



## **From X-ray CT to finite element models: A fully automated pipeline for mesoscale modelling of as-manufactured textile composites**

Downloaded from: <https://research.chalmers.se>, 2026-02-27 13:19 UTC

Citation for the original published paper (version of record):

Friemann, J., Oddy, C., Mikkelsen, L. et al (2026). From X-ray CT to finite element models: A fully automated pipeline for mesoscale modelling of as-manufactured textile composites. *Composites Science and Technology*, 278. <http://dx.doi.org/10.1016/j.compscitech.2026.111561>

N.B. When citing this work, cite the original published paper.



## From X-ray CT to finite element models: A fully automated pipeline for mesoscale modelling of as-manufactured textile composites

Johan Friemann <sup>a</sup>, Carolyn Oddy <sup>a,c</sup>, Lars P. Mikkelsen <sup>b</sup>, Martin Fagerström <sup>a</sup>

<sup>a</sup> Department of Industrial and Materials Science, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

<sup>b</sup> Composites Manufacturing and Testing, Department of Wind and Energy Systems, Technical University of Denmark, DK-4000 Roskilde, Denmark

<sup>c</sup> GKN Aerospace Sweden, SE-461 81 Trollhättan, Sweden

### ARTICLE INFO

#### Keywords:

X-ray computed tomography  
Segmentation  
Machine learning  
Finite element modelling  
3D-textile reinforced composites

### ABSTRACT

Structural parts incorporating 3D-textile reinforced composites show great promise in high performance lightweight applications. For widespread industrial use, accurate predictions of mechanical properties are required. Experimental testing campaigns to generate these properties can be prohibitively expensive. To overcome this, meso scale models of the yarn architecture can be derived from X-ray Computed Tomography (XRCT), and computational homogenisation can be performed in the material's as-manufactured configuration. In this work a fully automated pipeline for the prediction of the full 3D elastic properties of 3D-reinforced textile composites from XRCT scans is presented. The proposed methodology enables the study of variations in the as-manufactured material properties from a single large field of view XRCT scan encompassing multiple unit cells, such that finite element homogenisation can yield statistical information. The pipeline includes a machine learning based segmentation model, finite element meshing and boundary condition assignment, and a material mapping procedure. For segmentation, a for textile reinforced composites completely novel 3D U-net architecture can be utilised, owing to the use of a fully synthetic automatically labelled training data set. An application of the pipeline on a 3D-reinforced material sample results in accurately predicted homogenised elastic stiffnesses, with a deviation from experiments of less than 6.5%.

### 1. Introduction

Composite components incorporating 3D-textiles are gaining attention partly due to their excellent impact and delamination resistance [1]. This, among other factors is important for safety in aerospace where bird strikes, fan blade-offs, and high dynamic loads are ever present risks. Therefore, textile composites with 3D-reinforcements are becoming more widespread in the aerospace industry. The commercially available LEAP engine, for example, utilises 3D-woven carbon fibre reinforced polymers (CFRP) in its fan blades and casing [2]. In order to increase the adoption of 3D-woven composites in the aerospace industry, accurate determination of elastic properties is paramount. For certification purposes it is further motivated to be able to give statistical measures on structural as-manufactured parts. A clear understanding of how material variations influence stiffness and strength, and how these properties are linked to the micro and meso scale material morphology, is essential for component manufacturers to reduce their safety factors. Reduced safety factors will lead to reduced use of material, making components more lightweight while saving raw material and energy.

To elucidate the effect of material morphology variations on mechanical properties, two approaches can be taken: experimental testing or numerical modelling. With physical testing, an excessive amount of samples would have to be manufactured and tested to obtain the necessary statistics. Furthermore, given that the material heterogeneity is typically in the order of tens of millimetres, standard test method for CFRP materials cannot be used. As a consequence, standard test methods have been adjusted such that the gauge regions accommodate a sufficiently large representative volume, cf. e.g. [3–6]. Still, these factors make numerical modelling a more feasible path forward.

In terms of modelling, the material could realistically be considered on two scales: the meso (yarn) scale or the macro (homogenised) scale. Aiming to model large structural parts with high fidelity meso scale models, where reinforced yarns are explicitly represented, are prohibitively computationally expensive. Therefore, a macro scale homogenised model becomes necessary for evaluating structural designs, especially in the prototyping stage of product development. For such models to be reliable, the macro scale properties and their spatial variation have to be used as input.

\* Corresponding author.

E-mail address: [johan.friemann@chalmers.se](mailto:johan.friemann@chalmers.se) (J. Friemann).

<https://doi.org/10.1016/j.compscitech.2026.111561>

Received 8 December 2025; Received in revised form 12 January 2026; Accepted 7 February 2026

Available online 9 February 2026

0266-3538/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Numerically, such equivalent properties are typically derived through finite element (FE) simulations of representative volume elements (RVE), which capture the characteristic periodic structure of the woven material's meso scale. For the purpose of performing RVE simulations, geometry that captures the material morphology accurately needs to be derived. Geometry can be derived through purely geometric means using tools such as TexGen [7] or WiseTex [8]. This method can, however, result in far too idealised geometries that do not capture the as-manufactured properties of the studied material. As a consequence, there has been extensive work on modelling the weave process itself to generate more realistic RVE's. This started with the Digital Element Method (DEM) through Wang et al. [9], refined by Zhou et al. [10] and has been further improved by several different research groups [11–14]. One drawback of this second approach is the extensive CPU time required per unit cell (UC), which makes larger testing campaigns challenging. An alternative high performance method utilising position based dynamics has recently been presented for yarn deformation modelling by Gao et al. [15]. For all these cases, it is however unclear how variability in the meso structure can be accounted for, even more so how it can be linked to variability in part geometry or manufacturing.

That said, in recent years it has become popular to derive computational geometry from x-ray computed tomography (XRCT) scans. A clear benefit is that computational homogenisation of textile composites then can be performed in the textile's as-manufactured state, yielding realistic estimates of the elastic material properties. In fact, components could be scanned at various locations, such that realistic material variability could be acquired and accounted for. Ewert et al. [16] have shown that such geometries can outperform comparable DEM derived geometries.

The main challenge of deriving meso scale RVE geometry from XRCT scans is the segmentation problem. Namely, categorising each voxel in the volume as its corresponding phase (pure matrix, weft yarn, warp yarn, or voids). Due to the significant material heterogeneity, and thereby to capture the entire UC, scans need to include a fairly large field of view. The demands on field of view increase even further if material variability statistics are to be extracted. For efficiency, scanned samples need to include several UCs. Having the desired field of view will inevitably result in poor contrast, making segmentation through classical methods, such as thresholding or watershed, difficult or even impossible [17,18]. Furthermore, assigning local yarn fibre volume fractions and local fibre orientations poses additional challenges.

As alternatives to thresholding or watershed, other researchers have utilised commercial fibre tracking software [19], structure tensor analysis [20–22], or variational methods [23–25] for segmentation. For particularly large field of view scans, which is necessary for cm scale UCs and/or multiple UC scans, these methods might fail. The reason is that they all rely on structural information that may not be resolved in a low contrast scan, or require very long computational times which could make industrial application difficult.

Machine learning based segmentation methods have become increasingly popular for woven composites, and are seeing widespread use [26–28]. A drawback of machine learning based methods is the need for large labelled data sets. There has been work to get around this limitation by using synthetic data. This has for example been done by creating “pseudo CT” by training a surrogate model [29], using “pseudo labels” derived from classical image processing methods [30], or fortifying the data set with synthetic data based on statistics from a manually labelled data set [31]. A fully synthetic data has been used for segmentation of a six-phase Al–Si alloy composite reinforced with ceramic fibres and particles by Tsamos et al. [32]. Furthermore, the current authors recently proposed a fully automated pipeline for training a machine learning based segmentation model entirely using synthetic data [33]. The initial application of this model to segment an XRCT scan of a layer-to-layer angle interlock reinforced composite material showed promise in terms of obtaining acceptable voxel-wise

accuracy. A drawback was, however, that continuity of reinforcement yarns could not be guaranteed, something that is being addressed in the current work.

With an accurate segmentation at hand, as-manufactured properties can be derived through computational homogenisation. In this work a fully automated and open source<sup>1</sup> [34] pipeline to generate the necessary finite element models directly from large field of view CT scans of 3D-woven composites is presented. The pipeline extends previous work in utilising fully synthetic training data for machine learning based segmentation of 3D-reinforced textile composites. In addition to segmentation, it now also handles model discretisation, mapping of local fibre orientation and volume fraction, FE input file generation, and elastic property homogenisation. The only manual labour required is the identification of individual UCs from a large field of view segmentation. To promote 3D continuity of the segmented volume, the segmentation problem is handled with a novel 3D-convolutional neural network. Extending the work by Jing et al. [35], who performed meso scale elastic homogenisation of orthogonal weaves using an XRCT guided but manually defined geometry, the current work directly utilises the segmented volume information as basis for the numerical models. Furthermore, the current pipeline is free from any restrictions on (partially) manual data labelling, as done e.g. by Ahmed et al. [36]. Thereby, the pipeline enables a fully automated and efficient generation of multiple UC models from a single scan.

Leveraging the opportunity to automatically generate multiple UC models for macroscopic property extraction, we demonstrate the pipeline on a layer-to-layer angle interlock composite. To study variability of properties and how these are linked to variations in the reinforcement architecture, six unique RVEs are derived from a single large field of view XRCT scan. We demonstrate that macroscopic elastic properties can be predicted in good agreement with experimental observations. As an additional finding, we observe that the variability in macroscopic properties is very small, even when noticeable variations in the reinforcement architecture are present.

## 2. Material sample and its XRCT scan

The studied material is a 3D-woven layer-to-layer angle interlock CFRP manufactured through resin transfer moulding (RTM). The yarns are made out of IM7 fibres [37] with a diameter of 5.2  $\mu\text{m}$ . The warp and weft yarns are illustrated in blue and red respectively in Fig. 1(a). Warp yarns run in the main weaving direction of the machine and are held in tension during the manufacturing process. The weft yarns on the other hand run transversely to the warp yarns and are inserted during the weaving process. The fibre count is 12k and 24k for the weft and warp yarns, respectively. The matrix consists of RTM6 epoxy [38]. It should be noted that the weave has a quite large UC when compared with the scale of the fibre diameter. Further, it should be noted that the UC of this material is shorter in the warp direction compared to the other two directions. As shown in Fig. 1(a), the UC is 6.5 mm wide in the warp (1) direction, 21.5 mm long in the weft (2) direction, and 4 mm thick in the out-of-plane (3) direction. Fig. 1(a) also denotes and visualises the orientation conventions used in this paper. Further, it should be mentioned that the material contains some glass fibre tracer yarns that were used as aid during manufacturing. Since glass has a much higher x-ray attenuation than carbon fibres and epoxy, this can cause artifacts during XRCT scanning.

For the scan, a sample was created from 3 pieces cut out of a cured panel and stacked on top of each other using double-sided tape, see Fig. 1(b). Each piece contains roughly 2 by 3 UCs. The stacking is performed to obtain a relatively low aspect ratio in the 13-plane of the sample,

<sup>1</sup> All data not explicitly referenced, that has been used in this manuscript can be generated with the provided code. If the original data (roughly 700 GB) is needed it will be made available on request.

thereby yielding a relatively uniform x-ray attenuation when scanning. As a benefit, this allows for a large field of view XRCT scan which encompasses as many UCs as possible in a smaller scanning volume.

### 2.1. XRCT scan and sample reconstruction

The sample was scanned with a Zeiss XRadia 410 Versa lab-based XRCT scanner. It uses a Hamamatsu Photonics L8121-03 micro focus x-ray source with a tungsten target. There is a documented 200  $\mu\text{m}$  beryllium filter inside the source, but initial investigations suggests that there may be an additional undocumented 1 mm aluminium filter inside the source [33]. No additional filter is used. The detector in the imaging system is an ANDOR iKon-L DW936N-BV-556 CCD detector, with 2048 by 2048 pixels with dimensions 13.5  $\mu\text{m}$  by 13.5  $\mu\text{m}$ . One special feature of the Zeiss XRadia 410 Versa is the presence of an optical magnification system between the x-ray scintillation panel and the detector. Here the 0.4 times magnification large field of view objective was used. The source to sample distance was set to 80 mm, while the sample to detector distance was 150 mm. An x-ray tube voltage of 40 kV and power of 10 W was used. The sample was rotated 360° where 3001 projections, each with an exposure time of 5 s, were taken at equal angular intervals. Binning 4 was performed, which implies that each pixel in the output projections consisted of the average of 4 by 4 pixels. The geometric and objective magnification combined with the binning results in an effective voxel size of 46.96  $\mu\text{m}$ .

The Feldkamp–Davis–Kress (FDK) technique [39], which is a generalisation of filtered back projection for cone beam scans, was used for tomographic reconstruction. The GPU implementation of FDK inside the ASTRA toolbox [40] was used to perform the reconstruction. The reconstructed scan and a manually annotated segmentation of the scan are hosted online [41].

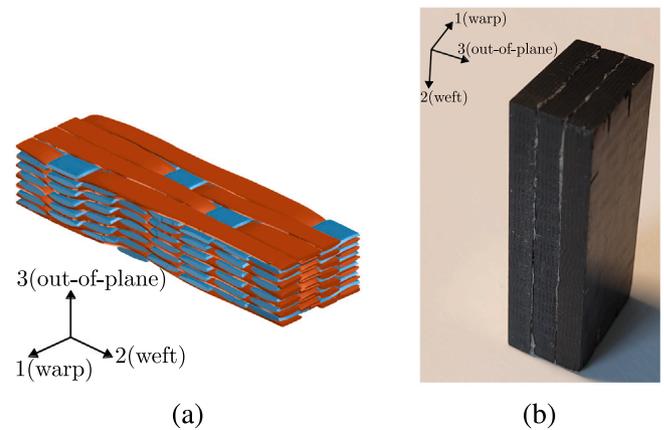
A selection of cross-sectional images of the reconstruction are shown in Fig. 4 (left) below. The combination of the large field of view XRCT scan, together with the lack of contrast between the impregnated yarns and matrix pockets, makes automated segmentation (i.e. the assignment of a material phase to each voxel) challenging using traditional methods. For this reason, a machine learning based segmentation algorithm provides a promising path forward.

## 3. Machine learning based segmentation of textile reinforced composites

The aim of the study is to develop a fully automated pipeline to predict the full elastic stiffness tensor of textile reinforced composites directly from CT images. As a first stage, segmenting the reconstruction is necessary. This process and its application to the sample introduced in Section 2 is discussed in more detail in the following subsections. It is further highlighted here that all novel contributions to this work are open source [34]. The training of the segmentation model has been carried out on a laptop with an Intel i9-13900H CPU, 32 GB of RAM, a 1 TB NVMe M.2 Performance SSD, and an Nvidia RTX 4080 Laptop GPU (12 GB VRAM).

### 3.1. Development and training of the machine learning model

Segmentation in this context entails assigning each voxel in the reconstructed XRCT scan to one of four phases: weft yarns, warp yarns, pure matrix, or air. For CFRP composites in particular, the similarity in x-ray attenuation between pure carbon fibres and hydro-carbon epoxy yields volumes with very poor contrast. As a consequence, classical segmentation strategies such as thresholding and watershed are often insufficient. Therefore, machine learning based segmentation models are gaining in popularity, as they can learn to effectively distinguish phase boundaries while ignoring noise and artifacts. The use of machine learning to carry out segmentation requires the generation of labelled



**Fig. 1.** The studied layer-to-layer angle interlock material is shown with the material coordinate system (1, 2, 3). A schematic of the UC is shown in (a) where weft yarns are shown in red, and warp yarns are shown in blue. Original graphics provided by courtesy of Oddy [6]. The specimen scanned in this study is displayed in (b).

3D training data, the construction of an appropriate neural network architecture and a strategy for the training process.

When it comes to the generation of labelled 3D training data, one major drawback and criticism of machine learning based approaches is the need for large data sets. This makes the use of traditional hand labelling challenging and exorbitantly time consuming. A possible alternative, and the one considered here, is the use of synthetic training data in which the constituents are automatically labelled. More specifically, in this work a method for the generation of synthetic, automatically labelled training data, recently proposed by the authors [33], is used. Quantitative similarity to real data was demonstrated through comparisons of intensity histograms. Further, the ability of a neural network trained on synthetic data to segment real tomogram slices was demonstrated with 88% pixel-wise accuracy. In contrast to previous work in the field, cf. e.g. Tsamos et al. [32], the synthetic data includes physically motivated attenuation features, noise, and artifacts. Using this method, a data set with 700 volumes used for training, 150 volumes for validation and 150 volumes for testing was generated. The exact same generation parameters used in the previous work were also used here.

Once the training set has been generated, the architecture of the neural network must be considered. In the context of segmenting scans of textile reinforced composites, 2D U-net architectures have been widely used with great success, cf. e.g. Mendoza et al. [29], Zheng et al. [31], Sinchuk et al. [42], Ying et al. [28], and Cao et al. [43]. However, 2D U-net architectures lack the ability to ensure continuity from one volume slice to the next. With the considered aim of generating volumetric model representation of the scanned samples a 3D-convolutional network based on the U-net architecture was considered here. This choice was further supported by the promising results by Tsamos et al. [32] in segmenting an advanced multiphase composite using synthetic training data and a U-net with 3D convolution kernels.

The considered neural network segmentation model is implemented in the PyTorch [44] framework. The implemented architecture is displayed in Fig. 2. The network is constructed from 5 main building blocks. The *Double convolution* block consists of two chained groups consisting of a 3D convolution layer, a batch normalisation layer, and a ReLU activation function. The convolution layers use a kernel size of 3 and a padding of 1. The *Max pool* block consists of a 3D max pool layer that utilises a kernel size of 2 and a stride of 2. The *Up convolution* block refers to a 3D transposed convolution with a kernel size of 2 and stride of 2. The *Concatenation* block concatenates the two inputs channel-wise. The final convolution block is a 3D convolution layer with kernel size

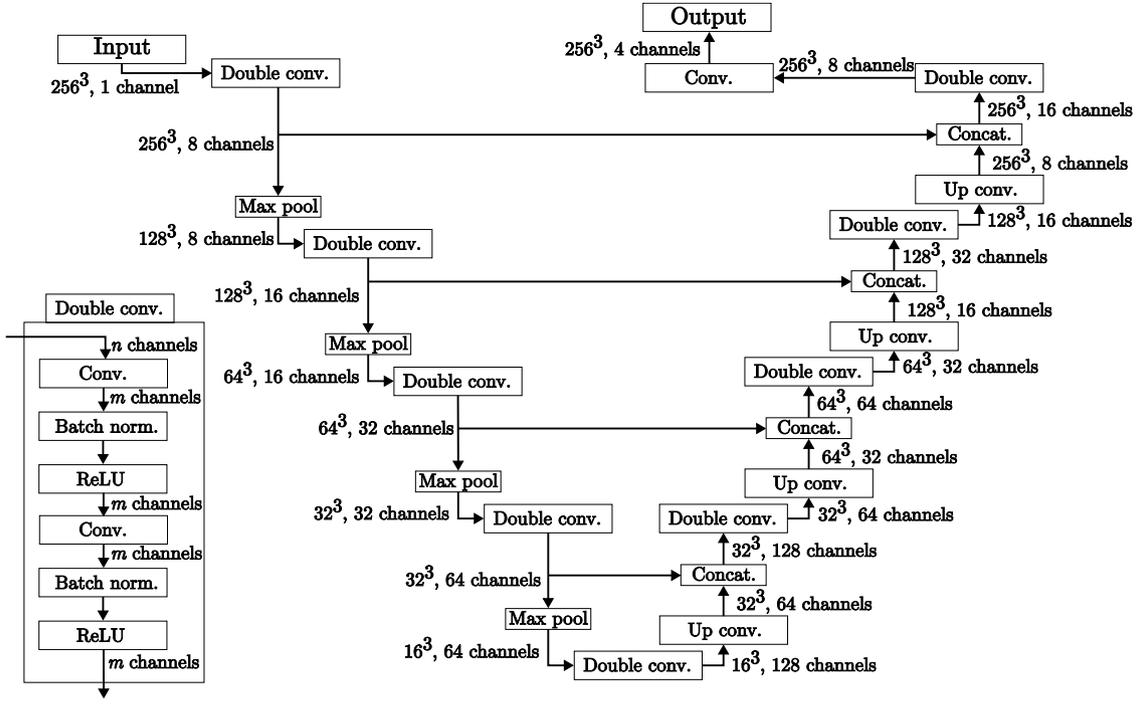


Fig. 2. The used 3D convolution U-net architecture is displayed. The input is a one channel (greyscale) XRCT image and the output with four channels represents the segmentation class probabilities. The structure of the *Double convolution* block is also shown.

1. Its purpose is to yield the desired number of output channels. All parameters not explicitly referred to use the PyTorch default values.

The training process itself requires further pre-processing of the training and validation data sets (splitting and normalisation), the definition of an appropriate loss function, and a selection of a set of important hyperparameters. The synthetic tomographic training volumes were 512 by 512 by 512 voxels in size. The volumes were divided into overlapping patches of 256 by 256 by 256 voxels, with a stride of 128. This means that each tomographic volume results in 27 data points. The input grey scale data is subject to z-score normalisation. This entail subtracting the mean and dividing by the standard deviation, where the statistics has been computed over the entire training set.

A loss function based on the Jaccard index (intersection over union) was used during training. Let  $A_{ijkl}$  be a four dimensional array representing a segmented image, which is equal to one if a voxel with coordinates  $i, j, k$  is of class  $l$  and zero otherwise. The loss for a batch of predicted segmentations  $A_{ijkl}^m$  and their corresponding arrays of the ground truths  $B_{ijkl}^m$  is defined as

$$J(A, B) = \frac{1}{n_b n_c} \sum_{m=1}^{n_b} \sum_{l=1}^{n_c} \left( 1 - \frac{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} A_{ijkl}^m B_{ijkl}^m + \epsilon}{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} (A_{ijkl}^m + B_{ijkl}^m - A_{ijkl}^m B_{ijkl}^m) + \epsilon} \right), \quad (1)$$

where  $n_c$ ,  $n_b$ ,  $n_x$ ,  $n_y$ , and  $n_z$  are the number of classes, the batch size, the image width, image height, and image depth respectively. The parameter  $\epsilon$  is here set to  $1 \times 10^{-6}$  to prevent division by zero for data where one or several of the classes are not present.

When training,  $A$  is obtained by passing the network output through a channel wise softmax function. This means that in practice the entries  $A_{ijkl}$  are not strictly 1 or 0. A training batch size of 2 is used for gradient update computations. The model is trained with the ADAM optimizer for 10 epochs. The validation loss is evaluated online, which means that the loss is computed from a batch of the evaluation set after each gradient update. The validation set is roughly 5 times smaller than the training set. This means that the validation set is repeated about 5 times during one training epoch. The model was trained with

a starting learning rate of 0.001. The learning rate is reduced linearly after each gradient update until it reaches 0 at the end of epoch 5. Thereafter, the learning rate is re-set to 0.001, whereafter it again is reduced linearly until it reaches 0 at the end of epoch 10. This “warm” restarting-type of learning rate scheduling has shown potential to improve convergence [45,46]. All other optimizer parameters are kept at the PyTorch default settings.

The model training and validation loss curves are displayed in Fig. 3. The loss does not show any signs of overfitting and reached stationary behaviour at the time of training termination. The model took approximately 73 h to train on the described laptop.

### 3.2. Segmentation results and discussion

Inference, i.e. segmentation of an entire sample (virtual or physical), took less than 20 s including all data input and output. To evaluate the segmentation performance of the model, three metrics were used. The *Jaccard index*<sup>2</sup> was computed on the testing set (composing synthetic but unseen data) using the same patch size and stride as for the Jaccard loss function during training and validation, but with a batch size of 1. The *Precision* metric  $P$ , which represents the number of true positives divided by the sum of the true positives and false negatives, was computed<sup>3</sup> as

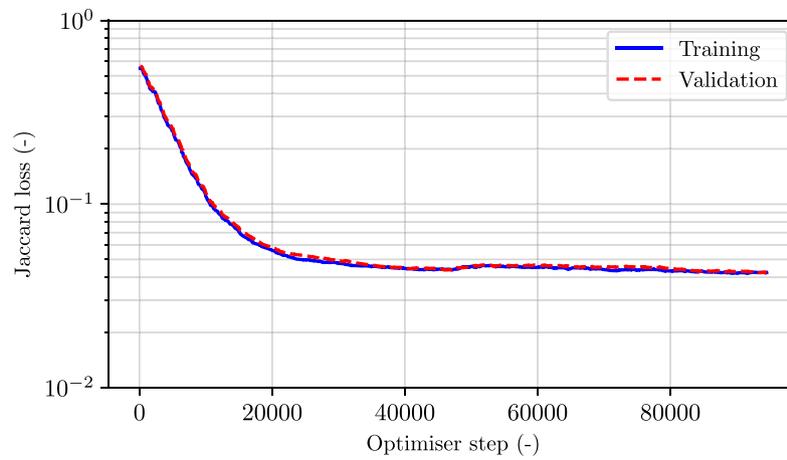
$$P(A, B) = \frac{1}{n_c} \sum_{l=1}^{n_c} \frac{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} A_{ijkl} B_{ijkl} + \epsilon}{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} A_{ijkl} (1 - B_{ijkl}) + \epsilon}, \quad (2)$$

whereas for the defined loss, the output of the network is passed through a softmax function. Finally, the *Recall* metric  $R$ , defined as the ratio between true positives and the sum of true positives and false negatives, was computed as

$$R(A, B) = \frac{1}{n_c} \sum_{l=1}^{n_c} \frac{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} A_{ijkl} B_{ijkl} + \epsilon}{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \sum_{k=1}^{n_z} (1 - A_{ijkl}) B_{ijkl} + \epsilon}. \quad (3)$$

<sup>2</sup> This is computed like in Eq. (1) removing ‘1-’ from the sum over  $n_c$ .

<sup>3</sup> using the same notation as in Section 3.1.



**Fig. 3.** A lin-log plot of the segmentation model loss is shown. Smoothing according to  $y[n+1] = 0.01x[n] + 0.99y[n]$  has been applied to the data, where  $y[n]$  is the smoothed data at step  $n$  and  $x[n]$  is the raw data.

**Table 1**

Evaluation metrics computed over the testing set. A higher number is better, where 1.0 is the maximum possible value.

Metric	Average	Minimum	Maximum
Jaccard index	0.958	0.897	0.977
Precision	0.978	0.948	0.988
Recall	0.978	0.942	0.988

The  $\epsilon$  parameter for all metrics uses the same value as in Section 3.1.

The average, the minimum, and the maximum of each metric are summarised in Table 1. It can be seen that the three metrics are very favourable. This indicates that the network has been sufficiently trained. The high minimums show that there is no overfitting that materialises through losing predictability on challenging outliers.

The XRCT scan of the physical sample was segmented<sup>4</sup> with the trained model, using the same patch size, stride and similar normalisation as during training. However, since the glass fibre tracer yarns caused some high contrast artifacts not present in the training data, a  $0.6 \text{ cm}^{-1}$  cut off value was used before normalisation (all voxels with an intensity value above 0.6 were set to 0.6). As overlapping patches were used to aid continuity and to prevent sharp borders between patches, the network outputs were summed up and the highest probability material class was selected when reassembling the patches. Refer to Fig. 4 for images of the XRCT data and the corresponding segmentations. The presented segmentations displays highly improved continuity of the weft yarns when compared with previously presented results in [33].

#### 4. Computational RVE homogenisation

Following the segmentation of the XRCT reconstructions the next step in the pipeline is the creation of FE models of the considered RVEs. These models are used to predict the as-manufactured macroscopic elastic properties of the material (with statistical variations) through computational homogenisation of multiple UCs. In this study, the commercial FE code LS-DYNA is used for the numerical analysis, with all FE simulations carried out on an HPC cluster. It is however reiterated, that all novel contributions in this work are open source, and that any reproduction of the study can exchange the particular finite element solver without altering the conclusions.

Generating the RVE models requires, in this case, three distinct steps. Each of these will be discussed in greater detail in the following sections. First, the segmented volume must be discretised into

<sup>4</sup> Before segmentation, the reconstructed volume was rotated such that the UC boundaries are aligned with the coordinate axes.

continuum elements, where each element is assigned one of four material classes (matrix, air, warp yarn, or weft yarn). Then, the general input file structure along with keywords for the specific FE-solver must be identified, and the appropriate periodic boundary conditions must be defined. Next, a micromechanical model must be selected to provide the material properties for the yarns. For this model, having the elastic properties of the carbon fibre and matrix as inputs (known from, e.g., technical data sheets), the local fibre volume fraction of the yarn must be approximated element-wise. Finally, the local orientation of the fibre reinforcements within the yarns must be predicted and mapped element-wise as well.

##### 4.1. Finite element model

As previously mentioned, the RVE FE analyses were performed using the commercial code LS-DYNA. More specifically, linear, static implicit analyses were considered. Due to the large number of elements in these models (see more details below), a direct solver is infeasible. Therefore, the iterative conjugate gradient solver with a local symmetric Gauss-Seidel pre-conditioner was used.

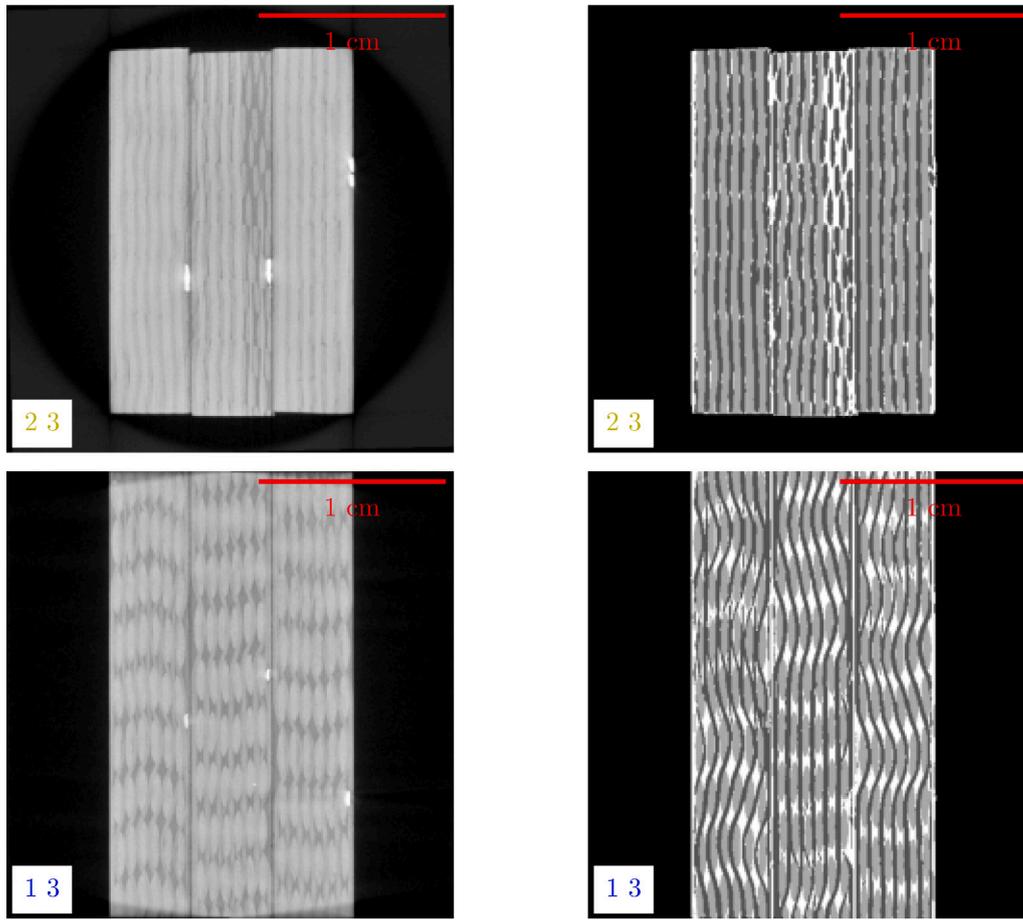
When meshing such RVEs, either non-conformal structured meshes or conformal meshes can be considered. Conformal meshes, which explicitly separate and follow the true geometry of the yarns and matrix pockets, have their benefits. However, they require specialised and intricate mesh generation tools. As an alternative non-conformal structured meshes are a favourable alternative, especially since the considered analysis in this work pertains strictly to the elastic regime. It has been demonstrated by Doitrand et al. [47] that the use of non-conformal hexahedral elements is sufficient as long as they are small enough to accurately represent the yarn fibre volume fraction.

One approach, and what is considered in this work, is to perform a one-to-one mapping from each voxel in the segmented RVE volume to an 8-node hexahedral element, with nodes positioned at the corners of the corresponding voxel. Specifically in this case, the RVEs are discretised with LS-DYNA's trilinear and under-integrated ELEMENT\_SOLID\_ORTHO elements (ELFORM = 1).

A benefit of the structured mesh is that the application of periodic boundary conditions becomes straightforward. The absolute formulation of the periodic boundary condition equations, as presented by Garoz et al. [48], was used. This formulation uses three corner nodes as so-called 'control nodes' and one fixed 'reference node', making it easy to impose the desired macroscopic infinitesimal strain tensor  $\epsilon$ . This is done by prescribing the displacement on the three control nodes (and fixing one node) as

$$\mathbf{u}_i = \epsilon \cdot \mathbf{X}_i \quad i \in \text{control node} \quad (\mathbf{u}_{\text{fixed}} = \mathbf{X}_{\text{fixed}} = 0), \quad (4)$$

where  $\mathbf{u}_i$  and  $\mathbf{X}_i$  are the node displacements and positions, respectively.



**Fig. 4.** Two perpendicular slices of the rotated XRCT reconstruction (left) is shown next to their corresponding segmentations (right). For the XRCT scan brighter values correspond to higher attenuation (note the very bright tracer yarns). For the segmentation, black refers to air, dark grey refers to weft yarns, light grey refers to warp yarns, and white refers to pure matrix.

By performing a uniaxial strain controlled load simulation for each of the six strain components, each column of the macroscopic Voigt stiffness matrix can be computed as the resulting macroscopic stress components scaled by the inverse of the applied control strain. In this context, the macroscopic stress tensor,  $\bar{\sigma}$ , is computed from the stress tensor in each integration point,  $\sigma$ , as

$$\bar{\sigma} = \frac{1}{V} \int_V \sigma dV, \quad (5)$$

where  $V$  is the total volume of the RVE. Note that as all elements are equally sized in the voxel mesh, and one integration point per element is used, it is possible to compute the macroscale stress tensor by a simple average over all elements.

The orthotropic linear elastic material model MAT\_ORTHOTROPIC\_ELASTIC is used for all elements. The pure matrix phase is simply assigned isotropic matrix material properties in the orthotropic linear elastic material model. In the case of the impregnated carbon fibre yarns, the transversely isotropic elastic properties are calculated using a micromechanical model that is discussed in greater detail in Section 4.2.

Finally, in this context, it should be noted that LS-DYNA computes out-put stresses in a finite deformation setting even if the linear momentum balance was solved for infinitesimal strains. To make this discrepancy negligible, a control strain of 0.01% was used for all load cases.

#### 4.2. Micro-mechanical model and constituent properties

Since the FE models resolve individual yarns, but not individual fibres, the yarns are modelled as homogeneous and transversely

isotropic. To predict the linear elastic, transversely isotropic properties of the yarns, the Chamis micro mechanical model is used [49]. The model was initially developed for predicting UD laminate properties, but has been successfully used for the yarns of textile composites by several authors, cf. Stig et al. [50], Pierreux et al. [51], and Koptelov et al. [52]. To be specific, the yarn properties ( $\bullet^y$ ) are obtained in terms of transversely isotropic fibre properties ( $\bullet^f$ ), isotropic matrix properties ( $\bullet^m$ ) and the (local) volume fraction of fibres ( $V^f$ ) in each yarn as

$$\begin{aligned} E_{xx}^y &= V^f E_{xx}^f + (1 - V^f) E_m & v_{xy}^y &= V^f v_{xy}^f + (1 - V^f) v^m = v_{xz}^y \\ E_{yy}^y &= \frac{E^m}{1 - \sqrt{V^f} \left( 1 - E^m / E_{yy}^f \right)} = E_{zz}^y \\ G_{xy}^y &= \frac{G^m}{1 - \sqrt{V^f} \left( 1 - G^m / G_{xy}^f \right)} = G_{xz}^y \\ G_{yz}^y &= \frac{G^m}{1 - \sqrt{V^f} \left( 1 - G^m / G_{yz}^f \right)} & v_{yz}^y &= \frac{E_{yy}^y}{2G_{yz}^y} - 1 \end{aligned} \quad (6)$$

Here  $x$  denotes the direction parallel to the yarn (or fibre), while  $y$  and  $z$  denotes the two transverse directions.

The calculation of the yarn fibre volume fraction is discussed in more detail in Section 4.3. The constituent mechanical properties used in this work are given in Table 2. For the fibres, the longitudinal modulus is taken from the manufacturer's data sheet [37], and the remaining properties are taken from the third world-wide failure exercise [53]. For the matrix, the (tensile) modulus of RTM6 is also taken

**Table 2**  
The material constituent mechanical properties.

Property	Value	Reference
$E_{xxf}$	276 GPa	[37,53]
$E_{yxf}$	19 GPa	[53]
$G_{xyf}$	27 GPa	[53]
$G_{yzf}$	7 GPa	[53]
$\nu_{xyf}$	0.2	[53]
$E_m$	3.1 GPa	[38,54]
$\nu_m$	0.36	[54]

from the manufacturer's data sheet [38], which aligns well to what was (independently) measured by Exner et al. [54]. Further, Exner et al. [54] also reported the Poisson's ratio of the RTM6, which was used in this work.

#### 4.3. Fibre orientation mapping

As the next step, the local yarn orientations need to be estimated and assigned to each element. To determine the local fibre orientation it is assumed that the fibres can be estimated to have no in-plane misalignment. Based on this assumption, and speaking broadly, the local fibre orientation is inferred with a structure tensor analysis of the segmented RVE, carried out in 2D on a per-slice basis. The Python package `structure-tensor` [55] was used for this analysis. For both yarn types respectively, the volume was sliced along the axis that is perpendicular to both the yarn direction and the out-of-plane direction. An example is visualised in Fig. 5.

The binary nature of the segmentation yields little gradient information inside the yarns. This makes the structure tensor difficult to apply directly. Therefore, the Euclidean distance transform was computed for each slice, in which each segmented yarn voxel in the slice is assigned its distance to the closest voxel of any other class in the same slice. This creates ridge-shaped structures that follow the yarns, creating a large gradient perpendicular to the yarns and a near zero gradient along the yarns, which improves the accuracy of the structure tensor analysis, cf. Fig. 5a.

Furthermore, numerical issues arise when performing structure tensor analysis of large regions of constant value. As a remedy, uniform random noise between 0 and 0.1 (the higher value corresponding to a 10th of a voxel) was added to the distance value of each pixel. Thereafter, the image was smoothed with a Gaussian filter with kernel size 1. The noised, smoothed, distance transformed slice, cf. Fig. 5c, was analysed with a  $\rho$  parameter (integration kernel size) of 2.0 and a  $\sigma$  (gradient kernel size) parameter of 0.25.

Finally, the eigenvectors resulting from the structure tensor analysis were examined. In the structure tensor analysis the eigenvector associated with the smallest eigenvalue of the structure tensor represents the direction of minimal variation in the underlying image [56], in this case corresponding to the local orientation of the yarns. As an example, considering an image with stripes. For this case, the gradient is large when going from a stripe to background, while small when following the stripes.

To visualise the outcome of the structure tensor analysis, the weft yarn orientations of one UC are shown in Fig. 6. Analysing the direction of eigenvectors, it can be found that the level of crimp (the magnitude of the out-of-plane component of the fibre orientation vectors) in the weft is very consistent, even though it was not strictly enforced. The average weft yarn unit orientation vector out-of-plane component magnitude was determined to be 0.135 with a standard deviation of only 0.003. In addition, the warp yarns display very low crimp in all the identified UCs and are not visualised. The average warp out-of-plane component magnitude was found to be 0.047 with a standard deviation of 0.003.

#### 4.4. Fibre volume fraction mapping

For the purpose of computing the local fibre volume fraction of the yarns, it should be noted that in the case of the warp yarns, the segmentation is able to clearly differentiate between one yarn instance and the next throughout the volume. This, in simple terms, would make the calculation of the local yarn fibre volume fraction straightforward given the total area of the yarn, the number of filaments in the yarn and the diameter of each filament. One challenging aspect to this material, however, is that in some slices the weft yarns are highly compacted and appear to merge with one another, cf. the side pieces in the 2–3 plane in Fig. 4. This makes it difficult to differentiate one weft yarn instance from the next. Therefore, to estimate the local fibre volume fraction in each element of the yarns the a priori known number of yarns and the number of fibres per yarn was exploited.

It is assumed that, for each yarn type, the fibre volume fraction can be approximated as constant per cross section of the segmented volume along each respective yarn weave direction. As such, the volume fraction of each yarn type is estimated by slicing the segmented UC along the axis parallel to the respective yarn direction. Refer to Fig. 7 for a representative schematic. A first estimate of the slice-wise average volume fraction is given by the ratio between the theoretical total area of the fibres in a yarn type in a cross section and the area covered by voxels segmented as the same yarn type in the same cross section. The theoretical area covered by the fibres in a cross section is computed as

$$A_{\text{fibre}} = N_{\text{fibre}} N_{\text{yarn}} \frac{\pi d_{\text{fibre}}^2}{4}, \quad (7)$$

where  $d_{\text{fibre}}$  is the fibre diameter, and  $N_{\text{fibre}}$  and  $N_{\text{yarn}}$  are the number of fibres per yarn and number of yarns per cross section, respectively (which depend on the yarn type). This first estimate is accurate if the yarns are perpendicular to the cross section in question. In reality, however, the yarns are not straight. This was compensated for by scaling the area of each voxel with the unit orientation vector component in the yarn direction. That is, voxels where yarns are perpendicular to the cross section received a scale factor of 1, while other voxels representing up-wards or down-wards crimping yarn segments obtained a scale factor less than 1. In summary, the fibre volume fraction was thus estimated slice-wise by computing the ratio between the theoretical area covered by the fibres and the orientation-compensated area of the segmented voxels.

Finally, it is noted that, to avoid generating too many material cards, the volume fractions are binned into 20 bins that are equally spaced between the minimum yarn fibre volume fraction and the maximum yarn fibre volume fraction. For illustrative purposes, the fibre volume fraction inside both the weft and warp yarns is displayed in Fig. 8 for the same UC shown previously. As expected, the fibre volume fraction of the weft yarns is highest as it is squeezed between the rows of warp yarns. Further, it should be noted that the yarn fibre volume fractions vary between 64%–78%.

#### 4.5. Results, discussion and comparison to experimental results

From the segmented data, six non-overlapping UCs were identified by hand. The UCs were selected where the segmentation was qualitatively the most accurate. It is here emphasised that this was the only manual labour performed in the entire pipeline. After the initial selection, the described segmentation, meshing and material mapping procedures were performed on all six UCs, respectively. Separate FE analyses of each of the identified UCs were then performed to estimate the elastic properties and their statistics. As previously discussed, this involves the application of six distinct load cases using periodic boundary conditions, each of which represents a non-zero strain component. By carrying out first order computational homogenisation the entire stiffness tensor of the homogenised material sample can be computed.

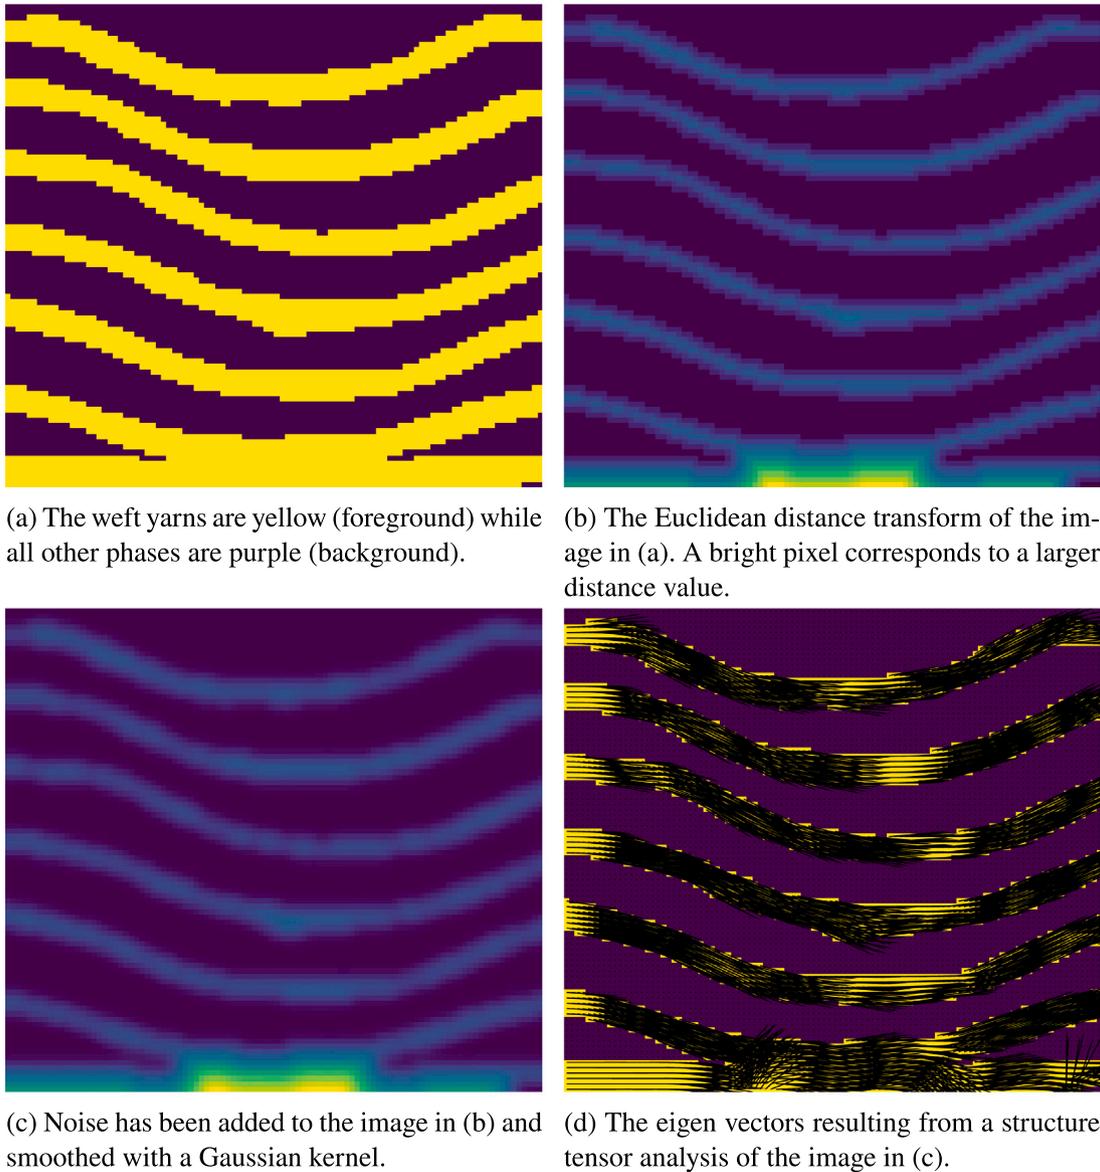


Fig. 5. The orientation analysis of a weft yarn is demonstrated. A small section of a slice along the warp direction is shown.

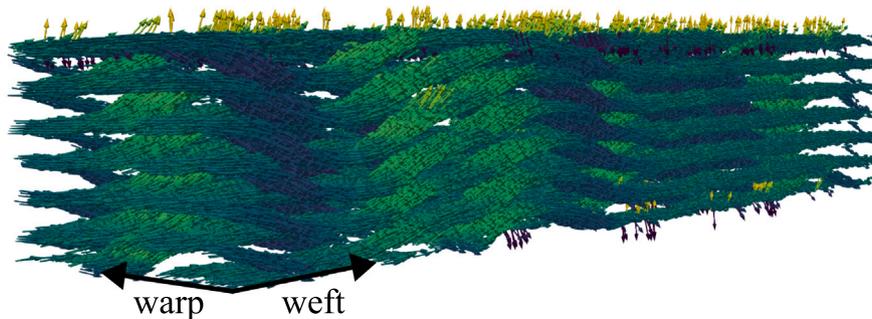


Fig. 6. The weft yarn orientation vectors in one of the identified UCs are visualised. The colouring aids visualisation contrast, where bright colours represent a positive out-of-plane component and dark represents negative.

For each of the 6 UCs, 3 tensile (along weft, along warp, and out-of-plane), and 3 pure shear (in-plane, weft out-of-plane, and warp out-of-plane) tests were performed. For each load case, the macroscopic stress scaled by the macroscopic strain resulted in one column of the

Voigt stiffness matrix derived for the corresponding UC. Fig. 9 shows the normal stress in the weft direction, for the load case of pure elongation in the weft direction. The figure concerns the UC (no. 1) that is displayed in Figs. 6 and 8. The values of the homogenised stiffness

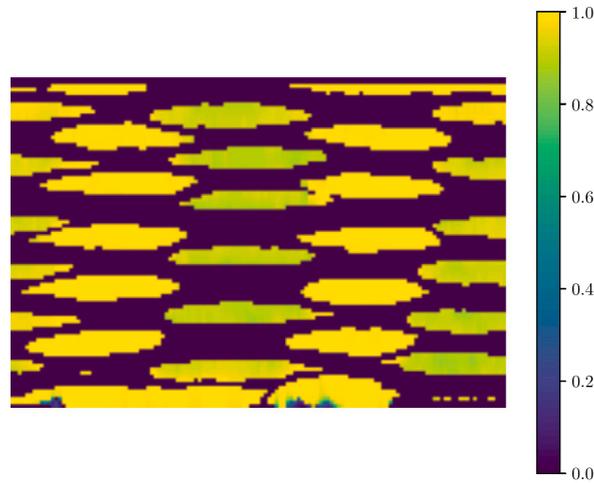


Fig. 7. The weft cross sections are shown in a slice along the weft direction. The colour bar refers to the voxel area scale factor due to the fibre direction.

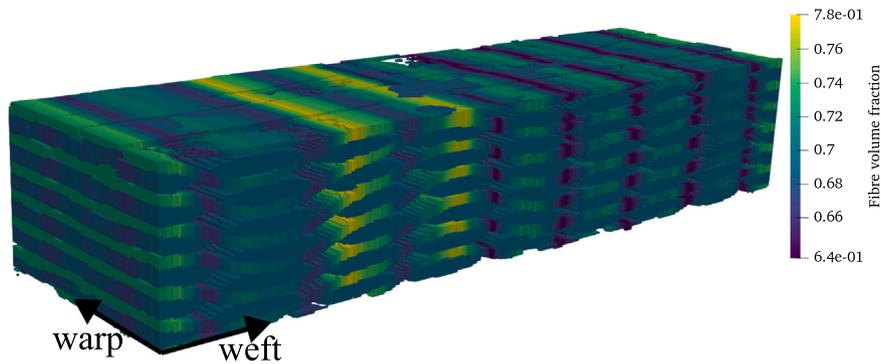


Fig. 8. The fibre volume fractions of one of the identified UCs are displayed.

Table 3

The computed homogenised mean and standard deviation of the macroscopic engineering tensile and shear moduli for the six UCs are shown together with the measured in-plane experimental values. The values are given in GPa.

Parameter	$E_{11}$	$E_{22}$	$E_{33}$	$G_{12}$	$G_{23}$	$G_{13}$
Experiment [6]	68.65	61.43		3.91		
Sim. mean	71.92	57.75	9.10	4.11	3.60	3.28
Sim. std dev	0.62	1.11	0.04	0.08	0.08	0.06

tensor of the same UC are given as

$$C_1 = \begin{bmatrix} 72.9544 & 5.1803 & 4.1161 & 0.0155 & 0.0014 & -0.3598 \\ 5.1798 & 60.1845 & 4.9221 & -0.0947 & -0.0275 & -0.0392 \\ 4.1161 & 4.9221 & 9.6476 & 0.0023 & 0.0042 & -0.0158 \\ 0.0046 & -0.0228 & 0.0007 & 4.1954 & -0.0066 & -0.0008 \\ 0.0004 & -0.0066 & 0.0012 & -0.0066 & 3.5235 & -0.0006 \\ -0.0896 & -0.0097 & -0.0038 & -0.0008 & -0.0006 & 3.2249 \end{bmatrix} \text{GPa}, \quad (8)$$

where  $C_1$  is the Voigt matrix representation thereof. It is evident that the resulting stiffness tensor demonstrates slight asymmetry. However, this is not unexpected, and given the relative magnitude of stiffness components, it has still been deemed acceptable. In addition, the lower left and upper right quadrants of the Voigt matrix representation of the stiffness tensor are not exactly zero as required by orthotropy. It is however deemed that, since the values in these quadrants are orders of magnitude smaller than the true non-zero components, the assumption of elastic orthotropy is still reasonable.

Table 4

The computed mean and standard deviation of the homogenised macroscopic Poisson's ratios for the six UCs.

Parameter	$\nu_{12}$	$\nu_{21}$	$\nu_{23}$	$\nu_{32}$	$\nu_{13}$	$\nu_{31}$
Mean	0.0509	0.0409	0.4922	0.0776	0.3954	0.0500
Std dev	0.0031	0.0034	0.0034	0.0019	0.0032	0.0005

To compensate for the minor asymmetry of the stiffness tensor, the extraction of engineering macroscopic properties was performed from the symmetrised Voigt stiffness matrix given by

$$C_{\text{sym}} = \frac{1}{2} (C + C^T). \quad (9)$$

Thereafter, the nine macroscopic elastic engineering parameters of orthotropy (using the coordinate convention in Fig. 1(a)) are easily identified from the compliance Voigt matrix through

$$S_{\text{sym}} = C_{\text{sym}}^{-1} = \begin{bmatrix} \frac{1}{E_{11}} & -\frac{\nu_{21}}{E_{22}} & -\frac{\nu_{31}}{E_{33}} & 0 & 0 & 0 \\ -\frac{\nu_{12}}{E_{11}} & \frac{1}{E_{22}} & -\frac{\nu_{32}}{E_{33}} & 0 & 0 & 0 \\ -\frac{\nu_{13}}{E_{11}} & -\frac{\nu_{23}}{E_{22}} & \frac{1}{E_{33}} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{G_{12}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{G_{23}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{G_{13}} \end{bmatrix}. \quad (10)$$

The macroscopic elastic parameters, computed on all of the analysed UCs, are displayed in Fig. 10. The average values of the tensile and shear moduli, and their standard deviations, are shown in Table 3. For comparison, also experimentally obtained tensile and shear moduli

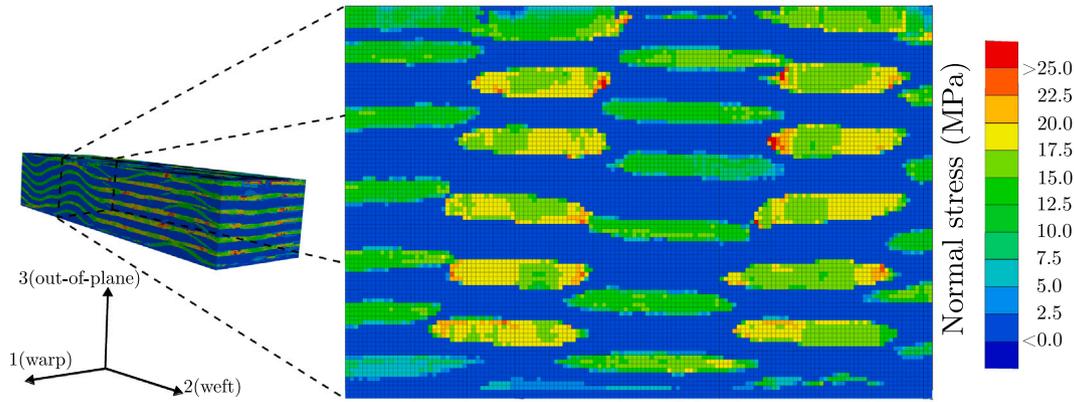


Fig. 9. The normal stress in the weft (2) direction is shown for the pure elongation in the weft direction load case. Note the one-to-one correspondence between elements and segmentation voxels. The figure illustrates how the weft yarns carry the stress.

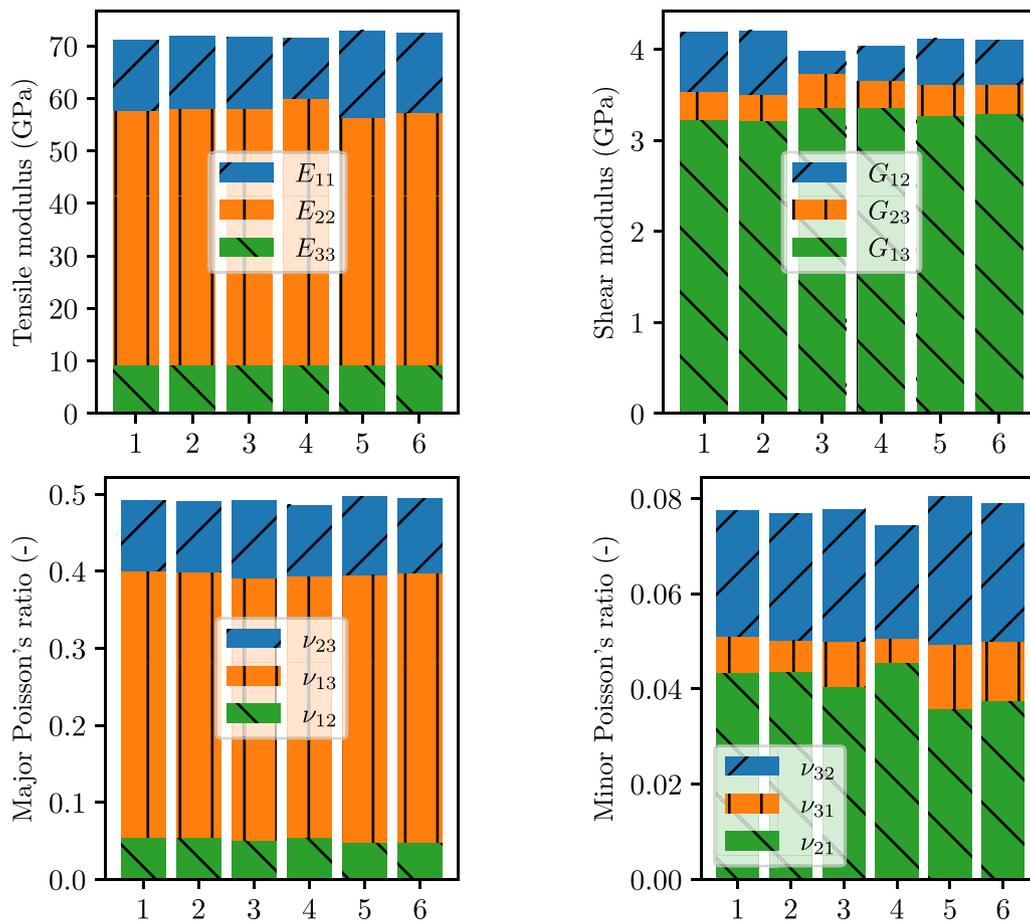


Fig. 10. The macroscopic elastic properties for each UC are displayed.

are presented in the same table. These moduli have been extracted for the exact same material from the cyclic test data obtained by Oddy et al. [6]. More specifically, each separate modulus has been determined via a linear least square fit of the stress–strain data of the earliest possible load cycle including the full strain range  $\Delta\epsilon = \epsilon_{max} - \epsilon_{min} = 0.003 - 0.001^5$  (as stipulated by *ASTM D3039* [57]). Furthermore,

<sup>5</sup> The test data in [6] was obtained through cyclic loading with a continuously increasing strain range for each cycle. The 'earliest possible load cycle' is thereby the first load cycle that contains the whole strain range

the average Poisson's ratios and their standard deviations are shown in Table 4.

From these results, it is observed that the predicted average of the tensile modulus in the warp direction differs with less than 4.7% compared to the experimental value. Similarly, the predicted average tensile modulus in the weft direction and the predicted average shear modulus show deviations of less than 6.4% and around 5%, respectively.

The small deviations from experimental values demonstrate the capability of the proposed pipeline in predicting the as-manufactured elastic properties of a 3D-woven CFRP. Moreover, owing to the fact

that several UCs can be investigated from only one XRCT scan, the variability in a material can be easily investigated.

It was surprising that the amount of out-of-plane fibre misalignment was as consistent as observed. At the same time, this explains the consistency in the predicted elastic properties. Still, as seen in the bottom left part of Fig. 4, there is a noticeable difference in warp yarn configuration between the three stacked pieces. The two pieces on each side have their warp yarns arranged in a U-shape, while the centre piece displays more straightly stacked warp yarns. It is likely that these morphological differences would have a larger impact if deformations increase (requiring finite strain simulations). As Oddy et al. [6] reported from their experiments, the material under study displayed yarn realignment for large strains. It is, however, evident that to extract linearly elastic properties (relevant for small strains), as long as the fibre volume fraction and orientation distribution is captured accurately enough, the analysis of a single UC is enough. This indicates that the samples under study have been manufactured in a fairly consistent manner.

## 5. Conclusion and future outlook

In this work, a fully automated and open source pipeline for the prediction of the elastic properties of 3D-woven composites in their as-manufactured state has been presented. The pipeline allows for automated segmentation of challenging large field of view scans, automated meshing and assignment of material orientations and properties as well as automated definitions of periodic boundary conditions for computational homogenisation. It includes a novel method for estimating the local yarn orientation in cases where instance segmentation is challenging or even impossible. In more detail, the 3D U-net architecture showed great promise in segmenting large field of view XRCT scans of CFRP composites.

Applying this pipeline to a resin-infused layer-to-layer angle interlock weave, it has been demonstrated that macroscopic elastic properties can be predicted in good agreement with experimental observations. Highly accurate predictions (less than 6.5% deviation from experiments) of the full 3D elastic properties were achieved. This indicates that the pipeline can be an efficient tool for predicting the elastic properties directly from large field of view XRCT scans of 3D-woven composites.

The execution of the pipeline is computationally efficient. When the segmentation model has been pre-trained, the time from a reconstructed volume to an FE input file is less than a minute and fully automated. Being automated and efficient, the pipeline allows for several UCs to be generated from a single scan, which unlocks the possibility to estimate statistical information of the investigated material properties. Interestingly, however, for the material under study, it was found that the elastic properties had very little statistical variation between UCs, even though the meso scale geometry had noticeable differences. Future works should investigate how large the variation in the meso scale geometry needs to be to yield significant variation in elastic properties.

Having validated the ability of the pipeline to accurately predict the elastic properties of as-manufactured samples, a relevant follow-up is the investigation of large deformations, plasticity and damage. In order to expand the analysis to these domains, the use of a voxel mesh becomes infeasible due to the very large number of elements. Furthermore, a voxel mesh may yield artificial stress concentrations [58] that can erroneously initiate plasticity and damage [59]. As a consequence, future works should incorporate a conformal meshing approach. A meso scale model with conformal meshing of the yarns would however require instance segmentation of individual yarns. This has been done by Blusseau et al. [30], Mendoza et al. [29], and Fourier et al. [22] among others. Future work could take inspiration from these approaches and generalise to allow for segmenting very low resolution, large field of

view scans studied in this work, and to eliminate remaining steps where manual input is necessary.

It is a challenging undertaking to generalise conformal meshing procedures to many large UCs. A possible alternative could be to encode not only phase ground truths in the training data, but also encode the yarn orientations and phase boundaries. Additionally, instead of predicting the segmentation as an intermediate step it would be interesting to investigate inferring geometry descriptions directly in the neural network model. This would require finding an encoding for the geometry information that can handle arbitrary number of yarns, and arbitrary yarn shapes.

## CRedit authorship contribution statement

**Johan Friemann:** Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Carolyn Oddy:** Writing – original draft, Supervision. **Lars P. Mikkelsen:** Writing – review & editing, Supervision, Resources. **Martin Fagerström:** Writing – review & editing, Supervision, Resources, Funding acquisition, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

The authors acknowledge funding from Horizon Europe through the MSCA Doctoral Network RELIANCE: Reliance: REal-time characterisation of ANisotropic Carbon-based tEchnological fibres, films and composites, grant no. 101073040. The X-ray scans were captured at DTU Center for Advanced Structural and Material Testing (CASMaT) located at the 3D imaging Centre (3DIM). The FE simulations were performed on resources at the Chalmers Centre for Computational Science and Engineering (C3SE) provided by the Swedish National Infrastructure for Computing (SNIC). The authors want to thank Dr. Lei Liu for help with the formulation of periodic boundary conditions.

## Data availability

All data not explicitly referred to in the manuscript will be made available on request.

## References

- [1] K. Jing, S. Xie, Y. Zhang, H. Zhou, H. Yan, Impact resistance of 3D woven fabrics and composites: A review, *Thin-Walled Struct.* 213 (2025) 113262.
- [2] M.N. Saleh, C. Soutis, Recent advancements in mechanical characterisation of 3D woven composites, *Mech. Adv. Mater. Mod. Process.* 3 (1) (2017).
- [3] P. Tan, L. Tong, G. Steven, T. Ishikawa, Behavior of 3D orthogonal woven CFRP composites. Part I. Experimental investigation, *Compos. Part A: Appl. Sci. Manuf.* 31 (3) (2000) 259–271.
- [4] Y. Mahadik, S. Hallett, Effect of fabric compaction and yarn waviness on 3D woven composite compressive properties, *Compos. Part A: Appl. Sci. Manuf.* 42 (11) (2011) 1592–1600.
- [5] K. Leong, B. Lee, I. Hertzberg, M. Bannister, The effect of binder path on the tensile properties and failure of multilayer woven CFRP composites, *Compos. Sci. Technol.* 60 (1) (2000) 149–156.
- [6] C. Oddy, M.y. Song, C. Stewart, B. El Said, M. Ekh, S.R. Hallett, M. Fagerström, Cyclic behaviour of 3D-woven composites in tension: Experimental testing and macroscale modelling, *Compos. Part A: Appl. Sci. Manuf.* 187 (2024) 108354.
- [7] L.P. Brown, A.C. Long, Modeling the geometry of textile reinforcements for composites: TexGen, in: P. Boisse (Ed.), *Composite Reinforcements for Optimum Performance*, second ed., in: Woodhead Publishing Series in Composites Science and Engineering, Woodhead Publishing, 2020, pp. 237–265.
- [8] S.V. Lomov, Modeling the geometry of textile composite reinforcements: Wise-Text, *Composite Reinforcements for Optimum Performance*, second ed., Elsevier Ltd., 2020, pp. 199–236.

- [9] Y. Wang, X. Sun, Digital-element simulation of textile processes, *Compos. Sci. Technol.* 61 (2) (2001) 311–319.
- [10] G. Zhou, X. Sun, Y. Wang, Multi-chain digital element analysis in textile mechanics, *Compos. Sci. Technol.* 64 (2) (2004) 239–244.
- [11] B.S. El Said, S. Green, S.R. Hallett, Kinematic modelling of 3D woven fabric deformation for structural scale features, *Compos. Part A: Appl. Sci. Manuf.* 57 (2014) 95–107.
- [12] D. Durville, I. Baydoun, H. Moustakas, G. Périé, Y. Wielhorski, Determining the initial configuration and characterizing the mechanical properties of 3D angle-interlock fabrics using finite element simulation, *Int. J. Solids Struct.* 154 (2018) 97–103.
- [13] Z. Yang, Y. Jiao, J. Xie, L. Chen, W. Jiao, X. Li, M. Zhu, Modeling of 3D woven fibre structures by numerical simulation of the weaving process, *Compos. Sci. Technol.* 206 (2021) 108679.
- [14] T. Shao, S. Zhang, W. Liu, R. Liu, W. Xu, High-fidelity architecture modeling and compressive strength prediction of 3D woven composite material, *Compos. Struct.* 354 (2025) 118775.
- [15] Z. Gao, L. Chen, S. Zhao, Z. Zhang, A deformation procedure using position-based dynamics to optimize the geometric model of woven fabrics, *Appl. Math. Model.* 137 (2025) 115650.
- [16] A. Ewert, B. Drach, K. Vasylevskiy, I. Tsukrov, Predicting the overall response of an orthogonal 3D woven composite using simulated and tomography-derived geometry, *Compos. Struct.* 243 (March) (2020) 112169.
- [17] S.U. Lee, S. Chung, R.H. Park, A comparative performance study of several global thresholding techniques for segmentation, *Comput. Vis. Graph. Image Process.* 52 (2) (1990) 171–190.
- [18] J.B. Roerdink, A. Meijster, The watershed transform: Definitions, algorithms and parallelization strategies, *Fund. Inform.* 41 (1, 2) (2000) 187–228.
- [19] R.M. Auenhammer, L.P. Mikkelsen, L.E. Asp, B.J. Blinzler, Automated X-ray computer tomography segmentation method for finite element analysis of non-crimp fabric reinforced composites, *Compos. Struct.* 256 (August 2020) (2021) 113136.
- [20] N. Naouar, E. Vidal-Sallé, J. Schneider, É. Maire, P. Boisse, Meso-scale FE analyses of textile composite reinforcement deformation based on X-ray computed tomography, *Compos. Struct.* 116 (1) (2014) 165–176.
- [21] R.M. Auenhammer, N. Jeppesen, L.P. Mikkelsen, V.A. Dahl, B.J. Blinzler, L.E. Asp, Robust numerical analysis of fibrous composites from X-ray computed tomography image data enabling low resolutions, *Compos. Sci. Technol.* 224 (2022) 109458.
- [22] G. Fourrier, A. Rassineux, F.H. Leroy, M. Hirsekorn, C. Fagiano, E. Baranger, Automated conformal mesh generation chain for woven composites based on CT-scan images with low contrasts, *Compos. Struct.* 308 (August 2022) (2023) 116673.
- [23] J. Bénézech, G. Couégnat, Variational segmentation of textile composite preforms from X-ray computed tomography, *Compos. Struct.* 230 (August) (2019) 111496.
- [24] Y. Sinchuk, P. Kibleur, J. Aelterman, M.N. Boone, W. Van Paepegem, Variational and deep learning segmentation of very-low-contrast X-ray computed tomography images of carbon/epoxy woven composites, *Materials* 13 (4) (2020).
- [25] V. Pidou-Brion, Y. Le Guilloux, Active yarn meshes for segmentation on X-ray computed tomography of textile composite materials at the mesoscopic scale, *Compos. Struct.* 281 (December 2021) (2022).
- [26] Y. Pannier, P. Coupé, T. Garrigues, M. Gueguen, P. Carré, Automatic segmentation and fibre orientation estimation from low resolution X-ray computed tomography images of 3D woven composites, *Compos. Struct.* 318 (April) (2023) 117087.
- [27] C. Tang, J. Zou, Y. Xiong, B. Liang, 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, *Compos. Sci. Technol.* 252 (April) (2024) 110616.
- [28] Z. Ying, H. Chen, X. Pan, Z. Wu, L. Peng, X. Cheng, J. Ou, Geometric reconstruction of mesoscopic models for three-dimensional woven composites using micro-CT technology, *Compos. Struct.* 371 (2025) 119466.
- [29] A. Mendoza, R. Trullo, Y. Wielhorski, Descriptive modeling of textiles using FE simulations and deep learning, *Compos. Sci. Technol.* 213 (2021) 108897.
- [30] S. Blusseau, Y. Wielhorski, Z. Haddad, S. Velasco-Forero, Instance segmentation of 3D woven fabric from tomography images by Deep Learning and morphological pseudo-labeling, *Compos. Part B: Eng.* 247 (March) (2022).
- [31] K. Zheng, C. Wu, H. Chen, X. Zhang, Z. Wang, Z. Pan, B. Qiu, Z. Wu, Improved 3D image segmentation for X-ray tomographic data of biaxial warp-knitted composites, *J. Reinf. Plast. Compos.* 44 (7–8) (2023) 344–357.
- [32] A. Tsamos, S. Evsevlev, F. Fiorelli, F. Faglion, G. Bruno, Synthetic Data Generation for Automatic Segmentation of X-ray Computed Tomography Reconstructions of Complex Microstructures, *J. Imaging* 9 (2) (2023).
- [33] J. Friemann, L.P. Mikkelsen, C. Oddy, M. Fagerström, Synthetic, automatically labelled training data for machine learning based X-ray CT image segmentation: Application to 3D-textile carbon fibre reinforced composites, *Compos. Part B: Eng.* 305 (2025) 112656.
- [34] J. Friemann, Textomos [software], 2025, URL: <https://github.com/Johan-Friemann/textomos>.
- [35] K. Jing, H. Zhou, H. Wang, H. Yan, S. Xie, Prediction and adjustment of elastic properties in three-dimensional orthogonal woven composites through dual-scale modeling, *Polym. Compos.* 45 (6) (2024) 5430–5448.
- [36] A. Ahmed, I. Ud Din, R. Ramachandran, M. Sikandar Bathusha, W. Cantwell, K.A. Khan, Micro-CT driven characterization and validation of constituents' properties used in origami-inspired foldable composites fabricated via different manufacturing processes, *Compos. Struct.* 322 (2023) 117382.
- [37] Hexcell, HexTow<sup>®</sup> IM7 Carbon Fiber, 2025, pp. 1–3, Product data sheet, [https://www.hexcel.com/wp-content/uploads/2025/12/IM7\\_HexTow\\_DataSheet.pdf](https://www.hexcel.com/wp-content/uploads/2025/12/IM7_HexTow_DataSheet.pdf).
- [38] Hexcell, HexFlow<sup>®</sup> RTM6/RTM6-2 180 °C Epoxy System for Resin Transfer Molding and Infusion Technologies, 2023, pp. 1–3, Product data sheet, [https://www.hexcel.com/wp-content/uploads/2025/12/RTM6\\_RTM62\\_HexFlow\\_DataSheet.pdf](https://www.hexcel.com/wp-content/uploads/2025/12/RTM6_RTM62_HexFlow_DataSheet.pdf).
- [39] L.A. Feldkamp, L.C. Davis, J.W. Kress, Practical cone-beam algorithm, *J. Opt. Soc. Amer. A* 1 (6) (1984) 612–619.
- [40] W. van Aarle, W.J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabrovolski, J. De Beenhouwer, K.J. Batenburg, J. Sijbers, Fast and flexible X-ray tomography using the ASTRA toolbox, *Opt. Express* 24 (22) (2016) 25129–25147.
- [41] J. Friemann, X-ray CT scan of layer to layer angle interlock composite sample with hand segmentation [dataset], 2025, URL: <https://zenodo.org/records/14891845>.
- [42] Y. Sinchuk, Y. Wielhorski, A. Mendoza, S. Blusseau, S. Velasco-Forero, Automatic yarn path extraction of large 3D interlock woven fabrics with confidence estimation, *Compos. Part A: Appl. Sci. Manuf.* 186 (2024) 108396.
- [43] Y. Cao, Z. Nie, F. Sun, Q. Zhao, X. Yang, C. Li, FabricFlow: Automatic segmentation of 3D closely packed woven fabric in low-contrast CT images through gradient flow tracking, *Mater. Today Commun.* 48 (2025) 113294.
- [44] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in: *Advances in Neural Information Processing Systems* 32, Curran Associates, Inc., 2019, pp. 8024–8035, URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [45] I. Loshchilov, F. Hutter, SGDR: Stochastic gradient descent with warm restarts, 2016, URL: <https://arxiv.org/abs/1608.03983>.
- [46] Y. Liu, Y. Ge, R. Pan, A. Kang, T. Zhang, Theoretical analysis on how learning rate warmup accelerates convergence, 2025, URL: <https://arxiv.org/abs/2509.07972>.
- [47] A. Doitrand, C. Fagiano, F.-X. Irisarri, M. Hirsekorn, Comparison between voxel and consistent meso-scale models of woven composites, *Compos. Part A: Appl. Sci. Manuf.* 73 (2015) 143–154.
- [48] D. Garoz, F. Gilabert, R. Sevenois, S. Spronk, W. Van Paepegem, Consistent application of periodic boundary conditions in implicit and explicit finite element simulations of damage in composites, *Compos. Part B: Eng.* 168 (2019) 254–266.
- [49] C.C. Chamis, Mechanics of composite materials: Past, present, and future, *J. Compos. Technol. Res.* 11 (1) (1989) 3.
- [50] F. Stig, S. Hallström, A modelling framework for composites containing 3D reinforcement, *Compos. Struct.* 94 (9) (2012) 2895–2901.
- [51] G. Pierreux, D. Van Hemelrijck, T.J. Massart, Automated generation of 3D orthogonal woven composites RVEs including yarn cross-section variations, *Compos. Sci. Technol.* 176 (2019) 90–102.
- [52] A. Koptelov, A. Thompson, S.R. Hallett, B. El Said, A deep learning approach for predicting the architecture of 3D textile fabrics, *Mater. Des.* 239 (2024) 112803.
- [53] A. Kaddour, M. Hinton, P. Smith, S. Li, Mechanical properties and details of composite laminates for the test cases used in the third world-wide failure exercise, *J. Compos. Mater.* 47 (20–21) (2013) 2427–2442.
- [54] W. Exner, R. Hein, T. Mahrholz, P. Wierach, H.P. Monner, M. Sinapius, Impact of nanoparticles on the process-induced distortions of carbon fiber reinforced plastics: An experimental and simulative approach, *J. Appl. Polym. Sci.* 136 (5) (2018).
- [55] N. Jeppesen, L. Mikkelsen, A. Dahl, A. Christensen, V. Dahl, Quantifying effects of manufacturing methods on fiber orientation in unidirectional composites using structure tensor analysis, *Compos. Part A: Appl. Sci. Manuf.* 149 (2021) 106541.
- [56] I. Straumit, S.V. Lomov, M. Wevers, Quantification of the internal structure and automatic generation of voxel models of textile composites from X-ray computed tomography data, *Compos. Part A: Appl. Sci. Manuf.* 69 (2015) 150–158.
- [57] A. International, ASTM D3039 Standard Test Method for Tensile Properties of Polymer Matrix Composite Materials, 2025, [http://dx.doi.org/10.1520/d3039\\_d3039m-08](http://dx.doi.org/10.1520/d3039_d3039m-08).
- [58] G. Fang, B. El Said, D. Ivanov, S.R. Hallett, Smoothing artificial stress concentrations in voxel-based models of textile composites, *Compos. Part A: Appl. Sci. Manuf.* 80 (2016) 270–284.
- [59] T. Zheng, L. Guo, Z. Tang, T. Wang, Z. Li, Comparison of progressive damage simulation of 3D woven composites between voxel and conformal discretization models, *Mech. Mater.* 158 (2021) 103860.