



Statistical Image Analysis of Tomograms with Application to Fibre Geometry Characterisation

Emerson, Monica Jane

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Emerson, M. J. (2018). *Statistical Image Analysis of Tomograms with Application to Fibre Geometry Characterisation*. DTU Compute. DTU Compute PHD-2017 No. 466

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

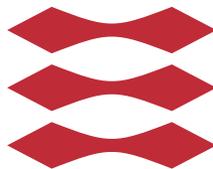
- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Statistical Image Analysis of Tomograms with Application to Fibre Geometry Characterisation

Monica Jane Emerson

DTU



Kongens Lyngby 2017

Technical University of Denmark
Department of Applied Mathematics and Computer Science
Richard Petersens Plads, building 324,
2800 Kongens Lyngby, Denmark
Phone +45 4525 3031
compute@compute.dtu.dk
www.compute.dtu.dk

Summary (English)

The goal of this thesis is to develop statistical image analysis tools to characterise the micro-structure of complex materials used in energy technologies, with a strong focus on fibre composites. These quantification tools are based on measuring geometrical parameters defining structures from 2D and 3D images, especially acquired through X-ray computed tomography (CT).

Fibre composites are extensively used in transportation and energy technologies such as wind turbines. It is of high importance to characterise composites accurately and to understand their behaviour under load to ensure efficiency and longevity of these technologies.

Imaging with X-ray CT has been the foundation of the thesis. This technique enables analysis in 3D and at the micro-scale, where individual fibres are distinguishable. By combining ultra-fast X-ray CT and in-situ loading environments, it is possible to image composites with high spatial and temporal resolution to capture the small and sudden micro-structural changes caused by loading.

This thesis demonstrates that statistical image analysis combined with X-ray CT opens up numerous possibilities for understanding the behaviour of fibre composites under real life conditions. Besides enabling characterisation of material properties, estimating individual fibre centre lines and diameters allows the quantification of small micro-structural changes with a high degree of accuracy, as it is possible to follow how each individual fibre changes across data-sets acquired under progressive loading conditions. Finally, the thesis demonstrates the precision to which fibre geometry can be characterised through X-ray CT and the developed data analysis tools.

Summary (Danish)

Målet med dette projekt er at udvikle statistiske billedanalyseteknikker, som kan anvendes til at karakterisere mikrostrukturer i komplekse materialer, med særlig fokus på fiberkompositter som ofte anvendes i energisektoren. Disse kvantificeringsværktøjer anvender geometriske parametre til at beskrive strukturer i 2D, 3D og 4D, baseret på røntgentomografidata.

Fiberkompositter anvendes i vid udstrækning i transport- og energiteknologier eksempelvis til vindenergi. Karakterisering af fiberkompositter er essentielt til at forstå deres opførsel under belastning og for at sikre effektivitet og levetid af materialerne.

Imaging med røntgentomografi er grundlaget for denne afhandling. Røntgentomografi tillader analyse på mikroskala i 3D i en opløsning, hvor individuelle fibre kan adskilles. Desuden gør ultrahurtig røntgen-CT det muligt at afbilde i høj tidlig og rumlig opløsning, og herved kan mikrostrukturelle ændringer observeres i materialer, der belastes in-situ.

Denne afhandling demonstrerer, at statistisk billedanalyse kombineret med røntgen-CT åbner adskillige muligheder for at forstå opførslen af fiberkompositter under realistiske forhold. I tillæg til karakterisering af materialeegenskaber tillader målingen af individuelle fibercentre og -diametre karakterisering af mikrostrukturelle ændringer med høj præcision, da det er muligt at følge ændringer i fibre mellem datasæt optaget under øget belastning. Endeligt demonstrerer afhandlingen højpræcisionsmålinger af fibergeometri, som kan opnås gennem karakterisering med brug af røntgen-CT og de udviklede analyseværktøjer.

Preface

This thesis was prepared at the Department of Applied Mathematics and Computer Science (DTU Compute) of the Technical University of Denmark in fulfilment of the requirements for acquiring a Doctor of Philosophy (PhD) degree in Applied Mathematics and Computer Science, with emphasis on Statistical Image Analysis.

A third of the funding originated from DTU Compute Graduate School and two-thirds were funded by the alliance for imaging and modelling of energy applications (CINEMA), a research project granted by Innovation Fund Denmark with the purpose of investigating the micro-structure of materials used in energy technologies. The thesis work was supervised by Associate Professor Anders BJORHOLM DAHL and co-supervised by Professor KNUT CONRADSEN. Assistant Professor VEDRANA ANDERSEN DAHL also contributed in high degree to the development of the methodology.

The research presented in this thesis includes methods for extraction and quantification of structures from images, mainly 3D and acquired through X-ray micro-tomography. The work was carried out in close collaboration with partners from the CINEMA alliance. The advanced statistical image analysis tools developed here can provide insight that will aid the understanding of the relationship between the micro-structure of a material and its macroscopic behaviour under real life conditions.

The thesis consists of an introduction, a description of the theoretical background, a discussion on the major scientific contributions and five original manuscripts prepared during the course of the project.

Lyngby, 11-December-2017

A handwritten signature in black ink. The name 'Monica' is written in a large, cursive script. Below it, the name 'Emerson' is written in a smaller, more compact cursive script. A long horizontal line extends from the end of the signature to the right.

Monica Jane Emerson

Acknowledgements

I would like to thank all the people that have surrounded me during the course of my PhD studies. I have enjoyed very much my life during these three years, travelled a lot, met amazing people and most importantly developed to be a better and wiser person.

I have very much enjoyed learning about the Danish culture and working at DTU Compute, with all the knowledge sharing tradition and friendliness amongst colleagues. Very nice technical and personal discussions have arisen, and I would like to especially thank my office mates (Tuan, Dolores, David, Florian, etc.) for the great working environment and support. I am also grateful to the old colleagues that became friends (Hildur, Sophie, Anne-Sofie, Federica, Alexandra, etc.) and to the new people that bring energy and joy to the group, like Mariam, Wail and Shihav.

A very special thanks to my supervisors for all the valuable advice. It has been a great pleasure to be supervised by Anders, young, ambitious and motivating, and Knut, with his immense expertise and wisdom, always open for personal counselling as well. Last but not least, without being an official supervisor, Vedrana has dedicated plenty of time in helping me develop this thesis, even during tough and busy moments which turned out to be very productive.

Thanks to my collaborators Lars, Kristine, Ying and Philip, for making my work meaningful by solving real world problems. I am also very grateful to everyone that contributed to the quality of this thesis with proofreading.

Finally, a very very special thanks to my biggest support in Denmark, Ugne, and

to my family for the encouragement all the way through. My dad, as always, an excellent asset in any team. He is able to take any role and has this incredible ability to immerse into a case and identify the problems so as to find feasible solutions. I hope I have inherited that!

Well, I should also thank my flatmates for looking after me while I was busy and sick at the end of the thesis. A very special thanks to Stefanos for making me feel like I have family also at my Danish home.

Contributions

Papers included in this thesis

- A** Emerson, M. J., Jespersen, K. M., Jørgensen, P. S., Larsen, R. & Dahl, A. B. (2015). Dictionary based segmentation in volumes. In *Scandinavian Conference on Image Analysis* (pp. 504-515). Springer, Cham. (Published)
- B** Emerson, M. J., Jespersen, K. M., Dahl, A. B., Conradsen, K. & Mikkelsen, L. P. (2017). Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials. *Composites Part A: Applied Science and Manufacturing*, 97, 83-92. (Published)
- C** Dahl, V. A., Trinderup, C. H., Emerson, M. J., & Dahl, A. B. (2017). Content-based Propagation of User Markings for Interactive Segmentation of Patterned Images. *IEEE Transactions on Image Processing*. (To be submitted)
- D** Emerson, M. J., Dahl, V. A., Conradsen, K., Dahl A. B. & Mikkelsen L. P. (2017). Statistical validation of individual fibre segmentation from tomograms and microscopy. *Composites Science and Technology*. (To be submitted)
- E** Emerson, M. J., Wang, Y., Withers, P. J., Conradsen, K., Dahl, A. B. & Dahl, V. A. (2017). Evolution in Fibre Orientation during Axial Compression of a Composite through Time-lapse X-ray Imaging and Fibre Tracking. *Scientific Reports*. (In submission)

Other conference and journal contributions by the author

- Emerson, M. J., Dahl, V. A., Mikkelsen, L. P., Dahl, A. B. & Conradsen, K. (2017). Geometrical characterisation of individual fibres from X-ray tomograms. In *Proceedings of 30th Nordic Seminar on Computational Mechanics*. (Abstract and talk)
- Emerson, M. J., Dahl, A. B., Dahl, V. A., Conradsen, K. & Mikkelsen, L. P. New approach for validating the segmentation of 3D data applied to individual fibre extraction. (2017). In *3rd International Conference on Tomography of Materials and Structures*. (Extended Abstract and talk)
- Emerson, M. J., Wang, Y., Jespersen, K. M., Mikkelsen, L. P., Withers, P. J., Conradsen, K., Dahl V. A. & Dahl, A. B. (2017). Unidirectional fibre composite characterisation from X-ray tomography. In *TMS 2017*. (Poster)
- Mikkelsen, L. P., Emerson, M. J., Jespersen, K. M., Dahl, V. A., Conradsen, K. & Dahl, A. B. (2016). X-ray based micromechanical finite element modeling of composite materials. In *Proceedings of 29th Nordic Seminar on Computational Mechanics*. (Extended abstract)
- Emerson, M. J., Jespersen, K. M., Dahl, A. B., Conradsen, K. & Mikkelsen, L. P. (2016). Segmentation of individual fibres in a uni-directional composite from 3D X-ray computed tomography data. In *3rd International Congress on 3D Materials Science*. (Poster)
- Einarsdóttir, H., Emerson, M. J., Clemmensen, L. H., Scherer, K., Willer, K., Bech, M., R. Larsen, B. K. Ersbøll & Pfeiffer, F. (2016). Novelty detection of foreign objects in food using multi-modal X-ray imaging. *Food Control*, 67, 39-47.

Abbreviations

2D Two-dimensional.

3D Three-dimensional.

ANOVA Analysis of variance.

ASC Ammonia slip catalyst.

CFRP Carbon fibre reinforced polymer.

CINEMA The alliance for imaging and modelling of energy applications.

CT Computed tomography.

DTU Danmarks Tekniske Universitet.

ESRF European Synchrotron Research Facility.

FBP Filtered back projection.

FEM Finite element modelling.

FoV Field of view.

FPR False positive rate.

FVF Fibre volume fraction.

GFRP Glass fibre reinforced polymer.

GLM General linear model.

GUI Graphical user interface.

KU Københavns Universitet.

NBI Niels Bohr Institute.

OM Optical microscopy.

PCA Principal component analysis.

PDF Probability density function.

RoI Region of interest.

SEM Scanning electron microscopy.

SOFC Solid oxide fuel cell.

TPR True positive rate.

UD Unidirectional.

Contents

Summary (English)	i
Summary (Danish)	iii
Preface	v
Acknowledgements	vii
Contributions	ix
Abbreviations	xi
1 Introduction	1
1.1 Scope of the Project	1
1.2 X-ray Imaging of a Material's Micro-structure	3
1.3 Materials for Energy Technologies	4
1.3.1 Characterisation of Fibre Composites	5
1.3.2 Understanding Damage in Fibre Composites under Load	7
1.4 Thesis Objectives	8
1.5 Thesis Overview	9
I Methodology	11
Introduction to Methodology	13
2 Imaging Modalities	15
2.1 X-ray Imaging	15
2.1.1 Synchrotron Light	16

2.1.2	Computed Tomography	18
2.1.3	Noise and Artefacts	18
2.2	Optical Microscopy	20
2.3	Scanning Electron Microscopy	20
3	Dictionary-based Probabilistic Segmentation	21
3.1	Motivation and Background	22
3.2	The Basic Algorithm	29
3.2.1	The Training Process	29
3.2.2	The Classification Process	31
3.2.3	Tuning the Parameters	32
3.3	Enumeration of Variants	36
3.3.1	Calculation of Cluster Centres	36
3.3.2	Implementation of Cluster Search	37
3.3.3	Description of the Relationship between Image and Dic- tionary	37
3.3.4	Implementation of Image-dictionary Mapping	38
3.3.5	Amount of Input Required for Training	41
3.3.6	Features for Clustering	41
3.3.7	Dimensionality	42
4	Extraction of Geometrical Parameters Defining Structures	45
4.1	Centre Point Estimation	46
4.1.1	Thresholding-based Methods	46
4.1.2	Blob Detection	50
4.1.3	Circular Hough Transform	53
4.2	Diameter Estimation	53
5	Statistical Characterisation of Structure Geometry	55
5.1	Exploratory Data Analysis	55
5.1.1	Summary Statistics	55
5.1.2	Empirical Probability Density Functions	55
5.1.3	Spatial Distributions	57
5.2	General Linear Models	58
5.2.1	Linear and Quadratic Dependence	58
5.2.2	Analysis of Variance	58
II	Analysis Pipelines and Applications	61
6	Fibre Extraction Pipeline	63
6.1	Basic Pipeline	64
6.1.1	Detection of Fibre Cross-sections	64
6.1.2	Detection of Fibres Trajectories	66

6.2	Improvements and Extensions	68
7	Applications of the Fibre Pipeline	73
7.1	Characterisation of Real Structures	73
7.1.1	Fibre Orientation	74
7.1.2	Image-based Micro-mechanical Models	77
7.2	Evolution of Structures under Changing Load	78
8	The 3D Validation Challenge	83
8.1	Initial Validation of Geometry Extraction	84
8.2	Statistical Validation of Analysis Methods and Imaging Modalities	84
	Conclusions	90
	Bibliography	93
III	Publications	99

Introduction

1.1 Scope of the Project

The focus of this thesis is on developing statistical image analysis pipelines primarily to characterise fibre geometry at the micro-scale from X-ray computed tomography (CT) images. The main tasks are geometry extraction and characterisation, illustrated in Figure 1.1.

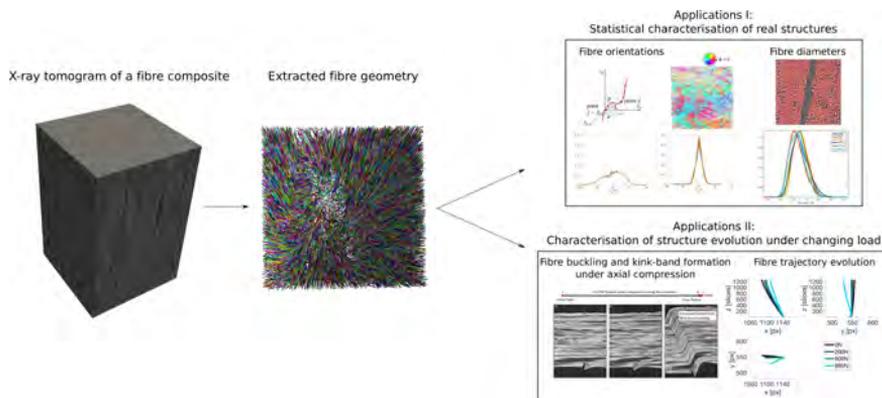


Figure 1.1: Statistical image analysis for fibre geometry characterisation.

The work presented in this thesis is part of the alliance for imaging and modelling of energy applications (CINEMA), a research project granted by Innovation Fund Denmark [DSF-Grant No. 1305-00032B]. The CINEMA project is a five year multi-disciplinary project comprising various industrial and academic partners with different scientific backgrounds including material scientists, mechanical and software engineers, chemists, physicists, mathematicians and data scientists. The academic partners are the departments of Physics, Energy, Compute and Wind Energy from Danmarks Tekniske Universitet (DTU); the Chemistry department and Niels Bohr Institute (NBI) from Københavns Universitet (KU); Manchester X-ray Imaging Facility; Northwestern University and the MAX IV Laboratory. The industrial partners are Haldor Topsøe, Amminex Emissions Technology, LM Wind Power, Xnovo Technology and Rockwool.

CINEMA's aims are to develop methods for characterising the three-dimensional (3D) micro-structure of complex materials employed in energy technologies under realistic conditions and in real time. This will enable improved understanding of their macroscopic behaviour, as the performance of such materials is strongly related to their multi-scale structure right down to the micro- and nano-scales.

The two major topics under investigation are flow and damage. Flow properties, relevant for porous materials, are of high importance in determining the performance of, for example, catalysts or fuel cells. Damage induced structural changes, such as crack or fracture nucleation and propagation, are relevant in determining the lifetime of, for instance, wind turbine blades or fuel cells.

The ambition of CINEMA is to cover the complete pipeline from experimental design to characterisation of the micro-structure. The pipeline involves a set of stages where different abilities are required and that is why a lot of the work presented in this thesis has been developed in close collaboration with CINEMA partners. The stages in this pipeline are experimental design, image acquisition, image reconstruction, image segmentation and structural modelling. It should be noted that modelling can refer to either finite element modelling (FEM) for simulation or statistical modelling for quantification.

This thesis covers the stages of image segmentation and statistical modelling necessary for the extraction of 3D structures from the tomographic reconstructions, and the posterior quantification of the extracted structures. Out of the broad range of complex materials used in energy technologies, this thesis focuses on unidirectional (UD) fibre composites, and briefly touches upon solid oxide fuel cells (SOFCs) and ammonia slip catalysts (ASCs).

1.2 X-ray Imaging of a Material's Micro-structure

X-ray imaging, discovered more than a century ago (Röntgen, 1898), enables inspection of a sample's interior in a non-destructive manner thanks to the penetration ability of X-rays. While X-ray imaging is mostly known from medical radiography, it is now widely used for both commercial and research purposes. Examples include security screening, in-line defect and quality control in the manufacturing industry or engineering of new materials and products.

Absorption X-ray imaging, which will be explained in more detail in Chapter 2, is the most common modality obtained with X-rays. As it is based on the attenuation properties of the sample, it is appropriate for differentiating materials that are significantly different in density. As aforementioned, the thesis work has mainly focused on characterising the structure of fibre composites at the micro-scale. The first step in analysing the X-ray scans is to separate fibres from the other material phases in the composite, referred to as background.

To segment individual fibres these need to be somewhat resolved in the scan. Thus, an adequate spatial resolution and a good enough contrast between fibres and background is necessary. The contrast between fibres and background is mostly determined by the difference in density between the two. The fibre phase can be represented by natural or human-made fibres. Examples of natural fibres include hemp, banana, cotton, wood and minerals. With regards to human-made fibres, these can be semi- or fully synthetic including Kevlar[®], carbon, glass, metal and plastic fibres. We considered glass and carbon, which have densities of 2.6 g/cm^3 and 1.8 g/cm^3 respectively. As to the background, it is usually comprised of a polymer and possibly manufacturing defects such as air bubbles or gases from the resin curing process. The polymer is commonly epoxy, which has a density of 1.2 g/cm^3 ; three orders in magnitude larger than the density of air. The difference in density and chemical composition between carbon fibres and epoxy is such that these can be challenging to segment, especially when the quality of the scan is low. Concerning the spatial resolution of the scan, the composites must be imaged at the micro-scale so that individual fibres can be resolved, as we have been looking into fibres with average diameters ranging from $7 \mu\text{m}$ to $17 \mu\text{m}$. Thus, X-ray micro-CT, which refers to 3D absorption X-ray imaging at the micro-scale, is an adequate imaging modality for investigating composites (Pyrz, 1999) at a resolution that allows segmentation of individual fibres (Requena et al., 2009).

In recent years, advances in the resolution that can be obtained using X-rays have enabled inspection of materials down to the micro- and nano-scale even in laboratory set-ups. In the past, such resolutions could only be obtained at large scale facilities such as synchrotrons (Stock, 1999), which are difficult to

access because of the cost or waiting time and the short allocated beamtime. The possibility of using laboratory set-ups for inspecting the 3D micro-structure of materials in a non-destructive manner means that the number of applications is growing and, consequently, also the amount of data that is being acquired. Moreover, advances in time-resolved X-ray imaging mean that in-situ studies to examine small and sudden changes in the structure under certain loading conditions have become possible, and add to the ever growing amount of data acquired. Although X-ray CT is transforming from a qualitative diagnostic tool to a quantitative one (Maire and Withers, 2014), a lot of this data is still being inspected visually because the current image analysis software - ImageJ (Schneider et al., 2012) and the WEKA tool (Hall et al., 2009), Thermo Scientific™ Amira-Avizo or VGStudio Max (Volume Graphics GmbH) - cannot provide a solution to every problem. The ever increasing volumes of data mean that it is not feasible to inspect this data manually, due to time constraints. Moreover, manual analysis introduces human errors which can create statistical biases, and is anyway often qualitative in its nature, as it is not practically possible, for example, to count the amount of particles or fibres even if examining a small volume. Therefore, there is a need for image analysis methods and pipelines that can provide quantitative information in a reasonable amount of time and in a statistically meaningful way, and that are appropriate for the problem under study, hence the need for tailored and/or flexible solutions.

1.3 Materials for Energy Technologies

Denmark wants all energy consumption to be based on renewable sources by 2050 (Government, 2013). To achieve this, it is important to develop renewable energy technologies to ensure both efficiency and an extended lifetime. As was mentioned above, characterising the 3D micro-structure of complex materials used in energy technologies, and its changes under certain real life conditions (e.g. changes in temperature, tension or compression stresses), can be essential for understanding the properties of these materials.

A big contribution of this thesis is a pipeline that can measure fibre geometry from tomograms of UD composites by segmenting individual fibres. The pipeline has been applied to investigate glass and carbon fibres, but could also be applicable to other synthetic or natural fibre-based materials. A large amount of studies on fibre composites rely on artificially generated fibre arrangements (Budiansky and Fleck, 1993; Fleck and Shu, 1995; Kyriakides and Arseculeratne, 1995; Jensen and Christoffersen, 1997), but in real composites there is a significant degree of randomness. The individual fibre segmentation pipeline combined with X-ray imaging opens up the possibility of studying real fibre arrangements. This

can provide information to optimise manufacturing processes and ultimately accelerate the development of cost-efficient components based on fibre composites.

An essential part of reducing the cost is in providing precise numerical and analytical models that can predict the macroscopic properties and behaviour of fibre composites based on their micro-structure. Existing models can be now tested over real fibre architectures with the aim of validating the predictions of these models. A higher confidence in these models can reduce the need of over-engineering components and performing expensive destructive mechanical testing. Additionally, the fibre extraction pipeline can be employed to follow and investigate the changes in fibrous materials under progressive loading conditions. Besides contributing to the understanding of damage mechanisms and their evolution from localised structural damage to catastrophic failure, valuable insight can be obtained on the precursors to the very complex damage mechanisms that affect fibre composites under real life working conditions.

Several case studies have been carried out in collaboration with CINEMA partners. Together with DTU Wind Energy we are working on characterising individual fibre orientations and diameters. Together with material scientists based at Manchester X-ray Imaging Facility we are observing, measuring and modelling the changes in fibre orientations and curvature caused by static loading in compression. We are also starting to look into the relationship between fibre geometry and fibre fracture nucleation and propagation under static loading in tension.

Last but not least, it should be noted that the dictionary-based probabilistic segmentation framework that will be presented in Chapter 3 has also been applied to segment material phases in phantom data of solid oxide fuel cells (SOFCs) and individual particles in ammonia slip catalysts (ASCs). This segmentation methodology combined with mathematical modelling and statistical analysis could potentially be applied to characterise other materials as well.

1.3.1 Characterisation of Fibre Composites

An improved understanding of fibre composites will enable the development of longer and stronger turbine blades, which will in turn allow for more cost-effective wind turbines and ultimately energy. In theory, the power output of a wind turbine is proportional to the blade length squared (Burton et al., 2001). The load carrying parts of long wind turbine blades are mainly made of UD fibre reinforced polymers where most of the fibres are aligned with the primary loading direction. Together with DTU Wind Energy we have investigated the micro-structure of these fibre composites. Understanding the micro-structure

is important because it is highly related to the macroscopic behaviour of these materials.

The key properties the fibre composite should have when designing long blades are high fatigue resistance, low weight, high stiffness and high compression strength (Brøndsted et al., 2005). Fatigue resistance is crucial when designing long lasting blades because the blades should withstand repeated loads during a long period of time. As it is a key factor limiting the design of long blades, a great deal of attention is being paid to fatigue damage evolution (Jespersen et al., 2016; Jespersen and Mikkelsen, 2017). The interactive segmentation tool in Paper C is able to segment fibre fractures with just a little bit of manual input, reducing the amount of time required for counting and determining the location of fibre fractures.

Our main collaboration with DTU Wind Energy has focused on characterising individual fibre orientations employing the fibre geometry extraction pipeline. It is in quantifying the stiffness and compression strength properties where fibre orientations are most relevant. While carbon fibres are both stiffer and lighter than glass fibres, they are much more expensive. The stiffness of the considered composites in the fibre direction is related to the fibre material and also to the alignment of the fibres inside the composite (Krenchel, 1964). As to strength, carbon fibres have a significantly higher tensile strength but only a slightly higher compression strength than glass fibres. In compression, the matrix phase of the composite plays an important supporting role by keeping the fibres straight while they carry the load. For this reason, the matrix material strongly determines the compression strength together with fibre misalignment, which influences the matrix ability of keeping the fibres straight because it introduces shear stresses at the angled fibre segments (Budiansky and Fleck, 1993). Hybrid materials that combine both glass and carbon fibres are being investigated to achieve a solution that provides a good compromise between stiffness and strength compared to price.

The production process of the fibre bundles and the manufacturing technique chosen for building up the composite have a high impact on the composite's performance. On the one hand, the fibre production process determines the distribution of fibre diameters inside a bundle. On the other hand, the manufacturing technique of the composite establishes the bundle arrangements and determines the alignment of individual fibres inside the bundles.

Various analytical and numerical models have been proposed to relate the properties of the material to its micro-structure. For example, the formula by Krenchel (Krenchel, 1964) relates fibre misalignment to stiffness whereas Budiansky's (Budiansky and Fleck, 1993) relates fibre misalignment to compression strength. These models are simplifications of much more complex relationships.

Although Krenchel's formula can provide quite accurate stiffness predictions, Budiansky's obviates a very important factor. That is the spatial distribution of the fibre orientations, which could notably influence the compression strength of the composite (Liu et al., 2004).

The pipeline for fibre geometry extraction has been combined with mathematical modelling and statistical analysis techniques to study the spatial distribution of individual fibre orientations and diameters at the micro-scale. Material property estimates can be obtained from these and compared to specifications provided by the manufacturer or to properties measured through, for example, mechanical testing. Additionally, the extracted fibre geometry can be imported into the FEM software ABAQUS which, besides predicting material properties, can perform simulations on how loading conditions would affect the experimentally measured image-based geometry.

To sum up, the access to individual fibres will enable a better understanding of how the arrangement in space of fibre orientations and diameters may influence the behaviour of the material, opening up the opportunity for incorporating spatial distributions in models that relate micro-structure with physical properties. In other words, understanding better how fibre orientations and diameters will influence the material's properties can accelerate the development and validation of numerical and analytical models that describe the composites behaviour.

1.3.2 Understanding Damage in Fibre Composites under Load

Damage mechanisms in composites are very complex. Following the changes in the micro-structure of fibre composites under progressive loading conditions will reveal the precursors to damage. Together with the School of Materials of the University of Manchester we are investigating the formation of kink-bands, the most common type of damage found in fibre reinforced polymers under axial compression (Dow and Gruntfest, 1960) - compression in the direction parallel to the primary fibre direction. Structural damage caused by static tensile loading is also of great interest for the composite researchers at Manchester.

Investigating the geometry of the final kink-band damage provides little clues on what gives rise to its formation. Therefore, it is of interest to study the fibre bending at the micro-scale, named fibre micro-buckling or kinking, which happens before the fibres break and form the kink band (Moran et al., 1995; Guynn et al., 1992). An experiment that combines time-lapse X-ray CT with in-situ loading in compression was designed by our collaborators to image a composite while loading it in axial compression until failure (Wang et al., 2017).

The hope was to capture the onset and progression of fibre micro-buckling in the steps leading up to kink-band formation, as well as the precursors to fibre micro-buckling. Although many researchers believe that initial fibre misalignment and waviness play a big role (Soutis, 1997; Wisnom and Atkinson, 2000), together with the constraints imposed by the matrix with regards to lateral movement of fibres (Pimenta et al., 2009), the initiation mechanism for fibre micro-buckling is still a subject of debate. The methodology developed for this collaboration is able to extract 3D fibre geometry at successive loading steps with fibre to fibre correspondence. Following the changes that each individual fibre undergoes as a consequence of loading enables the characterisation of small 3D micro-structural changes in the fibre architecture.

Predicting fibre failure under static tensile loading is also fundamental. Although different models have been developed towards this aim (analytical, statistical, numerical, etc), this is still a challenge. Combining physical-based observations (e.g. in-situ tensile experiments and X-ray CT) with statistical image analysis based on the image-measured geometry, can lead to insights into the relationship between fibre fractures and the as-manufactured fibre geometry. More specifically, the ultimate goal here is to understand how spatial location, alignment and diameters of individual fibres influence fracture nucleation and propagation. While it is expected that fibre orientations will not have a large influence on the tensile behaviour, the local arrangement of the fibres (e.g. fibre packed regions and fibres belonging to resin rich regions) could be crucial. In addition, the spatial distribution of fibre breaks and their progressive accumulation can provide a means to appreciate the influence of the stress concentration around a broken fibre and on how the load is redistributed in the surrounding unbroken fibres.

1.4 Thesis Objectives

The main thesis objectives are driven by the requirements of the partners inside the CINEMA alliance. These include:

- Demonstrating that X-ray tomography combined with advanced statistical image analysis can be employed to gain understanding about materials used in energy technologies.
- Developing analytical tools to extract parameters defining geometrical structures. This involves two- and three-dimensional image segmentation, feature detection and tracking, mainly applied to measure fibre geometry,

but also briefly touching upon other energy materials such as solid oxide fuel cells (SOFCs) and ammonia slip catalysts (ASCs).

- Measuring, quantifying and modelling in a statistically meaningful manner the geometry of fibre arrangements at the micro-scale, as well as the changes that the geometry undergoes under real life conditions.
- Validating quantitatively the accuracy of using X-ray tomography and our individual fibre segmentation pipeline for characterising the geometry of fibre composites.
- Providing a software toolbox which material scientists and mechanical engineers can apply to analyse their own data-sets.

1.5 Thesis Overview

The thesis is divided into three main parts: I. Methodology, II. Analysis Pipelines and Applications, and III. Publications.

Part I addresses the methodology that has been used in this thesis. Chapter 2 provides some background on the imaging modalities employed for acquiring the data analysed in this thesis work. Chapter 3 covers the dictionary-based probabilistic segmentation framework, which yields probability maps containing the likelihood of each pixel belonging to a set of pre-defined classes. Further processing is required to extract geometry in the form of coordinates and boundaries. Chapter 4 deals with geometry extraction, which can be performed directly over the acquired intensity image or from its probabilistic segmentation. Finally, Chapter 5 describes the statistical methodology applied for inspecting the parameters we use for characterising fibre geometry and its changes, together with the statistical models applied for validating the accuracy of the developed data quantification pipelines.

Part II serves as a link between the methodology presented in Part I and the contributions included in Part III. Chapter 6 explains the different blocks of the basic pipeline developed for fibre geometry extraction and the extensions and improvements it has gone through during the course of the thesis work. Chapter 7 discusses some of the applications of the pipeline, i.e. characterising fibre geometry through fibre orientations and diameters, and characterising the changes in individual fibre orientations under loading of the sample. Chapter 8 covers the topic of validation, where different data quantification pipelines are compared to assess the accuracy in measuring the real physical data.

Part III contains the main articles that have come out from the work of these three years.

Part I

Methodology

Introduction to Methodology

To harvest as much information as possible when investigating a material's micro-structure through imaging, it is essential to develop image analysis pipelines that are able to go all the way from data acquisition to data quantification and statistical modelling. These pipelines are comprised of several building blocks, which are based on the methodology that will be explained in this part of the thesis.

Imaging Modalities

This chapter describes the imaging modalities that have been employed to capture the data analysed in this thesis. As the main focus has been on X-ray CT, an explanation on the basics of absorption X-ray imaging is provided; together with a description of synchrotron facilities. The chapter continues by explaining the basics of CT and discussing the main quality deteriorating mechanisms in X-ray imaging. A brief explanation on scanning electron microscopy (SEM) and optical microscopy (OM) is also included, as some of the thesis work has also touched upon data acquired with these two-dimensional (2D) modalities.

2.1 X-ray Imaging

Absorption X-ray imaging is the most common modality for X-ray systems. A typical set-up consists of a source, a sample and a detector. In a laboratory set-up, the source will be an X-ray tube (or gun), whereas in a large scale facility the source will be synchrotron light. The source generates an X-ray beam with a certain intensity and energy profile (spectrum), and the detector measures the remaining X-ray signal after the beam has been attenuated by the sample and potential filters that lie in between source and detector.

Synchrotron radiation will be explained later in Section 2.1.1, so the focus here is on describing how X-ray tubes used in laboratory set-ups generate X-ray radiation. X-ray tubes consist of a filament-type cathode and a metallic target called anode (Karunakaran and Jayas, 2014) and are subjected to an accelerating voltage (kV) and a current (mA). When the filament is heated, electrons are generated by the cathode. These electrons are accelerated in vacuum and, when they strike the anode (typically made of Tungsten), X-ray photons are emitted. While the current determines the amount of electrons that hit the anode, and thus the number of photons emitted by the tube, the voltage determines the speed at which the electrons hit the anode, which will affect both the energy and the number of photons produced by the tube. The amount of photons is linear with respect to the amount of electrons that strike the anode whereas it is cubic with respect to their speed. Another important characteristic of the tube is the focal spot of the anode, an area defined by the size and shape of the electron beam when it hits the anode (Curry et al., 1990). The smaller the area of the focal spot, the finer the details that can be observed in the X-ray image.

As aforementioned, an X-ray beam is characterised by its energy and intensity I . The energy range is shaped by the target anode material and the voltage of the tube whereas the intensity, regulated by the number of photons emitted by the tube, is determined by the voltage and the current. When the X-ray beam passes through the sample, its intensity is reduced because of the interaction of the photons with matter. This loss is known as attenuation, which is caused by all phenomena removing photons from the X-ray beam path (absorption, scattering, etc.). Higher energy photons travel across the sample more easily than lower energy photons and that is why it is said that they have greater penetration power. The amount of radiation that passes through a material is given by the Beer-Lambert law,

$$I_{\text{out}} = I_{\text{in}}e^{-(\mu l)} \quad (2.1)$$

where μ [m^{-1}] is the linear attenuation coefficient and l [m] is the thickness of the material. The linear attenuation coefficient of a material μ depends on the energy of the X-rays and on the material's atomic number Z and density d .

2.1.1 Synchrotron Light

Synchrotron radiation is produced at large scale facilities which are able to provide a high flux and highly collimated beam that, compared to laboratory sources, is extremely coherent both in time and space (Huang, 2013). This allows for obtaining high quality images with good contrast, high signal-to-

noise ratio and fine spatial resolution, in much shorter time spans than what can be obtained in laboratory set-ups. Synchrotron facilities are highly superior to laboratory set-ups not only in terms of the source, but also for other reasons. These include the access to a broader range of imaging modalities and a greater deal of flexibility when designing experimental set-ups, for instance, with regards to the sample and its environment. It should be mentioned that I was fortunate enough to spend three months of my PhD studies at ID19, the micro-tomography beamline of the European Synchrotron Research Facility (ESRF).

A synchrotron source consists of a linear accelerator (Linac), a booster and a storage ring. Electrons are produced by an electron gun and accelerated by the Linac until they almost reach the speed of light. Once the electrons have entered the booster, they circulate inside this circumference until they have acquired a high enough energy to be inserted in the storage ring. The storage ring, which is actually a polygon, is made of straight sections and bending sections where the electrons run in almost perfect vacuum. Straight sections host focusing magnets to ensure that electrons remain close to their orbital path. On the other hand, bending sections host large dipole electromagnets, called bending magnets, that are used to create the turns in the trajectory of the electrons. When electrons turn, they lose energy and emit photons.

Immediately after most of the bending magnets there are beamlines. A few of the beamlines use the photons emitted by the bending magnets. However, in third generation synchrotrons, most beamlines use the X-ray radiation produced by insertion devices (wigglers or undulators), placed just before the bending magnets. In this case, bending magnets just serve the purpose of isolating the X-rays from the electrons. Beamlines that use radiation from bending magnets have a broad and fixed X-ray spectrum, while those that use radiation from insertion devices can either have a broad spectrum or discrete harmonics. The gap and the current of the insertion devices can be tuned to influence the energy spectrum and intensity. The beam resulting from insertion devices is more collimated and consequently more brilliant than the one produced with bending magnets.

Beamlines consist of an optics cabin, an experimental cabin and a control cabin. It is at the experimental cabin of the beamline where the sample, its environment and the detector are placed. In the optics cabin, the X-ray beam can be shaped further with filters and monochromators until the required energy spectrum is obtained. As to the control cabin, it is a place where researchers customise the necessary software to control the different devices in the optics and experimental cabins before and while running the data acquisition process.

2.1.2 Computed Tomography

When a sample is scanned from a specific direction, the information captured by the detector is in the form of X-ray projections, which provide the attenuation across the sample. In this work there is a need for a 3D X-ray image of the sample from which the 3D micro-structure of the material under study can be observed, measured and analysed.

Computed Tomography is the procedure which involves imaging a sample from different directions so as to obtain 3D information, and is not only applied to X-ray imaging but can also be applied to other imaging modalities. For this purpose, either the source and detector or the sample is rotated while a number of projections from different angles are taken (Michael, 2001). From these projections, the 3D attenuation can be inferred by solving a multi-dimensional inverse problem named tomographic reconstruction. The most common tomographic reconstruction technique is an analytical method called filtered back projection (FBP) (Herman, 2009) and it is based on the Radon inverse formula (Radon, 1917). FBP applies a convolution filter to overcome the blurring produced by the conventional back-projection algorithm (Dudgeon and Mersereau, 1984). A downside of FBP is that it would provide the exact solution if infinite projections with infinite resolution were to be available. For more realistic settings, where the data is discrete and noisy or fewer angles are available, iterative reconstruction algorithms (Beister et al., 2012) might be necessary. To our knowledge, all the data in this thesis was reconstructed by means of FBP.

2.1.3 Noise and Artefacts

The primary determinants of image quality are noise, artefacts and variations in spatial resolution. This section will cover the first two, noise and artefacts, because these are the ones that challenge the task of image segmentation, the core methodology of this thesis.

The term artefacts refers to undesirable image features that would occur again if the scan is retaken under identical conditions. In contrast, noise always appears at random and would therefore be different even if the scan is repeated under identical conditions. A sample comprising several uniform phases ought to have a high enough relative contrast between its different material phases if wanting to separate them one from another. High noise levels can reduce the relative contrast between phases, making them indiscernible from one another. Artefacts will also affect the relative contrast between material phases, but normally just locally whereas noise will usually affect globally.

There are several random processes in X-ray imaging that can lead to quantum noise. These include i) the number of photons that leave the source; ii) the number of photons that travel through the object unaffected, without being scattered or absorbed; iii) the number of photons that are captured by the detector; iv) the number of light photons generated per captured X-ray photon and v) the response of the imaging sensor when receiving light photons. Overall, the noise is approximately Poisson distributed.

The most typical artefacts are covered in (Boas and Fleischmann, 2012). Some of these are already present in the projections, for example, streak artefacts. These can be caused by multiple mechanisms including Poisson noise or beam hardening, where low energy photons are absorbed quicker by the material than those with higher energy. Streak artefacts can be reduced using iterative reconstruction techniques. Other common artefacts appear when performing the reconstruction from the projection data. If the sample has moved during the scan there will be motion artefacts in the form of blurring and, if the centre of rotation is not measured precisely, edge artefacts can also be present. The type of artefacts that were encountered more often during this work were ring artefacts. Ring artefacts occur when a feature is not properly corrected by the flat field or if there are dead pixels in the detector.

Flat-field correction is the standard procedure at synchrotrons and laboratory sources to obtain the beam intensity I_{in} , which is necessary to calculate the attenuation μl from the detected intensity I_{out} , as was shown in Equation 2.1. It is performed by taking an image with the beam on but without a sample (flat-field image). For the same experiment, the scan time can be much longer at laboratory sources than at synchrotrons. Thus, flat-field images are taken several times during an experiment in order to correct for temperature variations that cause changes in the intensity of the beam. At synchrotrons, flat field images might just be taken before and after scanning.

Besides correcting for non-uniformities in the incident X-ray beam, flat-field images will also correct for non-uniformities in the detector system. During the flat-field correction process, dark field images are also acquired. These are images with the beam off and are usually taken at the end of the scanning to capture fixed-noise patterns or systematic errors.

During the course of this project, collaborators have acquired most of the data and performed the reconstruction, noise and artefact reduction techniques themselves. The synchrotron data-set used in Paper D was acquired by myself at beamline ID19 of ESRF and there I did run the reconstruction and ring artefact correction myself. I did the correction post-reconstruction using polar coordinates to remove circular objects around a centre of rotation. For all other data-sets used in this thesis, we were aware that some artefact or noise re-

duction had been applied before, for example, beam hardening correction, and further removal of artefacts or noise has not been necessary. This is mainly due to applying analysis methods that are robust to noise and, if an artefact were to occlude a randomly picked fibre occasionally, it would not affect the resulting statistics.

2.2 Optical Microscopy

The optical microscope (OM), often referred to as light microscope, has its roots in the 17th century. The latest microscopes are digital and use visible light, a set of lenses and a digital camera to capture a 2D image of the surface of a sample. The maximum resolution that can be obtained is limited by the minimum wavelength of visible light ($\lambda = 0.4\mu\text{m}$), and is of approximately 200 nm.

2.3 Scanning Electron Microscopy

A detailed overview of the origins and development of scanning electron microscopy (SEM) until it was first marketed by the Cambridge Instrument Company is given by (McMullan, 1995). A scanning electron microscope is an electron microscope that scans a sample's surface with a focused beam of electrons and provides information about the composition and topography of the scanned surface. The magnification provided by SEM can be tuned over a range of around six orders of magnitude up to resolutions that lie under 1 nm. The sample preparation process is quite tedious, requiring grinding and polishing of the surface to be scanned until it becomes perfectly smooth. This often causes fractures in weak materials, such as in the glass phase of glass fibre composites.

Dictionary-based Probabilistic Segmentation

Image segmentation or pixel classification is the process of assigning a label to each individual pixel in an intensity image. There can be as many labelling options for each pixel as classes into which we choose to separate the image in. For every class, the resulting segmentation will be a binary image that has 1's in the pixels belonging to that class and 0's elsewhere. In contrast, probabilistic image segmentation does not decide whether a pixel belongs to one class or another, it produces a probability map for each class instead. In other words, an image containing for each pixel a value in the range $[0, 1]$, accounting for the probability that the given pixel has of belonging to a certain class.

This chapter starts in Section 3.1 by motivating the choice of the dictionary learning segmentation framework. Then, it continues by presenting in Section 3.2 the basic segmentation algorithm together with a discussion on tweaking its parameters. This algorithm is supervised, meaning that it requires a learning (or training) step, where the relationship between a raw image and a corresponding pixel-wise labelling is learnt and encoded in a dictionary. It should also be mentioned that this algorithm lies in the category defined above of probabilistic image segmentation algorithms. The chapter ends by describing in Section 3.3 the main variants to the basic algorithm, which has been in constant development during the course of this thesis work.

3.1 Motivation and Background

A dictionary of image patches is a compact manner of representing the information contained in an image using small patches that capture local differences at the patch scale. The dictionary needs to contain enough elements so that it can model all the different patterns present in an image, and the image needs to be somewhat regular to be described adequately by a compact dictionary that actually condenses all the information contained in the image. Sparse image coding (Elad, 2010; Olshausen and Field, 1997), applied to various image generation problems; and textons (Julesz, 1981; Leung and Malik, 2001; Shotton et al., 2008), for texture modelling, are both based on dictionaries. Dictionary learning has proven to be successful for texture segmentation as well (Mairal et al., 2008; Malik et al., 2001). There are various mathematical approaches for learning the dictionary and a good overview is given in (Etter et al., 2011). Dahl and Larsen (Dahl and Larsen, 2011) obtain state-of-the-art performance for texture segmentation using a simple weighted k -means approach to learn a dictionary of image patches.

Learning dictionaries of discriminative image patches (Dahl and Larsen, 2011) showed successful results for supervised image segmentation even when images were contaminated by noise. Actually, this segmentation approach provides successful results even if the pixel labelling used in the training phase of the supervised algorithm is of low quality. When the method was first presented, it was demonstrated on two data-sets. The first case involved classification of individual pixels in artificially created textured images into the different texture patterns. The second case dealt with segmentation and classification of cells from histopathological coloured tissues in different cell types. A year after it was also applied to tracking neural progenitor cells in phase contrast microscopy images for mitosis detection (Vestergaard et al., 2012a).

As was mentioned in Chapter 1, much of the work in this thesis has been on developing methods to aid characterisation of fibres, specifically those in UD composite materials. Inside a certain UD bundle, there is a high fibre volume fraction (FVF) and fibres are well aligned with the primary bundle direction. The cross-sections of UD fibre composites visualised at the micro-scale through X-ray imaging display closely packed circles corresponding to fibres in a darker background of epoxy resin, as shown in Figure 3.1. It seemed adequate to try out the dictionary learning algorithm to separate the individual fibres from the epoxy resin matrix for several reasons. The dictionary learning methodology i) showed promising results for segmenting cells, quite similar in structure to the fibre cross-sections, ii) is robust to noise and iii) the required labelling of the training images needn't be a perfect ground truth, nor performed by an expert from the application side.

On the one hand, Figure 3.2 exemplifies that the dictionary learning method is robust to noise. Even though the image to be segmented is noisy (Figure 3.2a), its probabilistic segmentation (Figure 3.2c) is adequate and contains much less noise. That is to say, the probabilistic segmentation provided by the dictionary has improved the separability of both classes significantly, as can be inferred from the probability density function (PDF) of its pixel values as well. The PDF in Figure 3.2d shows much better defined valleys than the PDF of the original intensity image, shown in Figure 3.2b. On the other hand, Figure 3.1 shows the segmentation resulting from thresholding the probabilistic segmentation obtained through dictionary learning (Figure 3.1d). It illustrates that the dictionary method can provide good segmentations even if the quality of the pixel labelling employed in the training step is low (Figure 3.1b).

Initially, the idea was to segment the fibre cross-sectional images into two classes, corresponding to the two material phases, namely fibres (made of glass or carbon) and matrix (made of epoxy resin). As illustrated in Figure 3.1, a segmentation that separates the image in two phases does not separate individual fibres one from another, due to the high density of fibres. Similarly, cells were not separated individually in (Dahl and Larsen, 2011). This two class fibre-matrix segmentation could be useful for characterising the composite quantitatively through, for example, the 2D fibre area fraction or the 3D fibre volume fraction (FVF) which quantify the amount of fibre material inside the composite sample. However, it would be more interesting for material scientists and mechanical engineers to measure the geometry of the individual fibres. If individual fibres are separated from each other, a larger range of properties can be quantified including number of fibres, individual fibre orientations, diameter distributions or contact points. Additionally, the geometry defining the individual fibres can be imported into a FEM software for simulating how real life stress conditions would affect that specific fibre geometry and material architecture.

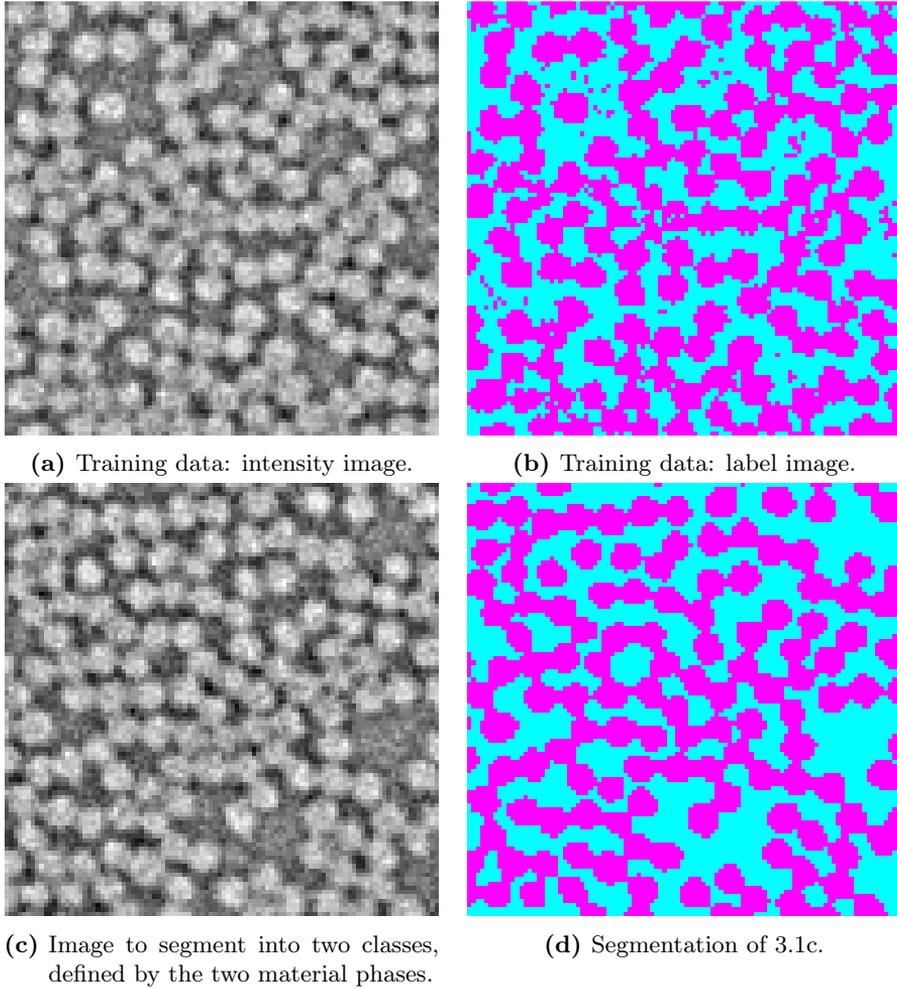
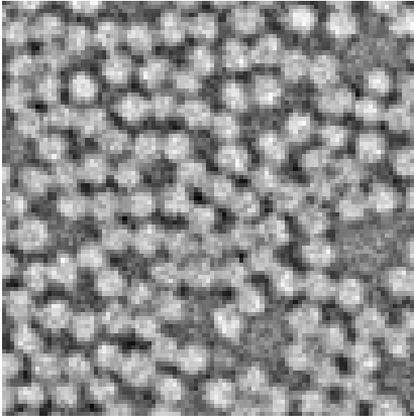
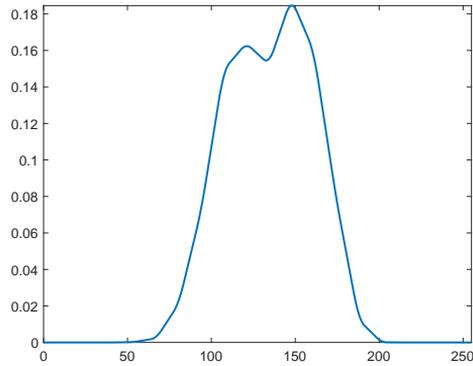


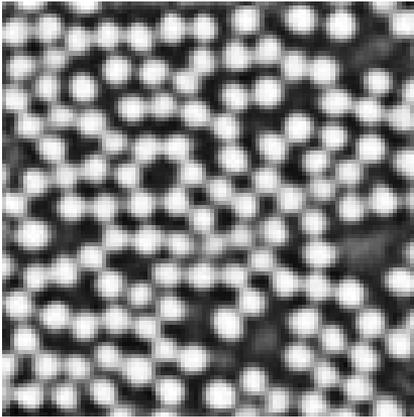
Figure 3.1: Example to show that 1) good segmentations can be achieved using low quality labellings for training the dictionaries and 2) individual fibres cannot be separated by segmenting the two material phases. In (a-b) the training data for learning the dictionaries, where (b) is a low quality labelling obtained by thresholding (a). In (c) the intensity image that will be segmented into the two material phases (carbon fibre and epoxy matrix). In (d) the segmentation resulting from thresholding the probabilistic segmentation granted by the dictionaries trained with (a) and (b).



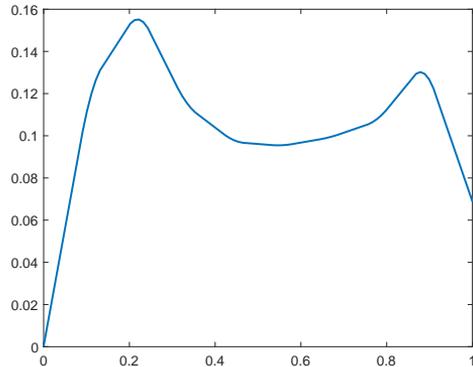
(a) Image to segment into two classes, defined by the two material phases.



(b) Probability density function of 3.2a.



(c) Probabilistic segmentation of 3.2a using dictionary learning.



(d) Probability density function of 3.2c.

Figure 3.2: Histograms of the intensity image and its probabilistic segmentation obtained through dictionary learning. Both histograms show two peaks and a valley. The valley is deeper and the peaks are further away in (d), meaning that noise has been significantly reduced and thus the separability of the two classes improved.

The individual fibre segmentation problem was solved by incorporating a third class to separate the pixels inside a fibre but close to the fibre centre from those closer to the fibre boundary, where the fibre is either in contact with other fibres or with the epoxy phase. This approach is illustrated in Figure 3.3, where the segmentation in Figure 3.3c has been decided by choosing for each pixel the class that gave the highest probability response. Note that a few fibre centres are missing in this segmented image, even though the fibre centres are visible in the probability map in Figure 3.3f. There are other ways of processing the probability maps which would result in finding the centre regions that are missing here, and these will be discussed later.

The first application of the fibre extraction pipeline was to measure individual fibre orientations. For this purpose, fibre diameters are not necessary. 2D centre coordinates estimated from the centre class probability map or segmentation are enough to form individual fibre ‘trajectories’. Fibre trajectories are 3D fibre centre lines connecting the 2D centre coordinates that belong to the same fibre, with one centre coordinate per fibre and per cross-sectional slice. As the centre class is enough, we decided to merge the other two classes into a background class, as shown in Figure 3.4. Now that there are only two classes, thresholding the probability map for the centre class with a $th = 0.5$ would grant the same segmentation as choosing the class that has the highest probability at a certain pixel, as was done for Figure 3.3. Instead, if we lower the threshold down to $th = 0.4$ and apply it to the probability map for the centre class, the centres which were not visible before, are now present in Figure 3.4b.

It should be noted that the segmentation algorithm has been in constant development throughout the thesis work. The principal has actually remained the same, and it is the implementation that has been changing towards becoming more efficient and more general. Currently, the segmentation algorithm can process an image almost in real time and can be applied to semi-supervised settings, lowering the amount of manual input required for the learning part of the method.

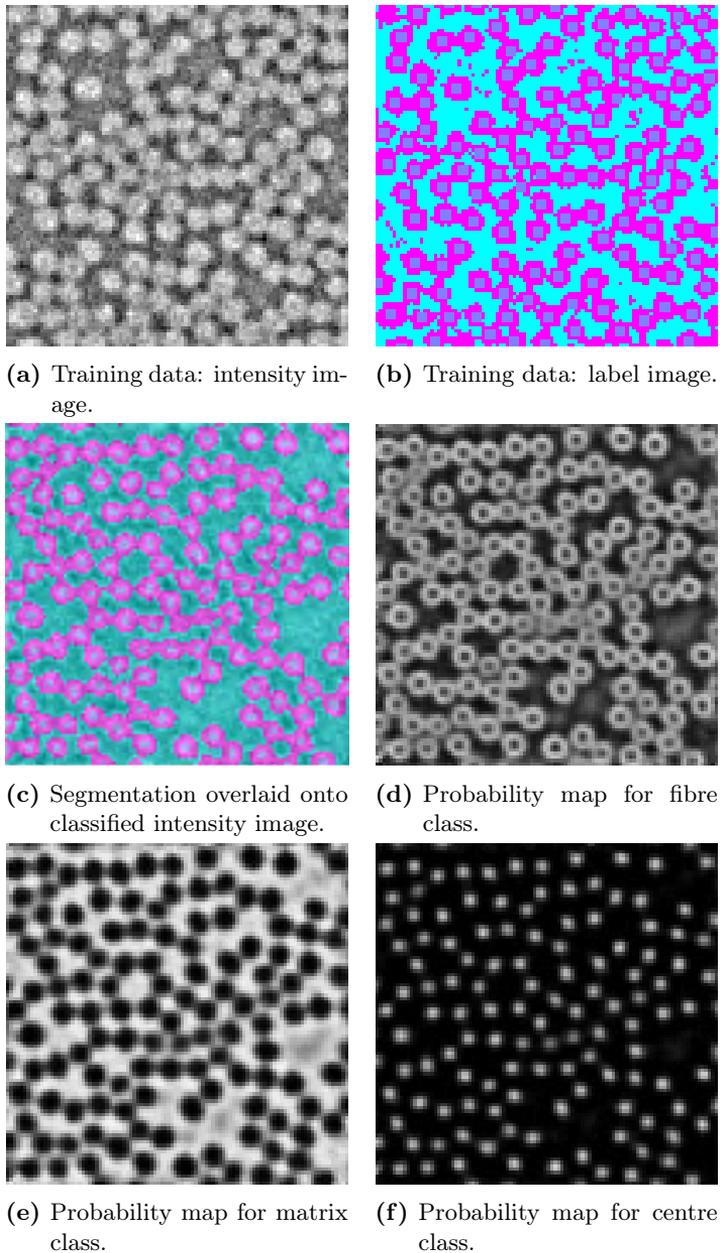
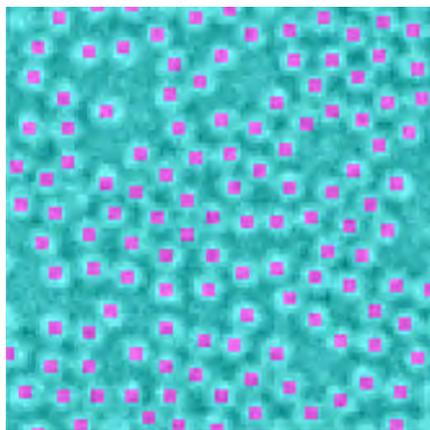
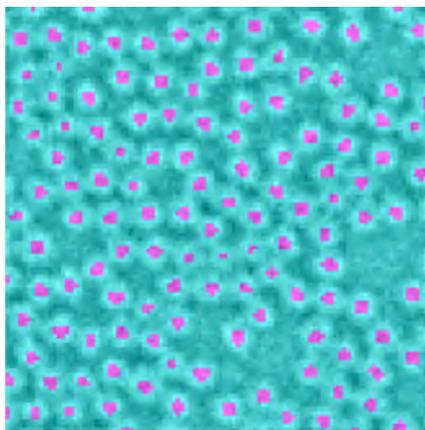


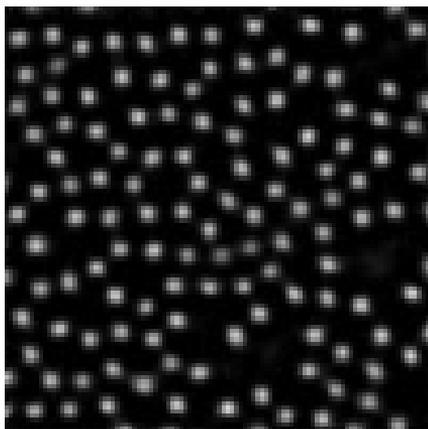
Figure 3.3: Probabilistic segmentation into three classes. In (a-b) the training data for learning the dictionaries, where (b) is the labelling of (a). In (d-f) the probabilistic segmentation given by the probability maps for each of the three classes and in (c) the segmentation obtained by choosing for each pixel the class label that has a highest probability in that pixel.



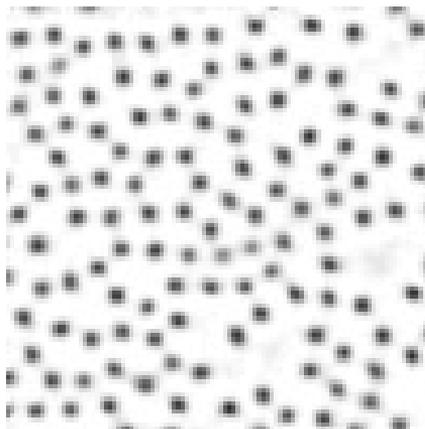
(a) Training data: intensity image and labelling overlaid.



(b) Segmentation overlaid onto classified intensity image.



(c) Probability map for centre class.



(d) Probability map for background class.

Figure 3.4: Probabilistic segmentation into centre and background classes. In (a) the training data for learning the dictionaries and in (b) the segmentation obtained by thresholding (c), the probability map for the centre class. In (d) the probability map for the background class.

3.2 The Basic Algorithm

The dictionary learning algorithm is able to segment an unseen image once it has learnt a dictionary of image patches. That is why there are two main steps in the algorithm, namely training and classification.

In short, the training process is when dictionaries are computed from the input training data. The input is comprised of an image, similar to those that will be classified, and a labelling, that associates every pixel in the image with one out of several pre-defined classes. The dictionary ought to have discriminative power so that the variety of local differences occurring at the patch scale can be fully represented with the dictionary elements-also called dictionary atoms. Actually, the dictionary of image patches is a set of dictionaries: an intensity dictionary and one label dictionary per class. Each atom inside the dictionary of intensities will have a corresponding atom in each dictionary of labels.

The classification process can compute probabilistic segmentations of images that look alike to the one used in the training. Every possible patch in the image to be segmented is matched with an intensity atom. Its corresponding label atoms are the probabilistic segmentation for that patch. The output from the classification process is a probability map for each class. A probability map encodes the likelihood of belonging to a certain class for each individual pixel in the image. Below each of the steps inside the training and classification process are detailed further.

3.2.1 The Training Process

The training process is illustrated in Figure 3.5 and includes the following steps:

- First** Set the required parameters. The number of classes c in which the image is going to be segmented, the size (in width) of an image patch M , the number of random patches Φ selected out of all overlapping patches extracted from the training data-set, the number of dictionary elements n and the number of iterations n_{iter} . Details on tuning these parameters are discussed in page 32.
- Second** Obtain the input data-set for the training process. The training data-set should be an intensity image that has been fully annotated to associate pixels with a set of pre-defined classes using one label per class. Inside the label image, every pixel should take a value in the range $l \in \{1, \dots, c\}$ where c is the number of classes. This is equivalent to having one label image

per class (as in Figure 3.5) where pixels $\in \{1, 0\}$ depending on whether they belong to class l or not.

Third Extract all overlapping image patches from the intensity and label images. Image patches of area $M \times M$ pixels are extracted from the intensity and label images. Patches are extracted starting on the top-left corner of the image and moving one pixel down every time until the bottom of the image has been reached. Then, the next patch will be one pixel to the right of the top-left corner.

Fourth Select a sub-set Φ of these image patches at random. The sub-set of corresponding image and label patches $\mathbf{v}_j = (\mathbf{v}_{Ij}, \mathbf{v}_{Lj})$ where $j = 1, \dots, \Phi$ should be large enough to capture the variety of patches in the training set and always significantly larger than the number of dictionary atoms n .

Fifth Compute the intensity dictionary $\mathbf{d}_I \in \mathbb{R}^{n \times M^2}$. Intensity patches \mathbf{v}_I are grouped using k -means clustering (MacQueen, 1967) where the resulting cluster centres will be the intensity atoms. This clustering algorithm is iterative with iterations $t = 0, \dots, n_{\text{iter}}$. In the first iteration, the n cluster centres \mathbf{d}_I^0 are randomly selected out of the Φ patches. Then, in each iteration t , patches are assigned to clusters by proximity in euclidean space,

$$\kappa^t(j) = \underset{i}{\operatorname{argmin}} \|\mathbf{d}_{Ii}^t - \mathbf{v}_{Ij}\|_2^2 \text{ for } j = 1, \dots, \Phi, \quad (3.1)$$

where κ^t indexes the clusters to which the intensity and corresponding label patches $\mathbf{v} = (\mathbf{v}_I, \mathbf{v}_L)$ are assigned to in iteration t . The size of the matrices \mathbf{v}_I and \mathbf{v}_L is $\Phi \times M^2$ and $\Phi \times M^2 \times c$. The n_i patches assigned to each cluster i are averaged to compute the i -th cluster centre for the next iteration,

$$\mathbf{d}_{Ii}^{t+1} = \frac{1}{n_i} \sum_{p \in \mathbf{S}_i} \mathbf{v}_{Ip}, \quad (3.2)$$

where \mathbf{S}_i is a vector of length n_i containing the set of indices j for the patches in cluster i , those for which $\kappa^t(j) = i$.

Sixth Compute the label dictionaries $\mathbf{d}_{L|l} \in \mathbb{R}^{n \times M^2}$, where each element i is calculated as

$$\mathbf{d}_{Li|l} = \frac{1}{n_i} \sum_{p \in \mathbf{S}_i} \mathbf{v}_{Lp|l} \text{ for } l \in \{1, \dots, c\}. \quad (3.3)$$

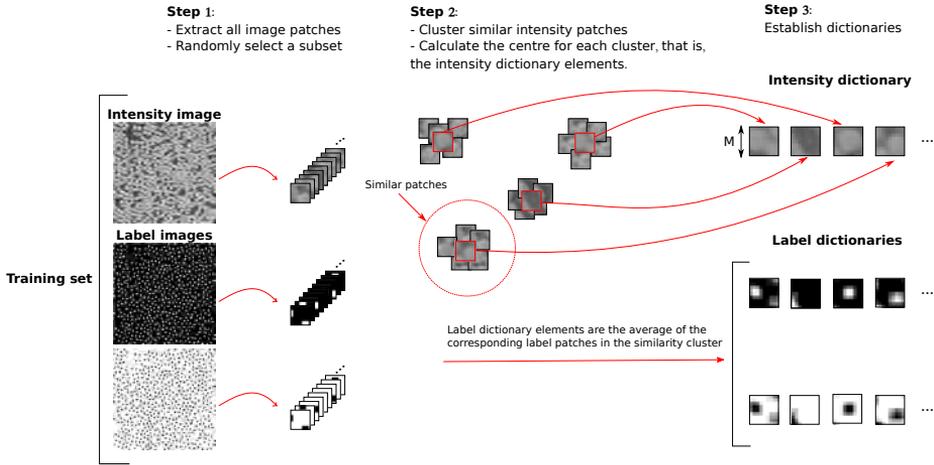


Figure 3.5: Illustration of the training process. First, patches of M pixels in size are extracted from each of the training images and a random sub-set Φ is selected. The n intensity atoms are found by clustering the intensity patches using k -means (MacQueen, 1967). The label atoms are the average of the label patches associated to the intensity patches in each cluster. Please note that $\Phi > n$ and large enough to capture the variety of patches in the training set.

3.2.2 The Classification Process

To find the probabilistic segmentation of an unseen image $U_I \in \mathbb{R}^{x \times y}$ the classification procedure illustrated in Figure 3.6 is carried out. Initially the c probability maps are empty, that is, every pixel is set to zero. Then, the first patch, located in the top-left corner, is extracted and matched to the closest intensity atom in Euclidean distance. The intensity atom has c corresponding label atoms which are added onto the probability maps for each class l at the same position from which the intensity patch was extracted. For this first patch, the top-left corner. The same procedure is carried out for all the patches travelling through U_I as explained in the training procedure for extracting overlapping image patches. Once the bottom-right corner has been reached, the probability maps are complete. The last step is to normalise the probability maps so that the value of each pixel across the c classes sums to one.

The dictionary learning method is robust to noise because the probabilistic segmentation for each pixel is calculated by averaging the label atoms assigned to all overlapping patches containing that pixel. Additionally, each label atom was also formed in the training step by averaging the label patches assigned to

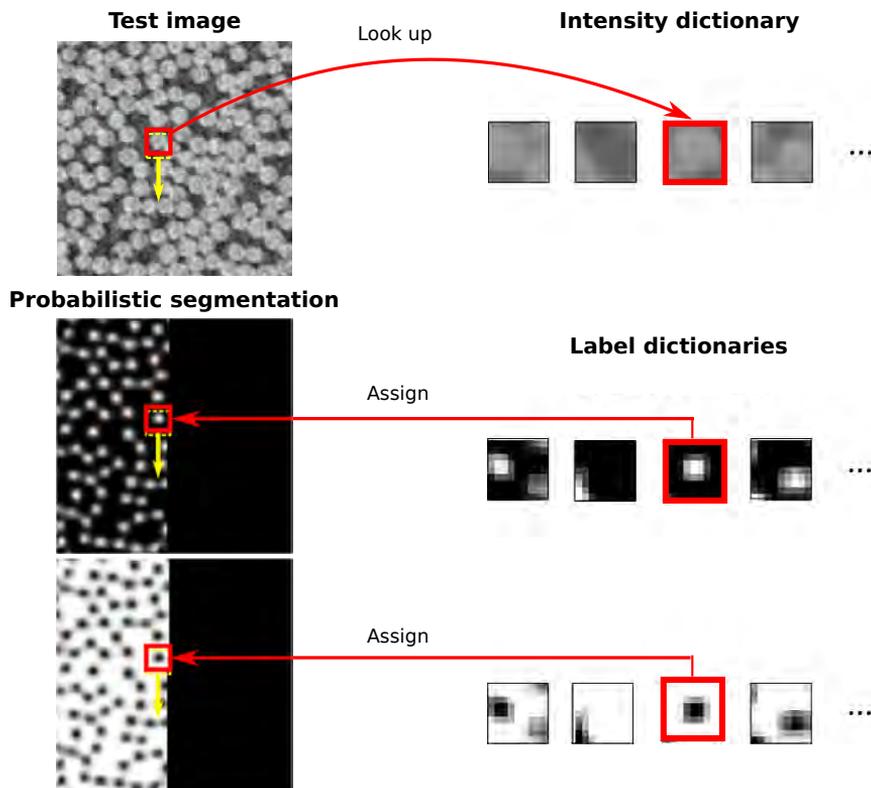


Figure 3.6: Illustration of the classification process. First, overlapping image patches are in turn assigned to the closest intensity atom in euclidean distance. Then, the c label atoms that correspond to the assigned intensity atom are added to the probability maps for their class l in the same position where the intensity patch had come from. And finally, the probabilistic maps are normalised so that the sum of the c probability maps yields an image with ones in every pixel.

the same similarity cluster.

3.2.3 Tuning the Parameters

The resulting probabilistic segmentation is most sensitive to the number of dictionary atoms n and the patch size M . It is also important to set a reasonable value for the number of randomly selected patches Φ for the clustering step

in the training process, but results are not sensitive to a fine tuning of Φ . The number of iterations in the clustering process n_{iter} does not require a fine tuning either.

The sub-set of patches to be clustered Φ should be large enough to capture the variety of patches present in the training image and always significantly larger than the number of dictionary atoms n . In our experience, a good value for structures that do not vary much in size and shape, such as the fibres here, is $\Phi \approx 10n$. If sizes and shapes vary a little more, it is recommended to increase Φ to, for example, $\Phi \approx 20n$. The higher Φ the more certain we can be of capturing all the possible variations in the data. However, there is a limit to how high one should go when setting Φ because the higher the longer the clustering step will take.

With regards to tuning the number of iterations in the clustering step n_{iter} . The problem of separating Φ patches in n clusters where each patch belongs to the cluster with the nearest centre is a computationally difficult problem (NP-hard). Heuristic algorithms are usually employed to quickly find a local optima, but there is no guarantee of finding the global optimal solution. In the considered k -means (MacQueen, 1967) iterative clustering approach we found that after $n_{\text{iter}} = 10$ results did not change significantly.

The number of dictionary atoms n and patch size M have both proven to be important. In Paper A we compared the performance of the dictionary-based fibre centre extraction while varying both n and M . The performance measures we used were true positive rate (TPR) and false positive rate (FPR) and the reference centre coordinates were computed by comparing to a manually annotated reference image, as explained in Paper A. Table 6.1 summarises the performance information for the different combinations of parameters. The gain from $n = 250$ to $n = 1000$ is not worth the extra computational time required for finding the closest intensity atom out of a larger set because this has to be done every time a patch ought to be classified.

The computational time also grows with the patch size M due to the nearest neighbour searches taking place in a higher dimensionality space. With regards to performance, the ideal patch size should be in the scale of the texture that ought to be captured. For obtaining a probabilistic segmentation into fibre centre regions and background, patches should have a similar size to fibres so that patches covering centre regions can contain both fibre phase and matrix phase. If the patch size is too small for covering both phases, it would not be possible to distinguish pixels inside a fibre and close to the centre from those inside a fibre and close to the boundary. This is because pixels around fibre centres and pixels around fibre edges cannot be discerned by their intensity, they have the same, it is the distance of the pixel to the transition between fibre

Table 3.1: Performance of the dictionary-based centre extraction with different combinations of patch sizes M and number of dictionary atoms n .

Dictionary parameters		Performance measures	
M [pixels]	n [atoms]	TPR ^{2D}	FPR ^{2D}
3	50	0.8563	0.1674
3	250	0.9082	0.1678
3	1000	0.9350	0.1721
5	50	0.9468	0.0399
5	250	0.9762	0.0248
5	1000	0.9799	0.0208
7	50	0.9866	0.0104
7	250	0.9873	0.0117
7	1000	0.9910	0.0144
9	50	0.9484	0.0171
9	250	0.9930	0.0100
9	1000	0.9953	0.0097

and matrix phases that is relevant. In contrast, if the patch size is excessively big, too much smoothing will occur because the probabilistic segmentation for a pixel will be an average of many patches, the number of overlapping patches covering a pixel is M^2 .

Figure 3.7 exemplifies what happens when patches are too small, of an adequate size or too big. Fibre centre regions are more enhanced and valleys in between fibre centres are deeper as the patch size reaches the optimal value (see Figure 3.7b). For lower values it just learns the intensity difference between the phases (fibres and matrix). However, because some of the pixels in the fibre class were annotated as belonging to the background class, the probabilities for the centre class are not too high (see colorbar in Figure 3.7a). When patch sizes are too big, probability maps are blurred (see Figures 3.7c and 3.7d) and it is not possible to separate individual fibres any longer. To sum up, larger patches are more discriminant than smaller patches but, if the patch size is too large, the notion of neighbourhood and distances is lost. Additionally, large patch sizes increase the computational time. Thus, the patch size M should be as small as possible to keep the computational cost low and large enough so that it is still discriminant.

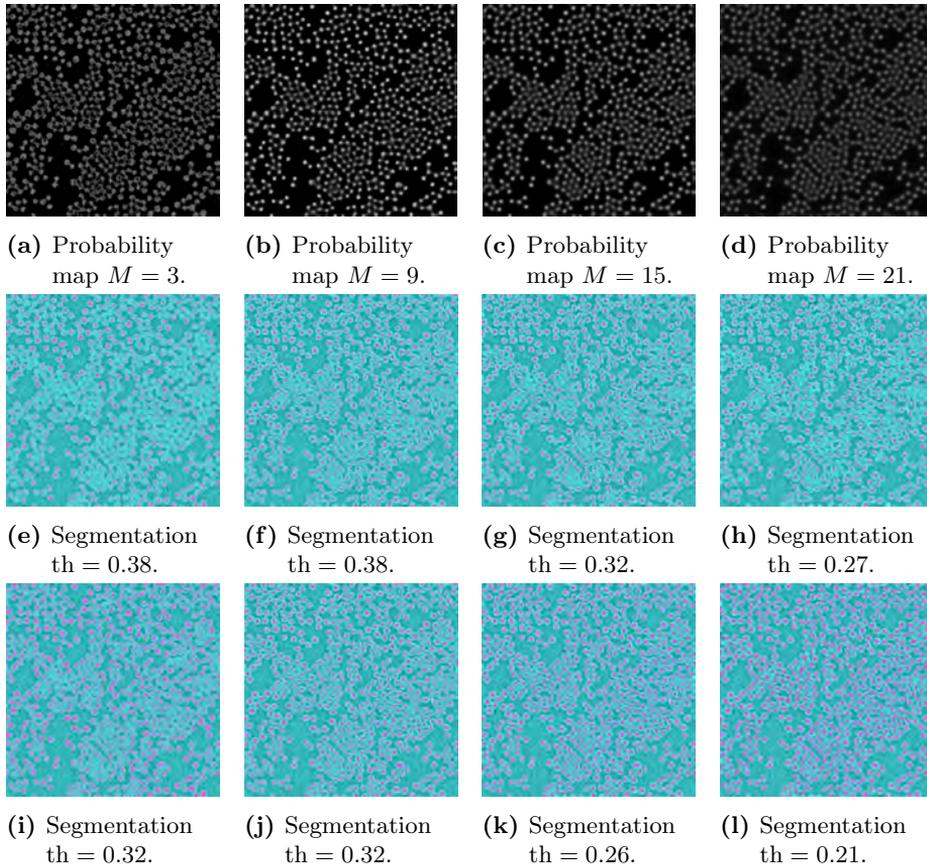


Figure 3.7: Comparison of fibre centre segmentations provided by the dictionary learning method for different patch sizes M . (a-d) show the probability maps for the centre class, (e-h) and (i-l) show thresholded versions of the probability maps for a high and a low threshold value respectively. (a,e,i) illustrate that small patch sizes do not separate the fibre centre regions from the rest of the fibre phase and thus thresholded segmentations show either fragmented or connected centres. (b,f,j) show that an adequate patch size provides high probabilities for the centre class with well-defined valleys in between fibre centres and thus (j) shows a segmentation where all fibre centres are captured without any connections amongst fibres. (c,g,k) and more so (d,h,l) show that, when the patch size is too high, centres from different fibres connect due to the probability maps becoming more blurred and the valleys in between fibres not being deep enough.

3.3 Enumeration of Variants

As mentioned above, the algorithm has been in constant development so this section focuses on describing the main variants of the method, which are listed below. Some of these variants were developed with the purpose of improving the accuracy of the segmentation (1, 6 and 7), another was developed to improve the computational time of the algorithm (2) and some others were necessary for developing the interactive tool presented in Section 6.2 and Paper C, which allows faster training requiring minimal user interaction (3, 4 and 5).

1. **Calculation of cluster centres:** From simple to weighted averaging.
2. **Implementation of cluster search:** Towards a faster matching between patches and clusters.
3. **Description of the relationship between image and dictionary:** From operational to graph-based.
4. **Implementation of image-dictionary mapping:** From sums and averages to a linear operator.
5. **Amount of input required for training:** From full to partial labelling.
6. **Features for clustering:** From intensity patches to more elaborate image features.
7. **Dimensionality:** From 2D to 3D.

3.3.1 Calculation of Cluster Centres

In the basic algorithm described in Section 3.2 intensity atoms were the cluster centres resulting from applying k -means (MacQueen, 1967) to the intensity patches. A variant of this approach called weighted k -means clustering (MacQueen, 1967) was implemented with the purpose of improving the discriminative power of the dictionary. The weights are optimised to separate dictionary atoms in the label space, and cluster centres are computed at each iteration as

$$\mathbf{d}_{I_i}^{t+1} = \frac{1}{\sum_{p \in \mathcal{S}_i} w_{pi}} \sum_{p \in \mathcal{S}_i} w_{pi} \mathbf{v}_{Ip}, \quad (3.4)$$

with $w_{pi} = 1 - \frac{1}{\sqrt{2M}} \|\mathbf{d}_{Li}^t - \mathbf{v}_{Lp}\|_2$ and \mathbf{S}_i is a vector indexing the n_i patches assigned to cluster i . Intensity patches that have similar label patches will have more influence on the cluster centre than those with dissimilar label patches.

In this approach label atoms also need to be updated at each iteration, and not only at the end, because they are influencing the weights in every iteration. Label atoms are estimated by simply averaging the label patches, for cluster i at iteration $t + 1$,

$$\mathbf{d}_{Li}^{t+1}|_l = \frac{1}{n_i} \sum_{p \in \mathbf{S}_i} \mathbf{v}_{Lp}|_l \text{ for } l \in \{1, \dots, c\}. \quad (3.5)$$

It could be said that the clustering in the basic algorithm is unsupervised whereas the variant introduced here is a supervised way of forming clusters. Our latest implementation uses the unsupervised scheme so that the mapping between image and dictionary space (see Sections 3.3.3 and 3.3.4) can be calculated prior to computing the dictionary of labels. This is necessary for running the graphical user interface (GUI) where the user can label the training image interactively until the segmentation results are satisfactory. For more details see Section 6.2 and Paper C.

3.3.2 Implementation of Cluster Search

In the process of assigning intensity patches to clusters, the exact solution would be to assign each intensity patch to the closest cluster centre, as was shown in Equation 3.1. This matching step would be very slow if the exact solution should be found. For this reason we either use the fast library for approximate nearest neighbours (flann) (Muja and Lowe, 2009) or build our own single search tree. For our application, experience has shown that it is not so important to get a match that is very close. Thus, the single search tree is employed in the latest implementations so as to get a simple and fast matching that is independent from external libraries.

3.3.3 Description of the Relationship between Image and Dictionary

The relationship between the image space and the dictionary space has been described in an operational manner but can also be described by a graph. Each

patch in the image space, whether it is the intensity image, the label image or any of the probability maps, has one match in the dictionary. If we look at the pixel level, each intensity pixel is related to M^2 intensity atoms, one per overlapping patch covering the pixel. In the label space, the probabilistic segmentation of a pixel is calculated summing up the M^2 label atoms matched to the overlapping patches covering the pixel. However, each overlapping patch contains the pixel in a different location and it is only that location of the label atom which contributes to the probabilistic segmentation of the pixel.

Thus, the relationship between image and dictionary spaces can be expressed with a graph as shown in Figure 3.8. Each pixel in the image is connected to M^2 dictionary atoms, one per overlapping patch containing the pixel. For each dictionary atom, the connection is to a different location and is determined by the location of the pixel in the overlapping image patch.

As will be explained in the next section, the graph relationship can be saved in a matrix format so that pixels values in one space can be mapped into the other space very fast using Matlab[®]. This enables fast propagation of updates from the labelling of the training image to the dictionary and back to the resulting segmentation. This feature is core in the implementation of the graphical user interface (GUI) developed for training the dictionary with minimal (manual) user input. For more details the reader is referred to Section 6.2 and Paper C.

3.3.4 Implementation of Image-dictionary Mapping

The operational description of the relationship between image and dictionary spaces was implemented with summations and averages. As these are linear operations, it is possible to find a linear mapping that expresses pixels in one space as a linear combination of pixels in the other space. No matter if these are label or intensity pixels.

If the image $U \in \mathbb{R}^{x \times y}$ and the dictionary $\mathbf{d} \in \mathbb{R}^{n \times M^2}$ are vectorised, the matrices $\mathbf{T}_1 \in \mathbb{R}^{nM^2 \times xy}$ and $\mathbf{T}_2 \in \mathbb{R}^{xy \times nM^2}$ can be computed from the adjacency matrix \mathbf{A} that represents the finite graph in Figure 3.8. The adjacency matrix is a square matrix of size defined by the number of vertices, which in this case is $xy + nM^2$. Its elements $a_{pt} \in \{0, 1\}$ take values depending on the connections

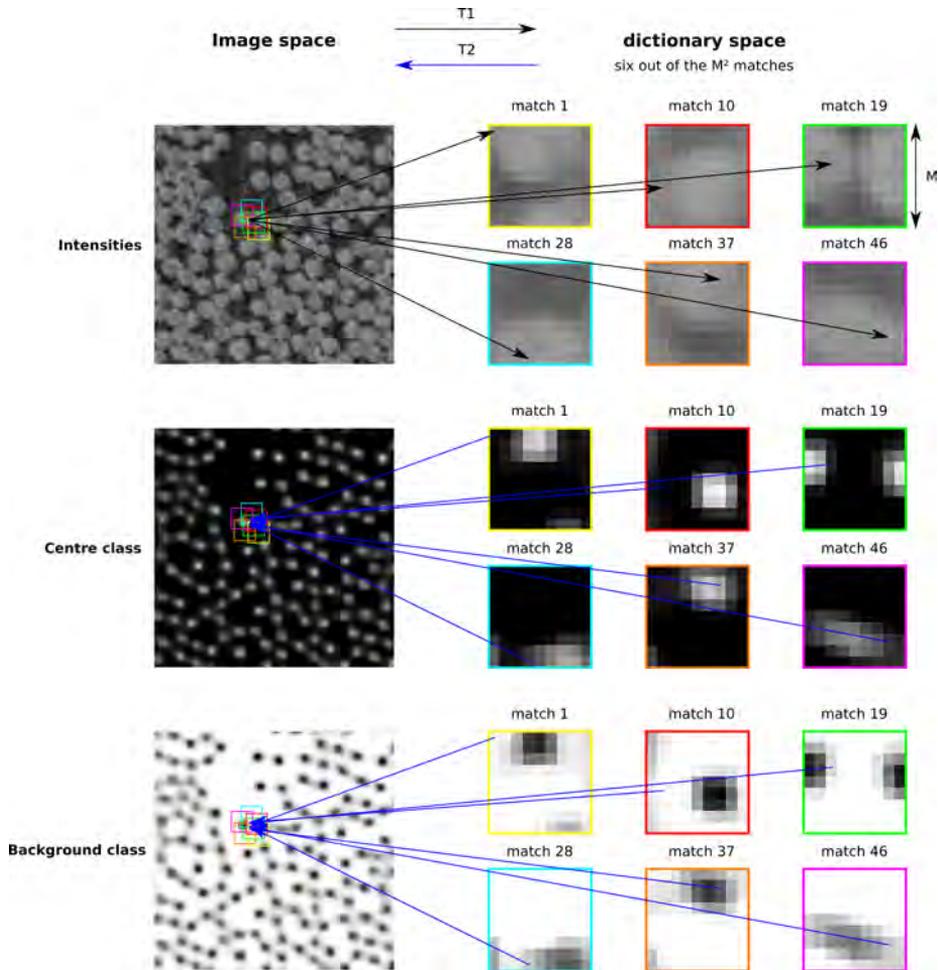


Figure 3.8: Graph description of the relationship between image and dictionary spaces. Each image pixel is connected to the M^2 dictionary atoms that have been matched to the M^2 overlapping patches covering that pixel. The link is to a specific pixel of the dictionary atom. The location of this pixel is determined by the position that the image pixel has inside the overlapping image patch.

between vertices,

$$a_{pt} = \begin{cases} 0 & \text{if } p = t \\ 0 & \text{if } p \text{ and } t \text{ not connected} \\ 1 & \text{if } p \text{ and } t \text{ connected} \end{cases} \quad (3.6)$$

for $p \in \{1, \dots, xy + nM^2\}$ and $t \in \{1, \dots, xy + nM^2\}$

with the connections from each node p in the image space to nodes t in the dictionary described by the graph in Figure 3.8 and explained in Section 3.8.

The adjacency matrix can also be expressed in terms of four building blocks,

$$\mathbf{A} = \left[\begin{array}{c|c} \mathbf{0}_{xy \times xy} & \mathbf{T} \\ \mathbf{T}' & \mathbf{0}_{nM^2 \times nM^2} \end{array} \right] \quad (3.7)$$

where the biadjacency matrix \mathbf{T} is a logical matrix of size $xy \times nM^2$. The top-left and bottom-right blocks are zeros because the graph is bipartite, that is, vertices can be divided into two sets such that there is no connections between vertices in the same set. In our case, the set of vertices representing pixels in the image space and the set of vertices representing pixels in the dictionary space. Additionally \mathbf{A} is symmetric because the graph is undirected, i.e. connections between pixels are bidirectional.

\mathbf{T}_1 and \mathbf{T}_2 are row-wise normalised versions of \mathbf{T}' and \mathbf{T} respectively. This normalisation is performed so that pixels in one space are calculated by averaging, instead of summing, the pixel contributions from the other space. Note that these matrices will be sparse because a pixel in one space is not connected to very many pixels in the other space. For example, a pixel in the image space is connected to only M^2 pixels out of the nM^2 pixels in the dictionary space. If connections are sparse in one direction they are also in the other because the graph is undirected.

Transformations between spaces by matrix multiplication can be performed efficiently because the algorithm is implemented in Matlab[®], which has been especially optimised for being fast at performing matrix multiplications. Actually, Matlab[®] has even optimised storing of sparse matrices in memory, so this computation is feasible even if the image is very large and the dictionary has many large dictionary atoms.

3.3.5 Amount of Input Required for Training

The training step has changed from being fully supervised to semi-supervised so as to reduce the amount of manual input required for creating the training data-set. In this context, semi-supervised means that not all pixels need to be labelled. Previously, the intensity image used in the training phase had to be fully annotated.

This is achieved by giving the unlabelled pixels an equal probability of belonging to each class $l \in \{1, \dots, c\}$. In each of the c label images used as input to the training step, labelled pixels take values 1 or 0 depending on whether they belong to that class or not whereas unlabelled pixels will be given a value of $\frac{1}{c}$.

Image and dictionary spaces are related as shown in the graph presented in Section 3.3.3. Thus, the labelling can be efficiently transferred to the dictionary and back to the probability maps by the linear mapping presented in Section 3.3.4, which Matlab[®] can perform efficiently using sparse matrix multiplications. To allow propagation of labels from the labelled pixels to the unlabelled the process needs to be done at least two times. The partial labelling, where unlabelled pixels take the value $\frac{1}{c}$, is propagated to the probability maps, which are then thresholded and serve as the new labelling of the training image. These ought to be propagated at least once more to the dictionary and back to the probabilities to ensure that the unlabelled pixels have taken a side. More details are given in Paper C.

3.3.6 Features for Clustering

In addition to the M^2 intensity features per pixel, captured by the image patch which is centred over the pixel, more elaborate image features can be included. This can be of interest when the amount of variety in the image is large.

Even though the amount of variety in the image might be captured adequately by a dictionary of image patches if the number of atoms is sufficiently large, introducing more features can still be beneficial. The advantage lies in the possibility of having a reduced dictionary that at the same time captures all the patterns present in the image. Figure 3.9 shows an example tomogram where scale-invariant features (SIFT) (Lowe, 2004) were employed to capture the variety in sizes and shapes of particles in this ammonia slip catalyst (ASC).

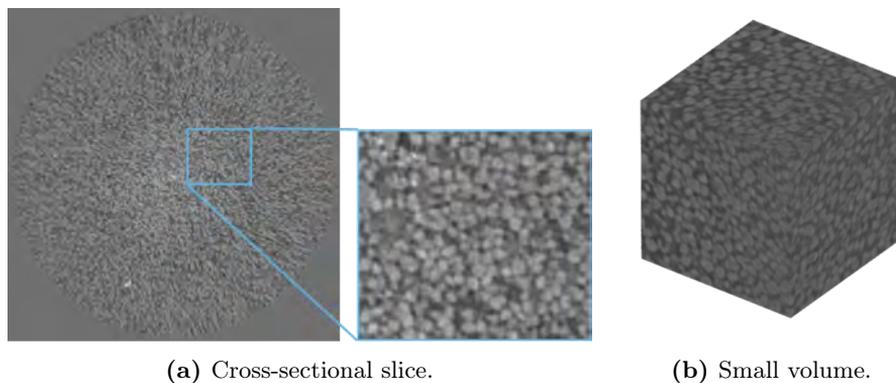


Figure 3.9: X-ray tomogram of an ammonia slip catalyst.

3.3.7 Dimensionality

Until now, this chapter has covered probabilistic segmentation of 2D images. However, the thesis work deals mostly with 3D images acquired with X-ray CT, which ought to be processed slice by slice when using the dictionary of image patches.

The slicing and posterior processing is usually performed in one or several of the three orthogonal directions of the scan. When structures have a preferred direction, like the fibres considered in this thesis, good segmentation results are achieved when processing slices in the direction orthogonal to the structure. In acquiring X-ray tomograms of UD composites, the depth direction is normally well aligned with the nominal fibre direction. Consequently, slicing in this direction will provide cuts that are almost orthogonal to the fibres. To obtain a probabilistic segmentation that is volumetric, the probabilistic segmentations for each slice are simply re-stacked.

On the contrary, for more isotropic structures like the particles inside the considered ammonia slip catalyst (ASC) displayed in Figure 3.9, the volume might be processed in three orthogonal directions. The combined probabilistic segmentation obtained by averaging the three volumetric probabilistic segmentations has enabled more accurate particle centre detections than the individual segmentations in just one direction.

It seemed obvious to extend the dimensionality of our intensity features to exploit the contiguity of our 3D structures. Vestergaard et al. (Vestergaard et al., 2012b) previously used 3D intensity atoms to profit from the contiguity in time

of 2D images when classifying radar data. Therefore, we decided to replace image patches with image cubes, illustrated in Figure 3.10. The hope was to obtain a more discriminative dictionary by exploiting the contiguity of the data in the third spatial dimension.

In the investigation carried out in Paper A, we did not observe a big improvement in the quality of fibre segmentations when comparing the performance of patches to cubes for different values of M and n , nor in segmenting a phantom solid oxide fuel cell (SOFC) data-set in three material phases (see Paper A for the performance measures). Meanwhile, the amount of required memory increased significantly and the computational time increased from an average of 3.2 seconds for the 2D dictionary to an average of 100.3 seconds for the 3D dictionary, when analysing fibres. As to the processing time for the solid oxide fuel cell (SOFC) phantoms, the computational time also increased in two orders of magnitude.

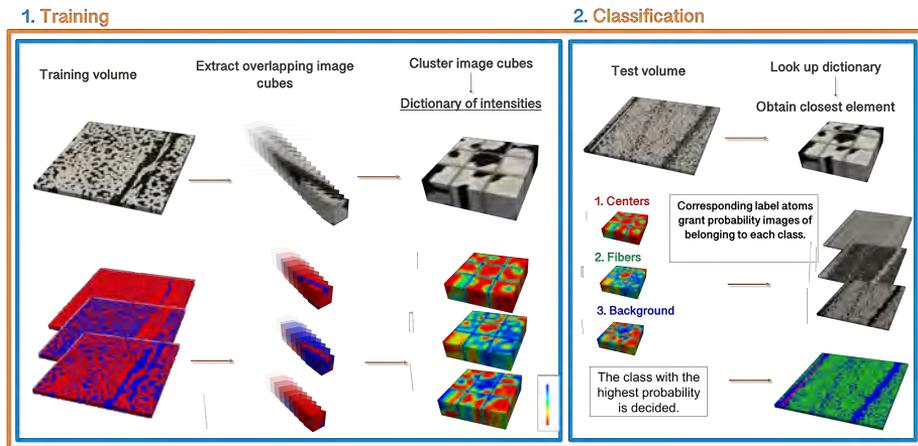


Figure 3.10: Volumetric segmentation with intensity cubes.

CHAPTER 4

Extraction of Geometrical Parameters Defining Structures

In the context of this work, geometry extraction refers to the process of defining structures and their changes through coordinates and boundaries. Geometry extraction in itself is often a task that requires several analysis blocks. These may involve:

1. Detection of 2D or 3D geometrical features from intensity images, which can be raw intensity images or their probabilistic segmentations.
2. Tracking of 2D features in space to obtain 3D structures.
3. Matching structures across data-sets to follow structural changes under external loading conditions or simply to compare geometries.

In 2D images of fibre cross-sections, each individual fibre is defined by a centre coordinate and a diameter. In 3D images, each fibre is defined with a diameter (assumed to be constant along the length of the fibre) and a 3D trajectory—the fibre centre line that connects the centre coordinates across cross-sectional

images. For the ammonia slip catalyst (ASC) in Figure 3.9, individual particles are defined through 3D centre coordinates and particle boundaries. Section 4.1 describes the methods for centre coordinate detection employed in this work while Section 4.2 describes fibre diameter determination.

Paper B explains the basic fibre tracking algorithm to form fibre trajectories from 2D centre coordinates and the improvements to this algorithm are detailed in Section 6.2.

In Paper E the evolution of fibre geometry is imaged under progressive loading conditions. The reader is referred to Paper E for details on fibre matching across loading steps. Besides quantifying structural changes, matching fibres can be necessary for comparing geometries extracted from different data-sets. In Paper D fibres were matched across data-sets with the purpose of validating imaging pipelines for measuring fibre geometry. A summary of this paper is contained in Chapter 8.

4.1 Centre Point Estimation

The task of detecting centre points in 2D images of fibre cross-sections or 3D images of particles can be solved with a variety of methods. However, not all methods provide a solution for every imaging modality or image quality. This is shown in Figure 4.1, where three different methods have been applied to find centre coordinates of fibres in cross-sectional slices of tomograms. The methods under comparison in Figure 4.1 are amongst those detailed in later this section.

Centre points can be detected directly from raw intensity images or alternatively from centre class probability maps. The raw intensity images might have been obtained with X-ray CT, optical microscopy (OM) or scanning electron microscopy (SEM) whereas the probability maps are always obtained through dictionary-based probabilistic segmentation. The methods for centre point detection that will be presented in the following are mostly demonstrated over 2D images but are expandable to 3D.

4.1.1 Thresholding-based Methods

Image thresholding designates the task of binarising an image $U \in \mathbb{R}^{x \times y}$ with a chosen threshold value $\text{th} \in \mathbb{R}$. In choosing the threshold value, inspecting the histogram of intensities and manually choosing a value is quite common. For

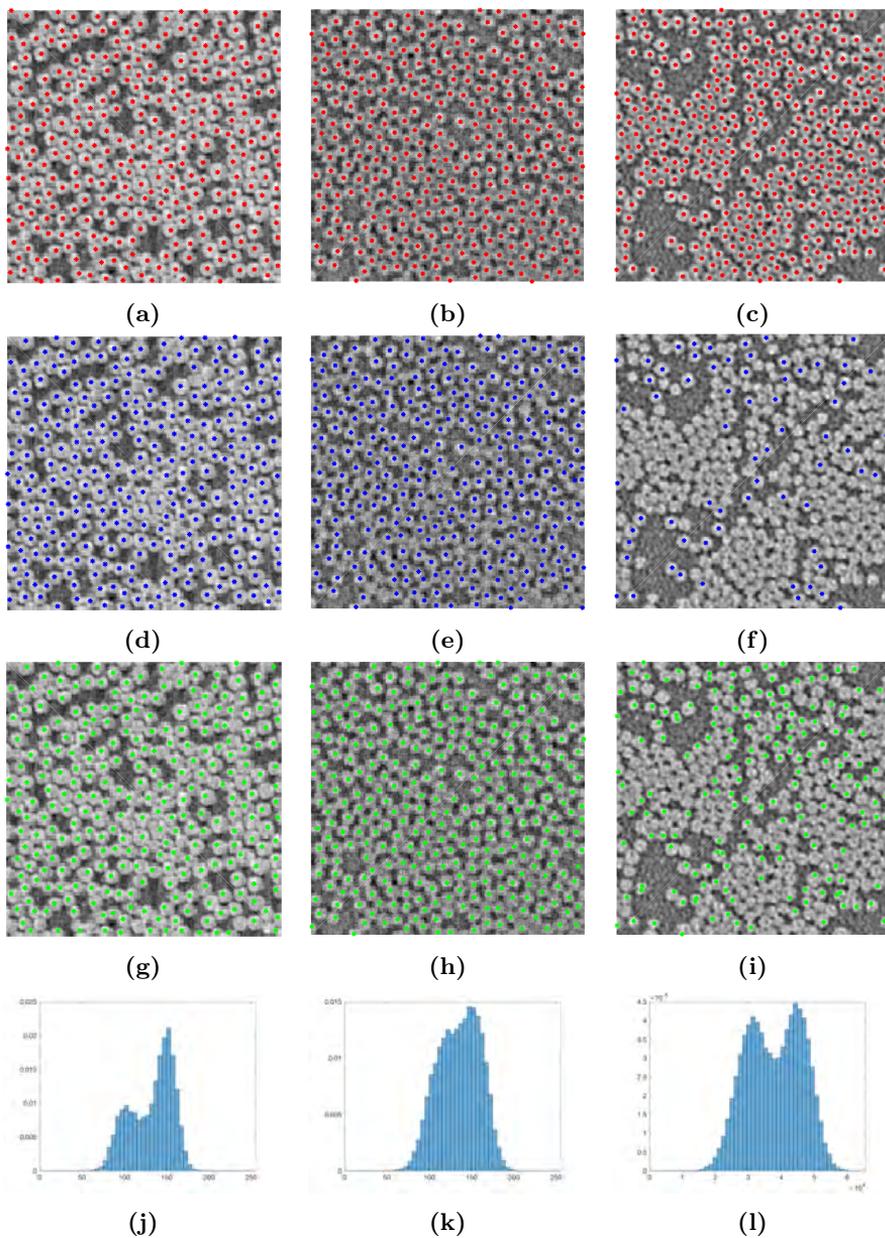


Figure 4.1: Comparison of centre detection methods over different data-sets of fibres. In (a-c) connected component analysis of the thresholded probabilistic segmentation obtained by dictionary learning, in (d-f) the circular Hough transform and in (g-i) a simple blob detector. In (j-l) the histograms for each of the three images. Even though the pixel size in (c,f,i) is the finest and the histogram shows a separation between material phases, there is a slight blur.

cases where there histogram displays a clear valley between two peaks, a popular method for automatically setting the threshold value is Otsu's method (Otsu, 1979).

Connected Component Analysis

When the dictionary-based probabilistic segmentation is applied to an X-ray tomogram, the probability map for the centre class can be binarised to obtain connected regions, each of them corresponding to individual particles or fibres. The centroids for each of these connected regions will give the centre coordinates for each particle or fibre cross-section as illustrated in Figures 4.2 and 4.3. The coordinates of the centroid or geometric centre for each connected region $\mathbf{x}_c \in \mathbb{R}^{\text{dim}}$ are calculated by averaging the coordinates of the connected pixels, where dim is the dimensionality of the euclidean space. Thereby, $\text{dim} = 2$ for fibre cross-sections and $\text{dim} = 3$ for particles.

$$\mathbf{x}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} \mathbf{x}_i \quad (4.1)$$

where n_c is the number of pixels inside a connected component and $\mathbf{x}_i \in \mathbb{N}^{\text{dim}}$ are the coordinates for the i 'th pixel of the connected component.

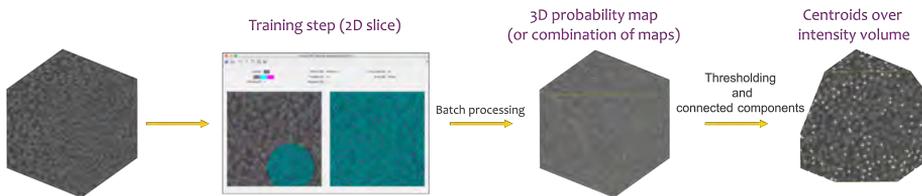


Figure 4.2: Pipeline for particle centroid detection. The dictionary is trained with the interactive interface presented in Paper C and using both intensity patches and SIFT features. The probability map is computed in three orthogonal directions and combined by averaging. The particle centres are found as the centroids of the connected components after thresholding the probability map.

This connected component approach is not directly applicable over a raw intensity image because, even if thresholding provides a good segmentation of the two material phases, it will not give disconnected components when individual elements (particles or fibres) are touching.

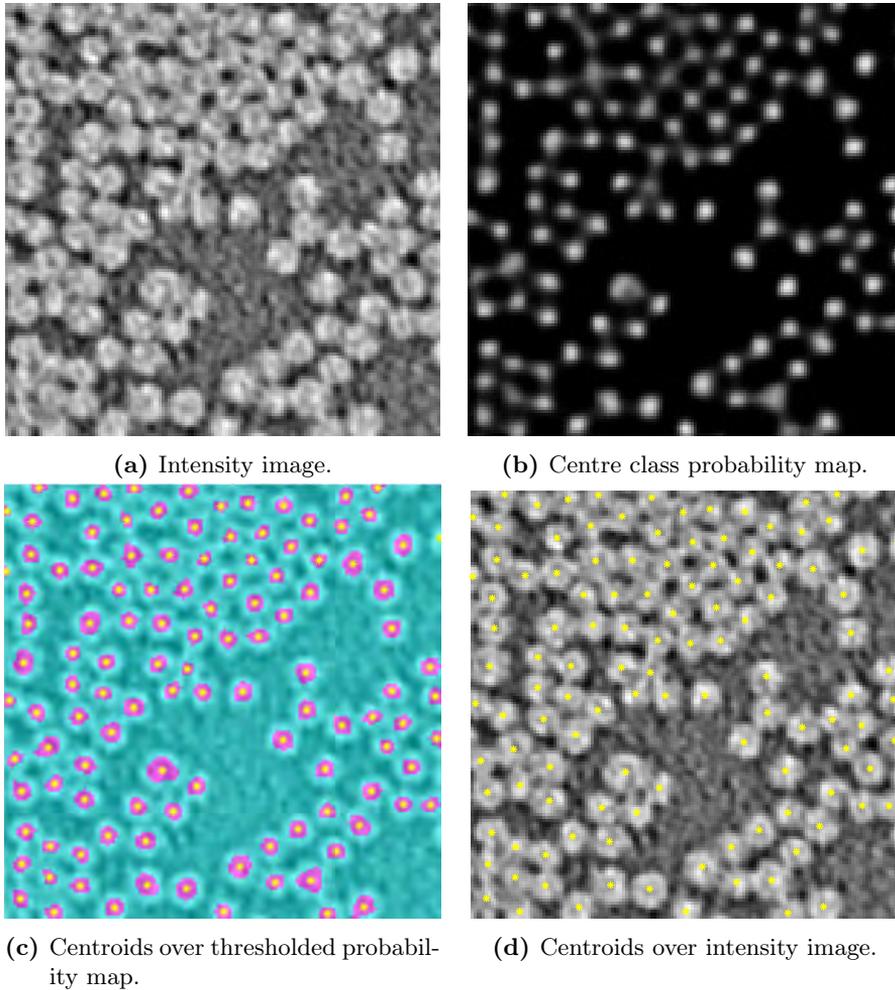


Figure 4.3: Centre point detection by connected component analysis of the thresholded dictionary-based probabilistic segmentation. The intensity image in (a) was acquired with X-ray tomography and the probability map in (b) was obtained through dictionary learning. The centre points in (c-d) were calculated by connected component analysis of the thresholded probability map.

There are methods that can be applied before the connected component analysis to separate individual elements. The distance field presented in the following is able to disconnect elements from high quality binarisations. These high quality segmentations of two material phases can be obtained from imaging modalities

where the contrast between material phases is excellent, such as SEM.

Distance Field

The distance field (Danielsson, 1980) can be applied to binarised images, where pixels take values $\in \{0, 1\}$ depending on whether they are background or foreground. The output will be an image of the same size as the input containing the distances from each pixel to the closest background pixel. Thus, foreground pixels will be assigned distances greater than zero and background pixels will be zero.

For our purpose, the euclidean distance measure is well-suited. In fibre cross-sectional images, pixels that designate fibres will take values in the range $[1, r]$, where r is the fibre radius in pixels. The lower limit of the range will be at the fibre-matrix boundaries and the upper limit at the fibre centres. The process of computing the distance field is clarified in Figure 4.4. From the distance field, centre coordinates are computed in the same way as from the probabilistic segmentation, i.e. by thresholding and connected component analysis. A threshold value slightly lower than r will disconnect the fibre phase and provide connected components from which fibre centroids can be computed.

It can be said that the distance field serves the purpose of enhancing the central region of the fibres, as did the dictionary-based probabilistic segmentation. However, the distance field cannot be applied over all imaging modalities and qualities, as it relies on a high quality binarisation of the image. This is demonstrated in Figure 4.4. The distance field applied to the tomographic slice will not yield correct centre points because the centre region is missing for four of the fibres (see Figure 4.4h). Of course further processing of the distance field could reduce this problem but the purpose here was to illustrate that, even if the contrast between material phases is high, a slight noise level in the image can already deteriorate the resulting distance field.

4.1.2 Blob Detection

Blob detection in grey-scale images is the process of finding a region in the image that differs in brightness compared to its surroundings. Hence, blob detection can be applied to extract centre coordinates from raw intensity images, as in Figure 4.1, or from the centre class probability map obtained through dictionary learning. Finding centre coordinates from probability maps using blob detection, instead of thresholding and connected component analysis, can be of interest in

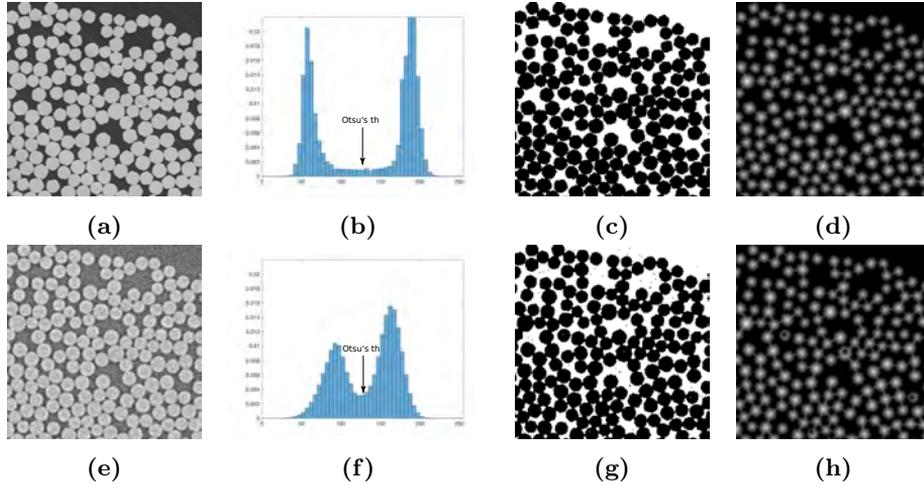


Figure 4.4: Signed distance field for fibre cross-sections imaged by two different modalities. In the *top* row scanning electron microscopy (SEM) and in the *bottom* row X-ray computed tomography (CT). In (a,e) the intensity images, in (b,f) their histograms showing the threshold value that is used for obtaining the binarised images in (c,g). In (d,h) the distance fields which are calculated from the binarisation of the intensities.

some cases. For instance, centre regions might not stand out in comparison to the global probabilities, so they cannot be detected with a global threshold (same for all the image), but might still stand out locally.

The most common procedure for performing blob detection is convolving the image with a function, often based on a Gaussian kernel. Depending on the derivative order of the function, the variety of methods that can perform blob detection may be classified in two big groups. The first group is based on spatial derivatives and the second group is based on local extrema.

We used one of the most simple and common blob detectors, the Laplacian of the Gaussian (LoG), based on local extrema. First, a representation of the input image $f(x, y)$ at scale t should be obtained by convolving the image with the Gaussian kernel $g(x, y, t)$ of scale $t = \sigma^2$ in the directions x and y ,

$$g(x, y, t) = \frac{1}{2\pi t} e^{-\frac{x^2+y^2}{2t}}. \quad (4.2)$$

The linear scale-space representation $L(x, y; t) = g(x, y, t) * f(x, y)$ of the image (Lindeberg, 1994) ensures that there is no aliasing artefacts when going from a finer scale to the coarser scale t where blobs will be detected. Then, blobs are detected from $L(x, y; t) \in \mathbb{R}^2$ by applying the Laplacian operator $\nabla^2 L = L_{xx} + L_{yy}$, which will provide the strongest responses for dark blobs of radius $r = \sqrt{2}t$. To capture light blobs the sign of the Laplacian operator should be inverted. It is up to the user to set a minimum response for considering a local optima as a blob detection. Local optima responses that fall underneath this minimum are discarded.

Blob detection relies on well-defined local optima. For this reason, it can be useful for finding fibre centre coordinates in images of a certain quality. However, results degrade when the image is considerably contaminated by noise or artefacts. These can lead to false centre detections or affect the precision of true fibre centre detections. Additionally, if the contrast between phases is not sufficient or the image is blurred, fibre centres will be missed as illustrated in Figure 4.1.

Modified Blob Detection

When the contrast between material phases is low, the response to the Laplacian operator will be weaker. In this case, the user will most likely need to lower the minimum response that determines when a local optima is considered as a detection. If this minimum response value is global (the same value applies to all the regions), lowering it might cause false detections. Due to this approach only accounting for relative intensities, the task of detecting fibre centre coordinates could suffer from false detections in background regions where there is noise or artefacts. This can be solved by having a local minimum response which is dependant on, for example, the intensities of the input image $f(x, y)$.

Multi-scale Blob Detection

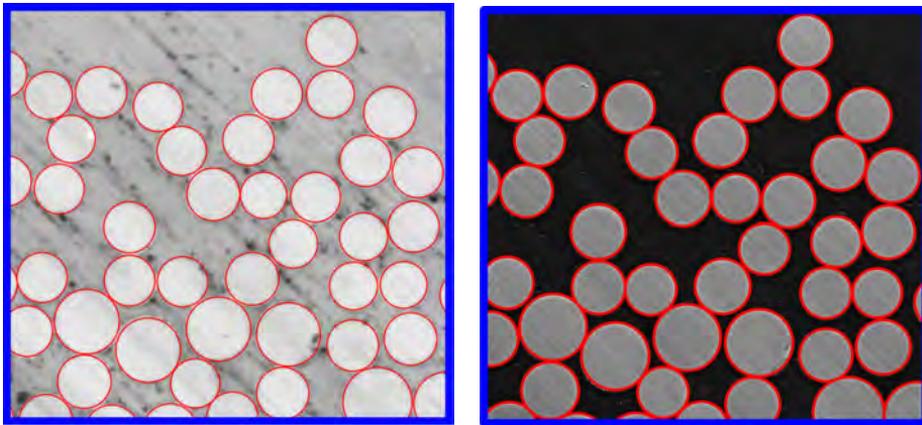
For finding centre coordinates of individual fibres, blob detection can be performed at one only scale because the fibre diameters are known and do not vary in size very much. In contrast, for detecting blob structures of unknown sizes or blob structures that may vary in size, a multi-scale approach is necessary. In this case, the scale-normalised Laplacian operator is applied $\nabla_{\text{norm}}^2 L = t(L_{xx} + L_{yy})$, where the scale-space representation of the image has now been computed for a range of scales $L(x, y, t) \in \mathbb{R}^3$. The blob locations (\hat{x}, \hat{y}) and scales \hat{t} are found as the local optima of $\nabla_{\text{norm}}^2 L$.

4.1.3 Circular Hough Transform

The circular Hough Transform is a two stage algorithm Atherton and Kerbyson (1999); Yuen et al. (1990). First an edge image is computed and then circles are matched to the edge image through a voting scheme. The algorithm we used in Paper D has three parameters including an edge parameter, determining when a pixel is considered an edge pixel, a sensitivity parameter, deciding when a circle is detected, and a parameter determining a circle diameter range.

As the circular Hough transform relies on well defined edges, its accuracy decreases with image blurring (see Figure 4.3) and image resolution (see Paper D).

For high-resolution modalities with well-defined edges such as SEM and OM the circular Hough transform provides precise centre coordinates, as well as diameters. This is illustrated in Figure 4.5.



(a) Optical microscopy (OM)

(b) Scanning electron microscopy (SEM)

Figure 4.5: Circular Hough transform (Atherton and Kerbyson, 1999) applied to OM and SEM to determine fibre centres and diameters.

4.2 Diameter Estimation

The methods chosen for fibre diameter determination are dependant on the method used for centre coordinate detection. The circular Hough transform gives the diameters for the circles at the same time as their centre coordinates. Through blob detection, fibre diameters can be calculated using the multi-scale

blob detection approach presented in Section 4.1.2 and, from the distance field, diameters can be calculated by interpolating the pixel values of each connected component.

While fibre diameters are practically given by the methods covered in the previous paragraph, diameters still need to be computed after detecting centre points from the probability maps by thresholding and connected component analysis. Since the UD fibres under consideration are densely packed, we can assume that fibres are always in contact with some other fibre. Thus, diameters can simply be computed as the distance to the nearest centre point for a 2D tomographic slice.

As was mentioned in the introductory part of this chapter, it can be assumed that diameters do not vary along the length of a fibre. Therefore, a fibre diameter per 3D fibre trajectory is computed from the diameters estimated from the 2D tomographic slices. While the assumption of touching fibres might not always hold for every 2D image, especially at the edges of fibre bundles, it is likely that fibres will be in contact at some point when considering a 3D image.

If a finite element model (FEM) is going to be created from the geometry, the minimum diameter should be selected because structures are not allowed to overlap in the FEM software ABAQUS. Instead, in Paper D we chose to average the diameters across the length of the fibre to prevent biases coming from unprecise centre coordinate detections.

CHAPTER 5

Statistical Characterisation of Structure Geometry

5.1 Exploratory Data Analysis

This section describes different approaches for inspecting data in a statistically meaningful manner.

5.1.1 Summary Statistics

Summary statistics (Mann, 2007) describe data in a simple form through a set of informative measures that can be calculated fast. Amongst others, these measures include mean, median, mode and standard deviation for univariate distributions plus correlations in the case of multivariate distributions.

5.1.2 Empirical Probability Density Functions

The probability density function (PDF) of a random variable is a function representing the likelihood of the variable falling within a certain interval. An

empirical PDF is when this probability has been measured from data.

Empirical densities can be obtained by interpolation from the normalised histogram, meaning the histogram of relative frequencies. A histogram represents the frequency of variable occurrences by counting how many times the variable lies within a specific range of values or interval. Intervals are consecutive, fixed to be equal in size and cover the whole variable range. Besides providing the distribution of the measured data, histogram inspection allows to discover variable occurrences that have an abnormal behaviour (outliers). Figure 5.1 exemplifies a histogram of relative frequencies with outlier values (circled in *red*).

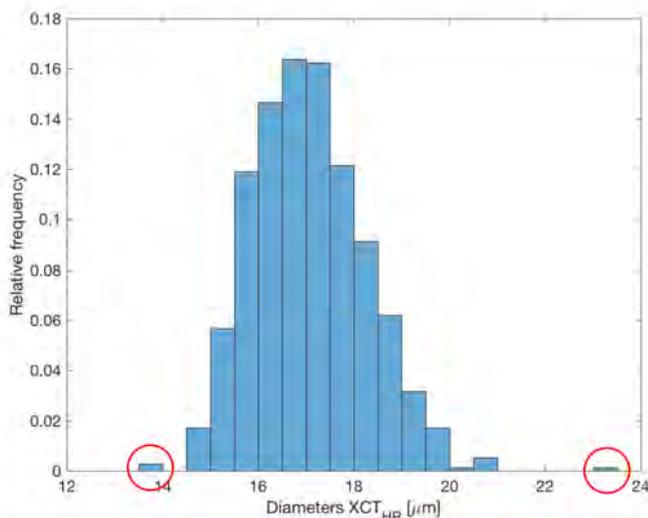


Figure 5.1: Normalised histogram of fibre diameters measured from a high resolution X-ray tomogram.

To obtain a smooth continuous probability curve, as opposed to what the histogram provides, a parametric curve can be fitted to the data, e.g. a normal distribution. With a parametric model a probability distribution is fitted to the data, for which a fixed set of parameters is estimated. For instance, normal distributions have the same bell shape which varies in position and width according to two parameters, namely mean and standard deviation.

Fitting a parametric function assumes a certain probability distribution for the data, so the fit will only be precise when the assumption is correct. For other cases, non-parametric methods can be employed to estimate the continuous probability density function. A common method is kernel density estimation (Hill, 1985) which adapts to the data using linear combinations of e.g.

Gaussian basis functions rather than selecting a specific PDF and fitting its parameters to the data.

5.1.3 Spatial Distributions

In analysing structures that are distributed in space the location of variable occurrences is also important. When characterising the data through summary statistics or empirical PDFs, this spatial information is obviated.

Plotting the value of variables in the actual spatial location can provide information about trends or interdependencies in space amongst individual elements. Figure 5.2 shows the spatial distribution of individual fibre diameters in a cross-section where there is a trend, diameters are smaller in the top-left region of the image. In Paper E we were able to extract valuable insights into the relationship between fibre misalignment and fibre clustering in space, as well as with regards to the collective behaviour of fibres under load (see Figure 7.6) by inspecting spatial distributions.

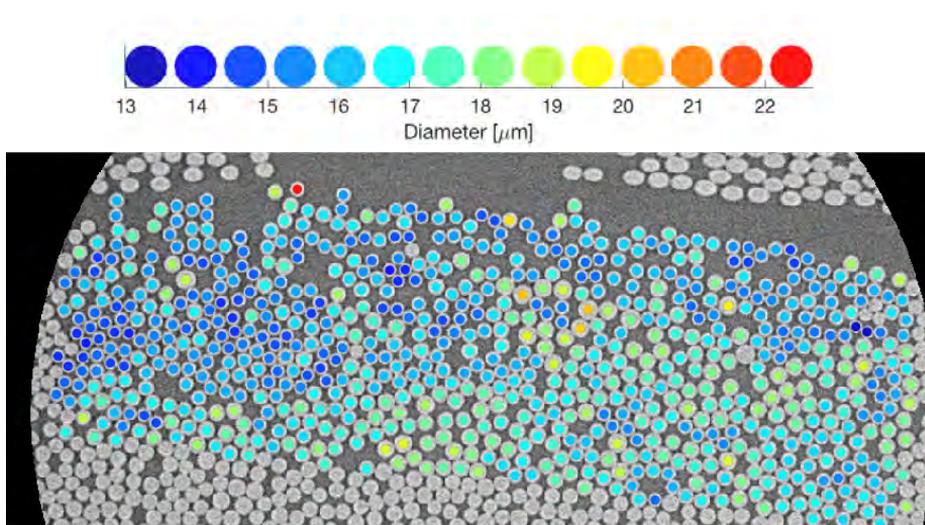


Figure 5.2: Cross-sectional slice with colour-coded diameters.

5.2 General Linear Models

The spatial variation of diameters in UD fibre composites can be modelled with the linear and quadratic models presented below. To decide which general linear model (GLM) fits best to the measured data an analysis of variance (ANOVA) can be carried out.

In this thesis work a two- and three-way ANOVA has been used as well to determine whether the measured diameters are independent from the scanning technique and/or image analysis method used in this two-step data quantification pipeline. The ultimate aim of this investigation is to ensure that measurements capture the physical diameters and thus the real variation in space.

5.2.1 Linear and Quadratic Dependence

The relationship between a dependent variable d and two independent variables x and y can be modelled by a function f , i.e. $\hat{d} = f(x, y)$. When the dependency is linear, the function f will be a plane defined by the parameters a , b and c as

$$f(x, y) = ax + by + c. \quad (5.1)$$

When the dependency is quadratic, the function f will be a curved surface defined by

$$f(x, y) = ax + by + cx^2 + dy^2 + exy + f. \quad (5.2)$$

The parameters of these models can be found by performing a multiple linear regression (Freedman, 2009) where the best fit is usually calculated by minimising the residuals ϵ_i in the least-squares sense, i.e. $\min \sum_{\forall i} \epsilon_i^2$. The regression model for the linear dependency is

$$d_i = ax_i + by_i + c + \epsilon_i \text{ for } i \in \{1, \dots, N\}, \quad (5.3)$$

where N is the number of occurrences of variable d .

5.2.2 Analysis of Variance

An ANOVA is useful in determining which GLM fits best to the data under study. The model that explains the largest amount of variance in the data is

usually selected, however simpler models are preferable (Blumer et al., 1987) if the variance explained by the more complex models is not significantly higher.

Besides deciding whether a linear or a quadratic dependence model is better, it is possible to investigate whether the parameters of the linear or quadratic model are dependant on the methods used for measuring the data.

Before setting up a spatial model with parameters that vary depending on the specific data quantification pipeline, the significant sources of variance will be tested.

Test for the Sources of Variance

We have a data quantification pipeline to measure fibre diameters. This data quantification pipeline (method) is comprised of two steps:

1. Scanning technique, which refers to a specific imaging modality and resolution (mode).
2. Image analysis method, which refers to the method for diameter quantification (algorithm).

We set up a two-way ANOVA to find out whether different methods provide different average fibre diameters,

$$d_{\nu k} = \mu + \theta_{\nu} + \gamma_k + \epsilon_{\nu k}, \quad (5.4)$$

where $\nu = \{1, \dots, Q\}$ and $k \in \{1, \dots, K\}$, with Q the number of methods and K the number of fibres.

We set up a three-way ANOVA to find out whether the measurements are dependant on the choice of mode and/or algorithm,

$$d_{ijk} = \mu + \alpha_i + \beta_j + \gamma_k + (\alpha\beta)_{ij} + \epsilon_{ijk}, \quad (5.5)$$

where $i \in \{1, \dots, N\}$ and $j \in \{1, \dots, M\}$ and $k \in \{1, \dots, K\}$, with N the number of algorithms, M the number of modes and K the number of fibres.

For instance, the ANOVA table for the model in Equation 5.5 will provide the amount of variation explained by the fibre term, the mode term, the algorithm

term and the interaction term between mode and algorithm. Additionally, the residual variation which cannot be explained by the model will also be provided. For each source of variation included in the model, a p -value will be given by the the ANOVA table. From this value we determine the significant sources of variation.

Spatial Regression

When setting up spatial regression models, we start with the most complex model. That is, the quadratic dependence model with its parameters varying according to all the significant sources of variation,

$$d_{\nu k} = \gamma(\nu) + \delta(\nu)x_j + \rho(\nu)x_j + \eta(\nu)x_k^2 + \kappa(\nu)y_k^2 + \gamma(i)x_k y_k + \epsilon_{\nu k}, \quad (5.6)$$

where $\nu \in \{1, \dots, Q\}$ and $k \in \{1, \dots, K\}$, with Q the number of methods.

We can then reduce the complexity of the model by eliminating the dependency of the parameters δ , ρ , η , κ and γ with the method,

$$d_{\nu k} = \gamma(\nu) + \delta x_k + \rho x_k + \eta x_k^2 + \kappa y_k^2 + \gamma x_k y_k + \epsilon_{\nu k}, \quad (5.7)$$

where $\nu \in \{1, \dots, Q\}$ and $k \in \{1, \dots, K\}$, with Q the method and K the number of fibres. This model assumes that all methods capture the same spatial variation in fibre diameters and that the quantification pipeline simply introduces a systematic bias in the mean fibre diameter.

We can reduce the complexity even more and test a model where the intercept γ does not vary with the method,

$$d_{\nu k} = \gamma + \delta x_k + \rho x_k + \eta x_k^2 + \kappa y_k^2 + \gamma x_k y_k + \epsilon_{\nu k} \quad (5.8)$$

where $\nu \in \{1, \dots, Q\}$ and $k \in \{1, \dots, K\}$, with Q the method and K the number of fibres.

To choose a model, the residual variation obtained for the different models is compared. As mentioned above, if the residual variation between two models is very similar, the simplest model should be selected. If all the parameters of the quadratic terms (η , κ and γ) turn out to be not significant, linear dependency models should be investigated in the same manner, from complex to simple.

Part II

Analysis Pipelines and Applications

CHAPTER 6

Fibre Extraction Pipeline

The following chapter explains how the methodology presented in Chapters 3 and 4 is combined to form an analysis pipeline. This pipeline can accurately extract the geometry of individual fibres from tomographic data of unidirectional (UD) fibre composites imaged at the micro-scale.

As was detailed in the Introduction (Chapter 1), measuring fibre geometry from 3D images acquired through X-ray CT enables investigation of real fibre arrangements. Besides enabling characterisation of real geometries, it provides the possibility of studying the very complex damage mechanisms that affect UD fibre composites. During the course of this thesis work we have applied the pipeline for the investigations discussed in Chapter 7.

Paper A demonstrates the applicability of the dictionary-based segmentation framework for detecting individual fibres in 2D slices of fibre cross-sections. Paper B goes one step further and introduces the complete pipeline for determining the 3D fibre centre lines (called fibre 'trajectories'), by adding a fibre tracking step.

The pipeline has been in constant development throughout the thesis work. The main changes include: i) an increase in processing speed, ii) a reduction in the amount of manual input required for learning the dictionary, iii) a graphical user interface (GUI) for interactive dictionary training, iv) an increase in track-

ing accuracy and v) the possibility of determining fibre diameters. The basic pipeline is presented in Section 6.1 and the improvements and extensions are covered afterwards in Section 6.2.

6.1 Basic Pipeline

6.1.1 Detection of Fibre Cross-sections

In Paper A we presented the dictionary-based probabilistic segmentation framework as a method for supervised volumetric segmentation. We demonstrated volumetric segmentation first using the basic algorithm covered in Section 3.2, which works with 2D image patches. In this case the tomographic volume needs to be sliced beforehand to obtain 2D images of fibre cross-sections that are then segmented individually and combined at the end to obtain a volumetric segmentation. In addition, the paper presents an extension to the basic methodology, which uses 3D image cubes and processes the whole volume at the same time (see Section 3.3.7). In the clustering process for building the dictionary, please refer to Equation 3.1 for the matching between intensity patches and cluster centres, as there was a typo in the corresponding equation in Paper A.

As was shown in Figure 3.7, the dictionary framework smooths the boundaries between classes, more so the larger the patch size. This smoothing especially affects the segmentation of small features, such as the fibre centre regions we wish to segment here. To overcome the excessive smoothing that comes from the averaging (performed both in the training and classification steps of the algorithm), a weighting scheme was incorporated to enhance the fibre centre class in the probabilistic segmentation. This approach enabled a more fair comparison of the dictionary performance across varying patch sizes M , and was not employed in future papers/implementations where the optimal patch size had already been set.

Paper A compares the performance of a simple blob detector over the intensity image (explained in Section 4.1.2) to both the dictionary of image patches and image cubes, where centre coordinates are obtained from the probabilistic segmentation by thresholding and connected component analysis (see Section 4.1.1). The performance of the three fibre detection methods is evaluated by comparison to a manually annotated reference and is reported in Table 6.1. The table compares the true positive rate (TPR) and false positive rate (FPR) for the three methods. Both dictionary methods outperform the blob detection method, and the 3D dictionary method slightly enhances the results compared

to the 2D method. For more details the reader is referred to Paper A.

Table 6.1: Performance in detection of fibre centres (Paper A).

Performance Measure	Method		
	Blob detector	2D dictionary	3D dictionary
TPR	0.8339	0.9953	0.9980
FPR	0.0737	0.0097	0.0087

Figure 6.1 shows the detected fibres over a small part of the cross-sectional area used for evaluating the performance. The *red* circles mark a couple of the fibres that are missed using blob detection while the *green* circle shows that two fibres are detected as one by the 2D dictionary. This happens when the valley between two neighbouring fibre centres is not deep enough in the probability map for the centre class. In other words, fibres are not detected individually if the boundary of the central fibre region is not well defined in the probability map for the centre class. In some cases, a slight increase of the threshold value will disconnect the two fibres. In other cases, increasing the threshold to disconnect two fibre centres will mean that other fibres will be missed, even if their fibre centres are well-defined. It can happen that centres are well-defined locally but the probability values are low relative to the probabilities at other centres, and thus compared to the global threshold value. The solution will be discussed in Section 6.2.

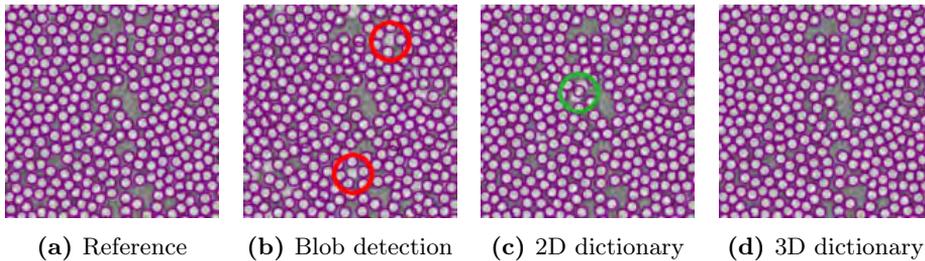


Figure 6.1: Manual reference and circles around the detected fibre centres for the three methods under comparison in Paper A.

To sum up, the main contributions of Paper A are:

- Application of the dictionary-based probabilistic segmentation framework to segment complex image structures. Specifically the basic version of the algorithm (Section 3.2) with the weighted clustering explained in Section 3.3.1.
- Extension of the segmentation methodology in (Dahl and Larsen, 2011)

from a dictionary of image patches to a dictionary of image cubes (see Section 3.3.7).

- Demonstration of the accuracy in segmenting individual fibre centres from tomograms with the dictionary framework when these UD fibres are densely packed inside the composite.
- Evaluation of the computational time and accuracy of the two dictionary methods for different patch/cube sizes and number of dictionary elements (see Paper A for details).
- Comparison of the two dictionary methods against the simple blob detection method in Section 4.1.2.

It should be noted that we have also compared our centre detection methodology to the circular Hough transform 4.1.3. A visual comparison is included in Figure 4.1, where fibre centres have been detected from three tomograms differing in quality with the dictionary framework, blob detection and the Hough transform. Additionally, a quantitative comparison to Hough transform across micro-CT resolutions is included in Paper D, and summarised in Section 8.2.

6.1.2 Detection of Fibres Trajectories

In Paper B we decided to use the dictionary of 2D image patches (same version as in Paper A) for its high accuracy and the low computational time and memory requirements compared to the 3D version based on image cubes. As mentioned above, when working with 2D patches the volume is sliced and slices are segmented individually.

When segmenting particles inside ammonia slip catalysts (ASCs) (see Figure 3.9) or material phases in solid oxide fuel cells (SOFCs) (see Paper A), it is important to merge the segmented slices into a volume before computing particle centre coordinates or quantifying material structures (volume fraction of each phase, contact points between phases...). In analysing UD fibres this is not necessary because we extract individual fibre centre coordinates from the 2D probability maps. After determining 2D fibre centre coordinates, these are connected through the volume to form individual fibre trajectories, as explained in Paper B. The complete pipeline is illustrated for a very small volume in Figure 6.2.

We evaluated the accuracy of centre coordinate detection by comparing the estimated points to those obtained from a manually annotated test image. The

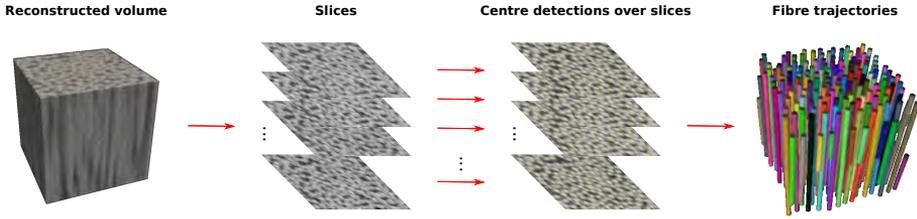
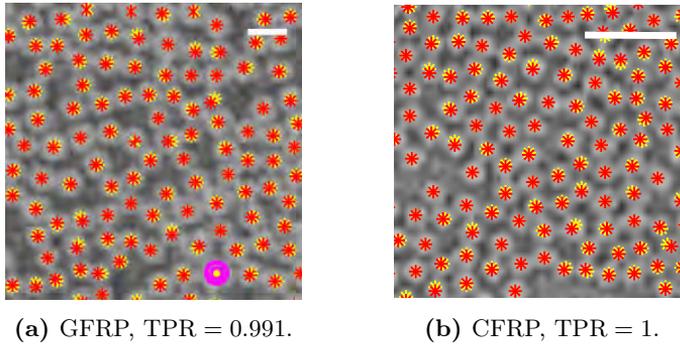


Figure 6.2: Description of the individual fibre segmentation pipeline introduced in Paper B. First, the volume is sliced. Then, fibre centre coordinates are found in every slice. Finally, the centre points corresponding to the same fibre are connected through the volume by tracking.

test image, which was not used in the training nor the validation steps of the algorithm, had a width of $674.4\ \mu\text{m}$ for the glass fibre reinforced polymer (GFRP) and $289.2\ \mu\text{m}$ for the carbon fibre reinforced polymer (CFRP). A small region of this test image is shown in Figure 6.3, where the reported accuracy is measured as the TPR for the whole test image.



(a) GFRP, TPR = 0.991.

(b) CFRP, TPR = 1.

Figure 6.3: Centre points over a region from the test image with a scale bar indicating $40\ \mu\text{m}$. In *yellow* the reference points and in *red* the detected ones, missed detections are circled in *pink*.

To sum up, the main contributions of Paper B are:

- Development of a pipeline for automated fibre geometry extraction that can be applied to segment UD fibres in low contrast and noisy tomograms containing densely packed fibres.
- Evaluation of the centre detection accuracy through comparison to a manually annotated reference.

- Characterisation of fibre geometry through the orientation of individual fibre trajectories, detailed in Section 7.1.1.

6.2 Improvements and Extensions

The basic blocks of the pipeline for individual fibre extraction developed in Paper B are shown in *blue* in Figure 6.4. The basic blocks have been improved and the pipeline has been extended to incorporate the possibility of estimating individual fibre diameters.

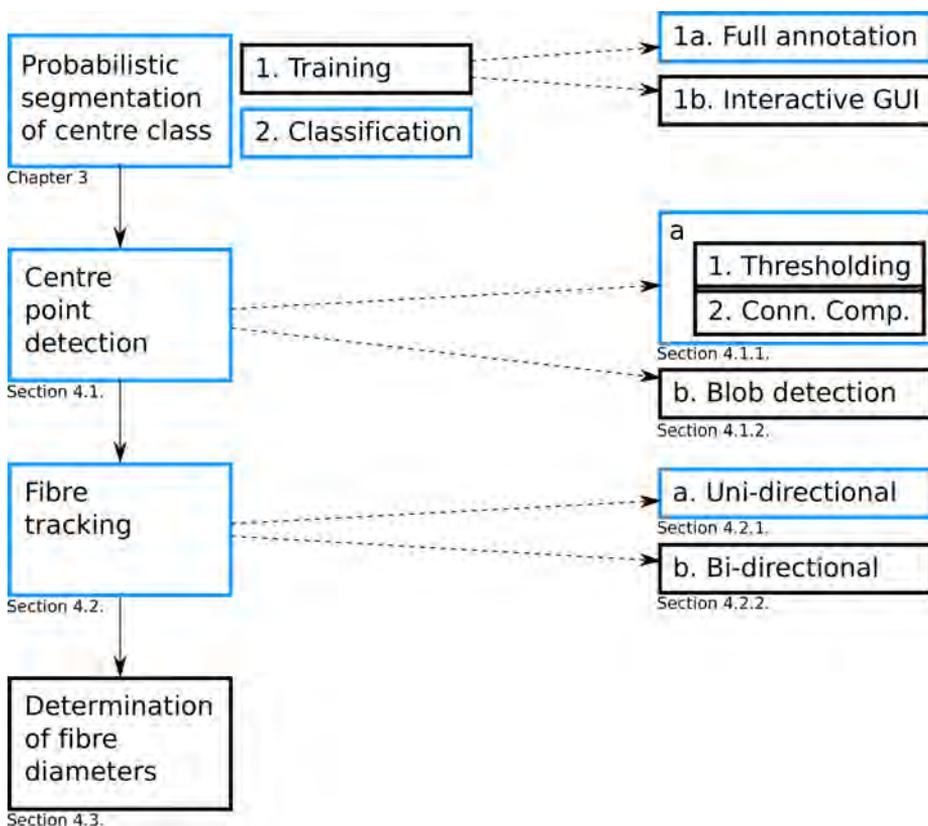


Figure 6.4: Schematic description of the basic pipeline for individual fibre segmentation presented in Paper B (in *blue*) with the developed improvements and extensions in *black*, which are applied/developed in Papers C, D and E.

In the basic pipeline the dictionary training step requires fully annotated label images as explained in Section 3.2 and the dictionary is computed by weighted k -means clustering (Section 3.3.1). An improved version, employed in Papers C, D and E, provides a faster training which requires minimal user input and is user friendly. To make the training process more user friendly a graphical user interface (GUI) was developed. The user can partially label a training image on the left window and immediately obtain a full segmentation of this image on the right-hand side window, as shown in Figure 6.5. The user can continue labelling on the left until the segmentation on the right is satisfactory. At this point, the dictionary can be exported and a batch of slices can be processed.

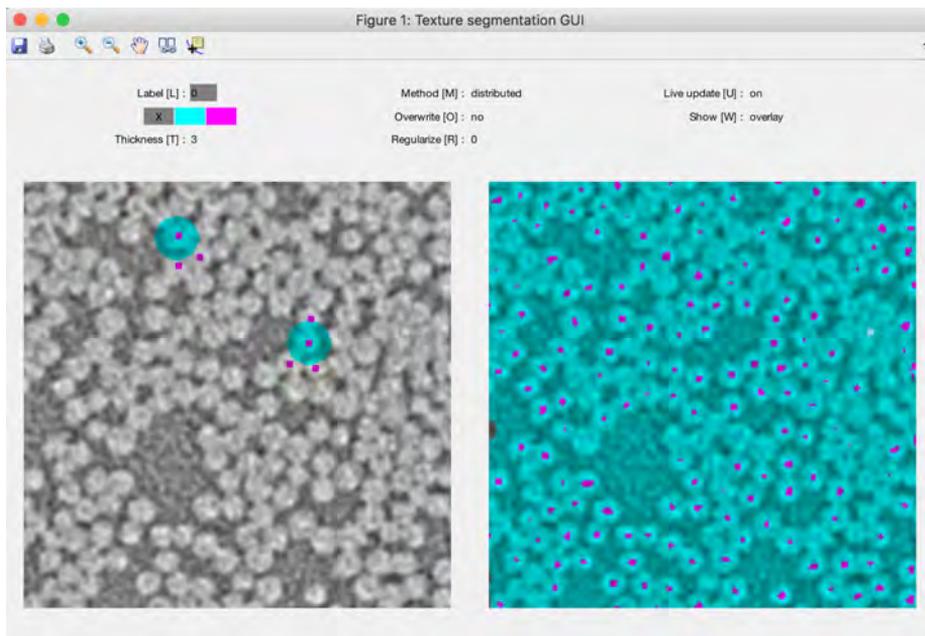


Figure 6.5: GUI for dictionary-based probabilistic segmentation.

This interactive dictionary training is possible by computing the dictionary of intensities (simple k -means) and finding the linear mapping between image and dictionary space prior to running the GUI. The linear mapping, implemented in Matlab[®] as a matrix multiplication (see Section 3.3.4), enables fast propagation of the labelling updates from the left window to the label dictionary and back to the segmentation window on the right. Pixels on the left that have not been labelled are given an equal probability of belonging to each class, as explained in Section 3.3.5.

It could be said that now the training process is semi-supervised because a par-

tial annotation of the training image can already provide an adequate dictionary. The great advantage of the GUI lies in the possibility of seeing immediately how the changes in the labelling affect the segmentation. This means that the user can easily adapt the amount of manual input to the specific data-set. For high quality data-sets annotating a couple of fibres will be enough whereas for lower quality data-sets with noise, artefacts and/or blurring fibres will not be well-resolved and more user input will be necessary.

The label propagation idea for image segmentation is presented in Paper C. The main contributions of this paper are:

- Development of a flexible approach for characterising image content where the user is able to define the structures of interest.
- Design and implementation of an interactive GUI where the user can annotate pixels interactively with a real-time feedback. This allows for semi-supervised segmentation with minimal user interaction.
- Demonstration of the applicability of this label propagation methodology for segmentation of patterned images originating from imaging systems for materials science and medicine.

The next improvement in the pipeline was created to solve the problem mentioned in the previous section, where two fibres were detected as one (see Figure 6.1c). A solution to this problem is to calculate the centre coordinates from the probability maps using blob detection instead of thresholding and connected component analysis.

As for fibre tracking, the detection of 2D fibre centre coordinates is so accurate that a simple tracking straight down the volume in the same direction for all fibres and volume depths provides accurate fibre trajectories. In Paper B we had a unidirectional tracking algorithm where we matched centre coordinates from one slice to the previous slice by a nearest neighbour search in one direction (unidirectional fibre tracking methodology detailed in Paper B). We realised afterwards, through the visual validation explained in Section 8.1, that sometimes two fibres were merging into one if the centre coordinates had not been placed very precisely due to a lower data quality or coarser resolution. For this reason, we incorporated a bidirectional tracking that only connects centre coordinates from two consecutive slices if there is a nearest neighbour match both ways.

The tracking was also updated to handle double detections. Sometimes a fibre is detected as two because the centre of the fibre is slightly black. This is simply handled by merging centre detections if these are closer than the mean fibre radius (reported by fibre manufacturers). Until now, more improvements of the tracking have not been needed. However, we did try replicating the Kalman filter approach for fibre tracking by (Czabaj et al., 2014), which we might apply in the future. It is definitely important to somehow incorporate the direction of the fibre in the tracking search. This will enable precise tracking of fibres that are significantly misaligned due to, for example, compression loading of the composite (see Section 7.2).

Finally, an important extension to the basic pipeline is the determination of diameter sizes. These are computed as explained in Section 4.2 and provide the possibility to have a more complete description of fibre geometry from which contact points between fibres can be estimated and more complete finite element models (FEM) can be created.

Applications of the Fibre Pipeline

We have applied the fibre extraction pipeline covered in Chapter 6 to study the geometry of real UD fibre reinforced polymers. These composite material systems are being employed increasingly in aero (Soutis, 2005), ground transportation and environmentally sustainable energy production systems such as wind turbines (Brøndsted et al., 2005).

The case studies covered in this chapter can be divided into characterising the geometry of real structures and studying the changes in geometry under progressive loading conditions.

7.1 Characterisation of Real Structures

As was explained in Chapter 1, it is important that fibre composites have good stiffness and strength properties. Even though the influence of idealised fibre orientation arrangements on the compression strength has been investigated analytically and numerically (Budiansky and Fleck, 1993; Fleck and Shu, 1995; Kyriakides and Arseculeratne, 1995; Jensen and Christoffersen, 1997), segmentation of individual fibres enables characterisation of real geometries, which have

a larger degree of randomness.

Additionally, extracting fibre diameters opens up possibilities for making more complete studies, where contact points might be important. In particular, examining the spatial distribution of fibre diameters can provide insights into the manufacturing process of fibre bundles.

7.1.1 Fibre Orientation

Fibre misalignment controls compressive failure and fatigue life under aligned cyclic loads (Dadkhah et al., 1995; Fleck and Budiansky, 1991). Additionally, the orientation of individual fibres is related to the stiffness and compression strength properties of the material. In stiffness it is the mean fibre orientation that is important (Krenchel, 1964) whereas in compression how the orientations are distributed spatially is also very important (Liu et al., 2004).

Fibre orientations in 3D can be estimated directly from the intensities of an X-ray tomogram. A good review is contained in (Wirjadi et al., 2016), where four algorithms for local orientation analysis are compared and applied to estimate orientations in fibre reinforced composites. However, the advantage of segmenting individual fibres beforehand is that the precision of the extracted geometry can be assessed prior to computing fibre orientations to ensure accuracy in these estimates. For more details on related work the reader is referred to Papers B and E.

In Paper B we model the orientation for each individual fibre locally by calculating the Δ 'th order difference at every point f of a fibre, as illustrated in Figure 7.1. The orientation at each point f is characterised with the azimuth angle ϕ and the inclination angle θ (called elevation in Paper B). These represent the direction of misalignment and amount of misalignment respectively in the reference coordinate system (x_p, y_p, z_p) , which was calculated as follows. The z_p axis was aligned with the mean fibre orientation (it could have also been aligned with respect to a certain structural direction), calculated as the local fibre orientation for the first point of the fibre and with a length Δ defined by the depth of the imaged field of view (FoV)—also called global fibre orientation. Afterwards, x_p and y_p were rotated so that x_p would be aligned with the direction of maximum variation in fibre orientation, calculated through principal component analysis (PCA) (Abdi and Williams, 2010). More details are given in Paper B.

Two filament wound UD composites (Prabhakaran et al., 2013) were analysed in Paper B, a glass fibre reinforced polymer (GFRP) and a carbon fibre reinforced

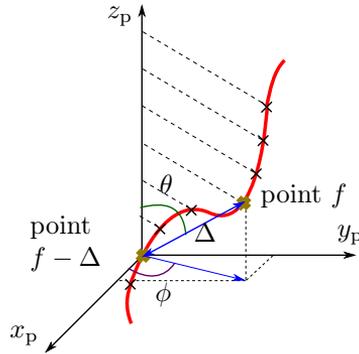


Figure 7.1: Local fibre orientation in point f with respect to the reference coordinate system, where $\theta \in [0, 90]$ and $\phi \in [0, 360]$.

polymer (CFRP) with fibre diameters of $17\ \mu\text{m}$ and $7\ \mu\text{m}$. Figure 7.2 shows a plane view of the extracted fibre trajectories. The length of the fibres in this view gives an idea of the amount of misalignment (θ angle), fibres appear as points if they are completely aligned, whereas the colour encodes the direction in which the fibres are misaligned (ϕ angle).

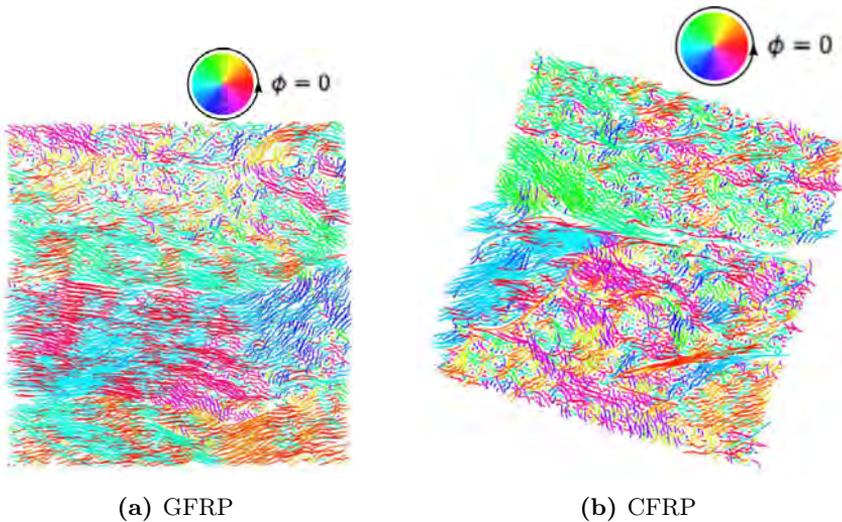


Figure 7.2: A plane view of the fibre trajectories. Fibre lengths increase with the amount of misalignment (angle θ) and colours indicate the direction of misalignment (angle ϕ).

For the GFRP the local orientation was investigated for different lengths Δ and it was found that varying Δ had no effect over the distribution of orientations. In other words, fibres are essentially straight over the considered depth of $D = 1.35$ mm, which is $D/d = 79$ times the size of the glass fibre diameter d .

In Paper B we also estimated the compression strength of the material using the formula by Budiansky (Budiansky and Fleck, 1993) and the mean fibre tilt $\bar{\theta}$ obtained for the imaged field of view (FoV). We found that it was in the same order of magnitude as the compression strength measured for the whole samples which had a gauge section of size $3 \times 10 \times 12$ mm³, compared to the analysed regions of size $1.35 \times 1.35 \times 1.35$ mm³ and $0.58 \times 0.58 \times 0.87$ mm³.

To sum up, the main contributions of Paper B in terms of fibre orientation characterisation are:

- Extraction of individual fibre trajectories, enabling characterisation of orientations for individual fibres.
- Characterisation of local fibre orientations for individual fibres, defined at every point in a fibre by an inclination and an azimuth angle.
- Quantification of orientations through summary statistics and histogram inspection, as well as visualisation of spatial distributions (methodology covered in Section 5.1).
- Demonstration of the method over a GFRP and also a CFRP. The orientations from both materials were compared through by means of exploratory statistics (see Section 5.1).
- Estimation of material properties for both composites from their orientations. In particular the compression strength through the formula by Budiansky. (Budiansky and Fleck, 1993)

The compression strength estimation in Paper B was pretty rough because neither the individual fibre orientations nor the spatial distribution of fibre orientations were considered in the estimation. For more precise estimates of the compression strength see Section 7.1.2.

This will not be the case for the longitudinal stiffness of the sample which can be predicted quite accurately for UD fibre composites from the fibre orientations obtained in Paper B. The rule of mixtures for Young's modulus can be combined

with Krenchel's efficiency factor (Krenchel, 1964) to calculate the longitudinal stiffness of the sample as,

$$E_{\text{composite}} = \eta_o E_f V_f + E_m V_m, \quad (7.1)$$

where V_f and V_m are the volume fractions for the fibre and matrix material; and E_f and E_m are the Young's modulus, with $E_f \gg E_m$. Krenchel's efficiency factor $\eta_o \in [0, 1]$ is 1 when all fibres are completely aligned with the nominal fibre direction. Krenchel's efficiency η_o corrects for the efficiency of the fibre reinforcement which decreases with fibre misalignment and is calculated as,

$$\eta_o = \sum_{\forall i} a_i \cos(\theta_i)^4, \quad (7.2)$$

where $a_i = 1/N$ with N being the number of fibre segments. In this equation we can include the amount of misalignment for each individual fibre as opposed to just using the average fibre misalignment.

7.1.2 Image-based Micro-mechanical Models

In the estimation of the compression strength carried out in Paper B we only used the mean fibre tilt and obviated a very important factor, the spatial distribution of fibre orientations. To obtain more accurate predictions of the compression strength, the extracted fibre geometry can be imported into the FEM software ABAQUS.

The possibility of estimating the transverse stiffness from a 2D FEM model produced using X-ray CT and our pipeline for individual fibre trajectory and diameter extraction is introduced in (Mikkelsen et al., 2016). We extracted individual fibres and diameters from three different bundles of the uni-directional non-crimp fabric epoxy composite illustrated in Figure 7.3a and our collaborator from DTU Wind Energy made a 2D micro-mechanical model for the 45° degree bundles to predict the transverse stiffness. Then, this prediction is compared to analytical estimates obtained with, amongst others, the inverse rule of mixtures.

Besides predicting material properties, simulations can be performed to investigate how certain loading conditions would affect the real geometries we measure. Figure 7.3b shows the distribution of the stresses in the cross-section of a transverse loaded 45° backing bundle.

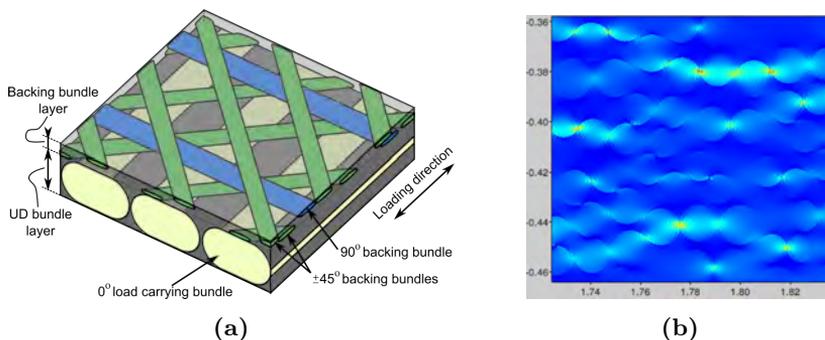


Figure 7.3: Image-based micro-mechanical model of a non-crimp fabric (Mikkelsen et al., 2016). In (a) a schematic of the 3D structure of the non-crimp fabric, where the cross-section of each UD bundles is around 2 mm^2 . In (b) the stress contour plot of a transverse loaded 45° backing bundle for a section of size $0.1 \times 0.1\text{ mm}^2$.

7.2 Evolution of Structures under Changing Load

Paper E brings together high speed X-ray CT with advanced image analysis to follow the changes in UD fibrous materials under progressive loading conditions. By establishing correspondence between the sequence of time-lapse X-ray CT images we are able for the first time to accurately quantify the changes in each individual fibre trajectory in 3D.

We applied the methodology to study a GFRP composite rod under uninterrupted compression loading of the sample in the axial direction—also called nominal fibre direction. When a composite is loaded in compression fibres start to buckle until the matrix material is not able to support the fibres any longer, at this moment a kink-band is formed. In Paper B we demonstrate that we can accurately measure fibre reorientations and determine the progressive deflection of the composite ensemble and of the individual fibres in the steps leading up to fibre micro-buckling and kinking.

The data acquisition was conducted by our collaborator from the School of Materials of the University of Manchester. The scanning was carried out at the TOMCAT synchrotron beamline of the Swiss Light Source where the sample was loaded continuously from 0 N until failure using an in-situ rig, as displayed in Figure 7.4. A scan was taken before loading (0 N), two interrupted scans were acquired at loads 200 N and 600 N, after which continuous scans were taken until sample failure. This occurred at some point in between 895 N and 900 N. We analysed the region of interest (RoI) that was common to the four loading steps

selected for the study in Paper E, shown in Figure 7.5.

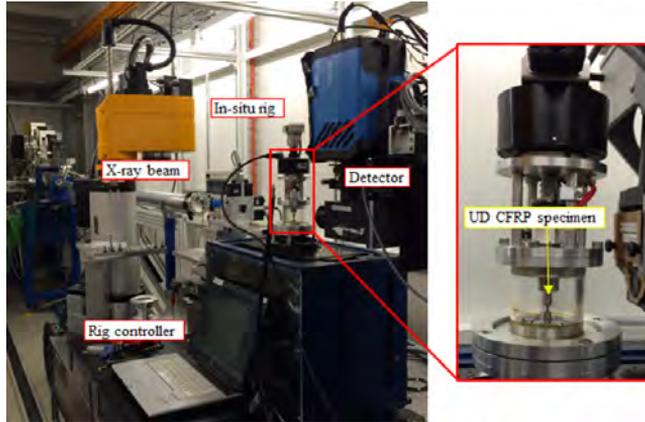


Figure 7.4: Experiment set-up at the TOMCAT synchrotron beamline with in-situ loading rig and ultra-fast imaging (Wang et al., 2016).

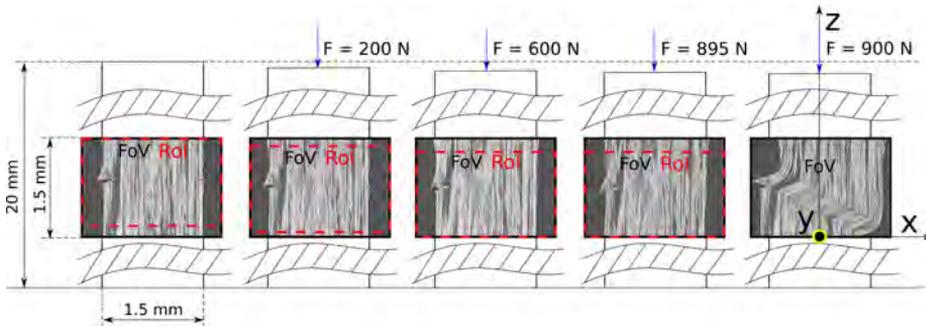


Figure 7.5: X-ray CT longitudinal slices from the centre of the composite rod. The analysed region of interest (RoI) has a height of approximately 1.4 mm and the loading steps selected for the analysis are: i) 0N, ii) 200N, iii) 600N and iv) 895N. The coordinate system is defined according to the loading direction and the kink-band plane.

Fibres trajectories are computed for the four data-sets and in Paper E each of these is approximated by a straight line that connects the points where the fibre enters and leaves the RoI. Fibre curvature will be the subject investigated in our next publication. As was mentioned in Section 7.1.1, in Paper B the z -axis had to be aligned with the nominal fibre direction. This was not necessary in Paper E because the loading was practically aligned with the mean fibre direction already. However, we did rotate the xy plane around the z -axis so as to characterise the orientations and deflections with respect to the ultimate

kink-band location (see Figure 7.5).

Our results show that fibre misalignment is grouped in clusters and these seem to coincide with the spatial clustering of fibres. Additionally, we observe that the whole composite rod tilts towards the direction of the ultimate kink-band plane already at 25% of the failure load, as shown in the *top row* of Figure 7.6. Additionally, we observe that fibres twist around the centre of the composite rod (see Figure 7.6). For more details see Paper E.

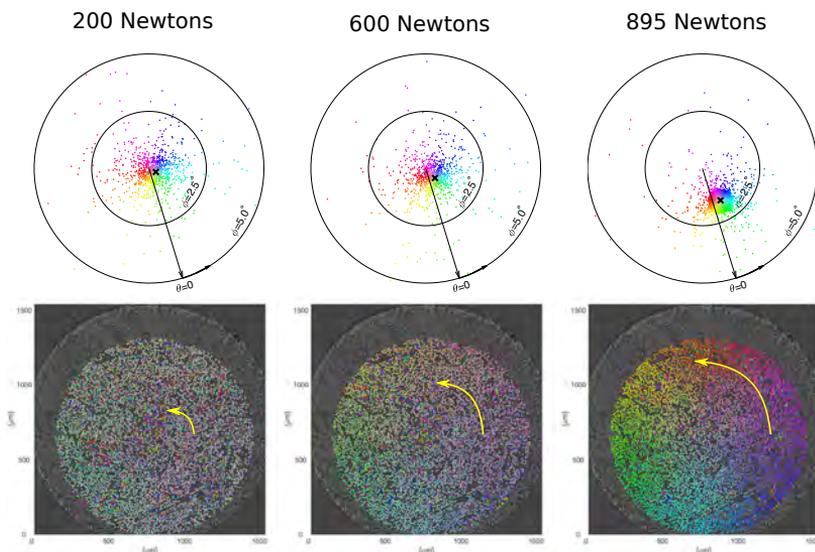


Figure 7.6: Evolution of fibre deflections under load. In the *top row* the scatter plots of the deflection vectors in the horizontal plane. In the *bottom row* the X-ray CT cross-section for the middle slice of the ROI with the deflections displayed at the actual spatial location. The colour-coding has been set relative to the deflection of the composite rod, indicated by the *black* crosses. The *yellow* arrows indicate the twist observed around the centre of the sample, note that these do not represent the absolute magnitude of the twist, which is 1.06° at 895 N.

Summing up, the main contributions of Paper E are:

- Development of a methodology that combines the fibre geometry extraction pipeline with high speed X-ray CT to follow individual fibre reorientations under progressive loading conditions.
- Interpretation of the spatial correlations between fibre orientations and

fibre density in a glass fibre reinforced composite.

- Analysis of fibre deflection progression under uninterrupted compressive loading of a composite rod in the nominal fibre direction.
- Demonstration of the insights that can be obtained by combining individual fibre tracking, high speed X-ray CT and in-situ loading.

Last but not least, it should be noted that Paper E does not focus on explaining the mechanical causes of the observations. For example, we just mention that the twist could be a result of either a preferential direction of misalignment in the as-manufactured sample or a cause of a compressive loading that is not perfectly aligned with the nominal fibre direction.

Thanks to the possibility of extracting individual fibre trajectories, instead of just computing orientations directly from the tomograms, the latter explanation could be discarded by creating a FEM with the geometry extracted from the 0N data-set. The effect of loading the sample in axial compression could be simulated in ABAQUS where we can be sure that the loading is perfectly aligned with the nominal fibre direction.

The 3D Validation Challenge

This chapter discusses the topic of validation, which refers to the process of determining the geometric accuracy of objects. In our case, we would like to assess the accuracy of the pipeline for geometry extraction developed in this thesis work, as well as the accuracy of X-ray CT in providing images that represent the real fibre geometry.

A common approach for validating geometry extraction/segmentations is based on comparing to a reference segmentation of the image, often times created manually by an expert in the field. In the UD fibre geometry considered in this thesis, a manual reference can be created by ourselves because it is easy to determine where are the fibre centres by just looking at the images, without being experts in material science. While it is feasible to make a manual reference of a significant size in 2D, it would take too much time to create a 3D manual reference that is of a size which will be representative. For this reason, qualitative validation based on visual assessment is quite common.

The advantage of a quantitative validation lies in showing the accuracy through quantification, which makes the validation process more believable as well as allowing for comparison. Section 8.1 explains the validation approaches employed in Papers A and B whereas Section 8.2 summarises Paper D, where we

took a statistical approach for validating both our image analysis pipeline for extraction of fibre geometry (Chapter 6) and X-ray CT, the imaging modality for which it was developed.

8.1 Initial Validation of Geometry Extraction

In Paper A we assessed the accuracy in detecting fibre centres from tomographic slices quantitatively by using a manual reference.

In Paper B we demonstrated the accuracy of the centre point detections in the same way as in Paper A whereas we checked the accuracy of the 3D fibre trajectories qualitatively by visual inspection. For this purpose, a Matlab[®] script was developed, the script would slice through the volume showing the estimated centre coordinates. To check that the fibres had been tracked correctly, a few fibres were randomly picked and the centre coordinates belonging to those fibres were given different colours. Using this approach, we found that our tracking algorithm was sometimes merging fibres. This motivated us to improve the tracking as explained in Section 6.2.

Another approach for validating the accuracy of the pipeline is to compare material properties estimated analytically from the fibre orientations to properties measured through mechanical testing. In Paper B we compared the compression strength but we could have also compared the stiffness, which can be obtained numerically as detailed in Section 7.1.1 or by finite element modelling (FEM) as explained in Section 7.1.2.

8.2 Statistical Validation of Analysis Methods and Imaging Modalities

Paper D is concerned with assessing the accuracy of using X-ray CT coupled with image analysis for characterising the micro-structure of UD fibre composites. We perform a statistical analysis where we compare different quantification pipelines (methods) for measuring fibre diameters. Each **method** consists of two steps:

1. **Scanning technique**, which refers to a specific imaging modality and resolution (**mode**). We compare micro-CT at different resolutions (from

synchrotron and laboratory sources) and microscopy, namely optical microscopy (OM) and scanning electron microscopy (SEM). For details on these imaging modalities the reader is referred to Chapter 2.

2. **Image analysis method**, which refers to the method for diameter quantification (**algorithm**). We compare our pipeline for fibre trajectory and diameter determination (see Chapter 6) to the circular Hough transform (see Section 4.1.3).

As explained in Section 2.1.3, there are several random processes in X-ray imaging that can lead to quantum noise. Additionally, the volumetric image is reconstructed from multiple measurements (see Section 2.1.2 on X-ray CT). This indirect way of imaging 3D structures could introduce artefacts and variations in spatial resolution across the volume, as mentioned in Section 2.1.3. Instead, the microscopy modalities directly depict the surface of the sample.

We use the microscopy modalities as a reference for this reason, and also because they were acquired at a very high spatial resolution. Table 8.1 shows the resolutions used for the six modes under comparison. As explained in Paper D, the pixel/voxel sizes given by the instruments deviate a little from those measured relative to the OM pixel size through the image registration employed to align the different scans. To compare our diameter measurements to the manufacturer reported average fibre diameter of $17\ \mu\text{m}$, we transform the diameters from pixel/voxel sizes into μm using the estimated pixel/voxel sizes.

Table 8.1: Instrument reported and measured pixel sizes for the different scans.

Modality	Reported	Estimated
SEM	$0.1852\ \mu\text{m}$	$0.1882\ \mu\text{m}$
OM	$0.2908\ \mu\text{m}$	reference
SRCT	$0.6500\ \mu\text{m}$	$0.6440\ \mu\text{m}$
XCT _H	$1.0376\ \mu\text{m}$	$1.0356\ \mu\text{m}$
XCT _M	$1.6856\ \mu\text{m}$	$1.6835\ \mu\text{m}$
XCT _L	$2.8059\ \mu\text{m}$	$2.8082\ \mu\text{m}$

Fibre centres and diameters are detected from the field of view (FoV) shown in Figure 8.1. We compute the fibre centres and diameters in 2D for all six modes using the circular Hough transform (see Section 4.1.3). For the 3D scans (micro-CT) we also use our pipeline. We process 60 slices in a $\sim 60\ \mu\text{m}$, where we determine one diameter for each fibre trajectory by averaging the 60 diameters estimates, calculated as explained in Section 4.2. Paper D contains an illustration where we plot the fibre centre and diameter estimates over the image, so one can visually assess how well these estimates fit the data.

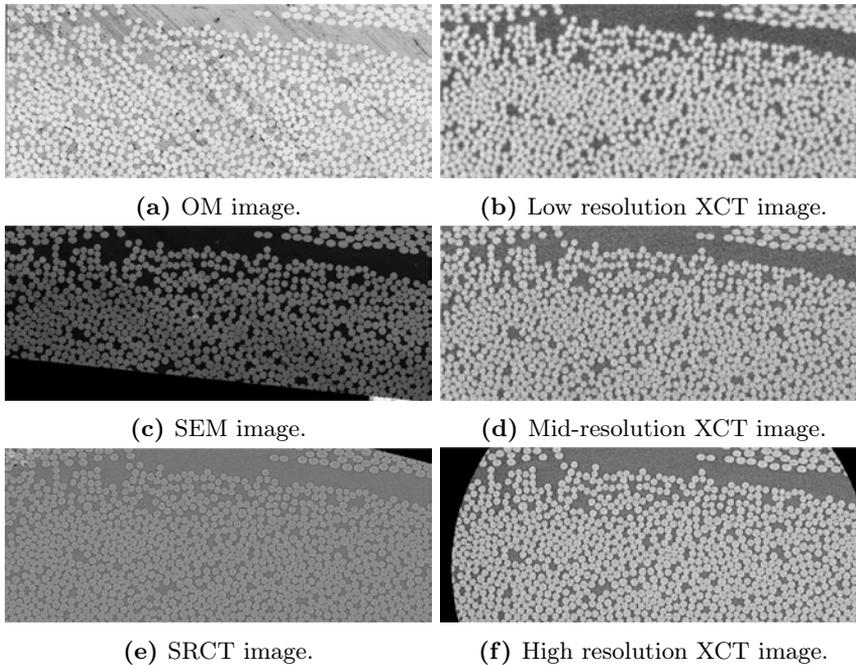


Figure 8.1: All six scans aligned after rotating and translating. Note that scaling has not been applied to show the actual resolution of the images.

We segmented $N = 757$ corresponding fibres for all $n = 10$ different quantification pipelines (methods). As mentioned above, the micro-CT modes are analysed with both methods (4 modes \times 2 methods) and the microscopy modes are analysed with one method (2 modes \times 1 method).

We consider a univariate approach and a multivariate approach, for more details see Paper D. In the univariate approach we describe the variation between fibre diameters as deterministic contributions from the methods (or algorithm and modes) adjusting for the different fibre contributions.

If we want to distinguish between algorithms and modes we use a three-way factorial design model and, if we just focus upon the different methods, we use a two-way layout. These models are written in Section 5.2.2 and Paper D.

With the two-way approach we test the difference between the average diameters given by each quantification pipeline, this resulted in a grouping of the methods. The groups that are closest to the production specification $d = 17 \mu\text{m}$ are formed

by the quantification pipelines listed below. The group that is furthest away is formed by the XCT_L analysed with the Hough transform. For the other groups cf. Table 3 in Paper D. From the three-way approach we found that both mode and algorithm have an effect over the measured diameter.

List of methods inside the group that is closest to the fibre production specification:

1. SEM + Hough.
2. SRCT + our pipeline.
3. SRCT + Hough.
4. XCT_H + our pipeline.
5. XCT_M + our pipeline.

The dependence of fibre diameters with spatial coordinates was investigated by regression analysis on the spatial coordinates using an ANOVA to choose the model that fits the best, as explained in Section 5.2. We found that all methods capture a spatial variation in diameters, so it must be that the variation is in the physical properties of the sample. Regression on spatial coordinates could therefore be an adequate method for learning about the fibre manufacturing process.

In the multivariate approach, amongst others, we compared corresponding fibre diameters across modes through scatter plots. In the first place, we compared the high resolution modalities (4 modes) from which we found that OM has a bias towards smaller diameters (see Paper D). In the second place, we compared micro-CT modalities (4 modes). From the comparison across micro-CT modalities in Figure 8.2 we can confirm that there is a good correspondence between measurements obtained with the same algorithm whereas there is a larger spread when comparing measurements obtained with different algorithms. While the correspondence between measurements obtained with our algorithm is still good for low resolutions, the correspondence between measurements obtained with Hough transform decreases rapidly when the resolution of micro-CT is decreased.

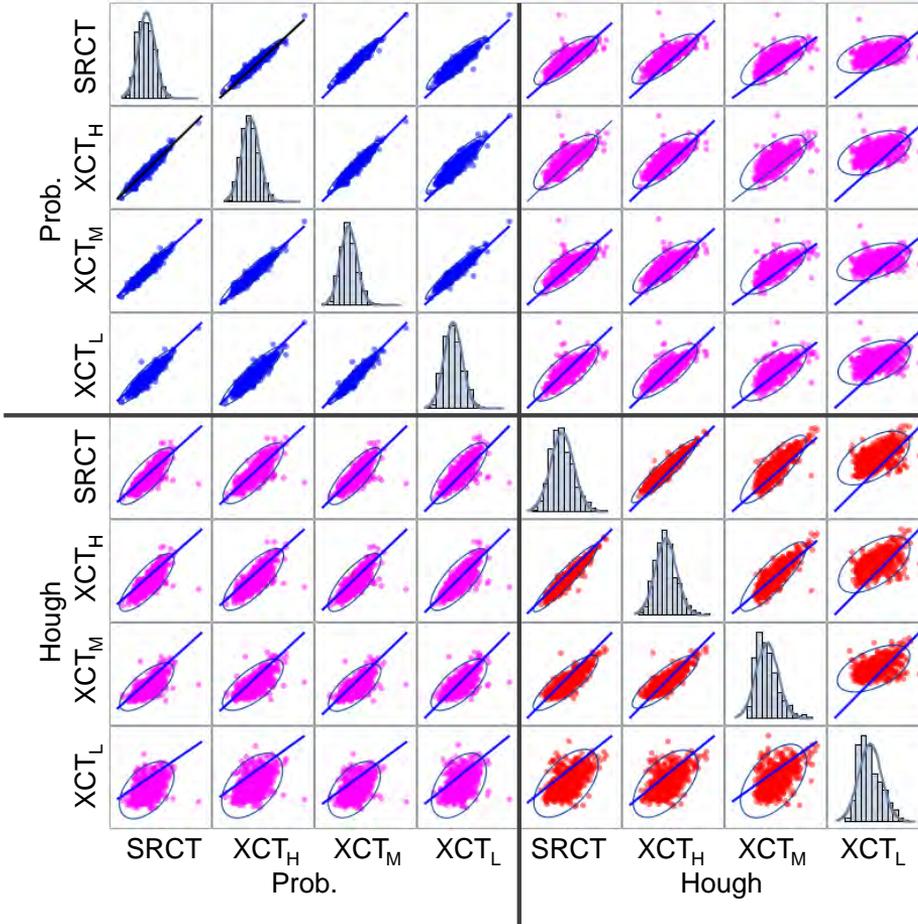


Figure 8.2: Scatter plot for diameters obtained for the X-ray imaging. Methods employed row-wise are SRCT, and high, mid, and low resolution XCT, first with our pipeline and then with the Hough transform algorithm. The blue line is the identity line and the blue ellipse is the 99% prediction ellipse for a new observation.

To sum up, the main findings of Paper D are:

- Demonstration of the precision that our fibre trajectory and diameter determination methodology has in characterising UD composites with micro-CT, even at low resolutions. Low resolutions enable the analysis of larger volumes, which are more representative of the actual fibre composite.

- A regression on spatial coordinates is a good method for assessing physical diameter variations. This could provide insights into the fibre manufacturing process.

Now that we have confirmed the validity of our fibre trajectory and diameter determination method for characterizing UD fibre composites with low resolution X-ray CT, we could analyse a whole fibre bundle and estimate the amount of fibre material in the bundle. A further quantitative validation could be a comparison between this estimation and the manufacturer reported tex-value [g km^{-1}], which gives the amount of fibre material inside a complete bundle through the weight of the bundle per length.

Conclusions

This thesis has presented statistical image analysis tools mainly for characterising the micro-structure of fibrous materials in 3D. The methodology is based on extracting fibre geometry from tomograms acquired through X-ray imaging. It has been demonstrated that extracting the geometry of real fibre arrangements opens up numerous possibilities for characterising and understanding the behaviour of fibre composites under real life conditions. In addition, the accuracy in characterising UD fibre composites with the developed methodology and X-ray tomography has been proven quantitatively.

The core methodology of this work is a pipeline that can extract the geometry of individual fibres from X-ray tomograms. The strength of this methodology is its ability to extract individual fibres from volumetric scans where there is a high fibre volume fraction. The method is robust to noise and can analyse low contrast scans, meaning that long scan times can be avoided and a broad range of fibre materials can be investigated. Additionally, it has been demonstrated that the methodology is highly precise for analysing large volumes, as it can deal with low spatial resolutions. Due to its ability to accurately analyse representative volumes and a wide range of materials, the fibre extraction methodology has great potential in providing insights for understanding the relationship between fibre micro-structure and material properties.

The pipeline, which has been under development throughout the thesis work, detects individual fibre centres through a probabilistic segmentation framework based on a dictionary that is able to learn image patterns. We have demonstrated the flexibility and accuracy of the dictionary framework in segmenting complex image structures. The speed of the method has been improved during the course of the thesis work and currently the segmentation framework falls in

the semi-supervised setting, as the amount of user interaction is minimal and can be adapted to the specific data requirements. Therefore, the dictionary segmentation framework can be a very useful tool in quantifying structures in complex volumetric data, such as tomograms of energy materials.

The applicability of the pipeline has been show-cased over real composites in collaboration with material scientists and mechanical engineers. The first application entails characterisation of real geometries through fibre orientations and diameters. Material properties such as stiffness or strength can be computed analytically from individual fibre orientations. For more precise estimates, micro-mechanical models can be created from the image-based fibre geometry. Besides estimating material properties, the behaviour of the composite material under load can be simulated through finite element modelling.

The other main application of the pipeline is in studying the behaviour of fibre composites under real loading conditions. By combining ultra-fast X-ray CT, in-situ loading of a sample and the fibre geometry extraction pipeline, we have been able for the first time to accurately quantify changes in fibre trajectories fibre by fibre. Quantifying the changes in composite micro-structure with a high degree of accuracy both in space, time and load can provide valuable insight into the behaviour of materials under real-life loading conditions, which will lead to a better understanding of the very complex damage mechanisms that affect composite materials. In particular, valuable insight has been provided with regards to the behaviour of a UD fibre composite under compression loading. Our work shows for the first time that fibres start to deflect at very low loads and in a direction which is related to the final damage.

In conclusion, this thesis has demonstrated the potential in volumetric imaging based on X-ray CT for understanding the relationship between the micro-structure of a material and its macroscopic behaviour. With current developments in imaging methods, it will be possible to image material structures in higher spatial and temporal resolution, and the need for simple, computationally efficient, while also accurate methods, is urgent. The work in this thesis has demonstrated an approach to meet these needs.

Bibliography

- H. Abdi and L. J. Williams. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):433–459, 2010.
- T. J. Atherton and D. J. Kerbyson. Size invariant circle detection. *Image and Vision Computing*, 17(11):795–803, 1999.
- M. Beister, D. Kolditz, and W. A. Kalender. Iterative reconstruction methods in X-ray CT. *Physica Medica*, 28(2):94–108, 2012.
- A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Occam’s razor. *Information Processing Letters*, 24(6):377–380, 1987.
- F. E. Boas and D. Fleischmann. CT artifacts: causes and reduction techniques. *Imaging in Medicine*, 4(2):229–240, 2012.
- P. Brøndsted, H. Lilholt, and Aa. Lystrup. Composite materials for wind power turbine blades. *Annual Reviews Materials Research*, 35:505–538, 2005.
- B. Budiansky and N.A. Fleck. Compressive failure of fibre composites. *Journal of the Mechanics and Physics of Solids*, 41(1):183–211, jan 1993.
- T. Burton, N. Jenkins, D. Sharpe, and E. Bossanyi. Aerodynamics of horizontal axis wind turbines. *Wind Energy Handbook, Second Edition*, pages 39–136, 2001.
- T. S. Curry, J. E. Dowdey, and R. C. Murry. *Christensen’s physics of diagnostic radiology*. Lippincott Williams & Wilkins, 1990.
- M. W. Czabaj, M. L. Riccio, and W. W. Whitacre. Numerical reconstruction of graphite/epoxy composite microstructure based on sub-micron resolution X-ray computed tomography. *Composites Science and Technology*, 105:174–182, 2014.

- M. S. Dadkhah, B. N. Cox, and W. L. Morris. Compression-compression fatigue of 3D woven composites. *Acta Metallurgica et Materialia*, 43(12):4235–4245, 1995.
- A. L. Dahl and R. Larsen. Learning dictionaries of discriminative image patches. In *22nd British Machine Vision Conference*, 2011.
- P. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14(3):227–248, 1980.
- N. F. Dow and I. J. Gruntfest. Determination of most needed potentially possible improvements in materials for ballistic and space vehicles. *General Electric Company, Air Force Contract AF*, 4(647):269, 1960.
- D. E. Dudgeon and R. M. Mersereau. Multidimensional digital signal processing. *Prentice-Hall Signal Processing Series, Englewood Cliffs: Prentice-Hall*, 1984, 1984.
- M. Elad. From exact to approximate solutions. In *Sparse and Redundant Representations*, pages 79–109. Springer, 2010.
- V. Etter, I. Jovanovic, and M. Vetterli. Use of learned dictionaries in tomographic reconstruction. In *Proceedings of SPIE*, volume 8138, 2011.
- N. A. Fleck and B. Budiansky. Compressive failure of fibre composites due to microbuckling. In *Inelastic Deformation of Composite Materials/IUTAM Symposium, Troy, New York*, pages 235–73, 1991.
- N. A. Fleck and J. S. Shu. Microbuckle initiation in fibre composites: A finite element study. *Journal of the Mechanics and Physics of Solids*, 43(12):1887–1918, 1995.
- D. A. Freedman. *Statistical models: theory and practice*. Cambridge University Press, 2009.
- T. D. Government. The danish climate policy plan – towards a low carbon society. Technical report, Danish Energy Agency, 2013.
- E. G. Guynn, W. L. Bradley, and O. O. Ochoa. A parametric study of variables that affect fiber microbuckling initiation in composite laminates: Part 1—analyses, part 2—experiments. *Journal of Composite Materials*, 26(11):1594–1643, 1992.
- M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
- G. T. Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Science & Business Media, 2009.

- P. D. Hill. Kernel estimation of a distribution function. *Communications in Statistics - Theory and Methods*, 14(3):605–620, 1985.
- Z. Huang. Brightness and coherence of synchrotron radiation and FELs. Technical report, SLAC National Accelerator Laboratory, 2013.
- H. M. Jensen and J. Christoffersen. Kink band formation in fiber reinforced materials. *Journal of the Mechanics and Physics of Solids*, 45(7):1121–1136, 1997.
- K. M. Jespersen and L. P. Mikkelsen. Three dimensional fatigue damage evolution in non-crimp glass fibre fabric based composites used for wind turbine blades. *Composites Science and Technology*, 2017.
- K. M. Jespersen, J. Zangenberg, T. Lowe, P. J. Withers, and L. P. Mikkelsen. Fatigue damage assessment of uni-directional non-crimp fabric reinforced polyester composite using X-ray computed tomography. *Composites Science and Technology*, 136:94–103, 2016.
- B. Julesz. Textons, the elements of texture perception, and their interactions. *Nature*, 290(5802):91–97, 1981.
- C. Karunakaran and D. S. Jayas. X-ray imaging. In *Imaging with Electromagnetic Spectrum*, pages 33–55. Springer, 2014.
- H. Krenchel. Fibre reinforcement; theoretical and practical investigations of the elasticity and strength of fibre-reinforced materials. 1964.
- S. Kyriakides and R. Arseculeratne. On the compressive failure of fiber reinforced composites. *International Journal of Solids and Structures*, 32(617):689–738, 1995.
- T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43(1):29–44, 2001.
- T. Lindeberg. Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of Applied Statistics*, 21(1-2):225–270, 1994.
- D. Liu, N. A. Fleck, and M. P. F. Sutcliffe. Compressive strength of fibre composites with random fibre waviness. *Journal of the Mechanics and Physics of Solids*, 52(7):1481–1505, jul 2004.
- D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. Oakland, CA, USA., 1967.

- J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Discriminative learned dictionaries for local image analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- E. Maire and P. J. Withers. Quantitative X-ray tomography. *International Materials Reviews*, 59(1):1–43, 2014.
- J. Malik, S. Belongie, T. Leung, and J. Shi. Contour and texture analysis for image segmentation. *International Journal of Computer Vision*, 43(1):7–27, 2001.
- P. S. Mann. *Introductory statistics*. John Wiley & Sons, 2007.
- D McMullan. Scanning electron microscopy 1928–1965. *Scanning*, 17(3):175–185, 1995.
- G. Michael. X-ray computed tomography. *Physics Education*, 36(6):442, 2001.
- L. P. Mikkelsen, M. J. Emerson, K. M. Jespersen, V. A. Dahl, K. Conradsen, and A. B. Dahl. X-ray based micromechanical finite element modeling of composite materials. In *In Proceedings of 29th Nordic Seminar on Computational Mechanics*, 2016.
- P. M. Moran, X. H. Liu, and C. F. Shih. Kink band formation and band broadening in fiber composites under compressive loading. *Acta Metallurgica et Materialia*, 43(8):2943–2958, 1995.
- M. Muja and D. G. Lowe. Flann, fast library for approximate nearest neighbors. In *International Conference on Computer Vision Theory and Applications*, volume 3. INSTICC Press, 2009.
- B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311–3325, 1997.
- N. Otsu. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1):62–66, 1979.
- S. Pimenta, R. Gutkin, S. T. Pinho, and P. Robinson. A micromechanical model for kink-band formation: Part i—experimental study and numerical modelling. *Composites Science and Technology*, 69(7):948–955, 2009.
- R. T. D. Prabhakaran, T. L. Andersen, C. M. Markussen, B. Madsen, and H. Lilholt. Tensile and compression properties of hybrid composites—a comparative study. In *Proceeding of the International Conference on Composite Materials*, pages 1029–1035, 2013.
- R. Pyrz. Application of x-ray microtomography to the study of polymer composites. Technical report, SAE Technical Paper, 1999.

- J. Radon. Über die bestimmung von funktionen durch ihre integralwerte laengs gewisser mannigfaltigkeiten. *Berichte Saechsische Acad. Wissenschaft. Math. Phys., Klass*, 69:262, 1917.
- G. Requena, G. Fiedler, B. Seiser, P. Degischer, M. Di Michiel, and T. Buslaps. 3D-quantification of the distribution of continuous fibres in unidirectionally reinforced composites. *Composites Part A: Applied Science and Manufacturing*, 40(2):152–163, 2009.
- W. C. Röntgen. Ueber eine neue art von strahlen. *Annalen der Physik*, 300(1): 1–11, 1898.
- C. A. Schneider, W. S. Rasband, and K. W. Eliceiri. Nih image to imagej: 25 years of image analysis. *Nature methods*, 9(7):671–675, 2012.
- J. Shotton, M. Johnson, and R. Cipolla. Semantic texton forests for image categorization and segmentation. In *Computer vision and pattern recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- C. Soutis. Compressive strength of unidirectional composites: measurement and prediction. In *Composite Materials: Testing and Design, Thirteenth Volume*. ASTM International, 1997.
- C. Soutis. Carbon fiber reinforced plastics in aircraft construction. *Materials Science and Engineering: A*, 412(1):171–176, 2005.
- S. R. Stock. X-ray microtomography of materials. *International Materials Reviews*, 44(4):141–164, 1999.
- J. S. Vestergaard, A. L. Dahl, P. Holm, and R. Larsen. Pipeline for tracking neural progenitor cells. In *International MICCAI Workshop on Medical Computer Vision*, pages 155–164. Springer, 2012a.
- J. S. Vestergaard, A. L. Dahl, R. Larsen, and A. A. Nielsen. Classification of polarimetric sar data using dictionary learning. In *In Proceedings Image Signal Processsing and Remote Sensing XVIII*, volume 8537, pages 85370X–1, 2012b.
- Y. Wang, S. C. Garcea, T. Lowe, E. Maire, C. Soutis, and P. J. Withers. Ultra-fast time-lapse synchrotron radiographic imaging of compressive failure in CFRP. In *17th European Conference on Composite Materials, Munich, Germany*, 2016.
- Y. Wang, T. L. Burnett, Y. Chai, C. Soutis, P. J. Hogg, and P. J. Withers. X-ray computed tomography study of kink bands in unidirectional composites. *Composite Structures*, 160:917–924, 2017.

- O. Wirjadi, K. Schladitz, P. Easwaran, and J. Ohser. Estimating fibre direction distributions of reinforced composites from tomographic images. *Image Analysis & Stereology*, 35(3):167–179, 2016.
- M. R. Wisnom and J. W. Atkinson. Fibre waviness generation and measurement and its effect on compressive strength. *Journal of reinforced plastics and composites*, 19(2):96–110, 2000.
- H. K. Yuen, J. Princen, J. Illingworth, and J. Kittler. Comparative study of hough transform methods for circle finding. *Image and vision computing*, 8(1):71–77, 1990.

Part III

Publications

Dictionary Based Segmentation in Volumes

Published in *Scandinavian Conference on Image Analysis*, June 2015.

Dictionary Based Segmentation in Volumes

Monica Jane Emerson¹() , Kristine Munk Jespersen²,
Peter Stanley Jørgensen³, Rasmus Larsen¹, and Anders Bjorholm Dahl¹

¹ DTU Compute, Lyngby, Denmark
{[monj](mailto:monj@dtu.dk),[r1ar](mailto:r1ar@dtu.dk),[abda](mailto:abda@dtu.dk)}@dtu.dk

² DTU Wind Energy, Roskilde, Denmark
kmun@dtu.dk

³ DTU Energy, Roskilde, Denmark
psjq@dtu.dk

Abstract. We present a method for supervised volumetric segmentation based on a dictionary of small cubes composed of pairs of intensity and label cubes. Intensity cubes are small image volumes where each voxel contains an image intensity. Label cubes are volumes with voxel-wise probabilities for a given label. The segmentation process is done by matching a cube from the volume, of the same size as the dictionary intensity cubes, to the most similar intensity dictionary cube, and from the associated label cube we get voxel-wise label probabilities. Probabilities from overlapping cubes are averaged and hereby we obtain a robust label probability encoding. The dictionary is computed from labeled volumetric image data based on weighted clustering. We experimentally demonstrate our method using two data sets from material science – a phantom data set of a solid oxide fuel cell simulation for detecting three phases and their interfaces, and a tomogram of a glass fiber composite used in wind turbine blades for detecting individual glass fibers.

Keywords: Volume segmentation · Materials images · X-ray tomography · Learning dictionaries · Glass fiber segmentation

1 Introduction

High resolution volumetric scanning has become a widely used technique in areas like material science and medicine, and automated quantification methods are necessary in order to obtain size and shape measures from these data. We present a method for supervised segmentation of volumetric data. The method is trained from manual annotations, and these annotations make the method very flexible, which we demonstrate in our experiments.

Our method infers label information locally by matching the pattern in a neighborhood around a voxel to a dictionary, and hereby accounts for the volume texture. Texture segmentation has been widely addressed in 2D [4, 8, 19], whereas volumetric texture segmentation has received less attention [1]. This fact could be due to the extra computational effort introduced in 3D. However, 3D texture segmentation is highly appropriate for quantifying size and shape in 3D data. Applications in this paper are from energy material science.

Denmark is pursuing to obtain 100% of the energy from renewable resources by 2050 [6]. For this reason, it is important to develop renewable energy technologies, which involves investigating material properties to ensure efficiency and lifetime. Imaging the 3D micro-structure can be essential in characterizing and understanding such properties, e.g. the geometric configuration of a solid oxide fuel cell or the fiber geometry in a wind turbine blade material, which are the applications we investigate here.

Our applications contain three spatial dimensions, but data could also be two spatial dimensions and time as the third. If it were to be the latter, the evolution of the material's micro-structure under certain conditions (e.g. temperature or tension) could be investigated. Moreover, the flow of different elements (e.g. gas or liquid) through porous materials can be investigated [17]. Other applications are segmentation of anatomical structures in medical imaging [1, 16], seismic facies analysis [14] or crystallography [15]. Many of the volumetric texture segmentation algorithms are extensions of common 2D techniques [2] or 2D segmentation propagation approaches for segmenting 3D materials [18].

Wind turbines blades commonly use glass fiber composites for the load carrying parts of the blades, for which the fatigue damage mechanisms are not well understood. Wind turbine blades have long expected lifespans where they experience a high number of load cycles, which gives rise to fatigue damage evolution. In addition, the blade lengths are being increased because the power output of a wind turbine is proportional to the blade length squared. As fatigue is one of the main limiting factors of designing longer blades, improving the understanding of fatigue damage evolution in glass fiber composites is important [7, 13]. Here microstructure analysis using imaging is an important tool, both for material characterization and modeling.

Solid oxide fuel cells (SOFC) operate by oxidizing a fuel to produce electricity and heat. The electrodes of an SOFC are typically two-phase porous systems. The two solid phases are responsible for electron and ion conduction and the pores allow transport of gaseous reactants and products to and from the electrochemically active sites at the triple phase boundaries (TPB). The chemical reactions can only take place at the TPBs where there is access for electrons, ions and gases through the corresponding three phases. The performance of an SOFC is thus strongly dependent on the density of TPBs in the electrodes and on how easily, electrons ions and gases can be transported to and from the TPBs [10]. The microstructure of an SOFC can only be indirectly controlled through a complicated interaction between powder particle sizes, casting methods and sintering temperatures. 3D characterization of the micro-structure is thus becoming an increasingly important tool to correlate the characteristics of the micro-structure to the cell performance and the production recipes.

In this paper we extend the 2D segmentation algorithm in [5] to 3D. The method is based on a dictionary of image patches and corresponding label patches. Here we replace the image patches with volume cubes and investigate the effects on the segmentation. We see an improved performance in some of our experiments by extending to 3D, but at the expense of longer computation times. However, we obtain close to perfect segmentation of individual glass fibers

in wind turbine blades, and also high performance in segmenting solid oxide fuel cell data.

The paper is organized as follows. In Section 2, we provide a description of the algorithm. Section 3 explains the data sets and materials which are utilized for the comparative study. In Section 4, the focus is on the results, where three methods (method from literature, 2D dictionary and 3D dictionary) are compared for each data set. Finally, Section 5 draws some conclusions.

2 Method

In the process of extending [5] to 3D, we have however changed some parts and therefore we include a description of the entire method despite the overlap with [5]. The method is based on a dictionary of small intensity cubes coupled with label cubes, and the dictionary is learned using weighted clustering. First we describe the dictionary, then we explain the problem we are optimizing, and finally we describe an algorithm to compute the dictionary and the method for inferring label information to an unlabeled volume.

The dictionary is based on annotated training data, where the annotation assigns each voxel in the training volume to one class label. Given a volume $V_I : \Omega \rightarrow \mathbb{R}$ where $\Omega \subset \mathbb{R}^3$ and an annotation of that volume $V_L : \Omega \rightarrow \mathbb{N}$ with labels $l = 1, \dots, k$, we want to build a dictionary $\mathbf{D} = (\mathbf{D}_I, \mathbf{D}_L)$ consisting of the intensity dictionary $\mathbf{D}_I \in \mathbb{R}^{m \times n}$ and the associated label dictionary $\mathbf{D}_L \in \mathbb{R}^{km \times n}$. Each column in \mathbf{D}_I contains a vector representation of small intensity cubes of side length M where $m = M^3$, and n is the number of dictionary elements. The columns of \mathbf{D}_L contain vectorized label cubes represented as probabilities of labels. Therefore, each label vector contains km elements, where the first m elements are the probabilities of label 1, the next m elements are for label 2, etc. We get

$$\sum_{l=0}^{k-1} \mathbf{d}_{Li}(lm + \tau) = 1 \text{ for all } \tau \in \{1, \dots, m\},$$

where \mathbf{d}_{Li} is column i in \mathbf{D}_L .

An ideal dictionary would be

$$\hat{\mathbf{D}} = \arg \min_{\mathbf{D}} \sum_{\eta=1}^o \left(\hat{\lambda} \|\mathbf{d}_{Li(\eta)} - \mathbf{v}_{L\eta}\|_2^2 + \|\mathbf{d}_{Ii(\eta)} - \mathbf{v}_{I\eta}\|_2^2 \right), \quad (1)$$

where $\mathbf{v}_{I\eta}$ and $\mathbf{v}_{L\eta}$ are the η 'th intensity cube and label cube from V_I , V_L respectively with $\eta = 1, \dots, o$ where o is the number of intensity cubes in V_I and $\hat{\lambda}$ is a scaling factor. $i(\eta)$ is the index of the nearest intensity dictionary element

$$i(\eta) = \arg \min_i \|\mathbf{d}_{Ii} - \mathbf{v}_{I\eta}\|_2^2. \quad (2)$$

The first norm in (1) $\|\mathbf{d}_{Li(\eta)} - \mathbf{v}_{L\eta}\|_2^2$ measures the Euclidean distance between the label dictionary element, and label cube and the second norm in (1)

$\|\mathbf{d}_{I_i(\eta)} - \mathbf{v}_{I\eta}\|_2^2$ is the Euclidean distance between the intensity dictionary element and intensity cube. Our aim is to minimize both simultaneously because we hereby obtain a discriminative dictionary and good clustering properties. We will later show that this is advantageous when using the dictionary for segmenting an unlabeled image.

Finding a solution to (1) is a hard problem. Without the first norm concerning the labels it is a k -means clustering problem, which is NP-hard, and we have not found a solution with the addition of the label information. So, based on (1) we suggest a heuristic clustering algorithm which has given good performance in our experiments. The basic idea is to estimate a weight from the label information and use that in a weighted k -means clustering approach.

The algorithm for building the dictionary is based on iteratively updating a set of dictionary elements. Initially a random set of ξ associated intensity and label cubes are selected as $\hat{\mathbf{v}}_j = (\hat{\mathbf{v}}_{I_j}, \hat{\mathbf{v}}_{L_j})$ from the annotated training volume and vectorized, where $j = 1, \dots, \xi$. A subset of n patches are randomly selected as the initial dictionary \mathbf{D}^0 , where $n < \xi$. New dictionary elements are now estimated iteratively as

$$\mathbf{d}_{I_i}^{t+1} = \frac{1}{v_i} \sum_{\kappa \in S_i} (1 - \lambda \|\mathbf{d}_{L_i}^t - \hat{\mathbf{v}}_{L\kappa}\|_2) \hat{\mathbf{v}}_{I\kappa}, \quad (3)$$

where $\mathbf{d}_{I_i}^{t+1}$ is the intensity dictionary element at iteration number $t + 1$. S_i is the set of indices with intensity cubes closest to dictionary element i

$$\kappa \in S_i \text{ s.t. } \kappa = \arg \min_j \|\mathbf{d}_{I_i} - \mathbf{v}_{Ij}\|_2^2.$$

The normalization factor is estimated as

$$v_i = \sum_{\kappa \in S_i} (1 - \lambda \|\mathbf{d}_{L_i} - \hat{\mathbf{v}}_{L\kappa}\|_2). \quad (4)$$

This approach gives high weight to training samples with labels similar to the dictionary element and low weight to dissimilar samples. The label dictionary elements are estimated as average labels for the cluster as

$$\mathbf{d}_{L_i}^{t+1} = \frac{1}{|S_i|_0} \sum_{\kappa \in S_i} \hat{\mathbf{v}}_{L\kappa}, \quad (5)$$

where $|S_i|_0$ is the cardinality of S_i . In our experiments we have seen a satisfactory result with little change in the dictionary after approximately 10 iterations.

The dictionary is used for segmenting a volume by building a label probability volume. This is done by matching the intensity dictionary elements to the volume we want to segment, and adding the associated label dictionary elements to an empty label probability volume.

Given a volume $U_I : \Omega \rightarrow \mathbb{R}$ that we want to segment with $U_I \in \mathbb{R}^{x \times y \times z}$ we compute a label probability volume $U_L : \Omega \rightarrow \mathbb{R}$ with $U_L \in \mathbb{R}^{x \times y \times z \times k}$. Initially we set U_L to having all elements zeros. We can extract vectorized intensity cubes

\mathbf{u}_{Ih} of the same spatial size as the dictionary elements, where $h = 1, \dots, \rho$ is the number of possible cubes with side length M , so e.g. for an odd M we get $\rho = (x - M + 1)(y - M + 1)(z - M + 1)$, which is a little less than the number of voxels due to the volume boundaries. Each \mathbf{u}_{Ih} is matched to the nearest intensity dictionary element using Euclidian distance. For a given intensity vector \mathbf{u}_{Ih} we get the nearest dictionary element

$$i(h) = \arg \min_i \|\mathbf{d}_{Ii} - \mathbf{u}_{Ih}\|_2^2. \quad (6)$$

From this we take the corresponding label dictionary element $\mathbf{d}_{Li(h)}$ and add it to the label volume U_L at the coordinates of h 'th cube extracted from U_I for each of the k labels. When the probabilities are added we weigh them using a Gaussian weight function with standard deviation σ centered at the cube. After adding the probabilities in U_L up, we simply normalize by dividing the sum over label probabilities for each voxel to make the probabilities sum to one over all k labels.

Some smoothing at boundaries occurs which especially affects small features. In order to account for that, we estimate label-wise weights on an annotated validation set where we minimize the difference between the obtained probability volume and the annotation. We are given an annotated volume Q_L and the computed label probability volume P_L using a trained dictionary. Then we rearrange these volumes to \mathbf{Q}_L and \mathbf{P}_L such that each row contains the voxel-wise probabilities and each column represents the labels with each row summing to 1, i.e. $\mathbf{Q}_L(r, l) \in \{0, 1\}$, $\mathbf{P}_L(r, l) \in [0, 1]$ and $\sum_{c=1}^k \mathbf{Q}_L(r, l) = \sum_{c=1}^k \mathbf{P}_L(r, l) = 1$, where $\mathbf{Q}_L(r, l)$ and $\mathbf{P}_L(r, l)$ are elements from \mathbf{Q}_L and \mathbf{P}_L at row r and column l respectively. We want to find the weight matrix $\mathbf{W} \in \mathbb{R}^{k \times k}$ that minimizes

$$\mathbf{W} = \arg \min_{\mathbf{W}} \|\mathbf{Q}_L - \mathbf{P}_L \mathbf{W}\|_2^2,$$

where the solution is found as

$$\mathbf{W} = (\mathbf{P}_L^T \mathbf{P}_L)^{-1} \mathbf{P}_L^T \mathbf{Q}_L.$$

The voxel-wise probability of the final segmentation is obtained as

$$\tilde{\mathbf{u}}_L(\mathbf{x}) = \mathbf{u}_L(\mathbf{x}) \mathbf{W}, \quad (7)$$

where $\mathbf{u}_L(\mathbf{x}) \in \mathbb{R}^k$ is a vector of label probabilities of the voxel from the spatial position $\mathbf{x} = (x, y, z)^T$ in the intensity volume U_I .

3 Materials

Two data sets are employed for the comparative study including:

1. Real data from glass fiber used for wind turbine blades.
2. Phantom data of solid oxide fuel cells.

3.1 Glass Fiber

The data set was obtained through 3D X-ray computed tomography imaging. The scanned sample is a cut-out of a fatigue test specimen and the dimensions of approximately $5 \times 5 \times 10$ mm. The material considered is a uni-directional (UD) glass fiber/polyester composite used in the load carrying beam of a wind turbine blade. Uni-directional in this case means that the fibers are aligned in one main direction, making the composite strong in one direction, and weak in other directions. In order to hold the UD fiber bundles in place during manufacturing, they are stitched to a thin layer of transverse backing fiber bundles. As the backing only contributes lightly to the mechanical properties of the material, the main focus in this study is on segmenting the UD fibers.

In Figure 1, we see one of the training slices and its corresponding annotated labeling where three different classes are defined: centers, fibers and matrix.

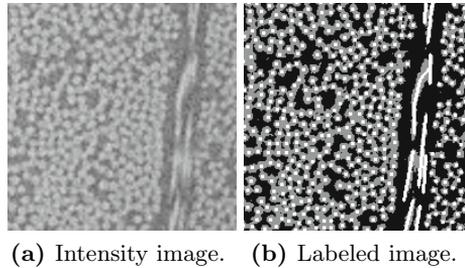


Fig. 1. One slice of the fiber training data. In *white*: centers, *gray*: fibers and *black*: matrix.

3.2 Solid Oxide Fuel Cell Phantom

We simulated a 3D SOFC phantom using a random process employing a simple curvature minimization scheme. Here the target phase fractions of 0.24, 0.38 and 0.38 were used for the pore, ion conducting and electron conducting phases respectively. This phantom creation scheme greatly simplifies many aspects of the evolution of a real micro-structure during sintering and as such the created micro-structure is not suited for comparison of higher order properties to a real SOFC. However, the scheme provides 3D micro-structure data that qualitatively matches the structure observed in real SOFC data sets [9]. The scheme is thus well suited as ground truth for segmentation. After the phantom creation of the artificial 3D micro-structure data the X-ray projections and reconstruction were simulated using slice wise radon and inverse radon transform using AIR tools¹. 3% Gaussian noise was added to the radon transformed data before reconstructing it.

In Figure 2, we can see one of the training slices and its corresponding annotated labeling, where three classes have been defined, corresponding to each of the three phases.

¹ <http://www2.compute.dtu.dk/~pcha/AIRtools/>

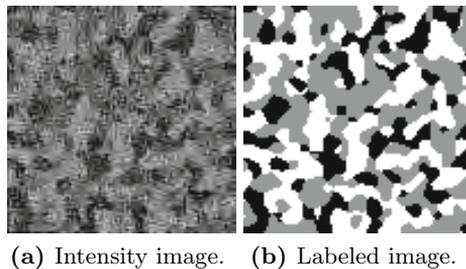


Fig. 2. One slice of the fuel cell training data. In *white*: phase one, *gray*: phase two and *black*: phase three.

4 Results

In this section we provide the results from our algorithm and compare it to well established image analysis methods. Glass fiber detection is compared to 2D scale space blob detection [12] and SOFC segmentation is compared to Markov random field (MRF) segmentation solved using graph cuts with alpha expansion [3, 11]. For both methods we chose parameters that minimized the segmentation error. It should be noted that the parameters for the MRF are not learned, i. e. they are known.

For the dictionary algorithms, several parameters need to be set. It is known that the segmentation is most sensitive to the change in atom size M , i.e. the side length of the cubes, and number of dictionary elements n , so segmentation errors will be calculated for different parameter settings so as to select the optimal.

4.1 Glass Fiber

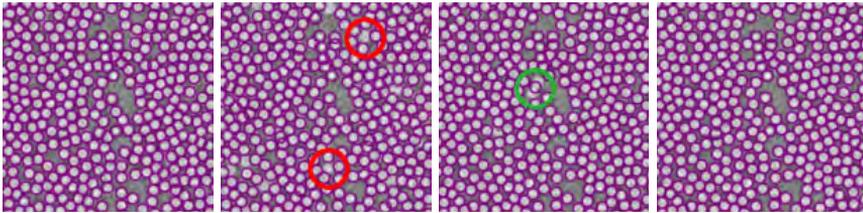
For the glass fiber, the dictionaries have been trained with 9 slices of size 200×200 pixels, the validation set also contains 9 slices of the same size. The performance is calculated over only one test slice of size 500×500 pixels.

The performance measure is computed object-wise over the center class, and represents the true positive rate and false positive rate. We segment the center part of the fibers, and the individual fibers are found as the center of mass using connected component analysis on the center segments. Comparison to the manual marked ground truth is done by counting the number of true matches, which are found as points that have each other as nearest neighbors as well as a distance of less than 4 voxels. The 4 voxels are chosen because the average radius of the glass fibers is around 4 voxels. Results are shown in Table 1. The computational time, defined as the time to train, validate and classify the one test slice, has also been calculated.

The computational time (C_{time}) grows as the number of elements in the dictionary increases (n) or the size of the volume cube (M) becomes bigger. The average C_{time} is 100.3 seconds for 3D, whereas it is 3.2 seconds for the 2D dictionary.

Table 1. Performance measure dictionaries with different parameters

M	n	TPR^{2D}	FPR^{2D}	TPR^{3D}	FPR^{3D}
3	50	0.8563	0.1674	0.9099	0.1202
3	250	0.9082	0.1678	0.9511	0.1520
3	1000	0.9350	0.1721	0.9561	0.1165
5	50	0.9468	0.0399	0.9926	0.0177
5	250	0.9762	0.0248	0.9977	0.0104
5	1000	0.9799	0.0208	0.9963	0.0121
7	50	0.9866	0.0104	0.9896	0.0104
7	250	0.9873	0.0117	0.9973	0.0080
7	1000	0.9910	0.0144	0.9980	0.0087
9	50	0.9484	0.0171	0.9461	0.0107
9	250	0.9930	0.0100	0.9940	0.0104
9	1000	0.9953	0.0097	0.9977	0.0100

**(a)** Ground truth. **(b)** Blob detection. **(c)** 2D dictionary. **(d)** 3D dictionary.**Fig. 3.** Ground truth and circles around detected fiber centers

In Figure 3, we see a zoomed image of fibers with circles around each of the detected centers for each of the three methods (blob detection, 2D dictionary and 3D dictionary) and the ground truth image. Circles are plotted with radius 4 for illustration purposes. As can be seen qualitatively from the images, the performance of the dictionary methods is substantially better than the one obtained through blob detection. The centers in the blob detection are found less accurately compared to those in the ground truth and some centers are not found (e.g. red circles in Figure 3). The displacement of the centers in the 2D and 3D dictionary with respect to the ground truth is very small. Moreover, the 3D dictionary performs slightly better than the 2D dictionary, as it finds all the centers, whereas in the 2D method results there are two centers which have been detected as one only center (blue circle Figure 3). This is because there are some pixels connecting these two centers.

Quantitatively the three methods can be compared using the performance measures TPR and FPR.

1. **Blob detection:** $\text{TPR} = 0.8339$ and $\text{FPR} = 0.0737$.
2. **2D dictionary** ($M = 9$, $n = 1000$): $\text{TPR} = 0.9953$ and $\text{FPR} = 0.0097$.
3. **3D dictionary** ($M = 7$, $n = 1000$): $\text{TPR} = 0.9980$ and $\text{FPR} = 0.0087$.

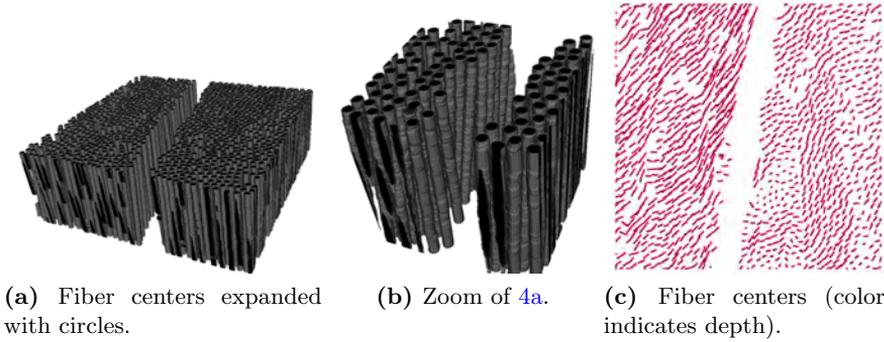


Fig. 4. Segmentation over a stack of 100 slices

In Figure 4, a stack of 100 slices of 300×300 pixels was segmented using the 3D dictionary ($M = 7$, $n = 1000$). In these images we present the results in 3D, where the fiber orientation can be visualized (straight fibers appear as points in 4c). Note how accurate the fiber orientation can be visualized using this approach.

4.2 Solid Oxide Fuel Cell Phantom

For the fuel cell data, the measure used to evaluate the performance is defined as the percentage of pixels which are classified correctly over all three classes. Results are shown in Table 2. The dictionaries are trained with a volume of $50 \times 100 \times 100$ pixels, validated with another volume of $50 \times 100 \times 100$ and the performance is measured over a volume of $100 \times 100 \times 100$ voxels.

As for the glass fiber, the computational time is two orders greater when using the 3D dictionary. In Figure 5, we see one segmented slice from the test set for each of the three methods (MRF, 2D dictionary and 3D dictionary) compared

Table 2. Performance measure dictionaries with different parameters

M	n	Performance ^{2D}	Performance ^{3D}
3	50	0.8864	0.9115
3	250	0.8917	0.9197
3	1000	0.8913	0.9192
5	50	0.9029	0.8886
5	250	0.9103	0.9034
5	1000	0.9119	0.9060
7	50	0.8819	0.8227
7	250	0.8995	0.8453
7	1000	0.9069	0.8512
9	50	0.8407	0.7455
9	250	0.8743	0.7786
9	1000	0.8886	0.7803

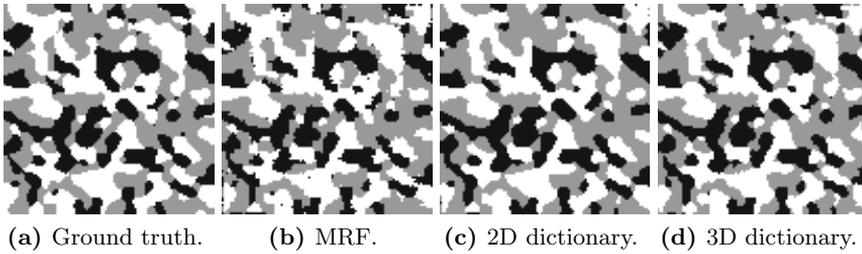


Fig. 5. Ground truth and segmentations of fuel cell phantoms

to the ground truth image. In this case, the benefit from the 3D expansion of the dictionary method is not significant due to the excessive smoothing introduced by the 3D dictionary. However, the dictionary methods do slightly outperform the MRF technique, as we can see from in the following quantitative results:

1. **MRF**: Performance = 0.9078.
2. **2D dictionary** ($M = 5$, $n = 1000$): Performance = 0.9119.
3. **3D dictionary** ($M = 3$, $n = 250$): Performance = 0.9197.

5 Conclusions

In this paper, we have presented a highly flexible and accurate method for 3D segmentation of complex image structures. The method is an extension of a 2D segmentation method and we have investigated if adding a new spatial dimension improved the segmentation performance, as this third dimension provides extra contextual information. We have conducted an investigation to demonstrate the flexibility of the proposed method. Our investigation included experiments for segmenting individual glass fibers in X-ray computed tomography data and for segmenting three phases in phantom data of solid oxide fuel cell data. With the glass fiber data we obtained close to perfect segmentation of the fiber centers. For the fuel cell phantoms we almost see no improvement going from 2D to 3D, which might be due to the excessive smoothing introduced by the 3D algorithm. The fiber detection was compared to scale space blob detection and the solid oxide fuel cell data was compared to Markov random field segmentation using alpha expansion. In both cases both the 2D and 3D methods outperformed these standard methods.

In the current implementation the computational cost is high, but the method has not been optimized for speed. Especially the 3D version is computationally expensive with two orders of magnitude higher computation time and our plan to optimize the method for computational speed in our future work. With the ease of training and very high performance, we believe that the proposed methods can be a very useful tool in quantifying structures in complex volumetric data like tomograms of material samples.

References

1. Aldasoro, C.C.R., Bhalerao, A.: Volumetric texture segmentation by discriminant feature selection and multiresolution classification. *IEEE TMI* **26**(1), 1–14 (2007)
2. Blot, L., Zwiggelaar, R.: Synthesis and analysis of solid texture: application in medical imaging, pp. 9–12 (2002)
3. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(11), 1222–1239 (2001)
4. Cremers, D., Rousson, M., Deriche, R.: A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International Journal of Computer Vision* **72**(2), 195–215 (2007)
5. Dahl, A.L., Larsen, R.: Learning dictionaries of discriminative image patches. In: 22nd BMVC (2011)
6. Government, T.D.: The danish climate policy plan - towards a low carbon society. Technical report, Danish Energy Agency (2013)
7. Hansen, J.Z., Brøndsted, P., Jacobsen, T.K.: The effects of fibre architecture on fatigue life-time of composite materials. Ph.D. thesis, Technical University of Denmark, Risø National Laboratory for Sustainable Energy (2013)
8. Ilea, D.E., Whelan, P.F.: Image segmentation based on the integration of colour-texture descriptors a review. *Pattern Recognition* **44**(10), 2479–2501 (2011)
9. Jørgensen, P.S., Yakal-Kremiski, K., Wilson, J., Bowen, J.R., Barnett, S.: On the accuracy of triple phase boundary lengths calculated from tomographic image data. *Journal of Power Sources* **261**, 198–205 (2014)
10. Jørgensen, P., Ebbenhøj, S., Hauch, A.: Triple phase boundary specific pathway analysis for quantitative characterization of solid oxide cell electrode microstructure. *Journal of Power Sources* **279**, 686–693 (2015)
11. Kolmogorov, V., Zabin, R.: What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(2), 147–159 (2004)
12. Lindeberg, T.: Feature detection with automatic scale selection. *International Journal of Computer Vision* **30**(2), 79–116 (1998)
13. Nijssen, R.P.L.: Fatigue life prediction and strength degradation of wind turbine rotor blade composites. Contractor Report SAND2006-7810P, Sandia National Laboratories, Albuquerque, NM (2006)
14. Randen, T., Monsen, E., Signer, C., Abrahamsen, A., Hansen, J.O., Sæter, T., Schlaf, J., Sønneland, L., et al.: Three-dimensional texture attributes for seismic data analysis. In: 70th Annual International Meeting, Society of Exploration Geophysics Expanded Abstracts, pp. 668–671 (2000)
15. Tai, C., Baba-Kishi, K.: Microtexture studies of pst and pzt ceramics and pzt thin film by electron backscatter diffraction patterns. *Textures and Microstructures* **35**(2), 71–86 (2002)
16. Tu, Z., Zhou, X.S., Comaniciu, D., Bogoni, L.: A learning based approach for 3D segmentation and colon detagging. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *ECCV 2006*. LNCS, vol. 3953, pp. 436–448. Springer, Heidelberg (2006)

17. Ushizima, D., Parkinson, D., Nico, P., Ajo-Franklin, J., MacDowell, A., Kocar, B., Bethel, W., Sethian, J.: Statistical segmentation and porosity quantification of 3d x-ray microtomography. In: SPIE Optical Engineering Applications, pp. 813502–813502. International Society for Optics and Photonics (2011)
18. Waggoner, J., Zhou, Y., Simmons, J., De Graef, M., Wang, S.: 3d materials image segmentation by 2d propagation: a graph-cut approach considering homomorphism. *IEEE TIP* **22**(12), 5282–5293 (2013)
19. Zhang, J., Tan, T.: Brief review of invariant texture analysis methods. *Pattern Recognition* **35**(3), 735–747 (2002)

Individual Fibre Segmentation from 3D X-ray Computed Tomography for Characterising the Fibre Orientation in Unidirectional Composite Materials

Published in *Composites Part A: Applied Science and Manufacturing*, June 2017.



Contents lists available at ScienceDirect

Composites: Part A

journal homepage: www.elsevier.com/locate/compositesa

Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials



Monica J. Emerson^{a,*}, Kristine M. Jespersen^b, Anders B. Dahl^a, Knut Conradsen^a, Lars P. Mikkelsen^b

^a DTU Compute, Richard Petersens Plads, Building 324, 2800 Kgs. Lyngby, Denmark

^b DTU Wind Energy, Frederiksborgvej 399, 4000 Roskilde, Denmark

ARTICLE INFO

Article history:

Received 6 May 2016

Received in revised form 27 October 2016

Accepted 30 December 2016

Available online 11 January 2017

Keywords:

A: Polymer-matrix composites (PMCs)

B: Strength

D: Non-destructive testing

Misalignment

ABSTRACT

The aim of this paper is to characterise the fibre orientation in unidirectional fibre reinforced polymers, namely glass and carbon fibre composites. The compression strength of the composite is related to the orientation of the fibres. Thus the orientation is essential when designing materials for wind turbine blades. The calculation of the fibre orientation distribution is based on segmenting the individual fibres from volumes that have been acquired through X-ray tomography. The segmentation method presented in this study can accurately extract individual fibres from low contrast X-ray scans of composites with high fibre volume fraction. From the individual fibre orientations, it is possible to obtain results which are independent of the scanning quality. The compression strength for both composites is estimated from the average fibre orientations and is found to be of the same order of magnitude as the measured values.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

The load carrying parts of large wind turbine blades are mainly made of unidirectional (UD) fibre reinforced polymers where the fibres are aligned with the main load direction. A low weight together with a high stiffness, a good fatigue resistance and a high compression strength are the key properties that the fibre composite should have when designing large wind turbines [1]. Even if the load carrying composite is stiff, strong and fatigue resistant, a small fibre misalignment would introduce shear stresses into the much weaker and softer polymeric matrix material between the fibres. Especially in compression this is critical [2], as the compression strength of a UD composite relies on the ability of the matrix material to keep the fibres straight [3]. Therefore, experimental characterisation and numerical predictions of the compression strength depend on a good knowledge of the fibre orientation. The influence of different idealised fibre orientation arrangements on the compression strength has been investigated analytically and numerically [3–6]. Nevertheless, in real composites, the fibre orientation will often have a large degree of randomness and some work addresses this as well [7,8].

Several image-based methods for determining the distribution of the fibre orientation in composites have been proposed. These are applied to 2D and 3D image data that can be acquired both

through destructive and non-destructive methods. 2D microscopy provides high-resolution images of individual fibres, and has been used for analysing fibres from different angles using 2D methods [9,7,10].

Destructive methods for analysing the structure in 3D have been investigated [11,12]. However, non-destructive 3D analysis opens the possibility for monitoring changes in the fibre architecture [13], e.g. during loading of a sample [14]. Therefore, it is of interest to consider non-destructive 3D characterisation techniques, which have become possible due to the recent increase in the resolution that can be obtained using different scanning techniques. Confocal laser scanning microscopy was used to determine the 3D fibre orientation distribution of different types of glass fibre composites in [12]. Nevertheless, this technique has some limitations such as small sample size, outlined in [15].

Recent studies also include using 3D X-ray computed tomography (CT), which enables imaging of volumes containing many fibres. If a lab scanner is used and the scan time is short, the resolution for CT will be lower than the one that can be obtained in other microscopic systems. Therefore, it can be hard to segment the individual fibres as they might be barely resolved. Fibre orientations in 3D were obtained for a carbon fibre composite in [16] and for a stitched non-woven fabric in [17], without extracting the individual fibres beforehand. Instead, the fibre orientations are estimated by applying filters to a binarisation of the 3D image which is assumed to be correct. Consequently, the results for the fibre orientation are highly dependant on the quality of the

* Corresponding author.

E-mail address: monj@dtu.dk (M.J. Emerson).

binarisation, which means that it is hard to assess for which kind of data the results will degrade. In [18], cellulose fibres were tracked individually; manual annotation had to be provided in one slice to enable the segmentation of each fibre. Short fibres in composites were also tracked in [15], using a line detection approach. However, the spacing between the fibres in the considered material was fairly large, making it possible to separate fibres even at a coarse resolution. Additionally, segmentation of carbon fibres was carried out in [19,20]. In the first, a lab scanner was used and the quality of the scans was similar to the one in this study, resolution and projection-wise. However, this method requires a tedious manual validation of the fibre tracks needed to deal with the large amount of false centre detections obtained by template matching. Therefore, it is not feasible for large volumes. In [20], the method is applied to phase contrast synchrotron CT data with well-defined boundaries and low noise.

To sum up, previous methods for automated fibre segmentation of composites that contain a high fibre volume fraction have only been able to perform individual fibre segmentation for low contrast scans in very small volumes. As well as enabling fibre orientation characterisation, individual fibre segmentation opens up the opportunity to acquire more knowledge in relation to the fibres and their arrangement, e.g. diameters, distance between fibres and contact points. The method that is proposed in this paper can segment individual fibres in large and noisy data sets extracted from standard compression test samples. The method is presented for filament wound UD glass and carbon fibre composites, which contain a fibre volume fraction in the range of $V_f = [53; 59]\%$. This serves as a good initial case study since they are purely UD with no secondarily oriented fibres and therefore should be considered as an idealisation of the actual non-crimp fabric used in wind turbine blades. The purpose of the idealisation is to exhibit the stiffness and strength potential of the material system and not to represent the actual architecture inside the blades. Furthermore, the applicability to both glass and carbon fibre material systems is highly relevant as they are commonly used composites.

2. Materials

This study is carried out on two frequently used composite materials, namely a UD glass fibre reinforced polymer (GFRP) and a carbon fibre reinforced polymer (CFRP). Both sample manufacturing and scanner settings influence the quality of the recorded data, so they are described in detail here.

Great care was taken to keep all manufacturing details as identical as possible when the fibre composite samples were produced. However, the fibre orientation inside the bundles is given by the fibre supplier and will therefore vary between the two materials. The UD composites are manufactured using a filament winding technique [21] with a small filament winding angle of 0.4° . The fibres are vacuum infused with an epoxy resin and subsequently cured and post-cured under controlled conditions. More details regarding the manufacturing of the laminates can be found in [21,22].

The compression strength is mainly controlled by the fibre orientation and matrix behaviour, which are related as modelled in Eq. (3). Therefore, quite similar compression strengths were expected for the two materials. Nevertheless, [21,22] reported that the compression strength was 10% higher for the CFRP than for the GFRP. This could be due to different fibre orientation distributions inside the fibre bundles. The present paper focuses on the characterisation of this orientation using small cut-outs where the orientation angles will be calculated with respect to a reference coordinate system defined by the mean fibre orientation, see later

Table 1
Materials and scans.

Material	GFRP	CFRP
Average fibre diameter d	17 μm	7 μm
Matrix density	1.2 g/cm^3	1.2 g/cm^3
Fibre density	2.6 g/cm^3	1.8 g/cm^3
Average fibre volume fraction	59%	53%
Optical Magnification	3.98(4 \times)	3.98(4 \times)
Source to sample distance	10 mm	10 mm
Detector to sample distance	20 mm	60 mm
Exposure time	6 s	15 s
Accelerating voltage	60 keV	60 keV
Power	4.98 W	5.04 W
Number of projections	3201	3201
Voxel size v_x	1.124 μm	0.482 μm
Binned voxel size v_{ds}	2.248 μm	0.964 μm
Field of view on detector	2248 μm	964 μm
Analysed width W	1348.8 μm	578.4 μm
Analysed depth D	1348.8 μm	867.6 μm

in Section 4.1. Note that the method would also work on full test samples.

The cut-outs come from a gauge section of size $3 \times 10 \times 12 \text{ mm}^3$ and were scanned in a laboratory micro-focus X-ray CT system (Zeiss Xradia Versa 520) with the settings listed in Table 1. Please note that the scan time can be estimated from the number of projections and the exposure time; and that other parameters such as tube current are set automatically by the scanner to match the voltage and power. As it is a region of width W and depth D that will be analysed, and not the whole field of view, these values are given in Table 1, instead of the cut-out sizes. The 3D image reconstruction was performed using the software XMReconstructor Cone Beam 10 by Zeiss, the binning was performed using ImageJ and the numerical experiments were carried out using Matlab.

3. Automated fibre segmentation method

A direct way of characterising the degree of fibre misalignment is by analysing the orientation of individual fibres. This requires the computation of each individual fibre 'trajectory', a 3D line characterising the fibre orientation that is computed as a set of connected centre points. The trajectories are obtained in two steps. First, a method for detecting individual fibre centre points in cross-sectional images from the reconstructed volumes is described in Section 3.1. Then, from these detections the fibre trajectories are computed through tracking, described in Section 3.2.

The process of obtaining the fibre trajectories is illustrated in Fig. 1. The volume is sliced and the segmentation is done over the 2D slices. Since the fibres are unidirectional, each slice gives an image showing the fibre cross-section, which is perpendicular to the axial direction of the scanner.

Despite the individual fibres appearing clear in the image, it is challenging to detect them individually since they are lumped together. Consequently, it was not possible to separate individual fibres using simple segmentation techniques such as thresholding with Otsu's method [23], illustrated in Fig. 2. More complex segmentation techniques such as blob detection or the template matching scheme used in [19] are outperformed by the dictionary segmentation method presented here. These methods can find most of the centres, but at the cost of having a much higher degree of false centre detections. If there should be no false centre detections, lots of true centres would be missed, leading to a significantly lower accuracy compared to the dictionary method, as demonstrated in [24] for blob detection.

Instead, a supervised segmentation technique is proposed, that is, a segmentation method that requires a training step. This

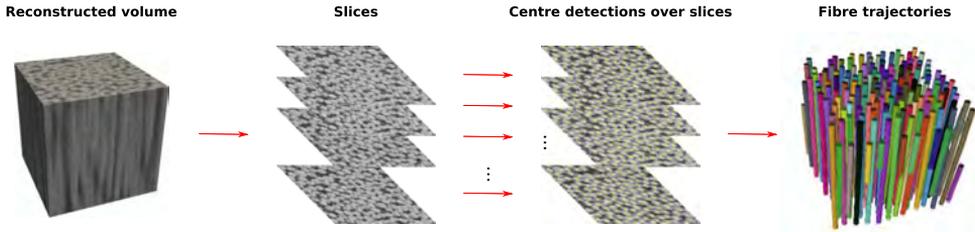


Fig. 1. Description of the whole process for individual fibre segmentation. First, the volume is sliced. Then, each slice is processed so as to obtain the fibre centre points in every slice. Finally, the centre points need to be tracked through the volume, i.e. the centres corresponding to the same fibre need to be connected. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

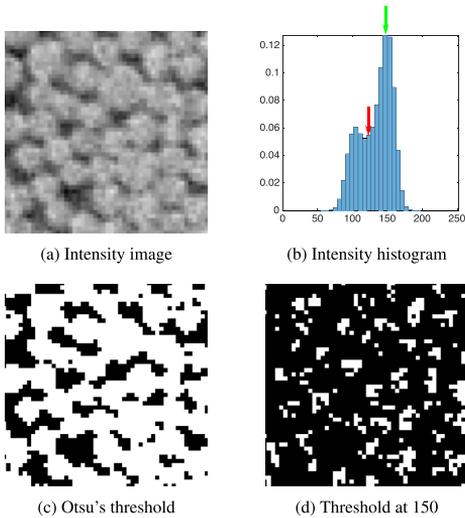


Fig. 2. Segmentation attempt of the intensity image in (a) using thresholding. This intensity image corresponds to a part of a slice from a composite volume. (b) shows the intensity histogram, (c) the thresholded image in the valley of the histogram at intensity level 127 (obtained by Otsu's method and marked in red in (b)). This image illustrates the problem of connected fibres. (d) shows a threshold at a higher intensity 150 (marked in green in (b)) illustrating the problem of fragmented parts. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

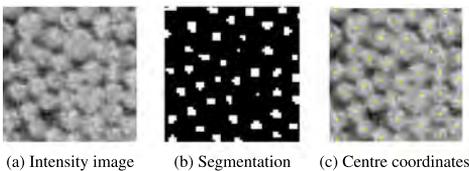


Fig. 3. Example of fibre centre point detection in an intensity image, which is a part of a slice from a composite volume. (a) shows the intensity image, (b) the segmentation of class A, and (c) the computed centre points displayed on top of the intensity image.

technique is based on 'learning' a so-called dictionary of image patches and corresponding label patches [24]. The dictionary is a compact representation of the information contained in a manually annotated intensity image. Looking up this dictionary will enable classification of an image, that is, to decide whether each pixel belongs to class A or class B. A pixel will be labelled as class A when it is in the area around a fibre centre, and class B if it is not.

3.1. Centre point detection

The two steps required for the centre point detection are illustrated in Fig. 3. First, a segmentation of the image is obtained. This corresponds to a labelling where each pixel is marked as white if it belongs to class A and as black if it belongs to B. The second step is to find the geometric centre for each set of connected white pixels. This will grant the centre points.

The process of segmenting an image based on the 'learnt' dictionaries is described in Section 3.1.2, but first the training step to obtain the dictionary will be presented in Section 3.1.1. These steps were also described in [24].

3.1.1. Training the algorithm

The training process is illustrated in Fig. 4, where the dictionaries are 'learnt' from a training set comprised by one single intensity image and a label image for each class, where the actual class is marked in white. The size of the intensity image used in this study covers a region of 200 × 200 pixels, which is equivalent to approximately 30 × d for both data sets. A training image should be chosen to represent the typical patterns in the image. However, for fibre segmentation the method was found to be very robust to the choice of training image. The label images were created manually by annotating the areas around the centres. By only labelling small areas around the centres instead of complete fibres, it is possible to obtain sets of centre pixels which are not connected one to another. Therefore, individual fibres can be segmented.

There are several parameters that need to be set when training the algorithm.

1. The number of classes k , which defines the amount of distinctive segments required. It is set to $k = 2$ in this specific case (class A and class B).
2. The size of an image patch $m = M^2$, where M is the width of the patch in pixels. Experience with the dictionary method has

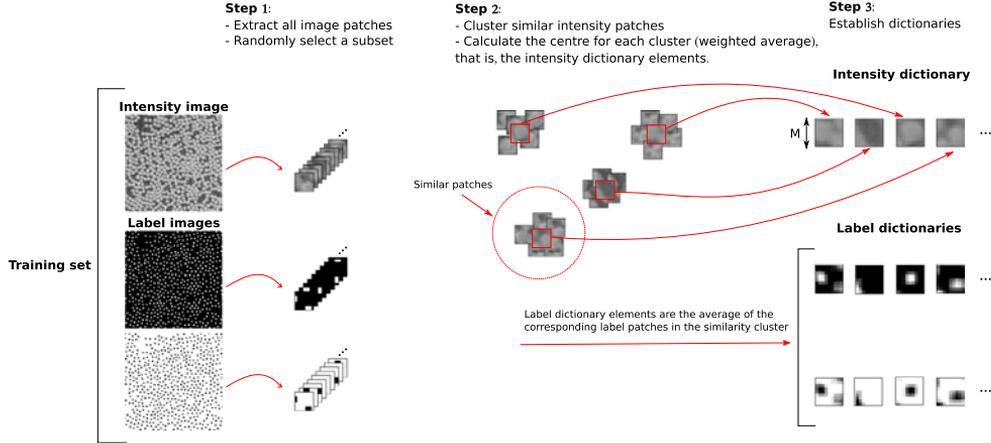


Fig. 4. Description of the training process. Overlapping image patches are extracted from the set of training images (intensity image and two corresponding label images for class A and B). A random set of ψ from the extracted image patches is selected and clustered using weighted k -means [25] to obtain the n intensity dictionary elements, where ψ is typically much larger than n . Each label dictionary element is the average of the label patches that correspond to the intensity patches used for estimating its corresponding intensity dictionary element. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

revealed that this parameter is important (see [24]). $M = 9$ pixels provides good results, that is $M \approx \frac{9d}{7v_{fs}}$, slightly bigger than the fibre diameter.

3. The number of dictionary elements n , where a dictionary element is comprised by an image patch and its corresponding label patches. This parameter has also proven to be important. $n = 250$ is selected because results do not change significantly when more elements are added (see [24]).
4. The number of iterations n_{iter} , which has proven to not be very important, $n_{iter} = 10$ is selected.

The training process in detail:

- First Obtain a training set. Select an intensity image and create the label images by manually annotating each class over the intensity image. In this case created by marking the area around the fibre centres.
- Second Extract overlapping patches from the intensity and label training images. This means that image patches will be extracted starting from the top left corner and moving one pixel down until the bottom of the image is reached. Then, the next patch extracted is the one that is one pixel to the right from the top left corner. Note that each intensity patch has $k = 2$ label patches associated.
- Third Select a random set of ψ corresponding intensity and label patches $v_j = (v_{Ij}, v_{Lj})$ where $j = 1 \dots \psi$. It should be noted that ψ should be large enough to capture the variety of patches in the image and always $\psi > n$. In this study $\psi = 3000$.
- Forth Initialise the n dictionary elements to a random selection of n patches out of the ψ .
- Fifth Follow the pseudo-code in Algorithm 1 to obtain the final dictionary elements.

Algorithm 1. Building the dictionary

Require: Initialised dictionaries d_i^0 and d_L^0 with randomly selected patches.

- 1: Set $n_{iter} = 10$
- 2: **for** $t = 1$ to n_{iter} **do**
- 3: Cluster the v_j patches in n clusters where the cluster centres are the intensity dictionary elements from the previous iteration (d_i^{t-1}) and the similarity measure is Euclidean.
- 4: **for all** clusters **do**
- 5: Define S_i : set of indices for the patches in cluster i .
- 6: Define n_i : number of elements in cluster i
- 7: Update the dictionary elements corresponding to cluster i
- 8:
$$d_{Li}^t = \frac{1}{\sum_{j \in S_i} w_{ji}} \sum_{j \in S_i} w_{ji} v_{Lj}$$

where $w_{ji} = 1 - \frac{1}{\sqrt{2M}} \|d_{Li}^{t-1} - v_{Lj}\|_2$
- 9:
$$d_{Ii}^t = \frac{1}{n_i} \sum_{j \in S_i} v_{Ij}$$
- 10: **end for**
- 11: **end for**

In Algorithm 1 both intensity and label patches take part in the process of building the dictionary. The intensity dictionary elements are a combination of the intensity patches in the corresponding cluster weighted to maximise the uniqueness of the label patches. That is, the smaller the difference between the label of the dictionary element d_{Li}^{t-1} and the label of the image patch v_{Lj} , the more influence the specific image patch will have over the dictionary element. Hereby a discriminative dictionary with good clustering properties is obtained.

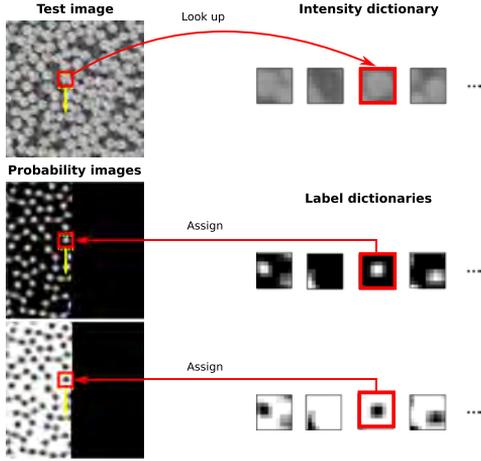


Fig. 5. Description of the classification procedure. First, overlapping image patches are extracted from the test image and each of these patches is associated to the closest element in the intensity dictionary, using Euclidean distance. Then, to create the probability images, the element from the label dictionary which corresponds to the closest element in the intensity dictionary is added to the probability image for that class in the position where the patch had been extracted from. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.1.2. Segmentation of an image

The classification procedure is illustrated in Fig. 5 and is based on assigning a probability of belonging to the classes A and B to each pixel in a test image, which is the image where the centre points should be detected. Each pixel in the test image will be classified as the class that is most likely, given by the probability images.

Given an image $U_l : \Omega \rightarrow \mathbb{R}$ that ought to be segmented with $U_l \in \mathbb{R}^{x \times y}$, label probability images must be computed $U_l : \Omega \rightarrow \mathbb{R}$ with $U_l \in \mathbb{R}^{x \times y \times k}$ where k is the number of classes. Initially U_l is set to having all elements zero. Overlapping image patches are extracted from U_l . These intensity patches \mathbf{u}_{hi} , which are of the same spatial size as the dictionary elements, are vectorised and then matched to the nearest intensity dictionary element using Euclidean distance. For a given intensity vector \mathbf{u}_{hi} , the nearest dictionary element is obtained as

$$i(h) = \arg \min_i \|\mathbf{d}_{hi} - \mathbf{u}_{hi}\|_2^2. \quad (1)$$

The corresponding label dictionary elements $\mathbf{d}_{i(i(h))}$ are added to the probability images U_l at the coordinates of the h 'th patch extracted from U_l for each of the k class labels. When all the dictionary elements have been added, the probability images are normalised to make the probabilities sum to one over all k labels.

Some smoothing occurs at boundaries and it especially affects small features. In order to account for that, label-wise weights are estimated on an annotated validation set comprised by an intensity image V and the label image for all classes Q_L . The difference between the probability images P_L obtained for V and its label images Q_L is minimised. The weights are found as

$$\mathbf{W} = (\mathbf{P}_L^T \mathbf{P}_L)^{-1} \mathbf{P}_L^T \mathbf{Q}_L, \quad (2)$$

where the matrices, in bold, are of size $n_x n_y \times k$. In this study $n_x = n_y = 100$ pixels.

3.2. Fibre tracking

The fibre tracking is the process of matching the centre points from one slice to the next so as to define connected fibres through the composite volume. Algorithm 2 describes this procedure, where a centre point in a specific slice will be matched to a centre point in an upcoming slice if the distance c between the centres is less than a decided threshold $\tau = 2$ pixels ($\approx \frac{2d}{7V_{fib}}$).

Algorithm 2. Fibre tracking

Require: Centre points in every slice.

- 1: Assign a unique fibre ID to each centre point in slice $p = 1$.
- 2: Add centres points in slice 1 to variable s , saved points.
- 3: **for** $p = 2$ **to** n_{slices} **do**
- 4: Create variable s_{new} , where new points will be saved.
- 5: Add centre points in slice p to s_{new} .
- 6: Nearest-neighbour search between points in s and s_{new} . All points in s_{new} will be assigned with a distance c to a point in s . However, some points in s might not have a corresponding point in s_{new} , but could have a match in the next slices, if just due to missed detections.
- 7: **for all** points $\in s$ **do**
- 8: **if** not matched **then**
- 9: Add centre point in s to s_{new}
- 10: **else**
- 11: **if** $c < \tau$ **then**
- 12: Old fibre continuing, assign to centre point in s_{new} the fibre ID from the associated centre point in slice s .
- 13: **else**
- 14: New fibre starting, assign new fibre ID to centre point in s_{new} .
- 15: **end if**
- 16: **end if**
- 17: **end for**
- 18: Clear s
- 19: $s = s_{new}$
- 20: **end for**

Once the fibre trajectories are obtained, they should be post-processed as shown in Fig. 6 and following the scheme described below. Please note that the voxel size is the same in all three directions of a volume, therefore the distance between pixels and slices is the same.

1. Split fibres where the detection has been missed in more than s_1 slices in a row. This is done so as to avoid certain tracking artifacts occurring at the borders. $s_1 = 5 \left(\approx \frac{5d}{7V_{fib}} \right)$ was chosen by visual inspection.
2. Eliminate fibres shorter than s_2 slices because, in general, these are due to artifacts from our tracking method. $s_2 = 10 \left(\approx \frac{10d}{7V_{fib}} \right)$ was selected by visual inspection, which corresponds to a length of 22.48 μm and 9.64 μm for the GFRP and CFRP respectively. The elimination of these fibres will not affect the overall fibre orientation measurements, as the percentage of eliminated centre points is of the order of 0.1% and 0.01% for the glass and carbon fibre respectively.
3. Where the detection has been missed in s_1 or less slices in a row, the missing centres are estimated by interpolation. $s_1 = 5 \left(\approx \frac{5d}{7V_{fib}} \right)$ was chosen because it was decided that a fibre should either be split or interpolated, that is, no fibres with missed detections should exist after post-processing.

4. Fibre trajectories are smoothed with a kernel of width equal to s_3 slices. This is done so as to eliminate the high frequency and low amplitude noise generated by the algorithm used for the centre point detection. This noise appears due to the lack of sub-pixel precision when computing the centre points. A value that eliminated the noise while preserving the natural movements of a fibre was found to be $s_3 = 30 \left(\approx \frac{30d}{7\lambda_{\text{obs}}} \right)$. This value corresponds to a ratio between the smoothing length and the average fibre diameter that is approximately equal to 4 for both data sets. The smoothing length is $67.77 \mu\text{m}$ and $28.92 \mu\text{m}$ for the glass and carbon fibre respectively.

4. Individual fibre orientation calculation method

This section describes how the local orientation is calculated with respect to a reference coordinate system.

4.1. Calculation of the reference coordinate system

The individual fibre trajectories have been obtained in the scanning coordinate system (x_s, y_s, z_s) . However, this coordinate system will not be used as a reference because it may not represent the objective orientation of the test samples. This is due to the scanned samples being a cut-out from the test samples. In order to be able to define the local fibre orientation in an objective way, the reference coordinate system (x_p, y_p, z_p) will be estimated statistically based on the data. The z_p axis will correspond to the mean fibre orientation (MF), the x_p axis to the direction that captures the maximum variation of the global fibre orientation in the plane that is orthogonal to MF and y_p will be defined given x_p and z_p .

The process of determining the reference coordinate system consists of two steps. First, MF is obtained as the average global

orientation of the fibres and is defined by θ_{MF} and ϕ_{MF} . In order to have a reliable estimate, only the fibres that run through the whole volume are taken into account. After rotating the data sets θ_{MF} so that the z axis corresponds to MF, the direction that captures the maximum variation of the global fibre orientation is computed through principal component analysis (PCA) [26]. Then, the data sets are rotated ϕ_{PC} .

The global orientation of a fibre can be fully characterised with the vector \mathbf{r} , which is determined by the projection of the fibre's end point onto the (x, y) plane, see Fig. 7a. It should be noted that the vector \mathbf{r} , and its projections r_x and r_y onto \mathbf{x} and \mathbf{y} respectively, will be normalised so as to indicate pixels of deviation per slice.

The magnitude of \mathbf{r} is proportional to the elevation angle θ . In other words, $|\mathbf{r}|$ encodes the amount of misalignment with respect to the z_p axis. The angle of \mathbf{r} determines the azimuth ϕ , which is the direction in which the fibre is misaligned.

A scatter plot of r_x against r_y for each fibre gives an idea of where MF lies. However, it should be interpreted carefully as some information is lost, namely the density, that is, the areas where there is a greater concentration of points. A histogram with contour plots can reveal better where the dense areas are located. Both scatter and contour plots are displayed in Fig. 8 for the GFRP.

In Fig. 9, the scatter plots of the global fibre orientation are shown for both data sets before and after the two rotations. Note that the rotations shift MF (marked in red) to the origin of the coordinate system, as well as aligning the direction of maximum variation with the x axis.

4.2. Local fibre orientation description

The local fibre orientation captures the variations within a fibre and is defined as the $\Delta\theta$ order difference. This is the simplest esti-

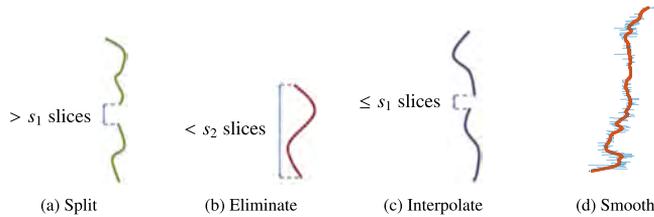


Fig. 6. Post-processing steps: split, eliminate, interpolate and smooth. Note that the waviness of the fibres seems greater than it actually is due to the axes being scaled differently. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

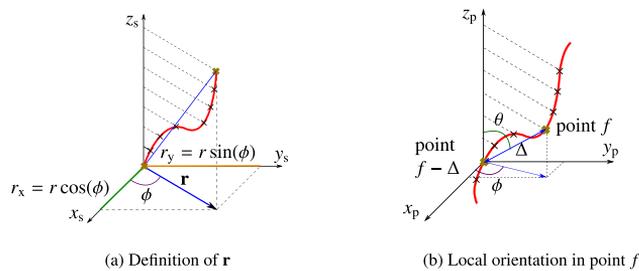


Fig. 7. In (a) the definition of \mathbf{r} , and its projections r_x and r_y , with respect to the scanning coordinate system, as the reference coordinate system has not yet been found. In (b) the local fibre orientation in point f with respect to the reference coordinate system, where $\theta \in [0, 90]$ and $\phi \in [0, 360]$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

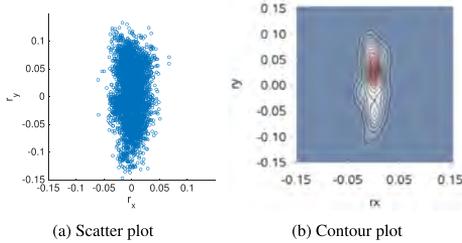


Fig. 8. The variation in the global fibre orientation of the GFRP with respect to the scanning coordinate system. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

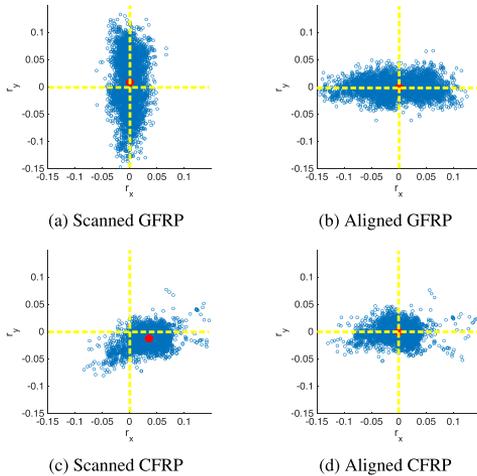


Fig. 9. Measure of the global fibre orientation for both data sets with respect to the scanning coordinate system on the left and with respect to the reference coordinate system on the right. In yellow the x and y axes and in red MF. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

mate of the gradient, and is illustrated in Fig. 7. An elevation θ and azimuth ϕ angle will be estimated for every point f in each fibre.

5. Results and discussion

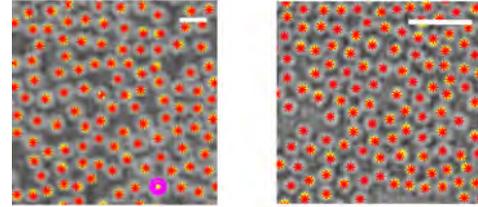
In this section, the results from the automated individual fibre segmentation are evaluated. Then, the mean fibre direction is provided and the reference coordinate system is defined for each data set. Afterwards, the fibre orientation distribution is presented and the compression strength predicted. Finally, a discussion on the performance of the algorithm is included.

5.1. Automated fibre segmentation

Here the accuracy of the centre point detection is estimated and the results from the tracking are visualised.

5.1.1. Centre coordinate detection

To evaluate the centre detection method, the automatically obtained centre points are compared to reference points that were



(a) GFRP, TPR = 0.991.

(b) CFRP, TPR = 1.

Fig. 10. Centre points over a region from the test image with a scale bar indicating $40\ \mu\text{m}$. In yellow the reference points and in red the detected ones, missed detections are circled in pink. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

created by manually annotating a test image. The test image corresponds to a part of a slice that has not been used for training nor validating the algorithm, the width of the test image is $674.4\ \mu\text{m}$ for the GFRP and $289.2\ \mu\text{m}$ for the CFRP.

The accuracy of the method is calculated as the true positive rate (TPR), which is the ratio between the amount of correctly detected centres and the total amount of reference centres. Fig. 10 shows a small region of the test image for both data sets. It should be noted that the centre coordinates will be less precise if the resolution of the scan is lower. In this study, the fibre cross-sections contain around 7 pixels per diameter, which leads to centre coordinate detections that are precise enough to characterise the fibre orientation. However, better resolution will lead to a more precise estimate, which might be necessary if one wants to calculate fibre diameters or create an accurate finite element model afterwards.

5.1.2. Fibre tracking

The fibre trajectories are shown with respect to the reference coordinate system in Fig. 11.

5.2. Mean fibre direction

MF is given by $\theta_{MF} = 0.53^\circ$ and $\phi_{MF} = 87.91^\circ$ for the GFRF, and by $\theta_{MF} = 2.14^\circ$ and $\phi_{MF} = -18.85^\circ$ for the CFRP.

5.3. Reference coordinate system

The rotations that are required to align the scanned volumes with the reference coordinate system are first $\theta_{MF} = 0.53^\circ$ and then $\phi_{PC} = 90.35^\circ$ for the GFRP; and $\theta_{MF} = 2.14^\circ$ followed by $\phi_{PC} = 16.60^\circ$ for the CFRP.

5.4. Individual fibre orientation

In this section, the orientation of the individual fibres is presented. Note that the projections θ_x and θ_y of θ on to the planes (x, z) and (y, z) can be calculated from θ and ϕ . The notation used for the estimated means will be $\langle \rangle$.

5.4.1. Global fibre orientation

The global fibre orientation is defined as the local orientation for $\Delta = D$, depth values reported in Table 1.

Fig. 12 contains a top view of the fibre trajectories shown in Fig. 11 which have now been coloured depending on the global azimuth angle. The azimuth encodes the direction in which the fibre is orientated in the (x, y) plane. It can be seen that in different regions the fibres tend to be orientated in a certain direction. The

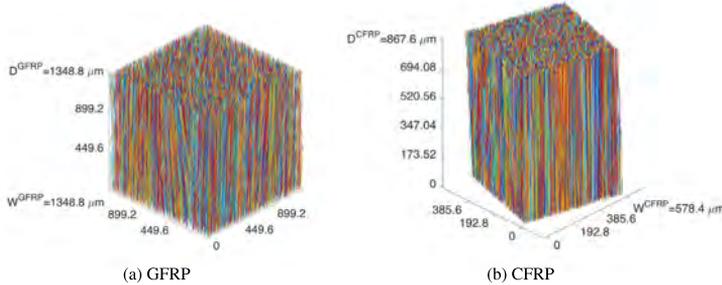


Fig. 11. 3D side view of the fibre trajectories as connected points, where the colouring of the fibres is random. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

elevation angle, which grants the amount of misalignment with respect to MF, is encoded in the length of the fibres, when these are visualised from the top. The longer a fibre, the more misaligned it is. When it appears as a point instead of a line, it means that it is completely aligned with MF.

5.4.2. Local fibre orientation

An adequate value for the parameter Δ (see Fig. 7b) must be determined. This value cannot be too small otherwise, the fibre orientation calculation will be affected by noise, if any is left after smoothing the fibre trajectories. In addition, the use of a fixed ratio R between the length Δ and the average fibre diameter d (see Table 1) will facilitate the comparison between data sets. In Fig. 13, the distributions of the local fibre orientation for the GFRP at ratios of $R = 4, 10$ and 79 are shown. Note that $\langle \theta_x \rangle$ and $\langle \theta_y \rangle$ are zero for $R = 79$ because this ratio corresponds to $\Delta = D$, the length used when aligning the data set with MF. Moreover, note that the influence of the length Δ on the calculated fibre orientations is small. Fig. 14 shows the distribution of the local fibre orientation in the planes (x, z) and (y, z) for the CFRP and $R = 4$. It can be seen that the glass fibres have a larger variation in fibre orientation, in addition to presenting a preferred misalignment plane.

5.5. Compression strength

Based on the measured degree of fibre misalignment, the compression strength can be predicted and compared with measured

strength values reported in [22]. Assuming an infinitely wide band of fibres with an initial misalignment $\hat{\theta}$ (measured in radians), Budiansky [3] shows that the compression strength can be estimated as

$$\sigma_c = \frac{G}{1 + \hat{\theta}/\gamma_y} \tag{3}$$

where it can be assumed that both the GFRP and CFRP have a shear modulus of $G = 4.0$ GPa and a shear yield strain of $\gamma_y = 0.006$. From the estimated mean of the absolute value of the elevation angles shown in Figs. 13 and 14, the compression strengths can be pre-

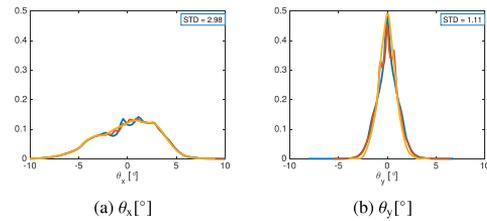


Fig. 13. Distribution of the local fibre orientation for the GFRP with $R = 4$ (in blue), 10 (in red) and 79 (in yellow), which correspond to lengths of $\Delta^{GFRP} = 68 \mu\text{m}, 170 \mu\text{m}$ and $1343 \mu\text{m}$. Note that the histograms are normalised, so as to approximate the probability density function. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

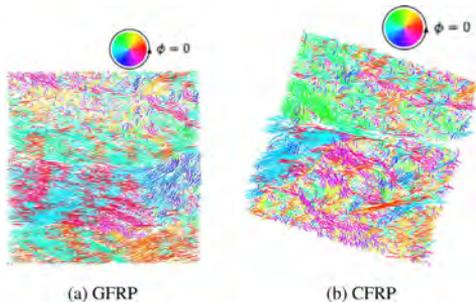


Fig. 12. A top view of the fibre trajectories as connected points. Fibres will appear as a point if they are completely aligned with MF, and the colour of each fibre encodes the fibre's azimuth angle ϕ , as shown in the colour wheel. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

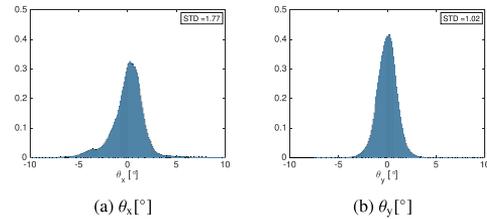


Fig. 14. Distribution of the local fibre orientation for the CFRP with $R = 4$, which means that $\Delta^{CFRP} = 28 \mu\text{m}$. Note that the histograms are normalised, so as to approximate the probability density function.

Table 2
Predicted and measured compression strengths.

Material	GFRP	CFRP
$\langle \theta \rangle$	2.75°	1.61°
$\langle \theta_x \rangle$	2.45°	1.23°
$\langle \theta_y \rangle$	0.86°	0.79°
$\bar{\sigma}$	0.44 GPa	0.70 GPa
$\bar{\sigma}_x$	0.49 GPa	0.87 GPa
$\bar{\sigma}_y$	1.14 GPa	1.21 GPa
σ_m	(0.82 ± 0.07) GPa	(0.90 ± 0.06) GPa

dicted $\bar{\sigma}$ and compared to the measured values σ_m from [21], as shown in Table 2.

The measured compression strength is found to be of the same order of magnitude as the predictions based on Eq. (3). As was stated in Section 2, the analysis here is performed just on a small region of the whole measured gauge section of size $3 \times 10 \times 12 \text{ mm}^3$. The analysed regions are of size $1.35 \times 1.35 \times 1.35 \text{ mm}^3$ and $0.58 \times 0.58 \times 0.87 \text{ mm}^3$ for the GFRP and the CFRP respectively, thus they only represent a 0.68% and a 0.08% of the whole tested gauge section. Despite this, the manufacturing process during filament winding is observed to result in rather homogeneous specimens. In the future, a more thorough investigation could be performed where either the whole tested gauge section is analysed or where it is shown that the analysed region is representative with respect to the fibre and bundle structure by inspecting a number of randomly chosen sub-regions. In addition, the lowest compression strength is found for the composite material with the largest degree of misalignment. Nevertheless, the complete spatial distribution of the fibre orientation should be used for a more accurate prediction of the compression strength. Readers are invited to try their own methods for estimating fibre orientation, as well as methods for investigating the effect of the non-homogeneous fibre orientation on the predicted compression strength. The files with the X-ray scans and the fibre centre lines are available online.¹ The fibre centre lines can be used as an example of fibre architecture in a filament wound composite material and, together with the average fibre diameters reported in Table 1, they can be used as a starting point for building up a micro-mechanical based finite element model.

5.6. Performance of the algorithm

The algorithm requires a manually annotated image for training the dictionary. Here an image of size 200×200 pixels was used and it took around half an hour to annotate. Afterwards, the processing of a full volume of size $1000 \times 1000 \times 1000$ pixels to extract the fibre trajectories takes approximately 3 h on a Macbook Pro 2.2 GHz Intel Core i7 and, since the method is highly robust, there is no need for manual corrections. Furthermore, it should be noted that the implementation in Matlab is not optimised for speed and the possibility for parallelisation was not used. That is, the most time consuming part is actually manually annotating the training image. Nevertheless, already trained dictionaries can be re-used for analysing samples with similar fibre volume fraction and relationship between fibre diameter and voxel size.

6. Conclusion

In this paper, a method that is able to extract individual fibres to calculate their orientation in a composite material has been

presented. The strength of this method is its ability to extract individual fibres from low contrast and noisy volumetric scans where there is a high fibre volume fraction, without requiring manual corrections due to its high robustness. Consequently, it provides accurate results independently of the quality of the grey scale images. Previous methods that calculated the fibre orientation directly from grey scale images suffered from noise contamination, granting different results for the same image with varying noise levels. The ability to deal with low contrast and noisy scans is very important, as it is not feasible to obtain high quality images in many occasions, due to, for example, long scan times.

Acknowledgements

Financial support from CINEMA: the allianCe for ImagiNg of Energy MAterials [DSF-Grant No. 1305-00032B] and DCCSM: the Danish Centre for Composite Structures and Materials for Wind Turbines [DSF-Grant No. 09-067212], both under the Danish Council for Strategic Research, are gratefully acknowledged. The material used has been manufactured in the project Blade King financially supported by the Danish National Advanced Technology Foundation.

References

- [1] Brøndsted P, Lilholt H, Lystrup HA. Composite materials for wind power turbine blades. *Annu Rev Mater Res* 2006;35:505–38.
- [2] Vogler TJ, Hsu SY, Kyriakides S. Composite failure under combined compression and shear. *International J Solids Struct* 2000;37(12):1765–91.
- [3] Budiansky B, Fleck N. Compressive failure of fibre composites. *J Mech Phys Solids* 1993;41(1):183–211.
- [4] Fleck NA, Shu JY. Microbuckle initiation in fibre composites: a finite element study. *J Mech Phys Solids* 1995;43(12):1887–918.
- [5] Kyriakides S, Arseculeratne R. On the compressive failure of fiber reinforced composites. *Int J Solids Struct* 1995;32(617):689–738.
- [6] Jensen HM, Christoffersen J. Kink band formation in fiber reinforced materials. *J Mech Phys Solids* 1997;45(7):1121–36.
- [7] Creighton CJ, Sutcliffe MPF, Clyne TW. Multiple field image analysis procedure for characterization of fibre alignment in composites. *Compos Part A: Appl Sci Manuf* 2001;32:221–9.
- [8] Liu D, Fleck N, Sutcliffe M. Compressive strength of fibre composites with random fibre waviness. *J Mech Phys Solids* 2004;52(7):1481–505.
- [9] Yurgartis S. Measurement of small angle fiber misalignments in continuous fiber composites. *Compos Sci Technol* 1987;30(4):279–93.
- [10] Kratmann K, Sutcliffe M, Lillheden L, Purz R, Thomsen O. A novel image analysis procedure for measuring fibre misalignment in unidirectional fibre composites. *Compos Sci Technol* 2009;69(2):228–38.
- [11] Paluch B. Analysis of geometric imperfections affecting the fibers in unidirectional composites. *J Compos Mater* 1996;30(4):454–85.
- [12] Clarke A, Archenhold G, Davidson N. A novel technique for determining the 3D spatial distribution of glass fibres in polymer composites. *Compos Sci Technol* 1995;55(1):75–91.
- [13] Buffiere J-Y, Maire E, Adrien J, Masse J-P, Boller E. In situ experiments with X-ray tomography: an attractive tool for experimental mechanics. *Exp Mech* 2010;50(3):289–305.
- [14] Garcea S, Sinclair I, Spearing S. Fibre failure assessment in carbon fibre reinforced polymers under fatigue loading by synchrotron X-ray computed tomography. *Compos Sci Technol* 2016;133:157–64.
- [15] Eberhardt CN, Clarke AR. Automated reconstruction of curvilinear fibres from 3D datasets acquired by X-ray microtomography. *J Microsc* 2002;206 (April):41–53.
- [16] Altendorf H, Jeulin D. 3d directional mathematical morphology for analysis of fiber orientations. *Image Anal Stereology* 2011;28:143–53.
- [17] Tausif M, Duffy B, Grishanov S, Carr H, Russell SJ. Three-dimensional fiber segment orientation distribution using X-ray microtomography. *Microscopy Microanal* 2014;1–10.
- [18] Axelsson M. Micromechanical investigation of fatigue damage in unidirectional fibre composites. In: International conference on image processing, p. 309–12.
- [19] Czabaj MW, Riccio ML, Whitacre WW. Numerical reconstruction of graphite/epoxy composite microstructure based on sub-micron resolution X-ray computed tomography. *Compos Sci Technol* 2014;105:174–82.
- [20] Requena G, Fiedler G, Seiser B, Degischer P, Di Michiel M, Buslaps T. 3d-quantification of the distribution of continuous fibres in unidirectionally reinforced composites. *Compos Part A: Appl Sci Manuf* 2009;40 (2):152–63.

¹ Supplementary material at <http://www2.compute.dtu.dk/abda/fibre/>. For use of this data, please make a reference to this paper.

- [21] Markussen CM, Andersen TL, Madsen B, Nielsen MW, Lilholt H, Mechanics M. Improved compression strength of carbon/glass/epoxy hybrid composites. In: 19th international conference on composite materials.
- [22] Prabhakaran RTD, Andersen TL, Markussen CM, Madsen B, Lilholt H. Tensile and compression properties of hybrid composites – a comparative study. In: 19th international conference on composite materials. p. 1029–35.
- [23] Otsu N. A threshold selection method from gray-level histograms. *Automatica* 1975;11(285–296):23–7.
- [24] Emerson MJ, Jespersen KM, Jørgensen PS, Larsen R, Dahl AB. Dictionary based segmentation in volumes. In: 19th scandinavian conference on image analysis. Springer; 2015. p. 504–15.
- [25] J. MacQueen. Some methods for classification and analysis of multivariate observations, in: Proceedings of the fifth Berkeley symposium on mathematical statistics and probability, vol. 1. Oakland, CA, USA; 1967. p. 281–97.
- [26] Jolliffe I. *Principal component analysis*. Wiley Online Library; 2002.

Content-based Propagation of User Markings for Interactive Segmentation of Patterned Images

To be submitted, *IEEE Transactions on Image Processing*.

Content-based Propagation of User Markings for Interactive Segmentation of Patterned Images

Vedrana Andersen Dahl, Camilla Himmelstrup Trinderup, Monica Jane Emerson and Anders Bjørholm Dahl

Abstract—Efficient and easy segmentation of images and volumes is of great practical importance. Segmentation problems which motivate our approach originate from imaging commonly used in materials science and medicine. We formulate image segmentation as a probabilistic pixel classification, and we apply segmentation as a step towards characterising image content. Our method allows the user to define structures of interest by marking a subset of pixels in an interactive manner. Thanks to the real-time feedback, the user can place new labelling strategically, depending on the current outcome. The final pixel classification may be obtained from a very modest user input. An important ingredient of our method is a graph that encodes image content. This graph is built in an unsupervised manner during initialisation, and is based on clustering of image features. Since we combine a limited amount of labelled data with the clustering information obtained from the unlabelled parts of the image, our method fits in the general framework of semi-supervised learning. We demonstrate how this can be a very efficient approach to segmentation through pixel classification.

Index Terms—Image segmentation, object detection, semi-supervised learning.

I. INTRODUCTION

IN this paper we propose an interactive method for probabilistic classification of pixels, which can subsequently be used for segmentation of 2D and 3D images. Our approach is especially advantageous for detecting patterns, a situation regularly occurring in images of materials and medical samples. Such images often show a collection of objects which are to be separated from the background. One example we extensively use in this paper is concerned with detection of fibres in volumetric data of composite materials, see Fig 1. Another example is detection and segmentation of cells in histological images.

When segmenting images showing a collection of similar objects, an established strategy involves extensive modelling the appearance of the objects, usually leading to a highly specialised method. Another common strategy is to learn the appearance of the objects from a large amount of pre-labelled data, often with high computational requirements during the training phase. Here we aim for a general method that requires limited computation.

Our method fits into the framework of semi-supervised learning, combining two ingredients: a model for image content created in an unsupervised manner from the image

features, and a modest input from the user. When a user marks a structure in the image as belonging to a class, our method propagates the marks to similar structures in the rest of the image. The output is a layered image which at every pixel position contains the probabilities of belonging to each of the defined classes. We call this output *pixelwise probabilities* of belonging to segmentation classes. From this output, segmentation is readily obtained by selecting the most probable class for each pixel. The method is highly flexible and captures the features which are of interest to the user, an example with various image features is shown in Fig. 2. This approach allows easy segmentation of complex structures, that would otherwise require the development of algorithms targeted at specific problems.

An important property of our model is real-time feedback, allowing the user to place new markings strategically, depending on the current result. For this to work without delay, the segmentation must be updated very fast. Our method relies on a very efficient update of the parameters used for pixel classification, and equally efficient update of the classification results. With results shown immediately, the user can continue adding marks until the desired outcome is learned by the algorithm. Having learned the desired outcome, the classification model can be applied to other images of the same type in an unsupervised manner, without additional user input. This has many applications, for example, in microscopy or when segmenting slices from a volumetric image.

Our prototype implementation, including a graphical user interface, is in Matlab, and we made the code available through MathWorks File Exchange (search for InSegt).

A. Related work

Benefits of user input with real-time feedback have been recognised in image segmentation. A comprehensive summary of interactive approaches can be found in Boykov [1]. Here we review some important advances to place our method in the existing framework, and to explain how our method differs from the current trends in interactive segmentation.

Early interactive techniques for segmentation of highly complex images include intelligent scissors [2] or live wire [3], where the user cuts out an object by placing markers along its boundary. Based on edge information, the algorithm traces the boundary by finding the shortest path in an edge-weighted graph. These algorithms are computationally cheap, but require a lot of user effort to obtain a segmentation. Less user input is required when using interactive graph cuts [4], [5], which often give very impressive results with only

V. A. Dahl, C. H. Trinderup, M. J. Emerson and A. B. Dahl are with the Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark (e-mail: vand@dtu.dk; ctri@dtu.dk; mojn@dtu.dk; abda@dtu.dk).

Manuscript received date; revised date

a few seeds provided by the user. The algorithm separates the foreground from the background based on the boundary and region properties of segments. In the GrabCut method [6] the user provides a bounding rectangle, often leading to very precise foreground-background separation. Optional editing using brush strokes can be carried out to correct finer details. Extensions of GrabCut include shape priors [7] and an improvement to graph cut energy representation [8]. An alternative to combinatorial graph-based solutions is the use of continuous representation of segmentation boundaries. Such interactive active contours often minimise a variational segmentation functional [9], [10].

Common to the described methods is the focus on segmenting relatively large foreground objects, which justifies using regularisation on the length or the curvature of the segmentation boundary. In some applications it is, however, not possible to use a strong regulariser. For example, when segmenting the circular fibres shown in Fig. 1, regularisation would remove or merge small regions. The need for segmenting a number of small objects is often seen in areas like microscopy imaging for life science or materials science. Appearance of such images can vary significantly, with texture as well as intensity carrying information useful for obtaining the desired segmentation. A specialist would use such clues to distinguish among structures, but automating the segmentation task typically requires highly sophisticated and problem-adapted methods. While there are situations which justify the development of a specialised method, in many cases a reasonable result with modest interactive effort would be preferred.

When segmenting small image structures, e.g. cells, a well-suited approach is classification of pixels. This is the basis for the *ilastik* segmentation tool [11], which employs a random forest classifier [12] trained on image features including colour, edges, orientation and texture. The features are computed from the image before starting the interactive labelling of image structures, while parameters of the random forest classifier are learned from the manual labelling. When a user updates the labels to improve the segmentation, the parameters of the classifier need to be re-learned, which is computationally costly and causes a noticeable delay in the feedback. Another specialised tool for segmentation of microscopy images is trainable *Weka* segmentation [13] (a part of the *Fiji* [14] distribution of *ImageJ*) which utilises a data mining and machine learning toolkit for pixel classification problems. A user can choose from a variety of image features and interactively re-train the classifier.

Frameworks using neural networks are increasingly popular in pixel classification, and often yield impressive results [15]. A neural network operates on features extracted locally from the image. This input is fed through a series of linear functions, with a non-linear activation between them, ending up in a probabilistic output. The weights of the linear functions need to be trained by optimising the performance on the usually large set of labelled data. This provides an extreme flexibility to the method and, provided an adequate training, neural networks may solve pixel classification problems as accurately as specialists. However, these methods are dependent on large

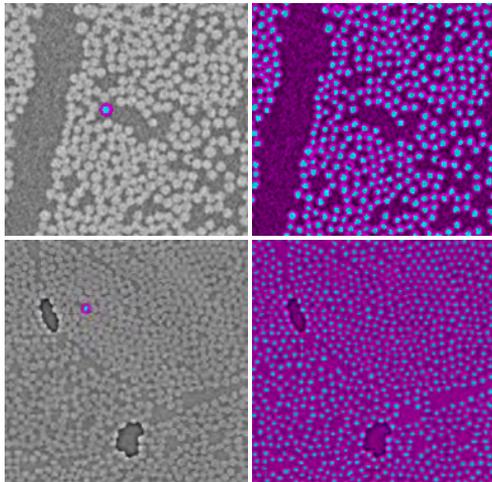


Fig. 1. Segmentation of glass and carbon fibers using our interactive pattern-based method. On the *left* input images and a very small subset of pixels manually labelled as either being close to fibre centre (cyan), or not being close to fibre centre (magenta). On the *right*, the result obtained by thresholding the probability image computed by using our method. The parameters used for these experiments are $M = 9$ and $K = 500$.

training sets and require computationally costly training, which makes them difficult to use for the task of segmenting a small set of images.

Our approach shares some similarities with neural networks. We also feed the input through linear functions with non-linear steps in between. However, our method uses the features extracted from the image to construct the linear functions in a preprocessing step. The functions are then kept fixed, while they operate on the interactive user input, resulting in a probabilistic output. Due to the fixed linear functions, our method is less adaptable than neural networks, and has limited use for semantic segmentation of photographs. Nevertheless, we achieve excellent results in case of patterned images, for our method does not require a large set of labelled data, and avoids performing a costly optimisation for each interactive update.

The foundation of our method is a linear operator encoding image content. This is done via a dictionary, by assigning image pixels to dictionary pixels. The relation between the image and the dictionary can be formulated as a bipartite graph and represented using a adjacency matrix. The approach has been used for evolving deformable models in [16], [17], [18]. In this work we use the image–dictionary relationship to propagate the brush strokes provided by the users.

II. METHOD

Our method combines two sources of information, the structure in the image and the user-provided partial labelling. The structure in the image is captured in the preprocessing step, namely clustering, which we describe in II-A. After that, in

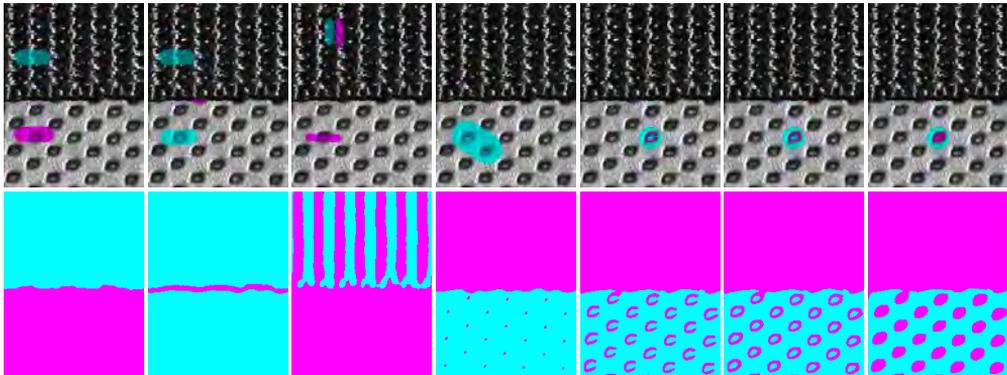


Fig. 2. An example demonstrating the flexibility of our method. In the *top* row, two manually labelled classes (cyan and magenta), corresponding to different image features. In the *bottom* row, a resulting image segmentation obtained by thresholding a probability image.

II-B, we explain how clustering is used for transforming the user-provided partial labelling into pixelwise probabilities of belonging to each of the classes. The central part of our segmentation, the interactive update, covered in II-C, is obtained by immediately displaying the result of the transformation and allowing the user to repeatedly improve the partial labelling.

Postprocessing choices, covered in II-D, are concerned with the output of the interactive update, which consists of two parts. The first and most obvious output is a probability image. The probability image can be thresholded for image segmentation, but other postprocessing methods may be utilised as well. For the second output, which we call *dictionary probabilities*, the user-provided partial labellings are propagated to the dictionary patches which were obtained in the preprocessing step. This encodes the learned information used for transforming the intensity image to the probability image, and can be used for subsequent automatic processing of similar images.

Our method comes in a range of flavours, governed by the features used for clustering and the strategy used for calculating pixelwise probabilities. In this section we only explain the simplest variant, the other possibilities are covered in Sec. III.

Notation. Throughout the paper we consider an image I defined on an X -by- Y image grid with pixel values in either grayscale or RGB colour space. During the interactive part, the user will be placing marks in the image grid, to indicate the pixels which belong to one of the C segmentation classes. We chose to represent this user-provided information with a layered *label image* L , where $L(x, y, c) = 1$ if the user indicated that pixel (x, y) belongs to class c , and 0 otherwise.

A. Clustering image patches

The aim of preprocessing is to find the structures in the image without considering the user-provided labels. In the framework of semi-supervised learning, a cluster assumption states that, if points are in the same cluster, they are likely

to be of the same class [19] (which does not imply that each class forms a single cluster). For our purpose, we assume that image features tend to form discrete clusters and that image features in the same cluster are more likely to share a class. However, we do not assume that each class is represented by only one cluster, so we will need much more clusters than classes. Therefore we create a multitude of clusters which capture the variety of features present in the image.

In Sec. III we will explain the implementation details and some more advanced ways of accomplishing clustering. In this section we outline the basic approach, which operates on intensity patches. For this case, only two parameters are required: the number of clusters K and the size of the patches M . The number of clusters should be large, measured in hundreds or thousands, and is roughly reflecting the variability in the image. The size of the patches should reflect the scale of the distinctive image features and could, for example, be 9 pixels.

For clustering, we extract M -by- M patches from the image, treat each image patch as a vector containing the pixel intensities and group those vectors into K clusters, e.g. using k-means clustering. The resulting collection of cluster centres represents the content of the image. As these basic elements are inferred from input data, we call the collection of K cluster centres a *dictionary*, and each of its elements (each cluster centre) is denoted *dictionary patch*. Every image pixel (x, y) in the centre of an M -by- M image patch is, by means of clustering, uniquely assigned to one cluster. We represent this using an *assignment image* A . For boundary pixels we define $A(x, y) = 0$.

B. Relation between image and dictionary

According to the cluster assumption, image patches assigned to the same dictionary patch are more likely to belong to the same class. Unique for our method is that we use this assumption on a pixel level. That is, if image patches are assigned to the same dictionary patch, their corresponding pixels (i.e.

the pixels at the same position in the patch) are more likely to belong to the same class. In other words, for every dictionary patch there is a certain (unknown) classification of its pixels, which all assigned patches are likely to share.

To exploit this assumption, we define a binary relation between pixels which we expect to belong to the same class. For example, a central pixel of an image patch assigned to a certain dictionary patch relates to central pixels of all other patches assigned to the same dictionary patch. Likewise, the pixel directly above the central pixel relates to corresponding pixels in other patches, and the similar relation extends to all positions in a patch. This results in M^2K cliques of pixels, one for every pixel in the intensity dictionary. Due to the overlap between image patches, every non-boundary pixel belongs to M^2 different cliques.

The central part of our method is concerned with transforming a user-provided partial labelling to pixel-wise probabilities. The transformation matrix we use has a very simple decomposition, which makes our method efficient and allows for immediate feedback to the user. The construction of the transformation matrix is therefore fundamental for our method. However, describing how this matrix is constructed provides little intuition about our method, so we start by motivating our approach.

As covered previously, the assignment image A , obtained in an unsupervised manner, contains information on clusters of structures in the image. We want to combine this information with the user-provided partial labelling L . We do this by creating a labelling dictionary to accompany our intensity dictionary. For each cluster $k \in \{1, \dots, K\}$ we use A to identify the locations of all image patches which are assigned to the cluster. From those locations in the image grid we extract corresponding patches of the labelling image L . For the set of labelling patches we compute a pixel-wise and label-wise average. As a result, every M -by- M dictionary patch has now a corresponding M -by- M labelling representation consisting of C layers.

When the image is fully labelled, the label image L sums to one in every pixel, as only one out of C classes has a label 1. Consequently, the labelling representation of every dictionary element also sums to one in every pixel. The values of this representation are not binary, they instead encode the normalised frequency of a dictionary pixel being labelled as belonging to class c in the current labelling image. For this reason, we think of this labelling representation as of pixelwise probabilities of belonging to class c , and we call them *dictionary probabilities*.

Dictionary probabilities can now be pasted back into an X -by- Y image grid, again using the location information from A , and again averaged in every pixel. This results in an X -by- Y probability image P consisting of C layers, where P is a diffused version of L . In other words, we use the self-similarity information encoded by A to diffuse the user-provided labelling information from L onto the rest of the image. Diffusion consists of two averaging processes, the first when producing dictionary probabilities and the second when producing image probabilities.

Now we turn to explaining the construction of the transfor-

mation matrices used for efficient computation of dictionary probabilities and image probabilities. Fundamental for this transformation is the relation between the X -by- Y image grid and the M -by- M -by- K dictionary grid. This relation will be encoded using an n -by- m biadjacency matrix \mathbf{B} , where $n = XY$ and $m = M^2K$. For this purpose, we need a linear (single) index for the pixels in the image and the pixels in the dictionary grid.

The linear index of an image pixel (x, y) is

$$i = x + (y - 1)X. \quad (1)$$

As for the dictionary grid, we use $(0, 0, k)$ for the central pixel of the k -th dictionary element, and coordinates of other pixels in the patch are defined in terms of within-patch displacements Δx and Δy , both from $\{-s, \dots, 0, \dots, s\}$ with $s = (M - 1)/2$. A dictionary pixel at coordinates $(\Delta x, \Delta y, k)$ has a linear index

$$j = (\Delta x + s) + (\Delta y + s)M + (k - 1)M^2. \quad (2)$$

Each assignment of an image patch centered around (x, y) to a k -th dictionary patch centered around $(0, 0, k)$ induces a relation between the M^2 image pixels and the M^2 dictionary pixels, see Fig. 3. The reason is that not only central pixels are related, but all pixels in the patch are. Using \sim for denoting a relation between image pixels and dictionary pixels gives

$$A(x, y) = k \Rightarrow \begin{array}{l} (x + \Delta x, y + \Delta y) \sim (\Delta x, \Delta y, k), \\ \text{for all } \Delta x \text{ and } \Delta y \end{array} \quad (3)$$

Since image patches are overlapping, every non-boundary image pixel relates to M^2 dictionary pixels. Image pixels in a boundary relate to less than M^2 dictionary pixels, and the four corner pixels relate to only one dictionary pixel. In total there are $(X - 2c)(Y - 2c)M^2$ relations between the image pixels and the dictionary pixels.

We represent the relations between n image pixels and m dictionary pixels using an n -by- m biadjacency matrix \mathbf{B} , with elements

$$b_{ij} = \begin{cases} 1 & i \sim j \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where i and j are linear indices of an image pixel and a dictionary pixel. The algorithm for constructing \mathbf{B} is summarised in Alg. 1

Algorithm 1 Construction of \mathbf{B}

- 1: Initiate \mathbf{B} as an n -by- m matrix with $b_{ij} = 0$
 - 2: **for** a non-boundary pixel (x, y) **do**
 - 3: Retrieve pixel assignment $k = A(x, y)$
 - 4: **for** within-patch displacement $(\Delta x, \Delta y)$ **do**
 - 5: compute i for $(x + \Delta x, y + \Delta y)$ using Eq. (1)
 - 6: compute j for $(\Delta x, \Delta y, k)$ using Eq. (2)
 - 7: assign $b_{ij} = 1$
 - 8: **end for**
 - 9: **end for**
-

The biadjacency matrix \mathbf{B} defines the linear mapping used to propagate the information from the image to the dictionary and vice versa. Consider a quantity defined on the image grid

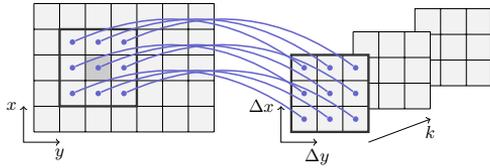


Fig. 3. A subset of relations between a 7×5 image and a $3 \times 3 \times 3$ dictionary caused by the framed patch centered around the darker pixel being assigned to the first dictionary patch.

(e.g. user-provided markings indicating pixels which belong to class 1) arranged into a length n vector \mathbf{v} such that the i -th element contains the value of the i -th image pixel. Propagating these values to the dictionary is carried out by calculating a length m vector

$$\mathbf{d} = \text{diag}(\mathbf{B}^T \mathbf{1}_{n \times 1})^{-1} \mathbf{B}^T \mathbf{v}, \quad (5)$$

where $\mathbf{1}$ denotes a column vector of ones, while $\text{diag}(\cdot)$ denotes a diagonal matrix with the diagonal defined by the argument. Each element of \mathbf{d} contains the value of the j -th dictionary element computed by averaging the values of the related pixels. The summation is accomplished by multiplying with \mathbf{B}^T while the diagonal matrix accomplishes the division with the total number of related pixels.

For this reason we define the m -by- n transformation matrix for mapping from the image to the dictionary as

$$\mathbf{T}_1 = \text{diag}(\mathbf{B}^T \mathbf{1}_{n \times 1})^{-1} \mathbf{B}^T. \quad (6)$$

Similarly, mapping from the dictionary to the image is given by the n -by- m matrix

$$\mathbf{T}_2 = \text{diag}(\mathbf{B} \mathbf{1}_{m \times 1})^{-1} \mathbf{B}. \quad (7)$$

Those two transformation matrices are fundamental for our method. The diffused user-provided labellings (as described in the motivational paragraphs) are computed as

$$\mathbf{P} = \mathbf{T}_2 \mathbf{T}_1 \mathbf{L}, \quad (8)$$

where \mathbf{L} is the user-provided labelling L arranged in a n -by- C matrix. The resulting n -by- C matrix \mathbf{P} can also be arranged back into an layered image P , which is a diffused version of L .

C. Interactive update

When equipping our method with the user-provided interactive update, we run into choices with regards to: i) the way in which we treat non-labelled pixels, ii) the number of applied diffusion steps, and the way of treating intermediate results between the steps, and iii) the possibility of changing the number of segmentation classes, the number of segmentation classes. After testing many types of the interactive updates, we kept three main versions. In all versions the number of classes C is chosen during initialisation and kept fixed during the update.

The way in which we handle pixels that have not been labelled by the user is also common to all versions. Such pixels

are assigned equal probability of belonging to each class. As a result, before the user places the first label, all probabilities are equal and no segmentation is possible.

The user starts the interaction by choosing a pencil corresponding to one of the C classes and scribbles over some pixels. The partial labelling information is immediately transformed to the probability image and shown to the user as an image segmentation, such that every pixel is placed in the class with the highest probability. Only one class will have values larger than $\frac{1}{C}$ in the label image L , and the same applies for the probability image P computed using (8). Thus, after the first pencil stroke, some pixels will belong to the marked class and no pixels will be assigned to the classes that have not been labelled yet.

Thanks to the real-time feedback, a user can quickly improve the result by placing markings in misclassified regions (the regions that have been incorrectly classified). With many un-labelled pixels in L , the image P will typically have many values that only differ slightly from $\frac{1}{C}$. Those small deviations carry the information needed for inferring the class of the unlabelled pixels.

As for the number of applied diffusion steps, we use either one or two. When using two diffusion steps, instead of continuing to diffuse the (already diffused) probability image, we can apply additional operations in between the two diffusions. Very good results are obtained if we apply *binarisation* of the labels between the two diffusion steps. For binarisation we identify the class of the highest probability for each pixel, and apply $\{0, 1\}$ labelling. If there are pixels with no clear probability maximum, we let them retain their unresolved labels. Consequently, for the second iteration of the diffusion, many pixels act as labelled, and this improves the quality of the result.

Another additional operation for the two-step diffusion involves the subset of pixels which has been labelled by the user. After the first diffusion step, the $\{0, 1\}$ labelling of those pixels has probably dispersed, and some might even have a maximal probability in a class different from the markings indicated by the user. The operation of *overwriting* imposes the original user-provided labelling to all labelled pixels in between the diffusion step.

The options for the two-step diffusion, binarisation and overwriting, are implemented in our segmentation tool, such that the user can quickly switch between the variants of the method and decide which one yields the best results for the data at hand. Likewise, the user can quickly determine whether the quality of the results is sufficient or decide to place additional scribbles in the image.

The user can choose to see the output of the classification displayed as a final segmentation based on the resulting probability image. Alternatively, there is an option for inspecting the C probability images, which often gives a better insight into the quality of the result. In some cases the final classification may seem incorrect, but the probability images do contain useful information which can be postprocessed in some other way to obtain the desired result.

D. Postprocessing

Our approach allows for two postprocessing strategies. One strategy involves processing the probability image to obtain the segmentation. The other strategy involves detection of interesting features from the probability image. These operations are application-driven and examples are illustrated in the Sec. IV.

The second strategy involves reusing the information stored in the dictionary and the associated dictionary probabilities. The linear transformation (8), which is core to our method, first transforms the user-provided markings from L to the dictionary space (using matrix \mathbf{T}_1) and then back to the image space (using matrix \mathbf{T}_2). Consider the first product

$$\mathbf{D} = \mathbf{T}_1 \mathbf{L}.$$

This is an m -by- C matrix containing the pixelwise probabilities of the dictionary pixels (i.e. the dictionary probabilities) which can be useful for processing a previously unseen image similar to I .

Processing a new image \hat{I} requires extracting all M -by- M patches for every pixel of \hat{I} and assigning those patches to the *existing* dictionary, i.e. the dictionary created using patches from I . Just like before, this assignment defines an image-to-dictionary relation and the closely related transformation matrices. Here we are interested in the dictionary-to-image transformation $\hat{\mathbf{T}}_2$. To compute the probability image corresponding to the unlabelled image \hat{I} we therefore need to compute

$$\hat{\mathbf{P}} = \hat{\mathbf{T}}_2 \mathbf{D}.$$

and rearrange the result into \hat{P} .

This way of using our method fits into the framework of supervised learning. The original image I and the computed labelling L can in this context be seen as a labelled training set (ignoring the fact that the labelling is computed in a semi-supervised way). Our method is then capable of producing the probability image \hat{P} for the new, unlabelled image \hat{I} . The approach will work as long as the initial clustering captures the features present in both I and \hat{I} . This will hold for similar images.

III. IMPLEMENTATION DETAILS

A multitude of implementational choices have been made when developing our method. The part of the method concerning the update eventually converged and only involves a few of the choices mentioned earlier (running the diffusion step twice and discretising between diffusion steps).

As for the clustering step for preprocessing the data, our experience leads to two conclusions. First, our method is rather robust to the *quality* of the clustering, so using an approximate result will generally not deteriorate the output. Second, the features used for clustering need to reflect the distinction in the appearance of the classes we want to separate. For many types of images, an intensity-based approach as sketched in Sec. II will perform reasonably well. However, in challenging cases, more elaborate image features might provide better results. In this section, we briefly touch upon different possibilities.

For simplicity, we always assume that the size of the image patches M is odd and patches are centred around the central

pixel, but we foresee a straightforward generalisation to patches of different sizes. When clustering intensity patches, different similarity measures can be employed for grouping, but in our experiments we use either squared Euclidean distance or normalised cross correlation. Furthermore, depending on the problem at hand, it might be advantageous to normalise all image patches prior to clustering. This is true when the distribution of the intensities, rather than the absolute value, is carrying the information we want to utilise for segmentation.

With the method being robust with respect to the quality of the clustering, we focus on efficiency when building the dictionary. Therefore, we choose to use a k-means tree [20], which is built from consecutive k -means clusterings. In this implementation, the size of the dictionary is defined in terms of the branching factor b and the number of layers t . Since each node in the tree makes up a dictionary element, the total number of dictionary elements is given by $K = \frac{b^{t+1}-1}{b-1}$.

Our experience is that good performance is obtained also without running the k-means until convergence for each layer, and therefore a fixed number of iterations is chosen, e.g. 10 iterations. Furthermore, in order to limit the computational burden and memory usage, we extract only a subset of M -by- M patches from the image when building the dictionary.

As for producing A given the dictionary represented by a k-means tree, the patch vector is compared with the nodes in the first layer to find the match. The patch vector is then compared to the children of this node, and the most similar node is again chosen. This process is repeated until a leaf node or an empty node is reached. The patch vector is assigned to the most similar node along this path.

Apart from clustering image patches, we also experimented with different image features. Some of the results we show in Sec. IV are based on SIFT [21], but other features can also be incorporated in our method. The approach is as following. First, image features represented by vectors are extracted from all pixel positions in the image and clustered in K clusters. For speed, it often suffices to consider only a subset of pixels for clustering, as long as we capture the variability in the image. Every position (x, y) from the image grid can now be uniquely assigned to one of the k clusters—the cluster that is closest to the feature vector extracted at (x, y) . This results in an assignment image A . The only additional information we need for building the transformation matrices is a value M , which earlier represented the size of the extracted image patches. The value M now determines the size of the overlap when linking the image to the dictionary. While we now freely chose M , it is reasonable to use a value which corresponds to the size of the extracted features.

IV. RESULTS

In Fig. 4 we show a three-class classification of a volumetric X-ray image of peripheral nerves appearing as tubular structures. Using a purely intensity-based approach to pixel classification, it would be difficult to differentiate between the bright background and the bright regions inside the dark tubes. Furthermore, a significant bias field makes it difficult to choose a global threshold. Our approach utilises a very limited user

input in just one slice of the volume to differentiate between three classes: background, tubes and inside. Furthermore, the learned classification can be used for automatic classification of all other slices in the volume, yielding a volumetric segmentation.

Fig. 5 shows an example of segmenting a volumetric image of a fibre composite into two classes: background class and fibre centre class. Using our method, a huge number of individual fibres can be segmented with a modest user input. The probability image of a fibre centre class precisely indicates a region for each fibre centre, and can readily be used in postprocessing for obtaining a quantitative information about the spatial distribution of fibres. In this example we also use the result of single-slice segmentation for batch processing of a whole volume stack. In principle, this yields the trajectory of each individual fibre. For comparison, we also show a result obtained by thresholding the intensity image. This nicely illustrates a challenge in segmenting densely packed fibres, when the image resolution does not suffice to clearly delineate the boundary of every individual fibre. In this case, a successful segmentation requires utilisation of the repetitive patterns in the image. Our method accomplishes this via clustering.

In Fig. 6 we show a three-class segmentation of onion cells. Since cell walls and nuclei both appear dark, a purely intensity-base method would not distinguish these two classes—a task which our methods successfully solves with only a modest user input.

In Fig. 7 we show the use of our method for counting cells in a stained microscopy image. Unlike other examples, this is a colour (RGB) image. To utilise colour information, the features extracted from every image patch contains three color channels concatenated in a single feature vector. Since the final goal is to count and measure the size of cells, we postprocess the probability images to obtain individual cell segmentation. We detect the center of each cell by computing the local maxima of the center-class probability image and we estimate the extent of each cell by considering both the centre-class and boundary-class probability images, coupled with the distance from the previously detected cell centers.

V. CONCLUSION

Instrumental for our method is a pair of closely related linear transformations which propagate the information from the image grid, via a dictionary, and back to the image. The transformations are constructed such that the propagation is strong between image pixels with similar appearance, captured by some extracted image features.

In this paper we present an algorithm for building an efficient matrix representation of those transformations, allowing a real-time processing. We demonstrate how this can be used for an interactive image segmentation by propagating user-provided labelling. Furthermore, such a segmentation of one image allows for subsequent automatic processing of similar images.

With only a modest user input, our method can yield good results when segmenting patterned images. We find this extremely useful for many tasks in materials and life sciences.

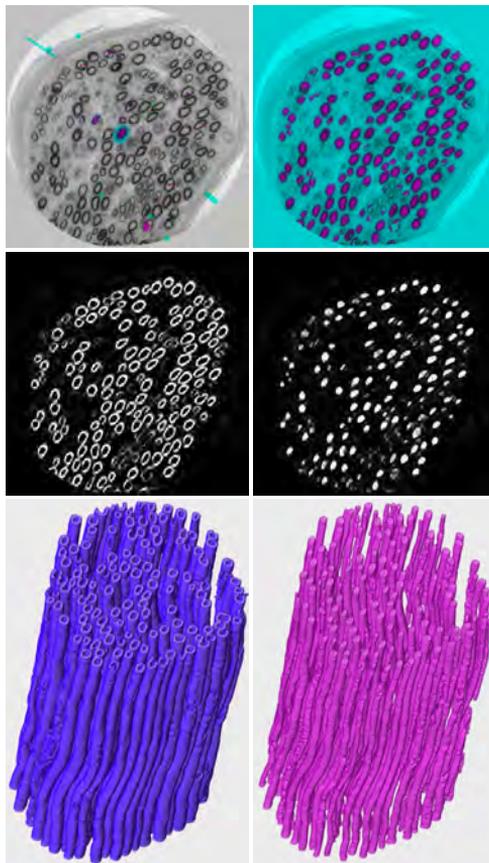


Fig. 4. Volumetric segmentation of peripheral nerves. In the *top* row, a slice from the volumetric data with overlaid limited user input and the three-class segmentation dividing the pixels into background (cyan), tubes (purple) and inside (magenta). The *middle* row shows two layers of probability images, corresponding to the tube class and the inside class. High intensity indicates a high probability of belonging to the class. The *bottom* row shows the 3D visualization of the volumetric data obtained by processing a full image stack and thresholding the resulting probability volumes. This experiment was performed using $M = 9$, $K = 4000$ and a clustering based on SIFT features. Small connected components of less than 10^4 voxels were removed during postprocessing.

REFERENCES

- [1] Y. Boykov, "Interactive segmentation," in *Computer Vision, A Reference Guide*, 2014, pp. 416–422.
- [2] E. N. Mortensen and W. A. Barrett, "Intelligent scissors for image composition," in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 1995, pp. 191–198.
- [3] A. X. Falcão, J. K. Udupa, S. Samarasekera, S. Sharma, B. E. Hirsch, and R. d. A. Lotufo, "User-steered image segmentation paradigms: Live wire and live lane," *Graphical models and image processing*, vol. 60, no. 4, pp. 233–260, 1998.
- [4] Y. Y. Boykov and M.-P. Jolly, "Interactive graph cuts for optimal boundary & region segmentation of objects in nd images," in *IEEE*

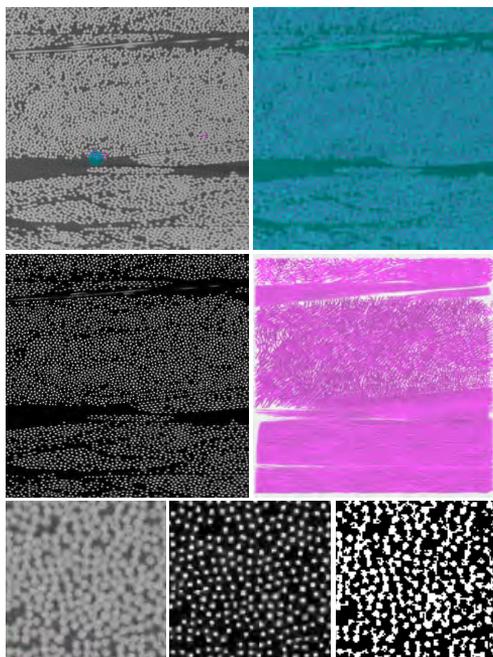


Fig. 5. Volumetric segmentation of glass fibres. In the *top* row a slice with manual input indicating fibre centers (magenta) on a background (cyan) class, together with a resulting two-class segmentation. The *middle* row shows a layer of a probability image corresponding to the fibre centres class. On the *middle* to the *right* is the output of processing a full volumetric stack and thresholding the probability volume. From the 3D visualisation it is evident that fibres form clusters of different orientations. The *bottom* row shows a zoom-in on the central part of the image slice, together with the corresponding probability image and (for comparison) a corresponding segmentation obtained by directly thresholding the image intensities. Settings used in this experiment are $M = 9$, $K = 4000$ and

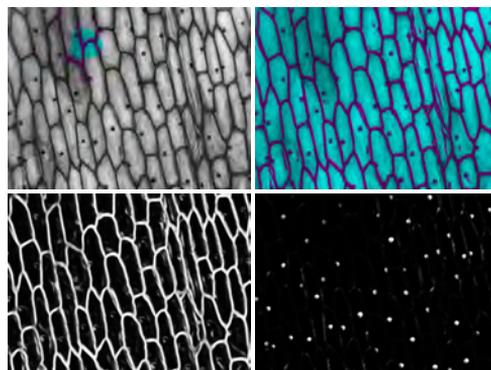


Fig. 6. A three-class segmentation of an image showing onion cells. In the *top* row an image with manual input and a segmentation into three classes: background (cyan), nucleus (purple) and wall (magenta). In the *bottom* row the probability images for the wall and the nucleus class. Settings used are $M = 9$ and $K = 4000$.

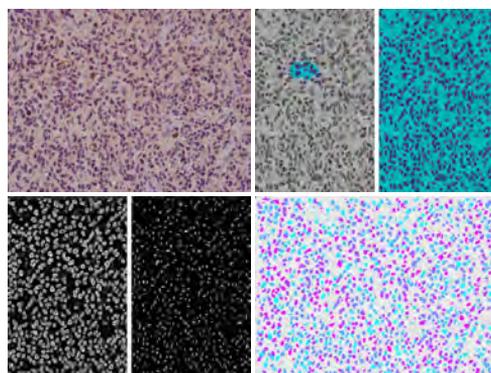


Fig. 7. A three-class segmentation of a histopathology image. In the *top* row an original colour image. The extent of the manual input and a corresponding segmentation into three classes are shown in the *right*, with a frame cropped to show the central column of the image. Likewise, in the *middle* row we show the probability images for the two classes for the central part of the image. In the *bottom right* is the final result obtained through additional postprocessing to distinguish individual cells. Settings used are $M = 5$ and $K = 4000$.

- International Conference on Computer Vision*, vol. 1. IEEE, 2001, pp. 105–112.
- [5] Y. Boykov and G. Funka-Lea, “Graph cuts and efficient ND image segmentation,” *International journal of computer vision*, vol. 70, no. 2, pp. 109–131, 2006.
 - [6] C. Rother, V. Kolmogorov, and A. Blake, “Grabcut: Interactive foreground extraction using iterated graph cuts,” *ACM Transactions on Graphics (TOG)*, vol. 23, no. 3, pp. 309–314, 2004.
 - [7] B. L. Price, B. Morse, and S. Cohen, “Geodesic graph cut for interactive image segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2010, pp. 3161–3168.
 - [8] M. Tang, L. Gorelick, O. Veksler, and Y. Boykov, “Grabcut in one cut,” in *IEEE International Conference on Computer Vision*. IEEE, 2013, pp. 1769–1776.
 - [9] M. Unger, T. Pock, W. Trobin, D. Cremers, and H. Bischof, “Tvsseg: interactive total variation based image segmentation,” in *Proceedings of the British Machine Vision Conference*, vol. 31. Citeseer, 2008, pp. 44–46.
 - [10] J. Santner, T. Pock, and H. Bischof, *Interactive multi-label segmentation*. Springer, 2011.
 - [11] C. Sommer, C. Straehle, U. Koethe, F. Hamprecht *et al.*, “ilastik: Interactive learning and segmentation toolkit,” in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2011, pp. 230–233.
 - [12] L. Breiman, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
 - [13] I. Arganda-Carreras, V. Kaynig, C. Rueden, K. W. Eliceiri, J. Schindelin, A. Cardona, and H. Sebastian Seung, “Trainable weka segmentation: a machine learning tool for microscopy pixel classification,” *Bioinformatics*, p. 180, 2017.
 - [14] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid *et al.*, “Fiji: an open-source platform for biological-image analysis,” *Nature methods*, vol. 9, no. 7, pp. 676–682, 2012.
 - [15] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
 - [16] A. B. Dahl and V. A. Dahl, “Dictionary snakes,” in *22nd International Conference on Pattern Recognition (ICPR)*. IEEE, 2014, pp. 142–147.
 - [17] ———, “Dictionary based image segmentation,” in *Scandinavian Confer-*

- ence on Image Analysis*. Springer, 2015, pp. 26–37.
- [18] V. Dahl and A. Dahl, *A Probabilistic Framework for Curve Evolution*. Springer, 2017, pp. 421–32.
 - [19] O. Chapelle, B. Scholkopf, and A. Zien, Eds., *Semi-supervised learning*. The MIT Press, 2006.
 - [20] D. Nister and H. Stewenius, “Scalable recognition with a vocabulary tree,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2. IEEE, 2006, pp. 2161–2168.
 - [21] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.

Statistical validation of individual fibre segmentation from tomograms and microscopy

To be submitted December 2017, *Composites Science and Technology*.

Statistical Validation of Individual Fibre Segmentation from Tomograms and Microscopy

Monica Jane Emerson^a, Vedrana Andersen Dahl^a, Knut Conradsen^a, Lars Pilgaard Mikkelsen^b and Anders BJORHOLM Dahl^a

^a*Department of Applied Mathematics and Computer Science, Technical University of Denmark.*

^b*Department of Wind Energy, Technical University of Denmark.*

Abstract

Imaging with X-ray computed tomography (CT) enables non-destructive 3D characterisations of the micro-structure inside fibre composites. In this paper we validate the use of X-ray CT coupled with image analysis for characterising unidirectional fibre composites. We compare X-ray CT at different resolutions to optical microscopy and scanning electron microscopy, where we characterise fibres by their diameters and positions. In addition to comparing individual fibre diameters, we also model their spatial distribution, and compare the obtained model parameters. Our study shows that X-ray CT is a high precision technique for characterising fibre composites and, with our suggested image analysis method for fibre detection, high precision is also obtained at low resolutions. This has great potential, since it allows larger volumes to be analysed. Besides analysing representative volumes with high precision, we demonstrate that studying complete bundles at the fibre scale can provide valuable insights about the fibre manufacturing process.

Keywords: Individual fibre segmentation A. Glass fibres C. Statistics D.

Non-destructive testing

*Monica Jane Emerson (monj@dtu.dk). Technical University of Denmark, Richard Petersens Plads, Building 324, 2800 Kgs. Lyngby, Denmark.

1. Introduction

Imaging is an important method for characterising the micro-structure of materials. Characterising a material requires measurements of the size and shape of material features, and this often involves an automated image analysis method. Therefore, the image-based characterisation of a material's micro-structure can be seen as a coupling of the imaging system and the image analysis method.

Some imaging modalities, like optical microscopy (OM) and scanning electron microscopy (SEM), directly depict the surface of the imaged sample. Other modalities measure the image indirectly, and X-ray micro-CT (computed tomography) is such a method. Here the attenuation of X-rays penetrating a sample is measured, and a volumetric image is reconstructed from multiple measurements. Since the reconstruction model can introduce artefacts, it is important to determine the accuracy of the measurements obtained from CT.

CT has for a long time shown to be useful for characterising the micro-structure of fibre composites [1], and the resolution of micro-CT images allows for studying fibre bundles or individual fibres in 3D [2]. Current micro-CT laboratory X-ray scanning systems produce volumetric images of $1k^3$ - $2k^3$ voxels, often depicting more than 10K distinguishable fibres, and synchrotron imaging often provides larger volumes (more voxels). Automated image analysis methods are required to obtain quantitative information from such images.

With sufficient resolution it is possible to characterise individual fibres, which typically involves a segmentation of the fibre to estimate its position and diameter. Individual fibre segmentation poses several challenges. In composites with high fibre volume fractions (FVF) the fibres are close together, so it can be difficult to separate them. Image noise also makes it challenging to automatically segment fibres, and reducing the noise in the image acquisition process requires longer scan times. These challenges become more pronounced when reducing the spatial resolution to include more fibres in the scanned volume.

Imaging with CT gives some freedom in choosing the spatial resolution of the

recorded data, within the limits determined by the scanning setup. However, there is a trade-off between the size of the imaged volume (field of view) and the size of the smallest distinguishable features (spatial resolution). Increasing the field of view will reduce the spatial resolution. For this reason, it is important to
35 investigate the accuracy of the combined CT acquisition and analysis pipeline across resolutions.

Assessing the precision of a scanning system can be done using a physical phantom reference object with similar material properties and known geometry. If a reference object is available, it should also have similar geometric features
40 to evaluate the image analysis method employed. An example of validating the precision of CT scanning with a physical phantom was shown in [3]. Their approach used a simple threshold-based analysis. In our investigation such a phantom should resemble geometry of fibres, and we do not have such an object. Alternatively, the size of the fibres can be measured with a system that is more
45 precise than micro-CT, which is why we use OM and SEM.

Employing virtual computerised phantoms, composed of artificially created images, will provide a ground truth reference for the analysis method, and this has been used in e.g. [4, 5, 6]. Typically, it is difficult to make a virtual phantom that resembles the imaging system accurately enough to assess its precision, in
50 this case the micro-CT scanning. Another method for validating a quantification pipeline (scanning technique and image analysis method) is to compare obtained measures to derived physical parameters, as shown in [7, 8, 9].

Studies for evaluating different tomographic techniques for characterising defects in composites have been demonstrated in [10, 5, 11]. These studies fo-
55 cused on porosity in composites, and have been carried out on images where individual fibres are not resolved. Nikishkov et al. [12] also perform void geometry quantification using X-ray CT, and they validated their approach through microscopy measurements.

For individual fibre characterisation, a segmentation method is typically
60 employed. Methods for segmenting individual fibres have been presented in [13] for a relatively small number of fibres and in [14] for a much larger number

of fibres. Both methods carried out finite element simulations based on the segmented fibres.

In [8] we employ a model for fibre segmentation that can handle large volumes of low quality images, which has shown to be highly robust to noise and allows precise segmentation of fibres from low resolution scans, as described in [4]. The method was used for extracting fibre orientations and relating these to the compression strength of the material [8]. In the current paper we use this segmentation method for estimating the diameter and position of fibres to compare the micro-CT scans to SEM and OM images fibre by fibre.

For validation, we compare the results obtained from laboratory X-ray CT (XCT) at three resolutions, synchrotron radiation CT (SRCT), OM and SEM images. The number of pixels covering a fibre in the SRCT and high resolution XCT is similar to the images analysed in [13, 14] whereas mid and low resolution scans have fewer pixels relative to the fibre size. This allows us to investigate how the loss in resolution affects the fibre diameter and position estimates, hereby assessing the degradation of the image analysis methods when increasing the imaged volume.

We have acquired OM and SEM images of very high spatial resolution with between two and fifteen times smaller pixels and, since these methods directly depict the fibres, we employ SEM and OM as reference for the validation. We do not have a direct measure of their precision, so we will make a relative comparison to micro-CT. SEM and OM images are obtained from a polished surface of the fibre sample, imaged orthogonal to the fibre direction, and from these images we can compute the fibre diameters with high precision.

Through validation we demonstrate the high accuracy of our method in measuring fibre diameters from micro-CT, also at the very low spatial resolution. Therefore, we are capable of analysing fibres from data which covers a large field of view. As such, we demonstrate a high precision method for measuring the fibre diameter distribution inside complete fibre bundles.

2. Materials

This study is based on imaging a non-crimp fabric (NCF) glass fibre reinforced polymer. A detailed description of this type composite is given in [15]. We could also have used other types of unidirectional (UD) composites such as
95 composites manufactured using filament winding, pultrusion or pre-preg, and even composites based on carbon fibres.

In the specific NCF, glass fibres are arranged in bundles. The load carrying part of the composite is made of UD fibre bundles that are stitched on backing bundles. Figure 1 shows the sample with the imaged surface marked in green
100 and the imaged volume delineated with a red box. The investigated fibres are from the UD part of the composite, and the product is specified to have approximately 4000 fibres per bundle, an overall fibre volume fraction (FVF) of 57% and an average fibre diameter of $17\ \mu\text{m}$.

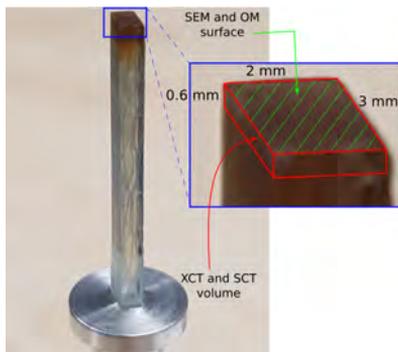


Figure 1: Glass fibre sample mounted for CT scanning. The volume outlined by the *red* box has been CT-scanned, and the top surface in *green* has been imaged using SEM and OM.

Microscope imaging based on OM and SEM was chosen to directly depict
105 the surface of the sample. For the SEM image to be acquired a thin layer of gold was added to the top of the sample, which was polished away after scanning. The XCT images were obtained by scanning the sample in a laboratory X-ray CT scanner (Carl Zeiss Versa 520) at three different resolutions, and finally the SRCT scan was acquired at the synchrotron ESRF (European Synchrotron

110 Radiation Facility, ID19). For XCT and SRCT volumetric images are obtained indirectly by computing the image using reconstruction methods. The reconstruction for the lab-based XCT was performed with the software provided by the scanner supplier (Carl Zeiss Versa 520) and for the SRCT image the filtered back projection algorithm was employed. The data files for the six scans can be
 115 found in [ref Data-in-Brief].

Table 1 lists the pixel/voxel sizes reported by the instruments, which deviate a little from those measured by comparing the images. Therefore, the image-based measured pixel/voxel sizes relative to the OM image are also given in Table 1, these were obtained through the image registration process described
 120 in Section 3.1.

Table 1: Instrument reported and measured pixel/voxel sizes for the different scans.

Scan	Reported	Reg. Scale
SEM	0.1852 μm	0.1882 μm
OM	0.2908 μm	0.2908 μm (reference)
SRCT	0.6500 μm	0.6440 μm
XCT _H	1.0376 μm	1.0356 μm
XCT _M	1.6856 μm	1.6835 μm
XCT _L	2.8059 μm	2.8082 μm

Images from the six scans employed in this study are shown in Fig. 2 after rotation and translation. To illustrate the actual resolution we did not resize the images. The size of the region that will be analysed is approximately the size covered by the images in Fig. 2, of length 1.02 mm and width 0.42 mm. For
 125 the three-dimensional modalities, 60 slices in a depth of 0.6 mm were analysed. This depth is shown in Fig. 1 to illustrate how much of the sample is included in the analysis.

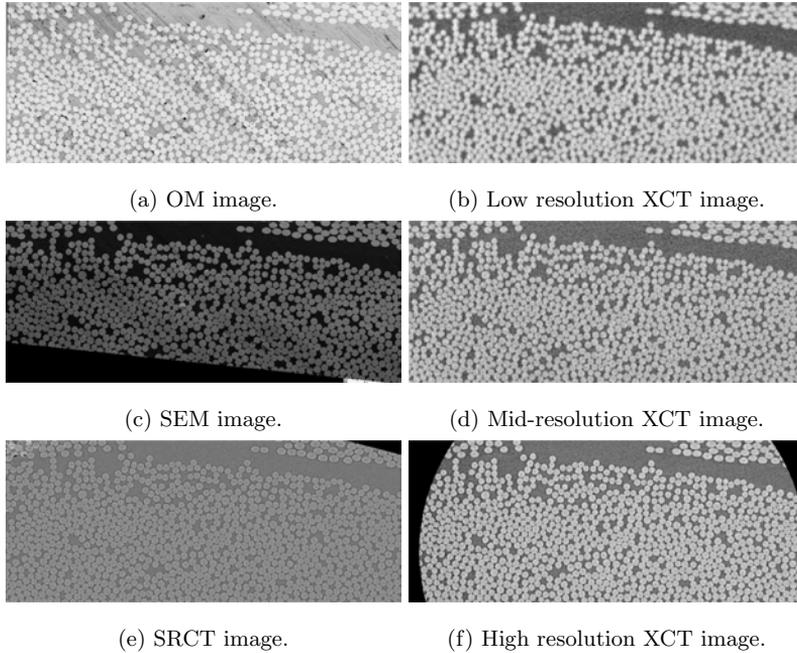


Figure 2: All six scans aligned after rotating and translating. Note that scaling has not been applied to show the actual resolution of the images.

3. Methods

To evaluate the use of CT imaging for fibre characterisation we have chosen to compute diameters and positions of individual fibres. The evaluation is based on comparing corresponding fibre diameters obtained from the six scans employing two different analysis methods.

3.1. Fibre segmentation and matching

Fibres are modelled as circles, each of them defined by a centre point and a diameter. These fibre parameters are calculated by segmenting the fibres as circles from the images shown in Fig. 2. To get an impression of the 3D structure of the sample, Fig. 3 shows a 3D rendering of the high resolution XCT scan.

Two methods were used for extracting fibre parameters, the circular Hough transform [16, 17] and our probabilistic feature labelling method [18], which we

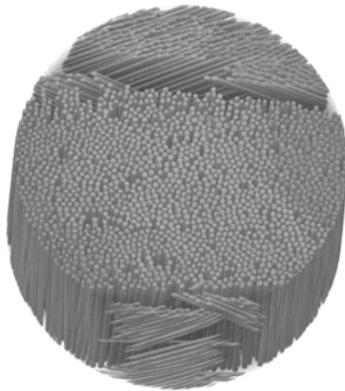


Figure 3: Rendering of fibres in the high resolution XCT scan.

140 have used for fibre detection in [8, 4]. In [8] we characterised fibres in 3D with fibre trajectories, centre lines connecting the 2D centre coordinates that belong to the same fibre. In contrast to that, here we compare micro-CT to 2D OM and SEM, so we only need to compute fibre parameters in 2D.

The circular Hough Transform is a two stage algorithm [16, 17]. In the first stage an edge image is computed, and in the second stage circles are matched to the edge image through a voting scheme. The algorithm we use has three parameters including an edge parameter, determining when a pixel is considered an edge pixel, a sensitivity parameter, determining when a circle is detected, and a parameter giving a circle diameter range. In the employed implementation, the diameter range is given as integers, which especially for low resolution images will give a coarse approximation of the fibre diameters. To obtain a higher precision in the diameter estimation, we up-scaled the CT-scanned images by linear interpolation. The SRCT and high resolution XCT were up-scaled with a factor 2, mid-resolution XCT with a factor 3 and low resolution XCT with a factor 4. Additionally, the sensitivity parameter was adjusted to find all fibres in the set used for the comparison.

The probabilistic feature labelling is a general method that we have developed for detecting repetitive image features [18, 8, 4]. The method detects

repetitive image features through a learning scheme based on limited user input,
160 and here we use it for detecting fibre centre points. This detection returns an
image where each pixel has a probability of being a centre point. To compute
the centre coordinates, the probability image is thresholded and the average
position of each connected component gives the centre position of the fibre.

After detecting the centre points, diameters must be computed. Since the
165 employed UD fibres are very dense, the diameter is computed as the distance
to the nearest centre point. To utilise our volumetric CT scan, we compute
the average of the fibre diameters in 60 slices, which gives extra robustness to
the measurements. Despite computing fibre parameters in 60 slices, we only use
the centre position from the top slice in our comparative study.

170 For a direct comparison, we must establish correspondence between the im-
aged fibres. Since the OM and SEM images are obtained from the top surface
of the sample and the micro-CT images are obtained just beneath this surface,
we assume a negligible change in fibre position. Therefore, the main difference
between the images can be modelled using a rotation, translation and scale.
175 The transformation is computed from $p = sRq + t$, where p and q are two set
of corresponding points, s is a scaling factor, R is a rotation matrix and t is a
translation. The parameters s , R and t are computed as the least squares fit
from four manually annotated points.

To establish correspondence between fibres, the detected fibre parameters
180 are transformed, i.e. their centre positions are rotated, translated and scaled,
and their diameters are scaled. Correspondence between fibres is then found
as the nearest neighbour using Euclidean distance. Examples of the segmented
fibres can be seen in Fig. 4.

3.2. Statistical analysis

185 We consider the fibre diameter measurements on $N = 757$ fibres by $n = 10$
different quantification pipelines (methods), listed in Tab. 2. Each method
consists of two steps: i) scanning technique (mode), which refers to a specific
imaging modality and resolution, and ii) image analysis method (algorithm).

The modes are CT scans including synchrotron, high, medium, and low resolution X-rays (called SRCT, XCT_H, XCT_M, XCT_L) and microscopy images, namely optical and scanning electron microscopy (called OM and SEM). The algorithms are denoted Prob. for the analysis based on the probabilistic feature labelling in 3D and Hough for the circular Hough transform on the top layer of the sample. We organise the measurements in two different ways: the univariate and the multivariate approach.

Table 2: Model parameters for statistical Model I and II. $D_{\nu k} = D_{ijk}$.

		Model I			Model II	
		Algorithm \times Mode \times Fibre			Method \times Fibre	
Abbrev.	Alg., $i = 1, 2$	i	j	D_{ijk}	ν	$D_{\nu k}$
SRCT	Prob.	1	1	$D_{(1,1,k)}$	1	$D_{(1,k)}$
XCT _H	Prob.	1	2	$D_{(1,2,k)}$	2	$D_{(2,k)}$
XCT _M	Prob.	1	3	$D_{(1,3,k)}$	3	$D_{(3,k)}$
XCT _L	Prob.	1	4	$D_{(1,4,k)}$	4	$D_{(4,k)}$
SRCT	Hough	2	1	$D_{(2,1,k)}$	5	$D_{(5,k)}$
XCT _H	Hough	2	2	$D_{(2,2,k)}$	6	$D_{(6,k)}$
XCT _M	Hough	2	3	$D_{(2,3,k)}$	7	$D_{(7,k)}$
XCT _L	Hough	2	4	$D_{(2,4,k)}$	8	$D_{(8,k)}$
OM	Hough	2	5	$D_{(2,5,k)}$	9	$D_{(9,k)}$
SEM	Hough	2	6	$D_{(2,6,k)}$	10	$D_{(10,k)}$

In the univariate approach we describe the variation between fibre diameters as deterministic contributions from the methods (or algorithm and modes) adjusting for the different fibre contributions.

If we want to distinguish between algorithms and modes we use a three-way factorial design model, symbolically written Algorithm \times Mode \times Fibre'. If we just focus upon the different methods we use a two-way layout written Method \times Fibre'. For the three-way layout (Model I) we consider the measured diameter

for fibre k , algorithm i and mode j ,

$$\begin{aligned} \text{Model I: } D_{ijk} &= \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \gamma_k + \epsilon_{ijk}, \\ &i = 1, 2, j = 1, \dots, 6 \text{ and } k = 1, \dots, N \end{aligned}$$

where μ is the general level, α_i is the deviation from the mean for algorithm i , β_j is the deviation from the mean for mode j , $(\alpha\beta)_{ij}$ is the interaction between
 205 algorithm and mode and γ_k is the k^{th} fibre's deviation from the mean. The error terms ϵ_{ijk} are assumed to be independent and normally distributed with the same variance, i.e. $\epsilon_{ijk} \sim N(0, \sigma^2)$. Note that the design is not balanced, i.e. we do not have measurements for all combinations of algorithm and mode
 210 since the 3D-based data analysis is not possible for the surface-based modes OM and SEM.

For the two-way model (Model II) we describe the measured diameters as

$$\text{Model II: } D_{\nu k} = \mu + \theta_\nu + \gamma_k + \epsilon_{\nu k}, \quad \nu = 1, \dots, 10 \text{ and } k = 1, \dots, N,$$

with similar interpretations and assumptions as in Model I.

The possible dependence on the spatial layout may be investigated by considering regression on spatial coordinates, i.e. we fit polynomial models to the
 215 size of the diameter as a function of the centre coordinates (x_k, y_k) for the k^{th} fibre. For the methods $\nu = 1, \dots, n$ we consider the models:

$$\text{A: } D_{\nu k} = \gamma(\nu) + \delta(\nu)x_k + \rho(\nu)y_k + \eta(\nu)x_k^2 + \kappa(\nu)y_k^2 + \lambda(\nu)x_k y_k + \epsilon_{\nu k},$$

$$\text{B: } D_{\nu k} = \gamma(\nu) + \delta x_k + \rho y_k + \eta x_k^2 + \kappa y_k^2 + \lambda x_k y_k + \epsilon_{\nu k},$$

$$\text{C: } D_{\nu k} = \gamma + \delta x_k + \rho y_k + \eta x_k^2 + \kappa y_k^2 + \lambda x_k y_k + \epsilon_{\nu k},$$

where the errors are independent and normally distributed $\epsilon_{\nu k} \sim N(0, \sigma^2)$. The first model corresponds to individual, quadratic surfaces, one for each method.
 220 In the second case we have the same coefficients to the spatial coordinates for all methods, only the intercepts $\gamma(\nu)$ are method dependent. In the last case we use the same surface for all methods.

The unknown parameters in the models above are estimated using ordinary least squares, and hypothesis on possible simplifications of the models are done by means of suitable Analyses of Variance (ANOVA). For a detailed description see for instance [19].

In the multivariate approach, we organise the observations in $N = 757$ ten-dimensional vectors

$$D_k = [D_{1,k}, D_{2,k}, D_{3,k}, \dots, D_{10,k}]^T, \quad k = 1, \dots, N .$$

Under the normality assumption, the multivariate distribution of the D_k 's is determined by the means and standard deviations of the coordinates (in this case measured diameters) in the observation vector defined above, and the correlations between those. The latter will be collected in the correlation matrix

$$Q = \begin{bmatrix} 1 & \cdots & \rho_{1,10} \\ \vdots & \ddots & \vdots \\ \rho_{1,10} & \cdots & 1 \end{bmatrix},$$

where $\rho_{r,s}$, $r = 1, \dots, 10$, $s = 1, \dots, 10$ is the correlation between the r 'th and the s 'th component in D_k . If we project D_k on the eigenvectors of Q we obtain the so-called principal components (PC) of D_k . They are uncorrelated and the first component describes the maximal variation amongst all linear combinations (with normed coefficients) of D -components. The following components maximise the variation subject to the constraint that they are uncorrelated with the previous. Such a Principal Component Analysis (PCA) [20] is thus a decomposition of the often very correlated coordinates in the observation vector into independent components describing decreasing amounts of the total variation.

4. Results and discussion

The field of view varies between scans, which constraints the amount of fibres that can be compared. Fig. 2 shows the fibres seen in all scans, and the overlap between these images defines the region of interest where we have carried out the analysis. Only the fibres in the lower part of the image have

been analysed (circled in *red* in Fig. 4(a)), because these are from a UD bundle that is orthogonal to the viewing direction. The upper fibres (not circled) are backing fibres that are at an angle of 45° relative to the UD part. This difference
 250 in orientation is easily seen in the upper part of the rendering in Fig. 3.

Examples of the fibre segmentation are shown for the different scans in Fig. 4. All scans have been analysed using circular Hough Transform, and the micro-CT modalities have in addition been analysed using our probabilistic feature labelling. This allows us to analyse the effect of scanning technique (mode) and
 255 the effect of the image analysis method for fibre segmentation (algorithm).

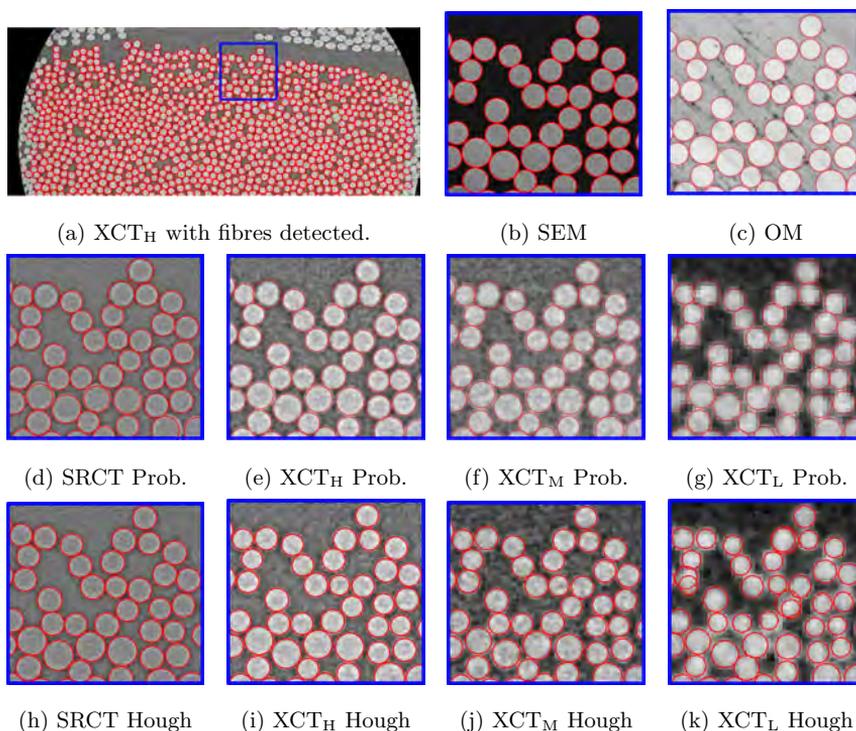


Figure 4: Diameter estimates over scans. In (a) the diameters for the high resolution XCT scan. In (b-k) zoomed versions of all scans showing the diameter estimates inside the blue square marked in (a) is shown. (k) shows some false positives using Hough Transform. Subscripts (H, M and L) are for high, mid, low resolution respectively. Prob. denotes probabilistic feature labelling and Hough denotes circular Hough Transform.

Fig. 4 shows that most fibres are found quite precisely for all scans using both algorithms. The only exception is the low resolution XCT analysed with circular Hough Transform, where some non-existent fibres (false positives) are detected, as shown in Fig. 4(k). False positives appear because the sensitivity parameter
260 must be set to a high value to detect all fibres. When all true positives are detected, some false positives are also detected. To include only one detection per fibre, we utilise the alignment of the images. In other words, after aligning the low resolution XCT image, we select the fibres closest to the fibres found in the high resolution XCT. The removal of these false positives obtained with the
265 Hough transform would not affect the precision of correct fibre detections. Thus, we decided not to focus on this for the current paper, where we are evaluating the precision of the Hough transform algorithm for fibre detection, and comparing it to our probabilistic fibre segmentation algorithm. Nevertheless, it should be noted that the Hough transform algorithm would not be directly applicable
270 for fibre segmentation at low resolutions when reference detections (obtained through other methods) are not available.

4.1. Univariate approach

The evaluation of X-ray imaging for fibre characterisation is based on comparing the measured fibre parameters. In the region of interest, that was determined by the overlap between the different scans, 757 fibres were segmented for
275 all combinations of modes and algorithms.

Average diameters and deviations from the mean are given in Tab. 3. We have applied a two-way analysis of variance to test the difference in average diameters between the 10 different methods, resulting in a grouping of the measurements. Group **a**, **b** and **c** are very close to the product specification of
280 $17\mu\text{m}$, and the high-resolution SEM is in this group. OM has a comparable pixel-size to SEM, but the diameters are on average $0.50\mu\text{m}$ smaller. For the micro-CT scans, the diameter estimates decrease with lower resolution, and this is especially pronounced for segmentation obtained using circular Hough
285 transform. Here the low resolution is significantly different from all other esti-

Table 3: Multiple comparisons between the means ($\hat{\mu} + \hat{\theta}$) for 10 methods using Model II. Average deviation is given by $\hat{\theta}_\nu$, and superscripts denotes groups in which the average diameters are not significantly different based on a two-way analysis of variance over Model II. The overall mean is $\hat{\mu} = 16.57 \mu\text{m}$.

Method	Hough		Probabilistic	
	Mean: $\hat{\mu} + \hat{\theta}_\nu$	Deviation: $\hat{\theta}_\nu$	Mean: $\hat{\mu} + \hat{\theta}_\nu$	Deviation: $\hat{\theta}_\nu$
SEM	17.05 $\mu\text{m}^{\text{a,b}}$	0.48 μm	-	-
OM	16.55 μm^{d}	-0.027 μm	-	-
SRCT	16.86 μm^{c}	0.29 μm	17.15 μm^{a}	0.57 μm
XCT _H	16.32 μm^{e}	-0.25 μm	17.06 $\mu\text{m}^{\text{a,b}}$	0.48 μm
XCT _M	16.26 μm^{e}	-0.32 μm	16.96 $\mu\text{m}^{\text{b,c}}$	0.39 μm
XCT _L	14.96 μm^{f}	-1.62 μm	16.58 μm^{d}	0.0026 μm

mates (classified in group **f**). A smaller decrease is observed for segmentation obtained using the probabilistic feature labelling method, where only the low resolution is significantly different from SEM. Segmentation from probabilistic feature labelling also gives very consistent standard deviations independently of scan resolution. Additionally, Tab. 3 indicates that both mode and algorithm have an effect on the measured diameters.

Spatial distributions of diameters

To further validate the findings of our analysis, we fitted to our data the three quadratic spatial models A, B, and C presented in Section 3.2. The diameter estimates employed here were obtained with the Hough method for the OM and SEM modes and with the Prob. method for the four modes. Since least square estimates are sensitive to extreme outliers, we made an outlier test on the 6×757 fibre measurements using the studentised residuals from a simple two-sided ANOVA with fibres and methods as factors. Residuals with an absolute value larger than 3.5 were considered extreme and discarded from the analysis leaving 738 fibres for further statistical investigations. The simplification from model A

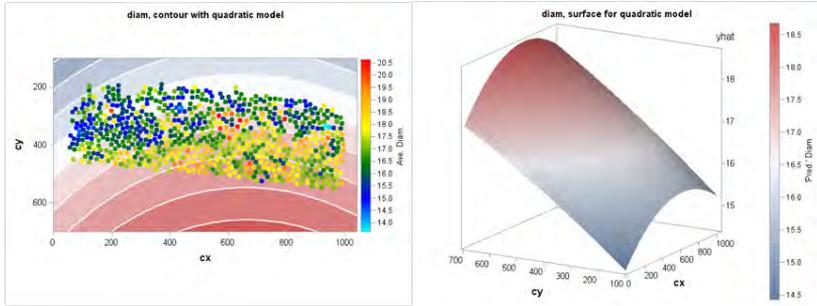
to model B is not statistically significant, i.e. we may assume that the coefficients to the linear and quadratic terms are independent of method. However, model C differs significantly from model B, i.e. we need individual intercept terms for the six methods, which correspond to the significant difference in the means for the methods depicted in Tab. 3. The fact that the spatial variation is the same for all methods indicates that the spatial dependence of the diameter values is not an estimation artefact, but very likely expresses a real phenomenon possibly due to inherent inhomogeneities in the fibre manufacturing process.

In Fig. 5(a) we show the spatial distribution of fibre diameters on the *left*, computed by averaging the estimates obtained through the six methods to which we fitted the spatial models (mentioned above), and the fitted model C on the *right*. Fig. 5(b) displays the diameter estimates for a complete UD bundle, computed with the method comprised by the low resolution XCT scan and the probabilistic algorithm. Note that there is an agreement between Fig. 5(a) *left* and Fig. 5(b), illustrating that it is possible to accurately measure the spatial distribution of fibre diameters for complete bundles. This computation requires under one minute and can, for example, be employed to obtain insights about the fibre bundle manufacturing process.

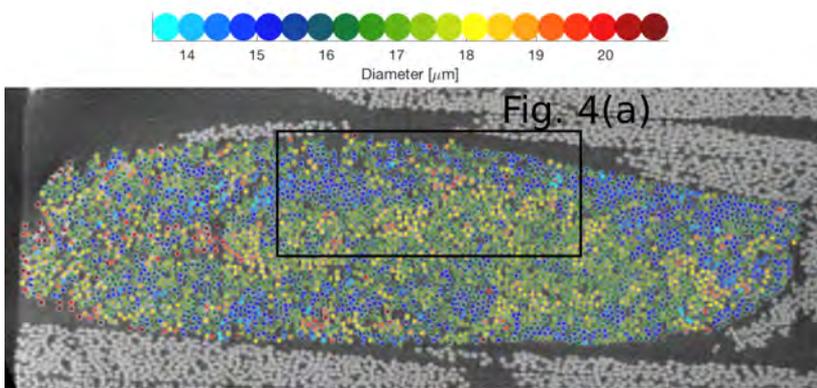
4.2. Multivariate approach

Fig. 6 shows scatter plots and diameter distributions of scans using the high resolution images (SEM, OM, SRCT and XCT_H). For SRCT and XCT_H both Hough transform and probabilistic feature labelling have been used. The plot is coloured according to three main groups, i.e. the SEM and OM (blue), XCT_H and SRCT using circular Hough transform (green) and probabilistic feature labelling (red). The interactions between groups are shown in other colours. It can be seen that OM has a general bias towards smaller diameters and that there is a good correspondence between segmentations obtained using the same algorithm whereas there is a larger spread between measurements obtained with different algorithms.

The coefficients for computing the principal components in Tab. 4 show that



(a)



(b)

Figure 5: Spatial distribution of diameter values. In (a) measured and predicted diameters inside the RoI. On the *left*, individual fibres are coloured according to the average of the diameter estimates across 6 methods. Contour lines (*left*) and surface (*right*) for the quadratic fit in model C (based on $N = 738$ fibres) give the predicted diameter values. In (b) diameters measured for the whole UD bundle from the low resolution XCT modality with our Prob. algorithm.

the first principal component is (approximately) proportional to the average of the values from all methods using high resolution images. The second principal component is a contrast (difference) between diameter values obtained by surface-based methods (Hough) and 3D methods (Prob.).

335

A comparison of the micro-CT scans and the two algorithms is shown in Fig. 7. The plot is divided into a part that is blue showing segmentation results

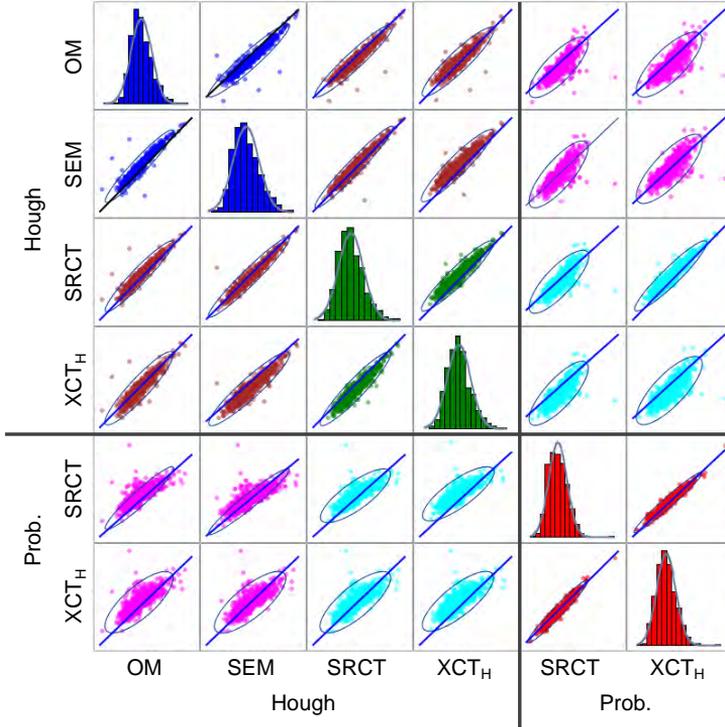


Figure 6: Scatter plot for diameters obtained for the high resolution images. Methods employed row-wise are OM, SEM, SRCT, and XCT_H , first with the Hough transform and then with probabilistic feature labelling. The blue line is the identity line and the blue ellipse is the 99% prediction ellipse for a new observation.

Table 4: Principal components based on the covariance of high resolution scans illustrated in Fig. 6.

Alg.	Hough				Prob.	
Mode	OM	SEM	SRCT	XCT_H	SRCT	XCT_H
PC1	0.4139	0.4178	0.4186	0.4130	0.3955	0.3898
PC2	-0.2977	-0.2827	-0.2825	-0.2461	0.5649	0.6099

obtained with probabilistic feature labelling, red shows circular Hough transform, and magenta is the interactions. Again, it is evident that the spread within group is smaller than between group, and especially the low resolution scan analysed with the circular Hough transform deviates very much from the other scans.

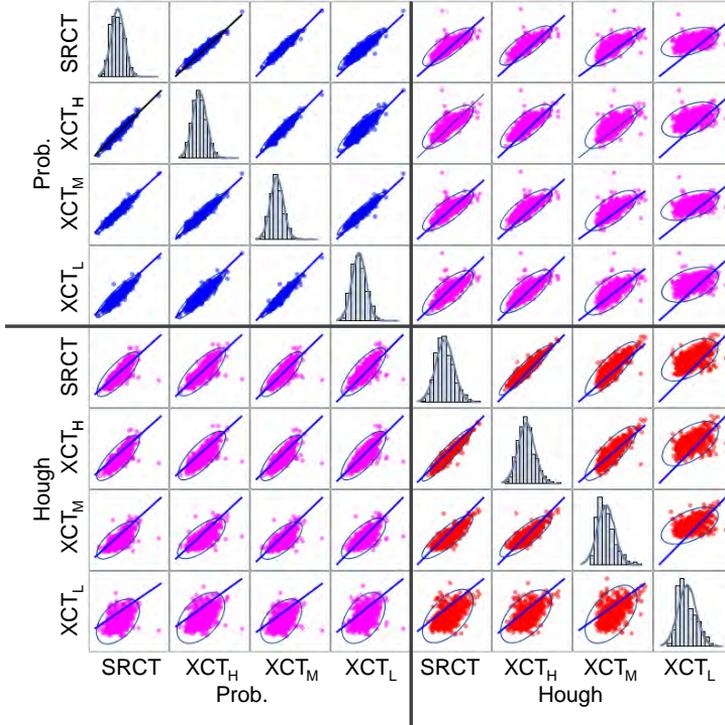


Figure 7: Scatter plot for diameters obtained for with micro-CT imaging. Methods employed row-wise are SRCT, and high, mid, and low resolution XCT, first with Hough transform and then with probabilistic feature labelling. The blue line is the identity line and the blue ellipse is the 99% prediction ellipse for a new observation.

The principal component coefficients in Tab. 5 show again that the first principal component is approximately proportional to the mean of all methods using synchrotron and micro-CT and that the second principal component captures the difference between the algorithms. Furthermore, it can be seen

that the coefficient for the Hough algorithm over XCT_L deviates considerably from the others. This corresponds to the rather low correlation between values obtained with the Hough transform and with the remaining methods as
 350 mentioned when discussing Fig. 7. A reasonable conclusion drawn from this is that we have reached the resolution limit for obtaining reliable results with the Hough transform algorithm.

Table 5: Principal components based on the covariance of micro-CT scans illustrated in Fig. 7.

Alg.	Prob.				Hough			
Mode	SRCT	XCT_H	XCT_M	XCT_L	SRCT	XCT_H	XCT_M	XCT_L
PC1	0.3777	0.3742	0.3777	0.3686	0.3650	0.3647	0.3352	0.2452
PC2	-0.2589	-0.2642	-0.2825	-0.2999	0.1932	0.1910	0.2777	0.7365

5. Conclusion

Our investigation of fibre characterisation using X-ray CT demonstrates that
 355 this scanning technique has similar precision as scanning electron microscopy (SEM) and optical microscopy (OM) for measuring fibre diameters and centre positions. We also showed that high precision can be obtained even if the resolution is lowered significantly, and we obtained results similar to SEM with X-ray CT images at nine times lower resolution and only slight degradation
 360 in performance at fifteen times lower resolution. However, this performance requires high precision methods for fibre segmentation, such as our probabilistic feature labelling method. This illustrates that image-based characterisation should consider the imaging modality in combination with the image analysis method. In addition to high precision, X-ray CT combined with the proposed
 365 segmentation method, opens the possibility for high precision analysis of large volumes, which has a great potential in giving insights for understanding the relation between fibre micro-structure and material properties.

References

- [1] R. Pyrz, Application of X-ray microtomography to the study of polymer
370 composites, Tech. rep., SAE Technical Paper (1999).
- [2] G. Requena, G. Fiedler, B. Seiser, P. Degischer, M. Di Michiel, T. Buslaps,
3D-quantification of the distribution of continuous fibres in unidirectionally
reinforced composites, *Composites Part A: Applied Science and Manufac-*
turing 40 (2) (2009) 152–163.
- [3] J. E. Little, X. Yuan, M. I. Jones, Characterisation of voids in fibre rein-
375 forced composite materials, *NDT & E International* 46 (2012) 122–127.
- [4] M. J. Emerson, K. M. Jespersen, P. S. Jørgensen, R. Larsen, A. B. Dahl,
Dictionary based segmentation in volumes, in: *Scandinavian Conference*
on Image Analysis, Springer, 2015, pp. 504–515.
- [5] A. Reh, B. Plank, J. Kastner, E. Gröller, C. Heinzl, Porosity maps–
380 interactive exploration and visual analysis of porosity in carbon fiber re-
inforced polymers, in: *Computer Graphics Forum*, Vol. 31, Wiley Online
Library, 2012, pp. 1185–1194.
- [6] K. K. Kratmann, M. Sutcliffe, L. Lilleheden, R. Pyrz, O. T. Thomsen, A
385 novel image analysis procedure for measuring fibre misalignment in uni-
directional fibre composites, *Composites Science and Technology* 69 (2)
(2009) 228–238.
- [7] Z. Yang, W. Ren, M. Mostafavi, S. A. McDonald, T. J. Marrow, Character-
isation of 3D fracture evolution in concrete using in-situ X-ray computed
390 tomography testing and digital volume correlation, in: *VIII International*
Conference on Fracture Mechanics of Concrete and Concrete Structures,
2013, pp. 1–7.
- [8] M. J. Emerson, K. M. Jespersen, A. B. Dahl, K. Conradsen, L. P.
Mikkelsen, Individual fibre segmentation from 3D X-ray computed tomog-

- 395 raphy for characterising the fibre orientation in unidirectional composite
materials, *Composites Part A: Applied Science and Manufacturing*.
- [9] D. Tsarouchas, A. Markaki, Extraction of fibre network architecture by
X-ray tomography and prediction of elastic properties using an affine ana-
lytical model, *Acta Materialia* 59 (18) (2011) 6989–7002.
- 400 [10] D. Bull, L. Helfen, I. Sinclair, S. Spearing, T. Baumbach, A comparison
of multi-scale 3D X-ray tomographic inspection techniques for assessing
carbon fibre composite impact damage, *Composites Science and Technology*
75 (2013) 55–61.
- [11] J. Kastner, B. Plank, D. Salaberger, J. Sekelja, Defect and porosity deter-
405 mination of fibre reinforced polymers by X-ray computed tomography, in:
2nd International Symposium on NDT in Aerospace, 2010, pp. 1–12.
- [12] Y. Nikishkov, L. Airoidi, A. Makeev, Measurement of voids in composites
by X-ray computed tomography, *Composites Science and Technology* 89
(2013) 89–97.
- 410 [13] M. W. Czabaj, M. L. Riccio, W. W. Whitacre, Numerical reconstruction
of graphite/epoxy composite microstructure based on sub-micron resolu-
tion X-ray computed tomography, *Composites Science and Technology* 105
(2014) 174–182.
- [14] R. Sencu, Z. Yang, Y. Wang, P. Withers, C. Rau, A. Parson, C. Soutis,
415 Generation of micro-scale finite element models from synchrotron X-ray CT
images for multidirectional carbon fibre reinforced composites, *Composites
Part A: Applied Science and Manufacturing* 91 (2016) 85–95.
- [15] K. M. Jespersen, L. P. Mikkelsen, Three dimensional fatigue damage evolu-
tion in non-crimp glass fibre fabric based composites used for wind turbine
420 blades, *Composites Science and Technology*.
- [16] T. J. Atherton, D. J. Kerbyson, Size invariant circle detection, *Image and
Vision computing* 17 (11) (1999) 795–803.

- [17] H. Yuen, J. Princen, J. Illingworth, J. Kittler, Comparative study of hough transform methods for circle finding, *Image and vision computing* 8 (1) (1990) 71–77.
- 425
- [18] V. A. Dahl, C. H. Trinderup, M. J. Emerson, A. B. Dahl, Content-based propagation of user markings for interactive segmentation of patterned images, In submission.
- [19] D. C. Montgomery, *Design and analysis of experiments*, John Wiley & Sons, 2017.
- 430
- [20] R. A. Johnson, D. W. Wichern, et al., *Applied multivariate statistical analysis*, Vol. 4, 2014.

Evolution in Fibre Orientation during Axial Compression of a Composite through Time-lapse X-ray Imaging and Fibre Tracking

In submission, *Scientific Reports*.

Evolution in Fibre Orientation during Axial Compression of a Composite through Time-lapse X-ray Imaging and Fibre Tracking

Monica Jane Emerson^{1,*}, Ying Wang², Philip John Withers², Knut Conradsen¹, Anders BJORHOLM DAHL¹, and Vedrana Andersen Dahl¹

¹Section for Image Analysis and Computer Graphics, Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kongens Lyngby, 2800, Denmark

²Henry Moseley X-ray Imaging Facility, School of Materials, The University of Manchester, Manchester, M13 9PL, United Kingdom

*monj@dtu.dk

ABSTRACT

Unidirectional fibre reinforced composites are widely employed for their high specific strength and stiffness. Analytical and numerical models that describe composite behaviour have been hard to verify due to a lack of experimental observation, particularly in 3D. The aim of this paper is to combine fast in-situ X-ray computed tomography (CT) with advanced image analysis to capture the changes in fibre orientation in 3D during uninterrupted progressive loading in compression of a UD glass fibre reinforced polymer (GFRP) composite. By analysing and establishing correspondence between a sequence of time-lapse X-ray CT images of the composite, we are able for the first time to follow each fibre and quantify the progressive deflection that takes place during axial compression in the steps leading up to fibre micro-buckling and kinking. Our results show fibre misaligned clusters of around 200 μm in diameter. Even at 25% of the failure load, fibres have started to tilt in approximately the direction of the ultimate kink-band. The rate of tilting increases as the load approaches collapse. More generally our fibre tracking approach can be applied to investigate the behaviour of a wide range of fibrous materials under progressive loading conditions.

Introduction

Composite materials are used extensively in advanced structures where properties such as high stiffness and low weight are required. In particular, carbon and glass fibre reinforced polymers (CFRP and GFRP) are being employed increasingly in aero¹ and ground transportation, as well as in environmentally sustainable energy production systems such as wind turbines². Therefore, it is crucial to understand their performance under realistic loading conditions with the ultimate goal of optimising performance through the use of analytical or numerical models³. Currently, the lack of confidence in these models means that expensive experimental mechanical testing is relied upon in the development of new composite structures, and is still the gold standard for quality control in industry⁴. What is more, this lack of confidence in the prediction of composite behaviour, together with high safety requirements, leads to an unnecessary over-engineering of the components.

The compressive strengths of unidirectional (UD) composites are typically significantly below their tensile strengths. This is because many of the mechanisms of compressive failure are dictated largely by the matrix properties, in contrast to the tensile properties which are dictated by the fibres. Fleck⁵ describes a number of failure mechanisms, the most important of which for the current study is fibre micro-buckling (see Fig. 1b-c). The former is a shear buckling instability where the matrix deforms in simple shear and the latter instability occurs at sufficiently large strains for the matrix to deform in a non-linear manner. Longitudinal splitting along the fibre direction is also a potential damage mode (see Fig. 1d). Fibre micro-buckling is highly sensitive to the initial fibre misalignment. As a result of the composite manufacturing process, the fibres are not generally perfectly aligned with the axial direction. This means that when compression is applied, some level of shear is induced in the region of the misaligned fibres. When the local shear stress exceeds the matrix shear yield strength, the matrix yields and fibres buckle. In other words, the compression strength of a UD composite relies on the ability of the matrix material to keep the fibres straight⁶. Fibres finally fracture when the strain on the compressive/tensile side of the fibre exceeds the fibre's compressive/tensile failure strain, and a kink-band is formed⁷. Although for the reasons discussed above many researchers believe that fibre misalignment plays a big role^{8,9}, and the degree of lateral constraint on the fibres from the matrix is important¹⁰, the initiation mechanism for fibre micro-buckling is still a subject of debate. This is largely because of the lack of direct evidence of the fibre movements in the moments leading up to micro-buckling and the point at which this becomes a

kink-band. The difficulty in capturing experimentally the precursors to failure is partly due to the small levels of misalignment required to trigger instability, and partly because of the sudden and catastrophic nature of the instability when the kink-bands form.

In recent years, optical and scanning electron microscopy have increased in image acquisition speeds. This has been exploited to study the development of kink-bands in composites loaded in-situ¹¹ by observing the surface of the composites. Guynn et al.¹¹ and more recently Moran et al.¹² have observed experimentally in 2D that fibres buckle before a kink-band is formed. However, the stage of fibre deflection prior to the onset of kink-band formation, which is expected to influence the kink-band geometry, has not yet been captured experimentally. The use of 2D characterisation techniques limits the development of accurate analytical and numerical models, as there is evidence that both micro-buckling and kink-band formation are 3D phenomena¹³. X-ray computed tomography (CT) enables 3D imaging of FRP specimens containing many fibres at a resolution where individual fibres can be resolved. Moreover, it is possible to scan a sample while it is loaded in-situ in compression so that the evolution in fibre orientation under load can be observed in 3D from zero load towards failure. Conventional tomography, even at synchrotron sources, is too slow to capture the onset of micro-buckling¹⁴. However, ultra-fast imaging ($\sim 500 - 10,000$ projections per second) is now possible on some beamlines which enables the events leading up to failure to be captured¹⁵.

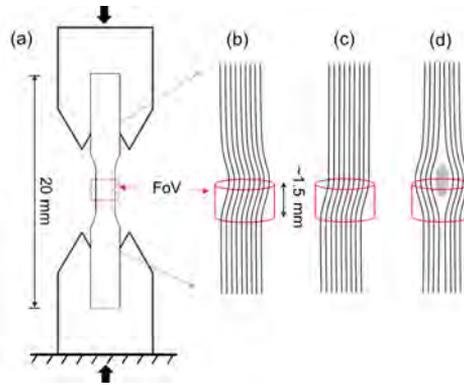


Figure 1. Schematic illustration of a loaded specimen and damage modes associated with axial compression. (a) Composite specimen geometry and loading condition, and potential damage modes associated with axial compression failure of unidirectional fibre composite, namely (b-c) fibre micro-buckling and (d) longitudinal splitting.

The focus of this paper is on combining in-situ X-ray CT with advanced image analysis with the purpose of providing an experimental description of the changes in fibre orientation during progressive loading in compression. Fibre orientations can be estimated from tomograms¹⁶⁻¹⁸ without previously segmenting individual fibres, but the accuracy of these estimates is highly dependent on the quality of the tomograms. However, if individual fibres can be segmented beforehand, the accuracy of the individual fibre trajectories can be ascertained prior to computing fibre orientations. Additionally, segmentation of individual fibres opens up the possibility of following the changes in trajectory for each individual fibre across successive loading steps. Methods that can segment densely packed individual fibres from tomograms have been developed and applied to extract fibre geometry for high quality data¹⁹ and also for lower quality scans²⁰⁻²². The method by Czabaj et al.²⁰ involves a time-consuming validation of the extracted fibre trajectories and is therefore not applicable to large volumes containing many fibres. Sencu et al.²² built a micro-mechanical model with some fibres extracted from bundles oriented in different directions whereas the focus of Emerson et al.²¹ was on determining individual fibre orientations.

Our proposed methodology is to characterise the trajectory of individual fibres in the as-manufactured condition in 3D from X-ray CT images and then to quantify, as compressive loading progresses, the deflections by the analysis of a time-lapse CT image sequence for the first time. The aim is to capture precursors to fibre micro-buckling in an end-tabbed UD glass fibre reinforced polymer (GFRP) rod specimen (see Fig 1a). As shown in Fig 1, the sample was scanned with a field of view (FoV) of 1.5 mm in height. To ensure that the material would fail within the FoV, a notch was introduced around the circumference of the reduced gauge section (diameter 1.5 mm) of the 20 mm long composite rod. The CT scanning was performed at the Swiss Light Source (SLS) while progressively loading the specimen in-situ in the axial direction from 0 N to failure (just before 900 N in this experiment).

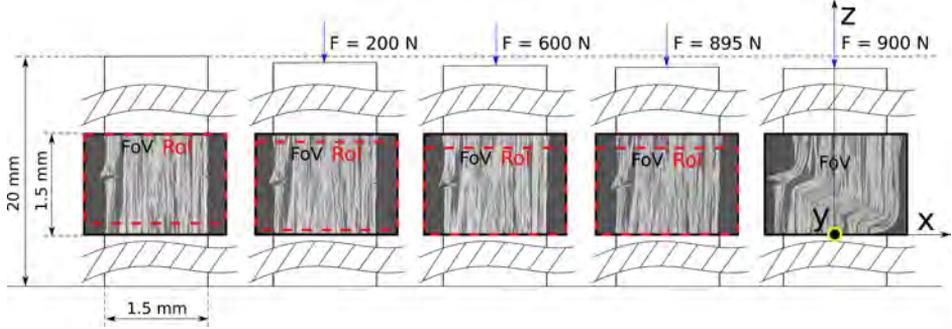


Figure 2. X-ray CT longitudinal slices from the centre of the composite rod in the direction of the kink-band plane. The composite rod of length 20 mm and diameter of 1.5 mm at the reduced gauge section was scanned with a field of view (FoV) of 1.5 mm in height during uninterrupted in-situ axial compression loading from 0 N to 900 N. Images corresponding to four loads are selected for the analysis: i) 0 N, ii) 200 N, iii) 600 N and iv) 895 N; and the analysed region of interest (RoI) has a height of approximately 1.4 mm. The coordinate system is defined according to the loading direction (z -axis, inclination $\phi = 0$) and the kink-band plane (x -axis, azimuth $\theta = 0$).

The evolution of individual fibre orientations is analysed in 3D at instances corresponding to four loads for the region of interest (RoI) shown in Fig. 2. The height of the RoI is defined by the region within the FoV that is common to all data-sets. As shown in Fig. 2, the overall length of the RoI along the fibre direction is 1.4 mm, which is short compared to the 20 mm long composite rod. As illustrated in Fig. 1, the small size of the FoV will mean that we can capture only a partial view of the fibre deflections along the sample height. The method by Emerson²¹ is employed here to track the individual fibres because of its ability to segment individual fibres in large volumes of densely packed fibres and its robustness to noise. The algorithm, presented in the Section "Methods", is able to delineate the fibre trajectories within the tomographic volumes and follow the evolution of each individual fibre from tomogram to tomogram as the load is increased. The fibre trajectories are shown in Fig. 3a-b for the as-manufactured condition and the evolution of a few randomly selected fibres is shown in Fig. 3c-d. It is evident from Fig. 3a-b that over the 1.5 mm high FoV the fibres are predominantly straight. Once individual fibre trajectories are obtained, individual fibre orientations are computed and individual fibre deflections can be calculated at every imaged loading step as the change in fibre orientation with respect to the as-manufactured fibre orientation.

In order to quantify the fibre trajectories and deflections during loading a series of vectors is employed. These are defined mathematically in the Section "Methods" and described pictorially in Fig. 4a. In essence, each slightly curved fibre trajectory can be approximated by a sequence of straight lines computed for different horizontal segments. Since the curvature of the fibres is small within the RoI, for our current analysis we approximate each fibre by a straight line through the RoI. The **fibre trajectory** for the n -th fibre under the i -th load state is captured by the **fibre vector** $\mathbf{v}_{i,n}^f$, which joins the points where the fibre enters and leaves the RoI. The corresponding vector for the composite is the average of all the fibre vectors and is denoted by the **composite vector** \mathbf{v}_i^c . The **fibre deflection** for the n -th fibre considers only the lateral deflection of the fibre end relative to its orientation in the 0 N load case, and is represented by the **fibre deflection vector** $\mathbf{d}_{i,n}^f$. The **composite deflection vector** \mathbf{d}_i^c captures the deflection of the composite rod given by the mean of all the fibre deflection vectors. Finally there is the **relative deflection**, which expresses the movement of the fibre relative to the deflection of the composite rod as a whole and is represented by the **relative deflection vector** $\mathbf{r}_{i,n}^f$. Since fibres are not free to move independently of one another, an additional important composite movement is given by the twist γ_c , as observed by Ueda et al.²³. The twist is proportional to the torque induced by the compression loading and involves the fibres rotating collectively around the central axis of the composite rod. For ease of discussion, the vectors are described using either Cartesian or spherical coordinates where i) the z -axis is aligned with the loading z -axis (inclination $\phi = 0$ and points in the direction opposite to the loading direction), and ii) the x -axis (azimuth $\theta = 0$) is aligned with the kink-band plane, pointing towards the shallowest part of the notch. For clarification, the coordinate system is defined in Fig. 2. In the following section these vectors for describing the orientation and deflection of both the individual fibres and the composite as a whole are used to examine the behaviour of the UD composite through uninterrupted compressive loading in the moments leading up to kink-band formation.

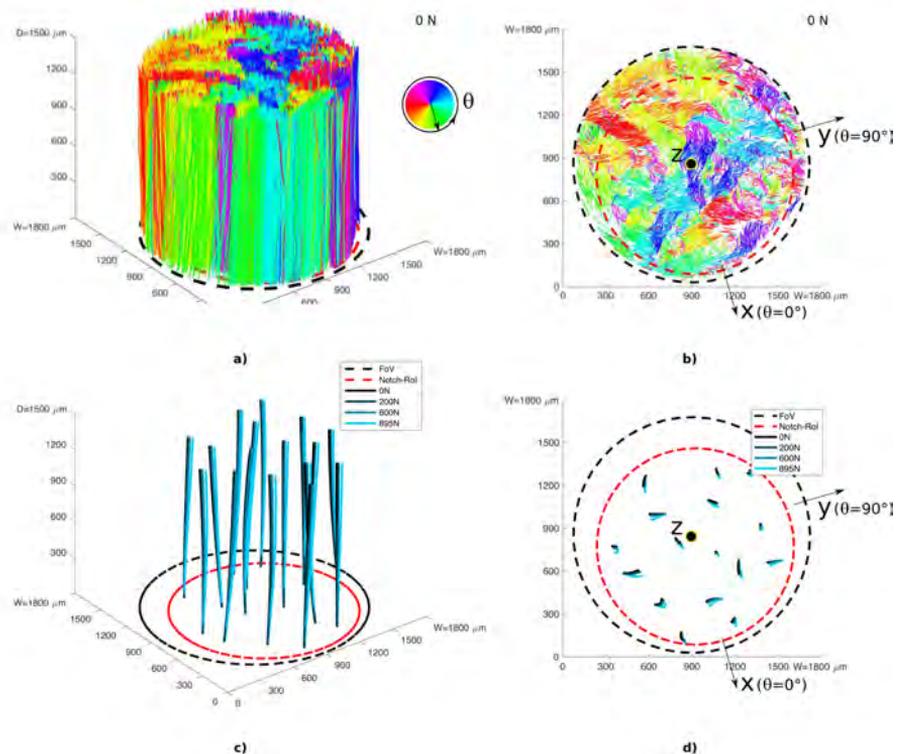


Figure 3. The as-manufactured fibre trajectories and their evolution under load. In the **top row**, a 3D side view (a) and plan view (b) of all the fibre trajectories extracted from the RoI for the as-manufactured test piece prior to the application of load, where the colour of each fibre encodes the azimuth θ of its fibre vector (defined in Fig. 4). The projected length in the plan view (b) shows the amount of misalignment, fibres appear as points when they are completely aligned with the z-axis. In the **bottom row** a side (c) and plan (d) view showing the evolution under compression of randomly selected fibre trajectories inside the RoI. The coordinate system, defined according to the loading direction and the kink-band plane is indicated in (b) and (d).

Results and discussion

Characterisation of fibre orientations before and under load

It has been widely accepted that fibre misalignment in the as-manufactured composite has a significant effect on the compressive failure load at which catastrophic kink-band forms⁵. Fig. 3a-b shows all the fibre trajectories for the as-manufactured data-set (0 N load). It is evident from the figure that the fibres are predominantly straight and that the orientation of the fibres is to a large extent grouped into clusters.

The spatial clustering of fibre orientations as a function of loading is shown in Fig. 5. The scatter plots (*top*) show the individual fibre vectors projected onto the horizontal plane where each point corresponds to the end of a fibre vector. The colour-coding represents the in-plane misorientation direction (represented by θ in Fig. 4) with a saturation that decreases with fibre alignment, i.e. with decreasing tilt (represented by ϕ in Fig. 4). When looking at the scatter plots one might underestimate

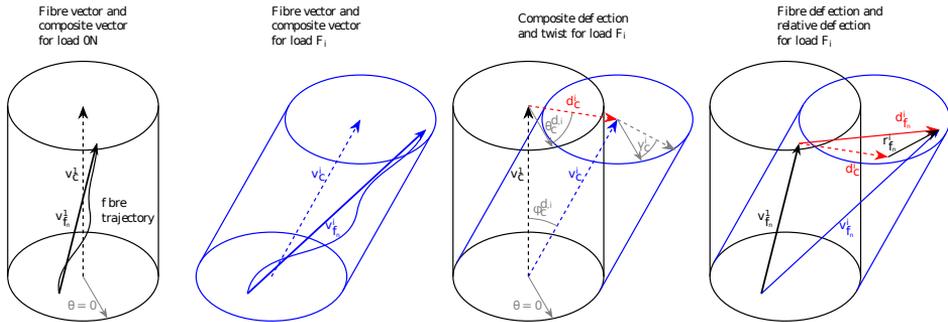


Figure 4. Definition of fibre and composite vectors, deflection vectors and relative deflection vectors.

the number of highly aligned fibres, as the scatter plots do not show the density of points when these are closely clustered together. That is, there are many more aligned fibres than it seems when looking at the scatter plots.

With regards to the fibre orientations displayed in Fig. 5, it can be concluded that both the degree of fibre misalignment relative to the loading axis (fibre tilt) and the direction of the misalignment (fibre direction) cluster spatially. It is noteworthy that the clustering of the fibre direction is correlated with the local fibre volume fraction. Where fibres are very close together they tend to be misaligned in the same direction whereas there tends to be a resin-rich region in between the differently oriented clusters of fibres. This is a reflection of the packing efficiency which drops when fibres are not similarly aligned. It is notable that while there is a gradual drift of the composite (as indicated by the *black* crosses in Fig. 5a) there are no significant changes in the clustering of fibre misalignments as the load is increased. This confirms that the fibres deform in a cooperative way in FRPs with relatively high fibre volume fraction. As to the evolution of the fibre vectors under progressive loading, a gradual deflection in the direction of the ultimate kink-band plane is observed in the scatter plots as the loading increases. This shift translates to an increase of *green* colour in the cross-sectional images under increasing load.

Evolution of fibre deflections under load

The spatial distribution of the deflection introduced by loading the sample in compression is analysed by looking at the deflection vectors and the relative deflection vectors.

As was mentioned previously, the deflection vectors capture the change in orientation that each individual fibre undergoes with respect to the 0N state as a consequence of applying a certain load F_i . The deflection vectors in the horizontal plane are shown in the scatter plots in Fig. 6a. They show that even at the 200N most of the fibres have already deflected towards the direction defined by the kink-band plane, although the magnitude of the deflection is still small.

The relative deflection vectors capture the deflection of individual fibres with respect to the mean fibre deflection, namely composite deflection \mathbf{d}_c^i (indicated with a black cross over the scatter plots in Fig. 6a). The deflection vectors in these scatter plots have been colour-coded according to the composite deflection. Thus, the cross-sectional slices in Fig. 6b show the spatial distribution of the relative deflection vectors. By looking at the spatial distribution of the relative deflection vectors it can be concluded that the deflection behaviour is not influenced by the initial orientation, as there is no similar clustering. Instead, all the fibres in the composite rod primarily undergo a cooperative deflection, as indicated by the crosses in *black* in Fig. 6a, and a twist in an anti-clockwise direction with respect to the centre of the rod, indicated by the arrows in *yellow* in Fig. 6b. The twist picks up the mean angle by which the relative deflection vectors turn around the centre of the rod.

Evolution of composite deflection under load

The histograms for the tilt of the fibre vectors are reported for all loads in Fig. 7a. Note that the spread of fibre inclinations increases somewhat as the load increases. The histograms for the direction of the fibre vectors are reported for all loads in Fig. 7b, it can be seen that fibres are not uniformly distributed around the circumference. Preferential directions exist and it is noticeable that there is a deficit of fibres at $-180/180^\circ$, this might be the reason for the kink-band forming in this direction. Additionally, the (all fibre averaged) composite behaviour is shown as a function of loading in terms of the composite deflection (inclination of the composite segment in the FoV (Fig. 7c)) and the twist of the composite rod about its central axis (Fig. 7d). The twist has been estimated by fitting a rotation to the relative deflection vectors and minimising the residual in the least squares sense²⁴. Unsurprisingly, the angle of inclination of the rod segment moves off axis faster as the load increases, but it is

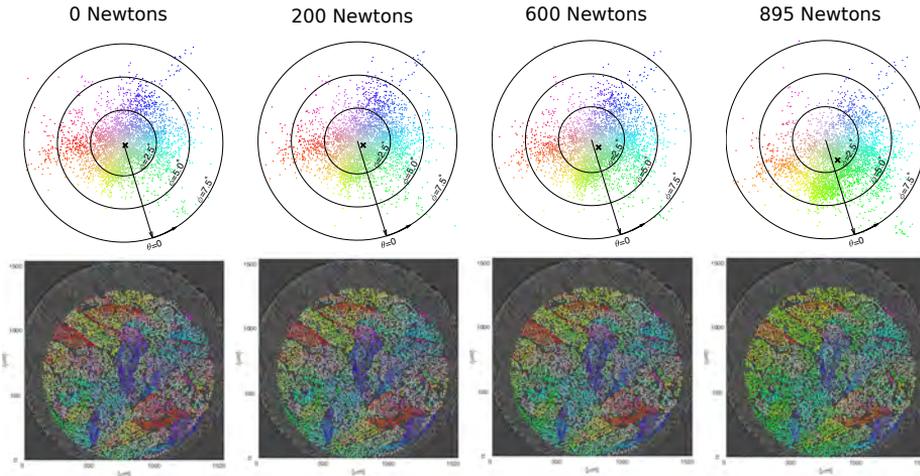


Figure 5. Spatial clustering of fibre orientations under load. The **top row** shows scatter plots of the projection of fibre vectors onto the horizontal plane for the four considered loading steps. Each point corresponds to a certain fibre, f_n , and is coloured according to: i) inclination ϕ (colour saturation) and ii) azimuth θ (colour hue). The **black crosses** indicate the composite orientation (average over all fibres). In the **bottom row** the X-ray CT cross-section of the fibres for the middle slice of the RoI is shown with the fibre vectors displayed at their actual spatial locations.

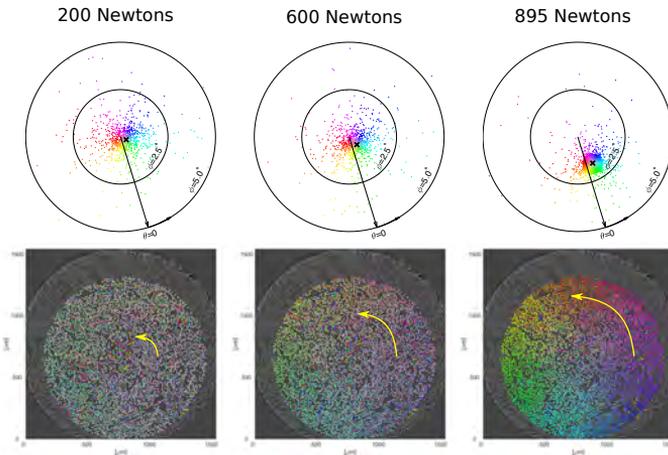


Figure 6. Evolution of deflection vectors and spatial distribution of relative deflection vectors under load. In the **top row** the scatter plots of the deflection vectors in the horizontal plane, $d_{f_n}^d$. The colour-coding has been set relative to the deflection of the composite rod, indicated by the **black crosses**. Thus, it is the relative deflection vectors $r_{f_n}^d$ that are shown over the X-ray CT cross-section (middle slice of the RoI) in the **bottom row**. The **yellow arrows** indicate the twist observed around the centre of the sample, note that these do not represent the absolute magnitude of the twist, which is 1.06° at 895 N.

perhaps surprising that the fibres deflect in the direction of the final kink-band (see Fig. 6 top) even at 25% of the kink-band collapse load.

The twist angle increases together with axial compression load, and amounts to 1.06° just prior to failure. On the one hand, this may arise because of a slight bias in the misalignment direction of the fibres in the as-manufactured test-piece. In histogram 5b fibres have a preferred alignment direction, as opposed to a random alignment direction which would give a uniform histogram for the azimuth. On the other hand, the twist could arise from non-ideal loading which could introduce some torque onto the specimen. While a rotational moment was not intentionally applied, it is possible that imprecise alignment of top and bottom jaws introduced some degree of rotational movement. Moreover, as the sample was end-loaded without being clamped, there is freedom of movement in the xy plane. Regarding the effect of this twist on the damage evolution, it is expected that the twist will result in some degree of fibre deflection out of the dominant kink-band plane, such that the collapse mechanism has a 3D nature.

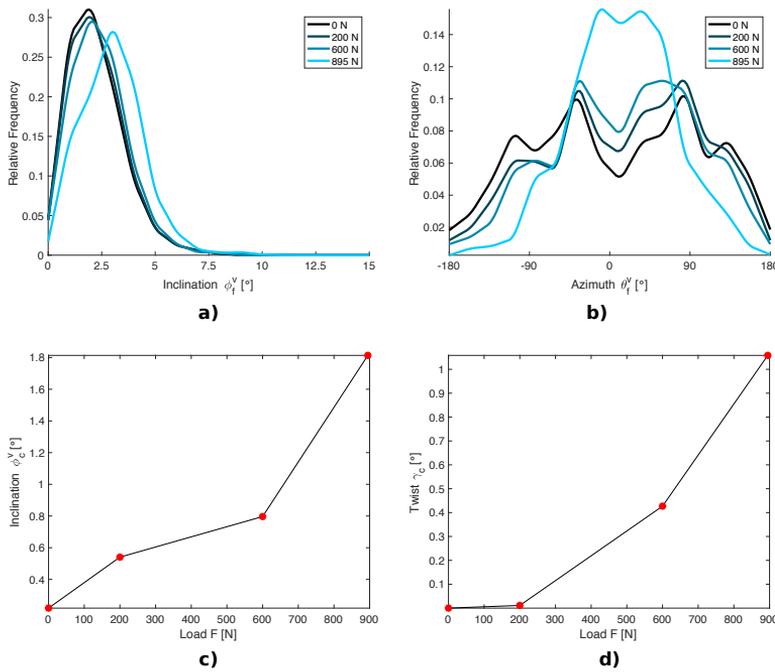


Figure 7. Progression of histograms capturing the magnitude and direction of fibre misalignment for the fibre vectors and the trend of the composite behaviour under load. In (a) and (b) the histograms for the inclination and azimuth of the fibre vectors. In (c) and (d) the trends for the composite tilt and twist.

Concluding remarks

In this study the 3D orientations of individual fibres under load have been measured and tracked experimentally for the first time by means of time-lapse X-ray CT and statistical image analysis. In this way we have been able to quantify the changes in fibre orientation to a very high level of accuracy, tracking the great majority of fibres in the composite. This approach has allowed to better visualise and understand precursors to fibre micro-buckling in a GFRP under compression loading. The main findings are considered below.

It is a good question to ask whether the ultimate kink-band location and failure direction is predestined from the start. This paper reports only the average orientations of individual fibres along the RoI and does not set out to measure the change in the curvature of the fibres. This will be the subject of a further paper. It is clear from Fig. 3a that the gauge section contains a number of misaligned clusters of around 200 μm in diameter. It is worthwhile to consider whether the gauge is large enough to be statistically representative of the composite as a whole. It seems that the test-piece contains around 30 – 40 clusters and so it is probably just large enough for it to be unlikely that a particular cluster would dominate the compressive failure behaviour. Nevertheless, it is useful to ask whether the average misalignment angle/direction of all the fibres affects the initiation of the micro-buckling behaviour. Certainly it was remarked upon earlier that there appears to be few fibres oriented at 180° to the kink direction (Fig. 7b) and this may help explain why the fibre tilting and eventual micro-buckling occurs in the direction it does. When considering the relative deflection vectors in Fig. 6 we can infer that there is no tendency for longitudinal splitting between fibres (no neighbouring clusters move in opposite directions laterally, as illustrated in Fig. 1d). Without information about the local curvature we are not able to identify whether the fibres buckle solely as depicted in Fig. 1c or whether there is an increase in the local curvatures as in Fig. 1b. Certainly our work shows for the first time that the deflection initiates at very low loads and is in a direction close to the ultimate kink-band direction. Furthermore, this deflection increase accelerates as load increases such that the composite segment in the RoI is inclined by 1.81° at 895 N (just prior to buckling instability). It is also clear that despite aiming to impose solely axial compression, there is a tendency for the composite to twist under increasing load, introducing another loading mode onto the composite. If the shear modulus of the composite is around 3 GPa (adapted from Hull and Clyne²⁵) the 1.06° of twist equates to a torque of around 0.02 N m. Ultimately the notch does ensure that the final failure sequence occurs local to it but the current fibre tracking method does not pick up any local reorientation of the fibres in the periphery of the gauge volume associated with the notch because it approximates fibres in the RoI with straight lines. Of course future measurements of the fibre curvature may provide indication of notch localised fibre curvature in the gauge volume. The quantification of fibre curvatures will be the topic of a follow up study. Finally, the presented methodology, which can follow changes in individual fibres under load, is applicable to a wide range of materials based on UD fibres.

Methods

Materials and experimental procedures

Manufacturing and preparation of composite sample

A modified resin infusion technique referred to as small-scale resin infusion, as reported by Wang et al.²⁶, was used to fabricate the 2 mm diameter UD E-glass fibre (12 μm in diameter)/epoxy (Huntsman Araldite LY564/XB3486) composite rod. Fig. 1a shows the specimen geometry for in-situ testing. Specimens with a reduced gauge section of around 1.5 mm in diameter over 3 mm in length were prepared from the manufactured composite rods. As mentioned in the Section "Introduction", in order to localise the damage site, a groove around the circumference was made using a razor blade giving a notch depth between 100 μm and 200 μm . The two ends of the specimen were then glued with an epoxy adhesive into chamfered steel end caps so that end-splitting damage could be avoided. This resulted in a length of the composite rod equal to 20 mm and an overall length of the specimen of ~ 28 mm.

Ultra-fast synchrotron X-ray computed tomography (CT) imaging under in-situ compression

Ultra-fast imaging was performed at the TOMCAT beamline from the Swiss Light Source (SLS). A monochromatic beam with mean energy of 20 keV was used and the specimen-detector distance was set so as to aid micro-crack detection with a level of phase contrast. 500 projections were acquired over 180° rotation and the exposure time for each projection was 2 ms. These settings equate to a fast acquisition at 1 scan per second. The data-sets were reconstructed at TOMCAT using the Paganin algorithm²⁷. The voxel size of the reconstructed CT data-sets is $1.1 \mu\text{m}^3$.

Specimens were end-loaded in a tension-compression in-situ loading rig developed at INSA-Lyon²⁸, which could be accommodated for dynamic imaging. The compression tests were conducted under displacement control at the rate of $1 \mu\text{m s}^{-1}$. For the sample presented in this paper, an initial static scan was performed at 0 N, followed by interrupted scans at intermediate loads (200 N and 600 N) followed by continuous dynamic scans as the expected failure load was approached. Each interrupted CT scan started at an arbitrary angle, so there is an angular difference between each data-set in the xy plane.

Individual fibre extraction and matching across data-sets

Extraction of individual fibre trajectories

UD fibres are essentially aligned with the z -axis, so in each horizontal slice through the volume fibres appear as circles, as shown in Fig. 8. This allows the detection of fibre centres in each cross-sectional slice and these are connected together afterwards to form individual fibres. Each individual fibre centre line is called a 'trajectory', which is the curve that connects the centres of a fibre through the volume. These trajectories are extracted in two steps: i) centre coordinates are detected for

every fibre in each 2D tomographic cross-sectional slice, obtained by slicing the 3D volume along fibre cross-sections and ii) 3D fibre trajectories are determined by tracking, i.e. connecting the 2D detected centres that belong to the same fibre.

The process of centre coordinate detection is illustrated and explained in Fig. 8. First, a dictionary-based probabilistic segmentation is applied so as to estimate pixel-wise probabilities of belonging to the centre region. Secondly, the obtained probability map is thresholded to obtain the individual centre regions. Finally, centre coordinates are computed as the centres of mass of the connected regions. For more details refer to Emerson et al.²¹.

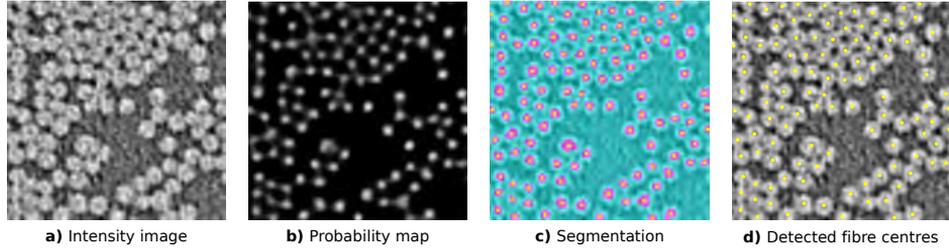


Figure 8. Steps in the 2D centre detection process. Fibre centre coordinates calculated using the method by Emerson et al.²¹. The intensity image in (a) is segmented into two classes (fibre centre regions or not). Using a dictionary-based probabilistic segmentation method the likelihood for each pixel belonging to the fibre centre class is obtained, shown in (b). This probability map is thresholded to obtain the segmentation in (c), which determines what are fibre centre regions (pink) or not (blue). The centre coordinates for the individual fibres in (d) are found as the centroids (yellow) of the pink connected regions.

A new implementation of the dictionary-based segmentation method is available²⁹. The main improvements are lower computational time for training the dictionary model as well as for the segmentation phase and, most importantly, the development of a graphical user interface (GUI) in Matlab[®] to wrap up the algorithm and provide a user-friendly tool that requires minimal user input for training the dictionary.

The process of fibre tracking for connecting the detected centres across the 2D slices is explained in Emerson et al.²¹. The fibre tracking used here is an improved version that can deal with the problems encountered in the previous version. This includes handling of i) double detections, by merging 2D fibre centre detections when these are closer than the manufacturer specified fibre radius, and handling of ii) merging fibres, by matching fibre centres from one slice to the next only if the match between the points is bidirectional, as opposed to utilising the unidirectional backward tracking approach in Emerson et al.²¹.

Matching corresponding fibre trajectories across loading steps

As mentioned above, there is an angular disregistry between each data-set in the xy plane and it needs to be corrected so that corresponding fibres can be matched across data-sets. Corresponding fibre trajectories are matched across the four microtomograms using a rigid registration for alignment, followed by a nearest neighbour search algorithm for matching corresponding fibres.

Tomograms are aligned manually, first in the z direction by visual inspection and then in the horizontal direction by choosing three corresponding points and calculating a rigid transformation (rotation plus translation). The corresponding points are selected as far as possible from each other and towards the edges of the cross-section to ensure precision in the registration.

Vectors of the Deflection Model

For the n -th fibre at loading step i the fibre vector $\mathbf{v}_{f_n}^i$ is calculated by joining the points where fibre f_n enters and leaves the RoI with $F_i \in \{0\text{N}, 200\text{N}, 600\text{N}, 895\text{N}\}$ and $n \in \{1, \dots, N\}$, where the number of fibres considered in this study is $N = 4957$ (see Fig. 4). The composite vector at load step i is computed by the mean of the fibre vectors at that step $\mathbf{v}_c^i = \frac{1}{N} \sum_{n=1}^N \mathbf{v}_{f_n}^i$.

Given the fibre vectors, the deflection vector $\mathbf{d}_{f_n}^i$ for fibre f_n at load step i is calculated as the difference between the fibre vector for f_n at load step i and the fibre vector for f_n at load $F_1 = 0\text{N}$, i.e. $\mathbf{d}_{f_n}^i = \mathbf{v}_{f_n}^i - \mathbf{v}_{f_n}^1$. The composite deflection vector at load step i is given by averaging the deflection vectors $\mathbf{d}_c^i = \frac{1}{N} \sum_{n=1}^N \mathbf{d}_{f_n}^i$. For the fibre and composite deflection vectors $F_i \in \{200\text{N}, 600\text{N}, 895\text{N}\}$.

The relative deflection vectors, which characterise how each fibre deflects compared to the mean fibre deflection (deflection of the composite rod), are given by $\mathbf{r}_{f_n}^i = \mathbf{d}_{f_n}^i - \mathbf{d}_c^i$.

Data availability

The data-sets analysed during the current study, the segmented fibre trajectories with correspondence across volumes and the code used for this study are available from the corresponding author on reasonable request.

References

1. Soutis, C. Carbon fiber reinforced plastics in aircraft construction. *Mater. Sci. Eng. A* **412**, 171–176 (2005).
2. Brøndsted, P., Lilholt, H. & Lystrup, A. Composite materials for wind power turbine blades. *Annu. Rev. Mater. Res.* **35**, 505–538 (2005).
3. Camanho, P. P. & Hallett, S. R. *Numerical modelling of failure in advanced composite materials* (Woodhead Publishing, 2015).
4. Sims, G. D. *Review of composites standardisation status and activities*. National Composites Network, UK (2007).
5. Fleck, N. Compressive failure of fiber composites. *Adv. Appl. Mech.* **33**, 43–118 (1997).
6. Budiansky, B. & Fleck, N. A. Compressive failure of fibre composites. *J. Mech. Phys. Solids* **41**, 183–211 (1993).
7. Dow, N. F. & Grunfest, I. J. Determination of most needed potentially possible improvements in materials for ballistic and space vehicles. *Gen. Electr. Company, Air Force Contract. AF* **4**, 269 (1960).
8. Soutis, C. Compressive strength of unidirectional composites: measurement and prediction. In *Composite Materials: Testing and Design, Thirteenth Volume* (ASTM International, 1997).
9. Wisnom, M. R. & Atkinson, J. Fibre waviness generation and measurement and its effect on compressive strength. *J. Reinf. Plast. Compos.* **19**, 96–110 (2000).
10. Pimenta, S., Gutkin, R., Pinho, S. & Robinson, P. A micromechanical model for kink-band formation: Part i—experimental study and numerical modelling. *Compos. Sci. Technol.* **69**, 948–955 (2009).
11. Guynn, E. G., Bradley, W. L. & Ochoa, O. O. A parametric study of variables that affect fiber microbuckling initiation in composite laminates: Part 1—analyses, part 2—experiments. *J. Compos. Mater.* **26**, 1594–1643 (1992).
12. Moran, P., Liu, X. & Shih, C. Kink band formation and band broadening in fiber composites under compressive loading. *Acta Metall. et Materialia* **43**, 2943–2958 (1995).
13. Waas, A., Babcock Jr, C. & Knauss, W. A mechanical model for elastic fiber microbuckling. *J. Appl. Mech.* **57**, 138–147 (1990).
14. Maire, E. & Withers, P. J. Quantitative X-ray tomography. *Int. Mater. Rev.* **59**, 1–43 (2014).
15. Garcea, S., Sinclair, I., Spearing, S. & Withers, P. Mapping fibre failure in situ in carbon fibre reinforced polymers by fast synchrotron X-ray computed tomography. *Compos. Sci. Technol.* (2017).
16. Altendorf, H. & Jeulin, D. 3D directional mathematical morphology for analysis of fiber orientations. *Image Analysis & Stereol.* **28**, 143–153 (2011).
17. Tausif, M., Duffy, B., Grishanov, S., Carr, H. & Russell, S. J. Three-dimensional fiber segment orientation distribution using X-ray microtomography. *Microsc. Microanal.* **20**, 1294–1303 (2014).
18. Wirjadi, O., Schladitz, K., Easwaran, P. & Ohser, J. Estimating fibre direction distributions of reinforced composites from tomographic images. *Image Analysis & Stereol.* **35**, 167–179 (2016).
19. Requena, G. *et al.* 3D-quantification of the distribution of continuous fibres in unidirectionally reinforced composites. *Compos. Part A: Appl. Sci. Manuf.* **40**, 152–163 (2009).
20. Czabaj, M. W., Riccio, M. L. & Whitacre, W. W. Numerical reconstruction of graphite/epoxy composite microstructure based on sub-micron resolution X-ray computed tomography. *Compos. Sci. Technol.* **105**, 174–182 (2014).
21. Emerson, M. J., Jespersen, K. M., Dahl, A. B., Conradsen, K. & Mikkelsen, L. P. Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials. *Compos. Part A: Appl. Sci. Manuf.* **97**, 83–92 (2017).
22. Sencu, R. *et al.* Generation of micro-scale finite element models from synchrotron X-ray CT images for multidirectional carbon fibre reinforced composites. *Compos. Part A: Appl. Sci. Manuf.* **91**, 85–95 (2016).
23. Ueda, M., Mimura, K. & Jeong, T.-K. In situ observation of kink-band formation in a unidirectional carbon fiber reinforced plastic by X-ray computed tomography imaging. *Adv. Compos. Mater.* **25**, 31–43 (2016).

24. Arun, K. S., Huang, T. S. & Blostein, S. D. Least-squares fitting of two 3-d point sets. *IEEE Trans. Pattern Anal. Mach. Intell.* 698–700 (1987).
25. Hull, D. & Clyne, T. *An introduction to composite materials* (Cambridge university press, 1996).
26. Wang, Y. *et al.* X-ray computed tomography study of kink bands in unidirectional composites. *Compos. Struct.* **160**, 917–924 (2017).
27. Paganin, D., Mayo, S., Gureyev, T. E., Miller, P. R. & Wilkins, S. W. Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object. *J. Microsc.* **206**, 33–40 (2002).
28. Maire, E., Le Bourlot, C., Adrien, J., Mortensen, A. & Mokso, R. 20 Hz X-ray tomography during an in situ tensile test. *Int. J. Fract.* **200**, 3–12 (2016).
29. Dahl, V. A., Trinderup, C. H., Emerson, M. J. & Dahl, A. B. Content-based propagation of user markings for interactive segmentation of patterned images. *In submission* (2017).

Acknowledgements

The authors would like to gratefully acknowledge Dr. Rajmund Mokso, former beamline scientist at TOMCAT (SLS), and Dr. Eric Maire at INSA-Lyon MATEIS. Financial support from CINEMA (the alliance for imaging of energy materials) [DSF-grant no. 1305-00032B] under the Danish Council for Strategic Research is acknowledged. P.J.W. is grateful to the European Research Council for the funding under CORREL-CT [grant no. 695638].

Author contributions statement

M.J.E. drafted the manuscript, performed the statistical image analysis of the tomograms and, together with V.A.D, A.B.D and K.C., decided on the geometric definitions for presenting the results of the statistical analysis. P.J.W. came up with the idea of investigating fibre micro-buckling and kink-band formation using X-ray micro-tomography. Y.W. designed the experiment, manufactured the sample, conceived the tomograms and contributed to providing the background. All authors reviewed the manuscript.

Competing financial interests statement

The authors declare no competing financial interests.