

Reinforcement Learning-Based Path Planning for Additive Manufacturing of Homogenised Grid Structures

Optimising Toolpaths and Defect Distribution in 2D
Voxel Grids

ME51035: MSc Thesis
Paul van der Valk

Reinforcement Learning-Based Path Planning for Additive Manufacturing of Homogenised Grid Structures

Optimising Toolpaths and Defect Distribution in
2D Voxel Grids

by

Paul van der Valk

Van der Valk

Paul

Instructor:	M.J. Mirzaali
Project Duration:	January 2025 - November, 2025
Faculty:	Faculty of Mechanical Engineering, Delft

Preface

This thesis was motivated by my interest in optimising material behaviour and my enthusiasm for additive manufacturing. Throughout my studies, I discovered that I particularly enjoy tweaking materials and exploring the computational methods behind them. Bringing together material optimisation, programming, and additive manufacturing allowed me to work on a topic that fits my interests perfectly. Learning to work independently on a project and noticing improvement in not only my results, but also in my delivered work, professional attitude and level of analytical thinking during my thesis. Developing a streamlined method that integrates several of these techniques has taught me to appreciate the value and complexity of each approach. I am grateful for the guidance and inspiring weekly discussions with my supervisor, Mohammad Mirzaali, which strongly contributed to the development of this work. I hope to continue working on projects in this field as I move forward in my professional career.

*Paul van der Valk
Delft, November 2025*

Summary

This thesis presents a methodology for designing and manufacturing highly tunable Repetitive volume element (RVE) with tunable isotropy and mechanical properties through homogenisation reinforcement learning-guided toolpath planning for material extrusion additive manufacturing.

Background and Motivation FDM is used to create highly adaptable objects with geometrically complex shapes and features. High anisotropic behaviour is measured in conventional methods of FDM, and by optimising the mechanical properties and a toolpath of a given RVE, the result becomes highly variable. Scaling the material properties using a combination of homogenisation and RL-based toolpath planning helps extend the uses of AM.

Methodology The approach consists of three main stages. First, 2D voxel-based grids with controlled infill densities of 63% to 93% are generated to create permeable structures. Homogenisation techniques are applied to determine effective mechanical properties, e.g. Young's moduli, shear moduli and Poisson's ratios, across multiple configurations varying in voxel size, 0.1-0.4 mm, and void distribution. Second, optimal grid designs are selected using sequential optimisation with a multi-objective function balancing mechanical stiffness, isotropy, shear performance, and geometric connectivity. Third, a Double Deep Q-Network (DDQN) reinforcement learning agent is trained on 50 diverse grids to generate void-aware toolpaths for a 0.4 mm nozzle. The agent learns to maximise material coverage while avoiding collisions with voids, minimising redundant motions, and ensuring printability.

Results The trained RL agent achieved 89.4% material coverage on test grids with infill ratios of 73%, demonstrating effective void avoidance and path continuity. The optimised grid design exhibited an improvement in isotropy compared to conventional raster patterns, with 1040% reduction in mechanical anisotropy. RVEs were fabricated via MEX and mechanically tested, validating the computational predictions within close accuracy based on the isotropic behaviour, yet tensile tests show lower Young's moduli than expected from the computational model.

Significance and Future Work This work demonstrates that reinforcement learning can effectively optimise toolpath planning for complex, defect-containing structures, enabling the fabrication of FDM-produced objects with improved mechanical isotropy. Future work should extend the methodology to 3D homogenisation to eliminate plane-stress/plane-strain approximations, incorporate multi-directional AM for true 3D toolpaths, and explore non-planar geometries that more closely replicate native tissue architecture. These advances would optimise the tunable properties already existing in AM.

Contents

Preface	i
Summary	ii
1 Introduction	1
2 Theory	3
2.1 Additive manufacturing	3
2.1.1 FDM and MEX	3
2.2 Anisotropy	4
2.2.1 Anisotropy in filament	4
2.2.2 Anisotropic behaviour in a single layer	4
2.2.3 Anisotropy in the bulk material	4
2.2.4 Anisotropy in conventional printing	5
2.3 Maximal stiffnesses for relative densities	5
2.4 Homogenisation	6
2.4.1 Plane stress and Plane strain	8
2.5 Toolpath optimisation and Multi-directional AM	8
2.5.1 Planar toolpaths	8
2.5.2 Multi-directional toolpaths	9
2.5.3 Path planning	9
2.5.4 Laplacian matrix and Eulerian paths	9
2.6 Reinforcement Learning and toolpath planning	10
2.7 Translating toolpaths to G-Code	11
2.7.1 Conversion to G-code	11
2.7.2 Reconstruction of nodes	11
2.8 Limitations of AM during production	12
2.8.1 Flow of filament	12
2.8.2 Sagging and collapse	12
2.8.3 Reduced shape fidelity	13
2.8.4 Adherence between layers	13
3 Method and Materials	14
3.1 Computational model	15
3.1.1 Generating grid and voxels	15
3.1.2 Material	16
3.1.3 Shape optimisation	16
3.2 Design requirements	17
3.2.1 Random defect placement trials	17
3.3 RVE design requirements	17
3.3.1 Finite Element Analysis constraints	17
3.4 Homogenisation in MATLAB	17
3.5 Multi-objective optimisation method	18
3.5.1 Maximising Stiffness	18
3.5.2 Optimising isotropy	18
3.5.3 Poisson ratio distribution	18
3.5.4 Connectivity	18
3.6 Optimisation Method	19
3.7 Data Filtering	20
3.7.1 Data Filtering Strategy	20

3.8	Toolpath optimisation	20
3.8.1	Reconstructing nozzle-sized voxels	21
3.8.2	Reinforced learning for path planning	21
3.8.3	Reward Function	22
3.8.4	Training Strategy	22
3.8.5	Manual modification	24
3.9	G-code preparation	24
3.10	FDM and MEX	24
3.11	Reproducibility on larger scales	24
3.12	Mechanical testing	24
4	Results	26
4.1	Computational model results	26
4.1.1	Optimisation results of varying voxel size	26
4.1.2	Examples of highly tunable parameters	28
4.1.3	Optimisation results of voxel sizing	30
4.1.4	Varying infill densities	32
4.1.5	Comparison of optimised to conventional structures	36
4.2	RL path planning	37
4.2.1	Reward optimisation	38
4.2.2	Test results of RL	39
4.3	Material extrusion	40
4.3.1	Optimised space compared to extruded grids	41
4.3.2	Adjusting printing properties	42
4.4	Mechanical testing	42
5	Discussion	44
5.1	Performance Assessment of the Homogenisation Method	44
5.1.1	General remarks	44
5.1.2	Representativeness of the voxel infill	44
5.1.3	Plane stress and plane strain	45
5.1.4	Removing heuristics in grid generating	45
5.2	Reinforcement learning and toolpath planning	45
5.2.1	Defect circumventing	45
5.2.2	Bridging	46
5.2.3	Coverage and efficiency trade-off	46
5.2.4	Generalised and specialised methods	46
5.3	Comparable toolpath panning methods	48
5.3.1	Limitations of Heuristic Methods	48
5.3.2	Travelling Salesman Problem Approaches	48
5.3.3	Shortest Path Algorithms	48
5.3.4	Advantages and Considerations for DDQN RL	48
5.4	Shape fidelity	49
5.5	Comparing experimental and computational results	49
5.6	Expansion to multi materials and 3D structures	50
5.6.1	Multi materials	50
5.6.2	3D homogenisation	50
5.6.3	3D path planning	50
6	Conclusion	52
6.1	Conclusion	52
6.2	Future work	52
	References	55
A	Simulation results	60
A.1	Horizontal grid	60
A.2	Spiral grid	60
A.3	Index numbers	60

B	Additional material for RL	62
B.1	Generalised training	62
B.2	Specialised training	62

List of Figures

2.1	a) Schematic of the principal stiffness tensors on an extruded line of a filament. The stiffness tensors are directed so that E_{11} is in the extrusion direction and E_{22} and E_{33} are directed in transverse directions. b) Visualisation of a toolpath with partial infill density stacked in 90° rotated angles using planar techniques in 2.5D AM [4]. c) Extrusion direction compared to loading directions for FDM-produced materials. Transverse loading results in delamination of the material, causing lower mechanical performance of the material [20]. d) The average values of smaller-sized voxels X_{ij-1k} to X_{i-1jk} are combined to reconstruct the shape of the grid for smoother surfaces with similar results to the simulated homogenisation results [21].	5
3.1	The following flowchart presents the iterative process of identifying an optimal grid structure that accommodates multiple objectives, whilst incorporating values that have been improved in at least one objective. In the following section, the subsequent stages of the process are to be examined, with particular reference to RL-based path planning and FDM.	14
3.2	Figures that helps visualise the flowchart in Figure 3.1. a) The local coordinate system of a single voxel is created to simulate the boundary conditions [41]. b) The interaction between a voxel and the adjacent voxels visualised using the Moore neighbourhood [42]. c) A complete RVE is created of 20×20 voxels with material class 1 (yellow) or 0 (blue). d) The optimal values are chosen after performing a multi-objective optimisation, where OF is the combined optimal score. e) RL-based toolpath planning is used to determine the optimal connection between newly created nodes within the voxels. f) FDM is performed to create a sample build-up using RVEs.	15
3.3	A flowchart showing the steps used in DDQN RL used to determine the new state while learning from its prior made decisions and double DQN to reduce bias [48].	21
3.4	A computer-generated image of the tensile test object. The tensile test specimen has a length of 120 mm, a width of 21 mm, and the specimen size between clamps is 40×12 mm. 10 layers are added to the specimen, resulting in a thickness of 4 mm.	25
4.1	Conventional methods recreated as a voxel grid with $l_{voxel} = 0.2$ mm and an infill density of 73%. a-b) The plot of horizontal lines used in conventional FDM and the stiffness plot showing high anisotropic values, where the E_{22} direction has a stiffness near the rule of mixture upper boundary. c-d) The spiral grid used in conventional FDM and the corresponding stiffness plot, where more balanced but lower weakness is seen compared to the original stiffness value.	27
4.2	The mechanical properties of randomised 10×10 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	27

4.3	The mechanical properties of randomised 20×20 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	28
4.4	The mechanical properties of randomised 40×40 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	29
4.5	Plots comparing the mechanical properties of the 100 highest-scoring results according to the OF of variable voxel size comparison to conventional methods. All solutions have an infill density of 73%. a-c) The normalised stiffnesses E_{11} , E_{22} and G_{12} set up per voxel. d-e) The Poisson ratios compared to the conventional methods. Due to high variance between Poisson ratios, both are shown in the plots. f) The Zener ratio compared to the grid size, with the optimal Zener ratio of 1 given as a dashed line. g) The connectivity value determined using equation 3.5.4.	29
4.6	5 grids chosen due to their differences in stiffness. Keeping the same infill density 73%, voxel size 0.2 mm and boundary conditions, it is still possible to generate grids with highly variable stiffnesses. Each red dot corresponds to the letter of the subfigure.	30
4.7	Examples of the influence placement of defects on the connectivity compared over the OF score for a 20×20 grid with infill density of 73%. Every red note indicates a grid with the corresponding subfigure attached to it.	31
4.8	The plots with the highest score based of the OF for each voxel size considered. Infill density is set at 73%. a-b) The optimal grid with 10×10 voxels with the corresponding stiffness plot in 2D. c-d) The optimal grid for a 20×20 grid with the corresponding stiffness plot. e-f) The optimal grid for 40×40 voxels with the corresponding stiffness plot in 2D.	31
4.9	The mechanical properties of randomised 20×20 grids with an infill density of 93%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	33
4.10	The mechanical properties of randomised 20×20 grids with an infill density of 83%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	33

4.11	The mechanical properties of randomised 20×20 grids with an infill density of 63%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.	34
4.12	Plots showing the impact of different infill densities on the mechanical properties of the grid using conventional methods as comparison, having a constant voxel size of 0.2 mm. The 100 highest-scoring grids according to the OF are used in this method to show the optimal performance. a-c) The normalised stiffnesses E_{11} , E_{22} and G_{12} set up per voxel. d-e) The Poisson ratios compared to the conventional methods. Due to high variance between Poisson ratios, both are shown in the plots. f) The Zener ratio compared to the grid size, with the optimal Zener ratio of 1 given as a dashed line. g) The connectivity value determined using equation 3.5.4	35
4.13	Figure of the optimal grids and the corresponding stiffness plot in 2D with variable infill densities. a-b) The optimal grid for an infill density of 93%. c-d) The optimal grid for an infill density of 83%. e-f) The optimal grid for an infill density of 73%. g-h) The optimal grid for an infill density of 63%.	35
4.14	Comparison bar plots showing the 100 test results of the generalised, specialised and manually modified results of the RL toolpath planning method. a) Total coverage of the grids after applying the generalised, specialised and manually modified methods. b-e) Nodal coverage, coverage distribution and path length are indicators the efficiency and are plotted over all 100 test results per method.	39
4.15	An overview of the RL-generated toolpaths. In the first row, the generalised toolpath is shown, in the second row, the specialised toolpath is shown, and in the third row, the manually modified toolpath is shown. All grids are a combined set of the toolpath in the given index plot, the super-voxel representation, the nodal visits per location and the coverage progress while completing the toolpath. The highest scoring grid of 20×20 with infill density 73% is used, but the image is mirrored in the y-axis.	40
4.16	FDM produced RVEs. a-b) length, width, thickness of a $10 \times 10 \times 10$ RVE block sized $40.5 \times 40.5 \times 4.1$ mm. d) Tensile test specimen, high printing quality. e) Tensile test specimen, low printing quality caused by poor adherence. f) A scale model of RVE structures. A single layer of thickness 0.4 mm scaled 5.0 times larger than the original design.	41
4.17	2 specimen after tensile testing. a-b) specimen Y_1 after tensile testing and a detailed image showing the partial fracture through the RVE. c-d) Specimen X_2 after tensile testing and a detailed image showing a clean cut between the RVEs.	42
4.18	Tensile test results given in strain over stress performed in two dimensions with the added Young's Modulus per sample attached in the legend. a) E_{11} oriented tensile test specimens corresponding to Y-numbered specimen during tensile testing. b) E_{22} oriented tensile test specimens corresponding to X-numbered specimen during tensile testing.	43
A.1	Geometric configuration and directional stiffness characteristics of the horizontal grid structure.	60
B.1	The episode Reward for an RL function using a DDQN. 3 layers are applied to the generalised function, where a set of 50 grids is used for training.	63
B.2	The Episode reward of the specialised training, occurring after a generalised training. Only 1 grid is used during the training of the specialised grid. The Episode reward increases at a lower rate than during 1500 episodes.	63

List of Tables

3.1	A set of variables within the optimisation method is maintained at a constant throughout the optimisation process.	16
3.2	DDQN reward structure for toolpath optimisation	23
3.3	DDQN input-reward-output summary	23
4.1	Results of the highest OF score for homogenisation and optimisation of a rectangle sized l_x by l_y , with variable voxel size l_{voxel} , including conventional methods for comparison. .	30
4.2	Results of grids with the highest-scoring values in the OF and their corresponding mechanical properties after optimisation with a variable infill density, while remaining the size of the voxels constant at 20×20 with $l_x = 0.2$ mm.	34
4.3	Mechanical properties retrieved from tensile testing of 7 specimens in two directions. . .	42
A.1	Index number to the corresponding best grids.	61

1

Introduction

Additive manufacturing (AM) encompasses a range of rapid prototyping methods capable of producing materials with highly adaptable properties [1]. By controlling structural parameters such as defect placement and infill density, the mechanical properties of printed materials can be systematically tuned to achieve specific performance characteristics. This tunability is particularly valuable when transforming highly anisotropic base materials into structures exhibiting more isotropic mechanical behaviour across multiple scales.

Fused Deposition Modelling (FDM) and Material Extrusion (MEX) offer highly tunable and adaptable manufacturing processes for rapid material production. These layer-by-layer fabrication methods enable precise control over structural geometry [2], yet they inherently produce anisotropic mechanical properties [3]. Conventional FDM-produced materials exhibit optimised strength in the extrusion direction whilst demonstrating reduced performance in perpendicular directions. This mechanical anisotropy, characteristic of layered manufacturing, limits the applicability of FDM components in applications requiring uniform mechanical response regardless of loading direction [3]. The high anisotropy observed in conventional layer-by-layer manufacturing is commonly addressed through interlayer solutions, such as rotating the extrusion direction between successive layers [4]. Whilst effective for simple geometries, this approach offers limited adaptability for complex multidirectional structures. Alternative strategies involving intricate print path optimisation show promise for reducing anisotropy, yet systematic methods for determining optimal toolpaths that enhance isotropy remain underdeveloped and require greater tunability.

This work addresses this challenge by developing a two-stage methodology to optimise material properties in FDM-produced structures. First, the mechanical properties of grid-based structures are evaluated to identify optimal geometries. Homogenisation techniques are used to computationally determine the placement of defects within grid structures, producing designs that exhibit near-isotropic behaviour. Optimal infill densities ranging between 63% and 93% are identified through this analysis. Second, these homogenised structures must be translated into feasible toolpaths executable by FDM nozzles. This translation process must account for practical constraints, including the minimisation of extrusion errors (over-extrusion and under-extrusion) and the reduction of non-productive travel movements.

The homogenisation process provides a computationally efficient approach to approximating the mechanical properties of complex geometries[5] [6]. By discretising structures into repetitive voxel units that interact solely with adjacent neighbours, this method significantly reduces computational cost whilst enabling systematic evaluation of diverse geometric configurations. This framework facilitates the identification of structures achieving optimal mechanical stiffness in multiple directions whilst maintaining near-isotropic behaviour through systematic variation of voxel resolution, void fraction, and defect distribution.

Translating optimised voxel grids into practical MEX toolpaths requires the determination of efficient trajectories between nodes. Research demonstrates that toolpaths exhibiting minimal discontinuities and retractions produce structures that more faithfully reproduce computational predictions. Reinforce-

ment Learning (RL), specifically Double Deep Q-Networks (DDQN), provides an adaptive framework for generating logically connected toolpaths from optimised grids. This approach demonstrates robustness when applied to structurally similar configurations, enabling automated toolpath generation across diverse geometries.

The central research question addressed in this report is: *How can the material properties of FDM-produced components be optimised and tuned using homogenisation and RL-based path planning?*

The Theory chapter establishes the material and mechanical property requirements for optimised structures and reviews current manufacturing technologies to identify design constraints for homogenised grid geometries. The Methods chapter presents a two-stage approach: first, optimal voxel grid geometries are generated through homogenisation-based structural optimisation, with mechanical performance evaluated across varying void ratios and resolutions. Second, a Double Deep Q-Network (DDQN) reinforcement learning agent is trained to generate continuous toolpaths for material extrusion from the optimised grids. The Results chapter evaluates this methodology through two stages: comparison of homogenised structures with conventional patterns in terms of mechanical isotropy and stiffness, followed by assessment of RL-generated toolpaths for coverage efficiency, print fidelity, and manufacturing feasibility relative to idealised simulation geometries. The Discussion chapter critically examines variances between predicted and achieved optimisation outcomes, analyses the strengths and limitations of the DDQN approach compared with traditional path planning algorithms, and identifies opportunities for methodological refinement in future work.

2

Theory

2.1. Additive manufacturing

Additive manufacturing (AM) is defined as the process of utilising rapid prototyping, whereby material is deposited in order to create an object. The utilisation of AM has numerous advantages, notably being the ability to generate intricate geometries, extensive personalisation, and expedited prototyping [1]. The selection of the most appropriate AM technique is contingent upon the dimensions of the design, the material to be utilised, and the intended application of the product. A plethora of prevalent techniques can be employed for AM. AM is classified in seven methods: Material Extrusion (MEX)[7], Direct Energy Deposition (DED), vat polymerization, powder bed fusion, material jetting, sheet lamination, and binder jetting [8]. The focus of this report is on MEX and especially on Fused Deposition Modelling (FDM), due to its accessibility, low cost, high control during the process, and high tunability [9].

FDM is a process whereby material is deposited in layers subsequent to its fluidisation, which is induced by either an increase in temperature, the application of pressure, or the extrusion of the material as droplets or a steady flow.

Some important influences on mechanical properties are: the infill density of a material, local defect formation, and anisotropic behaviour created by non-homogeneous bulk material.

Depending on the material used, there are different requirements that must be met. Polymers, such as poly-lactic acid (PLA), are known to undergo rapid hardening after extrusion. In order to ensure optimal interlayer connection, it is essential that the material is applied with greater force to the underlying layer.

2.1.1. FDM and MEX

FDM or Fused Filament Fabrication (FFF) is used to create layered structures using MEX. Extruded material creates a pattern, in conventional methods mainly focus on the repeatable motions, similar to the original form. A pattern of horizontal lines stacked to one another is commonly used (Figure 2.1b) to improve the mechanical performance unidirectionally, uses simple geometry, and is reproduced easily [10].

The use of FDM results in high anisotropy, whereby the scaling occurs from a single strand of extruded material to a layer and finally to the complete structure. A strand of extruded material, composed of a single material, shows increased stiffness in the longitudinal direction, corresponding to the extrusion direction [11]. The stiffnesses in the transverse directions are reduced compared to the longitudinal direction. It has been demonstrated that low isotropic abilities have the capacity to reduce the mechanical performances of a bulk material when stresses are applied in directions that are not equivalent to the layered directions.

In FDM we see material defects, either by not creating a maximum infill of the volume or experiencing poor adhesion quality of the material. The poor adhesion in the 2.5D layer FDM is most notable in the z-direction. Layers of extruded material cool down and adhere poorly to newly extruded layers [12]. Conventional methods of printing where horizontal lines are stacked next to each other show

low stiffness if forces are applied perpendicular to the strands seen in the structure. A method to reduce single-directional high stiffness is by rotating the extrusion direction per layer. Rotation of layers improves the mechanical properties of a 2D bulk structure.

Solving the problem of anisotropy created per layer on a layered basis has been proven by creating repeatable patterns that include multiple extrusion directions per shape. Examples of variable patterns are the spiral, grid, gyroid or honeycomb pattern [10].

By selectively distributing local defects to match a presented infill density, the mechanical properties can be tuned and optimised. Local defects seen in the model can be tuned in position to minimise the anisotropic behaviour. The distribution of defects is optimised using numerical simulation methods to create optimal materials, looking into a limited design space. By expanding outside of the conventional methods that have a bounded material setup, the optimal shape and defect distribution in a simple structure can be optimised. The bulk material is always used to compare the performance of the final result, testing the optimal shape.

2.2. Anisotropy

Anisotropy presents a major challenge for the mechanical performance of AM materials, including bioscaffolds [13], composite structures [14], and reinforced materials [15]. The layer-by-layer deposition process in FDM inherently induces directional dependencies in mechanical behaviour. Material extruded along one direction forms continuous strands with strong intralayer bonds, while transverse and interlayer connections remain discontinuous and weaker [11], [16]. This directional disparity leads to anisotropic stiffness and strength, which must be understood across multiple scales.

2.2.1. Anisotropy in filament

At the filament scale, FDM typically employs either homogeneous single-material filaments or composite blends. Pure polymers such as PLA are assumed to be isotropic at this scale, whereas multi-material or fibre-reinforced filaments, such as PLA reinforced with carbon fibre or biowaste inclusions, exhibit inherent inhomogeneity and direction-dependent stiffness [17]–[19]. These internal anisotropies are directly transferred to the printed structure, influencing local stress distribution and adhesion.

2.2.2. Anisotropic behaviour in a single layer

Within a single printed layer, the highest stiffness E_{11} occurs along the extrusion direction, where the continuity of the material is greatest (Figure 2.1a). In the transverse directions, E_{22} and E_{33} , the stiffness is lower because adjacent extruded lines are only connected through narrow contact regions [11]. If the spacing between lines is reduced, partial fusion improves load transfer and enhances stiffness in these transverse directions (Figure 2.1c). However, defects such as under- or over-extrusion interrupt strand continuity, creating unplanned defects that further amplify anisotropy. Minimising these process-related irregularities and optimising toolpaths to align with stress trajectories are essential to improving mechanical uniformity within the layer.

2.2.3. Anisotropy in the bulk material

At the bulk scale, anisotropy becomes more pronounced due to imperfect interlayer adhesion. Lower layers cool before subsequent deposition, leading to incomplete fusion at the interface and a noticeable reduction in stiffness along the build direction. Consequently, the stiffness components in the transverse directions E_{22} and E_{33} are significantly smaller fractions of the longitudinal stiffness E_{11} , and the material is prone to delamination under tensile or shear loads. While alternating the extrusion direction between successive layers can reduce anisotropy in the x–y plane, this approach does not fully resolve weak interlayer bonding.

Thermal gradients and residual stresses accumulating during printing further exacerbate this behaviour, particularly in large parts. Optimised, stress-oriented toolpaths and controlled layer cooling can mitigate these effects by promoting more uniform bonding.

Anisotropy in 2D and 3D structures can be quantified using the Zener ratio [22].




Figure 2.1: a) Schematic of the principal stiffness tensors on an extruded line of a filament. The stiffness tensors are directed so that E_{11} is in the extrusion direction and E_{22} and E_{33} are directed in transverse directions. b) Visualisation of a toolpath with partial infill density stacked in 90° rotated angles using planar techniques in 2.5D AM [4]. c) Extrusion direction compared to loading directions for FDM-produced materials. Transverse loading results in delamination of the material, causing lower mechanical performance of the material [20]. d) The average values of smaller-sized voxels X_{ij-1k} to X_{i-1jk} are combined to reconstruct the shape of the grid for smoother surfaces with similar results to the simulated homogenisation results [21].

$$A = \frac{C_{11} - C_{12}}{2C_{44}} = \frac{E}{2G(1 + \nu)} \quad (2.1)$$

where $A = 1$ corresponds to isotropy. Values approaching 0 or 2 indicate significant directional stiffness differences. Understanding and controlling anisotropy across filament, layer, and bulk scales are essential for designing FDM structures with predictable, near-isotropic mechanical behaviour.

2.2.4. Anisotropy in conventional printing

The anisotropy is seen in 2.5D layered conventional FDM used in conventional toolpath printing methods (Figure 2.1c). The stacked layers are prone to delamination if loaded in a traverse direction. By optimising the build direction layer-by-layer, this effect is reduced. The low adherence in the build direction reduces the mechanical properties of the bulk material, leaving the object anisotropic.

2.3. Maximal stiffnesses for relative densities

In FDM, it is common to use reduced infill densities. As a result, lower mechanical stiffness is expected and to determine the maximal stiffness, the Rule of Mixtures or Ashby-Gibson equation is used.

To determine the stiffness maximally achievable for a material that contains defects, the Rule of Mixtures is used. Intended to be used for composite materials to determine the stiffness of a material with reinforced fibres [23]. The equation for the rule of mixtures is written in equation 2.2. Stiffness E_{11} of the material is feasible in the direction parallel to the fibres placed in the material. This method can be used to determine the maximum achievable mechanical stiffness for a grid with known defects.

$$E_{11} = fE_f + (1 - f)E_m \quad (2.2)$$

Where f is the fraction of fibres, E_f is the stiffness of the fibres and E_m is the stiffness of the bulk

material. When opting for a E_f of 0.

On the contrary, when removing the fibre stiffness from the equation, the stiffness in the perpendicular direction E_{22} is written as seen in Equation 2.3.

$$\frac{1}{E_{22}} = \frac{f}{E_f} + \frac{1-f}{E_m} \quad (2.3)$$

Removing this factor gives an answer with no feasible outcome. For the unidirectional stiffness, this equation can approximate a maximum value for the material.

The relative density affects the stiffness of the material. Using the Gibson-Ashley power law, the relative density can be approximated for porous materials and lattice structures [24] as seen in Equation 2.4.

$$\frac{E^*}{E_s} = C \left(\frac{\rho^*}{\rho_s} \right)^n \quad (2.4)$$

Where C is the Ashby-Gibson constant, E^* equals the stiffness of the reduced density of the material, E_s equals the solid stiffness, ρ^* is the relative density and n is the load factor, either being 1 for strain-dependent loads and 2 for bending-dependent loads. The theoretical maximal limit is set for values $C = 1$ $n = 1$. Assumptions can be removed when the material is validated using tensile tests. A fitted curve for reduced infill densities can be created using experimental results.

It is evident that, when the theoretical maximal values are applied, the Ashby-Gibson equation is equivalent to the Rule of Mixtures in equation 2.2 when a single material is employed. Consequently, the two equations can be utilised interchangeably.

2.4. Homogenisation

Homogenisation is a computational method used to approximate the mechanical properties of complex structures by reducing them to equivalent homogeneous materials of simple geometry. This approach significantly reduces computational costs for both lattice structures [25]–[27] and composite materials [6]. The method relies on defining Representative Volume Elements (RVEs) that capture the microstructural properties of the material and can be scaled to represent the global domain [26], [28], [29]. By applying uniform strain or stress boundary conditions to these RVEs, material properties can be determined without requiring full-scale simulations.

The homogenisation process has been implemented in accessible computational frameworks. Ozdelik et al. [30] developed a Python-based homogenisation method that uses finite element analysis (FEA) to calculate the effective Young's modulus of materials with complex geometries. Their approach discretises the geometry into constant voxels, enabling efficient property calculations.

For AM materials, homogenisation is particularly relevant due to their inherent anisotropy. The extrusion process creates higher stiffness along the nozzle travel direction compared to transverse directions. At the macroscale, this directional dependency results in spatially varying effective material properties. Homogenisation theory [5], [31], [32] enables the characterisation of these lattice-like structures by examining local stiffness, strain, and elasticity tensors within each voxel element, thereby predicting macro-scale mechanical behaviour from microstructural properties.

The principle of homogenisation is shown in Equation 2.5 as defined by Hassani et al. [32].

$$E_{ijkl}^H = \frac{1}{||Y||} \int_Y (\varepsilon_{pq}^{0(ij)} - \varepsilon_{pq}^{*(ij)}) E_{pqrs} (\varepsilon_{rs}^{0(kl)} - \varepsilon_{rs}^{*(kl)}) \quad (2.5)$$

The elasticity equation that functions as the basis for Equation 2.5 is written in Equation 2.6.

$$\int_Y E_{ijpq} \varepsilon_{rs}^{*(kl)} \frac{\partial v_i}{\partial y_i} dy = \int_Y E_{ijpq} \varepsilon_{rs}^{0(kl)} \frac{\partial v_i}{\partial y_i} dy, \quad \forall v_i \in H_{per}(Y, R^3) \quad (2.6)$$

The strain tensor is given in Equation 2.7, where is the displacement field within the virtual work principle or the weak form of the equilibrium equation.

$$\varepsilon_{rs}^{kl} = \varepsilon_{rs}(x^{kl}) = \frac{1}{2}(u_{r,s}^{kl} + u_{s,r}^{kl}) \quad (2.7)$$

Following the stress-strain relationship, it must be concluded that the force or stress in both directions should be equal, since the strain seen in Equation 2.7 is equal in both directions. The local stiffness matrix is written as seen in Equation 2.8.

$$[E^H] = \frac{1}{||Y||} \int_Y (u^0 - u^*)^T [B]^T [E] [B] (u^0 - u^*) dY \quad (2.8)$$

The local strain over the spring stiffness of the material is determined using Equation 2.9, where E_{ij}^H is equal to C_{ij}^H , also referred to as the homogenised constitutive matrix.

$$E_{ij}^H = \frac{1}{||Y||} \sum_{e=1}^{N_e} \int_{y(e)} (u_e^{o(i)} - u_e^{*(i)})^T K^{(e)} (u_e^{o(i)} - u_e^{*(i)}) dY^{(e)} \quad (2.9)$$

Where Y is the total area, u is the displacement, and $K^{(e)}$ is the stiffness matrix.

The stress tensor is determined using Equation 2.10 [33].

$$\sigma_{ij} = E_{ijkl} \varepsilon_{kl} \quad (2.10)$$

For anisotropic stress-strain relationships, Equation 2.11 is used to determine the elasticity moduli, shear modulus, and Poisson ratio in 2D. Where the strain equation is written on the left side of the equation, the compliance matrix S is multiplied by the stress matrix on the right side of the equation [33]. All matrices and tensors are written in Voigt notation [34] similar to a 2D example in Equation 2.11.

$$\begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \gamma_{12} \end{bmatrix} = \begin{bmatrix} \frac{1}{E_1} & -\frac{\nu_{21}}{E_2} & 0 \\ -\frac{\nu_{12}}{E_1} & \frac{1}{E_2} & 0 \\ 0 & 0 & \frac{1}{G_{12}} \end{bmatrix} \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \tau_{12} \end{bmatrix} \quad (2.11)$$

Where ε is the strain and γ is the shear strain, calculated by multiplying the stress σ and shear stress τ by the compliance matrix S.

Derived from Equation 2.11, the factors S_{12} and S_{21} contain Poisson ratios specific to the loading direction. The Poisson ratio is approximated using Equation 2.12.

$$\nu_{12} = -\frac{S_{21}}{S_{11}} \text{ and } \nu_{21} = -\frac{S_{12}}{S_{22}} \quad (2.12)$$

The elasticity matrix [E] for each element is represented using the Lamé parameters for a system that is placed under loads written in Equation 2.13.

$$E^{(e)} = \lambda^{(e)} \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} + \mu^{(e)} \begin{bmatrix} 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.13)$$

Where the two Lamé parameters are determined according to Equation 2.14. The Lamé parameters are used to scale the geometric matrices that are the result of the process of homogenisation.

$$\lambda = \frac{Ev}{(1+v)(1-2v)} \text{ and } \mu = \frac{E}{2(1+v)} \quad (2.14)$$

Where E is Young's Modulus and ν is Poisson's ratio for the base material.

Homogenisation is used to approximate the material properties of complex but repetitive shapes. By focusing on smaller objects that can be repeated often, the run time of stiffness calculation is reduced drastically. Voxels are used as a simplification of the placement of material. Since the resolution of a voxel influences the accuracy of homogenisation, the size of voxel lengths is critical to determine the material used. Varying in the resolution of the voxel is used to determine the effectiveness of a shape to improve the accuracy of shape performance.

2.4.1. Plane stress and Plane strain

In two-dimensional homogenisation, it is necessary to assume either plane stress or plane strain conditions to represent the missing out-of-plane behaviour [35]. Since the model is restricted to a flat geometry, stresses and strains in the z -direction cannot be simultaneously evaluated. In a plane-stress formulation, the out-of-plane stress component is neglected ($\sigma_z = 0$), which is appropriate for thin structures or membranes that can deform freely through the thickness. Conversely, in a plane-strain formulation, the out-of-plane strain component is constrained to zero ($\varepsilon_z = 0$), representing thick or laterally confined bodies.

These assumptions significantly influence the calculated stiffness tensor. When plane-strain conditions are applied to a filled grid, the resulting effective stiffness often appears higher than the input material stiffness because the z -direction deformation is artificially restricted. In contrast, plane-stress conditions typically produce lower stiffness values due to the absence of constraint in the out-of-plane direction.

In homogenisation-based optimisation, this choice introduces numerical challenges. The 2D formulation yields either an over-constrained (plane-strain) or under-constrained (plane-stress) stiffness matrix, resulting in singularities or unrealistic eigenvalues when solving for equivalent material constants. This is a direct consequence of eliminating one physical dimension without a full 3D coupling term.

These issues disappear in 3D FEA, where all stress and strain components are explicitly defined and no assumptions on the z -direction are required. Nevertheless, the 2D approach remains attractive for its lower computational cost and for providing initial insights into anisotropy trends [35].

For reference, under plane-stress conditions, the first Lamé parameter can be expressed as:

$$\lambda_{ps} = \frac{E\nu}{1 - \nu^2} \quad (2.15)$$

2.5. Toolpath optimisation and Multi-directional AM

A key strategy to minimise anisotropy in AM is to generate multi-directional toolpaths that align with the principal stress directions of the part. By applying optimisation methods to define extrusion trajectories that mimic stress distributions, the need for support structures can be reduced, surface roughness minimised, and mechanical performance enhanced through more uniform load transfer.

Conventional toolpath generation methods rely on STL-based geometry, where the object is discretised into nodes (vertices) connected by struts or edges. These are used to represent simplified geometries in modelling software such as SolidWorks or computational tools like MATLAB. The mesh is then converted into a toolpath and subsequently into G-code for AM. However, this workflow focuses solely on geometric accuracy and neglects material and stress-related optimisation. As a result, the optimal toolpath produced from an STL file often fails to preserve desirable mechanical properties. In contrast, generating a physics-informed toolpath directly—then converting it to G-code—retains process-level control over stress and strain pathways, enabling more functional prints.

2.5.1. Planar toolpaths

It is imperative to examine the toolpaths that adhere to planar printing techniques in order to observe the impact of this method on anisotropic behaviour. The homogenisation method is employed, resulting in the struts being located solely on the planar level. In the z -direction, the layers are connected by weaker links. It is hypothesised that the largest strength components originate from the longitudinal

direction, thereby yielding the largest stiffness value in the same direction. The selection of an alternating direction for the principal stress, in conjunction with the establishment of connections between all optimal shapes, as observed in the process of grid structuring, results in the generation of a shape consisting of multiple points or vertices. The aforementioned points are connected by applying the Eulerian path planning method, as elucidated in the Theory section. The distance between the points is less than the radius. Consequently, it is only possible to establish a connection between a single point and two other points. However, there are exceptions to this rule. These exceptions include the starting and ending points, as well as points that are separated by a defect in the structure. These singularities must be avoided, since a reduction in internal stress is the result of this process [36].

A comparison of the connectivity between points with the simulation results is recommended, since this will reduce the mechanical properties of the output. A singularity can be regarded as a suboptimal outcome, as it signifies a point that is not optimally connected to the other points.

2.5.2. Multi-directional toolpaths

To resolve the anisotropic behaviour of planar toolpaths, multi-directional toolpaths can be used to increase the mobility range between nodes.

The feasibility of multi-directional AM is constrained by the DOF of the printer. Standard FDM systems operate with three translational DOFs, limiting printing to planar, layer-per-layer deposition. Multi-axis systems introduce additional rotational or translational DOFs, allowing non-planar and multi-directional deposition that can significantly reduce anisotropy. However, these systems require advanced motion planning to prevent nozzle–part collisions and to manage overhangs and sagging. Rotating either the printhead or the build platform can help maintain smooth toolpaths while avoiding interference.

2.5.3. Path planning

Path planning translates the digital model of a part into an optimised extrusion pattern for FDM. It defines the sequence and direction of material deposition to connect points efficiently while maintaining the desired infill density and mechanical properties. AM relies on numerically controlled (NC/CNC) systems, which require digitised toolpaths that minimise path length, reduce collision risk, and ensure consistent extrusion quality.

Path planning typically involves two main stages:

1. **Toolpath generation:** where geometric and physical constraints are used—often through graph theory, to identify efficient connections between nodes.
2. **Toolpath digitisation:** which converts these paths into machine-readable code, e.g., G-code, for execution by the printer.

2.5.4. Laplacian matrix and Eulerian paths

A method to create an optimal toolpath for AM is to build up the object into a set of nodes that are connected by edges. The number of nodes connecting to a single node is written in the Laplacian matrix as visualised in equation 2.16. By applying graph theory and Laplacian matrix analysis, an optimised toolpath can be generated that minimises disconnected segments and minimises high local stress distribution [37]. The Laplacian matrix quantifies the connectivity between nodes, where d_i represents the number of edges connected to node i , and -1 indicates a connection between nodes i and j :

$$L_{ij} = \begin{cases} d_i, & \text{if } i = j \\ -1, & \text{if } i \neq j \text{ and nodes } i, j \text{ are connected} \\ 0, & \text{otherwise} \end{cases} \quad (2.16)$$

By analysing the determinant of this Laplacian matrix, an Eulerian path can be created following the optimal shape within the given set of vertices. A true Eulerian path is constrained by multiple rules, but allowing the application of weakened constraints to the Eulerian path results in feasible outputs.

The Laplacian matrix is a combination of the Degree matrix and the Adjacency Matrix and can be written as $L = D - A$ where D is the degree matrix and A is the adjacency matrix. The degree matrix determines the number of connections per node. Edges are placed between two connecting nodes and the number of edges per node is determined by the adjacency matrix.

The Laplacian, as a result, gives the possible travel directions for each node compared to the others. Distances between nodes are not specified in this matrix and should be taken into account to check the feasibility of connecting two or several nodes. Having no directional constraints placed between nodes results in a symmetric Laplacian matrix, while a hierarchical directional matrix, where travel motions are solely restricted from one node to another, results in an asymmetric matrix.

Path planning is used to travel from one location to another while performing tasks. Using the Laplacian matrix, an Eulerian path can be formed connecting nodes while travelling from a starting node to a finishing node. Having constraints as visiting each node, avoiding objects and limiting the number of times a node and edge can be visited, helps to improve the toolpath to form a logical shape.

Combining RL and techniques to perform toolpath planning helps to create toolpaths that follow an optimised path, visiting as many nodes as possible, while being aware of other parameters that influence the mechanical properties of the extruded result. Combining real-world aspects with graph theory helps to find an optimal solution for path planning.

2.6. Reinforcement Learning and toolpath planning

Applying a Reinforcement Learning (RL) tool to optimise the toolpath helps to generate a printing path that follows multiple input parameters. The input parameters determine the state of the function, and based on the action, the state is changed. This is the function of a Double Deep Q Network (DDQN). Combining the state and the action determines the Q-value, $Q^* : State \times Action \rightarrow \mathbb{R}$, an estimate of each outcome that results in an output. Through hidden layers, the parameters are connected to each other, resulting in an optimised outcome, trying to expand the rewards using either exploitation or exploration. The reward function helps to indicate optimal choices performed by the DDQN by providing rewards or penalties.

Considering the use of RL with multiple layers that are fully connected, results in a solution that focuses on all possibilities included between layers [19]. For toolpath planning, the highest Q-value equates to the optimal reward of multiple parameters. A single reward can be tuned to either give a bonus or a penalty to the total outcome. By opting for different methods to optimise the outcome reward. In the current set-up, an optimal reward would be received when path planning is successful, by covering as much material as possible, while being aware of defects and keeping a similar material infill.

Qin et al. [38] and Twumasi et al. [39] have applied similar RL-based optimisation methods for toolpath planning to powder bed fusion. RL helps to determine multiple correct options with high variability while keeping track of previous decisions. Since each output results in an optimal movement using an RL for path planning the path can be expanded over time. As a result, the optimised shape is a guess on the currently available information classified as parameters. The parameters of the input functions determine the state and show all possibilities. Exploration is a term used in RL to determine the factor that wants to increase the reward by analysing new steps. By adding new steps to the reward per episode, the reward increases. Exploitation is a factor that uses already existing methods and tries to improve the reward by using as many explored nodes as possible. High exploitation factors reduce the number of exploration factors, resulting in poor performance, due to limited material being discovered.

On the other hand, high exploration factors result in low repetition and start to include high levels of randomness, since previously explored nodes are not exploited again. For a successful RL tool to create a toolpath that optimises both exploration in the form of nodal coverage, but exploits the function, focusing on high accessibility, filling as many nodes as possible, while holding on to similar results, helps to improve the current toolpath. Efficiency seen in the reduction of path length while visiting all nodes is seen to improve over training episodes using RL-based methods [39]. Therefore, it is important to increase the number of episodes to improve toolpaths after nodal exploration has been completed.

2.7. Translating toolpaths to G-Code

Toolpaths function as an algorithm that should be followed. Converting a toolpath to G-Code results in an algorithm that an FDM printer can follow to perform MEX. Additional information is added per line of G-Code to determine the flow rate, the printing speed, and additional machine instructions.

2.7.1. Conversion to G-code

Having determined the toolpath, it is possible to convert the planned path to G-code. G-code is often used by CNC machines, similar often used in AM to digitise the path algorithm. Giving simple instructions to a material extruder gives it the information it needs to perform AM. Travel motions, extrusion, and feed rates help to determine how much material can be deposited per distance between nodes. Machine instructions indicated with the letter M and the corresponding code can help to adjust the motor speed, fan speed, temperature of the nozzle, bed, and help indicate the start and end of the process. The feed rate F determines the velocity of the nozzle while travelling between two nodes.

After creating a toolpath, some steps need to be taken before FDM procedures can be performed. First, the toolpath must be constructed in a way that is shaped optimally for the size of a nozzle, using a reduced-size voxel demands reconstruction. Next, G-Code is applied to turn a toolpath into an algorithm that can be performed by an FDM printer. It is important that all additional information for path planning must be given in this algorithm to perform FDM.

2.7.2. Reconstruction of nodes

The process of homogenisation can be defined as the simplification of complex structures, which involves the division of the structure into smaller units. In order to achieve a more accurate estimation of the mechanical properties of a homogenised grid and create more complex geometries, it is possible to increase the resolution of the grid. This is achieved by increasing the number of elements in the block. Following the process of homogenisation, it is imperative that the structure is reconstructed at voxel sizes that can be reproduced. This process involves the convolution of voxels to generate larger voxels that approximate the diameter of a nozzle.

However, it is imperative to acknowledge that the dimensions of the nozzle significantly influence the thickness of the printing path in the AM process. It can thus be hypothesised that the value for the blocks can be averaged and used in reconstruction, as described by Ozdemir et al. [21] in their study of triply periodic minimal surface lattice structures. It is evident that the material has been segmented into blocks, with these exhibiting either a filled or an empty configuration. In Figure 2.17, the mean of the surrounding blocks is calculated, thus facilitating the creation of a more uniform surface.

$$\rho_{ijk} = \frac{\sum X_{ijk}^*}{N_{node}} \quad (2.17)$$

The process of voxelisation can be facilitated by the implementation of an algorithm that calculates the mean of a set of voxels. This method is advantageous in determining the infill density, yet the placement of defects can still present a challenge to the nozzle, necessitating the consideration of constraints to ensure precision. Due to the reconstruction, some local defects can be distorted in shape after being combined with material-filled voxels.

Either per point or per surface, a general score can be given following the analogy (Figure 2.1d). Reconstruction in the scale of the nozzle helps to create a toolpath useful for AM and meeting the lower constraints of a nozzle. This method of reconstruction is employed to enhance the surface quality of an object. It is imperative to note that the experience of defects, such as flow, pressure, or external forces, is particularly pronounced in this instance. To mitigate these issues, it is strongly recommended that the surface be smoothed.

The reconstruction of a grid pattern is necessary in order to recreate shapes that are generated from smaller-sized voxels and which are an interpretation of the original pattern. It is therefore essential to perform a comparative analysis in order to estimate the outcome of the reconstruction in relation to the original grid pattern. A thorough analysis of the loss of the original structure is necessary to determine whether reconstruction of a larger voxel size is beneficial. In the context of AM, it is imperative to

consider the shape fidelity of both the original grid and the created path, given the potential for smaller defects to either disappear or be covered by the extruded material during the process [40].

2.8. Limitations of AM during production

After a combined recreation of optimal infill patterns and an optimised version of the toolpath, FDM occurs. The limitations of FDM are commonly found, are caused by suboptimal flow, over- and under-extrusion, poor adhesion, temperature fluctuations, and low shape fidelity, resulting in a poorly manufactured bulk structure. By focusing on each factor that reduces optimal FDM, the key points can be pointed out.

2.8.1. Flow of filament

Focusing on the flow of filament that is extruded out of the nozzle onto the build plate, it is apparent that the material coming out of the nozzle must be equal to the expected volume deposited between points A and B. Since the length, thickness, and height of the material can be estimated for optimal shape quality, the material flow can be tuned to the flow. The material extruded over time is dependent on the volume of material that is extruded and the printing speed of the nozzle. High printing speed results in less material extruded over the same length and therefore must be taken into account. Poor flow can result in a set of problems: i) clogging, ii) under-extrusion and iii) over-extrusion. Optimal flow is necessary for optimal performance during FDM.

Clogging

Clogging is a severe form of under-extrusion when no material is extruded from the nozzle. If not sufficient material is extruded or the nozzle temperature is set too low, no material is properly extruded out of the nozzle, and no MEX occurs. By increasing the material flow and setting the nozzle temperature to the melting point of the material used, clogging is reduced. Physical blockage can also disrupt extrusion and cause clogging, and is often seen on the outside of the nozzle.

Under-extrusion

In the case where too little material is extruded, as a result of high printing speeds with low extrusion rates, too little material is extruded to perform FDM with the expected material volume. Local under-extrusion can occur in the outer diameter of a corner, since a larger volume is needed to extrude using the same extruded material. Under-extrusion leaves extruded material disconnected and reduces stiffness in the longitudinal direction, resulting in poor printing quality and mechanical performance.

Over-extrusion

Over-extrusion occurs if too much material is extruded between points A and B, resulting in a loss of shape of the filament, reducing the quality of the print. Over-extrusion is sometimes advised to improve adherence of the base-layer to the build plate, but it reduces the shape fidelity by increasing the infill density. Expecting a constant infill density is reduced by taking measures similar to this. Continuous extrusion of material during MEX is suboptimal when the toolpath follows a discontinuous path with varying gradients when comparing sections of the printing path. Commonly seen in sharp turns where the orientation of the turns is made that have a higher bend than 90° , excess material is deposited in the inner radius of the corner, while in the outer corner, due to the larger length of the turn, under-extrusion occurs. This results in poor shape fidelity and poorer mechanical properties.

Over-extrusion also occurs when a segment is twice extruded in the same layer. Minimising multiple extrusion motions on a single edge and including travel motions where no material is extruded to the same location is a solution to prevent this type of over-extrusion.

2.8.2. Sagging and collapse

Extruding partially fluid filaments results in the material sagging due to the effects of gravity. Especially when building thin layers in FDM have a chance to sag when left unsupported. Reducing the unsupported distance between two points so the internal forces in the material cancel the effects of gravity, minimises the effect of sagging and results, in improved layered building. The weight of the extruded material is a critical factor for sagging, therefore optimal material extrusion is necessary to prevent sagging.

For materials that are built using a single-sized extrusion strand, stacked material can fall over and collapse, reducing the shape fidelity.

2.8.3. Reduced shape fidelity

Due to the limitations of MEX, the shape of the extruded material can differ from the computationally simulated model. Poor flow and extrusion can either cause thinned or thickened extruded lines, resulting in poor results compared to the planned path. Due to the curing or hardening process, the material can deform, and it loses its original shape. Often seen in larger AM-manufactured products, the lower layers warp or curl due to differences in heat distribution, layer-per-layer curing processes.

Shape fidelity is critical to validate the result, since it is an indicator of incorrect infill density resulting in deviated bulk stiffnesses of the material.

2.8.4. Adherence between layers

Differences in temperature between the extruded material and completed layers should be taken into account to improve the adherence of layers. By pressing newly extruded materials into layers, the adherence is increased. Applying force to the filament is crucial during FDM to perform optimal interlayer adherence. Poor adherence can be improved by varying the pressure applied to the filament. Poor adherence can induce warping in larger objects due to shrinkage after gradient temperature variances in the material.

3

Method and Materials

The report is divided into 5 steps, focusing on different steps in the process. First, the numerical simulation is created, where homogenisation is used to determine an optimal composition for variable voxel sizes and infill ratios of the RVE. Next, the multi-objective optimisation is performed to give a value to the grids and compare for optimal performance. Next, reinforcement is applied to recreate an optimal grid for path planning. This would mean the current grid is recreated using FDM, where the focus lies on reproducibility and translation of optimal performance. Finally, by varying the FDM setting to recreate optimal and less optimal solutions is performed to give an additional feature is given to assure quality control. The method is given using the flowchart 3.1, and visual depictions of each step in the flowchart are given in Figure 3.2.




Figure 3.1: The following flowchart presents the iterative process of identifying an optimal grid structure that accommodates multiple objectives, whilst incorporating values that have been improved in at least one objective. In the following section, the subsequent stages of the process are to be examined, with particular reference to RL-based path planning and FDM.




Figure 3.2: Figures that helps visualise the flowchart in Figure 3.1. a) The local coordinate system of a single voxel is created to simulate the boundary conditions [41]. b) The interaction between a voxel and the adjacent voxels visualised using the Moore neighbourhood [42]. c) A complete RVE is created of 20×20 voxels with material class 1 (yellow) or 0 (blue). d) The optimal values are chosen after performing a multi-objective optimisation, where OF is the combined optimal score. e) RL-based toolpath planning is used to determine the optimal connection between newly created nodes within the voxels. f) FDM is performed to create a sample build-up using RVEs.

3.1. Computational model

The first part of the experiment is a simulation performed to estimate the stiffness tensor of a grid with a constrained size. By approximating the optimal placements of either defects in the material, the optimal mechanical performance is estimated. An optimisation function with multiple objectives is created to determine the best performance, while creating a material with close to isotropic behaviour. The optimised model is compared to conventional methods to determine the improved mechanical properties. By recreating a simplified structure of the material that behaves periodically, the periodic boundaries interact similarly in each direction. Creating a large array of highly tunable simple structures with a set infill is needed to approximate different material properties for different materials. By combining simple and repetitive structures, it is possible to locally combine different RVEs to increase the bulk properties of the material.

By creating large arrays of different planned defect placements in a set shape, the tunable properties of materials can be highlighted, which creates more optimal performance for a single structure. By efficiently placing defects over the material, local stress increases can be better distributed.

3.1.1. Generating grid and voxels

The objective is to produce a shape that is both repeatable and scalable. To this end, a square that is both straightforward to reproduce and scale is created. The grid structure is constructed from $N \times N$ voxels, which are characterised by a set height, length and width. The interaction between voxels for strain is in equilibrium, given in Equation 2.7 to relate strains. The voxels, as a form of grid or RVE of constant size, are constrained by the periodic boundary conditions and do not behave like floating objects [43].

To create defects in the RVE, voxels with a value of 0 are added to the RVE. These voxels do not interact with neighbouring blocks and have a stiffness of 0 MPa. All other blocks are classified with number 1, indicating that they are built up of the same material with the same material properties. Possibly more material can be added using multiple material constants and extra Lamé parameters.

Variable	Value and Unit
Grid length l_x, l_y and l_z	4 mm
Voxel length l_{voxel}	0.1, 0.2, 0.4 mm
Number of voxels per axis	10, 20, 40 (-)
Nozzle radius	0.2 mm
Young's Modulus E_s	3.6 GPa [44]
Poisson ratio ν	0.35 (-) [45]
First Lamé parameter λ	3.11 (-)
Second Lamé parameter μ	1.33 (-)
Infill density	73%
Maximum connecting value	8 voxels
Maximum trials	10^6

Table 3.1: A set of variables within the optimisation method is maintained at a constant throughout the optimisation process.

The defects of the grid are placed randomly in the grid, and a maximum number of attempts is given to minimise the running time for creating a grid. The grid was generated using the script `grid_maker.m` and `voxelize.m`.

The interaction between the constant blocks is summed up and divided over the total area or volume, depending on the dimension that is used in the equations to determine the homogenised Constitutive matrix C and stiffnesses in three Cartesian directions.

Constant input variables

3.1.2. Material

During MEX, PLA is used to create the RVE structure of the optimised grid. With a Young's Modulus E of 3.6 GPa [44] and a Poisson ratio of 0.35 [45]. PLA hardens out rapidly, resulting in a short curing time. The bulk properties of PLA are varied by increasing the number of defects within the obtained structure. Having relative densities that influence the infill of the structure shows a lower outcome compared to the expected result. This decrease in mechanical stiffness can be related to the rule of mixture, where the defects are shown to have material properties of 0 GPa.

PLA used in FDM is known to have lower output adherence to the build plate. Therefore, it is generally supported to increase the material flow while extruding the first layer of PLA. Lowering the nozzle closer to the build plate is also an option to improve adherence. The reduction of shape fidelity, tried to keep in the current set-up, is lost due to the extra material being extruded.

3.1.3. Shape optimisation

By optimising the location of the voxels in the RVE, a higher stiffness and shear modulus can be obtained. By randomly selecting locations for defects, the homogenisation method determines the constitutive matrix for each block. 10^6 RVEs are generated that follow the applied constraints to determine an optimal infill density and placement. To investigate if the optimisation method does not exceed the maximal stiffness obtainable with reduced material, the relative density ρ_{rel} is determined in Equation 3.1.

$$\rho_{rel} = \frac{(n_x \cdot n_y)(1 - v_f)}{n_x \cdot n_y} \quad (3.1)$$

Where v_f is the defect ratio and n_x and n_y is the number of voxels in each dimension. ρ_{rel} is equal to the infill density used in the process. Multiplying the Young's modulus by the relative stiffness gives the maximum stiffness in a single direction obtainable similar to the rule of mixture. This maximum stiffness functions as the upper limit for stiffnesses, and larger values, excluding scaled plane-stress values, are seen as computational errors.

The voxels at the edges of the RVE always need to be classified as material voxels, meaning that the design space for placing voxels is $[2, n - 1]$ where n is the total number of voxel elements.

The RVE is built up of voxels. The voxels are cube-shaped, with all sides having equal length l_{voxel} for l_x, l_y and l_z . All voxels are assumed to be uniformly shaped and stationary. The voxel is built up of 6 faces with 8 vertices. A local coordinate system is used to determine the shape of the voxel (Figure 3.2a). The grid has a physical constraint, where the outer voxels are always valued 1, for 10×10 and it has a single voxel thickness with a thickness of 0.4 mm, for the 20×20 grid it has a single thickness of 0.2 mm and for the 40×40 grid, the outer surface has a thickness of two voxels or 0.2 mm.

3.2. Design requirements

Looking at a homogenised voxel and grid, the design requirements need to be determined to create a shape that should be constrained to meet the criteria set for the experiment.

3.2.1. Random defect placement trials

Since the number of voxels and the voxel can be either filled or a defect, are already predetermined the design space can be determined using Equation 3.2.

$$D = m^{n^2} \quad (3.2)$$

Where m is the number of possibilities, 1 for filled, 0 for defect, in total 2, and n equals the number of voxels in 1 direction. Since the shape is a cuboid, n can be cubed.

When multiple materials are used in a single RVE, the value for m increases. Using a 20×20 grid, it can be concluded that calculating every possibility is too large a set of data and therefore the maximum trials is set at 10^6 . Limiting the maximal values is used to reduce the running time of the optimisation.

3.3. RVE design requirements

The shape used for the voxel should meet the criteria for the representative volume element (RVE). According to Hill [46], an RVE can be described as being entirely typical of the whole mixture on average, and the size of the material contains a sufficient number of inclusions for the apparent properties to be independent of the surface values of traction and displacement, so long as these values are macroscopically uniform.

The lower limit of an RVE is determined using the smallest continuous value that is used; for the current set-up, the smallest shape is the diameter of the nozzle, r_{nozzle} . Using dimensions for an RVE smaller than the nozzle diameter results in elements that have a variable content.

Using the grid as an RVE is considered to have a repeatable shape that has similar mechanical properties. The behaviour of a single grid is expected to be similar to every part of the structure. The dimensions of a single RVE are $4 \times 4 \times 0.4$ mm during material extrusion, where the height is constrained by the single layer extruded during 2.5D printing.

3.3.1. Finite Element Analysis constraints

During the homogenisation process, the object analysed is divided in smaller voxels that are connected to 8 voxels following the definition of a Moore Neighbourhood (Figure 3.2b) [42]. In this interaction, the material follows periodic boundary conditions [43] described by Andreassen et al. [6]. Using the periodic boundary conditions, no free movement of voxels is permitted and simplifies the computational model to have periodic DOFs linked to adjacent nodes and edges.

3.4. Homogenisation in MATLAB

Dong et al. [47] have created a homogenisation method in MATLAB to calculate the Young's Modulus of an RVE shape considered for lattice structures. The MATLAB script `homogenize.m` is openly accessible and is used to give an output that visualises the stiffness matrix E in 3 directions. Optimising the connections between the voxels determines the strength of the material in any direction. Using an FEA method as described in Theory, the material properties can be determined.

Ozdilek et al. [30] discuss that the material has an equal shear modulus for G_{12} , G_{13} and G_{23} , simplifying the anisotropic material to behave partially isotropic. Since the material properties are reduced in the

2D homogenisation method to only calculate G_{12} using the method provided by Andreassen et al. [6], no isotropic behaviour in Shear Moduli is expected.

After determining the C^H matrix for a single grid structure the S matrix, Zener ratio A, Poisson ratios ν_{12} and ν_{21} , Shear modulus and 4 stiffnesses, E_{11} , E_{22} , E_{12} and E_{21} are calculated, where $E_{12} = E_{21}$.

3.5. Multi-objective optimisation method

The material properties of a grid are compared by creating a multi-objective optimisation. This optimisation consists of multiple functions with respect to i) maximal stiffness, ii) optimal isotropic behaviour derived from the Zener ratio, iii) equal Poisson ratios and iv) maximising connectivity between voxels, while minimising direct connections between defects.

3.5.1. Maximising Stiffness

The first objective is to create a material with an optimal stiffness in both E_{11} and E_{22} directions. The maximal stiffness for a partially filled material converges to the rule of mixture [23] as its upper boundary. Since plane-strain is used to determine the mechanical properties, the rule of mixture can be exceeded and must be analysed during post-processing.

3.5.2. Optimising isotropy

Isotropic behaviour can be concluded from a Zener ratio approaching 1. Including the Zener ratio directly determines the optimal values for E_{12} and G_{12} after calculating optimal stiffness E_{11} , the Zener ratio in Equation 2.1 directly returns values E_{12} and G_{12} as optimal values.

3.5.3. Poisson ratio distribution

To ensure isotropic behaviour for both shear directions, the difference between Poisson ratios ν_{12} and ν_{21} must be minimised. To include all variables written in the compliance matrix S as shown in equation 2.10, the optimisation is subjected to multiple variables to validate isotropic behaviour and provide feasible results. A fully isotropic material follows the analogy $E_{11} = E_{22} = E$ in 2D. Similarly, the Poisson ratio of a fully isotropic material can be written as $\nu_{12} = \nu_{21} = \nu$. To ensure isotropic behaviour occurs in the grid Equation 3.5.3 d_{min} is used to determine the perpendicular distance between the point $[\nu_{12}, \nu_{21}]$ and the isotropic value line following $\nu_{21}(x) = \nu_{12} \cdot x$. Following this analogy, the Poisson ratio in all directions must also meet the isotropy results and is taken into account to determine the isotropic values

$$d_{min} = \frac{|\nu_{12} + \nu_{21}|}{\sqrt{2}} \quad (3.3)$$

3.5.4. Connectivity

The connectivity of a voxel is determined based on its placement compared to other voxels (Figures 3.2a and 3.2b). According to the Moore neighbourhood, a voxel is directly connected to 8 adjacent voxels with edge or nodal contact. The distance between the centres of two voxels for connectivity is set at $\sqrt{2} \cdot l_{voxel}$.

To determine the optimal connectivity for a complete grid, the maximal and minimal connectivity for each voxel should be determined. To meet the periodic boundary conditions of DOF, a voxel must be connected to at least 2 voxels or an external force and a voxel. Including connectivity in the optimisation function lowers the chance of larger defects starting to occur. In Equation 3.4, the summation of connectivity values of all material-filled voxels is written.

$$C_{i,j} = \sum_{p,q \in N_{ij}} g_{mat}(p,q) \quad (3.4)$$

Where g determines the connectivity in the Moore neighbourhood in domain p and q, and $g(p,q) \in 0,1$ [42]. Where N is the number of possible positions around voxel (i,j) varying between 3 and 8 for filled voxels.

$$C = \sum_{i,j} g_{mat}(i,j) C_{ij}^{mat} \quad (3.5)$$

The total number of voxels that can hold a defect has a range between $[2, n - 1]$ in both l_x and l_y direction.

C is currently written as a positive value, but in order to keep the outcome of the optimisation values scaled between 0 and 1, it is written as a penalty, whereas the total connectivity is subtracted from the optimal 1, resulting in a penalising value that improves scaling.

3.6. Optimisation Method

Structural optimisation was performed on an HP ZBook Studio G5 equipped with an Intel Core i7 processor (2.2 GHz) and 8 GB DDR4 RAM. The optimisation process applies homogenisation to randomly generated voxel grids, comparing the mechanical properties of $N = 10^6$ candidate structures through multi-objective optimisation. Voxel values are varied between 0 (defect) and 1 (material) while maintaining a constant infill density, allowing identification of optimal configurations through heuristic search. The optimisation procedure is implemented in MATLAB file `optimization_set.m`, lines 73-171.

For a voxel grid to be considered feasible, it must satisfy the following constraints: i) edge voxels are always assigned a value of 1 to maintain structural integrity, ii) the stiffness matrix remains non-singular, iii) mechanical properties remain below the rule of mixtures upper bound, and iv) the grid represents a unique geometry not previously generated.

The optimisation maximises the principal stiffness values E_{11} and E_{22} while promoting isotropic behaviour. From the stress-strain relationship in Equation 2.11, the compliance symmetry $S_{12} = S_{21}$ implies:

$$-\frac{\nu_{12}}{E_{11}} = -\frac{\nu_{21}}{E_{22}} \quad (3.6)$$

The optimisation follows the maximisation of the stiffness tensor 3.5.3 In the MATLAB code, a dominated value is used to determine the maximal output of every variable is compared to its predecessors. In the case that all outcomes of the optimised functions are larger than the predecessor, this score becomes the dominant score. This continues for N repetitions. The maximal theoretical stiffness is determined using the rule of mixtures and functions as the upper boundary for E_i . All other variables are used to achieve as much isotropic behaviour as possible and, therefore, are optimally set to 0 for optimal results.

The multi-objective function combines isotropic stiffness ratios (Equation 3.5.3), Zener anisotropy ratio (Equation 2.1), and connectivity (Equation 3.5.4) to evaluate mechanical performance in Equation 3.7

$$OF = \frac{|E_1 + E_2|}{|E_1| + |E_2|} - \frac{|\nu_{12} - \nu_{21}|}{|\nu_{12}| + |\nu_{21}|} - \frac{|A - A_{iso}|}{1 + |A - A_{iso}|} - \frac{1 - |C|}{|C|} \quad (3.7)$$

The first term rewards high combined stiffness values (approaching 1 when both E_{11} and E_{22} are positive and similar in magnitude), while the remaining three terms penalise deviations from ideal isotropic behaviour. Each deviation term is normalised using the general form in Equation 3.8.

$$F(x) = \frac{|parameter_i - parameter_j|}{\max(|parameter_i| + |parameter_j|)} \quad (3.8)$$

Each normalised term is bounded between 0 and 1. Under ideal conditions, high stiffness values with $\nu_{12} = \nu_{21}$, $A_{iso} = 1$ and $|C| = 1$, the objective function approaches its theoretical maximum of $OF = 1$. Lower values indicate suboptimal configurations with either insufficient stiffness or significant anisotropic behaviour. The objective function has no theoretical minimum, as increasingly poor designs yield increasingly negative values.

A Pareto-dominance approach identifies optimal configurations sequentially. At each iteration, a candidate solution is accepted only if it improves upon the previous best solution across all objectives

simultaneously. This strict dominance criterion ensures monotonic improvement toward the Pareto front. The relative improvement between successive accepted solutions provides a convergence metric, with diminishing returns indicating an approach to the optimum. Due to the finite iteration count (10^6), the true global optimum is approximated rather than guaranteed.

The rule of mixtures establishes the theoretical upper bound for E_i , ensuring physically realisable material properties. All other optimisation variables target values promote isotropy, with ideal convergence toward zero deviation from isotropic conditions.

3.7. Data Filtering

The optimisation process generates substantial data volumes, i.e. a single trial set produces arrays of dimension $20 \times 20 \times N_{\text{trials}}$, necessitating systematic filtering to manage storage requirements while preserving critical information.

3.7.1. Data Filtering Strategy

Optimisation datasets are stored in two formats: *complete* and *filtered*. The filtered dataset retains only configurations scoring highly on at least one objective criterion, enabling analysis of specific material properties while discarding designs that perform poorly across all metrics.

1. **Feasibility screening:** Grids failing homogenisation validity criteria are removed. Rejection rates increase substantially at lower infill densities due to connectivity constraints.
2. **Mean stiffness threshold:** The average of E_{11} and E_{22} is computed across all feasible designs. Configurations below this mean are discarded as mechanically suboptimal.
3. **Objective-specific top performers:** For each objective function component, the top 0.1% ($10^{-3} \times N_{\text{trials}}$) of designs are retained separately. This ensures representation of diverse optima across the multi-objective landscape.

The Pareto-style sequential optimisation tracking process used in the current process to determine values with an increased OF score compared to its predecessors, with a computational complexity of $O(M \cdot N(\log N)^{M1})$, where M is the number of objectives and N is the number of data points. By pre-selecting values to reduce the computational cost, the run time is reduced, and suboptimal results are removed. All values that meet the selection criteria are classified similarly to `connectivity_filtered`. All variables that are the result of the Pareto-style sequential optimisation method are classified using the variable `is_optimized`.

In order to perform a detailed structural analysis, the 100 configurations that have obtained the highest scores from each combination of voxel size and infill density are selected. This is based on the integrated objective function value (Equation 3.7). This score provides a unified metric quantifying how well each design satisfies the combined optimisation constraints.

In order to circumvent any potential bias towards a singular objective, the dataset meticulously preserves not only the aggregate scores but also the individual objective values and their corresponding grid indices. A second analysis of trade-offs between competing objectives and identification of Pareto-optimal solutions is performed using this method. Designs achieving sequential improvement during the optimisation process, i.e. those representing a new best solution at their iteration, are flagged with a marker in the dataset to facilitate convergence analysis.

3.8. Toolpath optimisation

Following the creation of an optimal grid, it is necessary to determine the toolpath that will be used with MEX. The employment of the Laplacian matrix in conjunction with an Eulerian path method facilitates the delineation of an optimal path. The restructuring of the grid into an adjacency and degree matrix serves to streamline the interconnectivity between nodes. The Laplacian matrix is used to determine optimal connections between unvisited nodes.




Figure 3.3: A flowchart showing the steps used in DDQN RL used to determine the new state while learning from its prior made decisions and double DQN to reduce bias [48].

3.8.1. Reconstructing nozzle-sized voxels

In order to ensure that the toolpaths are of equivalent size to the diameter of the nozzle, it is imperative to ascertain the distance and location of the new centre of each voxel. Path planning is constrained by a number of factors, including limited accessibility and sharp turns [19]. The relocation of the point of travel in the defect, without the specification of the true location, results in the creation of an area of reach, constrained by the borders of the voxel. In order to ascertain the most efficacious trajectory, it is possible to make minor adjustments to the location. The incorporation of the constant radius of the nozzle into the equation gives rise to certain difficulties when a voxel-sized $l_{\text{voxel}} = r_{\text{nozzle}}$ connected to multiple defects has a chance of not being reached, or when the area of the voxel is reduced.

3.8.2. Reinforced learning for path planning

A Double Deep Q-Network (DDQN) agent is employed to generate optimal toolpaths from homogenised voxel structures. The DDQN operates in a deterministic environment, learning a state-action value function that maps each state-action pair to an expected cumulative reward. The agent balances exploration and exploitation through ϵ -greedy policy, where ϵ represents the probability of selecting a random action rather than the current best estimate. Convergence to optimal Q-values is improved by increasing training episodes, adjusting the exploration rate decay, or expanding the experience replay buffer. In Figure 3.3, a flowchart is used to show the steps used in the DDQN process.

Environment Setup

The environment consists of a voxel grid with dimensions specified in Table 3.1. To accommodate the physical nozzle diameter, adjacent material voxels (value 1) are merged into "super-voxels" of size comparable to the diameter of the nozzle, while defect voxels (value 0) retain their original resolution. This reconstruction ensures the toolpath accounts for material deposition width during extrusion. Connectivity between super-voxels follows a Moore neighbourhood, allowing movement to any of the 8 adjacent positions. The grid boundaries define hard constraints preventing paths from extending beyond the RVE, ensuring the structure remains repeatable in x- and y-direction.

State Representation

The observation vector comprises 17 features characterising the current state:

- **Current node position:** Agent's location within the graph
- **Target node:** Goal position
- **Coverage metrics:** Material coverage ratio, node visitation ratio
- **Local connectivity:** Number of available neighbours from current node

- **Direction features:** Normalised vector toward target, unvisited neighbour count
- **Path efficiency:** Ratio of unique nodes visited to total steps
- **Grid properties:** Infill ratio, a normalised voxel size to r_{nozzle}
- **Step count:** Normalised episode progress
- **Edge coverage:** Ratio of traversed edges to total graph edges
- **Available degree:** Remaining connections from current node in dynamic adjacency

The starting node is selected from nodes with odd degree, preferably the minimum odd degree available. This selection aligns with Eulerian path theory: graphs with exactly two odd-degree nodes admit an Eulerian path starting at one odd node and terminating at the other. Although homogenisation constraints prevent degree-1 nodes, degree-3 nodes provide suitable starting positions. The adjacency matrix defines permissible transitions between nodes, with movement range limited to $\sqrt{2} \cdot l_{voxel}$, a Moore neighbourhood. This locality constraint promotes systematic coverage while preventing erratic long-distance jumps.

Network Architecture

The Q-network employs three fully-connected hidden layers with dropout regularisation:

- Input layer: 17 neurons, one per input parameter
- Layer 1: 256 neurons with 20% dropout
- Layer 2: 128 neurons with 20% dropout
- Layer 3: 64 neurons with 15% dropout
- Output layer: 8 neurons, one per action

This architecture follows the rule of thumb that hidden layer count should approximately equal $\log_2(n_{inputs})$ to prevent overfitting while maintaining representational capacity.

Action Space

The agent selects from 8 discrete movement directions corresponding to the Moore neighbourhood. Actions are constrained by:

- **Graph connectivity:** Only edges present in the dynamic adjacency matrix are traversable
- **Bidirectional edge removal:** Once an edge is traversed, it is removed in both directions to prevent immediate backtracking
- **Physical printing constraints:** Sharp directional changes ($> 90^\circ$) are penalised to avoid over or under-extrusion at corners

To account for nozzle diameter exceeding voxel size, sub-voxel positioning optimises node locations within each super-voxel to minimise defect overlap. The nozzle footprint is evaluated at each position, penalising paths where the circular nozzle cross-section significantly intersects defect regions.

3.8.3. Reward Function

The reward function balances multiple objectives to produce manufacturable, high-coverage toolpaths:

The reward structure prioritises edge traversal over node visitation, reflecting that material extrusion occurs along edges rather than at discrete points. Coverage milestones (80%, 90%, 95%) provide intermediate rewards to encourage complete structure printing. Additionally, reducing the number of travel motions to the bare minimum helps to improve efficiency. Travel motions also imply more singularities in the RVE.

3.8.4. Training Strategy

Two complementary training approaches are employed:

Behaviour	Reward	Rationale
Visit new edge	+300	Primary objective: maximise edge coverage
Revisit edge	$-500 \times count$	Severe penalty for material overlap
Visit new node	+50	Encourage spatial exploration
Revisit node (strategic)	-30 to -150	Allow junction revisits, penalize dead-ends
Coverage gain	$+400 \times \Delta$	Reward material deposition progress
Path smoothness	+20	Bonus for direction continuity $\cos \theta > 0.5$
Sharp reversal	-15	Penalty for 180° turns
Edge length efficiency	+10	Reward edges $> 0.8d_{nozzle}$ (low overlap)
Path overlap	$-25 \times count$	Penalize nozzle crossing previous paths
Premature dead-end	-50	Penalty if trapped before 75% coverage
Completion bonuses	+200 to +500	Scaled by final coverage (>90%)
Sparsity bonus	$250 - (\mu_{nodes} \cdot 200)$	Spread nodal visits over the entire grid
Boundary bonus	300	Promotes boundary visits to create a square RVE

Table 3.2: DDQN reward structure for toolpath optimisation

Generalised training: The agent is trained on 8 diverse grids spanning multiple infill densities and voxel sizes for 2000 episodes. This exposure teaches transferable strategies for navigating varied topologies, promoting robustness.

Specialised training: For production toolpath generation, the agent can be fine-tuned on a specific optimised grid for 1500 additional episodes, allowing adaptation to structure-specific features while retaining general navigation heuristics learned during broad training.

Training employs the following hyperparameters:

- Discount factor $\gamma = 0.995$ (long-term planning)
- Learning rate $\alpha = 5 \times 10^{-4}$
- Experience buffer: 10,000 transitions
- Minibatch size: 128
- ϵ -greedy decay: $\epsilon_{min} = 0.01$, decay rate = 0.995
- Target network smoothing: $\tau = 5 \times 10^{-4}$
- Maximum episode length: 3500 steps

Inputs	Rewards	Outputs
Optimised grid	+300 for new edge	Continuous toolpath
Super-grid adjacency	-500 for edge reuse	Total episode reward
Node coordinates	-30 to -150 for node revisit	Unique node coverage
Starting node (odd degree)	$+400 \times \Delta$ for coverage gain	Material coverage ratio
Target node	-100 for premature dead-end	Edge traversal ratio
Grid metadata (size, infill density)	+20 for path smoothness	Path efficiency metric
Dynamic adjacency	-25 for path overlap	

Table 3.3: DDQN input-reward-output summary

During training, a dynamic adjacency matrix is created that is able to remove already visited nearby nodes when the maximum number of visits is reached. Doing so reduces the need for multiple nodes to visit the same nearby nodes. Considering that the exploitation factor in the RL tool focuses on exploiting already known data, the reduction of adjacency to already known nodes can be achieved.

Additionally, nodal visits outside of the already discovered domain become the preferable option to discover, reducing the effect of crossing edges in some parts of the domain. Since a crossing domain on its own is not a problem, high concentrations of crossed edges result in local over-extrusion and reduce the homogenous properties of the RVE.

3.8.5. Manual modification

An RL-based toolpath results in a stochastic output; the hidden layers reduce the tunability of the final result based on input parameters. By adding a feature to the RL toolpath planning function to manually modify edges between nodes, it is possible to remove nodes that travel over defects to improve the shape created using homogenisation and multi-objective optimisation.

Additionally, if material-filled voxels are not connected to the toolpath, manual modification incorporates the missing nodes into the toolpath. After performing manual modification, the reward is compared to the original created toolpath using either specialised or generalised toolpath planning. Important to note is that the manual modified version of a toolpath has little underlying proof to be more optimal according to the RL function and therefore might have reduced values, but are incorporated in the results to generate higher coverage and similar shape infill to the homogenised grid.

3.9. G-code preparation

In order to reproduce the toolpath directly, the toolpath is converted into G-Code in MATLAB. All motions are directly transferred into the G-Code, resulting in a reflection of the computational model.

The toolpath is translated into G-Code to digitise the coordinates. In G-Code the location of every node is given, the feed rate and the travel speed of the nozzle [49]. Since the coordinates are determined, the G-Code is completed by tuning it to the used machine. Additional instructions are used to update the extruder. The conversion to G-Code is performed in the script `Ultimaker_gcode_witer.m`.

The scalability and repetition of the shape must be important. To visualise the performance of the shape, a larger scale model can be recreated to give an output to the current representation of the grid. Repetitive RVEs can be plotted in the x-, y- and z-direction and spacing between the RVEs can be included if needed.

3.10. FDM and MEX

FDM is performed using an Ultimaker 2+. The RVE can be tessellated to create larger structures by repeating the unit cell in three dimensions. Following the visualisation of the optimised structure, the corresponding G-code is produced for material extrusion. Maintaining constant material deposition rates during the extrusion process is crucial in order to minimise the occurrence of clogging and other extrusion errors. It is imperative to note that, given the printhead traverses varying path lengths in equal time intervals, the feed rate (F) must be dynamically adjusted to ensure uniform material deposition. Variations in thickness can lead to stress concentrations and suboptimal mechanical performance.

3.11. Reproducibility on larger scales

While the current setup prints a single RVE, larger structures can be created by tessellating the grid pattern in the x-, y-, and z-directions. By arranging multiple RVEs adjacently, extended sheets and bulk structures can be fabricated. According to homogenisation theory, these larger assemblies should exhibit mechanical properties equivalent to those predicted from the single RVE analysis.

Creating RVEs that can be scaled is also favourable, currently mostly to visually inspect the extruded result. Size variations also have the benefit that the current designed RVE can be produced using a different type of AM, increasing the precision of the result. Increasing the applicability of the current RVE helps to further expand the use of RVE structures to build larger set-ups, simplifying the method of AM.

3.12. Mechanical testing

Mechanical testing is performed on a Lloyd LR5K to determine the peak force, maximum strain, stress-strain behaviour, Young's modulus and UTS of the optimised structure. ISO 527-4 is followed to perform tensile testing [50]. Tensile testing is selected based on material stiffness—materials with low Young's moduli typically yield more accurate results under compression. Tensile tests are performed at a strain rate of 2 mm/min. Strain was measured as a function of applied external forces to characterise the material's deformation response. For tensile testing, custom clamps were fabricated from

PLA with dimensions matching the test specimens. These clamps distribute gripping forces more uniformly, reducing stress concentrations near the clamping region and preventing premature mechanical failure. Tensile test specimens are created by combining the tensile clamps to the middle part using the `CombineGCodeFiles.m`. The middle part of the specimen is generated as a G-Code and can not be converted to an STL file without loss of the optimised toolpath. Therefore, multiple G-Codes are combined as layers.




Figure 3.4: A computer-generated image of the tensile test object. The tensile test specimen has a length of 120 mm, a width of 21 mm, and the specimen size between clamps is 40×12 mm. 10 layers are added to the specimen, resulting in a thickness of 4 mm.

In figure 3.4, the shape and dimensions of the tensile test object are given. With a depth of 4 mm.

4

Results

4.1. Computational model results

After performing the homogenization method in the methods are compared to common structures with similar shapes. First, a divide between 2.5D layered AM and multidirectional AM is made for comparison. All mechanical properties used to determine the OF are compared in the figures, where a red dot represents the highest OF score in all plots.

4.1.1. Optimisation results of varying voxel size

With an infill density of 73%, the horizontal grid is used to compare the mechanical properties of optimised results to the conventional toolpath. The polar form of the Young's Moduli in E_{11} , E_{22} and G_{12} . In table 4.1, the mechanical properties of the grid with the corresponding highest score are given, while comparing the results to the conventional extrusion methods with similar infill densities.

Conventional methods

Considering the pattern where the printed line is layered in the x-direction and connected at the end to the next line is considered the conventional method of printing. The grid of this horizontal path printing is shown in Figure A.1a. Expected in the horizontal grid structure is that a single direction for the stiffness tensor approximates the stiffness value given by the rule of mixture. In the perpendicular direction, the value is 10 times smaller. High anisotropy is also given with a Zener ratio approaching 0. Due to minimised connectivity between nodes, the value is 0.18.

In the horizontal line grid, a large difference can be seen between ν_{12} and ν_{21} , resulting in poor isotropic behaviour in the horizontal grid and resulting in a low value for the OF.

The spiral path is a conventional method to generate a toolpath for AM. The spiral path has an infill density of 73% and is used to compare the mechanical properties of optimised structures to conventional extrusion methods.

The spiral path shows a more isotropic value, but is lower and approximately the same value as the perpendicular value of the horizontal grid. Having low stiffnesses in both the E_{11} and the E_{22} direction, the spiral pattern combined with a negligible shear modulus, the mechanical properties are reduced, since the material is poorly connected in the centre of the spiral, I

10 × 10 infill density of 73%

A set of grids with voxels sized $l_{voxel} = 2 \cdot r_{nozzle}$ is created to compare the mechanical properties of randomised defect placement with a size equal to the nozzle diameter used in FDM.

The results of the optimisation of the 10 × 10 grids with infill density of 73% using the multiple objectives of the OF are demonstrated in Figure 4.2.




Figure 4.1: Conventional methods recreated as a voxel grid with $l_{voxel} = 0.2$ mm and an infill density of 73%. a-b) The plot of horizontal lines used in conventional FDM and the stiffness plot showing high anisotropic values, where the E_{22} direction has a stiffness near the rule of mixture upper boundary. c-d) The spiral grid used in conventional FDM and the corresponding stiffness plot, where more balanced but lower weakness is seen compared to the original stiffness value.




Figure 4.2: The mechanical properties of randomised 10×10 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.




Figure 4.3: The mechanical properties of randomised 20×20 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.

20 × 20 with 73% infill density

To ensure a fair comparison, the 20×20 grid configuration was chosen, since it matches the voxel size and infill density of the conventional structures. This alignment minimises parameter variation and provides a consistent basis for evaluating optimisation performance. The results of the optimisation of the 20×20 grids with infill density of 73% using the multiple objectives of the OF are demonstrated in Figure 4.3.

40 × 40 infill density 73%

The smallest voxel size used to perform homogenisation is $0.5r_{nozzle}$, where 40×40 voxels are used per grid. Due to the size, only the optimised grids have been saved. All C^H matrices are saved. The results of the optimisation of the 40×40 grids with infill density of 73% using the multiple objectives of the OF are demonstrated in Figure 4.4.

Mechanical properties with variable voxel size

It can be concluded that a single voxel is not representative for each voxel when a voxel length of 0.1 mm is used. As demonstrated in Figure 4.8, a decrease in stiffness is observed for smaller voxel sizes.

The measured stiffnesses E_{11} , E_{22} and G_{12} are shown to decrease with 11% for E_{11} and E_{22} when comparing the median of a 10×10 grid to a 40×40 grid compared to E_0 , for G_{12} a reduced value of 8% measured to G_0 is seen comparing the results of a 10×10 grid to a 40×40 grid.

4.1.2. Examples of highly tunable parameters

The tunability of the grids is shown by comparing the results of grids with similar infill densities and voxel sizes. The mechanical properties can differ highly depending on the placement of defects.

In Figure 4.6, 5 grids are given in the domain of stiffnesses. Figure 4.6 b shows defects align in the y direction, corresponding to a larger E_{22} stiffness. Since not all defects are perfectly aligned in a single direction as is seen in the horizontal lines of Figure 4.1a, there is still a relevant amount of stiffness measured in the E_{11} direction, $0.50 \cdot E_s$. Conversely, Figures 4.6 d and f show defects that form parallel




Figure 4.4: The mechanical properties of randomised 40×40 grids with an infill density of 73%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.




Figure 4.5: Plots comparing the mechanical properties of the 100 highest-scoring results according to the OF of variable voxel size comparison to conventional methods. All solutions have an infill density of 73%. a-c) The normalised stiffnesses E_{11} , E_{22} and G_{12} set up per voxel. d-e) The Poisson ratios compared to the conventional methods. Due to high variance between Poisson ratios, both are shown in the plots. f) The Zener ratio compared to the grid size, with the optimal Zener ratio of 1 given as a dashed line. g) The connectivity value determined using equation 3.5.4.




Figure 4.6: 5 grids chosen due to their differences in stiffness. Keeping the same infill density 73%, voxel size 0.2 mm and boundary conditions, it is still possible to generate grids with highly variable stiffnesses. Each red dot corresponds to the letter of the subfigure.

Grid pattern	Infill Density	E_{11}	E_{22}	G_{12}	ν_{12}	ν_{21}	Zener Ratio	Connectivity
100% density grid	100	3.6	3.6	1.33	0.41	0.41	1.0	1.0
Optimal 10×10	73	1.69	1.64	0.32	0.33	0.35	0.59	0.93
Optimal 20×20	73	1.38	1.65	0.49	0.32	0.32	1.01	0.92
Optimal 40×40	73	1.65	1.44	0.49	0.30	0.30	0.99	0.90
Rule of mixture	73	2.63	2.63	15.28	-	-	-	-
Horizontal lines	73	2.82	0.21	0.77	0.02	0.02	0.04	0.18
Spiral	73	0.29	0.32	0.03	0.002	0.04	0.04	0.65

Table 4.1: Results of the highest OF score for homogenisation and optimisation of a rectangle sized l_x by l_y , with variable voxel size l_{voxel} , including conventional methods for comparison.

to the x-direction. Additionally, the grid that aligns the dashed line with examples c and e shows no defects in parallel directions, but shows defects attached in both x- and y-directions.

In Figure 4.7, the variance in connectivity is shown using 5 example grids, with highly variable OF scores. Little visual information can be retrieved from the grids compared to the connectivity scores, since the connectivity value is a summation of all connectivity values in the grid.

4.1.3. Optimisation results of voxel sizing

In Figure 4.8, the highest score received in the OF is given with the corresponding stiffness plot in 2D. The values given in Table 4.1 correspond to the given RVEs.

In table 4.2, the optimal result of the highest score is observed for all infill densities and their corresponding grids.

The material properties of the bulk material are highly variable. The random defect location has shown material properties, normalised to the bulk stiffness, with highly variable outcomes. Grids with infill that is similar to the horizontal line grid show high stiffnesses towards a single direction, whereas a randomly distributed infill shows comparable stiffnesses in the E_{11} and E_{22} direction, reduced compared




Figure 4.7: Examples of the influence placement of defects on the connectivity compared over the OF score for a 20×20 grid with infill density of 73%. Every red note indicates a grid with the corresponding subfigure attached to it.




Figure 4.8: The plots with the highest score based of the OF for each voxel size considered. Infill density is set at 73%. a-b) The optimal grid with 10×10 voxels with the corresponding stiffness plot in 2D. c-d) The optimal grid for a 20×20 grid with the corresponding stiffness plot. e-f) The optimal grid for 40×40 voxels with the corresponding stiffness plot in 2D.

to a single high performing shape as shown in the horizontal line structure. Due to the similar shape build-up, it can be concluded that for infill between 0.93 and 0.63% the Poisson ratios ν_{12} and ν_{21} are nearly identical and decrease with lower infill density. Conversely, the standard deviation for lower infill densities increases for the stiffness, Poisson ratio and the Zener ratio.

Combined with the higher anisotropy, the Zener ratio and connectivity value are reduced heavily for lower infill densities. Higher infill densities score for both the Zener ratio A and C near the optimal value 1. With a visible reduction when increasing the 0.63% infill.

The Poisson ratios are separated to compare with the OF and show the high variance between ν_{12} and ν_{21} in the conventional methods. After determining the similarities between ν_{12} and ν_{21} for the randomised grids, the Poisson ratios are considered as a single value and written as ν . This value is used to compare the other objective functions. Notably, compared to the conventional method is that the conventional method does not have a Poisson ratio that is nearly identical, and therefore, these values are separated in further research.

A clear decrease in stiffness is found when the infill density is decreased. Lower infill densities tend to have lower shear moduli, with the largest decline to G_0 seen. Using lower infill densities in the grid increases the run time. Since one defect is placed per turn, an increased number of randomly placed defects has a larger run time.

4.1.4. Varying infill densities

By varying the infill density between 93 and 63% over the 20×20 grid, the impact of defect distribution on the objective functions can be determined. 4 optimisations are performed with the only variable changing being the infill density with increasing values of 0.10 written as [63, 73, 83, 93] % first, the feasibility is checked over all solutions. Next, the maximal stiffness tensor is compared, followed by the anisotropy, connectivity and Poisson distribution. During the optimisation process, 4 optimisations are performed with different infill densities. Using a standard voxel size of 0.2 mm, the infill density varies between 93%, 83%, 73% and 63%. All values that are selected during the sequential selection method are encircled, where each variable must have a larger total OF value or outperform the current OF on a single value. These values are encircled in black.

20 × 20 with 93% infill density

In Figure 4.9, the mechanical properties for the 20×20 grid with an infill density of 93% corresponding to the OF are plotted for all values with stiffnesses E_{11} and E_{22} larger than the mean value.

Due to the low number of defects in the RVE structure, barely any losses are noticed, with maximal stiffnesses in a single direction approximating $0.97E_s$ in Figure 4.9. The connectivity stands out, due to the limited losses in connectivity, the results are layered in Figure 4.9I where the maximal value for connectivity contains almost all selected values in the sequential optimisation criterion shown with almost maximal value of connectivity of 0.93.

20 × 20 83% infill density

In Figure 4.10, the mechanical properties of the 20×20 grid with an infill density of 83% are shown. To categorise the performance, the results are coloured depending on the performance of multiple characteristics that are part of the OF. The dashed line indicates an isometric property, where the value is equal to the other value. The circled values follow the optimisation method, improving sequentially compared to its predecessor, either partially or on every function in the OF.

20 × 20 63% infill density

In Figure 4.11, the mechanical properties of all 20×20 grids with an infill density of 63% are shown with stiffnesses larger than the mean value of both E_{11} and E_{22} .

Variable infill density grid overall statistics

A selection of variable performing grids can be shown to compare the mechanical properties with a constant infill, while varying its location. Most notable is the large variation in mechanical stiffness. Clustering defects and therefore increasing the area of a defect are shown to reduce stiffnesses in directions, since fewer voxels are connected to redistribute stress and strain. Locally increased defects




Figure 4.9: The mechanical properties of randomised 20×20 grids with an infill density of 93%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.




Figure 4.10: The mechanical properties of randomised 20×20 grids with an infill density of 83%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.




Figure 4.11: The mechanical properties of randomised 20×20 grids with an infill density of 63%. a) E_{11}/E_s over E_{22}/E_s indexed by the connectivity score. b) E_{11}/E_s over E_{22}/E_s indexed by the Zener ratio. c) E_{11}/E_s over E_{22}/E_s indexed by the Shear modulus G_{12} . d) E_{11}/E_s over E_{22}/E_s indexed by the score of the OF. e) Poisson ratio ν_{12} over ν_{21} indexed by the connectivity score. f) Poisson ratio ν_{12} over ν_{21} over the Zener ratio. g) Poisson ratio ν_{12} over ν_{21} over the Shear Modulus. h) Poisson ratio ν_{12} over ν_{21} indexed by the OF. i) ν over E_{11}/E_s indexed by the OF. j) ν over E_{22}/E_s indexed by the OF. k) ν over Zener ratio indexed by the OF. l) ν over the connectivity indexed by the OF.

Infill Density	E_{11}	E_{22}	G_{12}	ν_{12}	ν_{21}	Zener Ratio	Connectivity
100	3.6	3.6	1.33	0.41	0.41	1	1
93	3.37	3.38	1.11	0.39	0.39	0.95	0.99
83	2.59	2.45	0.85	0.35	0.35	0.92	0.91
73	1.38	1.62	0.49	0.32	0.32	1.01	0.93
63	0.982	0.74	0.24	0.23	0.23	0.62	0.51

Table 4.2: Results of grids with the highest-scoring values in the OF and their corresponding mechanical properties after optimisation with a variable infill density, while remaining the size of the voxels constant at 20×20 with $l_x = 0.2$ mm.




Figure 4.12: Plots showing the impact of different infill densities on the mechanical properties of the grid using conventional methods as comparison, having a constant voxel size of 0.2 mm. The 100 highest-scoring grids according to the OF are used in this method to show the optimal performance. a-c) The normalised stiffnesses E_{11} , E_{22} and G_{12} set up per voxel. d-e) The Poisson ratios compared to the conventional methods. Due to high variance between Poisson ratios, both are shown in the plots. f) The Zener ratio compared to the grid size, with the optimal Zener ratio of 1 given as a dashed line. g) The connectivity value determined using equation 3.5.4




Figure 4.13: Figure of the optimal grids and the corresponding stiffness plot in 2D with variable infill densities. a-b) The optimal grid for an infill density of 93%. c-d) The optimal grid for an infill density of 83%. e-f) The optimal grid for an infill density of 73%. g-h) The optimal grid for an infill density of 63%.

show reduced mechanical stiffness compared to the material properties expected. Comparing the mechanical properties to an optimal set-up similar to the rule of mixture for single directions helps to estimate the trade-off being pursued for a single and multiple directions. It is important to always include the mechanical properties for each state.

In comparison, high stiffnesses in a single direction are seen when a grid is randomly generated, where the defective voxels are more commonly located in line structures.

It is important that in shapes with high shear moduli, the voxels are seen to connect also more often diagonally, less commonly seen in conventional methods, where edge-to-edge contact of a voxel, similar to a Neumann neighbourhood. In the horizontal line structure, the placement of material voxels is commonly interconnected to 5 other voxels, only varying at corners close to the edge of the RVE.

It is noteworthy that as the voxel size is increased, i.e. $l_{\text{voxel}} = 2 \cdot r_{\text{nozzle}}$, there is an observed decline in connectivity, with the formation of defects and the emergence of larger gaps. This has a detrimental effect on the average defect size created. This phenomenon transpires due to the infill density being calculated over the total material, as opposed to the inner material, within which the defects can be positioned. In the context of 10×10 grids, a total of 64 voxels are considered eligible to be designated as a defect. This results in 37 blocks being formed at the centre, with each block containing voxels with a value of 1. In the context of a 20×20 grid, the utilisation of 324 voxels is permissible, with 108 voxels designated as defects.

It is also worthy of note that defects have a tendency to cluster more readily in a 10×10 grid. This results in defects with a larger surface area, which consequently leads to greater permeability for objects of larger size. In grids with a reduced number of voxels, defects have a tendency to cluster with directly adjacent voxels. Given the reduced number of options for the even distribution of defects over voxels, there is an increased probability of adjacent voxels with larger voxel sizes.

It has been demonstrated that lower anisotropic values are obtained when the optimal shape is determined using larger voxel sizes. The utilisation of smaller grid sizes has been demonstrated to result in the formation of larger connected defect areas. Figure 4.2 is notable for the presence of a cluster of defects in the upper portion of the grid. The scarcity of material present on the surface gives rise to a structure that is reminiscent of a lattice. It has been determined that, due to the defects having the same diameter as the nozzle, there is a reduced probability of the toolpath travelling over the defects or replacing the centre. Consequently, the shape remains analogous to the optimised grid retrieved from the computational model.

Conversely, smaller defect sizes can create difficulties in the creation of toolpaths. Given that the nozzle is 4 times larger than a single voxel, partial defect overlap is observed to occur more frequently than would be expected. It is evident that there is a low level of similarity in comparison to the optimised grid, which can be attributed to the challenge of traversing defects. The effects of homogenisation to optimising a randomised RVE are lost, in the event that the nozzle is not able to avoid defects, resulting in poor shape quality of the toolpath on a 40×40 grid.

During the process of homogenisation, which involved grids with a 63% infill density, there was an occurrence of computational problems. The analysis revealed two data points that yielded results that were inconsistent with the original E_0 value, indicating an overestimation of the stiffness in a single direction. It has been demonstrated that inadequate outcomes in the computational process are exclusively observed in the grid with a lower infill density. It is evident that the random placement of defects within the grid results in a reduction in the feasibility of outputs, given that not all voxels can be periodically bounded.

4.1.5. Comparison of optimised to conventional structures

Considering the conventional methods simulated in Figure 4.1, the homogenised grids show mechanical properties that show higher isotropy, the Zener ratio is approaching 1 for the randomised grids, especially when considering the OF values, whereas the conventional methods have a Zener ratio smaller than 0.1.

Additionally, the Poisson distribution that originally has been separated into 2 variables, ν_{12} and ν_{21} , due to the high inequality between the ratios, is comparable in each method of homogenisation. The

values are approximated to be equal due to the OF favouring isotropic values.

Next, the distribution of defects in the grid is improved, large clusters of infill densities are penalised and as a result. The random locations of defects show that, parallel to a direction shaped clusters of defects have a closer resemblance to the conventional method of horizontal lines. Having no visible direct alignment to a stiffness direction shows to increase the performance according to the OF.

The horizontal grid and spiral grid are used to compare the optimised shapes. Since the infill densities and material properties are equal in the optimisation sets, the objective functions can be compared and treated as similar.

Compared to the conventional methods, the RVE created using the modified toolpath has a longer toolpath of 70.3 mm. The horizontal lines have a toolpath length of 29.4 mm, and the spiral path has a path length of 34.4 mm. An increase of 239% is seen compared to the horizontal grid, and 204% is seen compared to the spiral grid. The increase in path length can be explained by the increased nodal visits seen predominantly in the right top corner of the RVE in 4.15i and j.

For the horizontal grid, it is most noticeable that the stiffness tensor has a single large component in the E_{11} direction, where it approximates the maximal value given by the rule of mixture. On the other hand, the other variables E_{22} and G_{12} are in the range of 10 times smaller than the optimised values. For a single layer, this proves that the material is weak and will break apart easily if loaded in a transverse direction. To reduce the easy to break setup in non-optimal directions, the layers can be rotated compared to another, by either rotating every layer with 45° or 90° it is possible to reduce anisotropic behaviour in the bulk material. The spiral grid is shown to have a lower stiffness tensor compared to the optimised shapes for the three main directions, E_{11} , E_{22} and G_{12} . Due to the limited connection between voxels, between 3 and 7, mostly 5, it shows lower interaction, and therefore, a single layer of conventionally generated toolpaths is low performing under shear stress. An improved distribution of defects in the optimised grid shows that the connectivity of the average voxel increases. Therefore, the defect size of a clustered defect has decreased. The total defect size of a horizontal grid has a size of 1×18 voxels, an area of $3.6mm^2$. A total of 6 defects exist in the grid of 20×20 blocks. For the spiral grid, the defect size is a continuous shape of 1×108 voxels, an area of $21.6mm^2$, showing low diagonal connections in the spiral toolpath. Compared to the horizontal lines grid, all grids have a lower stiffness in a single direction, while the stiffnesses E_{11} and E_{22} are in the same range, compared to a factor 13 difference between E_{11} and E_{22} . Compared to the conventional toolpaths, optimised toolpaths show a 1000% increase in the shear modulus G_{12} .

For the 20×20 optimised grid with an infill density of 73%, a set with optimal shapes is determined to choose the distribution of defects that function as a defect. Tunability of the mechanical properties is heavily reliant on the defect distribution, seen in the 20×20 RVEs with an infill density of 73%, local directional alignment of the defects bears resemblance to the conventional grid distribution when comparing 4.6d and f and Figure 4.1a.

All shapes can be recreated as RVEs, and therefore, the scalability is ideal. The tunability is where the RL-based toolpath stands out. Here, the input parameters for optimal behaviour can be broadened to tune the material for ideal mechanical properties. Where conventional methods are limited due to their programmed shape, RL-based toolpaths are easier to recreate. Due to the fact that all RL-based toolpaths are created in MATLAB, some extra knowledge is needed to recreate shapes using this type of structure.

4.2. RL path planning

RL path planning is shown to produce results within a runtime of 2 hours. Using 17 input parameters results for both the generalised and specialised RL-based toolpath. After testing the training data 100 times per training agent, the results are compared between the generalised and specialised methods. Notably is that the specialised agent outperforms the generalised agent individually based on path length and nodal visits, but having a more robust set-up, the general performance, the average value is nearly identical after 100 trials.

By dividing the RL function into 2 sets for learning first, a generalised path following a learning process, where only a specialised grid is used for toolpath planning.

First, a 2000 episode training set is used to generate knowledge for randomly selected grids to optimise the toolpath planning, where defects are circumnavigated with different dimensions. A larger number of steps per episode is required to travel grids, due to the increased size of a 40×40 grid. After performing the first DDQN training set a specialised training is started. 1500 episodes occur during specialised training with a DDQN as agent.

The specialised method has a reduced MaxEpisode value of 1500, 500 fewer episodes to explore the grid. The learning rate of the specialised toolpath has been lowered, but starts with a higher expected reward gained from the generalised toolpath agent; a longer path can be created, improving the nodal visits heavily.

Partially combining the results of the coverage with the efficiency, it is shown that the coverage increase rate is declining when the total coverage is reached, nearing the final steps. Revisits of nodes become more apparent, and low nozzle awareness would result in double material extrusion or spreading of the material. As a result, the shape fidelity would be decreased, negatively impacting the physical result.

Favouring rewards by increasing the optimal value when pursued creates a highly tunable RL tool. Noticing that the coverage and path length are dependent, but also assure nodal revisits or crossing edges are common in earlier trials, but are removed after low output rates in FDM.

The results of the Episode reward for both generalised and specialised training are given in B. The final Q score after the generalised training process is 4810.40, and an average reward of 5206.01. For specialised training, the optimal scoring 20×20 grid with an infill density of 73% is chosen since the Stiffnesses in multiple directions are already calculated. The index number is 276883. The Q-score after the specialised training process is 8394.44, and the average reward of 10333.63.

Seen in the Appendix B, both specialised and generalised training show periodic drops in the reward after a certain number of episodes. The dead-end termination is included, where a path is shut down if no more nodes are accessible. Having larger penalties on a dead-end and missing out on coverage rewards reduces the value heavily, and encourages exploration of new paths.

4.2.1. Reward optimisation

The trade-off between giving rewards is heavily influencing the outcome of the RL tool. Having the main focus on including as many nodes as possible in the toolpath results in a stepwise increase in coverage. Still limiting the maximal coverage to the number of available nodes, where defect nodes should still be avoided as much as possible, is a must. The upper boundary for nodal coverage is set at 100% of available nodes, and percentages higher should be excluded.

The coverage is seen to increase initially at a constant rate, similar to the increase of path length, but near the 60-70% range, the nodal coverage starts to increase more slowly, since nodal revisits are needed to increase coverage in non-Eulerian paths. The adjacent nodes are dynamically removed to reduce revisits, but are included when a lack of options is available. Therefore, an increased path length is always supported. Hard constraints such as edge traversal of the RVE are also never overwritten, e.g. the environment is created to reduce the effect of any travel outside of the RVE. This is a guarantee for the reproducibility of the RVE in the x- and y-directions.

Conversely, soft constraints are allowed to be overruled. For instance, nodal revisits are allowed but penalised and local over-extrusion is discouraged. Nodal revisits help to improve total nodal coverage, but should be limited due to unnecessary motions included in the toolpath. Also to note, revisiting edges between nodes is still heavily penalised, since this will assure either local over extrusion, spreading of material, and low shape fidelity. Edge revisits are converted into travel motions in later iterations of the RL tool to prevent local over-extrusion. This helps to prevent the goals from remaining close to the original rate of infill.

Originally highly penalised, but reduced in the process, is the sharp turn penalty. While it is less desired to have a high rate of directional changes in the toolpath, it can be allowed to increase the coverage and reduce harsher penalties as edge revisit. Since most nodes have a higher adjacency than a limited value of 2, the sharp turns are not necessarily needed, yet are not desired. Local high coverage is always supported to ensure high coverage rates, yet the sparsity coverage bonus is included to optimise even distribution of nodal visits. This is an optimal output for having an even visit of nodes,




Figure 4.14: Comparison bar plots showing the 100 test results of the generalised, specialised and manually modified results of the RL toolpath planning method. a) Total coverage of the grids after applying the generalised, specialised and manually modified methods. b-e) Nodal coverage, coverage distribution and path length are indicators the efficiency and are plotted over all 100 test results per method.

while rewarding exploration over exploitation. As a result, difficult to reach areas as poorly connected voxels with low adjacency, edges and vertex nodes, are accessed more often.

Considering the average path score of the generalised and the specialised method of RL, the median path score is set to an equal, yet high outliers approaching 1 are seen in both methods. Since both coverage and path length play high roles in increasing the path score, these factors play a strong role. Yet, large path length can be the result of nodal revisits, edge crossing and local increases of material extrusion. The penalties given out for these results help to reduce the maximum length of a path to create more efficient paths.

4.2.2. Test results of RL

In Figure 4.14, the test results of the generalised, specialised and manually modified methods are presented after performing 100 tests. The total coverage is the result of complete coverage of the grid, and having an infill density of 73%, the total coverage exceeds this number since some defects are also visited during path planning. In 4.14e the path length is plotted over the coverage. This helps to indicate the total path performance, since high coverage can be the result of low efficiency due to having a heavily increased path length. Additionally, due to the dead-end statement, low coverage can be presented as high performance, since revisits occur more often when high nodal coverage is achieved.

In Figure 4.15, all RL-generated toolpaths are shown using the generalised, specialised and manually modified method. The generalised method has a training set built up of multiple grids containing different infill densities and all voxel sizes, where the specialised training set only uses the highest scoring grid for a 20×20 grid with an infill density of 73%. The manually modified toolpath is created by adding two extra nodes to the toolpath manually, and increasing the nodal visits, while circumventing defects. All methods contain the visualisation using the grid as background in subfigures 4.15 a,e and i, the visualisation of supervoxels that are scaled to the nozzle diameter in subfigures 4.15 b, f and i, show




Figure 4.15: An overview of the RL-generated toolpaths. In the first row, the generalised toolpath is shown, in the second row, the specialised toolpath is shown, and in the third row, the manually modified toolpath is shown. All grids are a combined set of the toolpath in the given index plot, the super-voxel representation, the nodal visits per location and the coverage progress while completing the toolpath. The highest scoring grid of 20×20 with infill density 73% is used, but the image is mirrored in the y-axis.

the nodal visits per supervoxel in subfigures 4.15 c, g and k and show the coverage over the path length in subfigures 4.15 d, h and l. The nodal visit heatmap lacks information in row and column 10, since the nodes are placed in the centre of the supervoxel, reducing the voxels to 9×9 . Including the outer radius of the nozzle, the size of the grid is equivalent to the originally planned 10 voxels, sized 4×4 mm.

The heatmap helps to visualise the visits per location, showing possible over- and under-extrusion that can be removed in post-processing the G-code to include travel moves.

In Figure 4.14, output values after the test function is performed are depicted for the generalised and specialised method, including the manually modified toolpath. The generalised method has $N = 100$ trials, and the specialised method has $N = 100$ trials. Low path coverage, with only unique visits to nodes, equates to high efficiency and therefore the parameters efficiency and coverage are compared in the scatter plot in Figure 4.14e.

4.3. Material extrusion

The material extrusion process is performed after creating the toolpath and gives a physical output of the results.

Using PLA as the material for FDM, it is seen that the first layer is poorly adhering to the baseplate. With a material flow at 100%, optimal adherence is seen to the baseplate. The material extrusion is determined to be equal to the volume of filament being used between two nodes. Using PLA filament with a filament radius of 2.85 mm.

The optimal results of a successful production of RVE structures (Figure 4.16a-d). A block of 10×10




Figure 4.16: FDM produced RVEs. a-b) length, width, thickness of a 10x10x10 RVE block sized 40.5x40.5x4.1 mm. d) Tensile test specimen, high printing quality. e) Tensile test specimen, low printing quality caused by poor adherence. f) A scale model of RVE structures. A single layer of thickness 0.4 mm scaled 5.0 times larger than the original design.

RVEs is created that has 10 layers. A successfully produced tensile test specimen with a thickness of 10 layers is produced (Figure 4.16d). A poorly designed tensile test structure with low print quality shows low adherence to the build plate, resulting in a useless specimen (Figure 4.16e).

A scale model of the RVE structure is created (Figure 4.16f), where the toolpath is visualised. The model is 5 times increased and has a line thickness of 0.4 mm. The toolpath can be recreated using extrusion, and it is shown that the toolpath can be replicated multiple times.

4.3.1. Optimised space compared to extruded grids

The test subject is a PLA MEX-produced block of $40 \times 40 \times 6.4$ mm. The test object is built up of $10 \times 10 \times 16$ layers using an Ultimaker 2+.

The material is extruded with high values over small distances, during turns while extruding blobs can form, since the nozzle remains in the same location for a prolonged period. This becomes apparent in the sharp turns. Stringing occurs due to limited contraction of the filament during travel motions between RVEs.

The size of a 10x10x10 RVE is projected as 39.9x39.9x4.0 mm. The extruded shape has 40.5x40.5x4.1 mm. To measure the infill density, the weight of the final result can be compared to the expected weight for a material with a 100% infill. A 100% infill object with a density for PLA of 1.25 g/cm^3 is expected to weigh 8 grams. With a weight of 6.7 grams for the FDM-produced RVE, an infill density of 79.7% is estimated. A deviation of 6.7% is seen compared to the planned infill. Nodal revisits and partial defect coverage can be appointed as causes of the increased rate of material.

Animations and video footage of the AM process have been attached to the additional material folder.




Figure 4.17: 2 specimen after tensile testing. a-b) specimen Y_1 after tensile testing and a detailed image showing the partial fracture through the RVE. c-d) Specimen X_2 after tensile testing and a detailed image showing a clean cut between the RVEs.

Sample	Peak load (N)	Young's Modulus (MPa)	UTS (MPa)	Max. extension (mm)	Strain(-)
X_1	819.80	265.43	23.40	2.58	0.0664
X_2	775.99	261.58	22.16	2.51	0.0649
X_3	710.55	311.88	20.28	2.16	0.0556
X_4	803.22	305.80	22.92	2.79	0.0720
Y_1	751.24	297.15	21.44	2.71	0.0696
Y_2	632.93	257.77	18.09	2.39	0.0609
Y_3	603.60	240.80	17.52	2.35	0.0604

Table 4.3: Mechanical properties retrieved from tensile testing of 7 specimens in two directions.

4.3.2. Adjusting printing properties

Optimising the FDM printing properties helps to improve the material flow to increase the printability. After tuning the optimisation and toolpath planning process, the adjustment of printing properties functions as manual tweaking.

Since the layer extruded on top of the build plate only experiences poor adhesion, other layers can be built using a lower material flow and have a closer proximity to be produced in a similar fashion to the result of the planned toolpath.

The FDM process of a tensile test specimen is estimated to take an hour; a $10 \times 10 \times 10$ RVE (sized $40 \times 40 \times 40$ mm) is printed in 2 hours. Since some voxels are revisited over 6 times (Figure 4.15), some areas experience over-extrusion. The local loss of defects increases the infill density. By removing extra extrusion motions, the target infill density can be achieved. Additionally, by lowering the extrusion speed to 1200 mm/min, the materials adhere more optimally to the build plate.

4.4. Mechanical testing

During tensile testing, 7 samples are used to determine the mechanical properties of the optimised structure (E_{11} oriented $N = 4$ and E_{22} oriented $N = 3$). The stress-strain curve shows similar results in tensile tests in each direction (Figure 4.18). Two testing directions are used, since the material is presumed to be anisotropic. The Mechanical properties retrieved from each tensile test are given in Table 4.3. It is noticeable that each specimen broke where RVEs connect. In specimen Y_1 , a partial fracture of an RVE (Figure 4.17a and b), which corresponds with an increased stress-strain relationship experiencing higher stresses on the specimen. Clean fractures that broke between RVEs (Figure 4.17c and d) show lower stress-strain results and a reduced Young's modulus.

Additionally, the average Young's modulus for E_{22} oriented tensile specimens is 286.25 MPa. The




Figure 4.18: Tensile test results given in strain over stress performed in two dimensions with the added Young's Modulus per sample attached in the legend. a) E_{11} oriented tensile test specimens corresponding to Y-numbered specimen during tensile testing. b) E_{22} oriented tensile test specimens corresponding to X-numbered specimen during tensile testing.

average Young's modulus for E_{11} oriented tensile specimens is 265.33 MPa. Similar to the computational results seen in Table 4.2 for an infill density of 73% corresponding to a computational E ratio of $E_{11} = 0.85E_{22}$, the test result E ratio $E_{11} = 0.93E_{22}$. However, the tensile test results show Young's modulus reduced tenfold compared to the computational model.

5

Discussion

5.1. Performance Assessment of the Homogenisation Method

The optimisation and homogenisation process has yielded notable features in the results. The discussion focuses on the utilisation of key values identified in the results.

5.1.1. General remarks

The optimisation process shows that the largest variability in mechanical properties can be found in RVEs with lower infill densities. Due to the larger number of defects, not all material-filled voxels can be fully connected. Additionally, due to the random placement of defects of the grid, most grids are filtered out for not being periodically bounded. Tunable mechanical properties can be seen in all infill densities and voxel sizes, yet a higher variety is seen in lower infill densities. For lower infill densities, it is expected to have lower output performance and random shape optimisation is not recommended due to its low efficiency. Creating lattice-like structures is recommended having low infill densities.

5.1.2. Representativeness of the voxel infill

A voxel should adequately represent the material characteristics at any point within the structure. However, due to voxel sizing, the average material infill within a single voxel can vary, potentially affecting the representativeness of the overall RVE. In the current method, three voxel sizes are considered, each scaled according to the nozzle diameter.

For a 20×20 grid, the voxel dimensions are equivalent to the nozzle radius, i.e. $l_x = l_y = l_z = r_{\text{nozzle}}$. The infill ratio of a single voxel is then given by $\rho_{\text{voxel}} = \frac{0.25\pi r_{\text{nozzle}}^2}{r_{\text{nozzle}}^2} = 0.25\pi$. In this context, both the 10×10 and 20×20 grids possess voxel dimensions that can be considered representative of the complete RVE.

Reducing the voxel length to $l_x = l_y = l_z = 0.5 r_{\text{nozzle}}$ introduces three possible values for the effective area of a voxel. The corresponding infill ratios are expressed in Equation 5.1.

$$\rho_{\text{voxel}} = \begin{cases} 1 & \text{for } [0, r/2] \times [0, r/2] \\ -1 + \frac{\sqrt{3}}{2} + \frac{\pi}{3} & \text{for } [0, r/2] \times [r/2, r] \\ 1 - \frac{\sqrt{3}}{2} + \frac{\pi}{3} & \text{for } [r/2, r] \times [r/2, r] \end{cases} \quad (5.1)$$

It follows that a single voxel of size $l_{\text{voxel}} = 0.5 r_{\text{nozzle}}$ cannot be fully representative of any portion of the grid. Nevertheless, an approximate representation can be achieved through voxel reconstruction.

While smaller voxels generally lead to improved homogenisation accuracy due to finer discretisation, they also result in substantially increased computational cost. The number of randomly generated voxels increases quadratically with grid refinement, directly scaling the runtime. According to the 2D homogenisation method described by [6], each voxel possesses eight DOFs, denoted by χ , which are

required to enforce periodic boundary conditions. Consequently, the quadratic increase in voxel count results in a quadratic increase in DOF, thereby significantly extending the computational runtime.

5.1.3. Plane stress and plane strain

It is important to note that the accuracy of the approximation is reduced, as evidenced by the observed change in the original stiffness. The maximal stiffness, determined to be 58.99 GPa, contrasts with the original output of 54 GPa. The discrepancy between the outputs for values that differ by 9.18% is indicative of a considerable degree of uncertainty.

The rule of mixture is determined using the following equation, which is derived from plane-stress methods. It is therefore evident that the horizontal method is an outlier of all results. It can thus be concluded that the outcome, in which the horizontal zigzag grid is presented in the results, can be explained on the basis that the in-plane stress of the material can be reduced and then falls within the range of the rule of mixture when considering $E_{defect} = 0$ MPa. The height of the voxel is compared to its length and width, with the length and width being expressed as $l_x = l_y = 20dz$, where $dz = r_{nozzle}$ for a cubic-shaped voxel. In accordance with the established principle, the consideration of plane-stress is mandated when the object's length is approximately 10 times greater than its thickness. It is evident that the requirement is being met in the current configuration. Consequently, this simplification can be applied. The material is considered to be of a sufficiently low thickness to be regarded as thin, and consequently, plane stress is employed.

As demonstrated in Figure 4.8, for an infill density of 93%, the average stiffness values exceed the rule of mixture. The utilisation of plane-stress instead of plane-strain enables the attainment of higher stiffness values, a principle that is governed by the rule of mixture. In selecting the plane-strain method, it is anticipated that the stiffness of high infill densities will fall below the threshold of the rule of mixture.

5.1.4. Removing heuristics in grid generating

The current method of homogenisation is based on heuristics; randomly generated RVEs are compared to determine an optimal shape. As a result, the optimum value can only be approached, where using a method of machine learning could benefit this process. Additionally, after creating the grids, optimisation was performed, and a lack of a converging value in the current code restricted the optimisation from early termination. Using a method that is aware of an upper limit and prioritises convergence would be recommended in future work. Similar to the toolpath planning method, using methods that learn from prior iterations can help to improve the overall structure and reduce the manual workload. Heuristics was currently chosen as a method to create large datasets, but in future work, this large set can again be reused, similar to the training sets for generalised and specialised RL training. Using a large quantity of similar training data can help to improve the shape of the RVE, optimising the placement of defects on a better scale.

5.2. Reinforcement learning and toolpath planning

The use of RL in path planning has a notable outcome. By constraining parameters, the final result can have minimal output, whereas soft boundaries are shown to improve coverage. Next, there is a connection to total surface coverage and the path efficiency. Finally, the generalised method and the specialised method are compared to determine the use of 2 training sets for optimal output.

5.2.1. Defect circumventing

During path planning, defects persist within the original voxel size. Conversely, the nozzle diameter is constructed from a combined set of voxels, i.e. the supervoxel. The process of reconstruction involves the combination of smaller voxels to form a supervoxel, thereby representing the average voxel value. In the event of a supervoxel being assigned a value of 0, it is to be excluded from the adjacency matrix. This is particularly advantageous for lower infill ratios, as it results in a reduction of the computational cost of the adjacency matrix. Subsequently, each voxel is analysed independently in order to minimise travel over defects during the process of path planning. The introduction of two parameters, which are cognisant of the nozzle diameter and which serve to minimise the travel of the toolpath over defects, enables the creation of a toolpath that is capable of traversing a larger object while circumventing smaller defects, thereby minimising defect coverage. It is evident that defect awareness is promoted in

the context of travel. Nevertheless, the implementation of complete defect avoidance is challenging, as evidenced by the fact that the total coverage exceeds the upper boundary for grid cover. It is permissible to mitigate by travelling partially over voxels classified as defects. However, this is penalised in order to ensure shape similarities to the homogenised grid. Furthermore, local minima in the form of partial defect coverage are permitted, provided that the total reward is increased.

5.2.2. Bridging

Some voxels are poorly connected to material voxels and reduce the overall performance of the toolpath. This problem is resolved by allowing a bridging motion between poorly accessible voxels. In figure 4.15i and j near the area [4,3] to [6,3], bridging occurs between larger voids. An exception is made to reduce the penalty of visiting defects, to improve the overall reward of the toolpath. By allowing bridging motions, no local minimum occurs, trapping the toolpath solely to toolpaths with minimal path length and still encourages minimising defect traversal. Bridging encourages impossible travel motions and is only specialised path planning. During bridging material extrusion is allowed, in contrast to standard travel motions where a longer distance between nodes is travelled than $\sqrt{2} \cdot d_{nozzle}$. Bridging can be regarded as an illustration of local minima not being outweighed by global optimal rewards. The utilisation of soft constraints has been demonstrated to enhance the overall performance of RL, thereby enabling the exploration of larger domains for other parameters.

5.2.3. Coverage and efficiency trade-off

Using the DDQN tool, it becomes apparent that there is a trade-off between coverage and efficiency. Every unique node that is visited increases the level of coverage. To get near-perfect visits in a non-Eulerian graph, nodes must be revisited to increase the coverage, reducing the efficiency. Similarly, high efficiency can be caused by a limited number of nodes being visited, and therefore, the shortest path analogy is used. The result is a toolpath of minimal length, with low coverage, while having a high efficiency rate. Coverage is a result of the exploration factor, as it allows an agent to improve the current knowledge about any action, leading to long-term benefit. High exploration is linked to lower exploitation, where the highest reward is pursued using the agent's current action-value estimates. Increasing the discount factor γ to approximately 1 helps to increase the level of exploration in the DDQN.

Following a near-Eulerian path while accessing all nodes gives problems with the coverage aspect. The maximal coverage rate is dependent on the efficiency. Visiting nodes multiple times is undesirable, but to achieve near complete coverage for all accessible nodes, while meeting the requirements to travel only to nearby nodes. On the other hand, complete coverage results in low efficiency as seen in prior trials. The total path length travelled during printing is while visiting close to 100% is heavily increased. A large penalty is placed on dead-end visits to reduce the runtime of poor-performing toolpaths. By terminating poor performance early, the mean value as shown in Figure 4.14e. is reduced in both stages, due to an upper boundary. To minimise the effects of the dead-end penalty, multiple grids are created during the testing phase to reduce suboptimal.

To determine an optimal toolpath, multiple factors are combined to evaluate its performance. Solely focusing on high coverage results in a lower efficiency, and including a factor for efficiency in the final test result shows an improved toolpath with a reduced number of redundant motions. Efficiency is already included in the RL reward function and is added to the test score as a small factor to prevent biased results towards efficiency.

$$S_{test} = 0.6 \cdot C_{area} + 0.35 \cdot C_{node} + 0.05 \cdot \eta \quad (5.2)$$

Where S_{test} is the final score after a test, C_{area} is the area coverage, C_{node} is the nodal coverage, and η is the efficiency determined by dividing the number of nodes visited by the total number of nodes visited. The weights were empirically chosen to balance coverage quality and motion efficiency.

5.2.4. Generalised and specialised methods

Using training data on a single set of grids, with the outcome to have an optimised toolpath for a single grid, is seen as a specialised method of using RL. For a single grid, a near-complete approximation is approached. The coverage and efficiency can be optimised; however, the specialised RL functions

show low robustness and fail if different grids are evaluated on them. Therefore, it is highly recommended for a single grid to be optimised using the specialised RL method. The generalised method shows to favour long distance nodal visits to increase total coverage, resulting in locally missed nodes. In figure 4.15a, large jumps are used to cover the grid maximally; however, the result would be translated into high rates over travel moves, where no extrusion occurs, minimising the quality of the FDM produced result.

On the other hand, generalised RL methods train using a set of grids, with more variation. By exposing the DDQN to multiple grids, a more robust outcome that is applicable to multiple grids is formed. For new information outside of the training data, a generalised RL function shows improved outcome.

The current goal of using RL in this report focuses on generating a toolpath near-optimal for multiple grids and, therefore, a generalised method is desired. Having a nodal 71% of supervoxels after training occurs often, and certain parts of the grid are less visited. After testing the results using multiple grids in the training data show varying success rates, i.e. lower coverage, low efficiency and multiple dead-ends, since the focus is broadened to perform as good as possible in a more diverse grid. The specialised method is able to cover 93% of the supervoxels and has no travel moves applied during its performance. Additionally, the manually modified version of the current toolpath avoids some defects and increases the nodal visits by 2. In general, the specialised version shows higher average coverage and path lengths than the generalised method, combined with a larger range of total coverage as seen in Figure 4.14e.

While in both methods grids of 20×20 voxels with an infill density of 73% are included in the training data and the path planning show coverage over 70% after performing 100 test runs, the specialised RL functions show better performance with a lower episode count.

On an individual basis, specialised tests outperform the generalised method of path planning, by nodal exploration and coverage. The efficiency of uniquely visited nodes falls in a similar range and is heavily influenced by the total number of nodes visited.

Combining the 2 methods results in optimal performance, since the specialised training data is prepared with a reward similar to an optimal result, reducing the need for exploration and allowing for local improvements based on efficiency, bridging and local nodal exploration.

To ensure converging Q values using a DDQN, it is recommended to determine the factor ϵ , since the stochastic results end up not being similar, following the analogy:

Where the expected variation ϵ is a small number, so $a_1 \neq a_2$. Since the environment and reward function are also created manually, the result should converge at a value close to the true Q^* value. In the current set-up, the environment is manually defined, and the state transitions (grid voxel connections) are not strictly Markovian, because Markovian sets should be independent from prior decisions, visited nodes and geometric constraints. This introduces non-stationarity and partial observability, both of which violate standard convergence assumptions. Thus, the Q-values might not converge smoothly, but instead oscillate or settle into quasi-stationary equilibria that depend on the random initialisation of weights and exploration rates. where ϵ represents the expected deviation due to stochasticity, function approximation errors, and environment uncertainty. In deterministic grid environments, ϵ ideally tends to zero, but in stochastic or manually tuned reward landscapes, a residual bias remains.

While the DDQN approach significantly improved toolpath generation for RVEs containing defects, convergence stability remains a key limitation. Because the environment and reward function are manually defined, the Markov assumption is only partially satisfied, leading to stochastic deviations in the learned Q-values.

$$Q^*(s, a_1) = Q^*(s, a_2) + \epsilon \quad (5.3)$$

The relation in Equation 5.3 highlights that convergence occurs within a bounded margin ϵ , where residual bias arises from stochastic exploration, correlated experience, and reward conflicts. To improve convergence reliability, future work should investigate adaptive exploration strategies, prioritised experience replay, and multi-objective normalisation schemes. Extending the RL environment to three

dimensions would further reduce plane-strain approximation errors and enhance generalisability to true 3D toolpaths.

5.3. Comparable toolpath panning methods

The proposed DDQN RL approach offers distinct advantages over conventional path planning methods for AM applications. Once trained, the DDQN agent can generate near-optimal toolpaths for environments similar to its training data. For grid structures with constant dimensions but varying voxel sizes or infill densities, the agent can adapt its learned policy to these slight variations, demonstrating generalisation capabilities that are valuable for AM applications with variable geometries.

5.3.1. Limitations of Heuristic Methods

Conventional heuristic approaches, such as greedy algorithms, determine travel motions by evaluating local cost functions at each node [51]. While computationally efficient, greedy methods are prone to becoming trapped in local minima, where all available moves result in highly penalised rewards—often due to necessary non-extrusion travel motions. A critical limitation of heuristic methods is their inability to reconsider previous decisions; they approximate optimal paths through limited forward-looking analysis without re-evaluating prior choices. In contrast, neural network approaches such as DDQN learn to anticipate and avoid locally suboptimal decisions through the training process [52].

5.3.2. Travelling Salesman Problem Approaches

The Travelling Salesman Problem (TSP) framework aims to minimise total path length while visiting each node, typically following an Eulerian path where each node is visited at least once [51]. While TSP solvers can efficiently generate optimal paths for small grids (e.g., 10×10), they suffer from poor scalability and reusability. Each new grid configuration requires a complete recalculation, making TSP-based methods impractical for highly adaptable or variable grid structures. The computational advantage of RL becomes apparent when considering repeated path generation across similar but non-identical geometries.

5.3.3. Shortest Path Algorithms

Graph-based algorithms such as Dijkstra and A* excel at finding minimal-distance paths between two points in weighted graphs, making them popular in robotics applications [53]. However, these algorithms fundamentally conflict with AM toolpath requirements. Their objective—minimising distance from point A to B—contradicts the goal of maximising node coverage while minimising non-extrusion travel. Consequently, they fail to generate efficient toolpaths for additive manufacturing applications.

5.3.4. Advantages and Considerations for DDQN RL

Well-trained RL agents demonstrate robust performance and adaptability to environmental variations within their training distribution. The DDQN approach generates an approximated Q-function that can generalise to similar environments, provided the training data is sufficiently diverse and representative. However, this method has important limitations. Performance degrades when applied to conditions substantially different from the training distribution, and biased or insufficient training data can lead to overfitting, where the agent learns spurious correlations that fail to generalise [54]. Therefore, generating diverse, representative training datasets is critical for optimal RL performance.

A notable challenge with RL methods is their interpretability. The hidden layers that process states and actions create a "black box" effect, making it difficult to diagnose errors or manually tune performance compared to traditional algorithms with adjustable parameters. Improving RL-based toolpaths requires expertise in deep learning and reinforcement learning principles, as external modifications must propagate through the network's hidden representations. Despite this added complexity, the adaptability and generalisation capabilities of DDQN RL justify its use for applications requiring repeated path generation across variable geometries [54].

A combination of both methods can be used, where the training data includes multiple grids to create a toolpath circumnavigating defects of various voxel sizes, while the test data focuses on performing optimal results on a single grid with optimal performance.

RL using a DDQN is ideal for considering path planning using multiple parameters. Not limiting the research to the physical path boundaries, as conventional methods show better results. As shown in the results, the combination of the modified toolpath has an optimal input of a heavily researched optimal path, and has manually added two extra nodes to increase the total coverage. The adaptability of this method shows how RL can be used in the design steps, and more complex decisions can be created to solve path planning. Multi-parameter optimisation has shown that a figure with defects smaller than the nozzle can be avoided while pursuing nodal coverage of 93 to 95%. Additionally, the lack of travel moves has shown that stresses are better distributed in the materials, and the number of singularities can be heavily reduced.

Other reinforcement learning approaches, such as Proximal Policy Optimisation (PPO), could be considered to mitigate the convergence instability observed in DDQN. PPO is a policy-gradient method that improves training robustness by introducing a trust region constraint on policy updates, effectively limiting the Kullback–Leibler (KL) divergence between successive policies [55]. This ensures smoother updates and prevents oscillatory behaviour often seen in value-based methods.

Unlike DDQN, PPO does not rely on an ϵ -greedy exploration schedule and therefore requires less hyperparameter tuning, which can improve reproducibility and training efficiency [56]. In the context of toolpath optimisation, PPO could be advantageous for handling continuous or quasi-continuous action spaces, such as nozzle direction and extrusion speed, where discrete state-action formulations are limiting.

Future work could integrate the homogenisation framework directly into the reinforcement learning environment, enabling stress-optimised grid designs to inform the policy itself. This would allow the agent to learn toolpaths that inherently follow principal stress directions. Compared to heuristic methods, the RL-based approach demonstrated convergence after approximately one thousand training episodes, particularly in generalised environments with higher reward scaling. The convergence behaviour suggests that policy-based extensions such as PPO may further enhance stability and facilitate the discovery of multiple near-optimal toolpath solutions.

5.4. Shape fidelity

During material extrusion, motions similar to the diameter of the nozzle occur. Depending on the material deposition, too much material can be extruded when the nozzle is not moving at a constant velocity. Low acceleration at the start results in higher material deposition rates during the initial step, forming local blobs, reducing the shape fidelity of the final result. Using PLA in MEX has a negative benefit of the material being pressed to the base plate or lower layers, which results in a spread-out line to create improved adherence, but reduces the shape fidelity due to increased thickness of the extruded lines.

The Ultimaker 2+ used has a nozzle diameter of 0.4 mm. In the current design, supervoxels have an equal size and include the changes in all directions after any node, resulting in high accelerations and a change of angle over small distances. The positioning precision of an Ultimaker 2+ on the x and y-axes is set at 12.5 microns, and 5 microns in the z-directions; thus, correct positioning can be expected using the current grid. Poor printability in the current setup can be caused by over-extrusion in the complete process. Since the RL function is aware of the nozzle diameter, less material should be extruded if the material is already covered by the nozzle. Reducing revisits grants a closer result to the planned toolpath.

During FDM, the optimal height for the first layer can be reduced to 0.27 mm, whereas all other layers are built on 0.4 mm, due to the limited adherence between the build plate and the PLA.

Due to the G-Code being created manually, the G-Code is highly tunable. Since scaling is possible, larger and smaller nozzle diameters can be used, while scaling or reducing the size of voxels compared to the current setup. Scaling the grid is easily achievable, making the output of the result optimal for reproduction and as a building block for more complex shapes.

5.5. Comparing experimental and computational results

To validate the computational model on tensile tests are performed to determine the stiffness for E_{11} and E_{22} . The experimental results show lower stiffnesses for the material than anticipated in the com-

putational model. Compared to the literature, the Yield stress and UTS are heavily reduced when considering a reduced infill of nearly 73% [57]. This can be explained due to the fracture occurring between the RVEs, where poor interlayer connections occur, due to reduced adherence between RVEs, causing RVE-by-RVE extrusion. Hardened material adheres poorly to newly extruded material. A recommendation to prevent this is to increase the size of the optimised path of an RVE to cover the complete area of the reduced thickness of the tensile test specimen. Due to the early rupture of the material, limited information is given to validate the computational model. While the specimens have a reduced Young's Modulus compared to the computational results, the results are in a similar range as the results of other PLA tensile tests of multiple experiments [20], [58]. This indicates that the computational model is overestimating, possibly by connecting all voxels, whereas in reality, a node is connected to only 2 other nodes. More research needs to be performed to validate this hypothesis.

5.6. Expansion to multi materials and 3D structures

In future processes, the current method used in this thesis can be expanded to cover more complex structures. By first including multiple materials in a single RVE, the tunable factors are expected to increase.

5.6.1. Multi materials

In the available code provided by Andreassen et al. [6] and Ozdilek et al. [30], multiple materials can be added. Adding multiple stiffnesses and Poisson ratios as a scaling factor is used by Andreassen et al. to include multiple materials in the linearised homogenisation method. Since the stiffness E and Poisson ratio ν are used as scale factors, adding multiple materials in a single RVE is achievable and can be included in future work to expand the tunability of AM-produced materials. A limitation of the voxel structure used to create a grid is that every voxel must be classified with a single number, so voxels that contain multiple materials result in an error. In future work, gradient distribution per voxel can be included to add multiple materials and also determine the interaction between materials due to the difference in stiffness. FDM is suitable for performing multi-material AM, and the extension of the current homogenisation method can be included to increase the mechanical properties of an RVE.

5.6.2. 3D homogenisation

The current research has focused on 2D homogenisation and is able to produce structures that are created using 2.5D FDM. By increasing the grid of voxels to the third dimension and creating a cubic shape, homogenisation can be performed in an improved version. Focusing on the same principles as for the homogenisation techniques, but not limited to the trade-off between plane-stress and plane-strain, the increased connectivity to a maximum of 26 voxels improves the simulation. By using different types of AM as multi-axis AM, it becomes feasible to create structures less dependent on flat base layers. The method of 3D homogenisation is supported by Ozdilek et al. [30] and Dong et al. [47], both theoretically as available in MATLAB and Python code. To continue the 3D homogenisation method, a MATLAB code has already been created to either build a layered set of 2D grids on top of each other or to create a 3D RVE out of voxels. 3D homogenisation could be incorporated with 3D or multidirectional path planning and AM. Multi-axis AM can be used to increase the mobility range of the nozzle, creating more complex shapes, with reduced limitations due to the range of motion.

5.6.3. 3D path planning

A current limitation in the path planning is remaining in planar toolpaths (2.5D). Extending the range of motion to start path planning in 3D using multiple DOFs would allow for the creation of more complex, spatially continuous structures. Multi-axis AM enables multidirectional extrusion, which reduces the need for redundant support material if overhanging or upper layers are self-supporting [59].

For multi-axis AM, additional travel motions must be included in the output layer. Additional functions should include self-supporting steps based on the extrusion angle compared to the nozzle, unsupported edge length, acceptable sagging rates and collision detection. Using AM in non-planar directions, the material can be extruded while being supported by lower layers, thereby reducing the downward gravitational forces acting on the newly deposited filament. Multi-axis AM also mitigates collision with previously printed features and enables the formation of continuous toolpaths, which reduces singularities

in motion control [36]. As a result, highly permeable and lightweight structures can be achieved.

FDM using Multi-axis AM 3D path planning introduces several geometric and mechanical constraints. Due to the nozzle shape and size, downward or negative extrusion motions are not feasible, as they risk collision with previously deposited material. Sharp downward movements during extrusion are particularly problematic. Consequently, 3D path planning requires hierarchical trajectory generation in the z-direction, rather than the layer-per-layer approach typically used in 2.5D AM. An asymmetric Laplacian matrix can be considered for use in 3D path planning.

A similar file has been created for `homogenize_3D.m` for homogenisation in 3D that can be combined automatically with the 3D optimisation file to produce a Young's Modulus in 3D for a cuboid with randomly generated voxels. By increasing the connectivity in 3D to 26 travel motions, the infill density can be further reduced while meeting the constraints for periodic boundary conditions. For the objective function, all values on the z-axis are incorporated.

In the `single_grid_reinforcement_planning.m` script, movement possibilities are currently limited to eight neighbouring nodes on the z-plane, each separated by $2\sqrt{2} \cdot r_{\text{nozzle}}$. Extending the model to full 3D motion with 26 potential nodes requires increasing the node spacing to $2\sqrt{3} \cdot r_{\text{nozzle}}$ and removing physically impractical transitions, resulting in a feasible motion space of 26 nodes. Incorporating the third dimension introduces additional parameters per position feature; in the current setup, at least three extra parameters are required to implement 3D motion in the input layer.

Future implementations should include additional kinematic and geometric constraints based on machine architecture and nozzle accessibility. The starting point should remain on the base layer to prevent unsupported material deposition. 3D nozzle geometry limits achievable movement paths, as collisions can occur between the nozzle and extruded material. The order of toolpath execution must also be considered, as accessibility of certain points depends on previously deposited layers. These aspects highlight that fully three-dimensional path planning demands a more complex hierarchical framework, integrating collision avoidance, support evaluation, and reachability constraints.

6

Conclusion

6.1. Conclusion

This research demonstrates that reinforcement learning (RL) can be successfully applied to optimise toolpath planning in material extrusion (MEX) and fused deposition modelling (FDM), producing structures with improved isotropy and mechanical performance. By simplifying complex geometries into repeatable voxel-based representations and applying homogenisation analysis, a framework was developed that identifies optimal grid configurations balancing stiffness, isotropy, and connectivity. The Double Deep Q-Network (DDQN) agent subsequently learned to generate near-Eulerian toolpaths that efficiently cover material domains while avoiding defects and redundant travel.

Compared with conventional raster-based path planning, the optimised RL-driven toolpaths exhibited smoother material distribution and improved shear performance, particularly at higher infill densities. The reinforcement learning process showed stable convergence across both generalised and specialised training sets, indicating that multi-objective path planning can be learned autonomously without exhaustive manual tuning. The results underline that computational optimisation across all stages of the design-to-manufacture workflow, from grid generation to extrusion planning, significantly enhances the predictability and performance of printed structures.

In practical terms, the research highlights the potential of data-driven toolpath optimisation to complement classical geometric design methods. The combination of homogenisation and reinforcement learning offers a scalable strategy for developing materials with tunable mechanical responses, applicable to both engineering and biomedical scaffolding contexts. Further refinement of printing precision and motion control could yield even greater mechanical consistency between simulated and manufactured results.

The tensile test performed to validate the mechanical properties shows limited results due to early fractures between RVEs caused by poor adherence between RVEs.

Overall, this thesis provides a foundation for the integration of RL into AM workflows, demonstrating that intelligent, adaptive path planning can reduce anisotropy and improve structural performance in printed materials.

6.2. Future work

Several extensions can further enhance the presented framework. Computationally, the current homogenisation and toolpath planning are limited to 2D and 2.5D extrusion paths. Future research should focus on implementing full 3D homogenisation to account for through-thickness stresses and interlayer bonding effects. Integrating these 3D material models directly into the RL environment would allow the agent to optimise across true volumetric design spaces.

From a manufacturing perspective, extending the method to multi-directional and non-planar printing using robotic or multi-axis AM systems could eliminate the layer-by-layer constraints that cause anisotropy.

Including self-supporting structure objectives within the reward function would encourage the agent to minimise unnecessary material use and produce lightweight yet mechanically robust geometries.

Algorithmically, alternative RL approaches such as Proximal Policy Optimisation (PPO) could be explored to further stabilise training by constraining policy updates through trust-region methods. Combining DDQN's discrete decision efficiency with PPO's policy stability may lead to hybrid agents capable of optimising both discrete voxel transitions and continuous motion trajectories.

The validation of the computational model should be reinvestigated, since the mechanical performance of the specimens was lower than anticipated. In future research, the computational model should be reinvestigated and possibly adjusted to compare with the physical results.

Finally, broader applications could include AM in tissue engineering, high-tech AM and metamaterial design, where stress-guided toolpaths and tunable anisotropy are critical. The integration of homogenisation, machine learning, and physical experimentation provides a promising path towards adaptive, closed-loop additive manufacturing systems capable of self-optimisation across design, simulation, and production.

Acknowledgments

In this thesis, a partial MATLAB code has been used, created by Andreassen and Andreassen and Abdulrahman Al-Sanea for homogenisation. Specifications for both codes are made to compare them more to my goal of random defect placement. For inspiration and improving the structure of the Reinforcement Learning code, large language models are used.

References


- [1] T. M. Tavares, G. M. D. Ganga, M. G. Filho, and V. P. Rodrigues, "The benefits and barriers of additive manufacturing for circular economy: A framework proposal," *Sustainable Production and Consumption*, vol. 37, pp. 369–388, May 2023, ISSN: 2352-5509. DOI: 10.1016/J.SPC.2023.03.006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352550923000507>.
- [2] X. Han, Y. Jiang, W. Wang, *et al.*, "Motion-driven neural optimizer for prophylactic braces made by distributed microstructures," *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, Dec. 2024. DOI: 10.1145/3680528.3687661. [Online]. Available: <https://dl.acm.org/doi/10.1145/3680528.3687661>.
- [3] J. Son, S. Yun, K. Park, S. Ryu, and S. Kim, "Isotropic 3d printing using material extrusion of thin shell and post-casting of reinforcement core," *Additive Manufacturing*, vol. 58, p. 102974, Oct. 2022, ISSN: 2214-8604. DOI: 10.1016/J.ADDMA.2022.102974. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214860422003670>.
- [4] D. X. B. Chen, "Extrusion bioprinting of scaffolds: An introduction," *Extrusion Bioprinting of Scaffolds for Tissue Engineering*, pp. 1–15, 2025. DOI: 10.1007/978-3-031-72471-8_1. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-72471-8_1.
- [5] "Operators in banach spaces," in *Non-Homogeneous Media and Vibration Theory*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1980, pp. 7–23, ISBN: 978-3-540-39317-7. DOI: 10.1007/3-540-10000-8_2. [Online]. Available: https://doi.org/10.1007/3-540-10000-8_2.
- [6] E. Andreassen and C. S. Andreasen, "How to determine composite material properties using numerical homogenization," *Computational Materials Science*, vol. 83, pp. 488–495, Feb. 2014, ISSN: 0927-0256. DOI: 10.1016/J.COMMATSCI.2013.09.006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0927025613005302>.
- [7] T. Liu, T. Zhang, Y. Chen, Y. Huang, and C. C. Wang, "Neural slicer for multi-axis 3d printing," *ACM Transactions on Graphics (TOG)*, vol. 43, p. 15, 4 Jul. 2024, ISSN: 15577368. DOI: 10.1145/3658212. [Online]. Available: <https://dl.acm.org/doi/10.1145/3658212>.
- [8] Unknown, *The 7 categories of additive manufacturing | additive manufacturing research group | loughborough university*. [Online]. Available: <https://www.lboro.ac.uk/research/amrg/about/the7categoriesofadditivemanufacturing/>.
- [9] A. Cano-Vicent, M. M. Tambuwala, S. S. Hassan, *et al.*, "Fused deposition modelling: Current status, methodology, applications and future prospects," *Additive Manufacturing*, vol. 47, p. 102378, Nov. 2021, ISSN: 2214-8604. DOI: 10.1016/J.ADDMA.2021.102378. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214860421005327>.
- [10] M. T. Birosz, D. Ledenyák, and M. Andó, "Effect of fdm infill patterns on mechanical properties," *Polymer Testing*, vol. 113, p. 107654, Sep. 2022, ISSN: 0142-9418. DOI: 10.1016/J.POLYMERTESTING.2022.107654. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0142941822001787>.
- [11] S. Garzon-Hernandez, D. Garcia-Gonzalez, A. Jérusalem, and A. Arias, "Design of fdm 3d printed polymers: An experimental-modelling methodology for the prediction of mechanical properties," *Materials & Design*, vol. 188, p. 108414, Mar. 2020, ISSN: 0264-1275. DOI: 10.1016/J.MATDES.2019.108414. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0264127519308524>.
- [12] J. P. Thomas and J. F. Rodriguez, "Modeling the fracture strength between fused-deposition extruded roads 16," 2000. DOI: 10.26153/TSW/2054. [Online]. Available: <https://hdl.handle.net/2152/74942>.

- [13] P. Datta, V. Vyas, S. Dhara, A. R. Chowdhury, and A. Barui, "Anisotropy properties of tissues: A basis for fabrication of biomimetic anisotropic scaffolds for tissue engineering," *Journal of Bionic Engineering*, vol. 16, pp. 842–868, 5 Sep. 2019, ISSN: 25432141. DOI: 10.1007/S42235-019-0101-9.
- [14] M. Somireddy and A. Czekanski, "Anisotropic material behavior of 3d printed composite structures – material extrusion additive manufacturing," *Materials & Design*, vol. 195, p. 108953, Oct. 2020, ISSN: 0264-1275. DOI: 10.1016/J.MATDES.2020.108953. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0264127520304871>.
- [15] G. Fang, T. Zhang, Y. Huang, Z. Zhang, K. Masania, and C. C. Wang, "Exceptional mechanical performance by spatial printing with continuous fiber: Curved slicing, toolpath generation and physical verification," *Additive Manufacturing*, vol. 82, p. 104048, Feb. 2024, ISSN: 2214-8604. DOI: 10.1016/J.ADDMA.2024.104048.
- [16] J. A. Reid, K. D. Dwyer, P. R. Schmitt, A. H. Soepriatna, K. L. Coulombe, and A. Callanan, "Architected fibrous scaffolds for engineering anisotropic tissues," *Biofabrication*, vol. 13, 10.1088/1758-5090/ac0fc9, 4 Oct. 2021, ISSN: 17585090. DOI: 10.1088/1758-5090/AC0FC9. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8686077/>.
- [17] M. Hassan, A. K. Mohanty, and M. Misra, "3d printing in upcycling plastic and biomass waste to sustainable polymer blends and composites: A review," *Materials & Design*, vol. 237, p. 112558, Jan. 2024, ISSN: 0264-1275. DOI: 10.1016/J.MATDES.2023.112558. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0264127523009747>.
- [18] M. Heidari-Rarani, M. Rafiee-Afarani, and A. M. Zahedi, "Mechanical characterization of fdm 3d printing of continuous carbon fiber reinforced pla composites," *Composites Part B: Engineering*, vol. 175, p. 107147, Oct. 2019, ISSN: 1359-8368. DOI: 10.1016/J.COMPOSITESB.2019.107147. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1359836818335686>.
- [19] Y. Huang, J. Ding, T. Chi-, *et al.*, "Learning based toolpath planner on diverse graphs for 3d printing," *ACM Transactions on Graphics (TOG)*, vol. 43, p. 16, 6 Dec. 2024, ISSN: 15577368. DOI: 10.1145/3687933. [Online]. Available: <https://dl.acm.org/doi/10.1145/3687933>.
- [20] S. Molazadeh, F. Diba, and A. Hosseini, "Anisotropic modeling of material behavior for additively manufactured parts made by material extrusion," *The International Journal of Advanced Manufacturing Technology* 2023 129:7, vol. 129, pp. 3453–3473, 7 Oct. 2023, ISSN: 1433-3015. DOI: 10.1007/S00170-023-12508-5. [Online]. Available: <https://link.springer.com/article/10.1007/s00170-023-12508-5>.
- [21] M. Ozdemir, U. Simsek, C. E. Gayir, K. Gunaydin, and O. Gulcan, "An integrated matlab code for homogenization-based topology optimization and generating functionally graded surface lattices for additive manufacturing," *Advanced Engineering Materials*, vol. 27, p. 2402567, 15 Aug. 2025, ISSN: 15272648. DOI: 10.1002/ADEM.202402567; REQUESTEDJOURNAL: JOURNAL: 15272648; PAGE: STRING: ARTICLE/CHAPTER. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/adem.202402567>20<https://advanced.onlinelibrary.wiley.com/doi/10.1002/adem.202402567>.
- [22] Z. A. Lethbridge, R. I. Walton, A. S. Marmier, C. W. Smith, and K. E. Evans, "Elastic anisotropy and extreme poisson's ratios in single crystals," *Acta Materialia*, vol. 58, pp. 6444–6451, 19 Nov. 2010, ISSN: 1359-6454. DOI: 10.1016/J.ACTAMAT.2010.08.006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1359645410005100>.
- [23] D. K. Y. Tam, S. Ruan, P. Gao, and T. Yu, "High-performance ballistic protection using polymer nanocomposites," *Advances in Military Textiles and Personal Equipment*, pp. 213–237, Jan. 2012. DOI: 10.1533/9780857095572.2.213. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9781845696993500104>.
- [24] M. Leary, "Design of lattice and zero-mean curvature structures," *Design for Additive Manufacturing*, pp. 123–163, 2020. DOI: 10.1016/B978-0-12-816721-2.00005-1.

- [25] R. Barretta, J. Somnic, and B. W. Jo, "Homogenization methods of lattice materials," *Encyclopedia 2022, Vol. 2, Pages 1091-1102*, vol. 2, pp. 1091–1102, 2 May 2022, ISSN: 2673-8392. DOI: 10.3390/ENCYCLOPEDIA2020072. [Online]. Available: <https://www.mdpi.com/2673-8392/2/2/72/htm%20https://www.mdpi.com/2673-8392/2/2/72>.
- [26] A. S. Phani, "Elastodynamics of lattice materials," *Dynamics of Lattice Materials*, pp. 53–92, Jul. 2017. DOI: 10.1002/9781118729588.CH3. [Online]. Available: [/doi/pdf/10.1002/9781118729588.CH3. \[Online\]. Available: https://doi/pdf/10.1002/9781118729588.CH3. \[Online\]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118729588.ch3%20https://onlinelibrary.wiley.com/doi/10.1002/9781118729588.ch3](https://doi/pdf/10.1002/9781118729588.ch3%20https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118729588.ch3%20https://onlinelibrary.wiley.com/doi/10.1002/9781118729588.ch3).
- [27] N. Pasvanti, A. Psarros, G. Korbetis, A. Vlahinos, and A. Mihailidis, "Lattice structures modeling: Introduction to homogenization," in *8th Before Reality Conference*, https://www.beta-cae.com/events/c8pdf/3D_1_PASVANTI.pdf, Thessaloniki, Greece, 2019, pp. —.
- [28] J. Yan, G. Cheng, S. Liu, and L. Liu, "Comparison of prediction on effective elastic property and shape optimization of truss material with periodic microstructure," *International Journal of Mechanical Sciences*, vol. 48, pp. 400–413, 4 Apr. 2006, ISSN: 0020-7403. DOI: 10.1016/J.IJMECSCI.2005.11.003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020740305002699>.
- [29] Z. Xia, C. Zhou, Q. Yong, and X. Wang, "On selection of repeated unit cell model and application of unified periodic boundary conditions in micro-mechanical analysis of composites," *International Journal of Solids and Structures*, vol. 43, pp. 266–278, 2 Jan. 2006, ISSN: 0020-7683. DOI: 10.1016/J.IJSOLSTR.2005.03.055. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020768305001460>.
- [30] E. E. Ozdilek, E. Ozcakar, N. Muhtaroglu, U. Simsek, O. Gulcan, and G. K. Sendur, "A finite element based homogenization code in python: Hompy," *Advances in Engineering Software*, vol. 194, p. 103674, Aug. 2024, ISSN: 0965-9978. DOI: 10.1016/J.ADVENGSOFT.2024.103674. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0965997824000814#b36>.
- [31] A. Bensoussan, J.-L. Lions, and G. Papanicolaou, *Asymptotic analysis for periodic structures*. American Mathematical Soc., 2011, vol. 374.
- [32] B. Hassani and E. Hinton, "A review of homogenization and topology optimization i—homogenization theory for media with periodic structure," *Computers & Structures*, vol. 69, pp. 707–717, 6 Dec. 1998, ISSN: 0045-7949. DOI: 10.1016/S0045-7949(98)00131-X. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S004579499800131X>.
- [33] D. Roylance, *Constitutive equations*, MIT OpenCourseWare, 3.11 Mechanics of Materials, Fall 1999, Lecture notes, https://ocw.mit.edu/courses/3-11-mechanics-of-materials-fall-1999/b51fefe25af90628d210e7029c7f744c_MIT3_11F99_const.pdf, Cambridge, MA, Oct. 2000.
- [34] L. L. Yaw, *Introduction to the virtual element method for 2d elasticity*, arXiv preprint arXiv:2301.11928, Version v3, updated 21 Sep 2023, 2023.
- [35] H. Sönerlind, *What is the difference between plane stress and plane strain?* COMSOL Blog, Accessed 2025-09-09, May 2021. [Online]. Available: <https://www.comsol.com/blogs/what-is-the-difference-between-plane-stress-and-plane-strain>.
- [36] T. Zhang, X. Chen, G. Fang, Y. Tian, and C. C. Wang, "Singularity-aware motion planning for multi-axis additive manufacturing," *IEEE Robotics and Automation Letters*, vol. 6, pp. 6172–6179, 4 Oct. 2021, ISSN: 23773766. DOI: 10.1109/LRA.2021.3091109.
- [37] T. Zhang, X. Chen, G. Fang, Y. Tian, and C. C. L. Wang, "Singularity-aware motion planning for multi-axis additive manufacturing," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6172–6179, 2021. DOI: 10.1109/LRA.2021.3091109.
- [38] M. Qin, J. Ding, S. Qu, X. Song, C. C. Wang, and W. H. Liao, "Deep reinforcement learning based toolpath generation for thermal uniformity in laser powder bed fusion process," *Additive Manufacturing*, vol. 79, p. 103937, Jan. 2024, ISSN: 2214-8604. DOI: 10.1016/J.ADDMA.2023.103937. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S221486042300550X>.

- [39] A. Twumasi, P. C. Roy, Z. Li, S. S. Bhattacharjee, and Z. Gan, *Laser scan path design for controlled microstructure in additive manufacturing with integrated reduced-order phase-field modeling and deep reinforcement learning*, 2025. arXiv: 2506.21815 [cs.CE]. [Online]. Available: <https://arxiv.org/abs/2506.21815>.
- [40] A. Schwab, R. Levato, M. D'Este, S. Piluso, D. Eglin, and J. Malda, "Printability and shape fidelity of bioinks in 3d bioprinting," *Chemical Reviews*, vol. 120, pp. 11 028–11 055, 19 Oct. 2020, ISSN: 15206890. DOI: 10.1021/acs.chemrev.0c00084.
- [41] P. Rajakareyar, M. S. ElSayed, H. A. E. Ella, and E. Matida, "Effective mechanical properties of periodic cellular solids with generic bravais lattice symmetry via asymptotic homogenization," *Materials*, vol. 16, 24 Dec. 2023, ISSN: 19961944. DOI: 10.3390/MA16247562.
- [42] M. Espinola, J. A. Piedra-Fernandez, R. Ayala, L. Iribarne, and J. Z. Wang, "Contextual and hierarchical classification of satellite images based on cellular automata," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, pp. 795–809, 2 2015, ISSN: 01962892. DOI: 10.1109/TGRS.2014.2328634.
- [43] W. Tian, L. Qi, X. Chao, J. Liang, and M. Fu, "Periodic boundary condition and its numerical implementation algorithm for the evaluation of effective mechanical properties of the composites with complicated micro-structures," *Composites Part B: Engineering*, vol. 162, pp. 1–10, Apr. 2019, ISSN: 1359-8368. DOI: 10.1016/J.COMPOSITESB.2018.10.053. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1359836818319838>.
- [44] W. Crupano, B. Adrover-Monserrat, J. Llumà, R. Jerez-Mesa, and J. A. Travieso-Rodriguez, "Investigating mechanical properties of 3d printed polylactic acid / poly-3-hydroxybutyrate composites. compressive and fatigue performance," *Heliyon*, vol. 10, e38066, 18 Sep. 2024, ISSN: 2405-8440. DOI: 10.1016/J.HELIYON.2024.E38066. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405844024140972>.
- [45] F. Aldosari, M. M. A. A. Khan, M. Asad, and F. Djavanroodi, "Finite element analysis of polylactic acid (pla) under tensile and compressive loading," in *IOP Conference Series: Materials Science and Engineering*, vol. 2468, IOP Publishing, 2023, p. 012 094. DOI: 10.1088/1742-6596/2468/1/012094.
- [46] R. Hill, "Elastic properties of reinforced solids: Some theoretical principles," *Journal of the Mechanics and Physics of Solids*, vol. 11, pp. 357–372, 5 Sep. 1963, ISSN: 0022-5096. DOI: 10.1016/0022-5096(63)90036-X. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/002250966390036X?via%3Dihub>.
- [47] G. Dong, Y. Tang, and Y. F. Zhao, "A 149 line homogenization code for three-dimensional cellular materials written in matlab," *Journal of Engineering Materials and Technology, Transactions of the ASME*, vol. 141, 1 Jan. 2019, ISSN: 15288889. DOI: 10.1115/1.4040555/368579. [Online]. Available: <https://dx.doi.org/10.1115/1.4040555>.
- [48] Y. Ji, Q. Chen, R. Wang, N. Ji, and H. Ma, "Quantum error correction for heavy hexagonal code using deep reinforcement learning with policy reuse," *Quantum Information Processing*, vol. 23, 7 Jul. 2024, ISSN: 15731332. DOI: 10.1007/S11128-024-04377-Y.
- [49] P. John, V. R. Komma, and S. P. Bhore, "Development of matlab code for tool path data extraction from the g code of the fused filament fabrication (fff) parts," *Engineering Research Express*, vol. 5, p. 025 018, 2 Apr. 2023, ISSN: 2631-8695. DOI: 10.1088/2631-8695/ACCC6F. [Online]. Available: <https://iopscience.iop.org/article/10.1088/2631-8695/accc6f%20https://iopscience.iop.org/article/10.1088/2631-8695/accc6f/meta>.
- [50] *Nen connect - iso 527-4:2023 en*. [Online]. Available: <https://connect.nen.nl/Standard/Detail/3678884?compId=10037&collectionId=0>.
- [51] A. Frieze and W. Pegden, "The bright side of simple heuristics for the tsp," Oct. 2023. [Online]. Available: <https://arxiv.org/pdf/2310.03222>.
- [52] A. Müller and L. Vollenkemper, "Reinforcement learning as an improvement heuristic for real-world production scheduling," *arXiv preprint arXiv:2409.11933*, 2024, Preprint. [Online]. Available: <https://arxiv.org/abs/2409.11933>.

- [53] R. Sedgewick and K. Wayne, *Algorithms, 4th edition — website*, <https://algs4.cs.princeton.edu/44sp/>, Section 4.4: Shortest Paths, 2025.
- [54] S. Casper, X. Davies, C. Shi, *et al.*, “Open problems and fundamental limitations of reinforcement learning from human feedback,” *Transactions on Machine Learning Research*, vol. 2023, Jul. 2023, ISSN: 28358856. [Online]. Available: <https://arxiv.org/pdf/2307.15217>.
- [55] J. Schulman, S. Levine, P. Moritz, M. Jordan, and P. Abbeel, “Trust region policy optimization,” in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ser. ICML’15, Lille, France: JMLR.org, 2015, pp. 1889–1897.
- [56] T. P. Lillicrap, J. J. Hunt, A. Pritzel, *et al.*, “Continuous control with deep reinforcement learning,” *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, Sep. 2015. [Online]. Available: <https://arxiv.org/pdf/1509.02971>.
- [57] S. Farah, D. G. Anderson, and R. Langer, “Physical and mechanical properties of pla, and their functions in widespread applications — a comprehensive review,” *Advanced Drug Delivery Reviews*, vol. 107, pp. 367–392, Dec. 2016, ISSN: 0169-409X. DOI: 10.1016/J.ADDR.2016.06.012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169409X16302058>.
- [58] C. Grabowik, K. Kalinowski, G. Ćwikła, I. Paprocka, and P. Kogut, “Tensile tests of specimens made of selected group of the filament materials manufactured with fdm method,” *MATEC Web of Conferences*, vol. 112, p. 04 017, Jan. 2017. DOI: 10.1051/mateconf/201711204017.
- [59] T. Liu, T. Zhang, Y. Chen, Y. Huang, and C. C. L. Wang, “Neural slicer for multi-axis 3d printing,” *ACM Trans. Graph.*, vol. 43, no. 4, Jul. 2024, ISSN: 0730-0301. DOI: 10.1145/3658212. [Online]. Available: <https://doi.org/10.1145/3658212>.




Simulation results


A.1. Horizontal grid

The horizontal grid or zigzag pattern is used to compare optimized grid patterns to existing shapes. The horizontal connected blocks are connected to their vertical components, therefore the stiffness tensor is the highest in the E_{11} direction. The stiffness for the horizontal grid is $E_{11} = 42.2523$ GPa, $E_{22} = 3.1664$ GPa and $G_{12} = 0.7707$ GPa. It has a Zener ratio of 0.0371, $\nu_{12} = 0.0233$ and $\nu_{21} = 0.3120$. The C^H

matrix for the horizontal grid is:
$$\begin{bmatrix} 38.9725 & 0.9081 & -0.0383 \\ 0.9081 & 2.9107 & -0.1225 \\ -0.0383 & -0.1225 & 0.7063 \end{bmatrix}$$



(a) Horizontal grid structure with an infill density of 73%.



(b) Polar plot of the Young's modulus of the horizontal grid. $E_{11} = 42.2523$ MPa, $E_{22} = 3.1664$ MPa and $G_{12} = 0.7707$ GPa.

Figure A.1: Geometric configuration and directional stiffness characteristics of the horizontal grid structure.


A.2. Spiral grid

The spiral grid is used to compare optimized methods to conventional printing techniques. $E_{11} = 4.3422$ GPa, $E_{22} = 4.7738$ GPa, $G_{12} = 0.0248$ GPa. The Zener ratio A is 0.0122. The Poisson ratios $\nu_{12} =$


0.0427 and $\nu_{21} = 0.0403$. The C^H matrix is
$$\begin{bmatrix} 4.6175 & 0.1969 & 0.0860 \\ 0.1969 & 4.8900 & 0.0560 \\ 0.0860 & 0.0560 & 0.0270 \end{bmatrix}$$

A.3. Index numbers

For figure 4.6 the IDs correspond to the 20x20 grid with infill of 73% are: b) 230067, c) 34327, d) 8978, e) 199543 f) 63345. The ID numbers correspond to the grid_all.mat file found in 2D voxel.



(a) Spiral grid structure with an infill density of 73%.



(b) Polar plot of the Young's Moduli of the spiral pattern.
 $E_{11} = 4.3422$ MPa, $E_{22} = 4.7738$ MPa and $G_{12} = 0.0248$ MPa.

Grid	Index Number
20 × 20 63%	31418
20 × 20 73%	57906
20 × 20 83%	54339
20 × 20 93%	17113
10 × 10 73%	61492
40 × 40 73%	576

Table A.1: Index number to the corresponding best grids.

For Figure 4.7 the IDs correspond to the 20x20 grid with infill of 73% are: b) 15785, c) 17016, d) 205746, e) 35077, f) 161831. The ID numbers correspond to the grid_all.mat file fund in 2D voxel.

B

Additional material for RL

B.1. Generalised training

The first step in RL is to train the tool using random grids, with varying grid sizes and voxel sizes. Starting with a reward of 0, the goal is to optimise the reward after each episode. The reward curve in Figure B.1 is increasing, and has partial drops between episodes, due to a dead-end or low coverage rate. The final Q value for the episode is 4810.3965. The runtime for 2000 episodes is 00:50:57. The combined table is a randomly generated set of grids that is created at the beginning of `RL_toolpath.m` and contains a chosen number of grids. For general testing, initially 50 grids were used, later slimmed down to 8 grids to improve repeatability and increase complete exploration.

The critic, agent and training options included for generalised training are:

- `LearnRate` = $5e^{-4}$
- `GradientThreshold` = 1
- `L2RegularizationFactor` = $1e^{-4}$
- `MiniBatchSize` = 128
- `DiscountFactor` = 0.995
- `ExperienceBufferLength` = 10000
- `TargetSmoothFactor` = $5e^{-4}$
- `UseDoubleDQN` = true
- `MaxEpisodes` = 2000
- `MaxStepsPerEpisode` = 3500
- `Verbose` = false

B.2. Specialised training

The specialised training starts after a generalised training set and evaluation. In the specialised training, a single grid is used for training to increase the coverage. After increasing the coverage, efficiency can also be increased. The runtime for 1500 episodes is 00:39:04. The training set is a single grid with index number 57906. And the initial starting point is randomly generated out of a list of minimal adjacent nodes. The critic, agent and training options included for specialised training are:

- `LearnRate` = $1e^{-4}$
- `GradientThreshold` = 1
- `L2RegularizationFactor` = $1e^{-4}$
- `MiniBatchSize` = 128




Figure B.1: The episode Reward for an RL function using a DDQN. 3 layers are applied to the generalised function, where a set of 50 grids is used for training.



Figure B.2: The Episode reward of the specialised training, occurring after a generalised training. Only 1 grid is used during the training of the specialised grid. The Episode reward increases at a lower rate than during 1500 episodes.

- DiscountFactor = 0.995
- ExperienceBufferLength = 10000
- TargetSmoothFactor = $5e^{-4}$
- UseDoubleDQN = true
- MaxEpisodes = 1500
- MaxStepsPerEpisode = 3500
- Verbose = false