CHALMERS
UNIVERSITY OF TECHNOLOGY

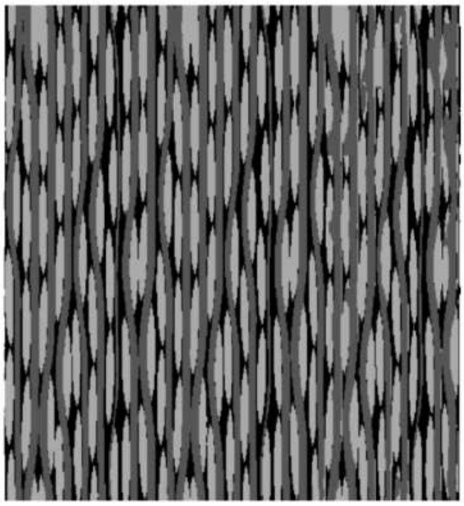

Automated CT-Based Segmentation and Modelling of Textile Composite Materials

Extension of an XCT-to-FE Modelling Workflow for Elastic Property Prediction of 2D Woven CFRP Composites

Master's thesis in Applied Mechanics

EMILY HUGHES, ANTHON RINDSÄTER

DEPARTMENT OF INDUSTRIAL AND MATERIALS SCIENCE
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Automated CT-Based Segmentation and Modelling of Textile Composite Materials

Extension of an XCT-to-FE Workflow for Elastic Property Prediction of 2D Woven CFRP Composites

EMILY HUGHES, ANTHON RINDSÄTER

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Industrial and Materials Science
Chalmers University of Technology
Gothenburg, Sweden 2026

Automated CT-Based Segmentation and Modelling of Textile Composite Materials
Extension of an XCT-to-FE Workflow for Elastic Property Prediction of 2D Woven CFRP Composites
EMILY HUGHES, ANTHON RINDSÄTER

Supervisors: Carolyn Oddy, GKN Aerospace AB
Johan Friemann, Department of Industrial and Materials Science
Examiner: Martin Fagerström, Department of Industrial and Materials Science

Master's Thesis 2026
Department of Industrial and Materials Science
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Schematic illustration of the XCT-to-FE workflow, showing a 5HS textile-reinforced composite XCT reconstruction, machine learning-based segmentation, and voxel-based finite element model.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Automated CT-Based Segmentation and Modelling of Textile Composite Materials
Extension of an XCT-to-FE Workflow for Elastic Property Prediction of 2D Woven CFRP Composites
EMILY HUGHES, ANTHON RINDSÄTER
Department of Industrial and Materials Science
Chalmers University of Technology

Abstract

Textile-reinforced carbon fibre composites are increasingly considered for aerospace applications due to their high specific stiffness and strength together with their ability to be tailored towards specific loading conditions. However, the mechanical behaviour of such materials is strongly influenced by the as-manufactured textile architecture, making accurate geometrical representations important for reliable prediction of homogenised material properties. X-ray computed tomography (XCT) combined with computational homogenisation provides a potential route for generating representative volume elements directly from physical composite samples.

The present work investigates the applicability of an existing XCT-to-FE workflow, previously developed for 3D layer-to-layer angle interlock composites, to a 2D five-harness satin (5HS) woven carbon fibre reinforced polymer composite architecture representative of aerospace applications. The workflow combines XCT imaging, machine learning-based image segmentation, and voxel-based finite element homogenisation in order to estimate elastic material properties from XCT reconstructions of as-manufactured composite samples.

Both virtual and physical investigations were carried out. A virtual 5HS composite sample generated in TexGen was utilised to evaluate XCT scan parameters and quantify segmentation performance against known ground truth data. In addition, physical XCT scans were performed on both resin-injected and dry fibre samples in order to assess the applicability of the workflow to manufactured materials. Experimental tensile testing combined with digital image correlation was further conducted in the principal material directions to provide reference elastic properties for comparison with the numerical predictions.

The results showed that the pre-trained segmentation model was capable of identifying the general textile architecture, but that additional transfer learning using synthetic 5HS training data was required to obtain segmentation quality suitable for reliable voxel-based finite element modelling. Following transfer learning, a volume-wide pixelwise agreement value of 94% was obtained for the evaluated virtual sample. The study further showed that attenuation contrast significantly influences the segmentation performance, with the dry fibre samples visually exhibiting substantially improved material constituent separation compared with the resin-injected sample.

Computational homogenisation of the segmented XCT reconstruction produced elastic stiffness predictions showing reasonable agreement with the experimentally measured tensile stiffness values, with a deviation of less than 12%. Although discrepancies between experimental and numerical results remained, the work demonstrates that the XCT-to-FE workflow can be extended to previously unseen 2D woven textile architectures through the introduction of additional target architecture-specific training data.

Overall, the results demonstrate the potential of combining XCT imaging, machine learning-based segmentation, and computational homogenisation for non-destructive characterisation of as-manufactured textile-reinforced composite samples.

Keywords: Textile-reinforced composites, Carbon fibre reinforced polymers (CFRPs), X-ray computed tomography (XCT), Computational homogenisation, Finite element analysis, Machine learning-based segmentation, Representative volume elements (RVEs), Five-harness satin weave (5HS), Voxel-based modelling, Elastic material properties.

Acknowledgements

We would sincerely like to thank all those who have made this thesis possible.

We would first like to thank our supervisor, Johan Friemann, for his guidance and support throughout the project. His valuable insights into XCT, the TexTomoS pipeline, and his willingness to guide us through challenges have been essential to this work. We are also grateful for the opportunity to use his computational resources and for the time he invested in helping us understand and work with his pipeline.

We are especially grateful to our supervisor and colleague, Carolyn Oddy, for her continuous support throughout the entire project. From assisting with the setup of the mechanical testing to providing valuable input during discussions and debugging, her contribution has been invaluable. Having had the opportunity to work with Carolyn over the past two years has been truly inspiring, and we deeply appreciate her dedication, encouragement, and support in all aspects of this work. We are particularly thankful for the significant amount of time and effort she has invested in our thesis, including proofreading and providing thoughtful and detailed feedback on the draft.

We would also like to thank Martin Fagerström for making this thesis possible, and for his insightful feedback and attention to detail, which have greatly contributed to the quality of this work. We further extend our gratitude to both Carolyn and Martin for their continuous support and guidance with any questions throughout the project.

We are grateful to Alexander Hughes for generously allowing us to use his private computing resources and for his support with dataset generation and neural network training. At GKN Aerospace, many people have contributed in various ways. We would like to thank Raul, Arda, and Johan from the inspection team for performing the XCT scanning of our composite samples, explaining the CT system, and accommodating repeated scans until satisfactory results were achieved. We also thank Kristoffer and Rickard for their assistance with the manufacturing of the physical dry fibre samples, and Monika for her support with the RWS workflow and LS-DYNA simulations. We would further like to acknowledge everyone at GKN Aerospace who contributed through support with waterjet cutting, mechanical testing, and valuable discussions related to neural networks.

Finally, we extend our gratitude to our friends and family for their continuous encouragement and support throughout our time at Chalmers. As a whole, this thesis would not have been possible without the contributions and support of all the people mentioned above.

Emily Hughes, Anthon Rindsäter
Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

2D	Two-Dimensional
3D	Three-Dimensional
5HS	Five-Harness Satin
ASTM	American Society for Testing and Materials
CFRP	Carbon Fibre Reinforced Polymer
CNN	Convolutional Neural Network
CT	Computed Tomography
DIC	Digital Image Correlation
FE	Finite Element
FOV	Field of View
GPU	Graphics Processing Unit
GUI	Graphical User Interface
LS-DYNA	Livermore Software Dynamic Analysis
ML	Machine Learning
ODD	Object-to-Detector Distance
PBC	Periodic Boundary Condition
ROI	Region of Interest
RTM	Resin Transfer Moulding
RVE	Representative Volume Element
SDD	Source-to-Detector Distance
SOD	Source-to-Object Distance
UC	Unit Cell
XCT	X-ray Computed Tomography

Nomenclature

Below is the nomenclature of indices, parameters, and variables used throughout this thesis.

Indices

$1, 2, 3$	Material directions: warp, weft, and through-thickness directions

Parameters

A_0	Initial cross-sectional area of tensile test specimen
b	Pixel binning factor
M	Geometrical magnification
P	Detector pixel pitch
P_{eff}	Effective detector pixel pitch after binning
V	Voxel size
V_f	Fibre volume fraction
μ	Linear X-ray attenuation coefficient
Δ	Denominator in the orthotropic stiffness matrix

Variables

E_1, E_2, E_3	Effective Young's moduli in the 1-, 2-, and 3-directions
G_{12}, G_{23}, G_{13}	Effective shear moduli in the 1–2, 2–3, and 1–3 planes
$\nu_{12}, \nu_{23}, \nu_{13}$	Effective Poisson's ratios in the 1–2, 2–3, and 1–3 planes
F	Applied force
σ_{eng}	Engineering stress
$\sigma_{11}, \sigma_{22}, \sigma_{33}$	Normal stress components
$\sigma_{23}, \sigma_{31}, \sigma_{12}$	Shear stress components
$\varepsilon_{11}, \varepsilon_{22}, \varepsilon_{33}$	Normal strain components
$\gamma_{23}, \gamma_{31}, \gamma_{12}$	Shear strain components
ε_y	Strain in loading direction obtained from DIC

Contents

List of Figures

List of Tables

1
Introduction

1.1 Background

Aero-engines account for a substantial part of an aircraft's overall mass, making weight reductions in the engines particularly effective for improving fuel efficiency. The increased pressure on the aerospace industry to reduce both fuel consumption and emissions has led to a growing need for alternatives to traditional aero-engine materials such as titanium and aluminium alloys. This has driven extensive research into advanced composite materials as potential replacements for the conventional metallic components. Among these possible replacements, carbon fibre reinforced polymers (CFRPs) stand out due to their impressive stiffness- and strength-to-weight ratios, which in many cases exceed those of metallic materials [1]. Hence, the integration of CFRP parts in aero-engines can remarkably lower the total weight of the engine, and subsequently help drive the reduction of fuel consumption. A notable example is the use of composite fan blades and ceramic matrix composites in the LEAP engine [2], which contributed to a 15% reduction in fuel consumption and CO_2 emissions compared to the best performing engine of the previous generation [3, 4].

Beyond the weight reductions, CFRPs offer additional advantages such as a high degree of design flexibility, as their mechanical properties can be tailored for specific load conditions, making them increasingly attractive for aerospace applications [1, 5].

One important subset of CFRPs is textile-reinforced composites. Compared with unidirectional reinforcements, textile reinforcements offer several advantages, particularly for components with complex geometries. For example, double-curved structures are difficult to manufacture using unidirectional reinforcements, whereas textile reinforcements exhibit greater drapability due to their architecture, enabling them to conform more effectively to complex shapes [6]. This makes textile reinforcements especially attractive for aerospace applications, where components such as fan blades, fan cases and engine frames often involve complex geometries. The specific reinforcement architecture investigated in this work is a five-harness satin (5HS) weave, which is introduced in more detail in Section 2.1.1. This textile architecture is commercially available from numerous suppliers with different fibre and binder types as well as with different yarn fibre counts.

To enable the utilisation of CFRPs in aerospace applications, comprehensive material characterisation is required. Composite materials are generally highly anisotropic, making purely experimental characterisation expensive and time-consuming. As an alternative, predictive modelling and numerical simulations can be employed to estimate homogenised material properties. To this end, representative volume elements (RVEs) are commonly used to model the mesostructure. For textile-reinforced composites, RVEs can be utilised to explicitly model the yarn architecture and the surrounding matrix material, enabling the prediction of homogenised material properties through computational homogenisation using finite element (FE) simulations. However, accurate geometric representations of mesoscale features are still required to ensure reliable results.

A simple way to create mesoscale geometrical representations is by describing the 3D textile architecture using modelling software such as TexGen [7]. However, this approach is typically overly simplistic, meaning that it is unable to represent as-manufactured geometries that can include ply nesting, tow shearing and/or misalignments, which significantly affect the resulting elastic material properties. As an alternative, X-ray computed tomography (XCT) is increasingly being used to generate RVE geometries inferred from three-dimensional reconstructions of scans of material samples [8, 9, 10].

Recent advances in XCT imaging and machine learning-based image segmentation have opened new possibilities for the creation of RVEs directly from experimental data. The word segmentation here refers to the act of labelling each voxel (a 3D-pixel) in an XCT reconstruction. In this work, it specifically aims to determine if each voxel is a warp yarn, a weft yarn, matrix or air. TexTomoS is an open-source XCT-to-FE workflow developed by Friemann [11], which predicts the full elastic stiffness tensor of textile-reinforced composites directly from XCT reconstructions [12, 13]. The name TexTomoS is a portmanteau of Textile, Tomography, and Segmentation. The considered "TexTomoS" pipeline is illustrated in Figure 1.1. It employs fully synthetic XCT data to train a machine learning-based segmentation algorithm. This trained neural network can then be used to segment XCT scans of as-manufactured composite samples which can then be directly converted into detailed RVEs suitable for property predictions using FE-modelling. The resulting predictive models enable the estimation of effective mechanical properties based on the yarn architecture and matrix distribution of the scanned material in its final, as-manufactured configuration.

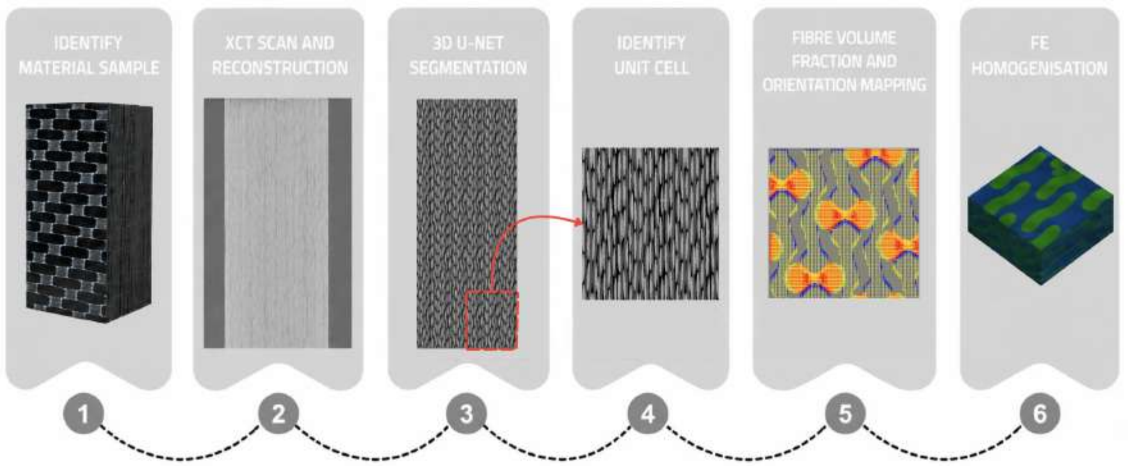

Figure 1.1: TexTomoS pipeline illustration.

This Master's thesis was carried out in collaboration with GKN Aerospace AB, which provided industrial supervision and guidance throughout the project. In this work, the TexTomoS pipeline will be utilised to generate RVEs directly from XCT scans of as-manufactured 2D textile-reinforced composite samples. In its current form, the pipeline is only trained on synthetic data representing 3D layer-to-layer angle interlock composite architectures. However, the main objective of this work is to investigate whether the methodology can be extended to other textile reinforcement architectures beyond those included in the original training data. The emphasis is placed on 2D textile-reinforced composite structures relevant to aerospace engine applications, where complex geometries and manufacturing strongly influence the resulting mechanical behaviour. By generating RVEs directly from the XCT data of as-manufactured materials, the proposed approach aims to enable direct estimation of homogenised elastic properties through computational homogenisation using the FE method.

At the current stage, such an approach is primarily applicable to relatively small material samples, due to limitations in the size of samples that can be scanned and segmented. Nevertheless, XCT-based RVEs allow non-destructive characterisation of mesoscale geometrical features, such as fibre volume fraction, yarn waviness, and defects, while providing a pathway towards linking these features to mechanical performance. As such, the proposed methodology supports the development of material characterisation techniques for complex textile-reinforced composites, with the potential for future extension towards non-destructive inspection and verification at the component level.

1.2 Project Scope

The present work is defined by a set of clearly stated objectives and is conducted within a well-defined scope, subject to practical and methodological limitations that govern the analysis and interpretation of results.

1.2.1 Goals

The aim of this Master's thesis is to evaluate and, where necessary, further develop an existing workflow for predicting the elastic mechanical properties of textile-reinforced composite materials from XCT data. The study focuses on 2D woven textile architectures representative of aerospace engine applications, where accurate prediction of stiffness properties is critical for design and material optimisation. In particular, the work builds upon the XCT-to-FE pipeline developed in [12, 13]. The pipeline combines XCT, machine learning-based segmentation, and computational homogenisation to generate RVEs from as-manufactured composite samples. This thesis applies this workflow to a new material system and assesses its robustness, with the additional aim of identifying whether modifications or extensions to the segmentation model are required to achieve sufficient accuracy for reliable mechanical property prediction. Specifically, this thesis aims to:

- Determine and document XCT scan parameters that enable clear differentiation of material constituents (warp yarn, weft yarn, matrix, air), such that the resulting segmented data is of sufficient quality for reliable voxel-based FE-modelling, while ensuring reproducibility for subsequent XCT reconstructions.
- Manufacture or obtain at least one representative 2D carbon fibre textile composite with an epoxy resin system, suitable for both XCT scanning and mechanical testing, ensuring consistent architecture and material properties between the phases.
- Carry out experimental tensile testing of the considered 2D textile-reinforced composite samples in at least two principal directions to measure the mechanical elastic properties.
- Evaluate the performance of the pre-trained segmentation model of XCT images on 2D textile-reinforced composites using a synthetically generated virtual 5HS sample as a known ground truth. Quantify the performance using pixelwise agreement and Dice similarity coefficients, with a target pixelwise agreement of at least 90%.
- If needed to achieve sufficient segmentation accuracy, generate additional training data to carry out transfer learning on the segmentation model.
- Perform XCT scanning of the manufactured composite samples, and process the reconstructed volumes using neural network-based segmentation within the TexTomoS pipeline to identify the material constituents.
- Apply the TexTomoS pipeline to both virtual and physical samples to generate voxel-based finite element models by mapping local yarn orientation and fibre volume fraction data to each voxel element for mechanical analysis.
- Perform computational RVE homogenisation on the voxel-based finite element models to predict the mechanical elastic properties and validate the results against the obtained experimental data.

1.2.2 Limitations

Due to the complexity of the problem, it is necessary to impose various limitations to ensure a conclusive outcome. The following limitations are therefore applied to each textile type:

- Due to time constraints, only one material system and one type of textile reinforcement will be considered. Additionally, only layups without severe misalignment, shearing or draping will be considered.
- Image segmentation will primarily rely on existing trained models. Retraining or extension of the training dataset will be limited and performed only if necessary.
- FE-analysis will be limited to linear elastic behaviour under small-strain assumptions. Nonlinear and instability effects, including plasticity, viscoelasticity, fatigue, and buckling, will not be considered.
- FE-analysis will be performed exclusively using voxel-based representations obtained from segmented XCT data. Alternative mesh representations will not be investigated.

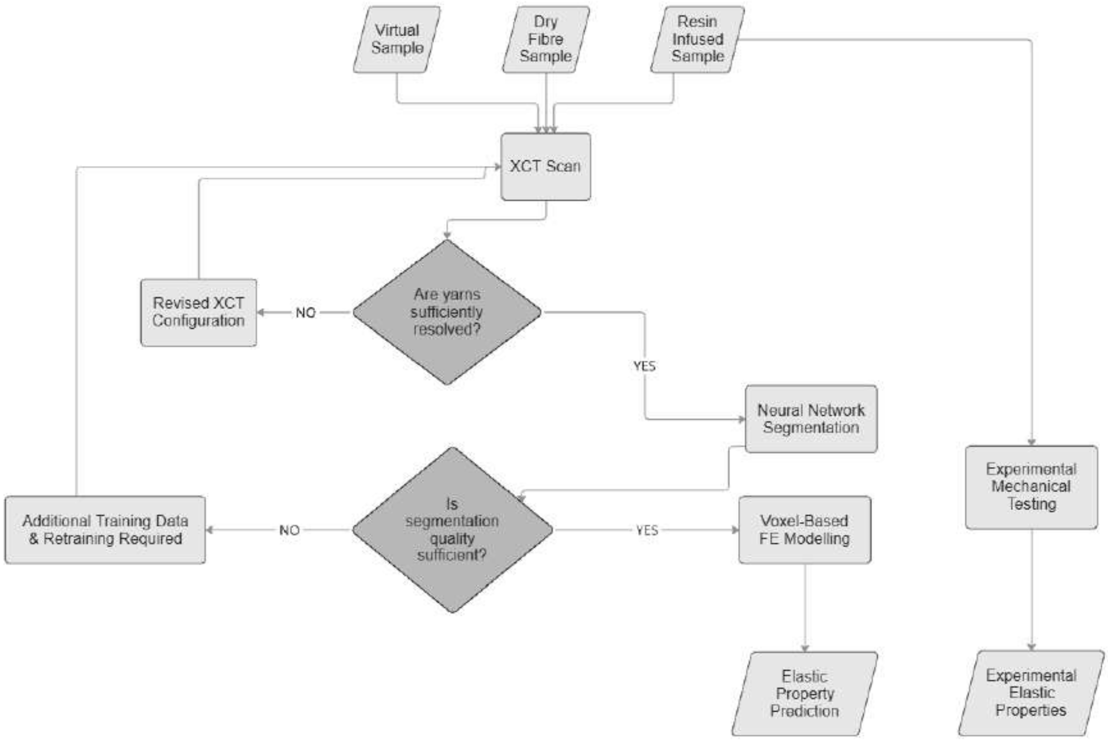

Figure 1.2: Overview of the workflow used in this work.

- Fibre orientation and fibre volume fraction will be determined from segmented image data, with accuracy limited by the resolution of the imaged data and the quality of the segmentation.

1.3 Workflow and Design Considerations

The ability to go directly from XCT reconstructions to as-manufactured material properties is an incredibly powerful tool. As previously mentioned, the existing pipeline, presented in [12, 13], was originally developed and trained using 3D layer-to-layer angle interlock composite architectures. In this work, the pipeline is instead applied to a 2D 5HS woven composite architecture. The network has never seen this type of woven architecture. It is of interest to know if the network has the ability to generalise and segment new composite architectures. Both virtual and physical investigations are considered to evaluate the transferability of the existing XCT-to-FE pipeline when applied to 2D textile-reinforced composite architectures. An overview of the workflow considered in this work is presented in Figure 1.2.

To assess the performance of the pipeline under different conditions, three material samples are considered: a virtual sample, a resin-injected sample, and a dry fibre sample. Each of these samples fulfils a specific role within the work:

- **Virtual TexGen model:** The virtual sample begins as a TexGen model of a representative weave geometry. An example of which is shown in Figure 1.3a. This sample is then virtually XCT scanned using a software package known as gVXR. An image of which can be seen in Figure 3.8. One benefit of this model is that it enables virtual XCT investigations where scan parameters can be fine-tuned directly without needing physical rescanning of the sample. Another direct benefit of this model is that the material phases can be labelled directly from the initial TexGen model. This provides a clear ground truth for evaluating how well the neural network segments an XCT reconstruction.
- **Resin-injected composite sample:** This is a real composite material sample, sectioned

from a composite plate and shown in Figure 1.3b. It has been scanned using parameters similar to those identified for the virtual model. The physical sample is subsequently used to evaluate whether this textile architecture can be segmented using the TexTomoS pipeline and whether representative mechanical properties can be obtained.
- **Dry fibre sample:** Shown in Figure 1.3c, this sample provides physical XCT imaging with increased attenuation contrast between the yarns and surrounding air using scan parameters corresponding to the virtual investigations. The dry fibre sample was prepared as a fallback if it proved not possible to achieve an XCT scan with sufficiently high attenuation contrast for the resin-injected sample.

(a) Virtual Sample **(b)** Resin-injected sample **(c)** Dry-fibre sample

Figure 1.3: Depiction of the three samples considered in this work.

1.4 Sustainability and Ethics

The environmental impact of the aerospace industry inevitably raises ethical questions regarding the necessity of air travel. While the development of lightweight structural components contributes to reductions in fuel consumption and associated emissions, it also highlights the debate regarding the long-term sustainability of aviation as a transport method. Simultaneously, advancements in lightweight structures could aid the development of future propulsion concepts, including the use of alternative energy sources such as sustainable aviation fuel and, in the longer term, electrified propulsion systems.

Related to this work, CFRPs play a significant role due to their high specific stiffness and strength, enabling substantial weight savings over conventional metallic materials. These weight reductions lead to improvements in energy efficiency over the lifetime of a component. However, CFRPs also introduce environmental challenges related to manufacturing. In particular, carbon fibre production is an energy-intensive process and the limited recyclability of CFRPs contributes to their environmental footprint. However, developments related to sustainable manufacturing methods and viable recycling techniques could help mitigate these issues.

This work is exclusively focused on the application of CFRPs in commercial aircraft engine components. It should be noted, though, that GKN Aerospace is active across both civil aviation and defence-related applications, and that advancements in CFRP technologies may have potential relevance to both sectors.

Artificial intelligence was used to assist with the enhancement of selected images where improvements in, for example, clarity or exposure were required. In all cases where such enhancements were applied, this is explicitly stated in the corresponding figure captions. The enhancements were performed using ChatGPT by OpenAI [14]. Artificial intelligence was also used to support consistency in writing style.

2

Materials and Textile Architecture

2.1 2D Woven Fabrics

A 2D woven fabric is constructed by a conventional weaving process consisting of interlacing two orthogonal sets of yarns forming a fabric architecture as illustrated in Figure 2.1. In woven fabrics, warp yarns are held under tension, while the weft yarns are inserted over and under the warp yarns. Due to the interlacing of warp and weft yarns during the weaving process, both yarns exhibit a certain degree of crimp. The magnitude of crimp depends on the specific weave architecture and manufacturing parameters, such as yarn tension and spacing. The four most commonly used weave styles are plain, twill, crowfoot and satin [15].

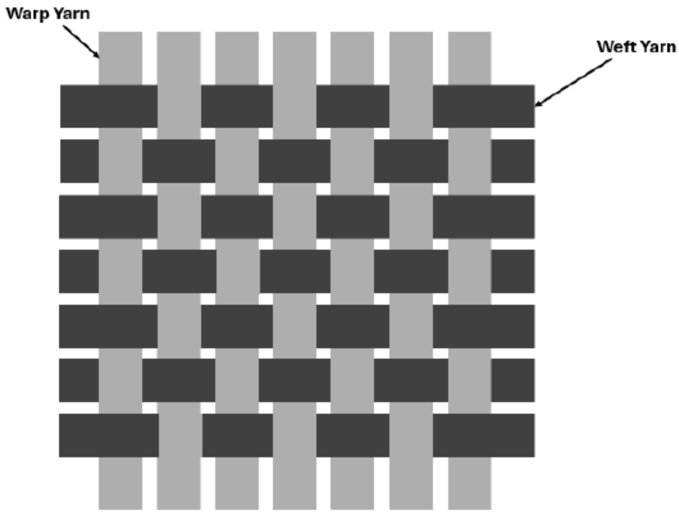

Figure 2.1: Illustration of weft and warp yarns in a 2D woven plain weave.

The type of weave style directly affects the properties of the fabric. For example, a plain weave, where warp and weft yarns are consecutively placed over each other, demonstrates good stability but also generates higher crimp levels in the structure, having a negative effect on the mechanical properties [16]. Satin weaves involve long yarn floats and few interlacements, creating a higher drapability of the fabric in contrast to plain weaves, and an enhancement in mechanical properties due to a decrease in crimp [15]. However, it should be noted that this increase in drapability leads to less stability in the material which can cause shearing of the yarns, especially in highly curved structures.

2.1.1 Five-Harness Satin Weaves

In a 5HS weave, warp and weft yarns are interlaced to form a coherent textile structure, as illustrated in Figure 2.2. Compared to plain weave architectures, satin weaves exhibit improved drapability, enabling the fabric to conform more easily to curved surfaces [17]. Several satin weave configurations exist, including 4-harness, 5-harness, and 8-harness variants, where the harness number denotes the number of warp yarns involved in the repeating pattern. In a 5HS weave, each weft yarn passes over four warp yarns and under one within a single repeating unit, which is depicted in Figure 2.2b. Woven fabrics can be manufactured with either the same or different number of carbon filaments in the warp and weft yarns. Typical examples of fibre count in the tows are 3k, 6k, and 12k, where "k" indicates thousands of individual carbon filaments.

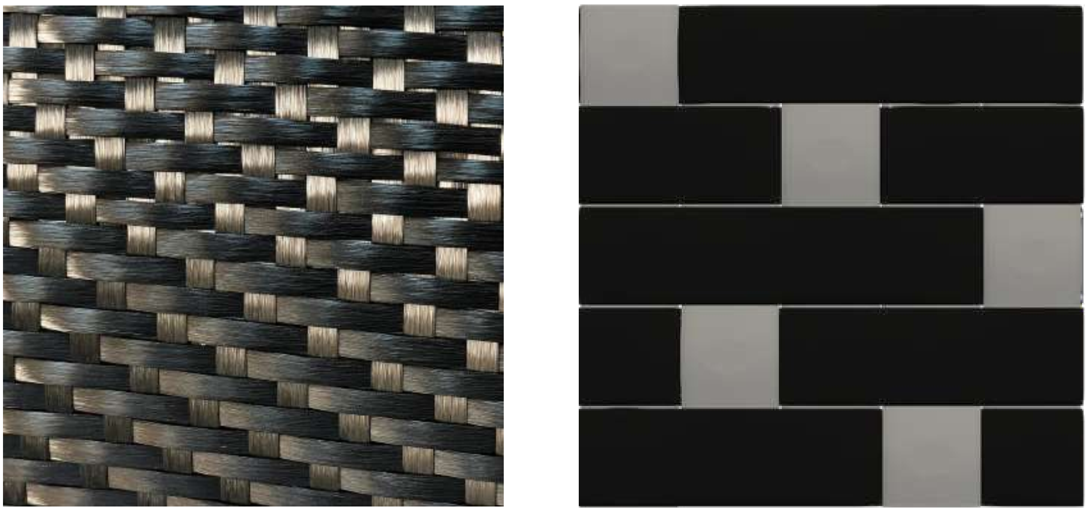

(a) Example of a 5HS fabric. Image clarity and brightness enhanced with ChatGPT.

(b) 5HS weave pattern: Weft (black) and warp (grey).

Figure 2.2: Five-harness satin (5HS) weave structure, showing an example of the fabric and the corresponding idealised weave pattern generated in TexGen.

2.2 Virtual Textile Modelling (TexGen)

TexGen is an open-source software licensed under the General Public License developed at the University of Nottingham for modelling textile structures. The software enables the creation of a wide variety of predefined textile structures, as well as the generation of custom structures through either the graphical user interface (GUI) or the use of Python scripts. In Figure 2.3, a 2D weave, a 3D weave, and a triaxial braid model generated in TexGen are presented. TexGen also provides methods for computing fibre volume fractions both within the yarns and within the whole domain. The created textiles can be exported in a variety of different formats, such as surface or voxelised meshes.

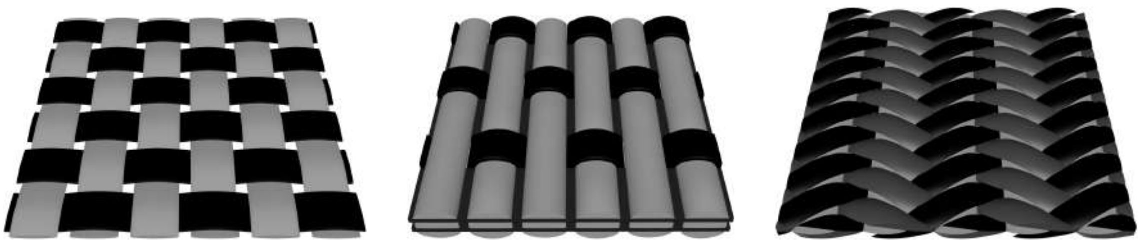

(a) 2D Weave **(b)** 3D weave **(c)** Triaxial braid

Figure 2.3: Different weave/braid types generated in TexGen [7].

2.3 Preparation of Physical Samples

The composite plates used to produce the physical samples in this study were manufactured using a resin transfer moulding (RTM) process, where a dry 5HS carbon fibre preform was infused with an epoxy resin system. The plates consisted of a 5HS 6k woven architecture with a $[0]_9$ layup, resulting in a nominal fibre volume fraction of approximately 57%. These plates form the basis for both the physical samples and the virtual model developed in this work.

The plates were waterjet cut into samples with dimensions suitable for the XCT system used in this study, along with larger specimens reserved for experimental mechanical testing, as shown in Figure 2.4a.

The smaller samples were then stacked and glued together to achieve the required out-of-plane thickness and aspect ratio for XCT scanning. The resulting resin-injected sample is shown in Figure 2.4b. The coordinate system shown in Figure 2.4b defines direction 1 as the warp, 2 as the weft, and 3 as the through-thickness direction. This convention is used throughout the study.

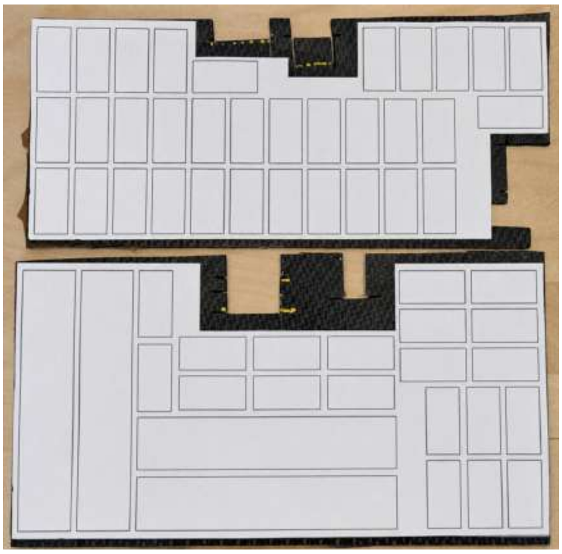

(a) Preparation of composite panels for waterjet cutting. Image brightness enhanced with ChatGPT.

(b) Stacked and bonded resin-injected composite sample.

Figure 2.4: Preparation of physical resin-injected composite samples for XCT scanning.

In addition to the resin-injected samples, two dry fibre preforms were also manufactured as backup samples in case segmentation of the resin-injected samples proved difficult due to low attenuation contrast in the XCT reconstructions. This was done by cutting out rectangles with two different sample sizes of the 5HS weave on the cutting table available at GKN. Two dry fibre sample dimensions were manufactured: 22×44 mm^2 to match the resin-injected samples, and 30×60 mm^2 to create a margin to the sample edges during XCT scanning due to fraying. After the material had been cut, 3D printed guides were utilised to stack 54 layers for each sample. The samples were sandwiched between two steel plates with shims, see Figure 2.5a, in order to reach a "cured" thickness of roughly 21.6 mm. The setup was then vacuum bagged and the samples were subsequently consolidated in the oven at 120°C for 20 minutes, see Figures 2.5b-2.5c.

(a) Dry fibre stacks setup before adding second plate on top.

(b) Setup placed inside vacuum bag.

(c) Consolidated dry fibre samples.

Figure 2.5: Manufacturing of dry fibre samples.

2.4 Virtual 2D Textile Generation

Based on the manufactured physical plate described previously, a virtual model was developed for use in virtual XCT simulations. In this study, a previously manufactured 2D textile-reinforced composite plate was used as the reference for modelling the virtual sample. A corresponding 5HS woven textile composite was generated in TexGen, with the aim of representing the physical architecture as closely as possible rather than relying on an idealised geometry. The model was, as previously stated, based on a 5HS 6k composite plate consisting of nine plies with a $[0]_9$ layup, the weft direction facing upwards, and a fibre volume fraction of 57%. To further improve the agreement between the virtual and physical architectures, microscopy data from the manufactured plate, shown in Figure 2.6, was incorporated into the virtual textile modelling.

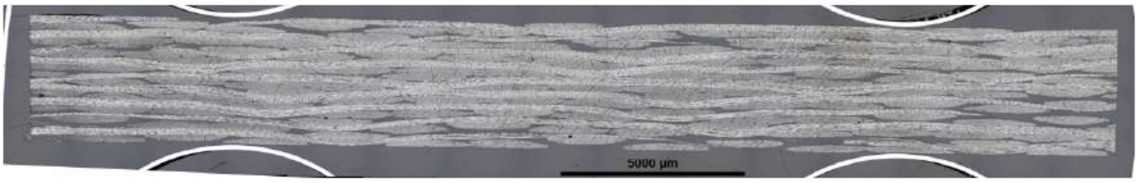

Figure 2.6: Microscopy image of physical resin-injected plate.

The modelling parameters extracted from the microscopy results are summarised in Table 2.1. Based on these parameters, the yarn cross-sectional shape was idealised as elliptical to provide a simple representation consistent with the observed geometry. These parameters were used to generate an RVE with in-plane dimensions of 11 × 11 mm. The RVE was replicated nine times in the out-of-plane direction to represent the $[0]_9$ plate layup. The resulting nine-ply stack was then further extended in the out-of-plane direction to improve the voxel aspect ratio during XCT scanning, giving an initial virtual sample size of $11 \times 11 \times \sim 10$ mm. The virtual model had a fibre volume fraction of 56.75%, which was close to the target value of 57%. Finally, the virtual sample was exported as three surface meshes corresponding to the weft, warp and matrix phases, enabling its use in subsequent XCT simulations with the pipeline developed in TexTomoS.

Table 2.1: Values used to model the virtual sample based on microscopy results.

Dimension	Value	Unit
Yarn Width	2.2	mm
Yarn Spacing	2.15	mm
Ply Height	0.37	mm

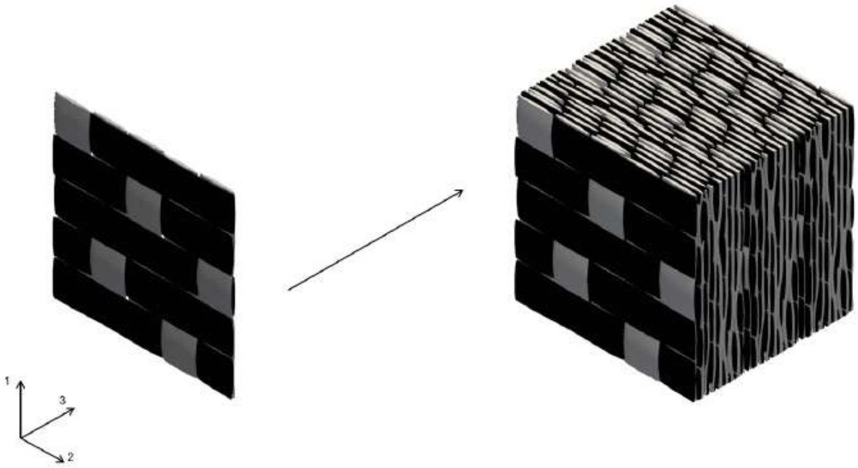

Figure 2.7: RVE generated in TexGen (left) and the replicated virtual sample representing the 5HS composite architecture (right).

2.5 Generating Experimental Material Data

In order to validate the elastic stiffness predictions made from the input files generated by the TexTomoS pipeline, it is important to have experimental stiffness data of the material. For this reason, a limited experimental testing campaign was carried out. Tensile testing was performed in both in-plane reinforcement material directions to provide experimental data for comparison with the numerical simulations.

2.5.1 Experimental Test Setup

As previously mentioned, the specimens used for mechanical testing were waterjet cut from an existing composite. In total, 4 samples were prepared for mechanical testing, where the first two were oriented in the warp direction and the latter two in the weft. All four specimens are shown in Figure 2.8a. The samples were 170 × 34 × 3.3 mm in dimension, where a gauge region of 80 mm in length was considered.

An image of a representative test setup is depicted in Figure 2.8b. The tensile tests were performed using an MTS 810 Material Test System machine with a load capacity of 100 kN. The samples were loaded at a machine displacement rate of 1 mm/min, with the aim of excluding any significant strain-rate effects. Furthermore, in order to measure the deformations of the specimens, digital image correlation (DIC) was used. DIC is a non-contact and non-destructive optical measurement method which aims to capture full-field surface displacements and strains by tracking specific points on the specimen during loading [18]. Specifically, a ZEISS ARAMIS [19] system with two stereo cameras was employed. An image sampling rate of 2 Hz was selected, meaning that two images were taken every second throughout the test. To ease post-processing, the force signal provided by the testing machine was synchronised to each image taken by the DIC system.

There are many methods that can be used to apply a speckle pattern to the test specimens. The method considered in this work was the use of white and black spray paint. An example of the pattern on one of the samples is shown in Figure 2.9. This method does not harm the specimen, does not require costly equipment and is easy to perform. For this reason, it is a popular choice amongst researchers [20]. The application of the paint involved first spraying a white base paint onto the specimen, followed by the black spray paint sprayed from a distance so that random speckles fall onto the specimen's surface. It should be noted that getting a high-quality speckle pattern is important and directly affects the accuracy of the results [20].

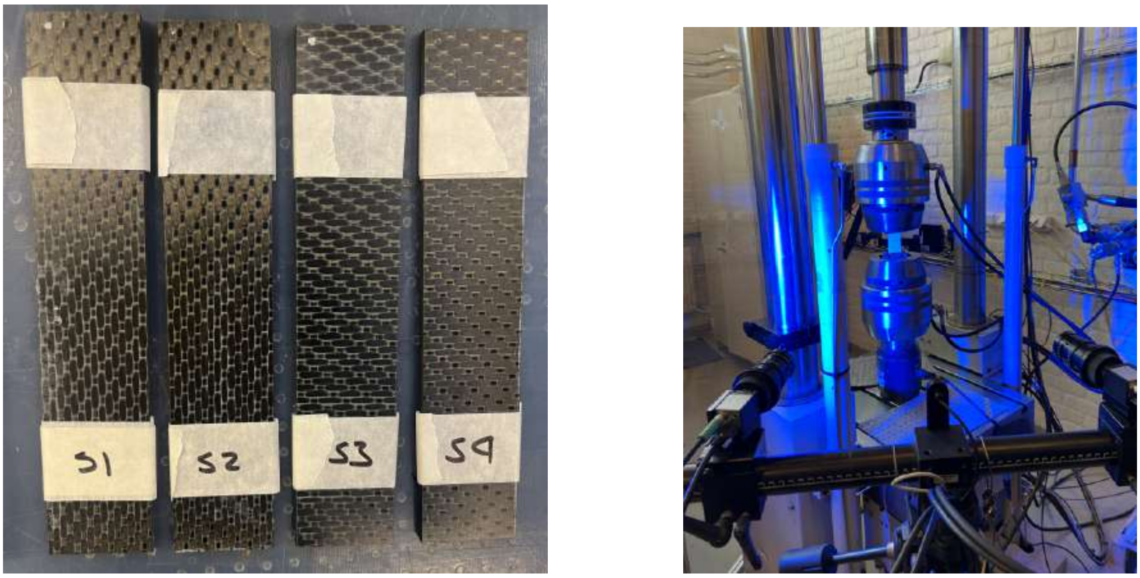

(a) Mechanical testing samples. **(b)** Representative test setup.

Figure 2.8: Experimental setup and samples used in mechanical testing.

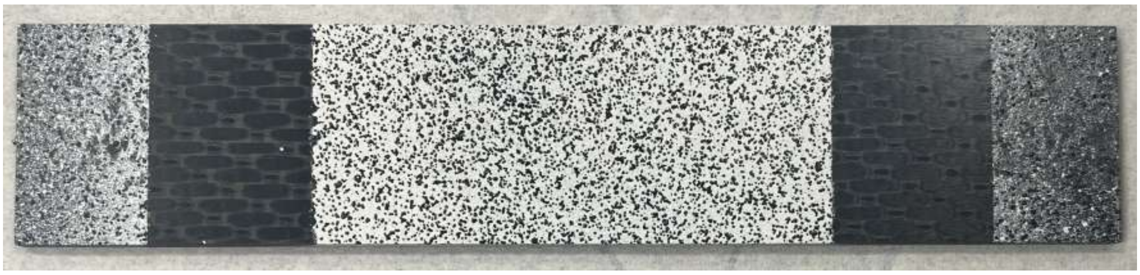

Figure 2.9: Example of artificially applied speckle pattern on specimen using spray paint.

2.5.2 Experimental results

Post-processing of the strain signal was conducted in ZEISS Correlate. The strain was calculated by applying virtual extensometers between two points in the sample. In order to compute a representative global strain value, rather than a local measurement, the virtual extensometer must be of a sufficient length and span multiple unit cells (UCs). A gauge length of approximately 60 mm was used for the virtual extensometer. Figure 2.10 illustrates the evolution of the vertical strain field (ε_y) for specimen S2. At an intermediate strain level of approximately 0.5%, a relatively uniform strain field is observed, with a noticeable concentration developing near the upper grip region. As the strain increases to approximately 1.16%, the majority of the specimen reaches a vertical strain value of around 1.1%, with only small regions remaining at lower strain levels. A pronounced localised strain concentration of approximately 1.35–1.45% develops, which corresponds to the eventual failure location, as confirmed by the post-mortem image. At the point of failure, the high dynamic loads caused some of the spray paint from the specimen's surface to debond.

The post-processed force-strain data was exported from ZEISS, and the force was converted into engineering stress through

$$\sigma_{\text{eng}} = \frac{F}{A_0}, \tag{2.1}$$

where F is the applied force and A_0 is the original cross-sectional area of the specimen. From the resulting stress-strain curve, the Young's modulus for each sample was extracted from the linear regions of the plots.

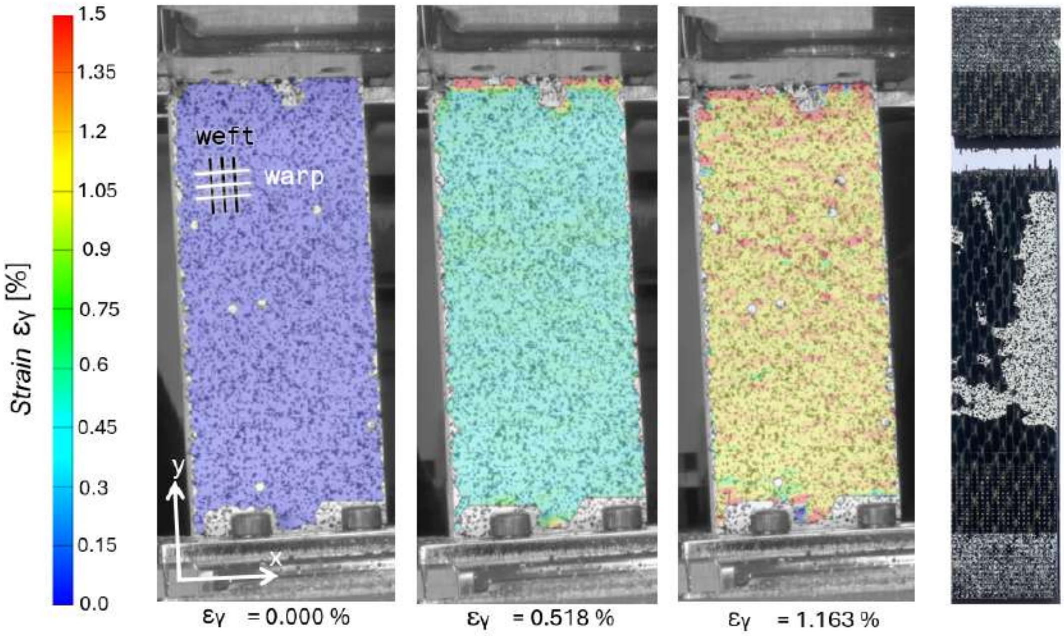

Figure 2.10: Images showing the full strain-field ε_y throughout the test as well as a post-mortem image of test sample S2.

A plot of the experimental stress-strain results from the experimental mechanical testing is presented in Figure 2.11. The first two samples, S1 and S2, were tested in the weft reinforcement direction, while S3 and S4 were tested in the warp reinforcement direction. A least-squares fit approach was taken when calculating the derivative of the stress-strain curve in order to reduce the sensitivity to experimental noise. In accordance with ASTM D3039/D3039M-17(2025), the elastic mechanical properties were extracted within the strain range 0.1% - 0.3%, illustrated by the grey area depicted in Figure 2.11. A corresponding table of the Young's moduli extracted from each specimen is presented in Table 2.2, together with the mean and standard deviation for each reinforcement orientation.

Table 2.2: Summary of tensile test results.

Specimen	**Orientation**	**E [GPa]**
Individual Test Results		
S1	Weft	63.44
S2	Weft	62.61
S3	Warp	61.75
S4	Warp	61.38
Overall Test Results		
S1/S2	Weft	63.03 ± 0.59
S3/S4	Warp	61.57 ± 0.26

The stress-strain responses presented in Figure 2.11 show a consistent mechanical behaviour across the four tested specimens. The overall shape of the stress-strain plot demonstrates an initially linear response followed by a gradual stiffening after a strain magnitude of approximately 0.4%. This stiffening effect is consistent with what is commonly observed in textile CFRPs. In particular, reinforcement reorientation and tow straightening are likely responsible for the observed stiffening [21]. At low strain levels, the waviness of the yarns reduces the effective stiffness. However, as the applied strain increases, the yarns progressively straighten, leading to a more efficient load transfer along the loading direction and consequently a higher effective stiffness.

To further visualise this non-linearity and assess the sensitivity of the extracted Young's moduli within the applied strain range, a sliding window analysis was also performed over overlapping strain windows ranging from 0.1% - 1.1%. The results of the sliding strain window analysis are presented in Figure 2.12. This analysis indicates that an increasing strain level leads to an increase in the extracted Young's moduli for all specimens, further confirming the non-linearity of the stress-strain response.

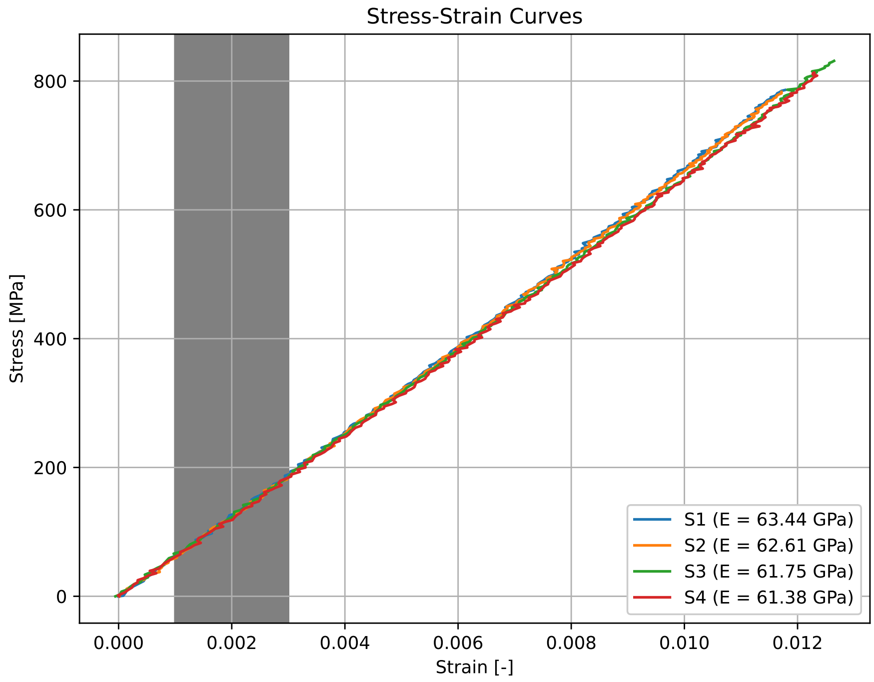

Figure 2.11: Stress-strain curves of all test specimens with corresponding stiffness properties. The grey highlighted area indicates the strain range from which Young's modulus was computed.

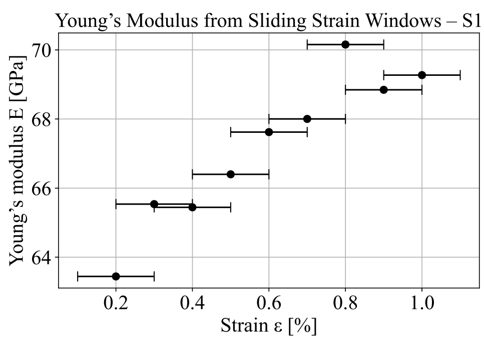

(a) Specimen S1 (weft).

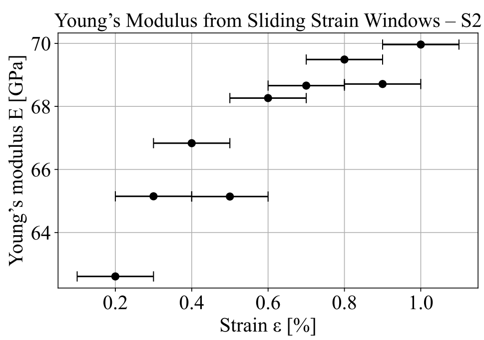

(b) Specimen S2 (weft).

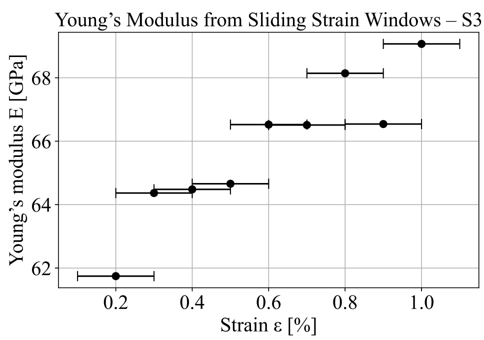

(c) Specimen S3 (warp).

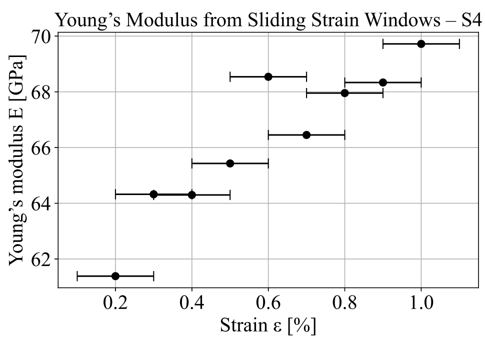

(d) Specimen S4 (warp).

Figure 2.12: Young's modulus obtained from sliding strain window analysis for all four specimens.

The Young's moduli reported in Table 2.2 are comparable among specimens tested in the same loading direction. For the weft direction, an average modulus of 63.03 ± 0.59 GPa was obtained, while the warp direction showed a slightly lower value of 61.57 ± 0.26 GPa. Here, the values after $\pm$ represent the standard deviation between the two specimens tested in each orientation. Although the low standard deviations suggest good repeatability of the experimental procedure and a relatively uniform material response within the tested specimens, the limited number of samples means that these values should not be considered statistically representative.

The difference between the stiffness in the weft and warp directions indicates that the 5HS composite exhibits a nearly balanced response. This is consistent with the 5HS architecture, where both the weft and warp yarns have the same fibre content. However, the slightly higher stiffness observed in the weft direction may be due to subtle differences in yarn crimp, packing density, or fibre alignment introduced during the weaving and manufacturing process.

It is, however, often expected that the warp direction exhibits a slightly higher stiffness compared to the weft direction. This is typically due to the manufacturing process, where warp yarns are held under tension during weaving, which may lead to an improved fibre alignment and reduced crimp. Since yarn crimp has been shown to influence the mechanical properties and stiffness of woven composites, a lower crimp level in the warp direction may therefore contribute to a slightly higher effective Young's modulus [22]. It is worth noting that this effect is not universal, but rather depends on the specific weave architecture and manufacturing conditions which may lead to the weft direction exhibiting a slightly higher stiffness compared to the warp direction, as is the case in Table 2.2.

It should also be stated that the experimental dataset is very limited, as only four specimens were tested, with two samples in each principal material direction. Furthermore, the measurements were restricted to in-plane tensile loading, meaning that no information regarding the out-of-plane response or shear properties could be obtained. Consequently, the experimental results should primarily be viewed as initial validation data points for comparison with the predicted mechanical elastic properties in Section 5.3, rather than a comprehensive mechanical characterisation of the material system.

3

X-Ray Computed Tomography

3.1 Principles of X-ray computed tomography

X-ray computed tomography (XCT) is a non-destructive evaluation technique that enables the 3D characterisation of internal structures within a material without damaging the specimen. In practice, however, the sample may need to be cut to a suitable size, which can make the process partly destructive. X-ray computed tomography captures a series of projectional radiographs, also known as 2D radiographs, by rotating the specimen, which are then reconstructed to display a 3D volume [23]. Feature detection within the volume of interest relies strongly on the contrast between different material phases, which arises from the different linear attenuation coefficients, μ, of the constituents [24]. The X-ray linear attenuation coefficient is a parameter that determines the extent to which photons are absorbed and scattered within a material, thereby influencing the resulting greyscale contrast [25]. Software such as VGStudio MAX, used in this case, can be used to reconstruct, visualise and analyse the resulting XCT data. Examples of two XCT slices of the considered material architecture are presented in Figure 3.1. The coordinate system introduced previously is used consistently throughout this report, where direction 1 corresponds to the warp, 2 to the weft, and 3 to the through-thickness direction.

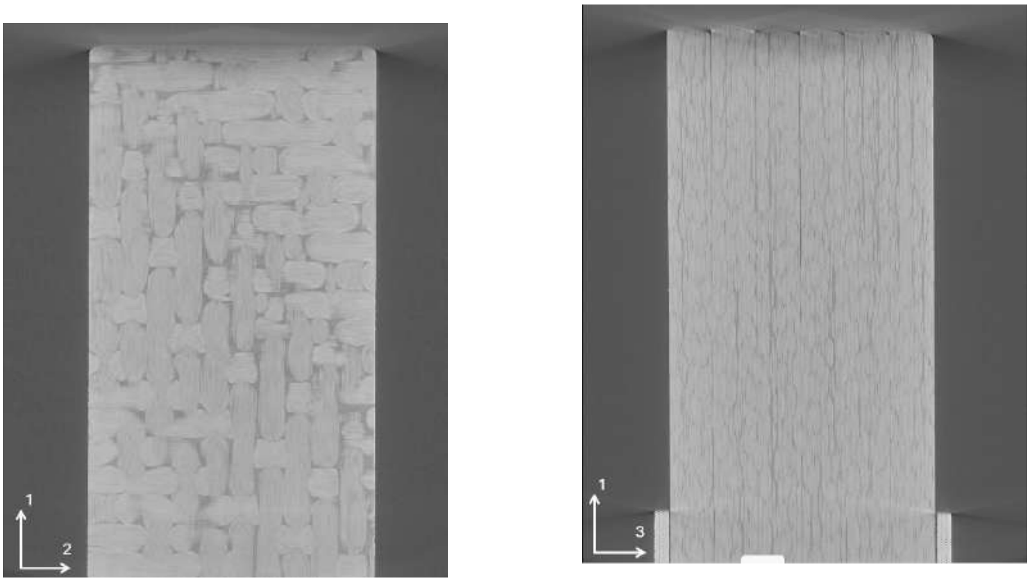

(a) Slice in the warp–weft plane (12).

(b) Slice in the warp–through-thickness plane (13).

Figure 3.1: Example XCT slices of the same resin-injected sample at different orientations. Direction 1 corresponds to the warp, 2 to the weft, and 3 to the through-thickness direction.

3.1.1 Scan Geometry

In XCT, the quality and resolution of the reconstructed volume is not only dependent on the material contrast, but also on the geometric arrangement of the XCT system. This arrangement, denoted the *scan geometry*, describes the relative positions of the X-ray source, specimen and detector. An illustration of the scan geometry is shown in Figure 3.2. Three key parameters must be considered when defining the scan geometry: the sample size, the field of view (FOV), and the voxel size.

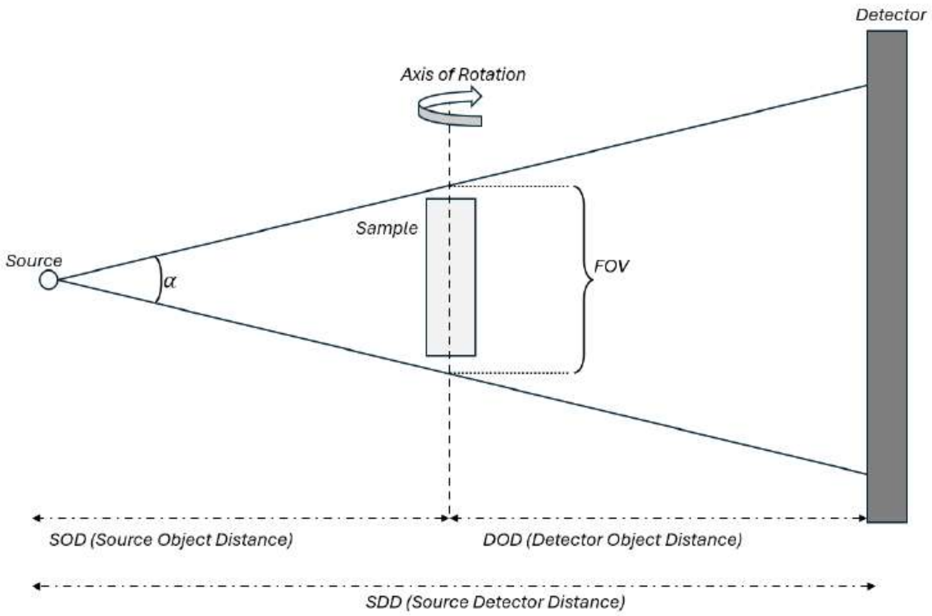

Figure 3.2: Overview of scan setup.

The size of the sample directly affects the amount of X-ray signal transmitted through the specimen, which in turn influences image quality. Larger samples attenuate more X-rays and therefore require longer scan times to maintain an adequate signal-to-noise ratio [26]. The signal-to-noise ratio is the strength of the signal compared to the background noise. The FOV represents the physical region of the sample that can be captured within a single scan. If the sample exceeds the FOV, it cannot be fully imaged within the cone beam angle, α, leading to truncated data and reconstruction artefacts within the visible region [27]. The FOV is further constrained by the voxel size [28], also known as a volumetric pixel [29]. A smaller voxel size provides higher spatial resolution, but also reduces the FOV. The voxel size is governed by the geometrical magnification of the scan, which depends on the source-to-object distance (SOD) and the detector pixel pitch, defined as the distance between the centres of adjacent detector pixels. For a detector with pixel pitch P and a target voxel size V, the magnification M and corresponding source-to-object distance (SOD) can be expressed as

$$M = \frac{P}{V}, \quad SOD = \frac{SDD}{M}. \tag{3.1}$$

where SDD is the source-to-detector distance. In addition to the detector pixel pitch, pixel binning can be applied to shorten the XCT scan time and reduce the dataset, at the cost of spatial resolution [30]. Pixel binning entails that neighbouring pixels are grouped together to form a larger effective pixel. For example, a binning factor of 2 combines four pixels into one and a binning factor of 4 combines 16 pixels into one, as illustrated in Figure 3.3. By grouping multiple pixels together, the detected signal increases, improving the signal-to-noise ratio. When binning is applied, the effective pixel pitch becomes $P_{\text{eff}} = bP$, where b is the binning factor. Substi-

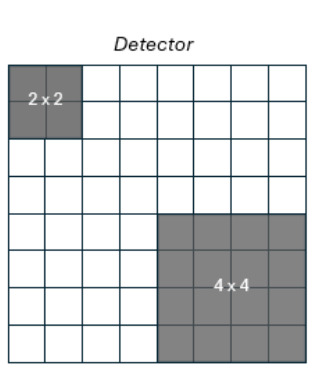

Figure 3.3: Illustration of pixel binning where multiple pixels are combined into one effective pixel.

tuting this into Eq. 3.1, the magnification can be rewritten as

$$M = \frac{P_{\text{eff}}}{V} = \frac{bP}{V}. \tag{3.2}$$

Once the SOD is defined, the FOV can be defined utilising simple trigonometry, and in extension the maximum allowable sample size in the vertical and horizontal directions. This geometric relationship shows how the scan geometry directly affects the voxel size, FOV, and sample size, emphasising the importance of considering them together prior to XCT scanning.

3.1.2 Opportunities and Challenges in XCT of CFRPs

X-ray computed tomography scans of composites can be performed to analyse a variety of aspects. This can include the quality of the manufacturing processes through the identification of voids, foreign object debris, cracks, delaminations, etc. It can also be used to understand failure morphology in terms of for example fibre kinking, fatigue crack development as well as impact damage [24]. Under the right conditions, high-resolution inspection can also reveal fibre orientation and distribution [31].

Although XCT of CFRPs has gained popularity in recent years, practical limitations still affect its application. A primary challenge associated with XCT of composites is obtaining a high contrast between yarns and epoxy due to their similar density and low atomic number, resulting in low linear attenuation coefficients, μ [32].

As illustrated in Figure 3.4, the contrast between dry and resin-injected samples differs significantly even at the same voxel size. Increasing the X-ray exposure time can improve image contrast, although this is limited by the XCT system. As a consequence, reduced contrast can obscure features and make it difficult to distinguish between the resin and the weft and warp yarns. In such regions, the greyscale values in the XCT images may not correspond clearly to a single material, particularly near yarn boundaries, where a voxel may contain both fibre and resin. This can produce intermediate greyscale values, making accurate phase separation difficult. These difficulties in separating the phases can affect both material classification and orientation mapping [33]. For example, incorrectly classified voxels may lead to errors in the estimated fibre, matrix, and void distributions. In addition, poorly defined yarn boundaries can make the local fibre orientations less reliable.

These errors can then be carried into FE-models, where the segmented phases and orientation fields are used to assign material properties. Since the elastic response of carbon fibre composites depends strongly on fibre direction and fibre volume fraction due to their highly anisotropic nature, incorrect phase or orientation assignment can lead to inaccurate effective stiffness predictions. For instance, if fibres are assigned the wrong local direction, they will contribute stiffness in the wrong direction. Similarly, if fibres and matrix are incorrectly classified, the estimated fibre volume fraction will change, leading to an inaccurate representation of the effective mechanical properties.

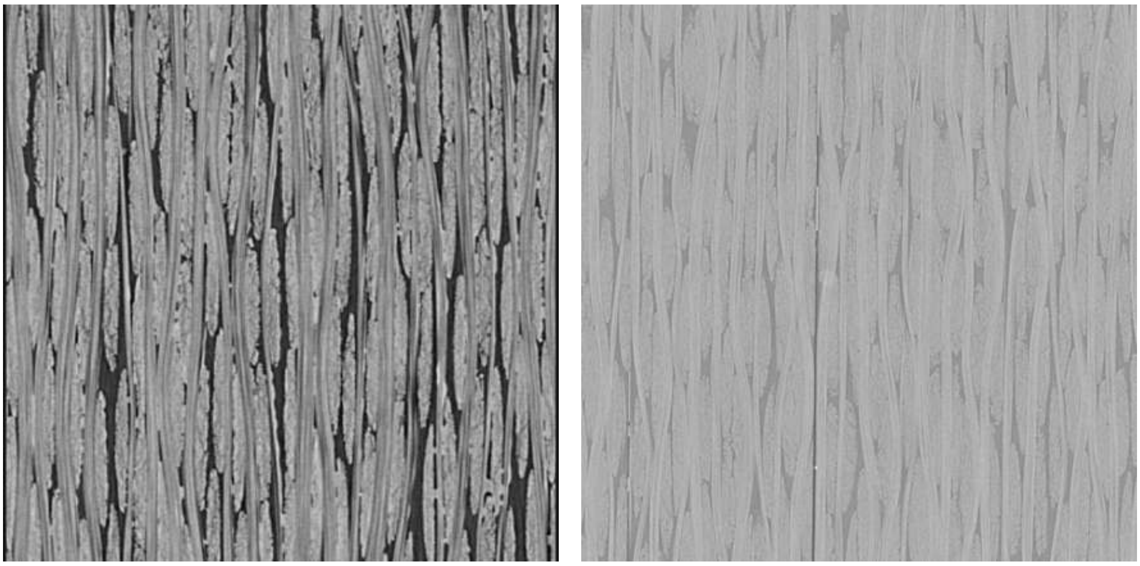

(a) Dry sample (23 μm). **(b)** Resin-injected sample (23 μm).

Figure 3.4: Comparison of XCT slices at equal spatial resolution, illustrating the difference in contrast between a dry and a resin-injected sample. The presence of resin decreases contrast between yarns and the surrounding epoxy matrix material.

Secondly, the spatial resolution of the image affects the level of detail that can be distinguished [33]. In composite materials, this is particularly important, as in some cases individual filaments within a tow may need to be resolved. An increase in spatial resolution, corresponding to a decrease in voxel size, allows finer details to be captured. However, this results in a trade-off of a reduced FOV, as previously explained in Section 3.1.1, limiting the maximum size of the sample that can be scanned and visualised. To address this trade-off between spatial resolution and FOV, varying approaches can be utilised. One possibility is to combine multiple scans into a single larger reconstructed image, enabling a broader field of view to be obtained while preserving a small voxel size. Another possibility is that a chosen region of interest (ROI) can be scanned at a higher resolution instead, allowing for detailed analysis of critical areas within a specific sample [33].

Figure 3.5 shows a comparison of XCT scans of the same material obtained at varying voxel sizes, illustrating how spatial resolution influences the level of detail that can be captured. Even when a smaller voxel size is used, the ability to carry out automated grey scale segmentation is challenging. Traditional thresholding-based segmentation algorithms, such as Otsu's method, Li's method, and the Triangle method ([34], [35], [36]), are insufficient for reliably distinguishing between warp yarns, weft yarns, and matrix in such scans. As one can imagine, the use of manual segmentation is also not a feasible solution and would require an exorbitant amount of labour to manually segment the hundreds of slices in each reconstruction. This work aims to overcome these challenges.

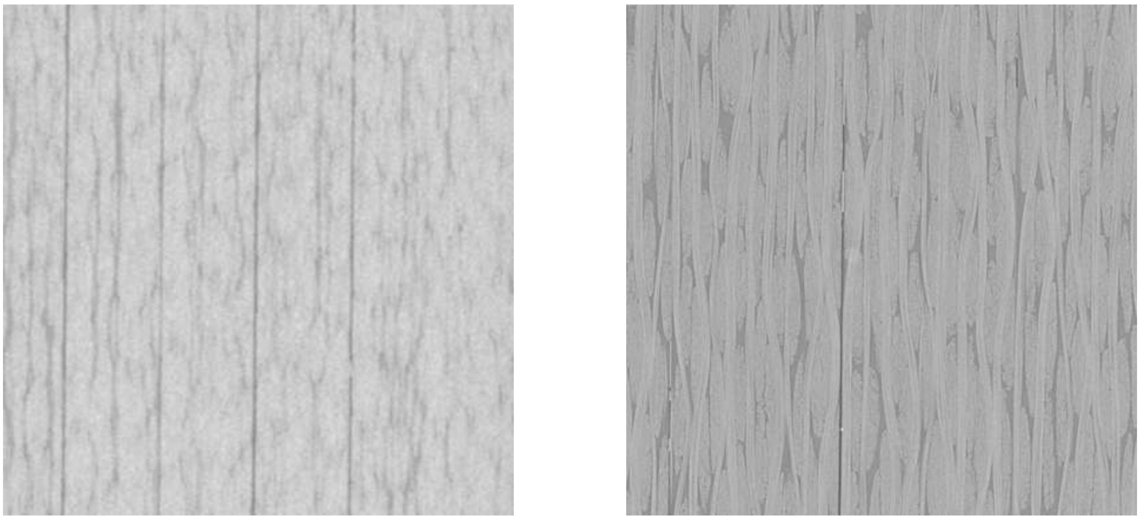

(a) Larger voxel size (46 μm). **(b)** Smaller voxel size (23 μm).

Figure 3.5: Comparison of spatial resolution in XCT scans of a resin-injected composite. A smaller voxel size improves the level of detail captured in the image.

3.2 Virtual XCT (gVirtualXray)

gVirtualXray (gVXR) is an open-source framework that utilises the Beer-Lambert law to simulate X-ray images in real time. The Beer-Lambert law describes the reduction in intensity, I, of monochromatic light as it passes through a material, and can be expressed as:

$$I = I_o \, e^{-\mu x} \tag{3.3}$$

where I_o is the intensity of the incident radiation, μ the linear attenuation coefficient and x the thickness of the material, as illustrated in Figure 3.6. The gVXR framework is implemented on the graphics processing unit (GPU) using the OpenGL shading language. SimpleGVXR is a smaller library built on top of gVXR which provides wrappers to Python, among many other programming languages, and is the library that will be utilised in this work. The framework has been applied in a wide range of applications, including real-time medical simulation, clinical radiographic imaging, fluoroscopy noise reduction studies, education in X-ray imaging and particle physics, as well as the prediction of image quality and artefacts in materials science [37]. In addition, gVXR has been used to produce a high number of realistic simulated images in optimisation problems and to train machine learning algorithms [37].

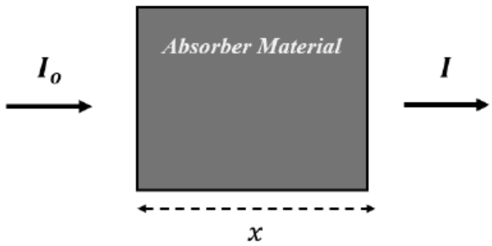

Figure 3.6: Illustration of reduction of monochromatic light passing through an absorber material.

3.2.1 Ring Artefacts in XCT Imaging

Computed tomography is a method which has many advantages compared to other 3D measurement methods, with the largest being the possibility to perform a non-destructive measurement of an object's inner geometry [38]. As computed tomography is a complex method with numerous parameters influencing the measurement results, some parameters can induce image artefacts that arise in the scanning and reconstruction process [38]. An artefact is an unintended feature visible in a CT image, originating from the scanning or reconstruction process rather than the physical object itself [38].

One common artefact that occurs in computed tomography, and is also observed in this work, is a ring artefact, which presents itself as a concentric or discontinuous circular band centred at the reconstruction isocentre in axial images [39]. Ring artefacts can be caused by detector calibration errors, mechanical misalignment or beam hardening related irregularities [39]. Two examples of a ring artefact are illustrated in Figure 3.7. Considering their primary causes, typical XCT artefacts can be generally classified into ring artefacts, metal-induced artefacts, scatter-related artefacts, and motion artefacts [39]. However, these will not be discussed any further, as ring artefacts are the most prominent in this work.

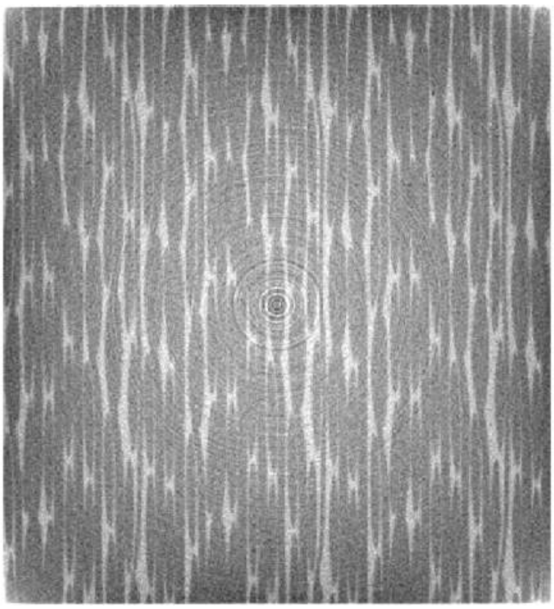

(a) Example of a ring artefact in virtually XCT scanned 2D woven composite.

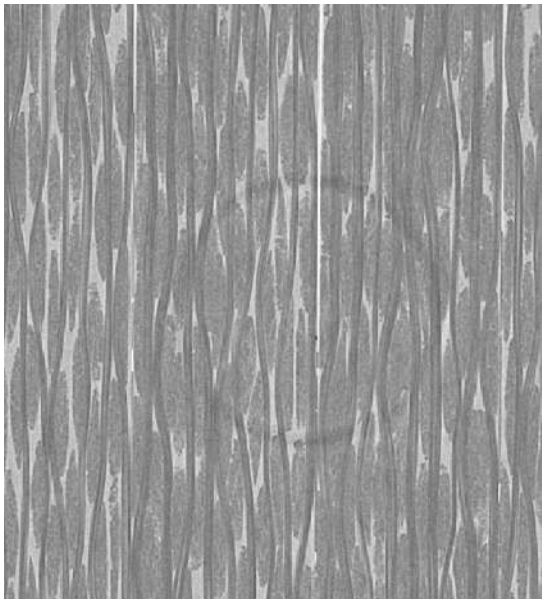

(b) Example of a ring artefact in physically XCT scanned 2D woven composite.

Figure 3.7: Examples of ring artefacts in virtually and physically XCT scanned samples.

3.3 XCT Scanning and Reconstruction of Virtual Samples

Following the theoretical background on XCT imaging and scan geometry, the established principles were applied to the virtual composite sample generated in TexGen presented in Section 2.4. The virtual XCT workflow, presented in Figure 1.2, aims to mirror that of a real experimental XCT setup. The workflow follows three main stages: generation of the virtual textile geometry in TexGen, virtual XCT scanning and reconstruction of the sample using gVXR, and image segmentation and subsequent mechanical property predictions using the TexTomos pipeline. Utilising a virtual model allows for direct comparison between segmented data and automatically labelled data, which is a key advantage of using a virtual model, as described in more detail in [12, 13].

The virtual XCT simulations were carried out using the TexTomoS [11] pipeline, which integrates the gVXR framework for X-ray projection simulation together with reconstruction tools. An illustration of the virtual scan geometry in gVXR is shown in Figure 3.8. Due to computational resource limitations, the virtual scans were restricted to a reduced model consisting of a single unit cell (UC) with three layers, as illustrated in Figure 2.7. Larger models were not feasible due to the memory constraints associated with both the reconstruction and surface mesh segmentation steps of the pipeline. Prior to performing XCT simulations, the scan geometry and acquisition parameters were defined, including the voxel size and FOV. The selected parameters for the virtual XCT simulations are summarised in Table 3.1.

To further reduce computational cost, a binning factor of 2 was utilised. This reduced the number of pixels by grouping them together into a larger pixel, decreasing memory requirements. However, as previously stated, this reduces the spatial resolution as well as the FOV, and the SOD

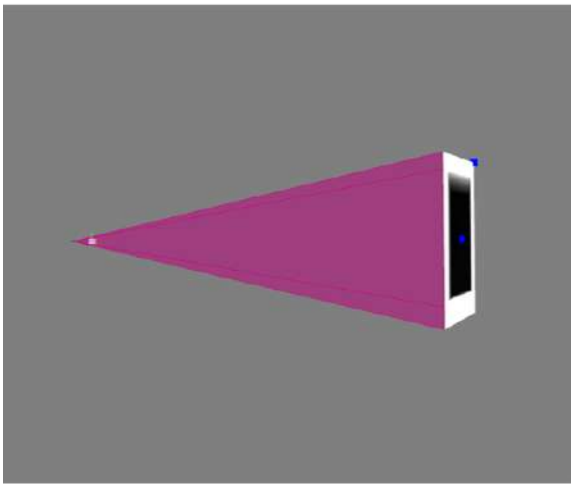

Figure 3.8: Scan Geometry in gVXR.

Table 3.1: Scan geometry and parameters implemented for the virtual XCT scan.

Parameter	Value	Unit
SOD	46	mm
ODD	754	mm
Objective	Micro-Focus	-
Binning	2	-
Source Tube Voltage	40	kV
Source Tube Power	10	W
Exposure Time	5	s
Scanning Angle	360	deg
Number of Projections	3001	-
Filter	Al	-
Filter Thickness	1	mm

therefore needs to be updated accordingly to preserve the correct voxel size.

The chosen voxel size was determined based on both the composite architecture investigated in this work and the requirements of the pre-trained neural network used in TexTomoS [12, 13]. As the physical plates examined in this study contained approximately half the number of filaments per tow compared to those in [12, 13], and the pipeline was trained on larger tow dimensions, a finer spatial resolution was required. Consequently, a voxel size of 23 μm was chosen to ensure both adequate yarn resolution and a comparable number of pixels per tow for this particular architecture studied.

Finally, the simulated projection data was reconstructed into a 3D volume within the TexTomoS pipeline utilising the ASTRA toolbox, which is an open platform for 3D image reconstruction in tomography [40, 11]. The reconstruction of the virtually scanned sample is presented in Figure 3.9. The main yarn regions are visible, and the overall warp and weft arrangement can be identified. However, the yarn boundaries appear diffuse rather than sharply defined, indicating that the selected effective voxel size is close to the lower resolution limit required for yarn-scale segmentation. Virtual ring artefacts are also present in the reconstruction, as different artefacts were included by the author in [12, 13] to better reflect artefacts commonly observed in experimental XCT scanning. As an initial evaluation step, the reconstruction was passed into the pre-trained neural network to assess whether the image quality and yarn visibility were sufficient for automated segmentation prior to any additional training.

At this stage, the concept of a ‘surface mesh segmentation’ should be clarified. In contrast to voxel-based segmentation, surface mesh segmentation does not involve any virtual scanning or reconstruction. Instead, it is based directly on the geometry generated from TexGen, where each

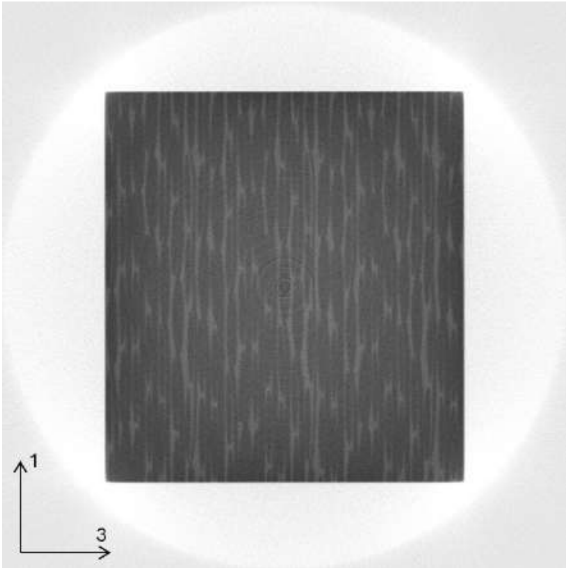

Figure 3.9: Slice of reconstruction in the 13-plane utilising an effective voxel size of 23 μm.

material constituent is automatically assigned the correct class. This provides an exact, geometry-based representation of the material architecture, which can be used as a reference for subsequent analysis. An example of this is shown in Figure 3.10, where light grey yarns correspond to the warp, dark grey yarns to the weft, and black regions to the matrix material. This automatically labelled surface mesh segmentation can then serve two purposes. Firstly, it acts as a ground truth to validate the segmentation performance of the neural network when segmenting the virtual model. It also acts as a way to automatically label training data, if transfer learning of the neural network is required.

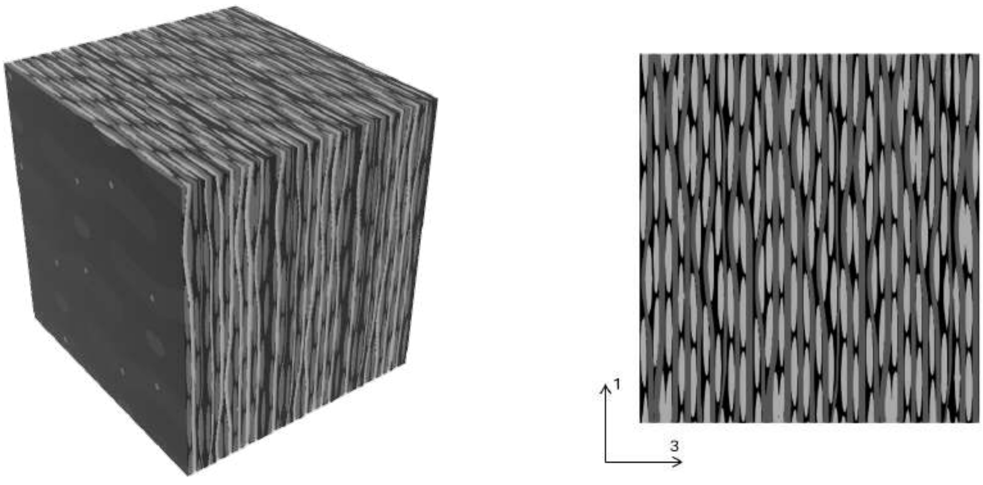

(a) Surface mesh segmentation of the full virtual volume.

(b) Surface mesh segmentation slice in the 13-plane.

Figure 3.10: Automatically labelled surface mesh segmentation obtained from the virtual XCT pipeline. Light grey yarns correspond to the warp, dark grey yarns to the weft, and black regions to the matrix material.

3.4 XCT Scanning and Reconstruction of Physical Samples

Following the virtual XCT simulations, the same general workflow presented in Figure 1.2 was applied to the physical composite samples discussed in Section 2.3. The resin-injected and dry fibre samples were scanned using a *Waygate Technologies Phoenix V|tome|x M240* XCT system [41]. The physical XCT scan setup is presented in Figures 3.11a and 3.11b for the resin-injected and dry sample, respectively.

(a) Resin-injected sample setup.

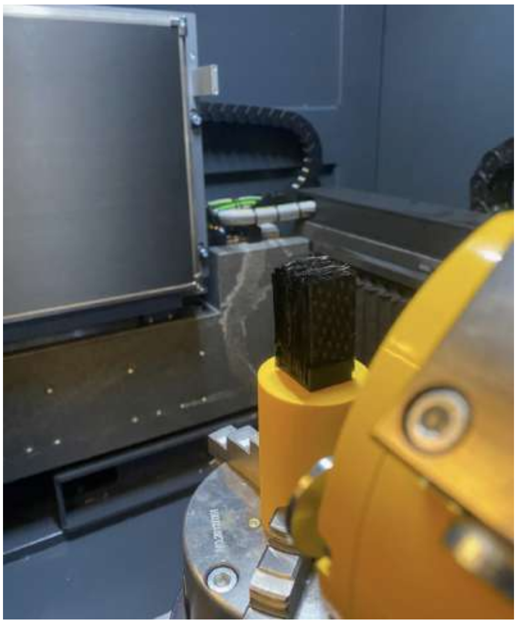

(b) Dry sample setup.

Figure 3.11: Physical samples positioned inside the XCT scanner.

The physical scans were performed using the same target voxel size as the virtual XCT simulations, corresponding to 23 μm. Since the physical XCT system differed from the virtual setup, the scan geometry and magnification parameters were adjusted using the equations presented in Section 3.1. The SOD and resulting FOV were chosen such that the relevant sample region could be captured while maintaining the same voxel size as in the virtual XCT scans. The selected parameters for the physical XCT scans are summarised in Table 3.2. Due to the maximum exposure time of 2 seconds available on the physical XCT system, the acquisition parameters could not be matched exactly to those used in the virtual scan. As a result, the physical scan parameters were adjusted according to CT experts at GKN Aerospace, to maximise the contrast while maintaining the voxel size.

Table 3.2: Scan geometry and parameters implemented for the physical XCT scans.

Parameter	Value	Unit
SOD	92.35	mm
ODD	707.65	mm
Objective	Nano-focus	-
Binning	1	-
Source Tube Voltage	40	kV
Source Tube Power	23	W
Exposure Time	2	s
Scanning Angle	360	deg
Number of Projections	3001	-

The projection data obtained from the physical scans was reconstructed using VGStudio MAX.

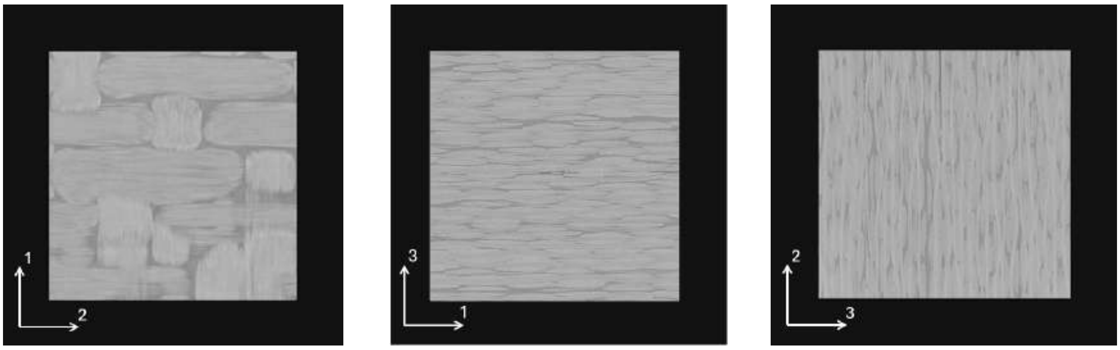

(a) 12-plane. **(b)** 13-plane. **(c)** 23-plane.

Figure 3.12: Representative reconstructed slices of a selected ROI within the resin-injected sample, shown in three orthogonal planes.

This differs from the virtual XCT reconstruction, which was performed using the ASTRA toolbox within the TexTomoS pipeline. Representative slices of the reconstructed volumes for both the dry and resin-injected samples are shown in Figure 3.13. The reconstructed slices show clear differences between the dry fibre and resin-injected samples. The dry fibre sample exhibits higher contrast between the yarns and the surrounding air, which is expected from the XCT imaging theory discussed in Section 3.1.2. This improved contrast makes the yarn architecture easier to identify, and the warp and weft yarns can be followed through the reconstructed volume at the selected voxel size of 23 μm.

For the resin-injected sample, the contrast between the yarns and surrounding matrix is lower. This is also expected, since the attenuation coefficient difference between the fibres and the resin is smaller than between the fibres and air in the dry scan. For this reason, the yarn boundaries are less sharply defined in the resin-injected reconstruction. However, the warp and weft yarn regions can still, to a large extent, be visually distinguished in the reconstructed slices.

Although a larger FOV was obtained due to using binning 1, and in extension a larger sample size, the final comparison was still limited to an ROI corresponding to one UC and three layers, in the same manner as the virtual model. As the reconstructed physical volumes were significantly larger than the corresponding virtual models, the ROI was extracted as a final preprocessing step before segmentation. This ensured consistency between the physical and virtual datasets processed in the subsequent stages of the pipeline. Slices of the extracted ROI in varying planes are shown in Figure 3.12. Similar to the virtual reconstruction, the reconstructed ROIs of the physical volumes were used as input to the pre-trained neural-network segmentation stage of the TexTomoS pipeline. The resulting evaluations are presented in the following chapter.

(a) 12-plane (resin).

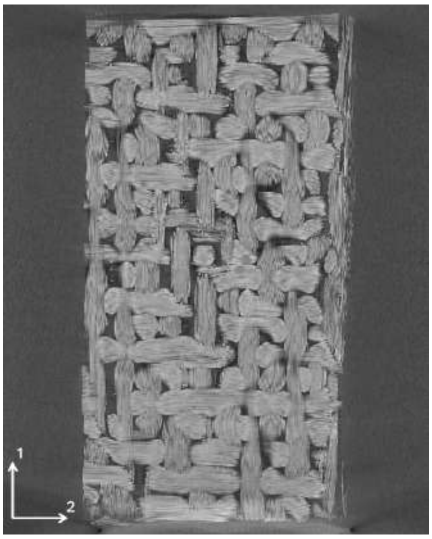

(b) 12-plane (dry).

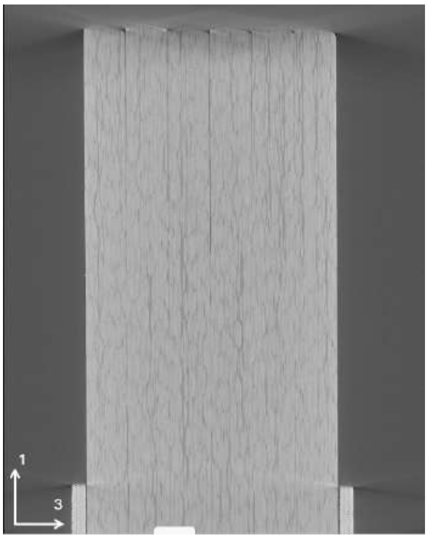

(c) 13-plane (resin).

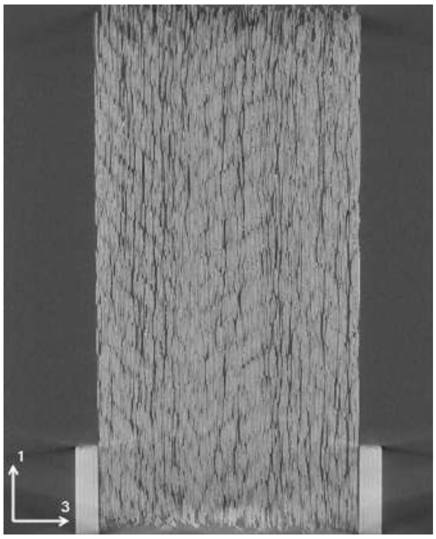

(d) 13-plane (dry).

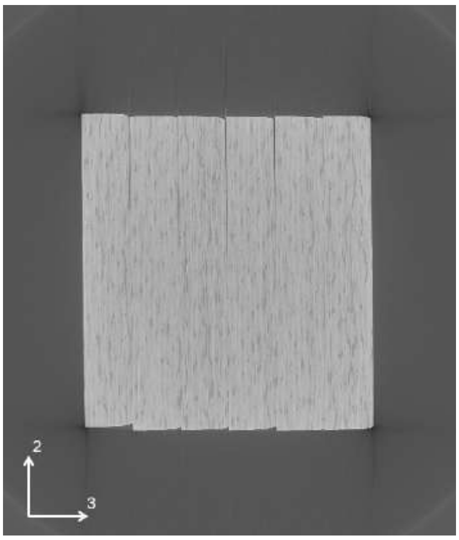

(e) 23-plane (resin).

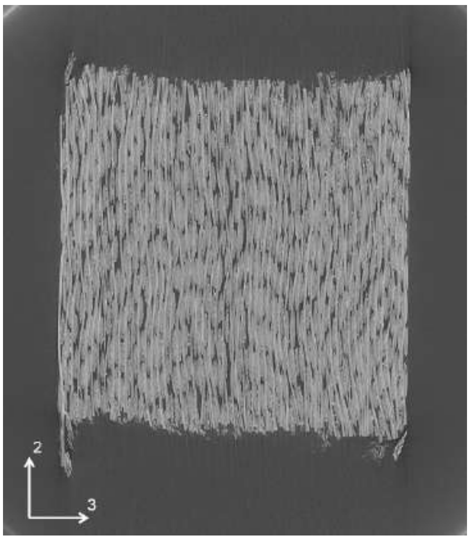

(f) 23-plane (dry).

Figure 3.13: Reconstructed slices of the resin-injected (left) and dry fibre (right) samples in three orthogonal planes.

4

Machine Learning-Based XCT Segmentation

As previously outlined, X-ray computed tomography reconstructions of textile composites consist of a series of volumetric greyscale images where each voxel represents the local X-ray attenuation within the material [42]. In order to utilise XCT data for voxel-based finite element modelling, the reconstructed volume must first be segmented into its corresponding material phases. In the case of textile-reinforced composites, these material phases typically consist of weft yarns, warp yarns, matrix and air or void regions.

Image segmentation can be performed either manually or automatically. While manual segmentation, which involves assigning material labels to individual regions or voxels by hand, can provide highly accurate results, the process is time-consuming, operator dependent and impractical for large three-dimensional XCT datasets. In addition, conventional grey scale segmentation is very challenging for CFRPs given the poor contrast between the constituents, as previously mentioned. For this reason, machine learning-based segmentation methods have become increasingly attractive for XCT analysis of composite materials [12, 13].

As previously discussed, this work makes use of the TexTomoS pipeline. This pipeline is an automated XCT-to-FE pipeline designed to automatically and efficiently predict the elastic stiffness tensor of textile-reinforced composites with minimal human intervention. The general steps of the pipeline are illustrated in Figure 4.1. This section will discuss the use of machine learning to carry out automated segmentation. Furthermore, it will expand on the type and structure of neural network used, how training data was generated for the training that was carried out, as well as how the model performed.

4.1 Machine Learning-Based XCT Segmentation

In order to carry out volumetric XCT segmentation the TexTomoS pipeline utilises a convolutional neural network (CNN). This network structure is known for its ability to automatically extract spatial features from image data. In fact, encoder-decoder architectures such as the U-Net have become widely used in medical imaging and materials science segmentation tasks [43]. The specific network architecture used by the TexTomoS pipeline is known as a 3D convolution U-Net and is illustrated in Figure 4.2.

The goal of such a neural network is to automatically segment an XCT scan, i.e. assign a label to each voxel of either a warp yarn, weft yarn, matrix or air. In order to train the model to do this, the network was trained on purely synthetic data of 3D textile-reinforced composites. In more detail, pairs of input images and corresponding labelled reference data, commonly referred to as ground truth data [44] were fed through the network. During training, the network learns to associate greyscale intensity patterns and geometric features in the XCT reconstruction with specific material phases. Once trained, the network can subsequently predict segmentation labels for previously unseen XCT data.

It's at this point relevant to discuss the training data in more detail. A major challenge in supervised machine learning is the requirement for large quantities of labelled training data [12, 13]. Manually creating segmented three-dimensional XCT datasets is highly labour intensive, and therefore not feasible. To reduce the need for manual segmentation, this work utilised synthetic XCT data generated in a very similar way to that discussed in Section 3.3. That is that hundreds

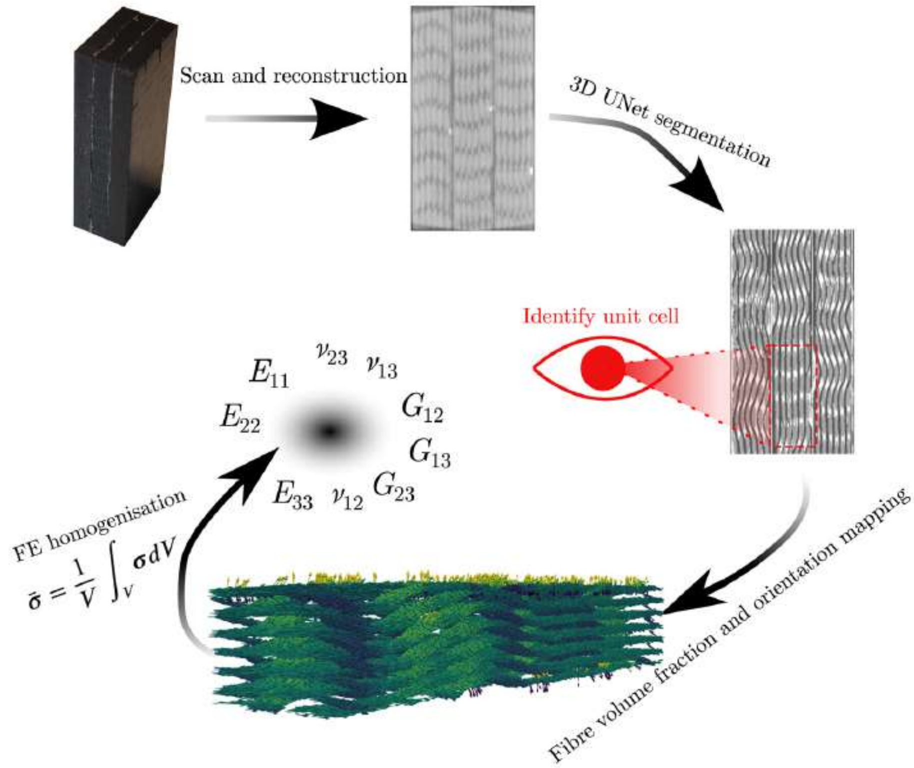

Figure 4.1: Illustration of TexTomoS pipeline. Reprinted with permission from [13].

of TexGen models of varying yarn architectures were generated and virtually scanned with gVXR. Since the material phases of virtual models from TexGen are already known, automatically labelled ground truth data can be generated directly without requiring manual labour.

4.2 Initial Performance of Pre-Trained 3D U-Net

As previously discussed, the 3D U-Net model utilised in this work had originally been trained exclusively on datasets containing 3D layer-to-layer angle interlock composite architectures. As a result, 2D 5HS weave reinforced composites represented a previously unseen architecture for the neural network. Consequently, the pre-trained model may not generalise well to this architecture.

Given the different architecture, the necessity of carrying out additional training, also known as transfer learning, depends on whether the pre-trained network can achieve sufficiently accurate segmentation results on the new 5HS data. To assess this, a quantitative evaluation of the segmentation performance was carried out using a synthetically generated virtual sample. As previously discussed, this virtual dataset provides a known ground truth enabling direct comparison with the machine learning-based segmentation. This was done using a pixelwise agreement analysis to assess the segmentation accuracy.

4.2.1 Quantitative Evaluation

Figures 4.3-4.5 present comparisons between the ground truth and the neural network output for representative slices in the three planes (*12*, *13*, and *23*). In addition to these slice-wise comparisons, Figure 4.6 shows difference overlays, highlighting regions of agreement and disagreement across all planes. For clarity, only results corresponding to a slice from the centre of the sample are presented here. Additional results, including supplementary evaluated slices across each plane (one slice in the beginning, centre, and end of the sample, respectively), are presented in Appendix A.1.1.

The quantitative results are summarised in Table 4.1, where both pixelwise agreements and Dice

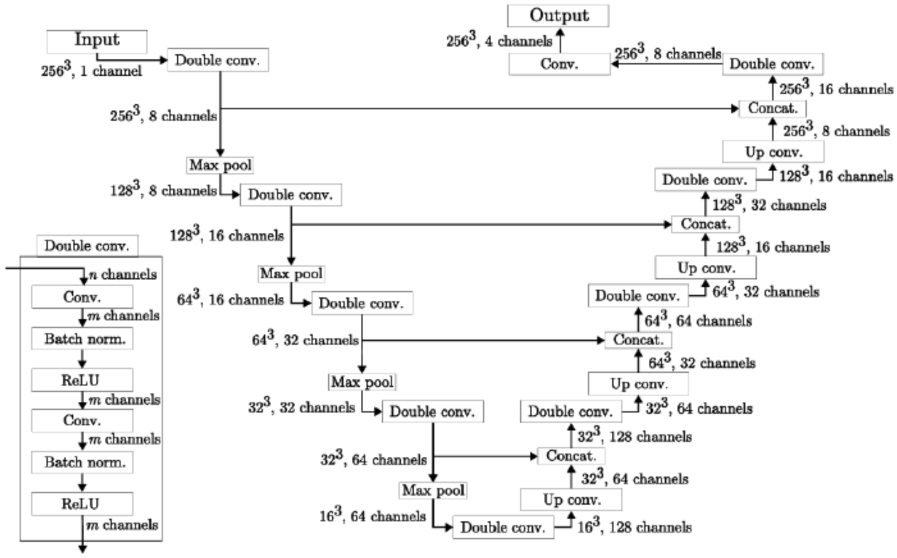

Figure 4.2: Overview of the 3D convolution U-Net architecture utilised in TexTomoS. Reprinted with permission from [13].

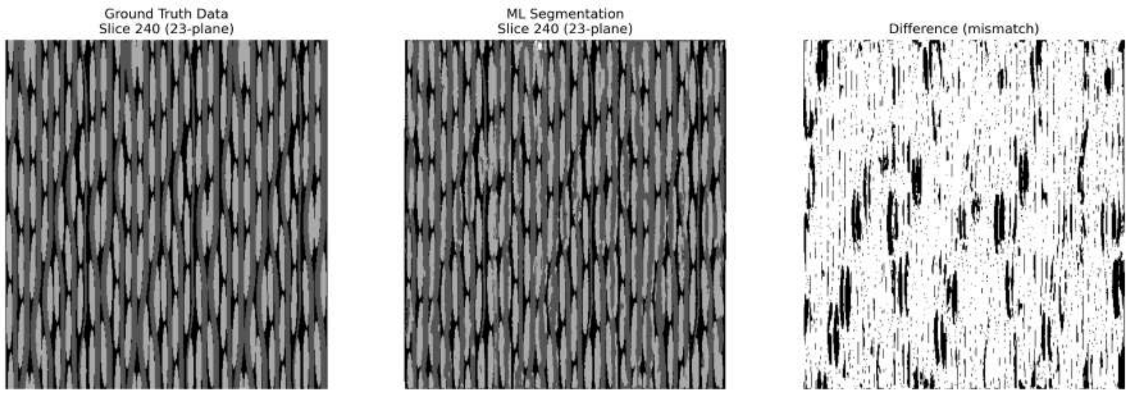

Figure 4.3: Machine learning-based segmentation performance evaluation on virtual sample, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image, white corresponds to agreement while black corresponds to differences.

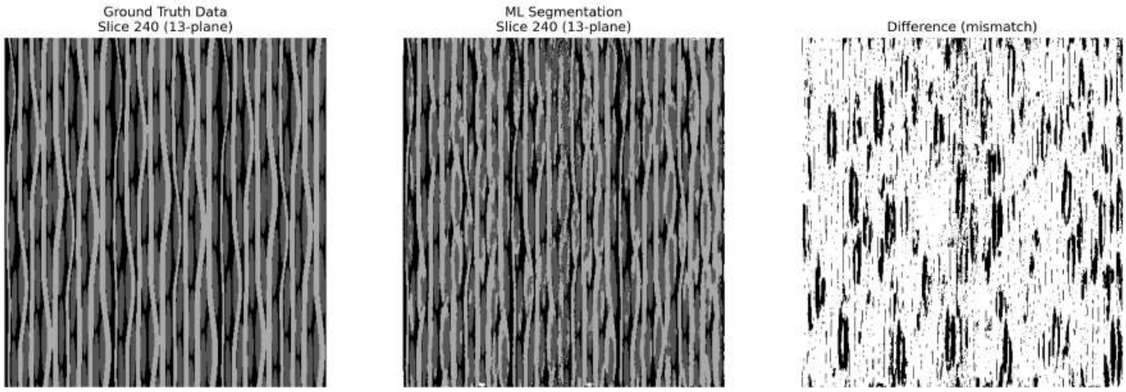

Figure 4.4: Machine learning-based segmentation performance evaluation on virtual sample, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image, white corresponds to agreement while black corresponds to differences.

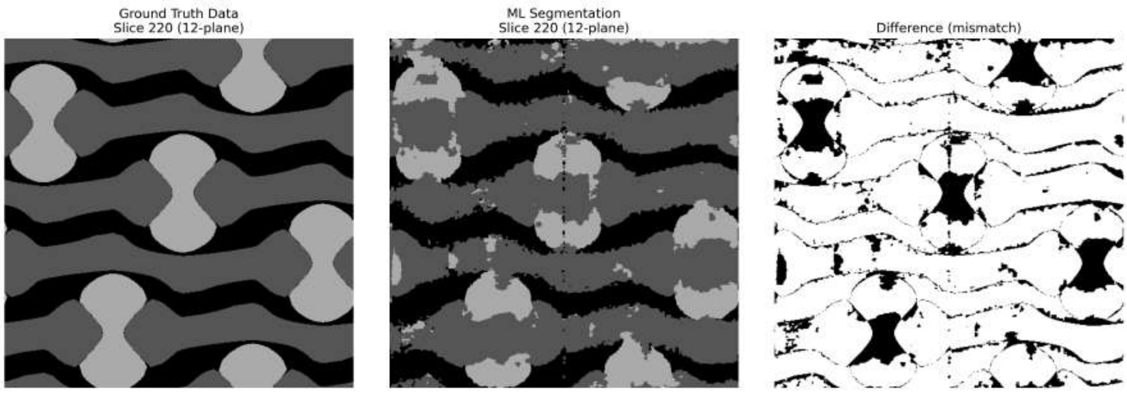

Figure 4.5: Machine learning-based segmentation performance evaluation on virtual sample, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image, white corresponds to agreement while black corresponds to differences.

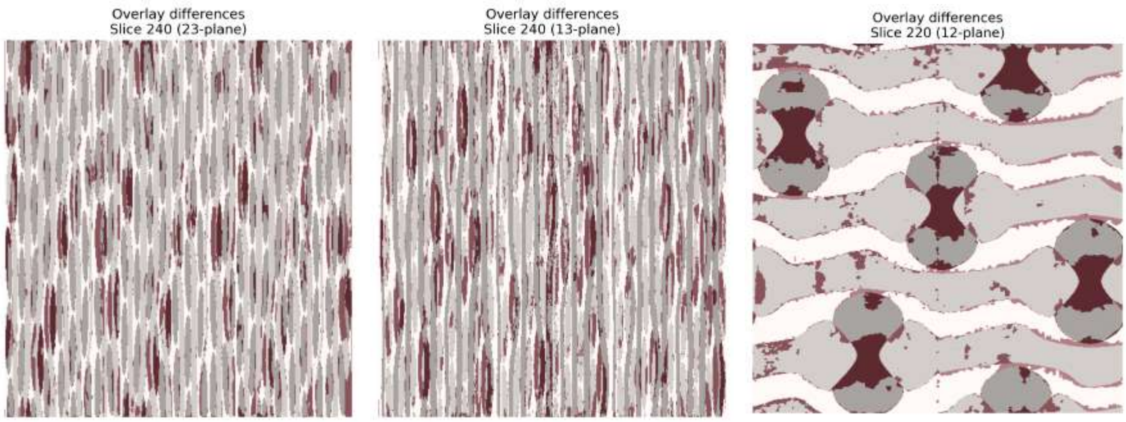

Figure 4.6: Machine learning–based segmentation performance evaluation: difference overlays shown for one slice per plane of the virtual sample. Red indicates areas of difference.

similarity coefficients are reported for multiple slices along each axis. In addition, the pixelwise agreement is reported for the entire volume to provide an overall measure of the segmentation performance. The pixelwise agreement reflects the proportion of correctly classified pixels, while the Dice similarity coefficient is a widely used metric for image segmentation that quantifies the spatial overlap between a predicted segmentation and the ground truth, with values ranging from 0 (no overlap) to 1 (perfect agreement) [45, 46]. Dice scores are reported separately for each class, where C1, C2, and C3 correspond to class 1 (warp), class 2 (weft), and class 3 (matrix), respectively. This enables class-specific evaluation of the segmentation performance.

Table 4.1: Segmentation comparison across selected slices and axes of virtual sample prior transfer learning. Pixelwise agreement and Dice scores are reported per slice.

Plane	Slice	Agreement (%)	Dice C1	Dice C2	Dice C3
		23-Plane			
23	20	82.78	0.8071	0.7518	0.9309
23	240	83.42	0.8055	0.8242	0.9032
23	460	82.22	0.7911	0.7660	0.9298
Mean		**82.81**			
		13-Plane			
13	20	83.50	0.7950	0.7977	0.9208
13	240	80.13	0.7873	0.7397	0.9154
13	460	80.09	0.7948	0.7551	0.8762
Mean		**81.24**			
		12-Plane			
12	20	85.21	0.9007	0.1023	0.9510
12	220	83.24	0.7095	0.8518	0.8859
12	420	86.71	0.5999	0.8834	0.9627
Mean		**85.05**			
		Volume-wide agreement			
Global		**83.30**			
Axis variability (STD)		1.70 / 2.21 / 6.47			

The mean agreement values range from approximately 81% to 85%, indicating a moderate overall correspondence between the neural network and the ground truth. However, a closer examination of the class-wise Dice scores reveals notable inconsistencies. While certain classes (e.g. C3) achieve high Dice scores exceeding 0.9 in several instances, other classes (e.g. C2) exhibit significantly lower and more variable performance, including cases such as a Dice score of 0.1023.

These discrepancies suggest that the network struggles to reliably distinguish between specific material phases, likely due to differences in feature representation between the training (3D textile) and target (2D 5HS textile) domains. Furthermore, the variation in performance across slices and axes indicates that the segmentation quality is not spatially uniform. This behaviour is more prominent towards the boundaries of the slices. A reason for this might be that the information available for the network during inference is limited. With the patch-based inference strategy employed by the 3D U-Net, areas close to the slice boundaries may experience misclassification due to the absence of sufficient surrounding context. Due to time limitations this was not investigated further or confirmed. The initial results indicate that the pre-trained network does not generalise sufficiently well to 5HS woven architectures, thus failing to reach the target value of 90% volume-wide pixelwise agreement (see Section 1.2.1), motivating the need for additional training of the neural network. A qualitative evaluation of the segmentation performance was also conducted on a real XCT scan of the resin-injected 5HS composite sample. The results of this evaluation are presented in Appendix A.2.

4.3 3D U-Net Transfer Learning for 5HS Woven Composites

To improve segmentation performance under new conditions, transfer learning can be employed as an effective strategy. In transfer learning, a neural network that has already been trained on a source dataset is reused as the initialisation point for a related target task, and subsequently fine-tuned using data that is more representative of the new application domain [47]. This approach enables the model to adapt previously learned feature representations while retaining knowledge from the original training process.

A key advantage of transfer learning is that it significantly reduces the amount of labelled data and computational resources required compared to training a model from scratch. In the context of textile composite XCT segmentation, transfer learning allows previously developed segmentation models to be efficiently adapted to variations in weave architecture, yarn geometry, fibre volume fraction, and XCT scan parameters such as resolution, noise, and contrast. Consequently, it provides a practical way of extending the applicability of existing models to new material systems without requiring a full retraining pipeline.

Following the outlined workflow presented in Figure 1.2, additional training data corresponding to the 5HS weave architecture was generated, and transfer learning was subsequently carried out on the pre-trained network. As previously mentioned, the main objective of this stage was to adapt the neural network's learned features from 3D layer-to-layer angle interlock architectures to the different geometric characteristics of 5HS woven architectures.

The transfer learning dataset consisted of 100 synthetically generated 5HS weave samples used for training, alongside 10 samples each for validation and testing. To enhance the generalisation capability, geometric variability was introduced during synthetic sample generation by offsetting the weave layers and applying deformation and rotations to the yarns. This was carried out to ensure that the model was exposed to a representative range of architectures and to avoid overfitting to a single idealised configuration.

The training was initialised using the weights of the pre-trained 3D U-Net, and the transfer learning was conducted using a low learning rate of 10^{-3}, which was linearly decreased to zero over 10 epochs. A low learning rate was employed to prevent forgetting, where newly learned information can overwrite previously learned features. Instead, this strategy allows the model to adjust higher-level feature representations while retaining the lower-level features learned from the original dataset.

4.3.1 Quantitative Evaluation

The performance of the transfer learning model was assessed using the same virtual 5HS dataset and evaluation methodology as previously described in Section 4.2. Figures 4.7-4.10 present the representative slice-wise comparisons between the predicted segmentation and the ground truth, along with difference overlays to highlight agreement and discrepancies. Once again, the complete set of results, including additional evaluated slices from each plane, is presented in Appendix A.2.1.

A clear improvement in segmentation accuracy can be observed across all planes following the transfer learning. This is confirmed in Table 4.2, where the global pixelwise agreement increases from approximately 83% prior to transfer learning to approximately 94% post transfer learning. Similarly, the mean agreement values for each plane exceed 93%, indicating a substantial enhancement in the machine learning-based segmentation performance.

The class-wise Dice similarity coefficients further illustrate the improvement in segmentation performance. In comparison to the pre-trained model, where significant variability between classes was observed, the post-transfer model achieves consistently high Dice scores across all classes, on average exceeding 0.94. In particular, class C2 now shows a significant improvement, with Dice scores stabilising around 0.94 across most slices. This indicates that the model has successfully learned to distinguish between constituents that were previously difficult to separate due to no prior training on 5HS textile architectures.

It is however important to note the presence of the outlier in the 12-plane (slice 420), where the obtained Dice score for classes C1 and C2 are as low as 0.6634 and 0.8922, respectively, resulting in a lower pixelwise agreement value of 88.01%. This local discrepancy contributes to the relatively high standard deviation observed along axis 3 of $\pm 8.68\%$. The results for the relevant slice are

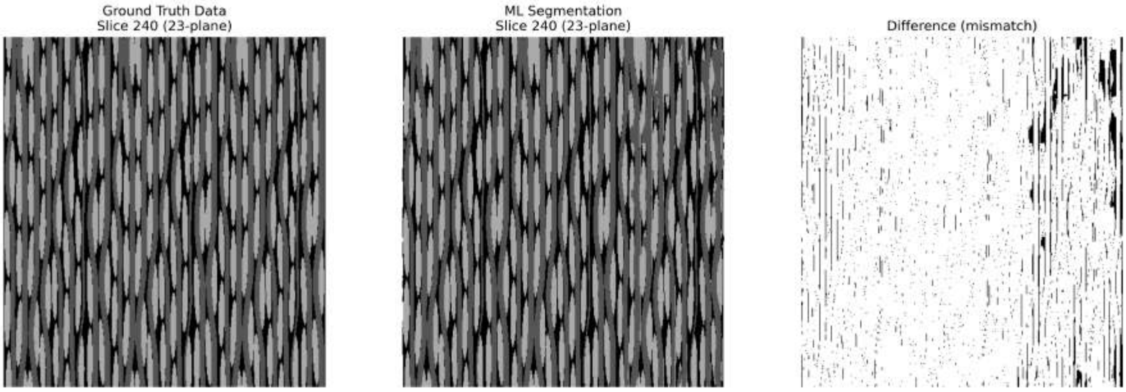

Figure 4.7: Machine learning-based segmentation performance evaluation on virtual sample post transfer learning, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

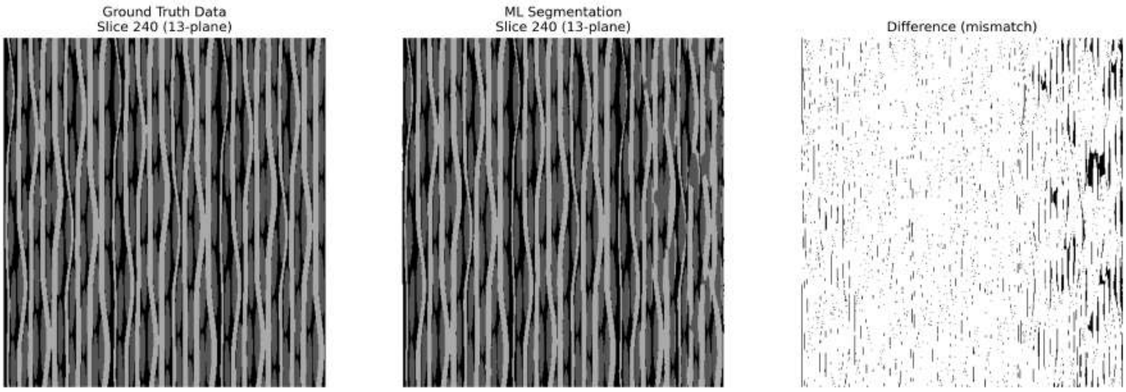

Figure 4.8: Machine learning-based segmentation performance evaluation on virtual sample post transfer learning, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

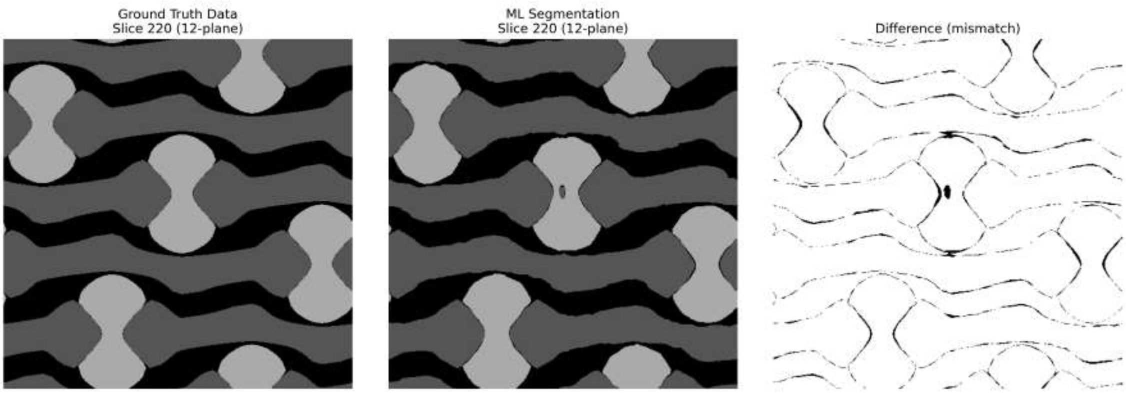

Figure 4.9: Machine learning-based segmentation performance evaluation on virtual sample post transfer learning, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

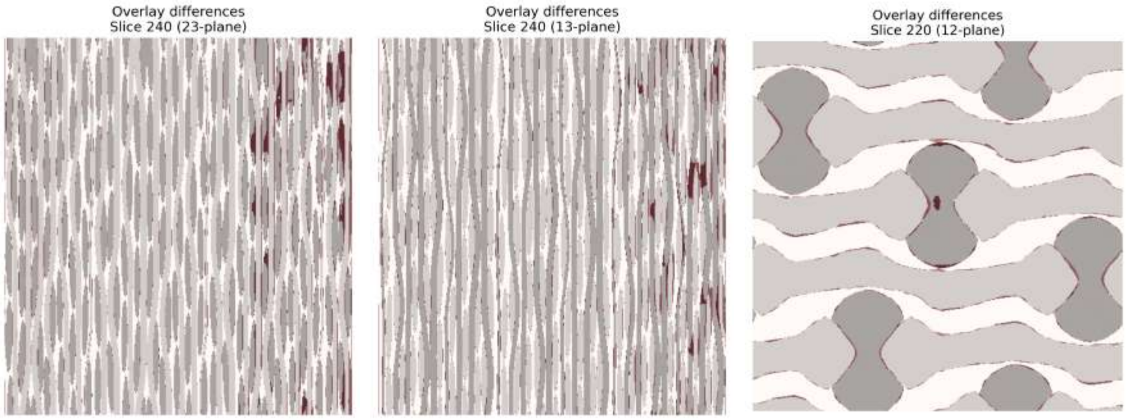

Figure 4.10: Machine learning–based segmentation performance evaluation: difference overlays shown for one slice per plane of the virtual sample. Red indicates areas of difference.

Table 4.2: Segmentation comparison across slices and axes of virtual sample post transfer learning. Pixelwise agreement and Dice scores are reported per slice.

Plane	Slice	Agreement (%)	Dice C1	Dice C2	Dice C3
		13-Plane			
23	20	94.78	0.9568	0.9497	0.9424
23	240	93.09	0.9350	0.9382	0.9182
23	460	93.35	0.9354	0.9315	0.9411
Mean		**93.74**			
		23-Plane			
13	20	93.35	0.9285	0.9344	0.9431
13	240	93.89	0.9476	0.9416	0.9276
13	460	93.87	0.9472	0.9395	0.9332
Mean		**93.70**			
		12-Plane			
12	20	97.13	0.9840	0.9711	0.9313
12	220	97.15	0.9778	0.9780	0.9549
12	420	88.01	0.6634	0.8922	0.9528
Mean		**94.10**			
		Volume-wide agreement			
Global		**94.00**			
Axis variability (STD)		0.0047 / 0.0060 / 0.0868			

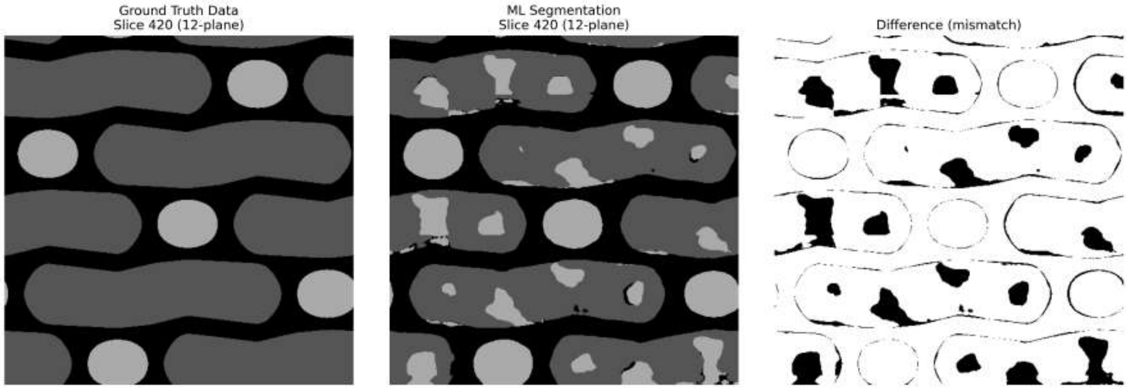

Figure 4.11: Machine learning-based segmentation performance evaluation post transfer learning, slice 420, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

presented in Figure 4.11.

As observed, the neural network exhibits misclassification between the weft yarns (dark grey) and warp yarns (light grey), in slices located near the overlap regions of the two yarn types. This behaviour is more prominent towards the boundaries of the volume, which can likely be attributed to the limited availability of surrounding contextual information for the neural network during the inference stage. An alternative explanation may involve the presence of complex geometric features or local variations in slices near the overlap regions, which remains difficult for the neural network to accurately learn given the limited dataset utilised for the transfer learning. However, since the misclassification is almost exclusively observed near the volume boundaries, it strongly indicates that insufficient contextual information is the primary contributing factor.

Despite this discrepancy, the overall increase in the performance of the machine learning-based segmentation model compared to the pre-trained model is significant. In order to enhance the performance of the segmentation model further, transfer learning with an extended dataset would have to be carried out. This would allow the neural network to learn more of the geometric features of the 5HS weave architecture, hence resulting in a better machine learning-based segmentation performance. However, due to the time limitations related to the present work, transfer learning with a larger dataset was not conducted.

4.3.2 Qualitative Evaluation on Physical XCT Data of Resin-Injected Sample

The quantitative evaluation is supported by a qualitative evaluation of the same physical resin-injected 5HS composite sample previously evaluated prior to transfer learning, as shown in Figure 4.12. Compared to the pre-trained model results presented in Appendix A.2, the transfer learning model produces cleaner and more coherent segmentation in the 13- and 23-planes. The yarn boundaries are more clearly defined, and the continuity of yarn paths is better preserved throughout the volume. These observations are consistent with the quantitative improvements presented in Table 4.2, where the transfer learning model achieved higher Dice similarity coefficients and pixelwise agreement values across all evaluated slices.

This suggests that the neural network has successfully adapted to the geometric characteristics of the 5HS weave architecture. In particular, the results suggest that the synthetically generated transfer learning dataset was sufficiently representative of the physical 5HS composite sample, at least for the 13- and 23-planes. This enabled the neural network model to generalise beyond the original layer-to-layer angle interlock architectures utilised during the initial training stage. This demonstrates the potential of transfer learning as an efficient strategy for adapting machine learning-based segmentation models to new textile architectures, eliminating the need for complete retraining of the neural network from scratch.

Although the segmentation performance in the 12-plane also appears to show some improvement, see Figure 4.12d, maintaining yarn continuity throughout the volume seems to remain a challenge for the neural network. This behaviour appears to be more pronounced in slices near the overlap regions between the different yarn types. A likely explanation is that these regions exhibit more complex geometric interactions, making them inherently more difficult for the model to segment consistently. This suggests that further improvements to the segmentation performance may be achieved through transfer learning with a more extensive dataset, or through increased dataset diversity.

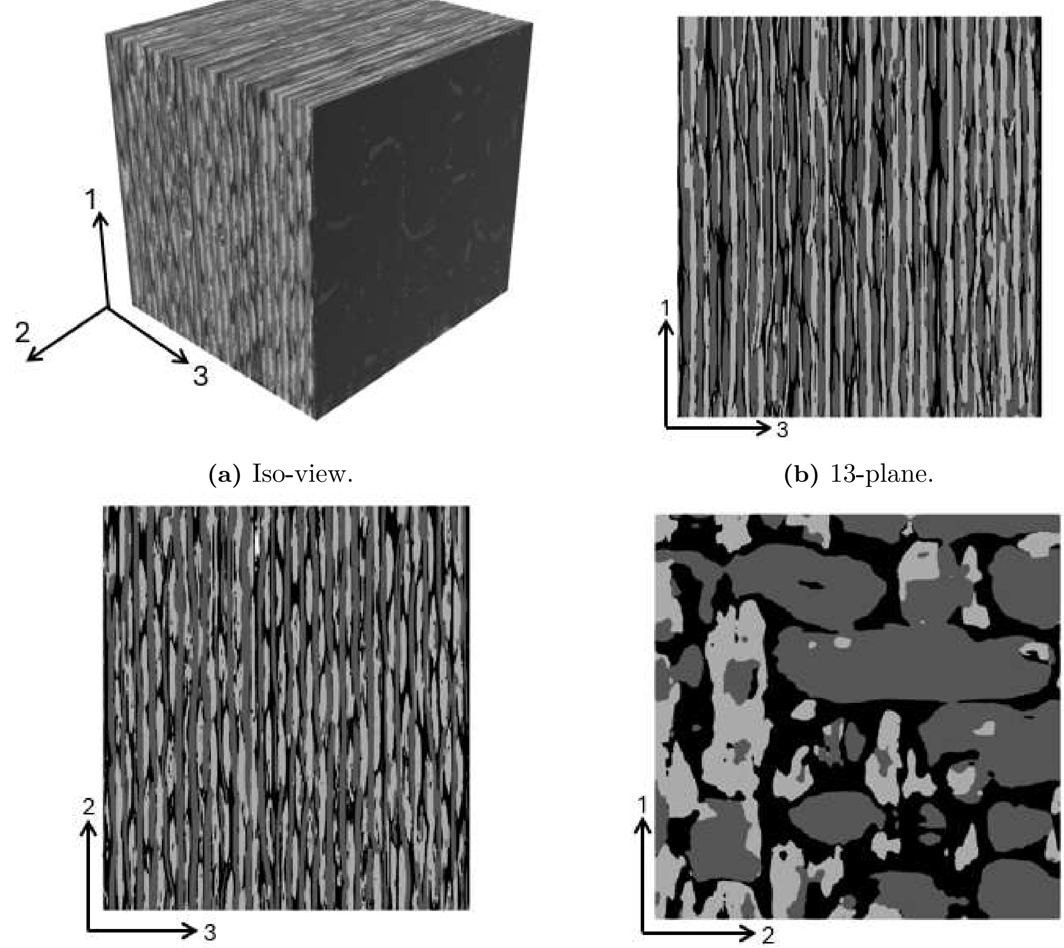

(a) Iso-view.

(b) 13-plane.

(c) 23-plane.

(d) 12-plane.

Figure 4.12: Machine learning segmented physical resin-injected sample post transfer learning. Black, dark grey, and light grey correspond to matrix, weft-yarn, and warp-yarn, respectively.

Maintaining yarn continuity in the segmented model is important because the segmentation is later used to generate FE models for computational homogenisation. Discontinuities or artificial breaks in the yarns can lead to an inaccurate representation of the reinforcement architecture, which may affect the predicted mechanical behaviour. As described later in Section 5.3, discontinuities in the segmented yarn paths may lead to an underestimation of the effective elastic properties. Therefore, achieving accurate yarn continuity throughout the segmented model is essential for accurately predicting the elastic properties of as-manufactured material samples.

Further improvements in segmentation performance would likely require transfer learning using a larger and more geometrically diverse synthetically generated dataset. This could enable the neural network to better adapt to the complex features present in the XCT data of the real sample. However, due to the time limitations associated with the present work, extended transfer learning with a larger dataset was not considered.

5

FE Modelling and Homogenisation

5.1 Multi-Scale Composite Material Characteristics

Woven composites exhibit a multiscale architecture that needs to be taken into consideration in mechanical modelling. Three scales are commonly distinguished: the microscale, mesoscale, and macroscale. In this context, the microscale corresponds to the individual carbon filaments and matrix at the μm scale. At this scale, the fibres themselves are highly anisotropic. However, depending on the modelling choices, fibres may be represented as anisotropic or approximated as isotropic or transversely isotropic in order to simplify the constitutive models. In this work, the material behaviour follows the constitutive assumptions implemented in the utilised TexTomoS modelling framework, namely that the fibres are represented as transversely isotropic.

The mesoscale represents the scale of the woven architecture where each yarn, consisting of thousands of carbon filaments embedded in the matrix, is explicitly described geometrically. At this scale, yarns are commonly modelled as transversely isotropic, characterised by one preferred fibre direction and one plane of isotropy perpendicular to the fibres [48]. The surrounding matrix is typically modelled as isotropic.

Finally, the macroscale describes the composite as a homogenised material, where the detailed yarn architecture is not explicitly represented. Instead, the woven composite is treated as an equivalent anisotropic material. For woven composites, this homogenised behaviour is commonly approximated as orthotropic. In Figure 5.1, the three considered material characteristics, namely isotropy, transverse isotropy and orthotropy, are illustrated.

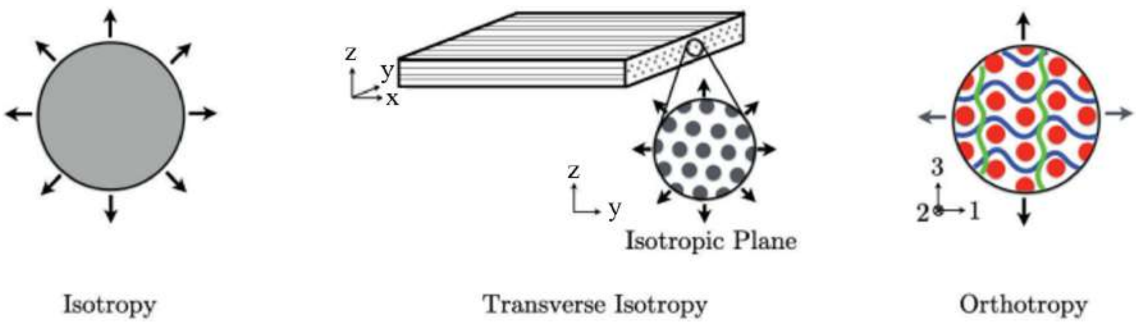

Figure 5.1: Illustration of isotropy, transverse isotropy and orthotropy. The orthotropic illustration represents a general 3D orthotropic material behaviour, while the weave considered in this work is 2D. Reprinted with permission from [48].

At the mesoscale, where yarns are modelled as transversely isotropic, the material response is described using the Young's moduli in the fibre and transverse directions, E_x, $E_y = E_z$, together with the shear moduli $G_{xy} = G_{xz}$. In addition, Poisson's ratios ν_{yz} and $\nu_{xy} = \nu_{xz}$ are required to fully define the elastic response of the yarns. The surrounding matrix is assumed to be isotropic, requiring only the Young's modulus E and Poisson's ratio ν.

The considered 5HS weave can, on the macroscale, be described as an orthotropic material. Orthotropic materials have three orthogonal material directions, each associated with unique elastic properties. Consequently, orthotropic materials require nine independent elastic constants to fully characterise their elastic response. These consist of the Young's moduli E_1, E_2, and E_3, the shear moduli G_{12}, G_{23}, and G_{31}, and the Poisson's ratios ν_{12}, ν_{23}, and ν_{31}. In this work, the material directions are defined such that direction 1 corresponds to the nominal warp direction, direction 2 corresponds to the nominal weft direction, and direction 3 corresponds to the out-of-plane direction of the weave, as illustrated in Figure 5.2. The orthotropic stiffness matrix, which is symmetric, can be expressed in Voigt form as

$$\begin{bmatrix}\sigma_{11}\\ \sigma_{22}\\ \sigma_{33}\\ \sigma_{23}\\ \sigma_{31}\\ \sigma_{12}\end{bmatrix} = \begin{bmatrix} \frac{E_1(1-\nu_{32}\nu_{23})}{\Delta} & \frac{E_1(\nu_{21}+\nu_{31}\nu_{23})}{\Delta} & \frac{E_1(\nu_{31}+\nu_{21}\nu_{32})}{\Delta} & 0 & 0 & 0\\ \frac{E_2(\nu_{12}+\nu_{13}\nu_{32})}{\Delta} & \frac{E_2(1-\nu_{31}\nu_{13})}{\Delta} & \frac{E_2(\nu_{32}+\nu_{12}\nu_{31})}{\Delta} & 0 & 0 & 0\\ \frac{E_3(\nu_{13}+\nu_{12}\nu_{23})}{\Delta} & \frac{E_3(\nu_{23}+\nu_{13}\nu_{21})}{\Delta} & \frac{E_3(1-\nu_{12}\nu_{21})}{\Delta} & 0 & 0 & 0\\ 0 & 0 & 0 & G_{23} & 0 & 0\\ 0 & 0 & 0 & 0 & G_{31} & 0\\ 0 & 0 & 0 & 0 & 0 & G_{12}\end{bmatrix} \begin{bmatrix}\varepsilon_{11}\\ \varepsilon_{22}\\ \varepsilon_{33}\\ \gamma_{23}\\ \gamma_{31}\\ \gamma_{12}\end{bmatrix}, \quad (5.1)$$

where Δ is defined as

$$\Delta = 1 - \nu_{12}\nu_{21} - \nu_{23}\nu_{32} - \nu_{31}\nu_{13} - 2\nu_{12}\nu_{23}\nu_{31}. \quad (5.2)$$

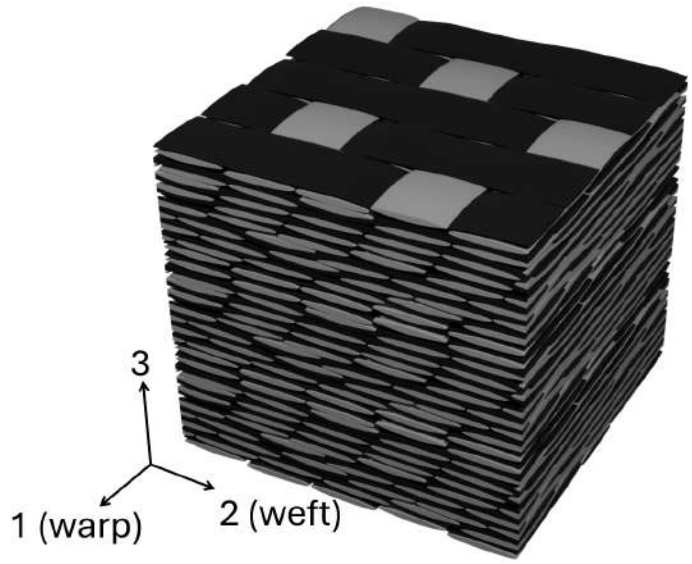

Figure 5.2: Virtual 5HS sample with coordinate system indicating the weft, warp, and out-of-plane directions.

5.1.1 Computational RVE homogenisation

To predict the effective macroscopic properties of heterogeneous composite materials, such as those with 5HS woven reinforcement, RVEs can be analysed using computational homogenisation. As the name suggests, RVEs must be chosen such that they are representative, i.e., they capture all of the essential features of the underlying mesostructure [5]. By then subjecting the RVE to a set of prescribed boundary conditions and averaging the resulting stress and strain fields, equivalent homogenised material properties can be obtained [49].

The underlying assumption behind RVE homogenisation is the existence of a separation of scales between the characteristic dimensions of the subscales and the structural scale of interest. In this setting, the heterogeneous material can be replaced by an equivalent homogeneous continuum whose constitutive response is obtained from the averaged behaviour of the RVE [49]. This enables detailed mesostructural effects to be incorporated into macroscale simulations without explicitly resolving the full textile architecture throughout an entire component.

In practice, the RVE is discretised using finite elements and subjected to a series of independent load cases. The resulting volume-averaged stresses are then related to the applied macroscopic strains through the homogenised compliance matrix. For periodic textile architectures, periodic boundary conditions (PBCs) are commonly employed to ensure periodic displacement compatibility between opposing RVE boundaries [50].

Computational homogenisation is particularly attractive for textile composites, where the complex yarn architecture strongly influences the global mechanical response. By deriving RVEs directly from measured or reconstructed mesostructures, the homogenised properties can capture important as-manufactured geometric effects that are difficult to represent using idealised textile or analytical models.

5.2 LS-DYNA Simulations

In order to predict the elastic mechanical properties of the scanned samples, simulations were carried out in the commercial FE code LS-DYNA using PBCs to extract the homogenised stiffness matrix of the RVE. Six independent load cases were simulated, consisting of three normal strain cases ($\varepsilon_{11}, \varepsilon_{22}, \varepsilon_{33}$) and three pure shear cases ($\varepsilon_{12}, \varepsilon_{23}, \varepsilon_{13}$). For each load case, the volume-averaged stress response divided by the applied macroscopic strain yielded one column of the homogenised stiffness matrix in Voigt notation.

A strain magnitude of 0.1% was used in the strain-controlled load cases. While Friemann et al. [13] used 0.01% to minimise the discrepancy between LS-DYNA's finite-deformation stress output and the infinitesimal-strain formulation, verification with 0.1% strain produced practically identical homogenised stiffnesses. Thus, the chosen strain magnitude was considered sufficiently small for linear elastic homogenisation.

The LS-DYNA keyword input files were generated using a Python-based workflow adapted from the code published in [11]. The Python script automatically converts segmented voxel-based tomography data into a structured finite element mesh with voxel-conforming hexahedral elements. Each voxel in the segmented image corresponds to one voxel element in the FE-model, thereby preserving the reconstructed yarn geometry and spatial distribution.

The mesh generation procedure first creates nodal coordinates and element connectivities for a structured three-dimensional grid. The nodal coordinates were generated directly from the voxel dimensions in the reconstruction images. In addition, periodic node pairs are identified on opposing faces, edges, and corners of the RVE to enable the application of periodic boundary conditions. These constraints ensure displacement compatibility between opposite boundaries and enforce periodic deformation of the RVE.

To account for the local transversely isotropic behaviour of the yarns, the fibre orientations were estimated directly from the segmented image data using a structure tensor analysis. The analysis was performed slice-wise for each yarn system independently. The structure tensor analysis was computed from a smoothed distance field, and the dominant eigenvector direction was extracted as the local yarn orientation. This produced a three-dimensional orientation vector for every voxel belonging to a yarn phase.

Subsequently, the local fibre volume fractions were estimated using a slice-wise approach. For each yarn system, the expected fibre cross-sectional area was distributed across the segmented yarn voxels while accounting for the local orientation of the yarns relative to the slicing direction. This yielded a spatially varying fibre volume fraction field throughout the RVE.

The effective orthotropic material properties assigned to each element were computed using the Chamis micromechanical model [51]. Based on the constituent fibre and matrix properties together with the locally estimated fibre volume fractions, the orthotropic elastic properties E_{11}, E_{22}, E_{33}, G_{12}, G_{23}, G_{13}, and the corresponding Poisson's ratios ν_{12}, ν_{13}, and ν_{23} were calculated for each material region. To reduce the number of unique material cards required in LS-DYNA, the continuous fibre volume fraction distribution was discretised into 20 bins, where each bin corresponded to a distinct homogenised material definition. Examples of the fibre volume fraction in virtual and physical resin-injected samples are presented in Figure 5.3a-5.3b, and the corresponding FE-model in LS-DYNA is presented in Figure 5.3c.

After simulating the six load cases, the homogenised stiffness matrix was assembled from the computed stress responses. The predicted Young's moduli, shear moduli, and Poisson's ratios were then obtained from the inverted compliance matrix.

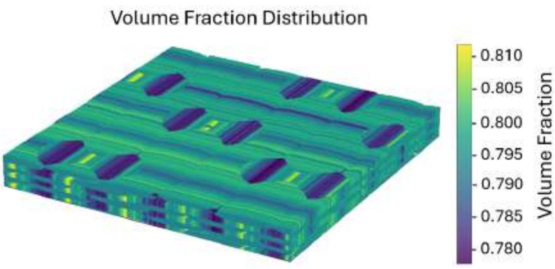

(a) Continuous fibre volume fraction distribution in virtual sample. The matrix constituent has been left out.

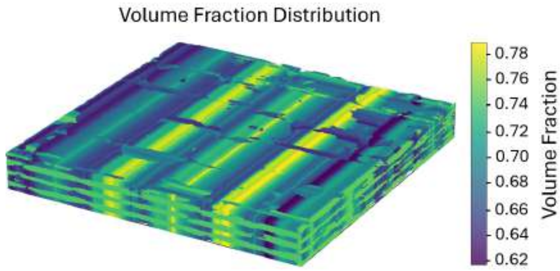

(b) Continuous fibre volume fraction distribution in resin-injected sample. The matrix constituent has been left out.

(c) FE-model of resin-injected sample. Different colours indicate fibre volume fraction bins, each with different material properties.

Figure 5.3: 3D plot of continuous fibre volume fraction in virtual and resin-injected sample, as well as corresponding FE-model in LS-DYNA.

5.3 Estimated Elastic Properties

Following the computational homogenisation strategy described in Section 5.1.1, the effective mechanical elastic properties of the 5HS composite samples were obtained from LS-DYNA simulations with an applied macroscopic strain of 0.1%. The extracted virtual and physical resin-injected RVEs utilised for the elastic property estimation are presented in Figure 5.4.

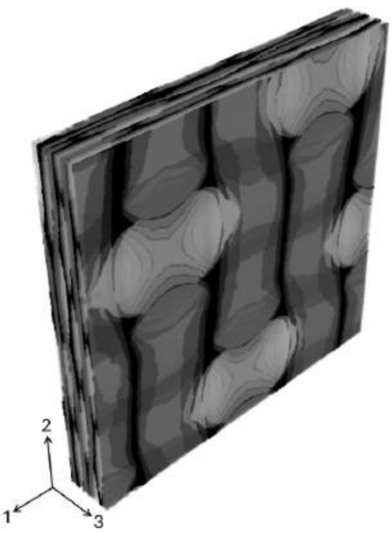

(a) Virtual sample RVE.

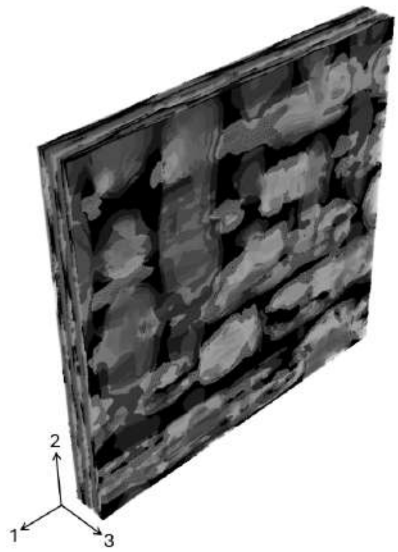

(b) Resin-injected sample RVE.

Figure 5.4: Virtual and resin-injected RVEs utilised for the estimation of elastic properties. Black, dark grey, and light grey correspond to matrix, weft-yarn, and warp-yarn, respectively.

The reason behind extracting smaller sub-volumes from the segmented models is mainly to reduce the computational complexity. This reduction is valid provided that the extracted volumes remain representative of the underlying mesostructure. In this context, the out-of-plane dimension of the volumes is not expected to significantly affect the estimated elastic properties as the volumes satisfy RVE size requirements and appropriate boundary conditions are applied.

5.3.1 Virtual Sample

The effective elastic properties obtained from the machine learning-based segmented virtual sample are summarised in Table 5.1. It should be noted that the virtual sample used for the elastic property prediction had not previously been seen by the neural network; this exact geometry was not included in the training, validation, or testing sets. The virtual model was, however, made to look as similar as possible to the resin-injected sample.

From the results in Table 5.1, it can be observed that the in-plane Young's moduli E_1 and E_2 are very similar, with values of 55.83 GPa and 56.63 GPa, respectively. This is consistent with the expected behaviour of balanced 2D woven composites, where the warp and weft yarns contribute similarly to the load-carrying capacity in the in-plane directions.

In contrast to the high values of the in-plane Young's moduli, the out-of-plane Young's modulus E_3 is significantly lower, with a value of 7.70 GPa. This highlights that the load transfer in the out-of-plane direction is primarily dominated by the matrix and the limited through-thickness reinforcement.

A similar trend can be observed for the shear moduli, where the in-plane shear modulus G_{12} with a value of 8.28 GPa is higher than the out-of-plane shear moduli G_{23} and G_{13}, with values of 5.77 GPa and 5.61 GPa, respectively.

Table 5.1: Effective elastic properties obtained from the ML-segmented virtual sample after transfer learning.

Property	**Value**
Young's Moduli	
E_1	55.833 GPa
E_2	56.631 GPa
E_3	7.700 GPa
Shear Moduli	
G_{12}	8.278 GPa
G_{23}	5.768 GPa
G_{13}	5.612 GPa
Poisson's Ratios	
ν_{12}	0.034
ν_{13}	0.353
ν_{23}	0.345

5.3.2 Resin-Injected Sample

The same modelling workflow was subsequently applied to the XCT scans of the resin-injected 5HS composite. The estimated effective elastic properties is summarised in Table 5.2. In contrast to the virtual sample, the resin-injected sample represents the as-manufactured composite sample and therefore includes geometric variations stemming from manufacturing, local yarn distortions, resin-rich regions, and potential imperfections captured directly from the XCT data.

The predicted in-plane Young's moduli for the resin-injected sample are $E_1 = 54.33$ GPa and $E_2 = 56.40$ GPa, respectively. Similarly to the virtual sample, the relatively small difference between the two in-plane directions is expected because both the weft and warp yarns contribute similarly to the load-carrying capacity in the in-plane directions, as previously mentioned.

Once again, the out-of-plane Young's modulus E_3 is significantly lower than the in-plane moduli, with a value of 8.11 GPa, as expected for this type of 2D woven architecture. The slightly higher value of E_3 compared to the virtual sample may indicate that the physical sample contains geometric features that increase the out-of-plane stiffness that are difficult to represent in the virtual sample.

Table 5.2: Effective elastic properties obtained from the ML-segmented resin-injected sample after transfer learning.

Property	**Value**
Young's Moduli	
E_1	54.332 GPa
E_2	56.404 GPa
E_3	8.106 GPa
Shear Moduli	
G_{12}	10.537 GPa
G_{23}	6.762 GPa
G_{13}	6.601 GPa
Poisson's Ratios	
ν_{12}	0.029
ν_{13}	0.366
ν_{23}	0.382

For the shear moduli, the resin-injected sample follows the same general trend as the virtual sample, with a higher in-plane shear modulus than the two out-of-plane shear moduli. The in-plane shear modulus G_{12} reaches 10.54 GPa, which is notably higher than the corresponding value for the virtual sample. The out-of-plane shear moduli G_{23} and G_{13} remain closer to the virtual sample, with values of 6.76 GPa and 6.60 GPa, respectively. The higher in-plane shear stiffness observed in the physical sample may be attributed to geometric features resolved in the XCT data that are not fully captured in the idealised virtual geometry.

The predicted elastic properties of the virtual and resin-injected samples are generally of similar magnitude and exhibit the same overall trends in orthotropic stiffness. However, the differences between the two sets of predictions should be interpreted with caution. As discussed further in Section 5.3.4, both models underestimate the experimentally measured in-plane Young's moduli. Therefore, the deviations between the virtual and resin-injected predictions cannot be attributed solely to manufacturing-induced geometrical effects, since they may also be influenced by uncertainties in the segmentation, RVE extraction, material assumptions, and homogenisation workflow.

5.3.3 Consistency Assessment of Homogenised Stiffness Matrices

To assess the numerical consistency of the homogenised stiffness matrices, a symmetry check was also performed both for the virtual and the resin-injected sample. This check is important due to the requirement from linear elasticity that the stiffness matrix must be symmetric. Deviations from symmetry therefore provide insight into numerical errors that might be introduced during segmentation, discretisation, and computational homogenisation. The results from the symmetry check are summarised in Table 5.3.

Table 5.3: Stiffness matrix symmetry check for the virtual and resin-injected samples.

Metric	**Virtual Sample**	**Resin-Injected Sample**
$\max \left\| \mathbf{C} - \mathbf{C}^{\mathrm{T}} \right\|$	40.08 MPa	201.54 MPa
Relative asymmetry	0.069%	0.349%

Regarding the virtual sample, the relative asymmetry of 0.069% confirms that the computed stiffness matrix is almost fully symmetric. The small asymmetry may be attributed to numerical errors stemming from geometry creation, finite element discretisation, solver tolerances, and post-processing.

The stiffness matrix symmetry check for the resin-injected sample further confirms the numerical consistency of the homogenisation procedure. Although the physical sample results in a larger relative asymmetry of 0.349% compared to the virtual sample, the stiffness matrix still remains close to symmetric. The increased asymmetry may be caused by the more geometrically complex architecture of the physical sample, leading to more segmentation imperfections, which worsen the symmetry of the obtained stiffness matrix.

5.3.4 Comparison with Experimental Results

To assess the accuracy of the predicted mechanical elastic properties, a comparison between the predicted values and the experimentally measured properties in Section 2.5 was carried out. As previously reported, the Young's moduli obtained from the tensile testing are approximately 61.57 GPa in the warp direction, i.e. E_1, and 63.03 GPa in the weft direction, i.e. E_2. The comparison of the experimentally measured and the numerically predicted effective elastic properties are summarised in Table 5.4.

Comparing the experimentally obtained values with the numerically predicted elastic properties for the virtual sample, it can be observed that the pipeline underestimates the in-plane stiffness. The predicted values $E_1 = 55.83$ GPa and $E_2 = 56.63$ are approximately 9-10% lower than the experimental values.

Similarly, the predicted elastic properties for the resin-injected sample are $E_1 = 54.33$ GPa and $E_2 = 56.40$ GPa. Once again, the predicted values underestimate the in-plane stiffness by

Table 5.4: Comparison of experimentally measured and numerically predicted effective elastic properties for virtual and resin-injected samples. Percentage differences are computed relative to experimental values.

Property	**Exp.**	**Virtual**	**Diff. (%)**	**Resin-Injected**	**Diff. (%)**
E_1 [GPa]	61.57	55.83	-9.32%	54.33	-11.75%
E_2 [GPa]	63.03	56.63	-10.15%	56.40	-10.55%
E_3 [GPa]	–	7.70	–	8.11	–
G_{12} [GPa]	–	8.28	–	10.54	–
G_{23} [GPa]	–	5.77	–	6.76	–
G_{13} [GPa]	–	5.61	–	6.60	–
ν_{12}	–	0.034	–	0.029	–
ν_{13}	–	0.353	–	0.366	–
ν_{23}	–	0.345	–	0.382	–

11-12% compared to the experimental values. This discrepancy between the predicted values and the experimental values could be due to several contributing factors related to the segmentation accuracy and the utilised FE homogenisation strategy.

One important factor to note is the remaining segmentation inaccuracies observed in the machine learning-based segmentation results, particularly in the 12-plane where yarn discontinuities and constituent misclassifications were still present, as discussed in Section 4.3.2. As the effective stiffness of the FE-models is highly dependent on continuous yarn paths throughout the volume, even small discontinuities in the segmented geometry may reduce the predicted effective stiffness. Interestingly enough, though, the virtual sample exhibits improved yarn continuity throughout the volume, while still showing a similar response to the physical sample. This indicates that other factors are present that significantly affect the predicted effective elastic properties, leading to an underestimation of the in-plane Young's moduli, compared to the experimental values. These factors may, for example, involve complex interactions between the yarns and the matrix which are difficult to capture with the current FE homogenisation model. However, due to the time limitations associated with this work, this possibility was not investigated further.

Another possible explanation for the underestimation of the in-plane stiffness is the out-of-plane thickness of the extracted FE model. Although this thickness should not significantly affect the predicted elastic properties if the volume is representative, accurately extracting an RVE from the 5HS weave architecture proved difficult. Increasing the out-of-plane thickness could therefore reduce the influence of local extraction inaccuracies by averaging them over a larger volume. This possibility was not investigated further due to time limitations.

Despite these inaccuracies, the predicted elastic properties lie reasonably close to the experimentally obtained values, as the predicted in-plane Young's moduli are within approximately 10% of the experimental values. This highlights that the evaluated machine learning-based segmentation and computational homogenisation pipeline is a promising tool that, with additional time and trials, may be capable of producing realistic effective elastic property estimates for the considered 5HS woven composite.

6

Conclusions and Future Work

6.1 Conclusions

The present work investigated the extension of an XCT-to-FE workflow, initially developed for 3D woven composites by Friemann et al. [12, 13], to a 2D woven textile-reinforced composite architecture representative of aerospace applications. More specifically, the work focused on evaluating whether the existing machine learning-based segmentation model and computational homogenisation framework, initially trained and applied to a 3D woven layer-to-layer angle interlock composite, could generalise and be directly applied to a 5HS woven CFRP architecture. This evaluation was carried out using both virtual 5HS models generated in TexGen, which provided access to known ground truth data, and physical XCT scans of manufactured material samples.

The virtual 5HS model was first virtually scanned and reconstructed, thereby serving as an initial benchmark case for evaluating segmentation performance against known ground truth data. In parallel, the virtual model was utilised to study the influence of XCT scan parameters and to support the planning of the physical XCT scans. The physical XCT scans of resin-injected and dry fibre samples were subsequently reconstructed and used as a second, more practically relevant benchmark case for assessing the applicability of the workflow to as-manufactured composite materials.

The results showed that the pre-trained segmentation model could partially generalise to the 5HS woven architecture, but that its accuracy was insufficient for reliable voxel-based FE-modelling. The main limitations were discontinuous yarn paths, inaccurate yarn boundaries, and misclassifications between warp and weft yarns. These errors are critical for voxel-based homogenisation, since the predicted stiffness depends on continuous load-carrying yarn paths, correct material assignment, and accurate mapping of local yarn orientations. For the virtual benchmark case, where ground truth data were available, the baseline model achieved a volume-wide pixelwise agreement of approximately 83%. Although this indicates that the model captured parts of the overall 5HS textile architecture, the slice-wise Dice scores revealed substantial class-dependent inconsistencies, particularly for the distinction between warp and weft yarns. For the physical resin-injected sample, where no ground truth segmentation was available, the assessment was limited to qualitative inspection. This showed that the pre-trained model could identify the approximate yarn locations and resin-rich regions, but the segmentation contained boundary inaccuracies, discontinuities and misclassifications, especially in regions with limited attenuation contrast between yarns and matrix.

Consequently, transfer learning was adopted. Additional synthetic training data representative of the considered 5HS architecture was generated and utilised for tuning the weights of the baseline segmentation model. Following transfer learning, a clear improvement in segmentation accuracy was observed for the virtual benchmark case, with the volume-wide pixelwise agreement increasing from approximately 83% to 94%, thereby exceeding the target value of 90%. The qualitative evaluation of the physical resin-injected sample also showed improved yarn continuity and clearer constituent separation, although some segmentation challenges remained, particularly in the 12-plane and near yarn overlap regions.

The performed investigations further demonstrated the strong dependency between XCT image quality and segmentation performance. In particular, comparison of the dry fibre and resin-injected XCT reconstructions showed that the dry fibre samples exhibited substantially improved attenuation contrast and clearer visual separation between material constituents. However, the dry fibre samples were not segmented using the neural network-based workflow in this work. Instead, they were used to assess the influence of resin on XCT image quality and attenuation contrast. The results therefore indicate that limited attenuation contrast remains one of the primary challenges when applying machine learning-based XCT segmentation to CFRP systems containing epoxy resin.

Following segmentation, voxel-based finite element models were generated from the segmented XCT reconstructions and subsequently utilised for computational homogenisation. The predicted elastic properties obtained from the XCT-to-FE workflow showed reasonable agreement with the experimentally measured effective stiffness in both principal material directions, with a maximum deviation of roughly 10% and 12% for the virtual and physical samples, respectively. Although differences between the numerical and experimental results were observed, the magnitude of the predicted elastic properties and the overall material response are considered sufficiently representative given the limitations associated with the segmentation accuracy, XCT resolution, and simplified material modelling assumptions adopted in the present work.

The experimental tensile testing further showed relatively consistent mechanical behaviour between the tested specimens, with only small differences observed between the warp and weft directions. However, it should once again be noted that the experimental campaign was limited to a small number of specimens and only considered in-plane tensile loading under linear elastic assumptions. Consequently, the experimental results should primarily be viewed as an initial validation of the developed workflow rather than a complete mechanical characterisation of the material system.

Overall, the results obtained in this thesis demonstrate that the XCT-to-FE workflow developed in [12, 13] can be extended to previously unseen 2D woven textile architectures through the introduction of additional training data representative of the target architecture. The work therefore shows the potential of combining XCT imaging, machine learning-based segmentation, and computational homogenisation for non-destructive characterisation of as-manufactured textile-reinforced composite materials.

6.2 Future Work

Although promising results were obtained in the present work, several aspects require further investigation and development before the methodology can be considered reliable for broader engineering applications.

A major challenge identified throughout the work was the limited attenuation contrast between the impregnated yarn regions and the surrounding matrix regions in the XCT reconstructions. Since the yarns consist of both carbon fibres and epoxy resin, their greyscale values can become closer to those of the surrounding matrix, making the separation of material constituents particularly challenging. Future work should therefore focus on improving the contrast and image quality of resin-injected composite samples through an investigation of XCT scan parameters, reconstruction settings, and filtering approaches. The influence of spatial resolution should also be assessed as part of a broader optimisation of the XCT imaging procedure, where voxel size, field of view, image noise, and material contrast are evaluated together for the considered material system.

Future work is also required regarding the segmentation methodology itself. Although transfer learning using synthetic 5HS data significantly improved the segmentation performance, the model remains limited by the range of geometrical variability represented in the training dataset. In this context, variability refers to differences in yarn geometry and manufacturing-induced features, such as yarn waviness, local distortions, tow spacing variations, resin-rich regions, and possible defects such as yarn misalignment. Future studies should therefore focus on generating more representative synthetic datasets that include a wider range of such variations within the considered 5HS architecture. This would extend the training space and make the model less dependent on extrapolating to unseen geometrical features, thereby improving its robustness when applied to as-manufactured 5HS composite samples with realistic manufacturing variability.

In addition, future work is required to quantify the sensitivity of the predicted elastic properties to uncertainties introduced throughout the workflow. This includes uncertainties related to XCT scan resolution, reconstruction artefacts, segmentation accuracy, and mapping of fibre orientation and local fibre volume fraction. Such investigations would be important in order to better assess the reliability of XCT-to-FE methodologies for engineering and aerospace applications.

The mechanical modelling approach considered in this work was limited to linear elastic voxel-based homogenisation. However, more advanced mechanical modelling should be introduced once the challenges related to XCT contrast and segmentation quality have been addressed, since reliable geometrical input is required for meaningful prediction of more complex material behaviour. Once these limitations have been reduced, the proposed pipeline should be extended to study non-linear material response, progressive damage, matrix cracking, and failure mechanisms. In addition, the computational efficiency of the workflow should also be investigated, particularly for larger RVEs or more computationally demanding solutions. Alternative mesh generation strategies beyond purely voxel-based discretisation should therefore be studied to reduce discretisation artefacts and improve computational efficiency.

The experimental testing performed in the present work was also limited. Future experimental testing campaigns should therefore include a larger number of specimens together with additional loading cases such as shear, compression, and out-of-plane loading. This would provide a more comprehensive validation of the predicted material properties and enable comparison with larger portions of the homogenised stiffness tensor.

Finally, while the present work focused on relatively small representative material volumes, future developments could investigate the applicability of the methodology to larger structures and component-level analysis. In the longer term, such developments could contribute towards XCT-based material characterisation and non-destructive verification methodologies for as-manufactured aerospace composite components.

Bibliography

[1] S. Hegde, B. Satish Shenoy, and K. Chethan, "Review on carbon fiber reinforced polymer (cfrp) and their mechanical performance," *Materials Today: Proceedings*, vol. 19, pp. 658–662, 2019. 1st International Conference on Manufacturing, Material Science and Engineering.

[2] Safran Group, "Leap-1a, a new-generation engine for single-aisle commercial jets," 2026. Accessed 21 Jan 2026.

[3] CFM International, "Leap overview," 2026. Accessed 21 Jan 2026.

[4] CFM International, "Cfm international reaches new milestones for leap engine maturity," 2025. Accessed 26 Jan 2026.

[5] L. Tong, A. P. Mouritz, and M. Bannister, *3D fibre reinforced polymer composites*. Elsevier, 2002.

[6] T. Gereke, O. Döbrich, M. Hübner, and C. Cherif, "Experimental and computational composite textile reinforcement forming: A review," *Composites Part A: Applied Science and Manufacturing*, vol. 46, pp. 1–10, 2013.

[7] L. Brown and A. Long, "8 - modeling the geometry of textile reinforcements for composites: Texgen," in *Composite Reinforcements for Optimum Performance (Second Edition)* (P. Boisse, ed.), Woodhead Publishing Series in Composites Science and Engineering, pp. 237–265, Woodhead Publishing, second edition ed., 2021.

[8] R. M. Auenhammer, L. P. Mikkelsen, L. E. Asp, and B. J. Blinzler, "Automated x-ray computer tomography segmentation method for finite element analysis of non-crimp fabric reinforced composites," *Composite Structures*, vol. 256, p. 113136, 2021.

[9] Y. Sinchuk, P. Kibleur, J. Aelterman, M. N. Boone, and W. Van Paepegem, "Variational and deep learning segmentation of very-low-contrast x-ray computed tomography images of carbon/epoxy woven composites," *Materials*, vol. 13, no. 4, 2020.

[10] C. Tang, J. Zou, Y. Xiong, B. Liang, and W. Zhang, "Automatic reconstruction of closely packed fabric composite rves using yarn-level micro-ct images processed by convolutional neural networks (cnns) and based on physical characteristics," *Composites Science and Technology*, vol. 252, p. 110616, 2024.

[11] J. Friemann, "textomos." `https://github.com/Johan-Friemann/textomos`. GitHub repository, accessed May 27, 2026.

[12] J. Friemann, L. P. Mikkelsen, C. Oddy, and M. Fagerström, "Synthetic, automatically labelled training data for machine learning based x-ray ct image segmentation: Application to 3d-textile carbon fibre reinforced composites," *Composites Part B: Engineering*, vol. 305, p. 112656, 2025.

[13] J. Friemann, C. Oddy, L. P. Mikkelsen, and M. Fagerström, "From x-ray ct to finite element models: A fully automated pipeline for mesoscale modelling of as-manufactured textile composites," *Composites Science and Technology*, vol. 278, p. 111561, 2026.

[14] OpenAI, "ChatGPT." `https://chatgpt.com`, 2026. Accessed: 2026-05-25.

[15] I. Chowdhury and J. Summerscales, "Woven fabrics for composite reinforcement: A review," *Journal of Composites Science*, vol. 8, 07 2024.

[16] M. S. Haque, "Processing and characterization of waste denim fiber reinforced polymer composites," 2014.

[17] ACP Composites, "Woven fabric style guide," 2011. Accessed 9 Feb 2026.

[18] B. Liu, S. Lan, J. Li, Q. Fang, Y. Ren, W. He, and H. Xie, "Digital image correlation in extreme conditions," *Thin-Walled Structures*, vol. 205, p. 112589, 2024.

[19] ZEISS, "ARAMIS High-Speed DIC Systems." https://www.zeiss.com/metrology/en/explore/topics/aramis-high-speed-dic-systems.html, 2026. Accessed: 2026-05-23.

[20] G. Quino, Y. Chen, K. R. Ramakrishnan, F. Martínez-Hergueta, G. Zumpano, A. Pellegrino, and N. Petrinic, "Speckle patterns for DIC in challenging scenarios: rapid application and impact endurance," *Measurement Science and Technology*, vol. 32, p. 015203, Jan. 2021.

[21] W. Tan and B. G. Falzon, "A crystal plasticity phenomenological model to capture the non-linear shear response of carbon fibre reinforced composites," *International Journal of Lightweight Materials and Manufacture*, vol. 4, no. 1, pp. 99–109, 2021.

[22] P. Potluri and V. Thammandra, "Influence of uniaxial and biaxial tension on meso-scale geometry and strain fields in a woven composite," *Composite Structures*, vol. 77, no. 3, pp. 405–418, 2007.

[23] J. M. Sietins, W. H. Green, and J. S. Jones, "Materials evaluation using x-ray computed tomography," in *Comprehensive Structural Integrity (Second Edition)* (M. H. F. Aliabadi and W. O. Soboyejo, eds.), pp. 159–199, Oxford: Elsevier, second edition ed., 2023.

[24] S. Garcea, Y. Wang, and P. Withers, "X-ray computed tomography of polymer composites," *Composites Science and Technology*, vol. 156, pp. 305–319, 2018.

[25] P. E. Kinahan, B. H. Hasegawa, and T. Beyer, "X-ray-based attenuation correction for positron emission tomography/computed tomography scanners," *Seminars in Nuclear Medicine*, vol. 33, no. 3, pp. 166–179, 2003.

[26] RX Solutions, "X-ray inspection for composite materials & fibers," 2020. Accessed 9 Feb 2026.

[27] Y. Huang, A. Preuhs, M. Manhart, G. Lauritsch, and A. Maier, "Data extrapolation from learned prior images for truncation correction in computed tomography," *IEEE Transactions on Medical Imaging*, vol. 40, pp. 3042–3053, Nov. 2021.

[28] S. Burgmann, M. Godehardt, K. Schladitz, and W. Breit, "Influence of voxel size for µct imaging of particles on measurement accuracy," *Construction and Building Materials*, vol. 289, p. 123148, 2021.

[29] Waygate Technologies, "Principles of ct operation: What determines the voxel resolution?," 2020. Accessed: 2026-05-25.

[30] R. M. Auenhammer, L. P. Mikkelsen, L. E. Asp, and B. J. Blinzler, "Automated x-ray computer tomography segmentation method for finite element analysis of non-crimp fabric reinforced composites," *Composite Structures*, vol. 256, p. 113136, 2021.

[31] M. Krause, J. M. Hausherr, B. Burgeth, C. Herrmann, and W. Krenkel, "Determination of the fibre orientation in composites using the structure tensor and local x-ray transform," *Journal of Materials Science*, vol. 45, no. 4, pp. 888–896, 2010.

[32] K. Naresh, K. Khan, R. Umer, and W. Cantwell, "The use of x-ray computed tomography for design and process modeling of aerospace composites: A review," *Materials Design*, vol. 190, p. 108553, 2020.

[33] A. Rashidi, T. Olfatbakhsh, B. Crawford, and A. S. Milani, "A review of current challenges and case study toward optimizing micro-computed x-ray tomography of carbon fabric composites," *Materials*, vol. 13, no. 16, p. 3606, 2020.

[34] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

[35] C. Li and C. Lee, "Minimum cross entropy thresholding," *Pattern Recognition*, vol. 26, no. 4, pp. 617–625, 1993.

[36] G. W. Zack, W. E. Rogers, and S. A. Latt, "Automatic measurement of sister chromatid exchange frequency," *Journal of Histochemistry and Cytochemistry*, vol. 25, no. 7, pp. 741–753, 1977.

[37] F. P. Vidal, S. Afshari, S. Ahmed, C. Atkins, É. Béchet, A. C. Bellot, S. Bosse, Y. Chahid, C.-Y. Chou, R. Culver, L. Dixon, J. Friemann, A. Garbout, C. Hatton, A. Henry, C. Leblanc, A. Leonardi, J. M. Létang, H. Lipscomb, T. Manchester, B. Meere, S. Middleburgh, I. Mitchell, L. Perera, M. Puig, and J. Tugwell-Allsup, "X-ray simulations with gVXR as a useful tool for education, data analysis, set-up of CT scans, and scanner development," in *Developments in X-Ray Tomography XV* (B. Müller and G. Wang, eds.), vol. 13152, p. 131520W, International Society for Optics and Photonics, SPIE, 2024.

[38] A. Horvatic Novak, B. Runje, Z. Keran, and M. Orošnjak, "Image artefacts in industrial computed tomography," *Tehnički glasnik*, vol. 14, pp. 434–439, 12 2020.

[39] K. Chen, Y. Wang, S. Ji, C. Meng, C. Pan, Z. Tang, and D. Chen, "A comprehensive review of ct artifact reduction: From traditional reduction techniques to deep learning methods," *Computerized Medical Imaging and Graphics*, vol. 129, p. 102728, 2026.

[40] W. van Aarle, W. J. Palenstijn, J. De Beenhouwer, T. Altantzis, S. Bals, K. J. Batenburg, and J. Sijbers, "The astra toolbox: A platform for advanced algorithm development in electron tomography," *Ultramicroscopy*, vol. 157, pp. 35–47, 2015.

[41] Waygate Technologies (Baker Hughes), "Phoenix v|tome|x s240 – ct inspection system," 2025. Accessed 26 Jan 2026.

[42] E. Maire and P. J. Withers, "Quantitative x-ray tomography," *International Materials Reviews*, vol. 59, no. 1, pp. 1–43, 2014.

[43] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*, 2015.

[44] J. Heaton, "Ian goodfellow, yoshua bengio, and aaron courville: Deep learning: The mit press, 2016, 800 pp, isbn: 0262035618," *Genetic Programming and Evolvable Machines*, vol. 19, 10 2017.

[45] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.

[46] K. H. Zou, S. K. Warfield, A. Bharatha, C. M. Tempany, M. R. Kaus, S. J. Haker, W. M. Wells, F. A. Jolesz, and R. Kikinis, "Statistical validation of image segmentation quality based on a spatial overlap index1: scientific reports," *Academic Radiology*, vol. 11, no. 2, pp. 178–189, 2004.

[47] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[48] C. Oddy, *Macroscale Modelling of 3D-Woven Composites: Inelasticity. Progressive Damage and Final Failure*. PhD thesis, Chalmers University of Technology, 2022.

[49] X.-Y. Zhou, P. Gosling, C. Pearce, Z. Ullah, and L. Kaczmarczyk, "Perturbation-based stochastic multi-scale computational homogenization method for woven textile composites," *International Journal of Solids and Structures*, vol. 80, pp. 368–380, 2016.

[50] Z. Xia, Y. Zhang, and F. Ellyin, "A unified periodical boundary conditions for representative volume elements of composites and applications," *International Journal of Solids and Structures*, vol. 40, no. 8, pp. 1907–1921, 2003.

[51] C. C. Chamis, "Mechanics of composite materials: Past, present, and future," *Journal of Composites Technology and Research*, vol. 11, no. 1, p. 3, 1989.

A

Appendix 1

A.1 Quantitative Machine Learning-Based Segmentation Performance Evaluation

A.1.1 Prior Transfer Learning

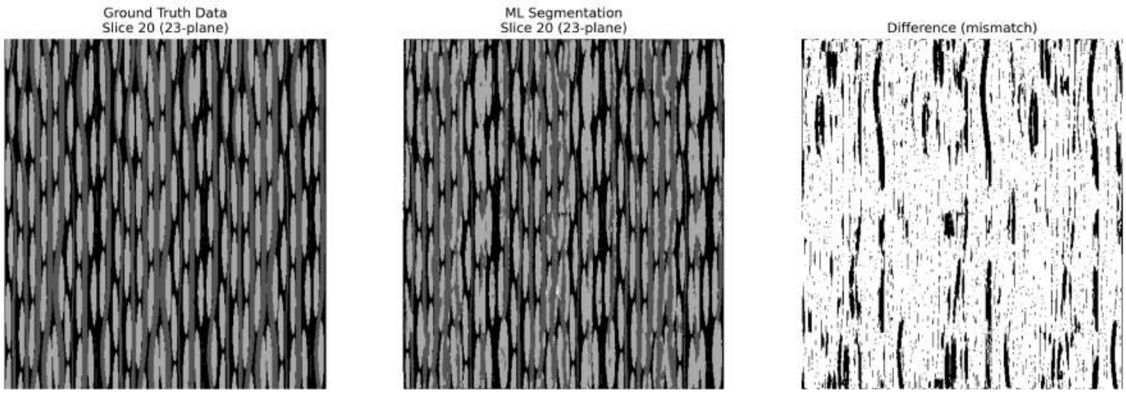

Figure A.1: Machine learning-based segmentation performance evaluation prior transfer learning, slice 20, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

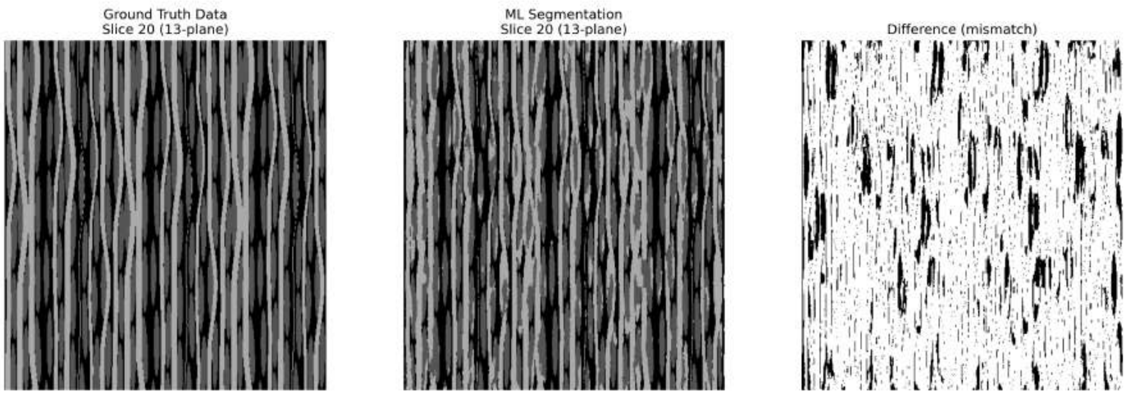

Figure A.2: Machine learning-based segmentation performance evaluation prior transfer learning, slice 20, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

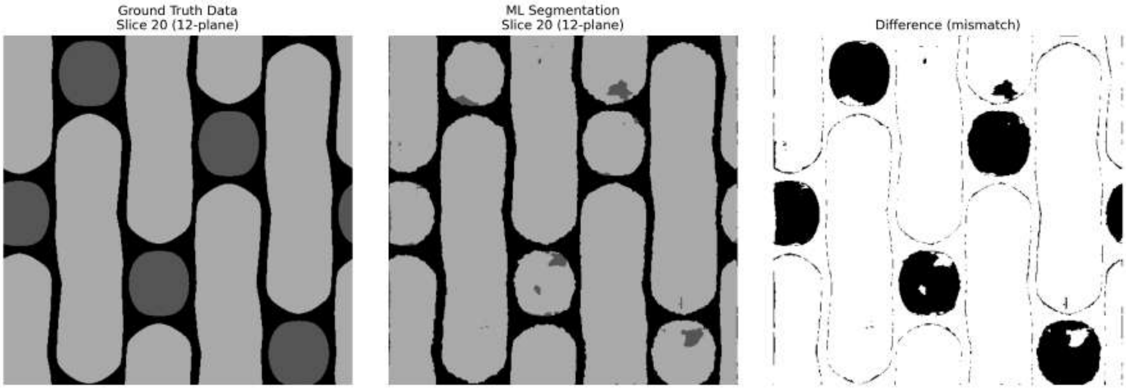

Figure A.3: Machine learning-based segmentation performance evaluation prior transfer learning, slice 20, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

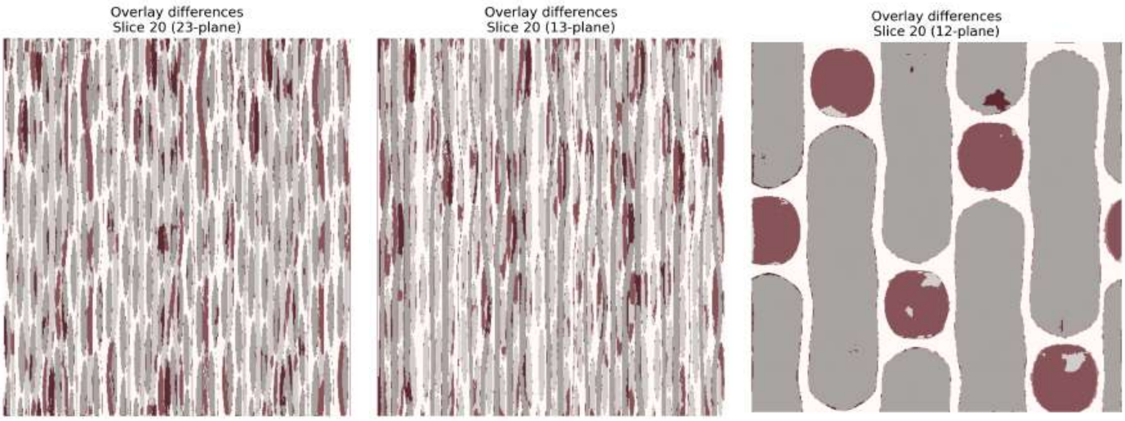

Figure A.4: Machine learning–based segmentation performance evaluation: difference overlays shown for slice 20 per plane of the virtual sample. Red indicates areas of difference.

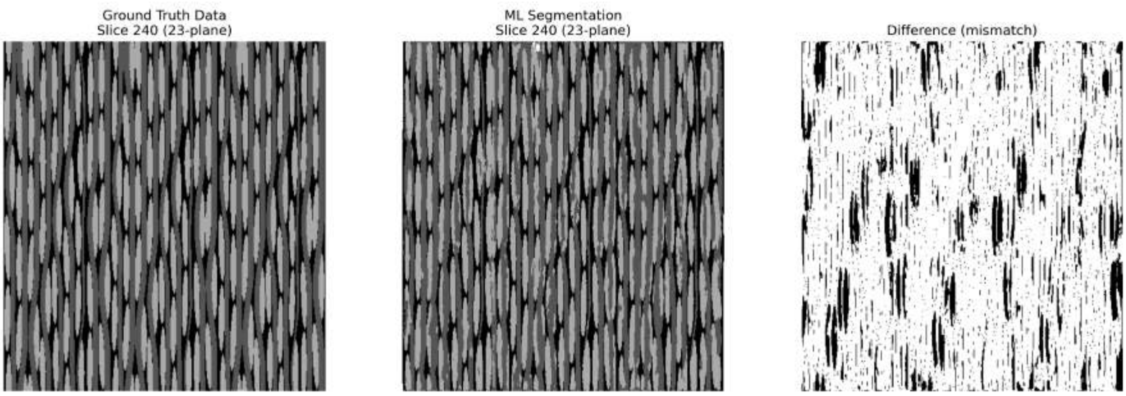

Figure A.5: Machine learning-based segmentation performance evaluation prior transfer learning, slice 240, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

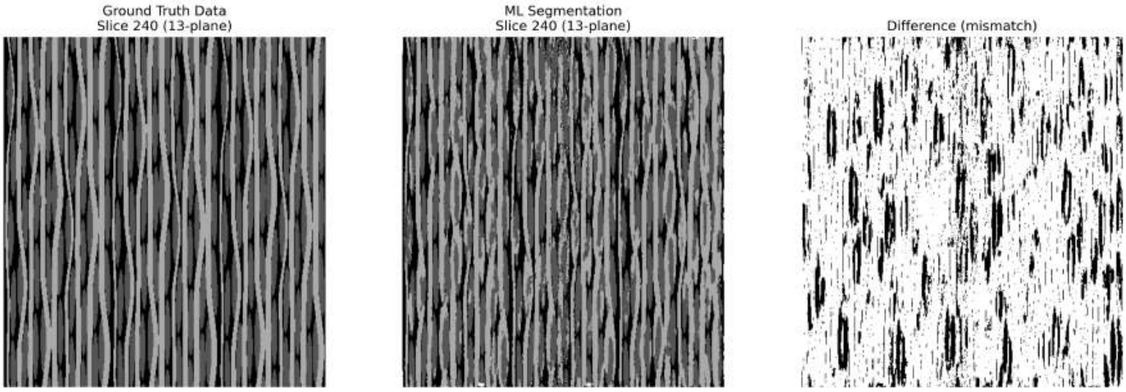

Figure A.6: Machine learning-based segmentation performance evaluation prior transfer learning, slice 240, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

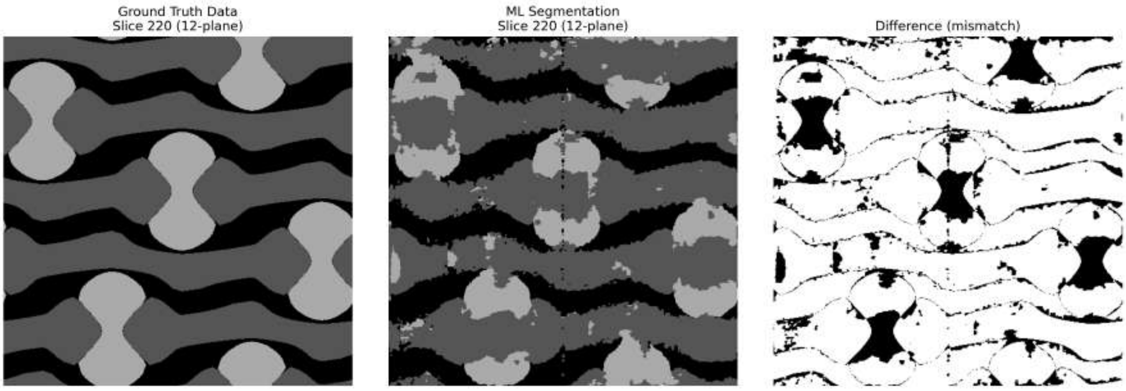

Figure A.7: Machine learning-based segmentation performance evaluation prior transfer learning, slice 220, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

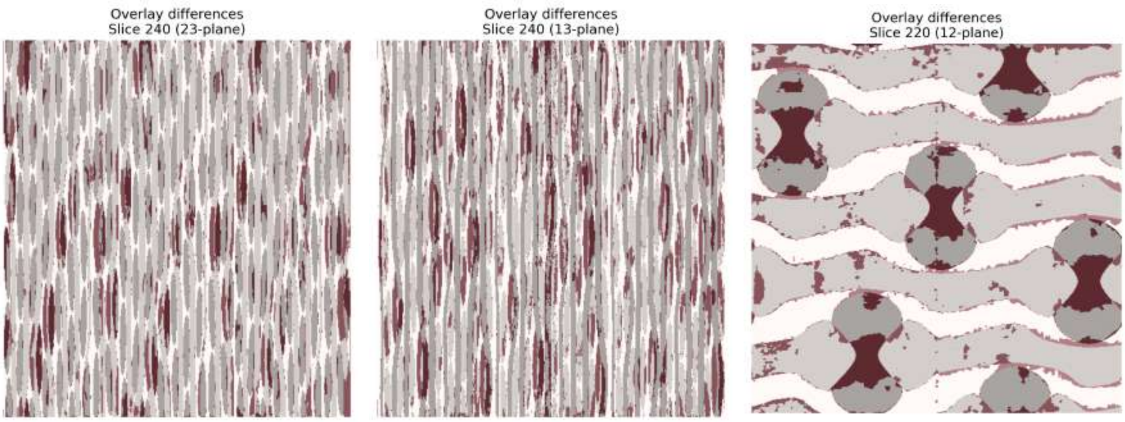

Figure A.8: Machine learning–based segmentation performance evaluation: difference overlays shown for slices 240 and 220 per plane of the virtual sample. Red indicates areas of difference.

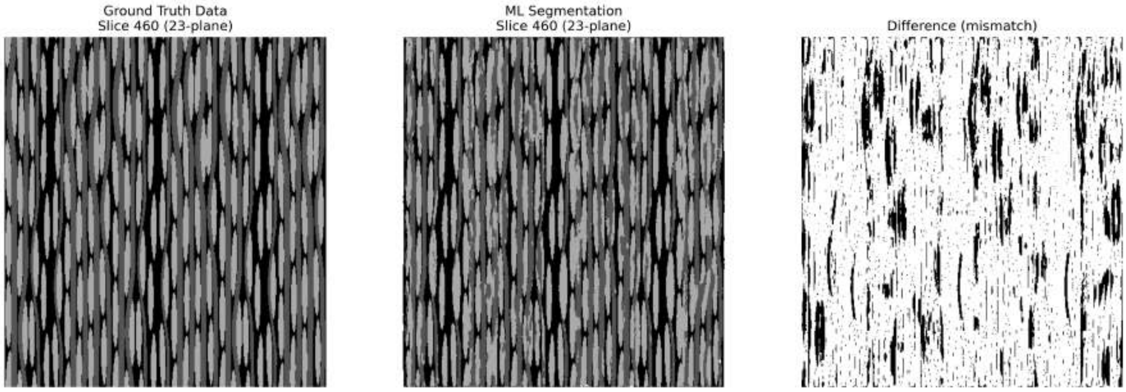

Figure A.9: Machine learning-based segmentation performance evaluation prior transfer learning, slice 460, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

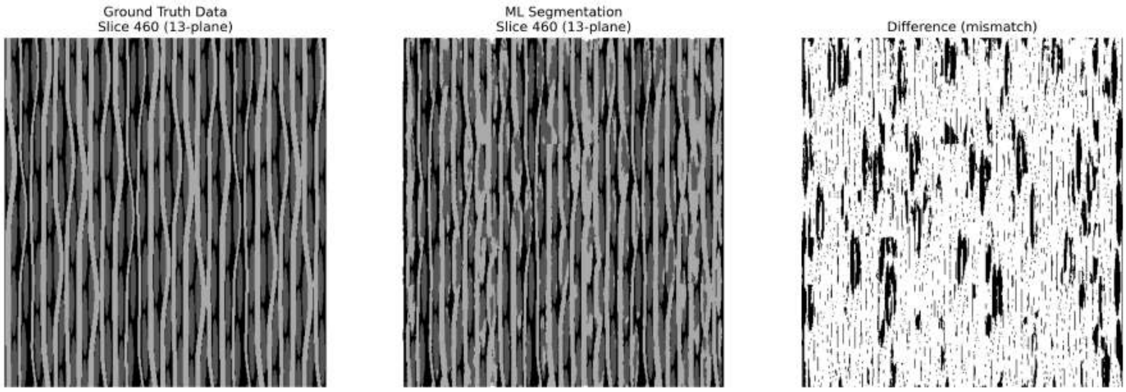

Figure A.10: Machine learning-based segmentation performance evaluation prior transfer learning, slice 460, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

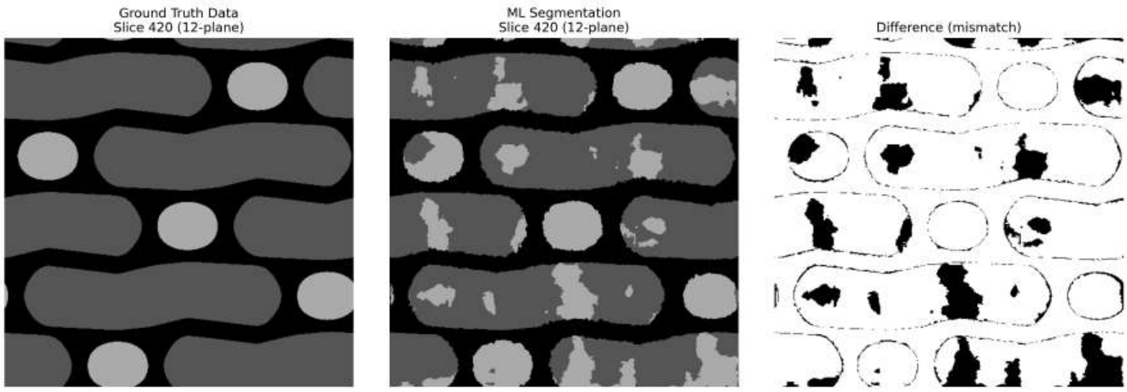

Figure A.11: Machine learning-based segmentation performance evaluation prior transfer learning, slice 420, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

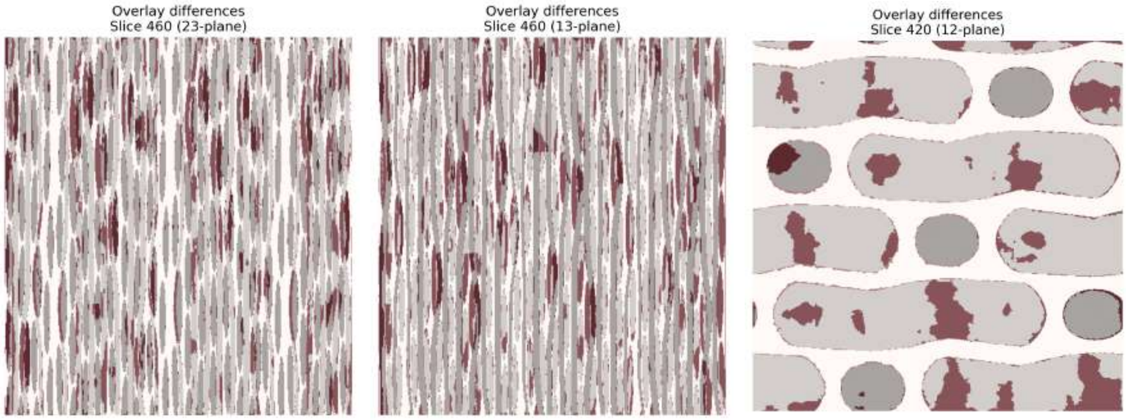

Figure A.12: Machine learning–based segmentation performance evaluation: difference overlays shown for slices 460 and 420 per plane of the virtual sample. Red indicates areas of difference.

A.2 Qualitative Evaluation on Physical XCT Data of Resin-Injected Sample

Figure A.13 shows the resin-injected sample segmented by the pre-trained network prior to transfer learning. The general weave architecture is captured, but boundary inaccuracies and local misclassifications remain, which may affect subsequent FE-based elastic property predictions.

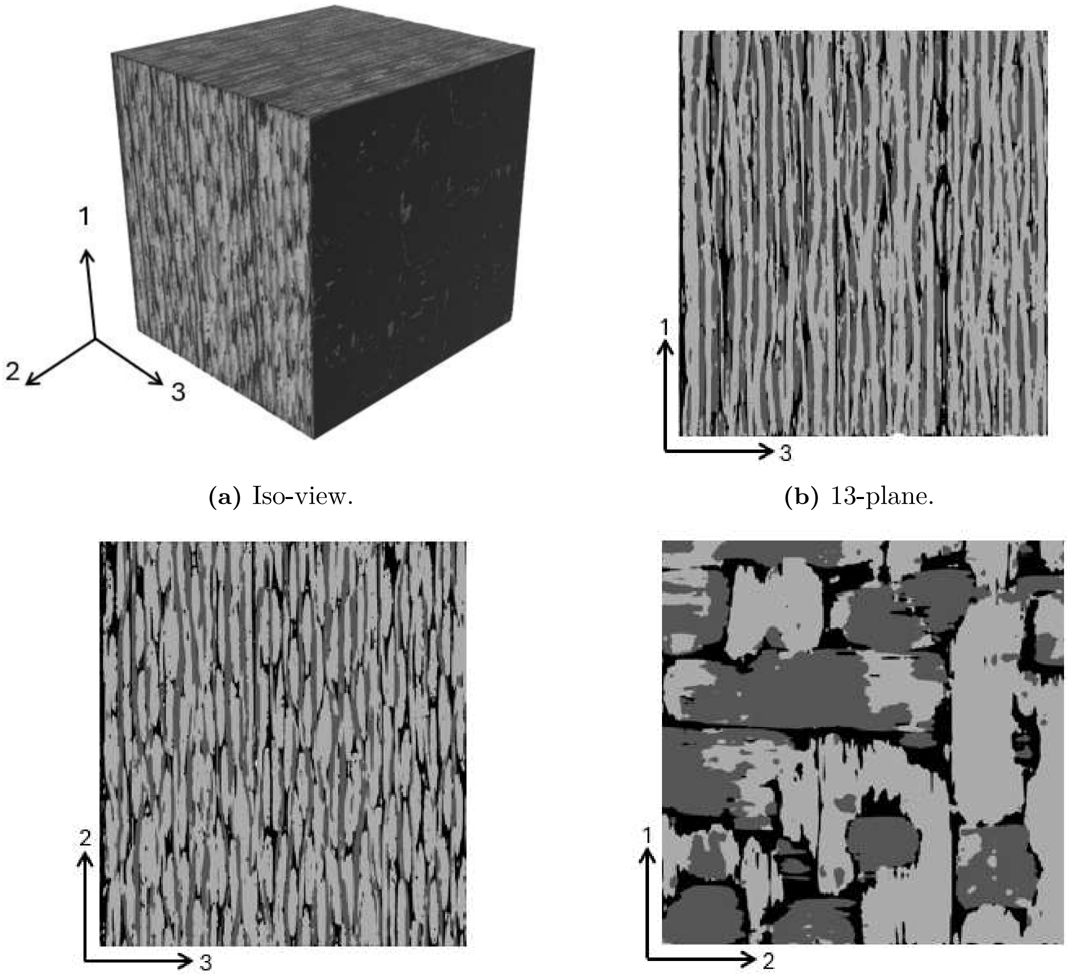

(a) Iso-view. **(b)** 13-plane. **(c)** 23-plane. **(d)** 12-plane.

Figure A.13: Machine learning segmented physical resin-injected sample prior to transfer learning. Black, dark grey, and light grey correspond to matrix, weft-yarn, and warp-yarn, respectively.

A.2.1 Post Transfer Learning

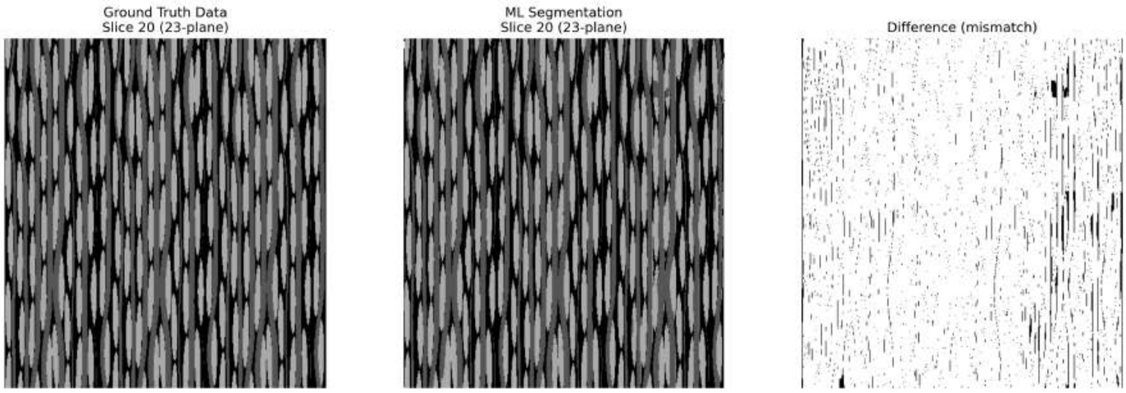

Figure A.14: Machine learning-based segmentation performance evaluation post transfer learning, slice 20, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

Figure A.15: Machine learning-based segmentation performance evaluation post transfer learning, slice 20, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

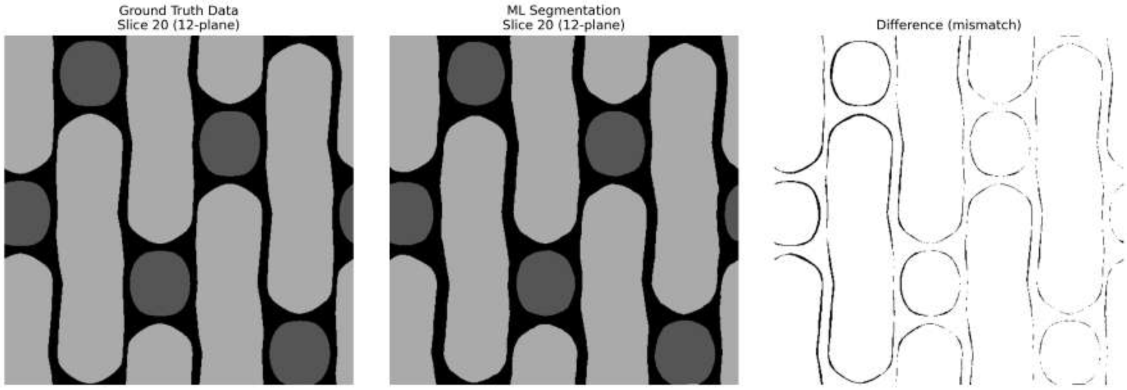

Figure A.16: Machine learning-based segmentation performance evaluation post transfer learning, slice 20, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

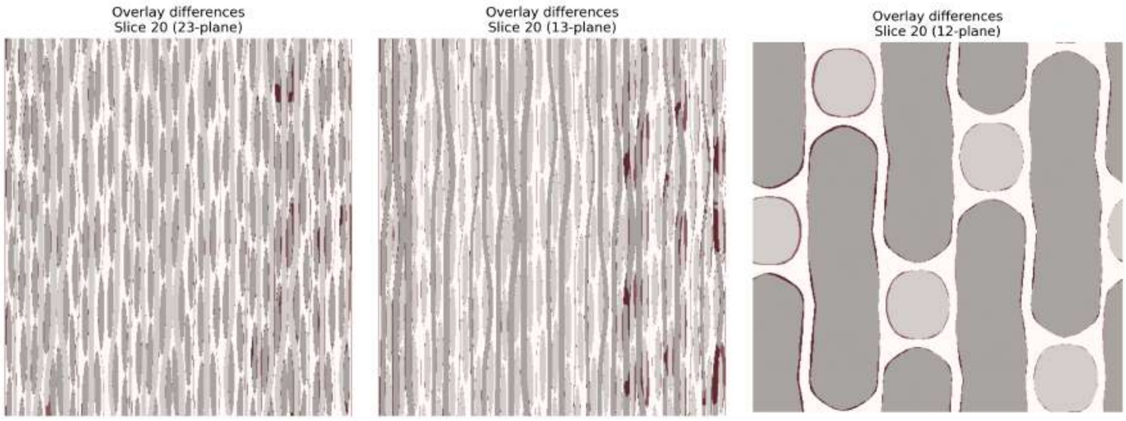

Figure A.17: Machine learning–based segmentation performance evaluation: difference overlays shown for slice 20 per plane of the virtual sample. Red indicates areas of difference.

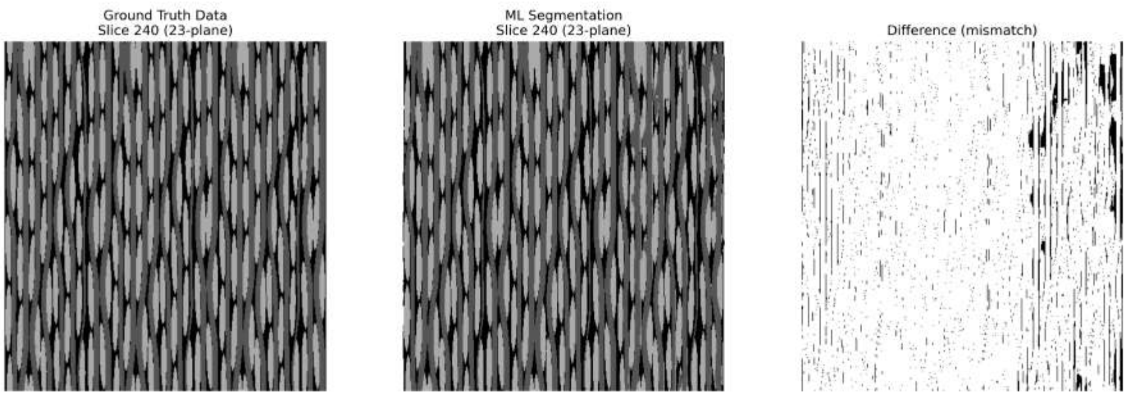

Figure A.18: Machine learning-based segmentation performance evaluation post transfer learning, slice 240, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

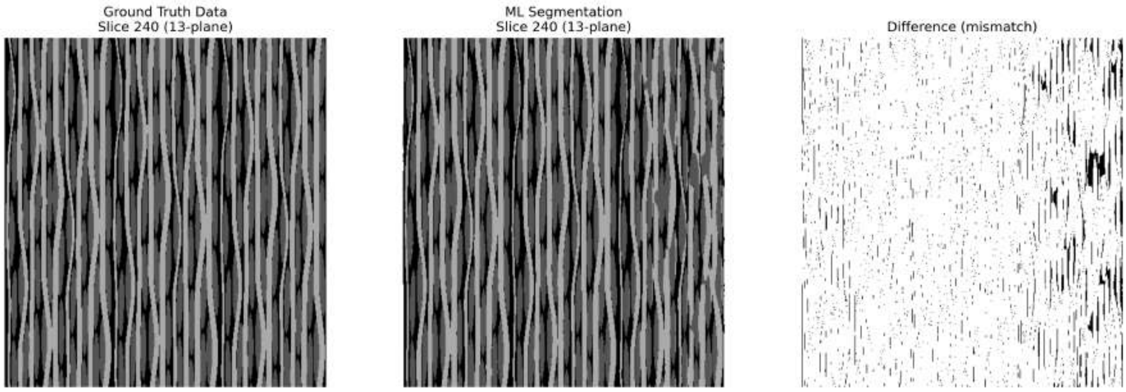

Figure A.19: Machine learning-based segmentation performance evaluation post transfer learning, slice 240, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

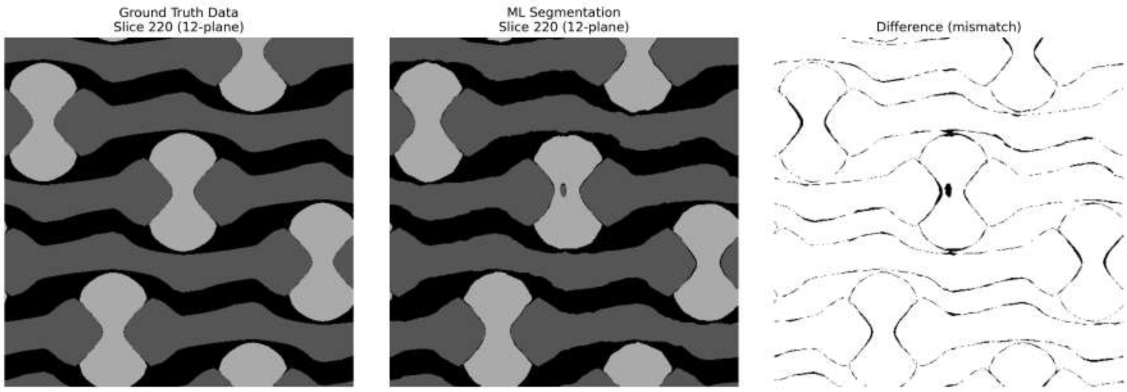

Figure A.20: Machine learning-based segmentation performance evaluation post transfer learning, slice 220, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

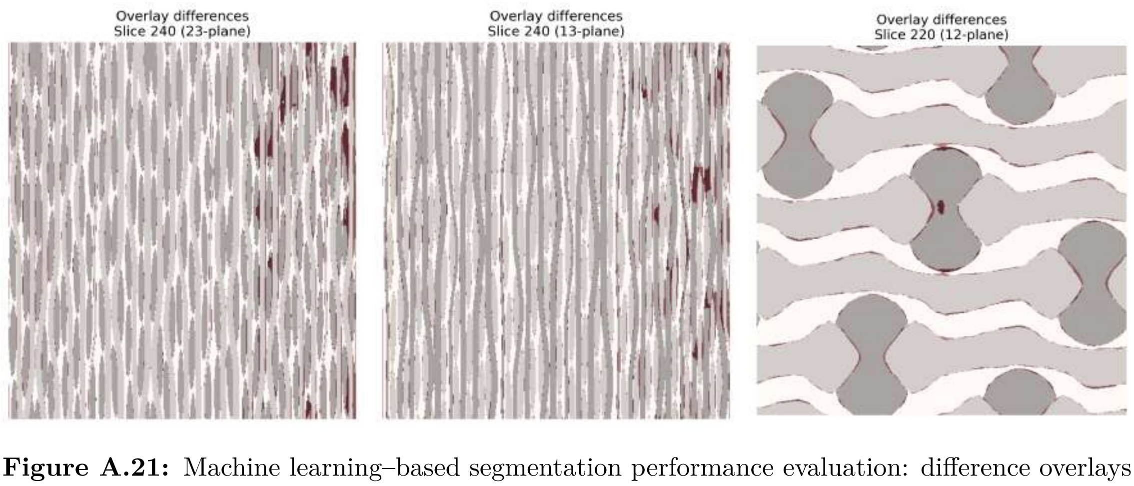

Figure A.21: Machine learning–based segmentation performance evaluation: difference overlays shown for slices 240 and 220 per plane of the virtual sample. Red indicates areas of difference.

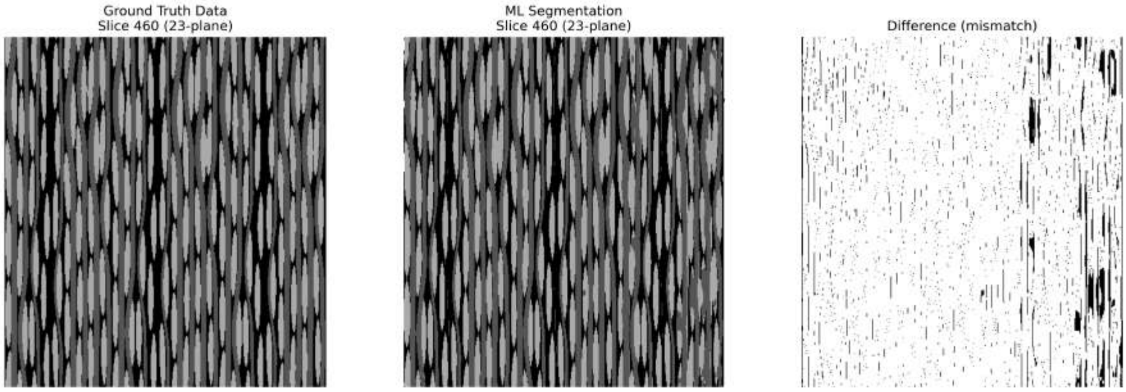

Figure A.22: Machine learning-based segmentation performance evaluation post transfer learning, slice 460, 23-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

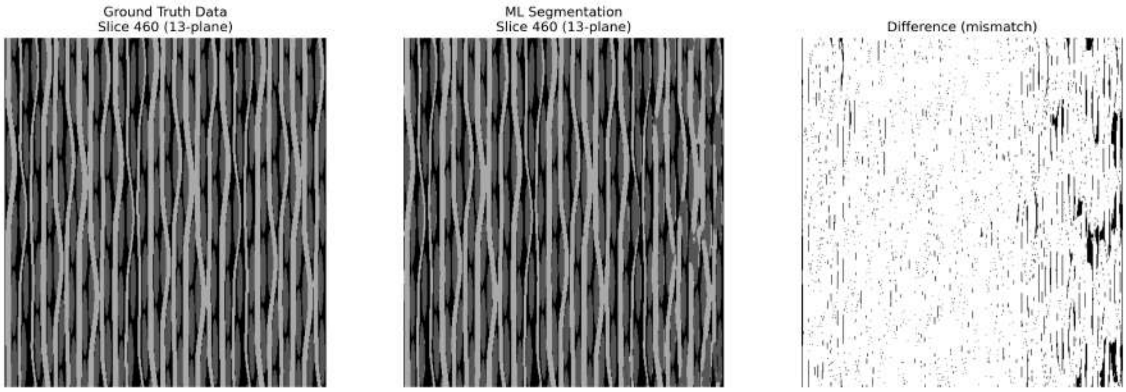

Figure A.23: Machine learning-based segmentation performance evaluation post transfer learning, slice 460, 13-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

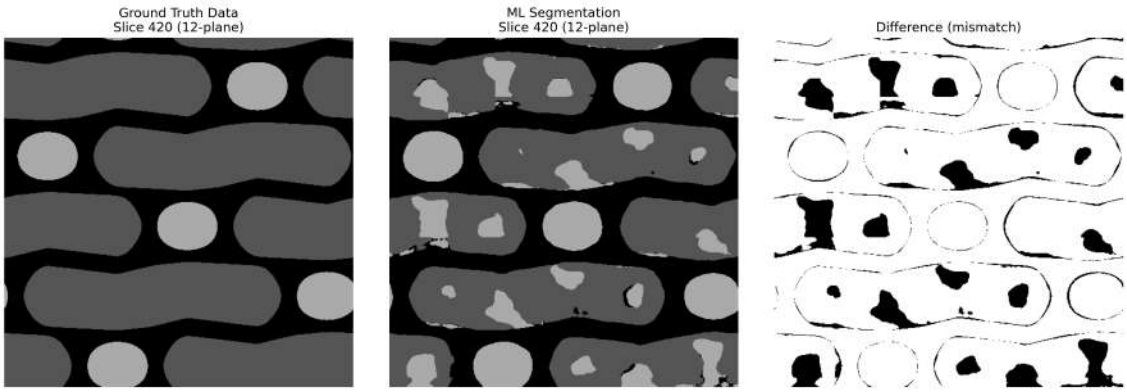

Figure A.24: Machine learning-based segmentation performance evaluation post transfer learning, slice 420, 12-plane. Black, dark grey, and light grey in the left-most two figures correspond to matrix, weft-yarn, and warp-yarn, respectively. In the right-most image white corresponds to agreement while black corresponds to differences.

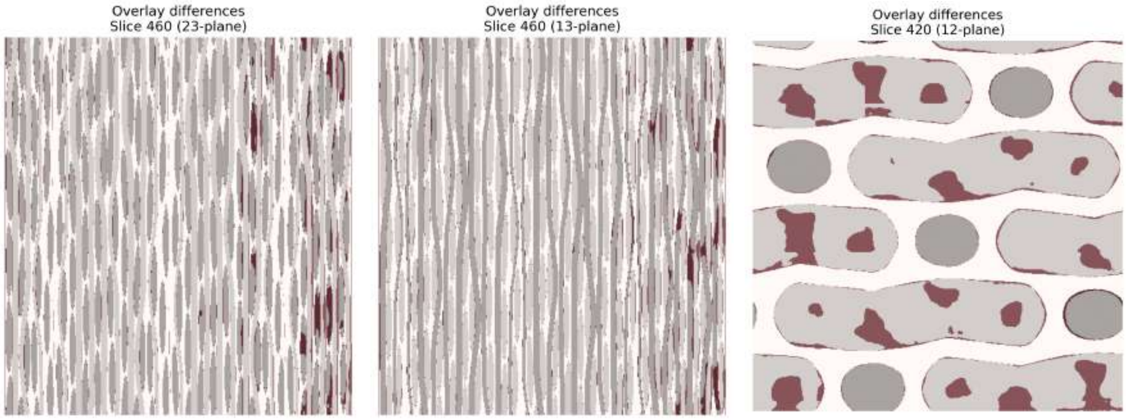

Figure A.25: Machine learning–based segmentation performance evaluation: difference overlays shown for slices 460 and 420 per plane of the virtual sample. Red indicates areas of difference.

DEPARTMENT OF INDUSTRIAL AND MATERIALS SCIENCE
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se