative Design* [Online]. Available from: https://www.digitalengineering247.com/article/generating-collaboration/gm_innovation_seat_bracket_concept_part/ [Accessed 4th September 2022].

[75] Briard, T., Segonds, F. and Zamariola, N., 2020. G-DfAM: a methodological proposal of generative design for additive manufacturing in the automotive industry. *International Journal on Interactive Design and Manufacturing (IJIDeM),* 14.

[76] *Using Generative Design To Create One Of The Lightest Wheels In The World* [Online]. Available from: https://www.autodesk.com/campaigns/generative-design/bac-mono [Accessed 4th September 2022].

[77] Perez, D., 2020. *Hyundai and Autodesk Collaborate on Concept Design for Innovative Walking Vehicle* [Online]. Available from: https://www.engineering.com/story/hyundai-and-autodesk-collaborate-on-concept-design-for-innovative-walking-vehicle [Accessed 4th September 2022].

[78] 2016. *Pioneering bionic 3D printing* [Online]. Available from: https://www.airbus.com/en/newsroom/news/2016-03-pioneering-bionic-3d-printing [Accessed 4th September 2022].

[79] Hawes, P., 2020. *Generative Design And Additive Manufacture Of Formula Student Suspension Uprights.* Bath.

[80] MacDonald, W.D. and Eagar, T.W., 1992. Transient Liquid Phase Bonding. *Annual Review of Materials Science,* 22(1), pp. 23-46.

[81] George, T., Deshpande, V.S. and Wadley, H.N.G., 2013. Mechanical response of carbon fiber composite sandwich panels with pyramidal truss cores. *Composites Part A: Applied Science and Manufacturing,* 47, pp. 31-40.

[82] Liu, X., Alizadeh, V. and Hansen, C.J., 2020. The compressive response of octet lattice structures with carbon fiber composite hollow struts. *Composite Structures,* 239, p. 111999.

[83] Jiang, S., Sun, F., Zhang, X. and Fan, H., 2017. Interlocking orthogrid: An efficient way to construct lightweight lattice-core sandwich composite structure. *Composite Structures,* 176, pp. 55-71.

[84] Cheung Kenneth, C. and Gershenfeld, N., 2013. Reversibly Assembled Cellular Composite Materials. *Science,* 341(6151), pp. 1219-1221.

[85] Zhao, Y., Liu, M., Zhang, T., Xu, B. and Zhang, B., 2020. Compressive properties of reversibly assembled lattice structures. *Journal of Reinforced Plastics and Composites,* 40(11-12), pp. 422-436.

[86] Ahn, B.Y., Shoji, D., Hansen, C.J., Hong, E., Dunand, D.C. and Lewis, J.A., 2010. Printed Origami Structures. *Advanced Materials,* 22(20), pp. 2251-2254.

[87] Yan, Z., Zhang, F., Wang, J., Liu, F., Guo, X., Nan, K., Lin, Q., Gao, M., Xiao, D., Shi, Y., Qiu, Y., Luan, H., Kim, J.H., Wang, Y., Luo, H., Han, M., Huang, Y., Zhang, Y. and Rogers, J.A., 2016. Controlled Mechanical Buckling for Origami-Inspired Construction of 3D Microstructures in Advanced Materials. *Advanced Functional Materials,* 26(16), pp. 2629-2639.

[88] Gurley, A., Beale, D., Broughton, R. and Branscomb, D., 2015. The Design of Optimal Lattice Structures Manufactured by Maypole Braiding. *Journal of Mechanical Design,* 137(10).

[89] Thomas, D. and Gilbert, S., 2015. *Costs and cost effectiveness of additive manufacturing: A literature review and discussion.*

[90] Melchels, F.P.W., Feijen, J. and Grijpma, D.W., 2010. A review on stereolithography and its applications in biomedical engineering. *Biomaterials,* 31(24), pp. 6121-6130.

[91] Jacobsen, A., Carter, W. and Nutt, S., 2007. Micro-Scale Truss Structures Formed From Self-Propagating Photopolymer Waveguides. *Advanced Materials,* 19, pp. 3892-3896.

[92] 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, p. 108336.

[93] Nyiranzeyimana, G., Mutuku Mutua, J., Ochuku Mbuya, T. and Robert Mose, B., 2022. Optimization of fused deposition modelling process parameters and the effect on residual stresses of built parts. *Materials Today: Proceedings*.

[94] Hofstätter, T., Gutmann, I., Koch, T., Pedersen, D., Tosello, G., Heinz, G. and Hansen, H., 2016. *Distribution and Orientation of Carbon Fibers in Polylactic Acid Parts Produced by Fused Deposition Modeling.*

[95] Hofstätter, T., Pedersen, D.B., Tosello, G. and Hansen, H.N., 2017. Applications of Fiber-Reinforced Polymers in Additive Manufacturing. *Procedia CIRP,* 66, pp. 312-316.

[96] Ivey, M., Melenka, G., Carey, J. and Ayranci, C., 2016. Characterizing short-fiber-reinforced composites produced using additive manufacturing. *Advanced Manufacturing: Polymer & Composites Science,* 3.

[97] Brooks, H. and Molony, S., 2016. Design and evaluation of additively manufactured parts with three dimensional continuous fibre reinforcement. *Materials & Design,* 90, pp. 276-283.

[98] Mori, K.-i., Maeno, T. and Nakagawa, Y., 2014. Dieless Forming of Carbon Fibre Reinforced Plastic Parts Using 3D Printer. *Procedia Engineering,* 81, pp. 1595-1600.

[99] Matsuzaki, R., Ueda, M., Namiki, M., Jeong, T.K., Asahara, H., Horiguchi, K., Nakamura, T., Todoroki, A. and Hirano, Y., 2016. Three-dimensional printing of continuous-fiber composites by in-nozzle impregnation. *Scientific Reports,* 6.

[100] Yamawaki, M. and Kouno, Y., 2017. Fabrication and mechanical characterization of continuous carbon fiber-reinforced thermoplastic using a preform by three-dimensional printing and via hot-press molding. *Advanced Composite Materials,* 27, pp. 1-11.

[101] *3D printing composites with continuous fiber* [Online]. Available from: https://www.compositesworld.com/articles/3d-printing-composites-with-continuous-fiber [Accessed 4th October 2021].

[102] *Markforged* [Online]. Available from: https://markforged.com/ [Accessed 4th October 2021].

[103] Zheng, X., Smith, W., Jackson, J., Moran, B., Cui, H., Chen, D., Ye, J., Fang, N., Rodriguez, N., Weisgraber, T. and Spadaccini, C.M., 2016. Multiscale metallic metamaterials. *Nature materials,* 15(10), pp. 1100-1106.

[104] Jang, D., Meza, L.R., Greer, F. and Greer, J.R., 2013. Fabrication and deformation of three-dimensional hollow ceramic nanostructures. *Nature Materials,* 12(10), pp. 893-898.

[105] Ajdari, A., Jahromi, B.H., Papadopoulos, J., Nayeb-Hashemi, H. and Vaziri, A., 2012. Hierarchical honeycombs with tailorable properties. *International Journal of Solids and Structures,* 49(11), pp. 1413-1419.

[106] 2021. Designing Hierarchical Honeycombs to Mimic the Mechanical Behaviour of Composites. *Journal of Composites Science,* 5(1), p. 17.

[107] Sun, G., Jiang, H., Fang, J., Li, G. and Li, Q., 2016. Crashworthiness of vertex based hierarchical honeycombs in out-of-plane impact. *Materials & Design,* 110, pp. 705-719.

[108] Tang, Y. and Yin, J., 2017. Design of cut unit geometry in hierarchical kirigami-based auxetic metamaterials for high stretchability and compressibility. *Extreme Mechanics Letters,* 12, pp. 77-85.

[109] Wu, Q., Gao, Y., Wei, X., Mousanezhad, D., Ma, L., Vaziri, A. and Xiong, J., 2018. Mechanical properties and failure mechanisms of sandwich panels with ultra-lightweight

three-dimensional hierarchical lattice cores. *International Journal of Solids and Structures,* 132-133, pp. 171-187.

[110] Yin, S., Chen, H., Li, J., Yu, T.X. and Xu, J., 2019. Effects of architecture level on mechanical properties of hierarchical lattice materials. *International Journal of Mechanical Sciences,* 157-158, pp. 282-292.

[111] Zhang, L., Hu, Z., Wang, M.Y. and Feih, S., 2021. Hierarchical sheet triply periodic minimal surface lattices: Design, geometric and mechanical performance. *Materials & Design,* 209, p. 109931.

[112] Shi, J. and Akbarzadeh, A.H., 2020. 3D Hierarchical lattice ferroelectric metamaterials. *International journal of engineering science,* 149, p. 103247.

[113] Al Nashar, M. and Sutradhar, A., 2021. Design of hierarchical architected lattices for enhanced energy absorption. *Materials,* 14(18), p. 5384.

[114] An, X. and Fan, H., 2016. Hybrid design and energy absorption of luffa-sponge-like hierarchical cellular structures. *Materials & Design,* 106, pp. 247-257.

[115] Li, W., An, X., Zheng, Q., Yang, F. and Fan, H., 2019. Hierarchical design, manufacture and crushing behaviors of CFRP tubular energy absorbers. *Thin-Walled Structures,* 140, pp. 416-425.

[116] Alan, L.B., Nancy, L.J. and Scott, R.W., 2006. Smart impact management devices: experimental validation of impact triggered rapid expansion of aluminum honeycomb. *Proc.SPIE*. Available from: https://doi.org/10.1117/12.641529 [Accessed 06/12/2021].

[117] Jiang, H., Zhang, Z. and Chen, Y., 2020. 3D printed tubular lattice metamaterials with engineered mechanical performance. *Applied physics letters,* 117(1), p. 11906.

[118] Le, C., Bruns, T.E. and Tortorelli, D.A., 2012. Material microstructure optimization for linear elastodynamic energy wave management. *Journal of the Mechanics and Physics of Solids,* 60(2), pp. 351-378.

[119] Nakshatrala, P.B. and Tortorelli, D.A., 2015. Topology optimization for effective energy propagation in rate-independent elastoplastic material systems. *Computer Methods in Applied Mechanics and Engineering,* 295, pp. 305-326.

[120] Asadpoure, A., Tootkaboni, M. and Valdevit, L., 2017. Topology optimization of multiphase architected materials for energy dissipation. *Computer methods in applied mechanics and engineering,* 325, pp. 314-329.

[121] Carstensen, J.V., Lotfi, R., Guest, J.K., Chen, W. and Schroers, J., 2015. Topology Optimization of Cellular Materials With Maximized Energy Absorption. *ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. V02BT03A014. Available from: https://doi.org/10.1115/DETC2015-47757 [Accessed 8/21/2022].

[122] Maskery, I., Aboulkhair, N.T., Aremu, A.O., Tuck, C.J., Ashcroft, I.A., Wildman, R.D. and Hague, R.J.M., 2016. A mechanical property evaluation of graded density Al-Si10-Mg lattice structures manufactured by selective laser melting. *Materials Science and Engineering: A,* 670, pp. 264-274.

[123] Norton, M.P. and Karczub, D.G., 2003. *Fundamentals of Noise and Vibration Analysis for Engineers.* 2 ed. Cambridge: Cambridge University Press.

[124] Sun, D., Qiu, L., Wang, B. and Xiao, Y., 2009. Novel Helmholtz resonator used to focus acoustic energy of thermoacoustic engine. *Applied Thermal Engineering,* 29(5), pp. 945-949.

[125] Kasliwal, R., Saxena, S., Jadhav, S., Kasliwal, R., Saxena, S. and Jadhav, S., 2021. Design & Development of Helmholtz Resonator for Low Frequency Exhaust Noise Reduction in Commercial Vehicles. SAE International. Available from: https://doi.org/10.4271/2021-26-0279 [Accessed 21/08/2022].

[126] Gorain, J. and Padmanabhan, C., 2021. Broadband low-frequency noise reduction using Helmholtz resonator-based metamaterial. *Noise control engineering journal,* 69(4), pp. 351-363.

[127] Guo, J., 2019. Low-frequency sound absorber based on checkerboard Helmholtz resonators with different extended necks.

[128] Baek, S.-H., Jang, J.-Y., Song, K.-J. and Park, S.-H., 2021. Design of flat broadband sound insulation metamaterials by combining Helmholtz resonator and fractal structure. *Journal of Mechanical Science and Technology,* 35(7), pp. 2809-2817.

[129] Guo, J., Zhang, X., Fang, Y. and Jiang, Z., 2020. A compact low-frequency sound-absorbing metasurface constructed by resonator with embedded spiral neck. *Applied Physics Letters,* 117(22), p. 221902.

[130] Cai, X., Guo, Q., Hu, G. and Yang, J., 2014. Ultrathin low-frequency sound absorbing panels based on coplanar spiral tubes or coplanar Helmholtz resonators. *Applied Physics Letters,* 105(12), p. 121901.

[131] Han, L., Ji, H. and Qiu, J., 2021. A broadband sound-absorbing panel based on the coiled coplanar absorber with multiple absorption peaks. *Physica Scripta,* 96.

[132] Maa, D.Y., 1975. THEORY AND DESIGN OF MICROPERFORATED PANEL SOUND-ABSORBING CONSTRUCTIONS.

[133] Tan, W.-H. and Ripin, Z.M., 2016. Optimization of double-layered micro-perforated panels with vibro-acoustic effect. *Journal of the Brazilian Society of Mechanical Sciences and Engineering,* 38(3), pp. 745-760.

[134] Sakagami, K., Nagayama, Y., Morimoto, M. and Yairi, M., 2009. Pilot study on wideband sound absorber obtained by combination of two different microperforated panel (MPP) absorbers. *Acoustical Science and Technology,* 30(2), pp. 154-156.

[135] Li, X., Wu, Q., Kang, L. and Liu, B., 2019. Design of Multiple Parallel-Arranged Perforated Panel Absorbers for Low Frequency Sound Absorption. *Materials,* 12(13).

[136] Mahesh, K. and Mini, R.S., 2021. Theoretical investigation on the acoustic performance of Helmholtz resonator integrated microperforated panel absorber. *Applied Acoustics,* 178, p. 108012.

[137] Huang, X., Huang, H., Wu, J., Yang, M. and Ding, W., 2020. Sound quality prediction and improving of vehicle interior noise based on deep convolutional neural networks. *Expert Systems with Applications,* 160, p. 113657.

[138] Iannace, G., Ciaburro, G. and Trematerra, A., 2020. Modelling sound absorption properties of broom fibers using artificial neural networks. *Applied Acoustics,* 163, p. 107239.

[139] Jeon, J.H., Yang, S.S. and Kang, Y.J., 2020. Estimation of sound absorption coefficient of layered fibrous material using artificial neural networks. *Applied Acoustics,* 169, p. 107476.

[140] Sun, X., Jia, H., Yang, Y., Zhao, H., Bi, Y., Sun, Z. and Yang, J., 2021. Acoustic structure inverse design and optimization using deep learning. *arXiv preprint arXiv:2102.02063*.

[141] Mahesh, K., Kumar Ranjith, S. and Mini, R.S., 2021. Inverse design of a Helmholtz resonator based low-frequency acoustic absorber using deep neural network. *Journal of Applied Physics,* 129(17), p. 174901.

[142] Crandall, I.B., 1926. *Theory of Vibrating Systems And Sound.* New York: Van Nostrand.

[143] Maa, D.-Y., 1998. Potential of microperforated panel absorber. *The Journal of the Acoustical Society of America,* 104(5), pp. 2861-2866.

[144] Beranek, L.L. and Mellow, T., 2012. *Acoustics: sound fields and transducers.* Academic Press.

[145] ASTM E-1050, Standard Test Method for Impedance and Absorption of Acoustical Materials Using a Tube, Two Microphones and a Digital Frequency Analysis System. American Society for Testing of Materials, (2008).

[146] ASTM E-2611, Standard Test Method for Measurement of Normal Incident Sound Transmission of Acoustical Materials Based on the Transfer Matrix Method, (2008).

[147] Chung, J.Y. and Blaser, D.A., 1980. Transfer function method of measuring in-duct acoustic properties. I. Theory. *The Journal of the Acoustical Society of America,* 68(3), pp. 907-913.

[148] Chung, J.Y. and Blaser, D.A., 1980. Transfer function method of measuring in-duct acoustic properties. II. Experiment. *The Journal of the Acoustical Society of America,* 68(3), pp. 914-921.

[149] Li, J., Wen, X. and Sheng, P., 2021. Acoustic metamaterials. *Journal of Applied Physics,* 129(17), p. 171103.

[150] Pedro, C., Cuesta, M. and Siguero, M., 2009. Comparison of models describing double layer microperforated absorbers. *Noise Control Engineering Journal - NOISE CONTR ENG J,* 57.

[151] Sakagami, K., Morimoto, M. and Koike, W., 2006. A numerical study of double-leaf microperforated panel absorbers. *Applied Acoustics,* 67(7), pp. 609-619.

[152] Park, S.-H., 2013. Acoustic properties of micro-perforated panel absorbers backed by Helmholtz resonators for the improvement of low-frequency sound absorption. *Journal of Sound Vibration,* 332, pp. 4895-4911.

[153] Magnani, A., Marescotti, C. and Pompoli, F., 2022. Acoustic absorption modeling of single and multiple coiled-up resonators. *Applied Acoustics,* 186, p. 108504.

[154] Catapane, G., Petrone, G., Robin, O. and Verdière, K., 2023. Coiled quarter wavelength resonators for low-frequency sound absorption under plane wave and diffuse acoustic field excitations. *Applied Acoustics,* 209, p. 109402.

[155] Carvalho de Sousa, A., Deckers, E., Claeys, C. and Desmet, W., 2021. On the assembly of Archimedean spiral cavities for sound absorption applications: Design, optimization and experimental validation. *Mechanical Systems and Signal Processing,* 147, p. 107102.

[156] ASTM D695, Standard Test Method for Compressive Properties of Rigid Plastics, (2023).

[157] ASTM D638, Standard Test Method for Tensile Properties of Plastics, (2022).

[158] Bieler, S. and Weinberg, K., 2025. Energy absorption in lattice-structured materials under impact loading. *European Journal of Mechanics - A/Solids,* 114, p. 105734.

[159] *trainbr, Bayesian regularization backpropagation* [Online]. Available from: https://uk.mathworks.com/help/deeplearning/ref/trainbr.html [Accessed 26th April 2024].