

Deep-learning image enhancement and fibre segmentation from time-resolved computed tomography of fibre-reinforced composites

Rui Guo^{*1}, Johannes Stubbe², Yuhe Zhang², Christian Matthias Schlepütz³, Camilo Rojas Gomez¹, Mahoor Mehdikhani¹, Christian Breite¹, Yentl Swolfs¹, Pablo Villanueva-Perez²

¹Department of Materials Engineering, KU Leuven, Kasteelpark Arenberg 44 box 2450, 3001 Leuven, Belgium

²Division of Synchrotron Radiation Research and NanoLund, Department of Physics, Lund University, Lund, 22100, Sweden

³Swiss Light Source, Paul Scherrer Institute (PSI), 5232 Villigen, Switzerland

*Corresponding author: R. Guo (rui.guo1@kuleuven.be), Tel.: +32486965754

Email addresses: J. Stubbe(jo0720st-s@student.lu.se), Y. Zhang (yuhe.zhang@sljus.lu.se), C.M. Schlepütz (christian.schlepuetz@psi.ch), C.R. Gomez (camilo.rojas@kuleuven.be), M. Mehdikhani (mahoor.mehdikhani@kuleuven.be), C. Breite (christian.breite@kuleuven.be), Y. Swolfs (yentl.swolfs@kuleuven.be), P. Villanueva-Perez (pablo.villanueva_perez@sljus.lu.se)

Abstract

Monitoring the microstructure and damage development of fibre-reinforced composites during loading is crucial to understanding their mechanical properties. Time-resolved X-ray computed tomography enables such an in-situ, non-destructive study. However, the photon flux and fibre-matrix contrast limit its achievable spatial and temporal resolution. In this paper, we push the limits of temporal and spatial resolution for the microstructural analysis of unidirectional continuous carbon fibre-reinforced epoxy composites by establishing a new pipeline based on CycleGAN for unsupervised super-resolution and denoising and U-Net-id for individual fibre segmentation. After illustrating the benefits of a 3D CycleGAN over a 2D one, we show that data enhanced by this pipeline can yield similar segmentation quality to that of a slow-acquisition, high-quality scan that took up to 200 times longer to acquire. This pipeline, therefore, enables more robust data extraction from fast time-resolved X-ray tomography, removing a critical stumbling block for this technique.

Keywords: A. Polymer-matrix composites, A. Carbon fibre, D. X-ray computed tomography, D. Non-destructive testing, Deep learning

1 Introduction

Fibre-reinforced composites (FRCs) are widely used in engineering for their excellent mechanical properties and low density. The damage development in FRCs is often a complex interplay between fibre breaks, matrix plasticity, fibre/matrix debonding at the microscale, and off-axis cracking and delaminations at the mesoscale. Monitoring the damage during loading is key to better understanding the mechanical properties and, eventually, developing improved materials [1]. Such non-destructive analyses are enabled through X-ray computed tomography (CT). Due to the much higher flux compared to lab-scale CT, synchrotron radiation CT (SRCT) is particularly suited for time-resolved analyses, as the typical SRCT scan times per volume range from 1 s to 5 min.

Despite the benefits in terms of time resolution, SRCT at scan times in the order of seconds comes with significant compromises in terms of image quality. The resulting images are often noisy, with limited spatial resolution and low contrast between fibre and matrix, especially for carbon FRCs. FRCs also have both high average (55-67% [2]) and local fibre volume fractions (up to ~85% within longitudinal fibre bundles [3]), making the fibres challenging to separate from the matrix with traditional segmentation methods like greyscale thresholding (Otsu's method [4]). Emerson et al. [5] developed a machine-learning-based tool called Insegt Fibre, which can accurately extract the centre of each fibre. This method is robust for relatively low-resolution images but still has some limitations [6, 7]. For example, when a fibre cross section in a CT slice is not fully circular, the extracted fibre centre can have multiple locations [6].

Deep-learning-based segmentation methods have gained popularity in recent years. Recently, they have been used to segment cracks [8] and textile structures for woven composites [9, 10]. Badran et al. [11] used FCDenseNet [12] based on Dragonfly software to successfully segment fibres, matrix and an environmental barrier coating, despite having minimal contrast between fibres and matrix. Their method, belonging to the class of semantic segmentation techniques, struggles with separating adjacent fibres from each other. Instance segmentation, how-

ever, can distinguish each object, making it an attractive method for individual fibre segmentation. Standard instance segmentation methods can be either i) two-stage segmentation, like U-Net plus Watershed [13], or ii) single-stage segmentation, like Mask R-CNN [14], Deep Watershed Transform [15]. Compared to Mask R-CNN, U-Net can obtain better results on the same datasets (higher Dice score) [16]. In contrast with these methods, U-Net-id [17], which is also an instance segmentation method based on the U-Net network, does not require further decision-making by the user. These deep-learning-based segmentation methods require the manual labelling of fibres as training data. However, in low-resolution and noisy images, it is difficult to label the pixels that belong to a fibre accurately. This results in flawed segmentation models for low-quality data. Therefore, high-quality and high-resolution images are required.

Deep-learning approaches have also demonstrated an unprecedented capability to enhance image quality by denoising [18] and retrieving super-resolution [19]. Most of these approaches require supervision [20], meaning that they expect paired datasets: each noisy or low-resolution image should have a corresponding low-noise and high-resolution image. However, obtaining these paired datasets for time-resolved experiments is almost impossible as most of the specimens are not reproducible, and the damage mechanisms have a stochastic nature. This is also the case for fibre-reinforced composites during loading.

Therefore, image enhancement approaches that do not require supervision or paired datasets are needed. An example of a deep-learning approach that does not require paired datasets is CycleGAN [21]. It provides models to translate image domains in an unsupervised manner by training or optimising between source and target domains without those images being paired. Thus, CycleGAN offers the opportunity to expand the temporal and spatial boundaries of time-resolved tomography by applying the image quality of slow-acquisition tomograms with low noise and high resolution to noisy and low-resolution fast-acquisition tomograms.

In this work, we explore how super-resolution and denoising can be exploited to improve the segmentation of individual fibres in SRCT images of unidirectional fibre-reinforced composites. We propose a new pipeline for high-quality segmentation based on CycleGAN and U-Net-id. We investigate the benefits of 3D CycleGAN over its 2D counterpart to generate high-resolution and denoised images from low-resolution and fast-acquisition datasets. Then, we compare the segmentation quality of the 2D U-Net-id for a slow-acquisition dataset, a fast-acquisition dataset, and a CycleGAN-enhanced dataset. The results show that our pipeline can extract the actual fibre microstructure information, which was impossible from the original low-resolution and noisy datasets.

2 Materials and methods

2.1 Materials and specimen preparation

TORAYCA T700SC-12K carbon fibres with 7 μm nominal diameter were sourced from Toray Carbon Fibre Europe. Manufacturing started by making pre-impregnated unidirectional sheets using drum winding. Beginning with dry rovings, which soak up resin while being guided through a resin bath, the fibres were wound onto a translating and rotating drum. We used the epoxy resin Sicomin SR8500 KTA313, which has a glass transition temperature of 120 $^{\circ}\text{C}$ and an E-modulus of 3.36 GPa [1]. The resulting 1 mm thick prepreg sheets of $\sim 300 \times 1900 \text{ mm}^2$ were then cut into smaller pieces of $300 \times 300 \text{ mm}^2$ and stacked in a $90^{\circ}/0^{\circ}/90^{\circ}$ layup. The layups were cured using a computerised autoclave under a combination of -0.7 bar vacuum pressure and 5 bar of overpressure. The curing cycle corresponded to the resin manufacturer's specification: a pre-curing stage of 60 min at 70 $^{\circ}\text{C}$ before the temperature was increased to the curing level of 120 $^{\circ}\text{C}$ for 90 min. After curing, 1 mm thick aluminium 2014-T6 end tab panels were bonded to the composites using 3M Scotch-Weld EC-9323 B/A structural adhesive. The adhesive was cured at 60 $^{\circ}\text{C}$ for 60 min. From these panels, specimens with a double T-shape at their ends and a double, circular notch in the middle were water jet cut [1].

2.2 Methods

2.2.1 Time-resolved computed tomography

We performed static and time-resolved dynamic tomography at a synchrotron radiation facility on the fibre-reinforced composite specimens during static mechanical loading. The experiments were conducted at the TOMCAT beamline of the Swiss Light Source (SLS) in Switzerland [22]. A Deben CT500 loading rig with a 500N load cell and poly(methyl methacrylate) protective tube was used. The rig was modified to be able to mount the T-shaped specimen ends. A small 10N load was applied after mounting the specimens to ensure stability during rotation.

Static tomography was used to obtain slow-acquisition tomograms, yielding high-quality images under static conditions. For that, we used X-rays at 15 keV from a multi-layer monochromator. The monochromatic X-rays illuminated our specimen, and its signal was detected by an indirect detector, which used a 20 μm thick LuAG: Ce scintillator to transform X-rays into visible light. We used a PCO.edge 5.5 camera with a square pixel size of 6.5 μm . The pixel size was demagnified using a 20 \times microscope, which provided an effective pixel size of 0.325 μm . Each tomographic acquisition with this system consisted of 2000 projections between 0 and 180 degrees and 0.25 s acquisition time per projection. Before tomographic reconstruction, we performed flat-field corrections to remove pixel-to-pixel noise due to systematic errors caused by the illumination and imaging setup. For that purpose, we collected: i) 30 dark current images, which are images without X-rays, and ii) 100 flat images, which are images with X-rays but without the specimen. 50 of the flat images were collected right before the tomographic acquisition and the rest right after. The tomographic reconstructions were performed using GridRec [23] together with a Parzen window. The obtained tomographic reconstructions were used for either reference or training the high-quality slow-acquisition tomograms. We refer to this as the slow-acquisition dataset.

Time-resolved dynamic tomography was used to acquire in-situ tomograms, yielding a high temporal resolution but also a high noise level and poor image resolution. For such acquisitions, we used the high flux provided by the white beam, which is the polychromatic radiation produced by the superbending magnet of TOMCAT. To increase the temporal resolution, we used the GigaFRoST detector [24] coupled to a high numerical aperture microscope by Optique Peter and a 100 μm thick LuAG: Ce scintillator. The effective pixel size of this system was 0.81 and 1.62 μm (compared to 0.325 μm for the slow-acquisition tomograms). Each tomogram was acquired with 500 projections between 0 and 180°. We used different acquisition times per projection (0.5, 1, and 3 ms) to study the effect of noise and evaluate the limits of the proposed approach. We will refer to them as the 0.5 ms (1.62 μm pixel size with 0.5 ms acquisition time per projection), 1 ms (0.81 μm pixel size with 1 ms acquisition time per projection), and 3 ms (1.62 μm pixel size with 3 ms acquisition time per projection) fast-acquisition datasets. To obtain the tomographic reconstructions for the fast-acquisition scans, we followed the procedure described for the slow-acquisition tomograms. The acquired tomograms with these settings will be used as the datasets on which we will apply our so-called 2 \times and 4 \times super-resolution using CycleGAN (see section 2.2.2).

2.2.2 CycleGAN for denoising and super-resolution

CycleGAN is a deep-learning algorithm that can learn to perform domain translations between two image domains. In our experiment, we have two domains: i) a *fast* domain, which contains low-resolution and noisy tomograms, and ii) a *slow* domain, which contains high-resolution and high-contrast tomograms obtained at slower rates. For simplicity, we call these two domains *fast* and *slow*. A schematic representation of the CycleGAN approach is provided in Fig. 1. On the one hand, the CycleGAN algorithm learns to enhance a low-quality tomogram from the *fast* domain to resemble a high-quality tomogram in the *slow* domain. On the other hand, CycleGAN also learns to transfer high-quality tomography from the *slow* domain to resemble the low-quality tomograms in the *fast* domain. Even though only the enhancing transformation (from *fast* to *slow* domains) is relevant for our purpose, both translations are required for the training process.

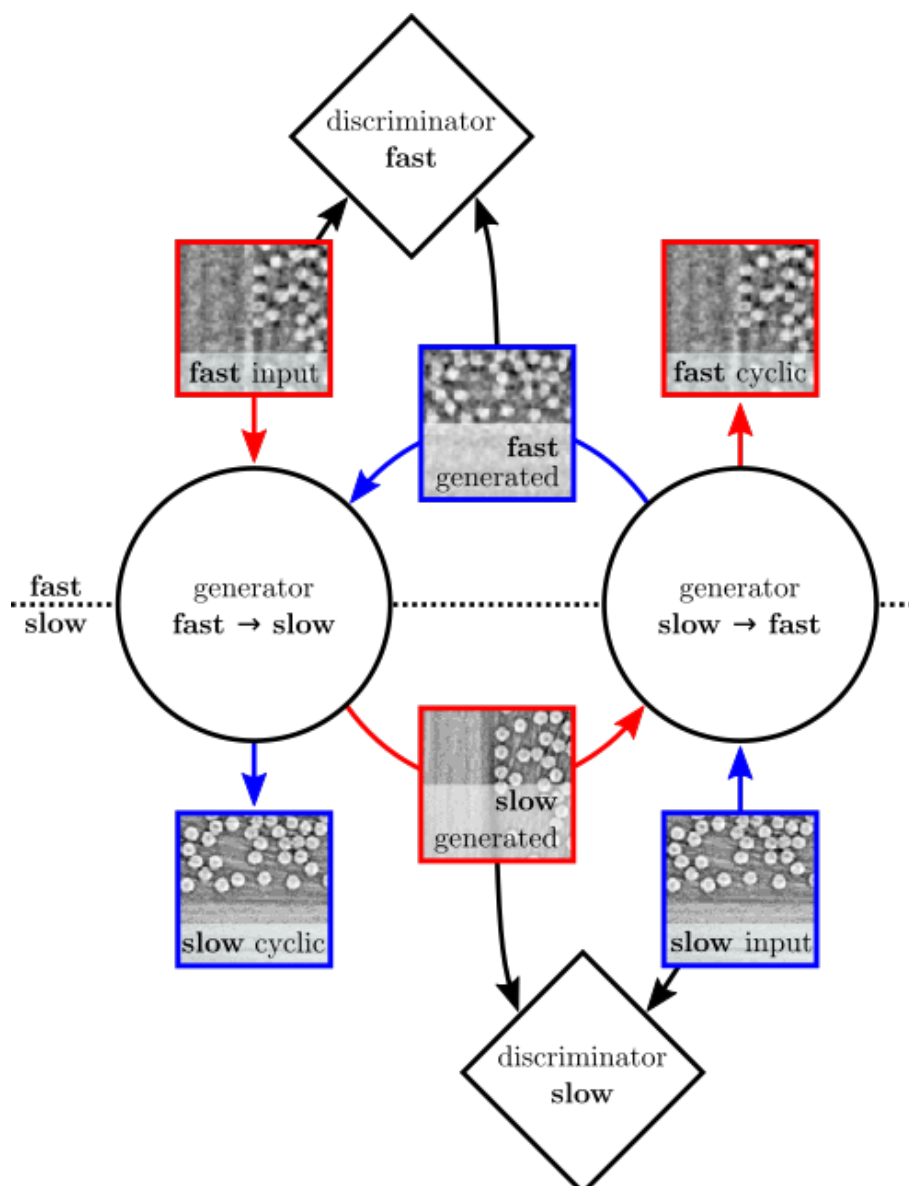


Fig. 1: Simplified flow chart of two cycles in CycleGAN. The fast-slow-fast (slow-fast-slow) cycle is outlined with red (blue) arrows. An image from the fast domain is passed to the generator 'fast → slow', which produces an enhanced image 'slow generated'. This image is evaluated by the discriminator slow and then passed to the generator 'slow → fast', which produces a new image called 'fast cyclic'. The cycle consistency requires the identity between 'fast input' and 'fast cyclic'.

By training both translations simultaneously, we ensure that when performing a full translation cycle from domain *fast* to *slow* and back to the *fast* domain, we retrieve the original image, and all information is conserved. This feature is called cycle consistency and is the reason for the name CycleGAN. The GAN refers to the generative adversarial network [25], a deep-learning framework that trains two neural networks, named generator and discriminator, which compete against each other. The discriminator tries to categorise whether its input is an original tomogram or was produced by the generator. In contrast, the generator tries to achieve a good transformation that fools the discriminator. CycleGAN uses two GANs, one for the translation from the *fast* to the *slow* domain and another one for the translation from the *slow* to the *fast* domain. Then, the cycle consistency requirement is applied, and the optimisation process or training is repeated until the results are satisfactory.

As a result of this cyclic approach, CycleGAN does not require a paired dataset with high-quality and low-quality images of the same specimen. This unpaired machine learning approach is sometimes called unsupervised learning. It is advantageous for our in-situ images as the specimens are destroyed after a loading test, and it is not possible to acquire slow and fast datasets on the same specimen. Furthermore, high-quality and low-quality tomograms can be acquired at different times and facilities. In our experiments, we captured slow and fast tomograms on different specimens and setups to train the CycleGAN, i.e., unpaired slow and fast datasets. We implemented

a 2D and a 3D version of CycleGAN. The 2D CycleGAN is trained with 2D tomographic slices, while the 3D version is trained with 3D cubes over the whole 3D reconstructed volume. For both approaches, we realised generators that retrieve super-resolution with scale factors $2\times$ and $4\times$. Since it enhances individual slices, 2D CycleGAN can apply super-resolution in two dimensions, while 3D CycleGAN increases the pixel count in all three spatial dimensions.

The slow-acquisition data were rescaled to an effective pixel size of $0.405\ \mu\text{m}$ to enable $2\times$ and $4\times$ super-resolution transformations of the $0.81\ \mu\text{m}$ and $1.62\ \mu\text{m}$ fast-acquisition tomograms, respectively. For this paper, we trained four CycleGANs: one 2D CycleGAN and three 3D CycleGANs, and Table 1 displays the attributes of the training datasets and the number of optimisation iterations with the training time for each network. The voxel sizes of the slow-acquisition datasets mentioned in Table 1 also determine the voxel sizes of the corresponding enhancements.

The algorithms were trained on a GPU cluster consisting of 4 NVIDIA Tesla V100 GPUs with $4\times 32\ \text{GB}$ VRAM. During this phase, the 2D version of CycleGAN could also be run on GPUs with smaller memory since it is less memory intensive and converges towards sound enhancements faster than the 3D version. Once the training is complete, only the enhancing generative network the “fast \rightarrow slow” (see Fig. 1) is used to improve the image quality of the fast-acquisition tomogram, and both algorithms produce enhancements of an $800 \times 1200 \times 400$ voxel volume within 2 minutes with the aforementioned hardware.

However, it is important to note that CycleGAN optimisation is unstable as it involves training four neural networks and hence relies on carefully chosen hyperparameters. Details on our hyperparameter choices can be found in Appendix A.

Table 1. Training details for 2D/3D CycleGAN enhancement algorithms.

Enhancement	Training data voxel size [nm^3]*		Optimisation iterations	Training time [h]
	domain <i>fast</i>	domain <i>slow</i>		
2 \times -1ms- enhanced (2D)	810 \times 810 \times 810		8000	4
	810 \times 405 \times 405			
2 \times -1ms- enhanced (3D)	810 \times 810 \times 810		32000	82
	405 \times 405 \times 405			
4 \times -3ms- enhanced (3D)	1620 \times 1620 \times 1620		12000	34 [†]
	405 \times 405 \times 405			
4 \times -0.5ms- enhanced (3D)	1620 \times 1620 \times 1620		25000	24 [†]
	405 \times 405 \times 405			

*The shape of the 3D array, including other parts in our paper, is depth, height, and width, where the depth is along the fibre direction.

†Dynamic hyperparameter updates during optimisation prolonged the training time and caused the contradictory number of iterations and training time.

2.2.3 Deep-learning-based segmentation

Fig. 2 provides a schematic overview of the segmentation process, consisting of three steps.

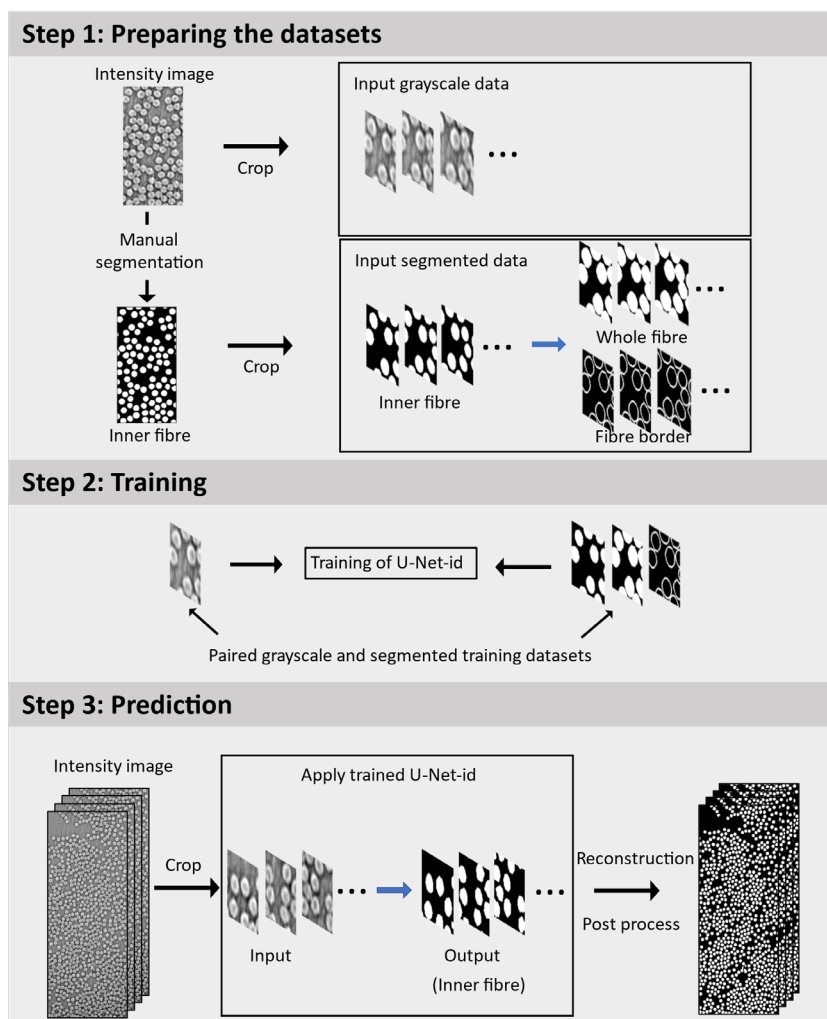


Fig. 2: Schematic of the individual fibre segmentation process using the U-Net-id network, from preparing the datasets to training and prediction.

In the data preparation stage (Step 1), a tomographic image is converted from floating point values to a greyscale image with values ranging from 0 to 255. As a supervised learning method, the training process for the U-Net-id model requires both greyscale input data and corresponding labels, i.e. inner fibres, whole fibres, and fibre borders. A single CT slice with sufficient details, including an adequate number of fibres, is chosen for the greyscale input data. The inner fibre is manually labelled using a manual annotation tool called Insegtpy, which is the Python version of Inseg Fibre [26, 27]. The greyscale data and manually labelled inner fibre data are then sampled into 64×64 pixels with a stride of 8 pixels. Afterwards, the whole fibres are obtained by two iterations of dilating the inner fibres with a cross kernel and the fibre borders are obtained by subtracting the inner fibres from the whole fibres, see Fig. 2.

In the training stage (Step 2), the greyscale input data is normalised to the range of 0 to 1 before being fed into the U-Net-id network [17], which is designed based on the U-Net for instance segmentation. The sum of binary cross-entropy and dice coefficient-related loss was used as the loss function [17], the Adam algorithm was used as the optimiser, and the learning rate was set to 0.001. During training, data augmentation was applied randomly to the input images to improve the network's generalisation [28]. This augmentation includes changes in brightness and contrast with the same factor chosen uniformly from 0.7 to 1.3 as well as Gaussian blur with a 5×5 Gaussian kernel and sigma from 0.7 to 1.3. The network is trained with 200 epochs and a batch size of 16. Finally, once the U-Net-id network is trained, large CT datasets can be segmented using a sliding window approach with the trained model (Step 3). However, the segmentation of fibres along the image borders is often suboptimal, which can negatively impact the final segmentation results. Significant overlap is ensured between two sliding windows to address this issue, allowing for more accurate fibre segmentation by focusing on extracting areas that do not include fibres at the image border. As the output of the algorithm is not binary, a threshold value of 0.5 is chosen to binarise the images. The network's direct outputs are fibre borders, whole fibres and inner fibres, which can be combined for further separation of fibres. However, in the current study, all further analysis is based on the inner fibres.

For this paper, we trained eight different U-Net-id models, one for each dataset (three fast-acquisition datasets, one slow-acquisition dataset, one 2D CycleGAN and three 3D CycleGAN-enhanced datasets). The physical size of each data is larger than $607.5 \times 601.8 \times 356.4 \mu\text{m}^3$. The resolution of all 3D CycleGAN enhanced data matches the slow-acquisition data, with cubic voxels of $0.405 \times 0.405 \times 0.405 \mu\text{m}^3$. For 2D CycleGAN, the slice resolution is the same, but along the thickness, it remains consistent with the 1ms-fast-acquisition data, which has a voxel size of $0.81 \times 0.81 \times 0.81 \mu\text{m}^3$. The voxel size for the 3ms- and 0.5ms-fast-acquisition data is $1.62 \times 1.62 \times 1.62 \mu\text{m}^3$.

For the datasets with voxel edges smaller than $1.62 \mu\text{m}$ (1ms-fast-acquisition, the enhanced, and the slow-acquisition datasets), a region of 300×150 pixels was selected for manual labelling. For the 0.5ms- and 3ms-fast-acquisition datasets, an image size of 200×100 pixels was chosen for training due to the larger pixel size and to limit the manual labelling work. As a result, we generated 330 patches from 300×150 pixels crops and 90 patches from 200×100 pixels crops, all patches being of a uniform size of 64×64 pixels. The total time required to train these segmentation models can be completed within around two hours using an Nvidia RTX A5000 GPU with 24 GB of VRAM.

It is worth noting that after segmentation, a limited number of artefacts are still present. Some artefacts can be removed - using an opening operator with a disk-shaped kernel (radius=4) for high-resolution images (slow-acquisition, 2D and 3D CycleGAN enhanced). However, this opening operator was not applied to the low-resolution images, as attempting to remove artefacts in such cases would also result in the removal of actual fibres.

2.3 Assessment metrics

To evaluate the performance of our CycleGAN - U-Net-id pipeline, we also acquired fast and slow tomograms of the same volume in a non-destructive manner. Specifically, we recorded tomograms of one intact specimen using the four previously described acquisition modes: slow-acquisition, 1ms-, 3ms-, and 0.5ms-fast-acquisition. These tomograms constitute a paired test dataset, and it is essential to highlight that they were not involved in the CycleGAN training and were only used to evaluate our algorithms.

Our evaluation process involves a comparison of the fast-acquisition tomograms from the paired test dataset, their CycleGAN enhancements, and the corresponding slow-acquisition reference, as well as their U-Net-id segmentations. To make a fair visual inspection and quantitative comparison of the same regions for all these datasets, we performed image registration based on the greyscale 3D images with Avizo and then extracted the calculated transformation values to transform the segmented data. However, such transformations (e.g., rotation, scaling) can lose the binary nature of the segmented data due to interpolation. To rebinarise these images, we used a threshold value of 125.

Geometric parameters such as the number, size, shape, and length of the fibres significantly influence the mechanical properties of fibre-reinforced composites. Therefore, these parameters form the primary criteria for evaluating the performance of our pipeline in terms of enhancement and segmentation.

The fibre sizes and shapes are determined using 2D slices. For such studies, the fibres at the edge of the image are excluded to avoid the influence of incomplete fibres. The equivalent fibre diameter D_f in μm can be calculated by:

$$D_f = \sqrt{\frac{4 * S_f}{\pi}} * pxs \quad (1)$$

where S_f is the total number of fibre-labelled pixels in the binary images, and pxs is the pixel size in μm . Each segmented fibre is then fitted with an ellipse using the Scikit-image library. The aspect ratio, which is the ratio of the major over the minor diameter of the ellipse, can be used to describe the fibre shape, where 1 corresponds to a perfect circle.

To calculate the fibre length in 3D, we used an effective tracking method based on the K-nearest neighbour algorithm [5, 7]. The Scikit-image library is used to locate the centre of each fibre in each slice. The fibre centres on each slice are connected to define fibre trajectories. The maximum distance to query the nearest fibre centre for each pair of adjacent slices is set to 1 pixel for all data. This setting allows us to detect specific segmentation errors effectively (see Fig. D.1 in Appendix D for examples). The fibre length can then be measured by the distance from the start to the end point of the tracked segment [29]. To maintain any possible error/artefact in fibre

detection, we did not perform further post-processing steps such as splitting, elimination, interpolation and smoothing [5], or blind and smart stitching [7] to quantify the impact of missing fibres and noise. This is an essential benefit of the proposed pipeline.

3 Results

This section first compares the 2D and 3D CycleGAN approaches by contrasting enhancements of the respective $2\times$ super-resolution networks applied to 1ms-fast-acquisition data of fibre-reinforced composites. Later, we demonstrate the full potential of the CycleGAN - U-Net-id pipeline for tomography enhancements by presenting results from the two-remaining 3D CycleGANs with $4\times$ super-resolution for 3ms and 0.5ms fast acquisition data. The subsequent evaluations exclusively rely on the paired test dataset after image registration.

3.1 Comparison of 2D and 3D CycleGAN

The 3D approach trains on cubes rather than 2D slices, offering a complete view of the fibre structure and providing an opportunity to learn features that are parallel to the tomographic axis.

To demonstrate the capabilities of 3D CycleGAN compared to its 2D counterpart, we used both algorithms to enhance 1ms-fast-acquisition tomograms with $2\times$ super-resolution and compared their results. Fig. 3 displays tomographic cross-sections from (a) the 1ms-fast-acquisition input, (b) the 2D CycleGAN $2\times$ enhancement, (c) the 3D CycleGAN $2\times$ enhancement, and (d) the slow-acquisition reference.

The fast-acquisition image in Fig. 3a is noisy, making it difficult to recognise individual fibres. Furthermore, their borders are blurred and non-circular. Fig. 3e presents a tomographic slice along the horizontal yellow line and parallel to the fibres. Intersecting fibres cause bright vertical lines.

Figs. 3b,f show the enhancements for the fast-acquisition images after using the 2D CycleGAN algorithm. The fibres are mostly round and distinguishable. However, the 2D CycleGAN enhancement in Fig. 3b displays fibre hallucinations that are not present in the slow-acquisition reference in Fig. 3d. The blue circles indicate such hallucinations that can extend through multiple slices along the fibres. For example, the blue arrow in Fig. 3f points to the position of a fibre hallucination, which was systematically introduced in multiple slices, mimicking the characteristics of an actual fibre. The occurrence of such systematic hallucinations impacts the authenticity and reliability of 2D CycleGAN enhancements of high-noise data. Additionally, the real 2D CycleGAN-enhanced fibres look frayed along their trajectory due to slightly inconsistent fibre centres across slices.

The number of hallucinations in the 3D CycleGAN enhanced dataset (see Fig. 3c and Fig. 3g) is limited, and none resemble fibres. The enhanced fibres are rounder, straighter and less frayed compared to those observed in the 2D CycleGAN enhancement.

In Appendix B, we evaluate the enhancement capabilities of the 2D and 3D CycleGAN algorithms by reporting the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index (SSIM) compared to the slow acquisition reference. We also quantify the occurrence of fibre hallucinations.

It is clear that the 3D CycleGAN outperforms the 2D CycleGAN in terms of image quality and reliability. The lack of fibre hallucinations in the 3D enhancement warrants using this computationally heavy algorithm over the 2D version. Consequently, all the results presented from now on are based on our 3D approach. Furthermore, we refer to 3D CycleGAN simply as CycleGAN in the rest of the paper.




Fig. 3: Comparison of cross-sectional slices of the (a,e) 1ms-fast-acquisition input, (b,f) 2D CycleGAN 2× enhancement, (c,g) 3D CycleGAN 2× enhancement and (d,h) slow-acquisition reference. The right-hand side images (e-h) display slices in the fibre orientation that intersect the left-hand side images (a-d) through the yellow lines. The blue circles and arrows highlight hallucinated artefacts which resemble fibres.

3.2 Individual fibre segmentation

Fig. 4 presents the 3D overlay of greyscale and segmented binary images with U-Net-id for the 1ms-fast-acquisition data, slow-acquisition data and associated 2× enhanced data. The dashed rectangles highlight regions with segmented artefacts that are not fibres, which exist in all of these cases. However, while the 1ms-fast-acquisition tomogram is riddled with segmentation artefacts, they are sparse in the slow-acquisition tomogram and the CycleGAN enhancement. The segmentation algorithm can detect some fibres in the low-quality 1ms-fast-acquisition images (see Fig. 4a). Compared with the segmentation results on the slow-acquisition data (see Fig. 4b), however, the segmented fibres along the thickness are rough. This is mainly because the segmentation algorithm cannot accurately capture the fibre information, such as fibre size, shape and position, at the low resolution and high noise level of the fast-acquisition tomogram. After the 2× super-resolution CycleGAN enhancement, the segmented fibres are much smoother than in the fast-acquisition data and similar to the slow-acquisition data.




Fig. 4: 3D overlay of greyscale and segmented binary images with U-Net-id for (a) 1ms- fast-acquisition data, (b) slow-acquisition data, and (c) associated $2\times$ enhanced data. The blue-coloured structures represent the segmented fibres, and the dashed black rectangles highlight artefacts that are not fibres.

Fig. 5 compares the segmentation results of the slow-acquisition dataset to the fast-acquisition and enhanced datasets. The key artefacts or errors are highlighted with ellipses:

- Red: touching fibres, where the algorithm did not separate the fibres.
- Blue: fibres that are not successfully segmented.
- Green: an imaging artefact in the slow-acquisition dataset that causes segmentation difficulties.
- Orange: falsely detected fibres, where there are no actual fibres.

Although segmentation errors are present across all datasets, including slow-acquisition, fast-acquisition, and enhancements, the displayed slices effectively represent the datasets concerning the segmentation quality and the types and quantities of errors. Touching fibres (red ellipses) frequently appear in the 1ms- and 0.5ms-fast-acquisition datasets. Although no touching fibres are detected for this particular slice in the 3ms-fast-acquisition data and all enhanced data, this does occur in other slices. Moreover, some fibres are not successfully detected in all fast-acquisition datasets (see the blue ellipses in Figs. 5d,f,h). The blue ellipse in Fig. 5b shows a missing fibre in the segmentation of the slow-acquisition data due to the presence of an imaging artefact (see the green ellipse in Fig. 5a). Orange ellipses in the 1ms-fast-acquisition data highlight some artefacts belonging to the matrix instead of fibres. The same artefacts are present in the enhanced 0.5 ms results where they occur due to faint fibre hallucinations that can also be found at the same positions in Fig. 5m.




Fig. 5: Comparison of the segmentation results of the (a,b) slow-acquisition data, (c-h) fast-acquisition data with acquisition times of 1 ms, 3 ms, and 0.5 ms, and (i-n) the corresponding enhanced data. The green ellipse indicates the existing artefacts in the slow-acquisition slice. Red ellipses highlight touching fibres that are not separated successfully, blue ellipses highlight missing fibres, and orange ellipses highlight wrongly detected fibres.

For the fast-acquisition data, the image quality is poor, resolution and contrast are low, and the size and shape of the segmented fibres are highly inconsistent. Comparing the 3ms- and 0.5ms-fast-acquisition datasets, these issues become significantly worse for the shorter acquisition times. After enhancement, however, these issues largely disappear. The 4x-0.5ms-enhanced dataset still shows some artefacts, but this is also the most challenging dataset of the three.

In terms of fibre shapes, the segmentation of the enhanced CT slices in Figs. 5j,l,n are much more homogeneous than the corresponding fast-acquisition segmentations. The fibres have similar sizes to the slow-acquisition reference and can predominantly be described as round. Most fibres that are missing in the fast-acquisition data (see blue ellipses) are successfully detected in the enhanced data. Fibre sizes, shapes and positions are similar in the enhanced and slow-acquisition data segmentation results.

3.3 Quantitative analysis

In this section, we present a quantitative analysis using quantities that are crucial to fibre analysis to evaluate our enhancement-segmentation pipeline's capabilities. We focus on fibre diameter, aspect ratio, fibre count per slice, and 3D fibre count as assessment metrics. Additionally, in Appendix C, we evaluate the enhancement and segmentation accuracy based on pixel-based metrics, such as PSNR, SSIM, and Intersection over Union (IoU). Appendix D demonstrates how 3D fibre count effectively captures segmentation errors without requiring manual labelling, unlike pixel-based segmentation metrics.

The quantitative analyses were performed on the same volumetric region for all studied datasets, measuring $401.76 \times 369.36 \times 152.28 \mu\text{m}^3$ and containing approximately 680 fibres. It is important to note that the fibre count per slice and 3D fibre count include fibres at the image border.

For comparison purposes, we manually corrected the segmentation errors in the slow-acquisition segmentation results and refer to this as ground truth. Fig. 6 shows letter-value plots [30] (or boxenplot [31]) of the equivalent diameter and aspect ratio of the fibres for this large number of datasets. The number of boxes to draw is based on the trustworthy method [31]. For the slow-acquisition data, outliers are clearly visible for the equivalent fibre diameter, compared to the ground truth. The small fibre diameters are mainly caused by small artefacts that were misidentified as fibres, whereas the large fibre diameters are mainly due to touching fibres.

The letter-value plots show that for the 1ms-fast-acquisition data, measured fibre diameters are significantly smaller than for the slow-acquisition data and ground truth. This is consistent with a visual inspection of Fig. 5. The aspect ratio is larger for the fast-acquisition than the slow-acquisition dataset and ground truth since the segmented fibre shape is less circular. This implies that the fast-acquisition datasets cannot be used to extract geometrical information of the fibres. Similar results can be seen for the 3ms- and 0.5ms-fast-acquisition datasets. As noted in Fig. 6, after the enhancement of the fast-acquisition dataset, the distributions of fibre diameter and aspect ratio become similar to those of the slow-acquisition dataset and ground truth, although some outliers remain.

The analysis of the average values, presented in Table 2, also leads to the same conclusion: the average fibre diameter and aspect ratio for all enhanced datasets are close to the values of the slow-acquisition datasets and ground truth. Thus, the enhanced images can retrieve fibre information that cannot be obtained from the original fast-acquisition image.




Fig. 6: Statistical distribution of the fibre geometrical parameters in different datasets: letter-value plot for (a) equivalent fibre diameter and (b) fibre aspect ratio.

Table 2: Average and standard deviation of the fibre geometrical parameters as well as the fibre count in different datasets.

Dataset name	Equivalent fibre diameter (μm)	Fibre aspect ratio	Fibre count (all slices)	3D fibre count
Ground-truth	6.08 ± 0.27	1.08 ± 0.04	684.65 ± 1.05	676
slow-acquisition	6.08 ± 0.29	1.08 ± 0.05	687.43 ± 2.01	620
1ms-fast-acquisition	3.45 ± 0.87	1.40 ± 0.38	655.10 ± 8.18	0
2x-1ms-enhanced	6.16 ± 0.30	1.09 ± 0.06	689.10 ± 1.96	616
3ms-fast-acquisition	3.60 ± 0.84	1.44 ± 0.39	610.76 ± 8.14	408
4x-3ms-enhanced	6.25 ± 0.27	1.09 ± 0.05	687.56 ± 1.40	659
0.5ms-fast-acquisition	3.25 ± 0.79	1.52 ± 0.47	549.68 ± 10.34	58
4x-0.5ms-enhanced	6.16 ± 0.34	1.11 ± 0.11	694.21 ± 3.88	548

Table 2 also presents the average and standard deviation of the fibre count for each slice along the volume depth ($401.76 \mu\text{m}$) and the 3D fibre count, which we defined as the number of segmented fibres that exceed a length

threshold of 364.5 μm . The 3D fibre count in the ground truth is approximately 1.3% lower than the average fibre count per slice. This discrepancy is primarily due to fibres at the border of the analysed volume that move in and out of the field of view. Comparing the slow-acquisition data to the ground truth, the 3D fibre count is reduced by 8.3%. This decrease can primarily be attributed to imaging artefacts, like the one shown in Fig. 5a, which prevent successful segmentation in individual slices and disrupt the otherwise intact fibres (See Fig. D.1 in Appendix D for possible cases of unsuccessful tracking). In the 0.5ms-, 1ms- and 3ms-fast-acquisition datasets, the number of detected 3D fibres is much lower than in the slow-acquisition data. This highlights the subpar quality of tomograms at fast acquisition speeds and the need for effective enhancements of such data.

For the enhanced datasets, significantly more 3D fibres were detected compared to the fast-acquisition datasets thanks to the CycleGAN enhancement. The number of detected 3D fibres in the enhanced images is also close to the ground truth, and the 3D fibre count in the 4 \times -3ms-enhanced data even beats the slow-acquisition 3D fibre count by 39 fibres. This can be explained by fewer imaging artefacts being present in this CycleGAN enhancement compared to the slow-acquisition data. In the most challenging 4 \times -0.5ms-enhanced dataset, we can report \sim 10 more detected single fibres, but a decrease of 18.9% of the 3D fibres compared to the ground truth. The overcount of average fibres per slice can be explained by the hallucinations that were also detected in Fig. 5m, while the undercount of 3D fibres results from the inferior image quality of this enhancement compared to the slow-acquisition reference. Some of these incomplete fibres can be recovered through additional error-correcting algorithms, such as those employed by Sosa-Rey et al. [7] and Emerson et al. [5]. Nonetheless, this is out of the scope of this paper and we consider the 4 \times -0.5ms-enhancement as the limit for tomography enhancements with our algorithm.

Based on the presented datasets, we can summarise that after super-resolution with unpaired data using CycleGAN, the geometric information of the fibres can be characterised with high accuracy. Moreover, 3D fibre trajectories can be reconstructed reasonably well if the fast-acquisition dataset is not of extremely low quality.

4 Conclusions

We propose a novel deep-learning-based pipeline for high-quality fibre segmentation of fast time-resolved computed tomography scans. This pipeline consists of two deep-learning methods. First, CycleGAN applies super-resolution to the fast-acquisition (low-quality) data and denoises them. Then, U-Net-id uses these enhanced images to segment individual fibres effectively. CycleGAN does not require paired datasets for training, which is a crucial advantage compared to traditional super-resolution methods.

2D CycleGAN can yield good-quality 2D slices but is prone to introducing fibre-shaped hallucinations and fails to properly denoise the third direction. On the other hand, 3D CycleGAN can access volumetric contextual information, and therefore (1) yields higher quality 2D slices with accurate fibre information, and (2) improves the resolution along the fibres. Segmentation of the fast-acquisition datasets yields many errors, including false fibre detections, missing fibres and merged fibres. The segmentation of the enhanced datasets, however, works and yields similar quality to the slow-acquisition data, as observed for both cross-sectional and transverse slices. The segmentation errors, which are observed for the enhanced datasets also appear in the slow-acquisition data.

Quantitative analysis revealed that the segmented fibre shape and diameter are similar for slow-acquisition and enhanced data, while they cannot be retrieved directly from the fast-acquisition data. The number of long 3D fibres that could be segmented from the fast-acquisition data was limited. After enhancement of the images, the number of long 3D fibres was closer to the number for the slow-acquisition data.

Future work will focus on the incorporation of the segmentation algorithm inside the CycleGAN to further improve the results. We will also study the use of a 3D version of the U-Net-id and apply the pipeline to a dataset with fibre breaks.

Source code and data repository

The codes utilised in this pipeline for the implementation of 3D CycleGAN and 2D U-Net-id can be accessed through the following publicly available repositories: <https://github.com/pvilla/3DCycleGAN> and <https://github.com/CMG-KULeuven/UnetID-FibreSeg>. The reconstructed fast- and slow-acquisition tomograms, as well as trained models for CycleGAN, are available at <https://zenodo.org/record/7632124>.

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

We acknowledge the Paul Scherrer Institut, Villigen, Switzerland for the provision of synchrotron radiation beam-time at the TOMCAT beamline X02DA of the SLS. R. Guo would like to thank F.H Wagner for the discussion of U-Net-id and acknowledge his PhD scholarship from China Scholarship Council (202006430010). C. Breite and M. Mehdikhani would like to acknowledge FWO Flanders for their postdoctoral fellowships, COCOMI (1231322N) and ToughImage (1263421N), respectively.

Supplementary materials

Appendix A. Hyperparameter configurations for CycleGAN

This section details the hyperparameters for CycleGAN, along with our hyperparameter choices for the algorithms presented in this paper.

The relevant hyperparameters for our four CycleGANs are displayed in Table A.1. ‘Training data shape’ refers to the size of the randomly sampled volumes from the training tomograms. Data shape and batch size were chosen to best fill the available memory of the GPU cluster consisting of 4 NVIDIA Tesla V100 GPUs with 4x 32 GB VRAM. The total training loss is computed using the weighted sum of the two GAN and two cycle consistency components with weights: i) λ_{gf} and λ_{df} for the GAN training of the slow to fast process, ii) λ_{gs} and λ_{ds} for the GAN training of the fast to slow process, and iii) λ_{cf} and λ_{cs} for the cycle-consistency constraints starting from the fast and slow domains, respectively. The learning rates for discriminators and generators determine the size of the optimisation steps.

Table A.1. Hyperparameter choices for model optimisation of 2D/3D CycleGAN enhancement algorithms.

Enhancement	Training data shape [voxels]		Batch size	Loss-weights			Learning rates
	Fast	Slow		λ_{gf}	λ_{df}	λ_{cf}	Discr.
				λ_{gs}	λ_{ds}	λ_{cs}	Gen.
2×-1ms- enhanced (2D)	1024×1024	16	10	1	50	1e-4	
	2048×2048		10	1	50	1e-3	
2×-1ms- enhanced (3D)	64×128×128	4	10	1	100	1e-5	
	128×256×256		10	1	100	1e-4	
4×-3ms- enhanced (3D)	64×64×64	4	10*	1*	100*	1e-4*	
	256×256×256		10*	1*	100*	1e-3*	
4×-0.5ms- enhanced (3D)	64×64×64	4	10*	1*	100*	1e-4*	
	256×256×256		10*	1*	100*	1e-3*	

* Initial hyperparameter choice. These hyperparameters were dynamically updated during optimization.

Appendix B. Quantitative evaluation of 2D and 3D CycleGAN image enhancement

In this appendix, we quantitatively evaluate the performance of 2D CycleGAN and 3D CycleGAN using the registered paired test dataset with the slow-acquisition data as a reference. The evaluation is based on Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). Higher PSNR values indicate better fidelity in the enhanced images, while higher SSIM values indicate a similar luminance, contrast, and structure. Additionally, we consider the average fibre count per slice since it is affected by hallucinated fibres.

Table B.1 presents the quantitative evaluation results for the same volume of size $401.76 \times 369.36 \times 152.28 \mu\text{m}^3$. The voxel size of the slow acquisition dataset matches that of the 3D CycleGAN enhancement. Nearest neighbour interpolation was employed for adjusting the voxel size in the 2D CycleGAN enhancement along one direction and for the fast acquisition dataset along three directions, ensuring equal voxel dimensions. We can see improvements over the original fast-acquisition image in the PSNR and SSIM for both the 2D and the 3D CycleGAN enhancement. However, both values suggest a significantly better image quality in the 3D enhancement. The most compelling argument for using the 3D algorithm over the 2D version is the fibre count though. In the fast acquisition image, we undercount the number of fibres by an average of 32 fibres per slice. In the 2D CycleGAN enhancement, we count 25 additional fibres, which can be attributed to the hallucinated fibres associated with 2D CycleGAN. In contrast, the number of fibres in the images enhanced by 3D CycleGAN shows a negligible deviation of 2 fibres compared to the slow-acquisition data.

Table B.1. Quantitative evaluation of image enhancement performance using 2D CycleGAN and 3D CycleGAN approaches for 1ms-fast-acquisition data

	PSNR	SSIM	Fibre count (all slices)
Slow-acquisition	-	-	687.43±2.01
Fast-acquisition	18.46	0.30	655.10±8.18
2D CycleGAN enhanced	19.74	0.55	712.03±6.24
3D CycleGAN enhanced	21.24	0.62	689.10±1.96

It should be noted that the image registration quality significantly influences the PSNR and SSIM, and the segmentation quality affects the fibre count.

Appendix C. Pixel-based evaluation of enhancement and segmentation accuracy

In this appendix, we provide pixel-based metrics to evaluate the enhancement and segmentation accuracy.

For enhancement performance evaluation, we employ PSNR and SSIM metrics. These measurements involve comparisons based on greyscale images and reference slow-acquisition datasets. To evaluate the segmentation accuracy, we calculate the Intersection over Union (IoU) with ground truth as a reference, which is manually corrected based on the slow-acquisition segmented data. The IoU value ranges from 0 to 1, with higher values signifying better segmentation results. Additionally, all fast-acquisition datasets are resized using the nearest neighbour interpolation to match the slow-acquisition data's size.

Table C.1. summarises the enhancement and segmentation evaluation values. This table shows that all enhancement datasets exhibit higher PSNR, SSIM and IoU values than the fast-acquisition datasets. This confirms that our enhancement algorithm successfully improves image quality and boosts segmentation performance. Notably, the slow-acquisition data has the highest IoU value, which is close to 1. This is because our ground truth is corrected based on the slow-acquisition segmented data, which has many completely overlapping regions.

The impact of acquisition time and voxel size on the enhancement performance is discussed. The 3ms-fast-acquisition data has the same voxel size as the 0.5ms-fast-acquisition data but with slower acquisition time, exhibiting higher PSNR, SSIM, and IoU values after being enhanced by our CycleGAN. Similarly, the 4×-3ms-enhanced data outperforms the 2×-1ms-enhanced data in terms of PSNR, SSIM, and IoU values, despite a larger voxel size. While the results of 2×-1ms-enhanced data exhibit similar PSNR and SSIM values compared to 4×-0.5ms-enhanced data, they present higher IoU values. These suggest that acquisition time affects our enhancement results more than voxel size. However, more thorough investigations would be required to precisely evaluate the effect

of time and spatial resolution, such as maintaining the same acquisition time but varying voxel sizes. Nonetheless, our developed CycleGAN can effectively improve image quality and segmentation quality for different time-resolved computed tomography, encompassing different temporal and spatial resolutions.

Table C.1. Comparative evaluation of enhancement and segmentation accuracy

	PSNR	SSIM	IoU
Ground-truth	-	-	-
slow-acquisition	-	-	0.99
1ms-fast-acquisition	18.46	0.30	0.30
2×-1ms-enhanced	21.24	0.62	0.81
3ms-fast-acquisition	17.85	0.28	0.27
4×-3ms-enhanced	22.95	0.67	0.84
0.5ms-fast-acquisition	18.32	0.25	0.20
4×-0.5ms-enhanced	21.95	0.62	0.77

All these pixel-based metrics are sensitive to the size, shape, and position of the segmented fibres and are therefore susceptible to the imperfection of image registration between fast, enhanced and reference tomograms, regardless of enhancement and segmentation performance.

Appendix D. 3D fibre count: A quantitative method for evaluating individual fibre segmentation

In this section, we demonstrate a new quantitative evaluation method called 3D fibre count for assessing the 3D individual fibre segmentation.

The 3D fibre count is the number of segmented fibres above a certain length threshold. The length of each fibre can be calculated after searching all slices for the nearest centre of each segmented fibre between adjacent slices within a given maximum query distance. A higher 3D fibre count indicates that more fibres across multiple slices have been successfully tracked and exceed the length threshold, suggesting a higher segmentation quality. In comparison, a lower count means more potential "broken fibres" due to segmentation errors. Increasing the maximum query distance can result in a higher 3D fibre count, but also means greater tolerance to segmentation errors.

The segmented images in Fig. D.1b-f are the possible segmentation errors, such as missing fibres, touching fibres, and unrealistically smaller, larger, and large-offset fibres, which are the reasons for the evaluation with IoU. By setting the maximum query distance to 1 pixel, these segmented errors can be identified as these segmented fibres cannot be tracked successfully, leading to a lower 3D fibre count. Therefore, this 3D fibre count can effectively evaluate 3D segmentation without manual labelling.

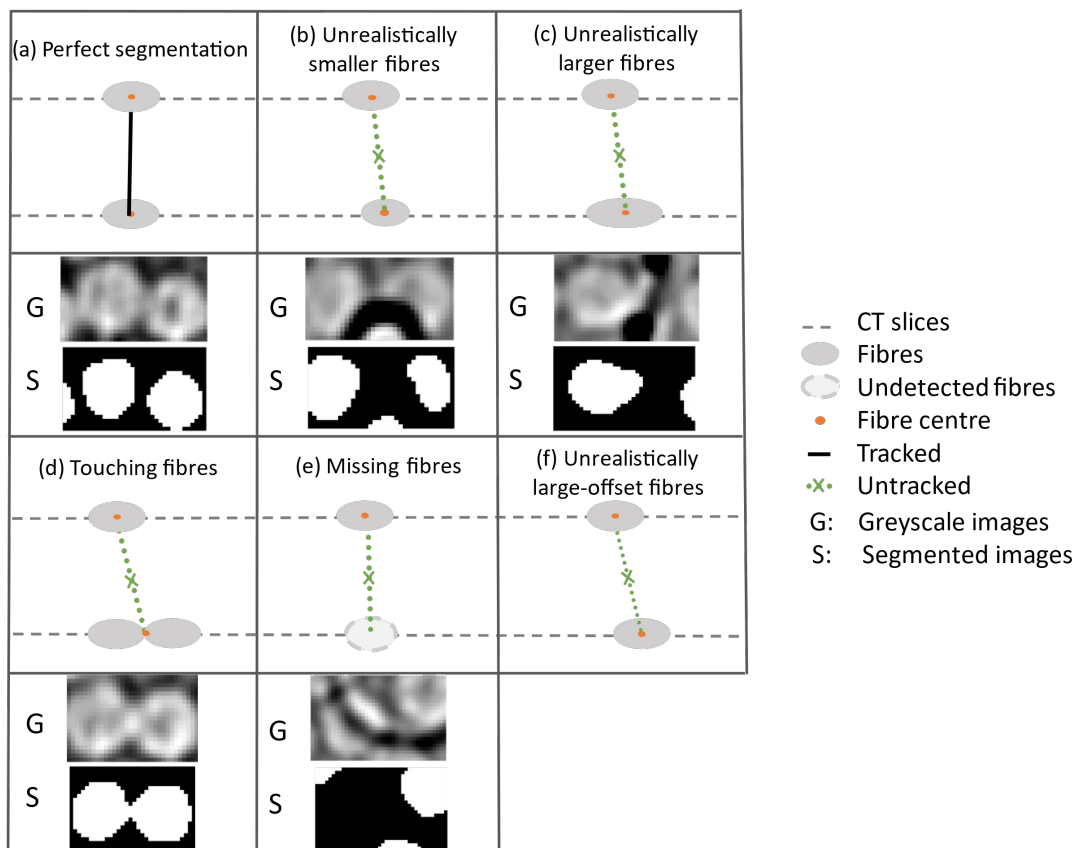


Fig. D.1 . A schematic depiction of 3D fibre tracking for different scenarios: (a) accurately segmented fibres, (b) unrealistically segmented smaller fibres, (c) unrealistically segmented larger fibres, (d) touching fibres, (e) missing fibres, and (f) unrealistically large-offset fibres. All of the greyscale and segmented images are for the slow-acquisition data, and some fibres were not successfully segmented, mainly due to artefacts. Consequently, cases (b-f) represent possible instances of unsuccessful segmentation, making it challenging to directly track these fibres using a maximum query distance of 1 pixel to find the nearest fibre centre for each pair of adjacent slices.

References

- [1] C. Breite, A. Melnikov, A. Turon, A.B. de Morais, C. Le Bourlot, E. Maire, E. Schoberl, F. Otero, F. Mesquita, I. Sinclair, J. Costa, J.A. Mayugo, J.M. Guerrero, L. Gorbatikh, L.N. McCartney, M. Hajikazemi, M. Mehdikhani, M.N. Mavrogordato, P.P. Camanho, R. Tavares, S.M. Spearing, S.V. Lomov, S. Pimenta, W. Van Paepegem, Y. Swolfs, Detailed experimental validation and benchmarking of six models for longitudinal tensile failure of unidirectional composites, *Compos Struct* 279 (2022) 114828.
- [2] M.W. Czabaj, M.L. Riccio, W.W. Whitacre, Numerical reconstruction of graphite/epoxy composite microstructure based on sub-micron resolution X-ray computed tomography, *Compos Sci Technol* 105 (2014) 174-182.
- [3] B.F. Sorensen, S. Goutianos, L.P. Mikkelsen, S. Faester, Fatigue damage growth and fatigue life of unidirectional composites, *Compos Sci Technol* 211 (2021) 108656.
- [4] N. Otsu, A Threshold Selection Method from Gray-Level Histograms, *IEEE Transactions on Systems, Man, and Cybernetics*, 1979, pp. 62-66.
- [5] M.J. Emerson, K.M. Jespersen, A.B. Dahl, K. Conradsen, L.P. Mikkelsen, Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials, *Compos Part a-Appl S* 97 (2017) 83-92.
- [6] F.S. Rasmussen, M.J. Emerson, Sonne, H. M.R., J.H., L.P. Mikkelsen, V.A. Dahl, Fiber segmentation from 3D X-ray computed tomography of composites with continuous textured glass fibre yarns, *International Conference on Tomography of Materials & Structures*, 2019.
- [7] F. Sosa-Rey, Y. Abderrafai, A.D. Lewis, D. Therriault, N. Piccirelli, M. Levesque, OpenFiberSeg: Open-source segmentation of individual fibers and porosity in tomographic scans of additively manufactured short fiber reinforced composites, *Compos Sci Technol* 226 (2022) 109497.
- [8] R. Helwing, D. Hulsbusch, F. Walther, Deep learning method for analysis and segmentation of fatigue damage in X-ray computed tomography data for fiber-reinforced polymers, *Compos Sci Technol* 230 (2022) 109781.

- [9] 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) 936.
- [10] 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 (2022) 110333.
- [11] A. Badran, D. Marshall, Z. Legault, R. Makovetsky, B. Provencher, N. Piche, M. Marsh, Automated segmentation of computed tomography images of fiber-reinforced composites by deep learning, *J Mater Sci* 55(34) (2020) 16273-16289.
- [12] S. Jegou, M. Drozdal, D. Vazquez, A. Romero, Y. Bengio, The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation, *Ieee Comput Soc Conf* (2017) 1175-1183.
- [13] Y. Al-Kofahi, A. Zaltsman, R. Graves, W. Marshall, M. Rusu, A deep learning-based algorithm for 2-D cell segmentation in microscopy images, *Bmc Bioinformatics* 19 (2018).
- [14] K.M. He, G. Gkioxari, P. Dollar, R. Girshick, Mask R-CNN, *Ieee I Conf Comp Vis* (2017) 2980-2988.
- [15] M. Bai, R. Urtasun, Deep watershed transform for instance segmentation, *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5221-5229.
- [16] A.O. Vuola, S.U. Akram, J. Kannala, Mask-RCNN and U-net ensemble for nuclei segmentation, 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), IEEE, 2019, pp. 208-212.
- [17] F.H. Wagner, R. Dalagnol, Y. Tarabalka, T.Y.F. Segantine, R. Thome, M.C.M. Hirye, U-Net-Id, an Instance Segmentation Model for Building Extraction from Satellite Images-Case Study in the Joanopolis City, Brazil, *Remote Sens-Basel* 12(10) (2020).
- [18] C.W. Tian, L.K. Fei, W.X. Zheng, Y. Xu, W.M. Zuo, C.W. Lin, Deep learning on image denoising: An overview, *Neural Networks* 131 (2020) 251-275.
- [19] Z.H. Wang, J. Chen, S.C.H. Hoi, Deep Learning for Image Super-Resolution: A Survey, *Ieee T Pattern Anal* 43(10) (2021) 3365-3387.
- [20] J. Hu, A Review of Super-Resolution Reconstruction Based on Supervised Learning, *International Conference on Artificial Intelligence and Advanced Manufacturing (AIAM)*, Dublin, Ireland, 2019, pp. 768-773.
- [21] J.Y. Zhu, T. Park, P. Isola, A.A. Efros, Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks, *Ieee I Conf Comp Vis* (2017) 2242-2251.
- [22] M. Stampanoni, A. Groso, A. Isenegger, G. Mikuljan, Q. Chen, A. Bertrand, S. Henein, R. Betemps, U. Frommherz, P. Bohler, D. Meister, M. Lange, R. Abela, Trends in synchrotron-based tomographic imaging: the SLS experience, *Proc Spie* 6318 (2006).
- [23] F. Marone, M. Stampanoni, Regridding reconstruction algorithm for real-time tomographic imaging, *J Synchrotron Radiat* 19 (2012) 1029-1037.
- [24] R. Mokso, C.M. Schlepütz, G. Theidel, H. Billich, E. Schmid, T. Celcer, G. Mikuljan, L. Sala, F. Marone, N. Schlumpf, M. Stampanoni, GigaFRoST: the gigabit fast readout system for tomography, *J Synchrotron Radiat* 24 (2017) 1250-1259.
- [25] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative Adversarial Networks, *arXiv* (2014).
- [26] V.A. Dahl, M.J. Emerson, C.H. Trinderup, A.B. Dahl, Content-based Propagation of User Markings for Interactive Segmentation of Patterned Images, 2020 Ieee/Cvf Conference on Computer Vision and Pattern Recognition Workshops (Cvprw 2020) (2020) 4280-4288.
- [27] github.com/vedranaa/insegtpy.
- [28] C. Shorten, T.M. Khoshgoftaar, A survey on Image Data Augmentation for Deep Learning, *J Big Data-Ger* 6(1) (2019).
- [29] C. Breite, *Aligning Fibre Break Models for Composites with the Observable Micro-Scale Material Behaviour*, KU Leuven, 2021.
- [30] H. Hofmann, H. Wickham, K. Kafadar, Letter-Value Plots: Boxplots for Large Data, *J Comput Graph Stat* 26(3) (2017) 469-477.
- [31] M.L. Waskom, seaborn: statistical data visualization, *Journal of Open Source Software* 6(60) (2021) 3021.