Universitat de Girona

Doctoral Thesis

Medical imaging techniques for the analysis and measurement of composite materials

Adrià Julià i Juanola

2022

Universitat de Girona

Doctoral Thesis

Medical imaging techniques for the analysis and measurement of composite materials

Adrià Julià i Juanola

2022

Programa de Doctorat en Tecnologia

Directors: Imma Boada i Oliveras, PhD
Universitat de Girona

Marc Ruiz i Altisent, PhD
Universitat de Girona

Memòria presentada per optar al títol de doctor per la Universitat de Girona

Medical imaging techniques for the analysis and measurement of composite materials
Programa de Doctorat en Tecnologia
June, 2022

Universitat de Girona
Graphics and Imaging Laboratory (GILAB)
Universitat de Girona, Edifici P-IV
17001 Girona
Catalonia

41.963175 N, 2.831735 E
adria.julia@udg.edu

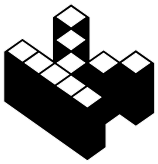

The information herein shall be free.

Josep Cornellà i Canals, MD, PhD
1950 - 2014

In memoriam

Contents

List of Figures

List of Publications

This thesis is a compendium of the following articles:

Article 1 (Chapter 2)

Marc Ruiz i Altisent; Adrià Julià i Juanola; Imma Boada i Oliveras, **Starviewer and its comparison with other open-source DICOM viewers using a novel hierarchical evaluation framework**, International Journal of Medical Informatics [1386-5056], *vol. 137 (2020): 104098*, 10.1016/j.ijmedinf.2020.104098, (accepted in February 2020)

SCImago Journal Rank (SJR): 1.124 (2020), Q1 Health Informatics, H-Index 114

Journal Impact Factor (JIF): 4.046 (2020), 26/107 Q1 Health Care Sciences And Services

Journal Citation Indicator (JCI): 1.13 (2020), 47/223 Q1 Comp. Science, Inf. Systems

Article 2 (Chapter 3)

Adrià Julià i Juanola; Marc Ruiz i Altisent; Imma Boada i Oliveras, **An efficient and uniformly behaving streamline-based μCT fibre tracking algorithm using volume-wise structure tensor and signal processing techniques**, Computer Methods in Applied Mechanics and Engineering [0045-7825], *vol. 394 (2022): 114898*, 10.1016/j.cma.2022.114898, (accepted in February 2022)

SCImago Journal Rank (SJR): 2.180 (2021), Q1 Computational Mechanics, H-Index 205

Journal Impact Factor (JIF): 6.756 (2020), 2/108 Q1 Mathematics, Interdisc. Appl.

Journal Citation Indicator (JCI): 2.33 (2020), 1/156 Q1 Mechanics

Article 3 (Chapter 4)

Adrià Julià i Juanola; Marc Ruiz i Altisent; Narcís Coll i Arnau; Imma Boada i Oliveras **A frequency-limited waviness and curvature measurement algorithm for noisy and irregularly sampled composite fibre trackings**, Computer Methods in Applied Mechanics and Engineering [0045-7825], (submitted in June 2022)

SCImago Journal Rank (SJR): 2.180 (2021), Q1 Computational Mechanics, H-Index 205

Journal Impact Factor (JIF): 6.756 (2020), 2/108 Q1 Mathematics, Interdisc. Appl.

Journal Citation Indicator (JCI): 2.33 (2020), 1/156 Q1 Mechanics

Abstract

IV

English

The purpose of this thesis is the research and development of methods, algorithms and tools for the analysis of composite materials. In order to improve their properties, one must employ specific techniques in order to gain insight of the structure of their internal reinforcement fibres. While these goals can be tackled in multiple ways, a common and established way is the employment of non-destructive means such as the micro-CT (micro-computed tomography). Due to its technical similarities with medical CT, first, we analyse the most relevant features across different state-of-the-art open-source medical imaging software using a novel hierarchical evaluation framework.

From this evaluation it has been observed that many of the 3D-based visualization features can be used for the analysis of composites reinforcement fibres. However, the measurement of their properties requires specialized methods to reconstruct them in order to quantify their orientations, curvatures, waviness, or other aspects. For this reason, we propose applying the following two methods respectively: (i) an efficient and uniformly behaving streamline-based micro-CT fibre tracking algorithm using volume-wise structure tensor and signal processing techniques; and (ii) a frequency-limited waviness and curvature measurement algorithm for noisy and irregularly sampled composite fibre trackings. The proposed methods have been implemented and integrated in the Starviewer platform. As these must remain understandable, explainable and controllable by an expert user, black-box approaches or the ones requiring training datasets have been discarded.

Despite the differences between the volumetric acquisitions and polygonal fibre reconstructions, the proposed methods consider the dataset as a valuable signal to preserve. It is progressively processed and filtered according to the signal processing principles in order to obtain a uniform 3D behaviour. The methods properties have been assessed with both real and synthetic datasets; and due to the complex waveform shapes required for their testing, some datasets were produced using a novel methodology based on Gaussian splatters. All methods proposed in this thesis and their corresponding articles have been mathematically defined and accompanied by a large amount of supporting material files in order to ensure their reproducibility.

Catalan

El propòsit d'aquesta tesi és investigar i desenvolupar algorismes, mètodes i eines per analitzar materials compostos. Per millorar-ne les propietats cal usar tècniques específiques que permetin conèixer l'estructura interna de les fibres de reforç. Tot i que aquests objectius poden assolir-se de diverses maneres, una forma usual de fer-ho és a través de tècniques no destructives com la micro-TC (micro-tomografia computada). A causa de la ressemblança amb la TC mèdica, primer s'han analitzat les característiques més rellevants dels principals programaris de codi obert per a la visualització de la imatge mèdica emprant un nou sistema d'avaluació jeràrquic.

A partir d'aquesta avaluació s'ha pogut observar com moltes de les funcionalitats de visualització tridimensionals poden usar-se per analitzar les fibres de reforç dels materials compostos. Tanmateix, per mesurar-ne les propietats calen mètodes especialitzats per reconstruir-les i així poder determinar-ne l'orientació, la curvatura, l'ondulació, o altres aspectes. Per aquesta raó s'han proposat dos mètodes que cal aplicar consecutivament: (i) un algorisme eficient i que es comporta uniformement per reconstruir les fibres volumètricament emprant el tensor d'estructura i les tècniques de processament del senyal; i (ii) un algorisme limitat freqüencialment per mesurar la curvatura i l'ondulació de fibres amb soroll i mostrejades irregularment. Els mètodes proposats han estat implementats i integrats a la plataforma Starviewer. Com que aquests han de ser comprensibles, explicables i controlables per un usuari expert, s'han descartat paradigmes de caixa negra o altres mètodes que necessitin dades d'entrenament.

Tant en el cas de les adquisicions volumètriques com en el de les reconstruccions de fibres poligonals, els mètodes proposats tracten tot el conjunt de dades com un senyal valuós a preservar. Aquest es processa i filtra eficientment segons els principis de processament del senyal per tal d'assolir un comportament uniforme des d'un punt de vista tridimensional. Llurs propietats han estat avaluades tant amb jocs de prova reals com sintètics. En aquest últim cas, i a causa de la complexitat de les formes d'ona requerides per realitzar-ne la seva avaluació, s'ha proposat una nova metodologia basada en *Gaussian splatters* per produir-los. Tots els mètodes proposats en aquesta tesi i en els corresponents articles han estat definits matemàticament i complementats amb una gran quantitat de material suplementari per tal d'assegurar-ne la reproductibilitat.

Spanish

El propósito de esta tesis es investigar y desarrollar algoritmos, métodos y herramientas para el análisis de materiales compuestos. Para mejorar sus propiedades es necesario emplear técnicas específicas que permitan conocer su estructura interna y la de sus fibras de refuerzo. A pesar de que estos objetivos pueden lograrse de distintas formas, una manera habitual de hacerlo es a través de técnicas no destructivas como la micro-TC (micro-tomografía computarizada). Debido a su semejanza con la TC médica, primero se han analizado las características más relevantes de los principales programas de código abierto de visualización de la imagen médica empleando un nuevo sistema de evaluación jerárquico.

A partir de esta evaluación se ha logrado comprobar como muchas de las funcionalidades relativas a la visualización 3D pueden ser usadas para analizar las fibras de refuerzo. Así mismo, para medir sus propiedades es necesario usar métodos especializados para reconstruirlas y así poder determinar su orientación, curvatura, ondulación, u otros aspectos. Por esta razón se propone aplicar los siguientes dos métodos consecutivamente: (i) un algoritmo eficiente y de comportamiento uniforme para reconstruir las fibras volumétricamente usando el tensor de estructura y técnicas de procesamiento de la señal; y (ii) un algoritmo limitado frecuencialmente para medir la curvatura y la ondulación de las fibras con ruido y muestreadas irregularmente. Los métodos propuestos han sido implementados e integrados en la plataforma de visualización Starviewer. Debido a que estos deben ser comprensibles, explicables y controlables por parte de un usuario experto, se han descartado los paradigmas de caja negra o aquellos que requieran datos de entrenamiento.

En el caso de las adquisiciones volumétricas, así como en el de las reconstrucciones de fibras poligonales, los métodos propuestos tratan todo el conjunto de datos como una señal valiosa a preservar. Esta es procesada y filtrada eficientemente según los principios de procesamiento de la señal a fin de lograr un comportamiento uniforme desde el punto de vista tridimensional. Estas propiedades han sido evaluadas a través de varios conjuntos de prueba reales y sintéticos. En este último caso, y debido a la complejidad de las formas de onda requeridas para realizar su evaluación, se ha propuesto generarlos a través de una nueva metodología basada en *Gaussian splatters*. Todos los métodos propuestos en esta tesis y sus correspondientes artículos han sido definidos matemáticamente y complementados con una gran cantidad de material suplementario a fin de asegurar su reproducibilidad.

1 Introduction

1.1 Motivation

We, the humans, are usually provided the ability to explore and interact with a physical object or specimen through our various senses. Vision, being the major contributor, provides vast amounts of information regarding the surface of an object; and, provided stereoscopic perception, one can even infer its volume and dimensions [15]. Combined with the tactile sense, one can be informed about the roughness and waviness of said surface [25]; and, with a little knock, augmented with the aid of our auditory system, the hollowness, solidness and other rough estimates about the consistency and mechanical properties can be extrapolated. While this could be deemed as the ultimate experience our senses can provide, note how almost all the information objectively gathered is only at the surface level; we lack the ability to properly sense the internal structure underneath. Fortunately, nowadays, a plethora of instrumentation is available to augment our senses. The purpose of this thesis is the development of algorithms and methods to analyse, visualize, and measure the internal structure of composite materials, which combine a reinforcement (see Figure 1.1) with a binder to make a much better material than its constituents alone.

(a) Woven glass mat with thermoplastic (GMT).

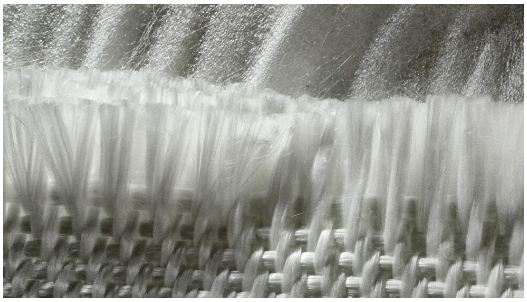

(b) Panel with its composite layers torn apart.

(c) Non-woven fibre glass.

(d) A foamed composite panel flexed and stressed.

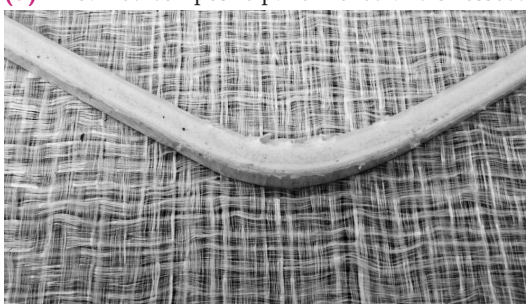

Figure 1.1.: Examples of different reinforcement materials and their applications.

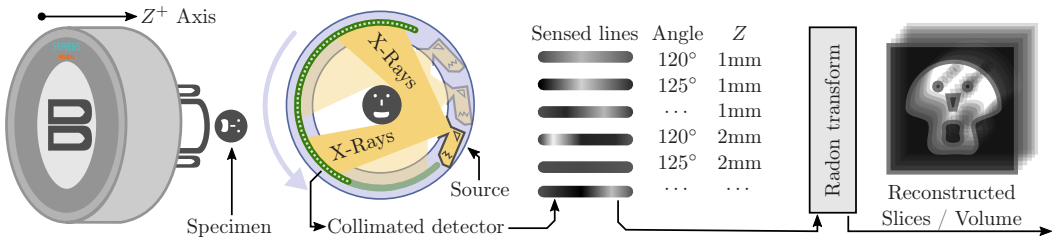

Figure 1.2.: Acquisition pipeline of a medical CT scanner from left to right where: (i) the detector and X-ray source assembly is rotated around the specimen; (ii) acquisitions are stored and annotated with an angle and z-depth metadata; and (iii) a computer performs the inverse Radon transform to obtain the final volumetric slices.

Composites are commonly found in sectors such as the automotive, aeronautical, energy generation, among others where a high mechanical performance is required while keeping a low weight.

1.2 Preliminary concepts

As per the compendium structure of this thesis, we will introduce several background concepts in order to complement and contextualize some of the aspects discussed in the corresponding articles. First, we explain the operating principles of computed tomography and how it is related to acquisition fundamentals for the purpose of non-destructively imaging the inner structure of a composite specimen. Second, the processing of volumetric datasets is contextualized in the medical imaging sector, as well as some common visualization techniques such as the volume rendering which enable a proper visualization and exploration of the acquisitions. Among others, we also present the principles of DTI (diffusion tensor image) modality which is used to image the orientations of neuron tracts in the brain. We do so due to its potential relevance in the analysis of composite materials.

1.2.1 Computed tomography

In that regard, one of the techniques enabling a proper insight underneath a surface is the computed tomography (CT) [48, 53]. The acquisition apparatus usually takes the form of an annular gantry where the object to explore (i.e. specimen) is placed at its centre. This device is then able to produce a reconstructed 2D image (i.e. slice) of the inner circular area passing through the specimen. Briefly summarized: each slice is imaged by the point emission of radiation and its consecutive measurement at a diametrically opposed sensing area. This requires the transmission of energy from the source to a sensing area following the shortest geometric path (i.e. a straight line). Depending on the specimen, said radiation can range from infrared, visible photons, and even electrical resistivity. Nonetheless, the most commonly used one is in the form of high energy photons emitted by X-ray sources. Despite its ionizing nature and the associated health hazards, this type of radiation exhibits a deep penetrative behaviour in a wide spectrum of substances ranging from metal to human flesh. For a proper reconstruction, aspects such as a linear absorption response, not cropping the object, or considering the scale of the features of interest relative to the radiation wavelength, are examples of relevant aspects to take into account for a successful acquisition. If we

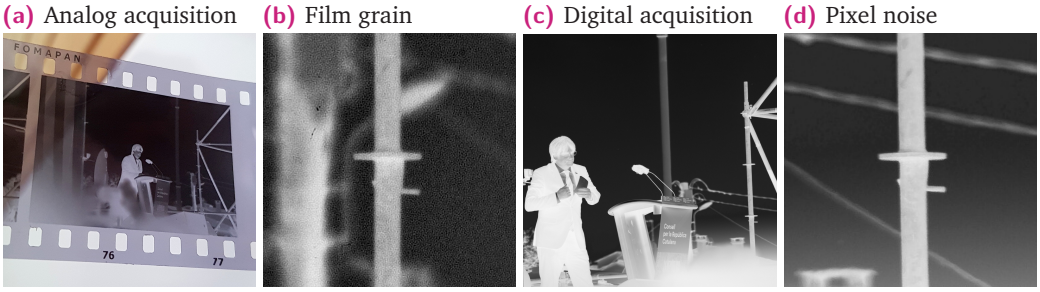

Figure 1.3.: A common scene and a detailed close-up (b, d) acquired using a 35mm ISO400 film (a, b), and a digital APS-C sensor of a full-spectrum Dynax 7D (c, d).

could experience vision close to the diffraction limit, about 500nm, we would observe a diffuse blob. But, if instead synchrotron light or other coherent source is employed one can overcome some of these physical barriers in order to, for example, image the tiny carbon fibres in a composite material whose scale is in the range of a few nanometres as shown in Figure 1.11d versus 1.11a, 1.11b, and 1.11c which are acquired with classic X-ray sources.

Unlike 2D radiographies on a chemically photosensitive emulsion as illustrated in Figure 1.3, CT is directly sensed through a digital photosite. If it were to be done with film, as illustrated in Figure 1.2, the equivalent acquisition would consist in the exposure of film strip (green) with a diametrically opposed fan-shaped beam X-Ray source (yellow). Repeating said process thousands of times at different angles would yield to a plethora of film strips labelled at the acquisition angle they were taken. All of them, despite jointly containing the complete information about the inner object structure, their direct visualization would be meaningless: a computer algorithm is required to perform the inverse Radon transformation [51] in order to obtain usable tomographic images. When the scanning process is also repeated in a sweeping movement, a 3D cylindrical field of view is obtained. This whole acquisition conforms a volume (i.e. 3D image) whose constituent parts are named voxels (i.e. 3D pixels).

1.2.2 Acquisition principles and limitations

> *Much of our knowledge must always remain uncertain. The* most *we can know is in terms of probabilities.*
>
> — **Richard P. Feynman, PhD**
> Feynman Lectures on Physics
> Vol. 1, Chapter 6

As illustrated in Figure 1.3 digital acquisitions are not immune to noise. This natural phenomena (whenever a measurement is taken) also applies in the case of digitally reconstructed acquisitions (such as CT). While it could be artificially cropped, provided an ADC (analog-to-digital converter) with enough dynamic range and resolution, noise will inevitably be captured in a noticeable manner. Its random behaviour has a major impact in the least significant bits (LSB) of a pixel.

Albeit this phenomena may be considered undesirable, it contains latent information nonetheless. For instance, should one average several noisy acquisitions of the very same scene, the final result would be a higher quality picture. Although trading spatial resolution, this averaging can be employed on consecutive (neighbouring) slices to obtain a less noisy image. Also, when an image is observed as a signal (i.e. frequency space), the amount of high-frequency noise will depend on the sampling (i.e the resolution) [68, 47, 61]. This means that the spectral energy of these random high-frequencies infers a fractal behaviour into the signal which shall be considered should any proper measurement be made [39].

1.2.3 The DICOM and other imaging modalities

Due to the vast amounts of information gathered, computers are the cornerstone making CT technology possible. Since its inception and development, CT, alongside PET (positron emission tomography), and MRI (magnetic resonance imaging) which employ other physics fundamentals, have made volumetric datasets a common structure for expressing and digitizing the inner features of objects. In the early days, axial reconstructions were laid in a grid (see Figure 1.4) and printed in a photographic medium such as film. Although being a

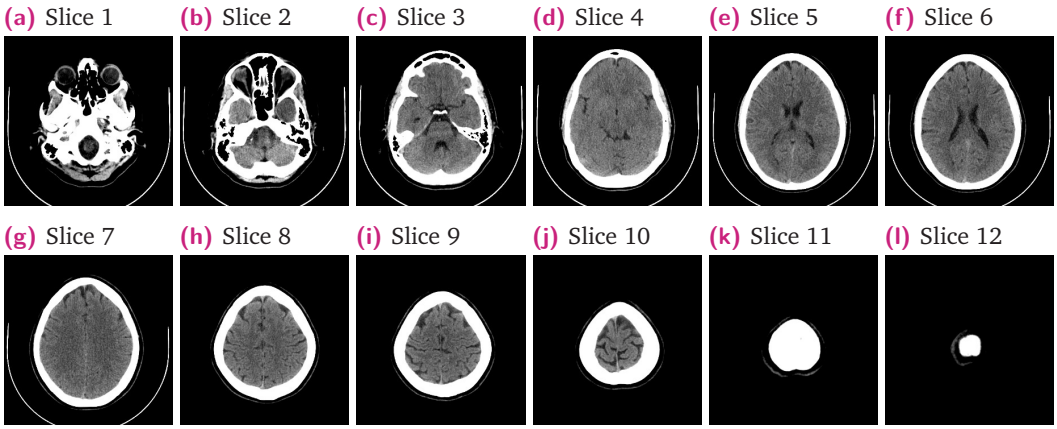

Figure 1.4.: Classical grid arrangement of brain slices acquired with a medical CT scanner.

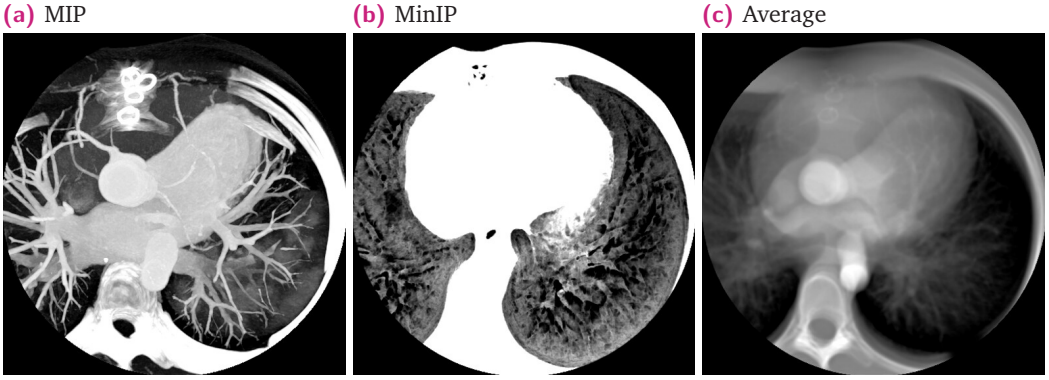

Figure 1.5.: Axial thick-slabs (i.e. many slices simultaneously) using three different projection strategies: (a) maximum intensity projection; (b) minimum intensity projection; and (c) average intensity projection. The circular perimeter is the CT cylindrical field-of-view observed from an axial viewport.

complete representation in the eyes of an expert, in order to unleash its full potential and enable an augmented perception, several specialised software tools had to be developed in order to satisfy the interactive visualization demands. In the medical field, these volumes or any other 2D modalities such as ultrasound, radiography, video, among other techniques dealing with raw signals such as ECG (electrocardiography), are stored and represented using the DICOM standard [45]. It stands for Digital Imaging and Communications in Medicine, and it is the leading standard ensuring the interoperability across medical applications and devices. It is used to capture, archive, and exchange data [24] and also defines a very detailed metadata structure storing aspects such as the acquisition parameters, spatial location, voxel sizes, etc. This information is essential for a metrologically correct interpretation of the data within a dataset.

1.2.4 Medical imaging software

Focusing on imaging software, whether for research or the daily clinical practice, it has become an essential component for the analysis of volumetric datasets [55, 67, 18, 6]. Several studies have been made comparing some of its main features [20, 44, 7, 50, 65, 21], most essential ones being: (i) slicing or the ability to interactively scroll through a volume; (ii) basic reslicing at orthogonal acquisition axes (i.e. axial, sagittal, coronal) as illustrated in Figure 1.6; (iii) measurement tools such as dimensions, area, volume, voxel intensity, statistics with a region of interest, among others; (iv) interactive adjustment of the transfer function (i.e. window-level/window-width) to brighten or enhance the contrast at relevant areas; and (v) the ability to spatially localize one acquisition respect others sharing a common frame of reference. In the medical sector the usage of these features will vary from user to user depending on the imaging modalities they visualise (especially respect 2D versus 3D datasets), and also on the pathologies and diagnosis procedures to be applied.

1.2.5 Other visualization techniques

In addition to the aforementioned ones, medical imaging software usually implements more advanced or refined visualizations. For example, when rendering a slice, one shall consider

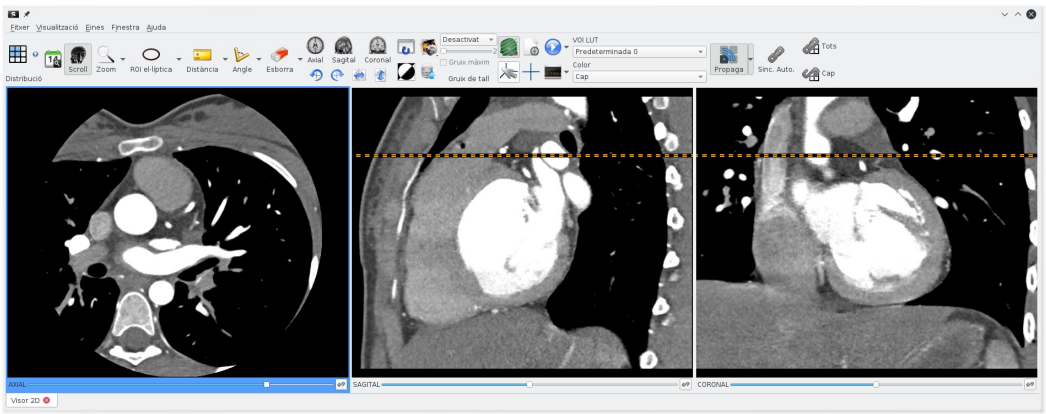

Figure 1.6.: Basic features of medical viewers (in this case Starviewer). A single volume is rendered at three different point of views: axial, sagittal, coronal respectively. The lines (orange) locate the selected viewport (blue) relative to the rest (gray).

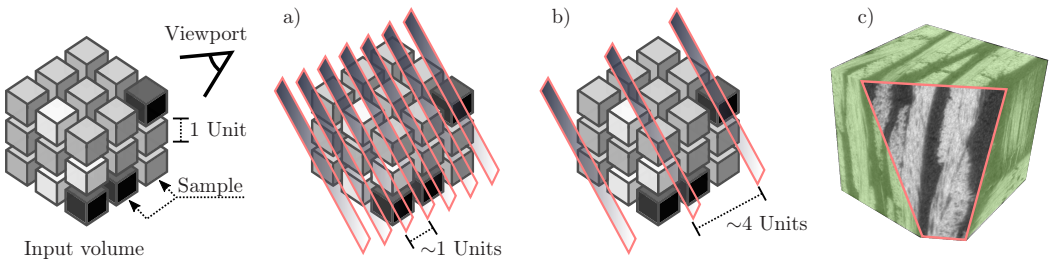

Figure 1.7.: Reslice of a volume respect a non-orthogonal viewport in two scenarios: (a) a slice separation similar to the original dataset; and (b) a bigger separation encompassing several samples which forces the use of a thick-slab approach. In (c) orthogonal views (green) are shown alongside a resliced section (red).

its inherent thickness. Under normal circumstances they are spaced (shifted) regularly, but one may want to combine the information of several slices into a single one. This technique is usually referred to as thick-slab, and requires the user to specify the desired amount of neighbours to "summarize". Then, at each voxel location, a list of values generated from each corresponding voxel respect its neighbouring slices has to be reduced into a new value using any of the three main strategies illustrated in Figure 1.5 which are: (i) *MIP* or maximum intensity projection which takes the greatest value in the list and thus is useful for removing occluding material such as air or water; (ii) *MinIP* which does the opposite by taking the smallest value; and (iii) *average* projection which reduces the noise by trading off spatial resolution, and also provides a similar result to the one obtained with a traditional 2D radiography. These operations can be performed interactively, or as a whole-volume operation by generating a new set of computed slices. Unlike more basic visualizations, note how the resulting level of abstraction becomes decoupled from the original dataset voxels. This technique can be meld into a more generalized form known as MPR (multi-planar reconstruction) or reslicing which produces visualizations not orthogonally aligned with the original acquisition axes. As illustrated in Figure 1.7, the input dataset is repartitioned, resampled, and interpolated into a new set of reoriented slices. Usually, the resolution and thickness match the original volume (in this way it can be made more computationally efficient), but should thicker slices be produced, then the aforementioned thick-slab technique is applied with the average mode. Despite its intensive cost, this operation can usually be performed interactively in modern hardware.

So far, these 2D visualization techniques allow the user to attain a whole volumetric orientation/perception by interacting with the viewport (i.e. changing slices). But, sometimes, a projected 3D view is more desirable for its ease of comprehension. A first approach is representing the volume as a set of surfaces which can be posteriorly rendered and accelerated in a GPU (graphics processing unit). In order to do this, first, certain voxels in a volume shall be converted to polygons using the marching cubes algorithm [38]. The result is an isosurface at a certain user-specified voxel value where a tiny surface is placed by identifying any of the corresponding cases in Figure 1.9. While this polygonal mesh is a lightweight representation of a hollow shell, and its render can perform very well, as shown in Figure 1.9 (right) the visual appearance is poor and lacks information to properly represent the volume with enough detail.

As with the adjustment of window-level/window-width in 2D visualizations; the equivalent feature in 3D is a bit more complex and corresponds to the alteration of the transfer function

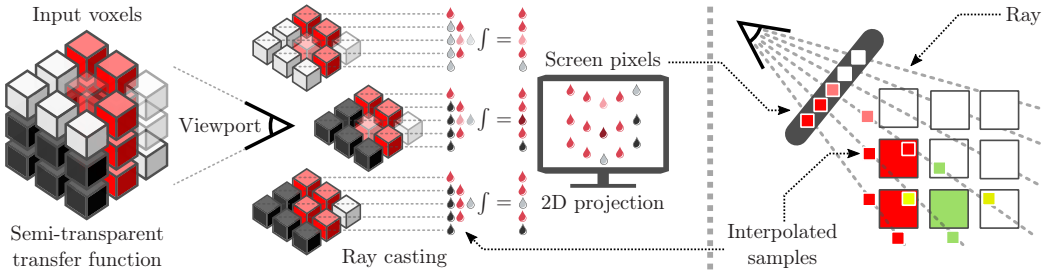

Figure 1.8.: The raycasting process traces piercing rays through volume from the viewport (left); these are then sampled and integrated into actual screen pixels (right).

transparency (i.e. the alpha channel). This function maps input volume scalars to multi-component screen/render values (e.g. red, green, blue, alpha). With it, the user can colourize certain intensities as illustrated in Figure 1.10a or perform interactive adjustments as shown in Figure 1.10c and Figure 1.10d where the volume is progressively peeled in order to explore its internal structure. Other strategies such as the MIP projection can be extrapolated into the 3D case as shown in Figure 1.10b.

Rendering such semi-transparent volume definition requires a specialized technique called raycasting [35]. Its simplified operating principles are illustrated in Figure 1.8 and it proceeds as follows: (i) for every pixel in the screen, a ray is pierced through the volume from the camera focal point (viewport); (ii) a list of interpolated values is produced by sampling the ray through the original volume; (iii) considering the transfer function, occlusion, etc., lists values are reduced into a single one; and (iv) the final RGB value for a specific pixel in the screen is obtained. While being much more computationally intensive than the rendering of a polygonal scene, as it can be observed in Figure 1.10 the quality and meaningfulness of the visualization is more appealing and richer in details.

1.2.6 Composite materials

Focusing in composite materials, the acquisition of volumetric datasets by means such as μCT (micro-computed tomography) is gaining traction as capture devices become more affordable and accessible [22]. Applying the practices and methods from the more mature medical field and adapting them to this new use case is beneficial from both sides. What makes composites special and at the same time complex is their combination of two or more constituent materials referred as: (i) the *reinforcement* which enhances the mechanical properties (e.g. carbon fibre); and (ii) the *matrix* or *binder* which is used to fill the gaps (e.g. resin, epoxy). This symbiotic combination produces a superior material than its constituents alone by combining their advantages. However, unlike isotropic solids such as metals, plastics, among others, composites are usually characterized for its anisotropic behaviour. This means that their mechanical properties vary across their axes depending on the local arrangement orientation of the reinforcement material. In Figure 1.1 one can observe an actual composite and why the reinforcement arrangement dictates its anisotropy. While finite element analysis (FEA) simulations can be crafted to take in account such complex scenario of varying mechanical properties respect the force orientation, circumstances such as fractures or cracks which emerge from a chaotically behaving system are in practice unsolvable in an exact manner. For example, a crystal lattice (i.e. glass) breaks at the weakest regions with more probability, however the effect of these very small and unpredictable variations is

greatly amplified when a failure occurs and thus ends up determining the overall behaviour of the macroscopic observable events. Other examples of a chaotic systems would be the path of least resistance in the atmosphere taken by a lightning, a turbulent flow, among others.

For this reason, checking the actual mechanical properties with real specimens is crucial to validate the performance of a composite [31, 19, 46, 64]. In that regard, the μCT technique fits this purpose as it is a very suited tool to capture the most subtle and finest details inside a specimen. Although the images by themselves (given a proper visualization software) enable a subjective analysis by an expert user, performing actual measurements requires specialized techniques and methods [32, 1, 70, 5, 23].

> *Chaos: When the present determines the future, but the approximate present does not approximately determine the future.*
>
> **— Edward Norton Lorenz, PhD**

As illustrated by the carbon fibre composite in Figure 1.11d; a synchrotron light μCT is able to capture the individual fibres. At the same time, in Figure 1.11a, Figure 1.11b and Figure 1.11c, while the pixel size is too big for this purpose and only the bundles are distinguishable, its local fibre orientation can be somewhat inferred by the texture. From an economic standpoint, the latter requires less expensive equipment but the voxel size is kept above the diffraction limit. Depending on the dataset to analyse, the measurement algorithms and methods will vary [17]. On the one hand, authors such as [34, 4, 54, 30, 62, 13, 49, 57, 27, 26] analyse the fibre orientations using image processing filters based on dictionaries, neighbour analysis, variance, segmentation, among others. Their results are then presented in such a way that the distributions of fibre orientations can be analysed and represented using plain data or more advanced visualizations such as 2D histograms. On the other hand, and pursuing similar goals, [8, 9, 12, 29, 52, 59, 14, 69, 10, 2, 40] employ a more fine-grained approach by processing and simplifying the whole volume into a polygonal representation where only the medial axis of each fibre is stored in a much simpler and separable entity from which independent measures and visualizations can be performed. In this scenario the acquisitions shall resemble the example in Figure 1.11d and have enough quality to properly separate each individual fibre path. This new representation

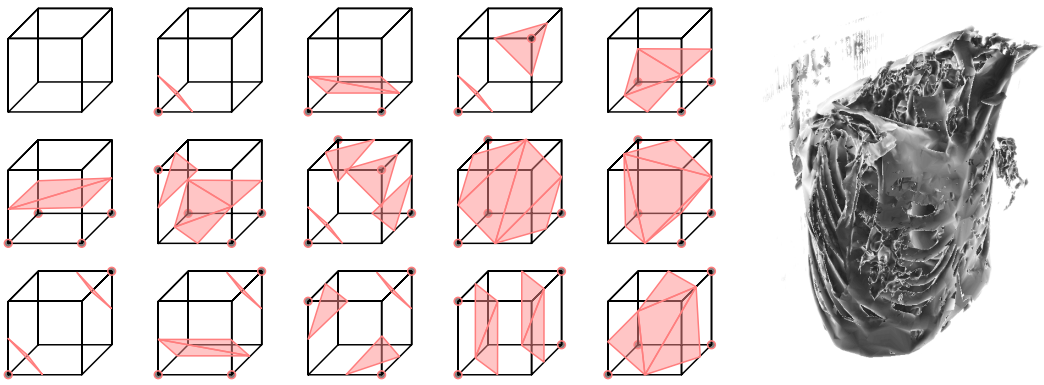

Figure 1.9.: All combinations of the marching cubes algorithm [38] and a rendered example.

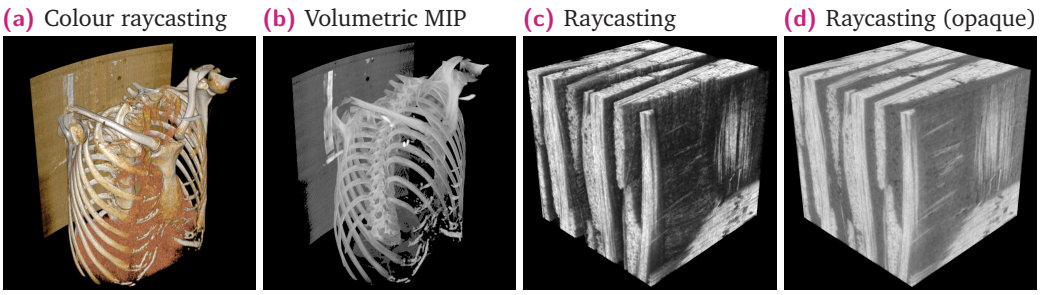

Figure 1.10.: Raycasted renders of a human torso and a glass fibre composite.

consisting of a connected list of segments (i.e. a VTK polyline [58]) may allow for a more advanced posterior analysis than pure image processing approaches. It is also worth noting that alternative acquisition and measurement techniques exist; albeit with less resolution they are usually capable of directly quantifying the anisotropy of a specimen by measuring an alternative physical property [42].

1.2.7 Tensorial volumes in medicine

Some of the aforementioned techniques employ tensors to express the local anisotropy around a position in space. So far, the volumetric datasets store a single scalar value corresponding to the intensity measured by the acquisition. Notwithstanding, there is no impediment to store multi-component values such as 3D vectors. While a volumetric vector field could be used to represent the directionality and magnitude (i.e. the presence) of fibres in each position in the acquisition, this is not a complete representation. For this reason, some of the aforementioned measurement techniques build a structure tensor from the image gradient (a vector field) in order to obtain a proper representation of the anisotropy within a local neighbourhood.

Processing and visualizing tensorial fields is a challenging task which requires specific techniques [66, 3, 43]. Fortunately, in the medical field, a well established and mature imaging modality in MRI is the DTI (diffusion tensor imaging) [33, 63, 41]. Although the physical principles measured by DTI do not apply in the aforementioned structure tensor description, this technique acquires a tensor at each measurement location expressing the Brownian diffusion freedom of a liquid constrained within a voxel. In the simplest form and in 3D space, the diffusion can be classified in three main categories: (i) *linear anisotropy*, found where the water diffusion is restricted linearly (e.g. aligned neuron axons); (ii) *planar anisotropy*, at the regions where neurons spread/diverge towards different directions; and (iii) *isotropy*, which can be found in pure and unrestricted water regions (e.g. cerebrospinal fluid). Note that these are not strict categorical classifications. Each one is quantified by a scalar map computed by analysing the eigenvalues of a tensor [66]. In the case of composites, despite not acquiring the tensor directly, applying some of these already established techniques in the medical field can be deemed of great value for the purpose of composites analysis.

(a) A stressed glass fibre composite with cracks. **(b)** Several layers of glass fibre fabric.

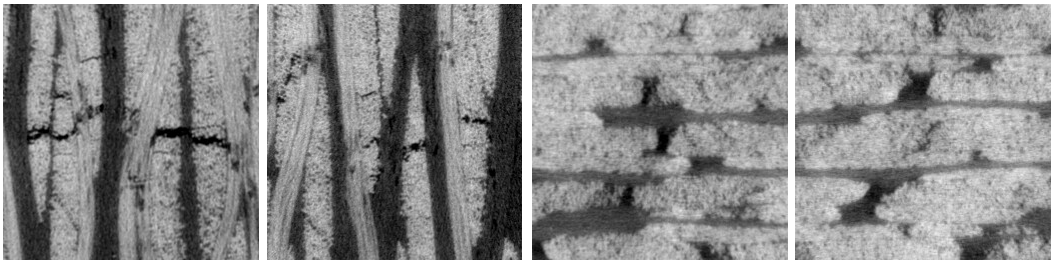

(c) Carbon fibres (smaller, less contrast). **(d)** Individual carbon fibres (synchrotron light).

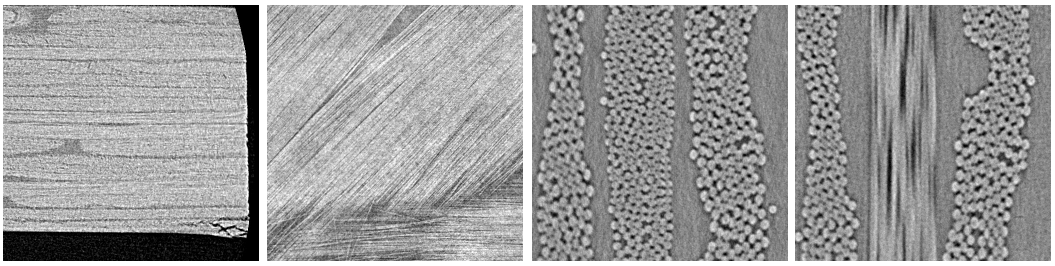

Figure 1.11.: A diverse set of μCT slices combined with different specimens: glass (top row), and carbon (bottom row). (d) Was acquired with synchrotron light due to the small pixel size (in the order of nanometres).

1.3 Objectives

The purpose of this thesis is researching and developing methods, algorithms and tools for the analysis of composite materials. Its nature requires advanced and specific techniques in order to achieve an understanding and ultimately lead to an improvement of the mechanical properties, characteristics, or features. While these goals can be tackled in multiple ways, a common and established way is the employment of non-destructive means such as the μCT technique which is capable of individually imaging each fibre in a specimen. In that regard, our goal is to analyse these vast amounts of information with explainable approaches which are controllable and understandable by an expert operator such as a researcher. Hence, the author explicitly sets as an objective the avoidance of unsupervised or pure black-box approaches. Its use would fundamentally undermine and conflict with the aforementioned principle. These type of tools, loosely referred as "artificial intelligence", employ techniques requiring training datasets such as machine learning, neural networks and genetic algorithms, among others. Despite this, its usage is not totally ruled-out for use cases such as assisting the user with recommendations or performing automations.

> *For a successful technology, reality must take precedence over public relations, for nature cannot be fooled.*
>
> **— Richard P. Feynman, PhD**
> Appendix F
> Personal Observations on Reliability of Shuttle

In the medical sector, the acquisition and processing of volumetric datasets from CT or MRI modalities is a very common and established practice for human diagnostic purposes. Despite being a different application, technically, μCT datasets are equivalent acquisitions from an image processing perspective. For this reason, we want to consider, analyse, and compare the available tools and its use cases among the medicine professionals in order to apply some of its key features in the field of materials analysis.

Briefly enumerated, the main pursued goals in this thesis are:

1. Study the usage and features of medical imaging tools dealing with volumetric datasets.
2. Applying these medical imaging approaches in the field of composite materials.
3. Developing specific methods and algorithms to achieve an understanding of the composite materials and their implications on mechanical properties.
4. Explicitly pursuing explainability, proper measurements, and avoidance of black-box approaches.

1.4 Thesis outline

This thesis has been organized in six chapters. After the introduction, motivation, preliminary concepts, objectives, and the outline herein described, the following chapters describe the contributions in an orderly manner conforming a continuum of concepts building upon the foundations set by the previous ones.

In Chapter 2 (Article 1) we present Starviewer, a DICOM viewer developed in C++ which employs several libraries, among them the VTK (Visualization toolkit) [28]. The viewer features have been assessed and compared with other viewers using a novel hierarchical evaluation framework according to the needs of different medical audiences. And the platform itself and its associated libraries have been a fundamental building block for Chapter 3 (Article 2) and Chapter 4 (Article 3).

While proven useful for the medical practice, in Chapter 3 we propose extending the use of Starviewer to the analysis of composite specimens imaged with μCT for its technical similarities to the datasets acquired with MRI or medical CT. However, these μCT datasets are larger and their analysis usually requires the employment of specialized tools for extracting the most relevant information. In that regard we propose a technique to reconstruct fibre paths (i.e. medial axis) as a simplified polygonal representation for its posterior analysis.

Taking these, in Chapter 4, we propose a novel curvature and waviness measurement method considering the varying fractality effect depending on the reconstruction sampling in order to obtain meaningful and comparable values when the acquisition varies. Note that Chapter 3 and Chapter 4 respectively conform a complete image processing pipeline starting from raw slices to useable measurements.

Finally the overall results are discussed in Chapter 5, and this thesis is concluded by Chapter 6 where the most relevant highlights are emphasized and the future is work presented.

2

Article 1
Starviewer and its comparison with other open-source DICOM viewers using a novel hierarchical evaluation framework

Marc Ruiz i Altisent; Adrià Julià i Juanola; Imma Boada i Oliveras, **Starviewer and its comparison with other open-source DICOM viewers using a novel hierarchical evaluation framework**, International Journal of Medical Informatics [1386-5056], *vol. 137 (2020): 104098*, 10.1016/j.ijmedinf.2020.104098, (accepted in February 2020)

SCImago Journal Rank (SJR): 1.124 (2020), Q1 Health Informatics, H-Index 114

Journal Impact Factor (JIF): 4.046 (2020), 26/107 Q1 Health Care Sciences And Services

Journal Citation Indicator (JCI): 1.13 (2020), 47/223 Q1 Comp. Science, Inf. Systems

International Journal of Medical Informatics 137 (2020) 104098

Contents lists available at ScienceDirect

International Journal of Medical Informatics

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

Starviewer and its comparison with other open-source DICOM viewers using a novel hierarchical evaluation framework

Marc Ruiz, Adrià Julià, Imma Boada*

Graphics and Imaging Laboratory (GILAB), Universitat de Girona, Edifici P-IV, 17001 Girona, Catalonia, Spain

ARTICLE INFO

Keywords:
DICOM viewer
Open-source
Software comparison
Medical imaging

ABSTRACT

Methods: The aim of the paper is twofold. First, we present Starviewer, a DICOM viewer developed in C++ with a core component built on top of open-source libraries. The viewer supports extensions that implement functionalities and front-ends for specific use cases. Second, we propose an adaptable evaluation framework based on a set of criteria weighted according to user needs. The framework can consider different user profiles and allow criteria to be decomposed in subcriteria and grouped in more general categories making a multi-level hierarchical structure that can be analysed at different levels of detail to make scores interpretation more comprehensible.

Results: Different examples to illustrate Starviewer functionalities and its extensions are presented. In addition, the proposed evaluation framework is used to compare Starviewer with four open-source viewers regarding their functionalities for daily clinical practice. In a range from 0 to 10, the final scores are: Horos (7.7), Starviewer (6.2), Weasis (6.0), Ginkgo CADx (4.1), and medInria (3.8).

Conclusions: Starviewer provides basic and advanced features for daily image diagnosis needs as well as a modular design that enables the development of custom extensions. The evaluation framework is useful to understand and prioritize new development goals, and can be easily adapted to express different needs by altering the weights. Moreover, it can be used as a complement to maturity models.

1. Introduction

DICOM, that stands for Digital Imaging and Communications in Medicine, is the leading standard for image data management in medical applications and is used to capture, exchange, and archive image data in Picture Archiving and Communication Systems (PACS) [1]. PACS can be seen as a centralized repository of all medical images. To diagnose, physicians require a DICOM viewer with features to connect with the PACS in order to retrieve and store images and with other functionalities to visualize, explore and analyse the information represented in these images [2].

There is a wide variety of privative and open-source DICOM viewers. Focusing on the latter, their functionalities range from simple 2D or 3D visualizations to more advanced techniques and image processing tools to measure the volume of lesions or combine information from different image modalities, among others. Recently, Valeri et al. [3] evaluated the most representative viewers for GNU/Linux, Windows, and macOS; beeing the best viewers 3D Slicer [4,5], medInria [6], MITK Workbench (MITK 3M3 before 2009) [7,8], VolView (currently deprecated) [9], VR Render (currently deprecated) [10], and OsiriX [11,12]. More recently, Haak et al. [13] did an evaluation centering not only on viewing functionalities, but also on platforms and interfaces. MIPAV [14,15] and Weasis [16] were identified as superior open-source tools. In the macOS community, the most popular open-source viewer was OsiriX [11] which is currently a proprietary solution. However, the last open-source version of OsiriX was forked into several projects by the community, being Horos [17] the most popular one.

The offer of open-source software is continuously growing in all the areas of research and practice and also in the medical one. Selecting the product that better fits the health environment and more specifically user needs becomes a complex task. In this context, under the premise that the quality of a software system is largely determined by the quality of the software process used to build it, capability/maturity models have been proposed. These can be represented in a two-dimensional space where a first dimension represents what is done and a second one of how well it is done, i.e. the capability/maturity dimension. These models have a set of levels ranging from the lowest, where no conception of maturity is defined, to the highest, representing total capability maturity. To determine the level of a software quality, attributes such as functionality, reliability, usability, efficiency,

* Corresponding author.
E-mail addresses: marc.ruiz@udg.edu (M. Ruiz), adria.julia@udg.edu (A. Julià), imma.boada@udg.edu (I. Boada).

https://doi.org/10.1016/j.ijmedinf.2020.104098
Received 3 February 2019; Received in revised form 18 December 2019; Accepted 10 February 2020

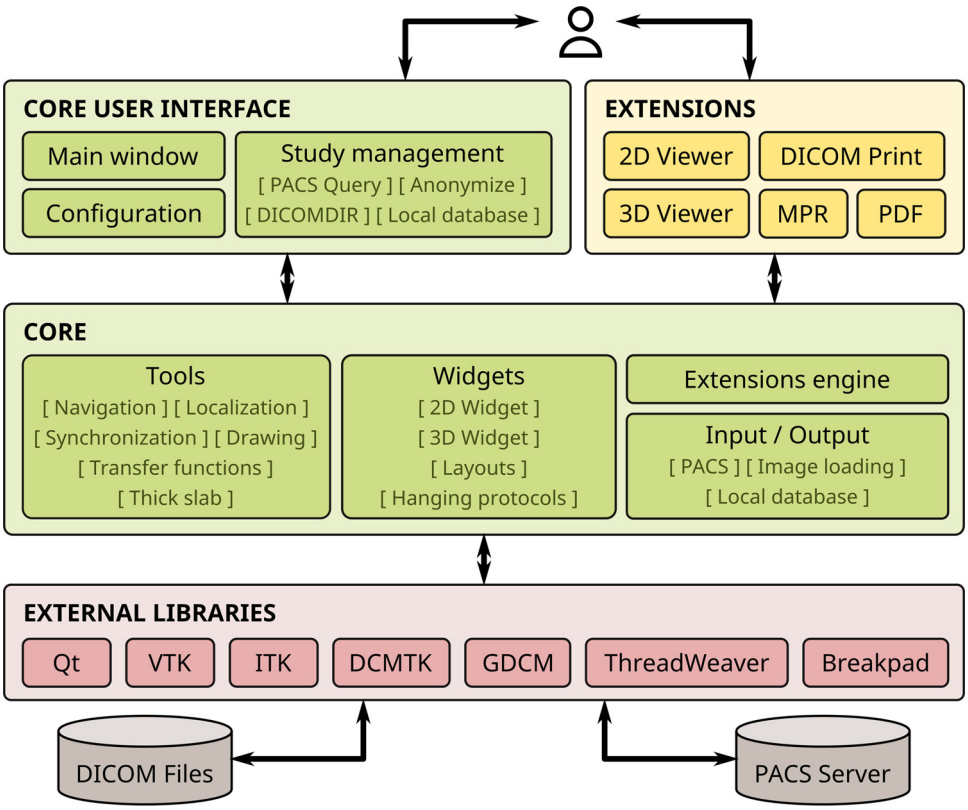

Fig. 1. Starviewer architecture block diagram.

maintainability, and portability are considered. In the healthcare domain, several maturity models have been proposed but they are still at an early stage of development [18–20]. For a review on this topic see [21] and for a comparison of maturity models on open-source software see [22]. Unfortunately, no consensus on the aspects or the models that have to be considered for an evaluation in a specific area exists. This leads, in some cases, the selection of a DICOM viewer to a process based on recommendation and past experiences [23,24].

A key point of maturity models relies on the selection of the parameters to be evaluated and the grading system used to evaluate them. In the case of open-source DICOM viewers, different parameters and grading systems have been proposed to compare them. On the one hand, there are authors that consider the most important features to be installation facilities, support, and documentation [25], while others consider the provided working functionalities more important [26,3,13]. On the other hand, grading approaches range from the simplest, where yes/no classification are considered [13,3] to the most elaborated grading scales [26]. With the aim of satisfying as many audiences as possible, we propose a fully customizable evaluation framework with no limits on evaluated criteria and with the possibility to modify the grading strategy. The framework uses a hierarchical structure to maintain information grouped in audiences, groups, criteria, and subcriteria. Weights are assigned for each node and are distributed across the sibling nodes for each level in the hierarchy. Scores range from 0 to 1 and are computed as a weighted sum which is propagated to the upper nodes providing intermediate weighted scores for each level in the hierarchy. For more details about the calculation refer to Fig. 2, the actual supporting material spreadsheet, and Appendix B. In this way, a rapid evaluation of the software can be obtained with more or less detail depending on the user interests. The proposed framework cannot be a substitute of a maturity model but a complement to determine the values of parameters required by the models.

Although our framework mainly focuses on technological aspects, the overall evaluation of free/open-source projects requires the consideration of other aspects such as production, distribution and support. Different models to evaluate these factors have been proposed. Kamseu and Habra [27] presented a three dimensional model where the development process, the community, and the project are evaluated. The *Qualification and Selection of Open Source Software* (QSoS) applies four independent and iterative steps aimed at defining, evaluating, qualifying, and selecting open-source solutions based on software support and technology [28]. The *Quality Platform for Open Source Software (QualiPSo)*, includes an evaluation framework regarding the trustworthiness of open-source projects [29,30]. Sung et al. adapted the ISO/IEC 9126 standard to free/open-source projects focusing on the quality of the products [31]. The *Evaluation Framework for Free/Open source projects* (EFFORT), is a framework that, once customized for a specific context, supports the evaluation of product quality, attractiveness, and community trustworthiness [32,33]. For more details on these methods see [33]. Note that unlike described methods our approach focuses on functionalities, audiences and technological aspects of medical imaging software but not on production, distribution and support. Therefore, we consider it as a suitable complement to the described methods and not a substitute.

Besides this introduction, the paper has been structured as follows. In Section 2, the architecture of Starviewer DICOM viewer is presented, as well as the proposed evaluation framework and the evaluation procedure that has been applied to compare different open-source radiological viewers with the proposed one. In Section 3, the main functionalities of Starviewer platform are illustrated and also the results obtained from the comparison of selected open-source radiological viewers. The obtained results are discussed in Section 4.

2. Materials and methods

2.1. Starviewer architecture

Starviewer was conceived as an integrable open-source solution for

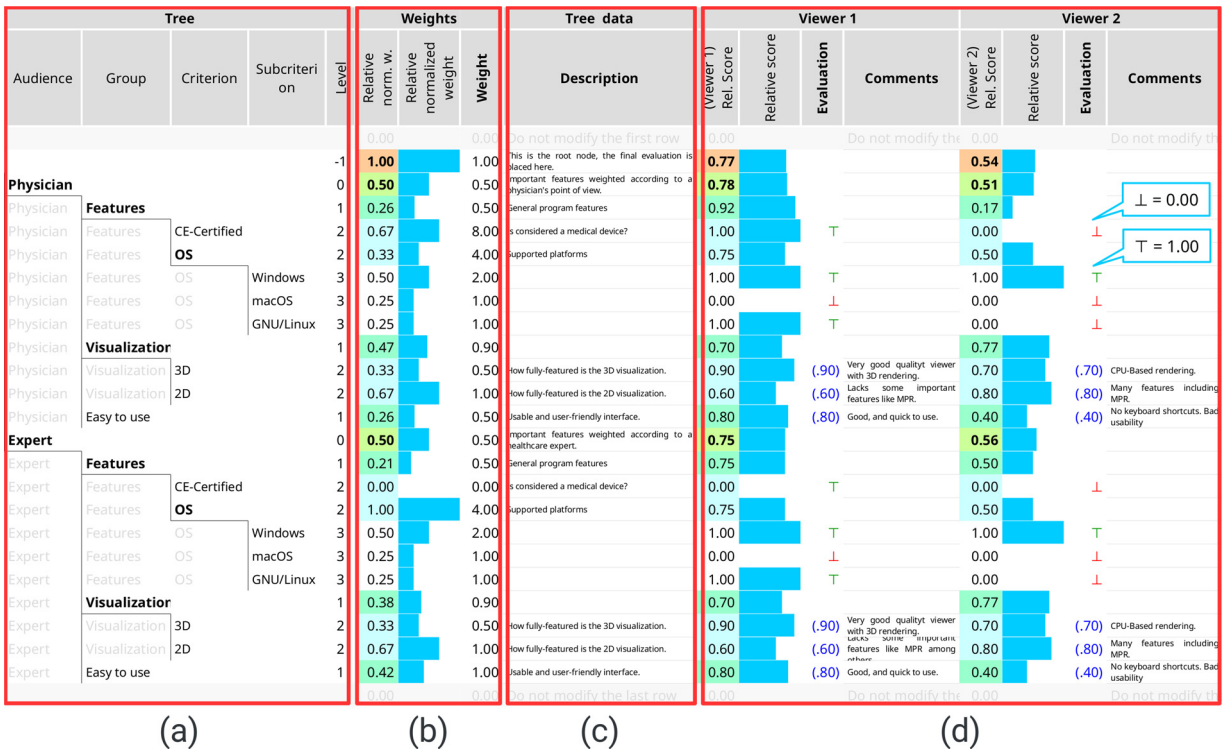

Fig. 2. The main blocks of the proposed evaluation framework where several auxiliary columns are hidden for the sake of readability. (a) Audiences, groups, criteria, and subcriteria are organized hierarchically; the tree structure is drawn automatically from the entered data. (b) The user enters weights in the bold column for each row in the tree; relative normalized weights and graph bars are automatically computed. (c) Description for each node of the tree. (d) Viewer scores: for each viewer a set of columns is added; the user enters the score in the "evaluation" column for each leaf node; relative scores and graph bars are automatically computed.

DICOM-based diagnosis. To design it we considered fundamental to support basic and advanced functionalities and also the possibility to extend the viewer in order to assist specific workflows required by experts. We designed the three-level architecture illustrated in Fig. 1 and described below.

2.1.1. External libraries

The first level of the Starviewer architecture contains the open-source libraries used by the platform. In particular, Qt [34] is used as the development framework, and VTK [35] and ITK [36] for image representation, processing, visualization, interaction, and rendering. DCMTK [37] is used to communicate with the PACS and as the primary choice when reading DICOM files. GDCM [38] is used by the anonymization feature and as a surrogate when reading unsupported DCMTK image codecs such as JPEG2000 (which is not included in the free version of DCMTK). ThreadWeaver [39], which simplifies the implementation of multi-threaded asynchronous algorithms using a thread pool, is used for image loading, PACS queries, downloads, and uploads. Finally, Breakpad [40] is a library used to create crash reports with memory dumps that are stored locally and then sent to our servers for further analysis.

2.1.2. Core

The second level of the Starviewer architecture has two main components, the core and the core user interface. The Starviewer core is made of algorithms, common user interfaces, and general purpose modules that conform the structure which can later be leveraged by the extensions. It has four main parts.

First, *input/output* modules and user interfaces consisting of: (i) the *local database*, (ii) the *PACS connection* module, and (iii) the *image loading* module.

Second, the *extensions engine*, that provides some classes and specifies a structured methodology to extend the program functionalities in order to satisfy new user requirements.

Third, *widgets* that can be divided in: (i) *2D widget*, focused on visualizing plain images, slices from volumes, and fusions between series; (ii) *3D widget*, targeting the visualization and rendering of volumetric datasets; (iii) *layouts*, whose task is to arrange a set of viewer widgets; and (iv) *hanging protocols*, which are a set of user-definable templates aware of the study context that have rules to decide when a hanging protocol is applicable, define a layout with custom-sized viewers, and have sets of rules to determine which contents shall be placed in each viewer.

Fourth, *tools*, a wide collection of components that can be associated with the basic 2D and 3D viewer widgets to add interaction and extend their functionality. They are used to implement or reuse extension-specific features and are divided in six categories: (i) *navigation* tools, (ii) *transfer function* tools, (iii) *thick slab* tool, (iv) *localization* tools, (v) *drawing* tools, and (vi) *synchronization* tools.

2.1.3. Core user interface

Besides core, in the second level of the architecture there is also the core user interface. This exposes common features that are not implemented by extensions but are required by them. Its main elements are: (i) the *main window*, (ii) the *configuration window*, and (iii) the *study management* window.

2.1.4. Extensions

In the last level of the architecture, there are the extensions. These make use of the underlying tools, widgets, and the Starviewer platform in general to develop specialized tools for specific use cases with custom user interfaces. An extension is implemented when a user requires a set of functionalities to deal with a determined protocol or situation. Although the functionalities of extensions are different, for the sake of simplicity for the user, they usually follow a standard layout pattern with two distinct regions: a *toolbar*, with shortcuts to main

functionalities, camera options, and tools; and a *views area* containing a layout of viewers, which can be 2D or 3D, or other specialized widgets to display information. A detailed description of how to create an extension is given in Appendix A.

2.2. The evaluation framework

Usually, software feature comparisons consist in gathering user needs, representing them as a set of criteria, and grading them following a basic methodology [3,13] or a more complex one which varies the requirements, adapts the weights, and uses fine-grained criteria [41,42]. Our evaluation framework expands on some of these ideas using a very fine-grained tree of hierarchically organized criteria where related features are grouped and scored. Every node of the tree is assigned a weight according to its relative importance in its context (sibling nodes); weights are automatically normalized at each context (sibling nodes) to sum 1. Then, each leaf node is evaluated and given an score in the range [0, 1] being the parent node score the weighted sum of them.

The proposed framework is presented in Fig. 2 where its main blocks have been highlighted. By default the evaluator has a template where the four first columns (see block (a)) represent the levels of the tree structure, where from left to right, each column represents: (i) *audiences*; (ii) *groups*, a high-level grouping of criteria not meant to be directly evaluated; (iii) *criteria*, specific items to evaluate; and (iv) *subcriteria*, present when a criterion is finely evaluated. The evaluator may alter the tree by adding, removing, or rearranging criteria, and adapting the weights of the rows to express his or her needs. Weights are presented in block (b). The user enters the weights in the last column and they are automatically normalized and shown both numerically and graphically in the first and second columns. In addition, there is a column to enter a textual description of the criteria (see block (c)). Finally, in block (d), the software products (in our case DICOM viewers) to be compared are placed. The evaluator enters the score in the "evaluation" column for each leaf node (criterion or subcriterion), and the relative score and graph bar are automatically computed. If a new software solution has to be evaluated, a new group of columns must be appended at the end.

2.3. The evaluation procedure

2.3.1. Evaluation framework definition

To fill the framework with criteria and weights specific to DICOM viewers, we restricted our study to functional requirements which where compiled from: (i) clinical experts feedback; (ii) our knowledge with the development of Starviewer; (iii) existing literature [3,13,26]; and (iv) a collection of the analysed viewers features. Although they are equally important, non-functional requirements such as security, privacy, availability, platform usability, or efficiency were not evaluated. However, they can be easily integrated in the framework. To select functional requirements we evaluated state of the art DICOM viewers (see Section 2.3.2) and we elaborated a list with their functionalities. We grouped these functionalities in seven groups: (i) *technical* features, (ii) *archive*, (iii) *workflow*, (iv) *visualization* and interaction with a volume or set of images, (v) *tools*, (vi) *other modalities*, and (vii) *processing* features. Then, we collected this information in a survey and asked radiologists to grade them according to importance in a range from 0 to 10. The survey includes a question asking the number of years each respondent has been diagnosing from medical images and their area of diagnosis. The survey was sent to 25 radiologists of different centers of Catalan Hospitals. From their answers, we created a first list of criteria and weights which were obtained from the mean grades. List items were transformed to the evaluation framework criteria. In addition, from the collected information we decided to consider three broad types of audiences: (i) *general*, which reflects the common needs of a physician, not exclusively focused on 2D or 3D; (ii) *2D-based* audience, that works with plain images like ultrasound, mammography, and X-ray; and (iii) *3D-based* audience whose datasets are mainly volumetric images from computed tomography (CT) and magnetic resonance (MRI) as well as more specific nuclear medicine modalities like positron emission tomography (PET) or single-photon emission computed tomography (SPECT). In the evaluation, these three audiences have the same criteria and evaluations, but their weights have been tuned to more accurately reflect their particular audience needs.

The filled framework was supervised by a group of experienced radiologists via personal interviews. They were asked to pinpoint missing criteria and check and fine-tune the weights. From this process we obtain the filled evaluation framework presented in Section 3.2.

2.3.2. Selection of viewers to compare

To select the open-source software candidates, a search focused on viewers for daily clinical practice was carried out. The search was conducted using: general purpose search engines like DuckDuckGo and Google; Debian, Fedora, and Ubuntu package archives; specialized websites like IDoImaging.com [43] and AlternativeTo.net [44]; and popular code forges like GitHub [45], Bitbucket [46], GitLab [47], and SourceForge [48].

A list of 54 candidates was produced from which viewers being unmaintained or not achieving the minimum requirements for daily clinical practice were discarded. This reduces the list to eleven pieces of software (Aliza [49], Amide [50], Ginkgo CADx [51], Horos [17], InVesalius [52], medInria [6], MIPAV [15], MITK [8], Slicer [5], Starviewer, and Weasis [16]) that were installed in a multi-platform testing environment with a PACS server for further analysis. After a first screening, six of them (Aliza, Amide, InVesalius, MIPAV, MITK, and Slicer), all excellent tools, were discarded because they were more research and processing focused. This leads to the final list with Ginkgo CADx, Horos, medInria, Starviewer, and Weasis.

Once the viewers to compare were selected and the evaluation framework was filled with criteria and weights, a group of three experts with experience on radiological viewers but not involved in the development of Starviewer carried out the comparison. The evaluation was performed taking in consideration the description of each item combined with the need to provide an individual justification for the more qualitative items. Each expert carried the evaluation individually; many items could be evaluated in a boolean fashion, however when this could not be done a grading between 0 and 1 was used instead. In a second stage the grades were shared in order to discuss a final consensus.

3. Results

3.1. Starviewer as a platform

Starviewer is developed in C++11 with support from several open-source libraries: Qt, DCMTK, ITK, VTK, GDCM, ThreadWeaver, and Breakpad. It is supported in Windows, macOS, and GNU/Linux operating systems. The source code is available at http://www.starviewer.org and https://github.com/starviewer-medical/starviewer licensed under GPLv3+ terms.

Starviewer core is extended with five stable extensions: *2D viewer* (with support for fusion), *3D viewer*, *MPR* (multi-planar reconstruction), *DICOM Print*, and *PDF*. To illustrate their functionalities, some screenshots are shown in Fig. 3.

Apart from the stable extensions, there are some experimental ones that are examples of research focused extensions: (i) *diffusion perfusion segmentation*; (ii) *edema segmentation*, which calculates the volume of automatically detected lesions in the brain (edemas and hematomas); and (iii) *rectum segmentation*. These extensions are not included in regular compiled binaries, but are available in the source code.

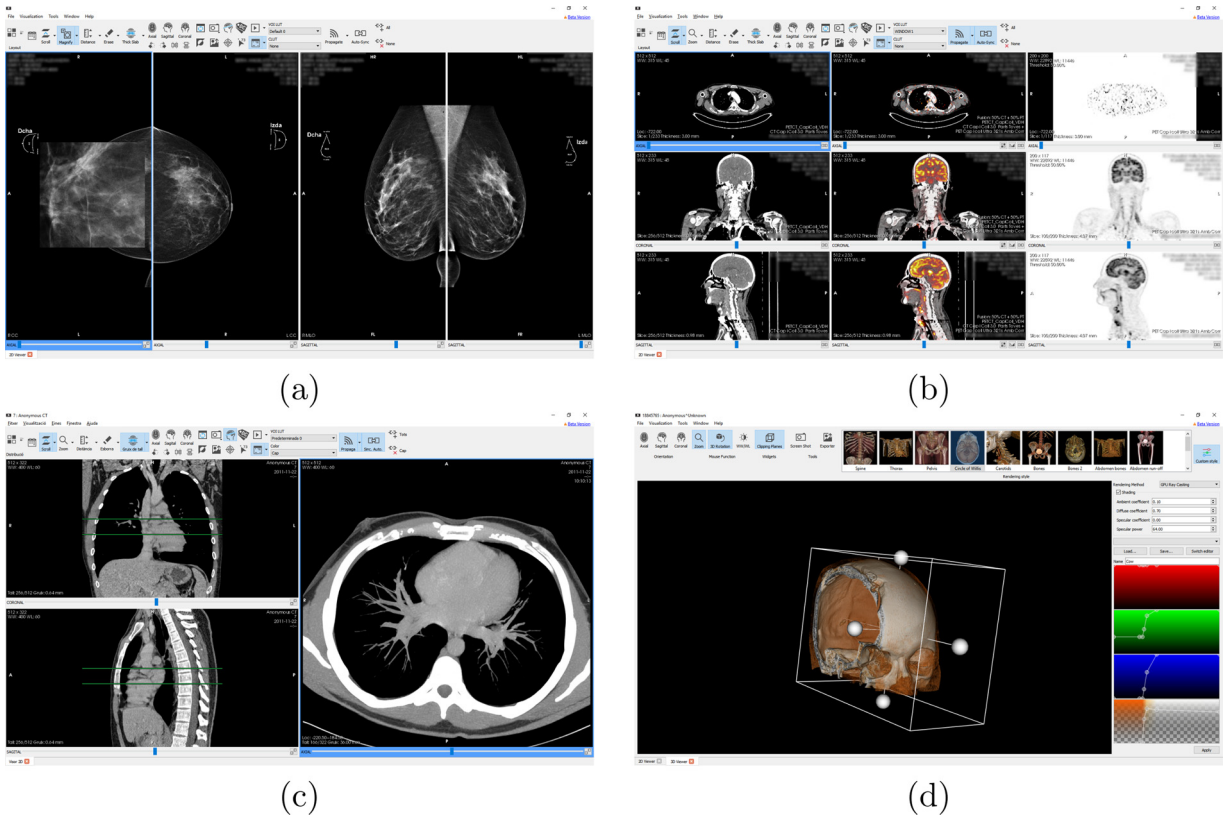

Fig. 3. Screenshots showing several features of Starviewer. (a) 2D viewer with the magnifying glass tool. (b) 2D viewer visualizing a PET-CT with axial, coronal and sagittal reconstructions in the rows combined with the CT, PET-CT fusion and PET modalities in the columns. (c) 2D viewer showing the reference lines tool with thick slab enabled; the green lines in the coronal and sagittal viewers show the location of the axial view and its thickness. (d) 3D viewer with clipping planes enabled and a sidebar to modify the current transfer function and illumination parameters. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.2. Evaluation

The results have to be interpreted as the capability of a viewer to address specific audience needs. A weighted and fine grained bottom-up evaluation of the features is performed in order to reduce the bias. Depending on the audience some needs may be more important than others. In the presented evaluation we expressed three broad types of audience that try to mimic the daily clinical practice needs.

The proposed evaluation framework has been implemented in a macro-free LibreOffice spreadsheet that automatically draws the tree, displays warnings, and has built-in help, among other advanced features. It is available as supporting material, and in https://github.com/starviewer-medical/dicom-viewers-comparison.

In Fig. 4 the results of the comparison of Starviewer with Ginkgo CADx, Horos, medInria, and Weasis using the proposed evaluation framework is presented. For a better comprehension, these results are shown only up to group level where each value is obtained from a weighted sum of the scores present at lower levels. The full evaluation can be seen in the spreadsheet in the supporting material, as illustrated in Fig. 5, where the “tools” group and the “annotations” criterion have been expanded.

From Fig. 4 we can observe that for a general audience the best score is obtained by Horos with 0.78, while Weasis and Starviewer have a similar score, 0.62 and 0.61, respectively. The worst positions are for Gingko CADx, and medInria, with 0.43 and 0.38, respectively. Focusing on 2D users, Ginkgo CADx (0.52) and Weasis (0.68) significantly increase their scores relative to the general audience. The other viewers show a slight increase, relegating Starviewer (0.63) to the third position. Finally, focusing on 3D users, rankings stay the same as in the general audience, but more 2D oriented viewers, Ginkgo CADx (0.37) and Weasis (0.56), decrease their scores significantly.

In addition, Fig. 6 shows the summarized results for the general audience in a graphical form. The right “weights” bar shows the maximum possible score for each category, thus helping to see where each viewer is good or has room for improvement.

4. Discussion

In recent years, open-source software has emerged as a potential solution to compete with privative products in medical applications. These products have many similarities in their approach but also present great differences in provided features and functionalities. For this reason, despite the interest from the research and medical community on open-source products, the lack of clear information about the advantages and disadvantages of one product over the others, makes their use in real scenarios quite complex. Aware of this situation, the aim of this paper has been two-fold. Firstly, we have presented Starviewer, a new open-source multi-platform DICOM viewer which despite its ability to satisfy and adapt to daily clinical needs has room for some feature improvements as described below. Secondly, we have introduced an evaluation framework that reduces the complexity of selecting a viewer that can: (i) be user-configured and adapted to any user specific scenario, and (ii) be used as a complement to maturity models or quality evaluation frameworks [33] in order to perform a feature analysis.

In our case the evaluation carried out presents two main limitations. Firstly, a reduced set of viewers has been evaluated because the evaluation of each viewer in such level of detail takes a considerable effort. However, with the proposed framework it is easy to add new viewers to the comparison. Secondly, it is difficult to attain a fully objective

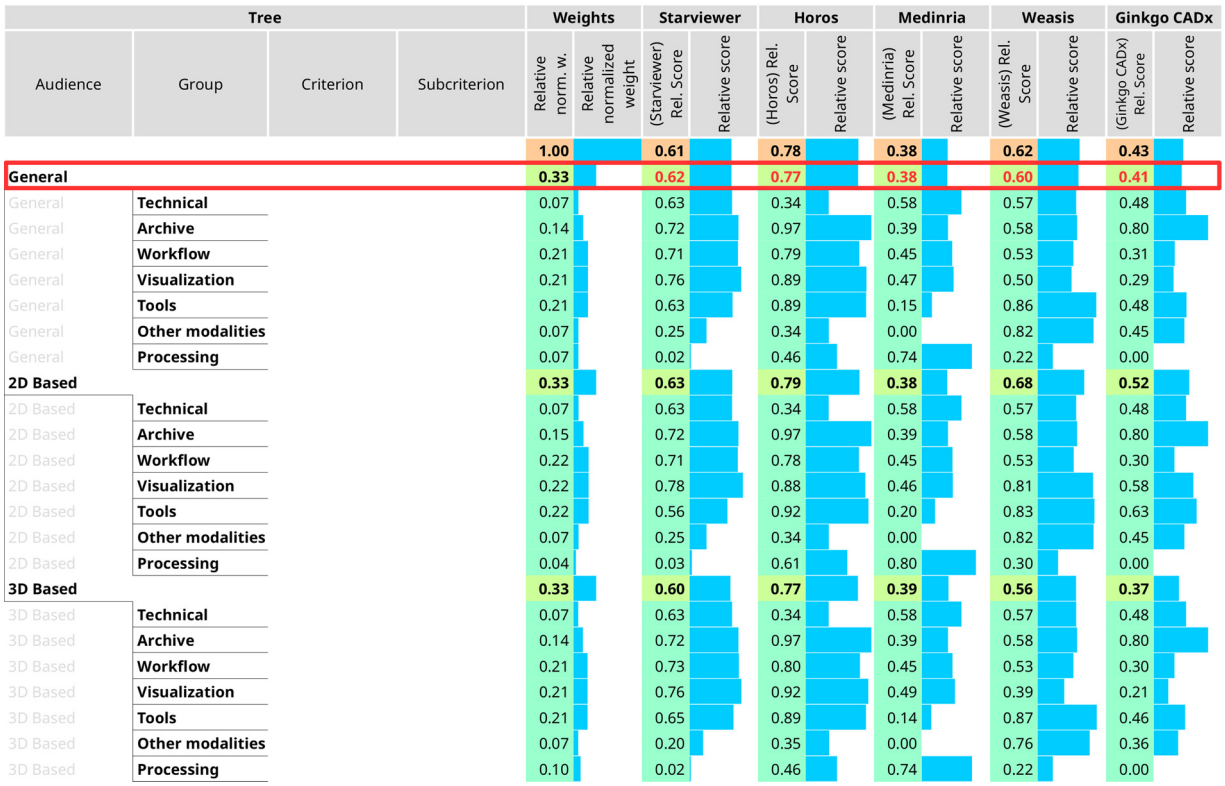

Tree				Weights		Starviewer		Horos		Medinria		Weasis		Ginkgo CADx	
Audience	Group	Criterion	Subcriterion	Relative norm. w.	Relative normalized weight	(Starviewer) Rel. Score	Relative score	(Horos) Rel. Score	Relative score	(Medinria) Rel. Score	Relative score	(Weasis) Rel. Score	Relative score	(Ginkgo CADx) Rel. Score	Relative score
				1.00		0.61		0.78		0.38		0.62		0.43	
General				0.33		0.62		0.77		0.38		0.60		0.41	
General	**Technical**			0.07		0.63		0.34		0.58		0.57		0.48	
General	**Archive**			0.14		0.72		0.97		0.39		0.58		0.80	
General	**Workflow**			0.21		0.71		0.79		0.45		0.53		0.31	
General	**Visualization**			0.21		0.76		0.89		0.47		0.50		0.29	
General	**Tools**			0.21		0.63		0.89		0.15		0.86		0.48	
General	**Other modalities**			0.07		0.25		0.34		0.00		0.82		0.45	
General	**Processing**			0.07		0.02		0.46		0.74		0.22		0.00	
2D Based				0.33		0.63		0.79		0.38		0.68		0.52	
2D Based	**Technical**			0.07		0.63		0.34		0.58		0.57		0.48	
2D Based	**Archive**			0.15		0.72		0.97		0.39		0.58		0.80	
2D Based	**Workflow**			0.22		0.71		0.78		0.45		0.53		0.30	
2D Based	**Visualization**			0.22		0.78		0.88		0.46		0.81		0.58	
2D Based	**Tools**			0.22		0.56		0.92		0.20		0.83		0.63	
2D Based	**Other modalities**			0.07		0.25		0.34		0.00		0.82		0.45	
2D Based	**Processing**			0.04		0.03		0.61		0.80		0.30		0.00	
3D Based				0.33		0.60		0.77		0.39		0.56		0.37	
3D Based	**Technical**			0.07		0.63		0.34		0.58		0.57		0.48	
3D Based	**Archive**			0.14		0.72		0.97		0.39		0.58		0.80	
3D Based	**Workflow**			0.21		0.73		0.80		0.45		0.53		0.30	
3D Based	**Visualization**			0.21		0.76		0.92		0.49		0.39		0.21	
3D Based	**Tools**			0.21		0.65		0.89		0.14		0.87		0.46	
3D Based	**Other modalities**			0.07		0.20		0.35		0.00		0.76		0.36	
3D Based	**Processing**			0.10		0.02		0.46		0.74		0.22		0.00	

Fig. 4. Summarized results of the comparison of Starviewer with Horos, medInria, Weasis, and Ginkgo CADx considering three different audiences (general, 2D-based, and 3D-based), and criteria grouped in seven categories. Criteria and subcriteria are hidden for comprehensibility. The final scores for the general audience (highlighted in red), are the ones we consider more relevant as they are in the most transversally weighted audience. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

analysis, but the hierarchical evaluation methodology and the different audiences approach forces the evaluator to perform a bottom-up analysis that helps to limit the subjectivity when setting scores, weights, and criteria. Scores were set according to the descriptions, but in some cases the assessment requires analysing and justifying the pros and cons of each candidate to set a fair final score, which may introduce some bias. Regarding weights, they were set following expert radiologists advice, but in a future work, we will consider making a survey in order

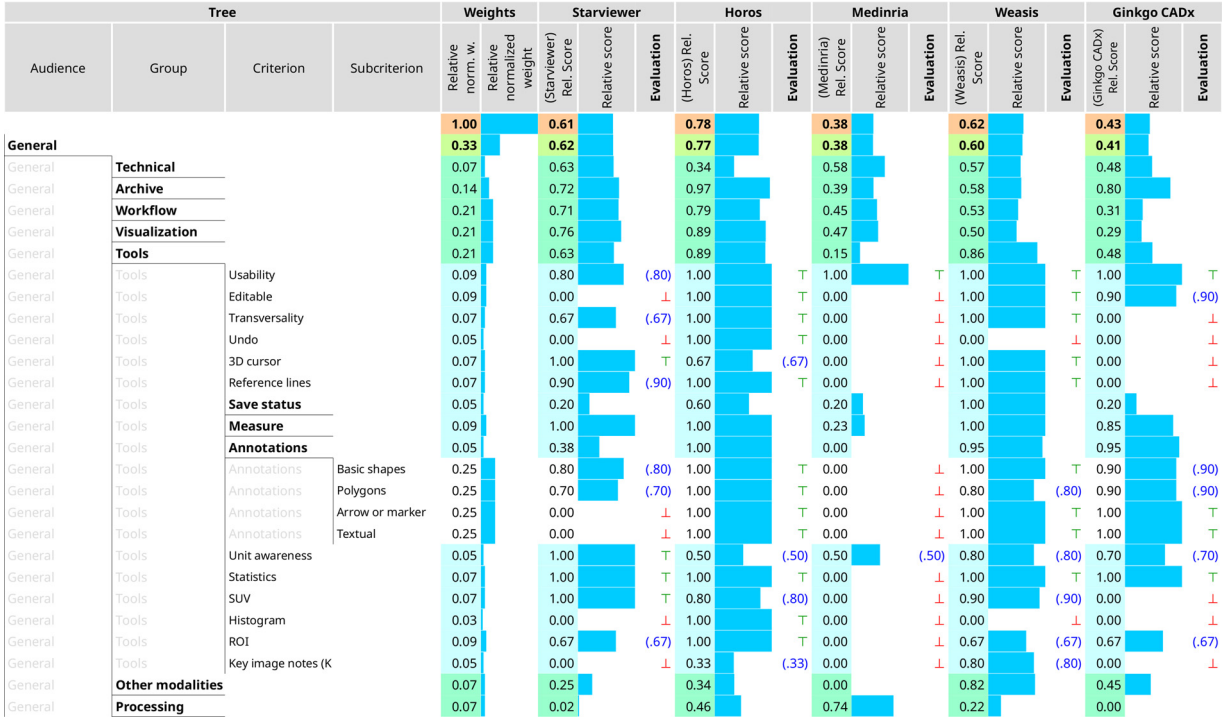

Tree				Weights		Starviewer			Horos			Medinria			Weasis			Ginkgo CADx		
Audience	Group	Criterion	Subcriterion	Relative norm. w.	Relative normalized weight	(Starviewer) Rel. Score	Relative score	Evaluation	(Horos) Rel. Score	Relative score	Evaluation	(Medinria) Rel. Score	Relative score	Evaluation	(Weasis) Rel. Score	Relative score	Evaluation	(Ginkgo CADx) Rel. Score	Relative score	Evaluation
				1.00		0.61			0.78			0.38			0.62			0.43		
General				0.33		0.62			0.77			0.38			0.60			0.41		
General	**Technical**			0.07		0.63			0.34			0.58			0.57			0.48		
General	**Archive**			0.14		0.72			0.97			0.39			0.58			0.80		
General	**Workflow**			0.21		0.71			0.79			0.45			0.53			0.31		
General	**Visualization**			0.21		0.76			0.89			0.47			0.50			0.29		
General	**Tools**			0.21		0.63			0.89			0.15			0.86			0.48		
General	Tools	Usability		0.09		0.80		(.80)	1.00		⊤	1.00		⊤	1.00		⊤	1.00		⊤
General	Tools	Editable		0.09		0.00		⊥	1.00		⊤	0.00		⊥	1.00		⊤	0.90		(.90)
General	Tools	Transversality		0.07		0.67		(.67)	1.00		⊤	0.00		⊥	1.00		⊤	0.00		⊥
General	Tools	Undo		0.05		0.00		⊥	1.00		⊤	0.00		⊥	0.00		⊥	0.00		⊥
General	Tools	3D cursor		0.07		1.00		⊤	0.67		(.67)	0.00		⊥	1.00		⊤	0.00		⊥
General	Tools	Reference lines		0.07		0.90		(.90)	1.00		⊤	0.00		⊥	1.00		⊤	0.00		⊥
General	Tools	**Save status**		0.05		0.20			0.60			0.20			1.00			0.20		
General	Tools	**Measure**		0.09		1.00			1.00			0.23			1.00			0.85		
General	Tools	**Annotations**		0.05		0.38			1.00			0.00			0.95			0.95		
General	Tools	Annotations	Basic shapes	0.25		0.80		(.80)	1.00		⊤	0.00		⊥	1.00		⊤	0.90		(.90)
General	Tools	Annotations	Polygons	0.25		0.70		(.70)	1.00		⊤	0.00		⊥	0.80		(.80)	0.90		(.90)
General	Tools	Annotations	Arrow or marker	0.25		0.00		⊥	1.00		⊤	0.00		⊥	1.00		⊤	1.00		⊤
General	Tools	Annotations	Textual	0.25		0.00		⊥	1.00		⊤	0.00		⊥	1.00		⊤	1.00		⊤
General	Tools	Unit awareness		0.05		1.00		⊤	0.50		(.50)	0.50		(.50)	0.80		(.80)	0.70		(.70)
General	Tools	Statistics		0.07		1.00		⊤	1.00		⊤	0.00		⊥	1.00		⊤	1.00		⊤
General	Tools	SUV		0.07		1.00		⊤	0.80		(.80)	0.00		⊥	0.90		(.90)	0.00		⊥
General	Tools	Histogram		0.03		0.00		⊥	1.00		⊤	0.00		⊥	0.00		⊥	0.00		⊥
General	Tools	ROI		0.09		0.67		(.67)	1.00		⊤	0.00		⊥	0.67		(.67)	0.67		(.67)
General	Tools	Key image notes (K		0.05		0.00		⊥	0.33		(.33)	0.00		⊥	0.80		(.80)	0.00		⊥
General	**Other modalities**			0.07		0.25			0.34			0.00			0.82			0.45		
General	**Processing**			0.07		0.02			0.46			0.74			0.22			0.00		

Fig. 5. Expanded evaluation results for the “tools” group and the “annotations” criterion. The user can interact to determine the information to be presented.

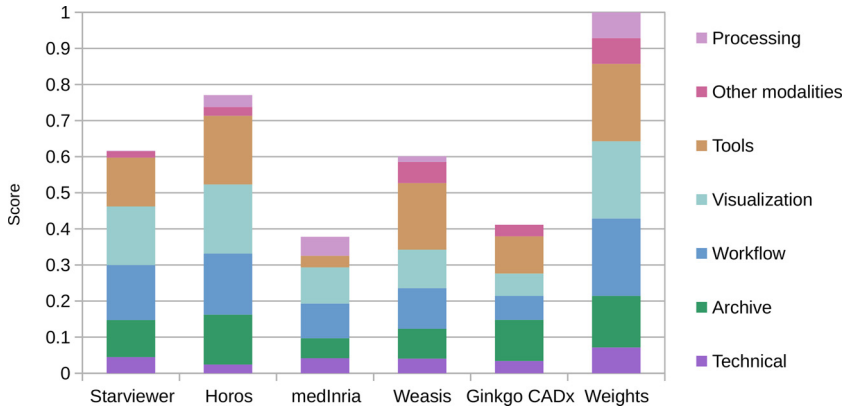

Fig. 6. Graph showing the summarized results of the comparison of Starviewer with the other viewers for the general audience. The right "weights" bar shows the maximum possible score for each category.

to get better measurable knowledge and contemplate more specific audiences. Fortunately, all these limitations can be overcome thanks to the adaptability of this evaluation framework by the end user.

The proposed viewer has been compared using the presented evaluation framework with state of the art open-source viewers such as Ginkgo CADx, Horos, medInria, and Weasis. Criteria have been grouped in seven categories and weighted according to user perspective. Focusing on the general audience, which takes into account all possible features, the obtained results are discussed below.

In the *technical* group, where aspects such as platform support, multi-language, and documentation are considered, Starviewer has the top score mainly due to the completeness of its documentation. Starviewer provides an extensive user and administrator documentation where all features and shortcuts are described while many of the analysed viewers do not provide it.

Regarding the *archive* category, where network protocols, local storage, and supported formats are considered, the maximum score is achieved by Horos, which supports almost all the features. The main limitations of Starviewer in this group are the lack of WADO support and its inability to convert from DICOM to non-DICOM formats and vice versa; for these reasons it is ranked in third position after Ginkgo CADx.

In the *workflow* group, where we consider features that improve user experience such as program customization, hanging protocols support, and synchronization, Horos is again the top-rated viewer mainly because it is fully customizable. Starviewer follows close in second position, being the best in synchronization and study comparison support; however it lacks a graphical hanging protocol editor.

Visualization is the group where all the 2D and 3D visualization and interaction features, and DICOM conformity are considered. Horos is first with great support for most of the criteria. Starviewer obtains a good score and is placed second, being its main drawbacks the lack of non-Latin characters support and the absence of full-screen visualization and true scaling (display of objects on screen in real size).

In the *tools* category is where we look for functionalities to measure, annotate, and comprehend datasets. Horos and Weasis obtain a very good score, with Horos slightly better due to its ability to undo actions. Starviewer is ranked third, with its main weaknesses being non-editable tools, non-savable annotations, and the lack of key image notes support.

In the *other modalities* group, where support for non-image modalities like structured reports and encapsulated documents is evaluated, the best score is achieved by Weasis thanks to its support for all special modalities. Starviewer is fourth because it only supports encapsulated documents.

Finally, in the *processing* group, which evaluates features that perform a complex analysis and usually generate new series, the top position is for medInria thanks to its extensive support for almost all the considered criteria. Starviewer obtains a bad score in this category with only some points in the segmentation criterion through the "magic ROI" tool, putting it in fourth position.

Focusing in the 2D-based audience, Weasis greatly increases its score ascending to the second position after Horos, which remains first. The main reason are the changes of the weights in the *visualization* group, where criteria that do not make sense in 2D, such as thick slab, MPR, fusion, and 3D, have their weights set to zero, thus increasing the relative weights of the other criteria.

Similarly, for the 3D-based audience, the weights of *visualization* criteria have been adjusted by increasing the ones that are more important for 3D datasets exploration. In addition the weight of the *processing* group has been slightly increased. In this case the positions stay the same as for the general audience but Weasis and Gingko CADx have significantly lower scores due to their lack of 3D support.

In our evaluation medInria has the lowest score. This is due to the fact that although it is an excellent processing tool it is not suitable for the daily clinical practice. From this evaluation we have seen that Horos is the best viewer overall, however it is only available for the macOS platform. In second position comes Starviewer, which despite having less features than Horos, is still a good viewer for daily clinical needs and has the advantage of being multi-platform. In addition, Starviewer is modular and extensible, thus new extensions and functionalities can be quickly developed according to user demands. This is a technically relevant characteristic to be taken into account, although it has not been considered in the evaluation because it is not as relevant from the end user perspective. Other viewers that have not been evaluated, like MITK [7] or 3D Slicer [4], provide specific APIs to manipulate and extend functionalities. In contrast, Starviewer extensions operate directly using VTK, reducing the degree of abstraction and enhancing VTK-based code portability.

From the aforementioned analysis, one can conclude that Starviewer main advantages are technical documentation, productive workflow, and extensibility. Main disadvantages are lack of WADO support, DICOM conversions, lack of hanging protocol editor, Unicode character support, non-editable tools, key image notes and lack of special modalities like EKG. All those shortcomings will be addressed in our future developments.

We are conscious that medical imaging software paradigm is changing to software-as-a-service where solutions run in the cloud. However there are still certain technological issues that need further development to provide the advantages of desktop applications [53]. On the one hand, implementing certain features such as interactive oblique reconstructions and 3D volume rendering on the client side require that all the images are available in the client memory and that requires powerful enough hardware to perform such operations, thus negating the advantage of requiring only cheap thin clients. On the other hand, if these and other features are implemented on the server

side, the viewer is highly dependent on and limited by the network bandwidth, latency and stability. In addition, a high client concurrency can lead to performance degradation in the server [54]. For these reasons, we consider that Starviewer can be of interest to the health community.

Summary points

What was already known on the topic:

- DICOM viewers are essential in many clinical processes, either to diagnose, plan operations, or follow up the evolution of pathologies.
- The interest in open-source DICOM viewers has increased considerably.
- Selecting the DICOM viewer that better fits user needs can be a complex task.

What this study added to our knowledge:

- A new open-source multi-platform DICOM viewer that supports main needs of image diagnosis experts and the creation of custom extensions for specific workflows.
- An adaptable evaluation framework that considers different audiences and criteria weighted according to user needs and that can be used as a complement to maturity models and quality evaluation frameworks.
- An application example where Starviewer and state-of-the art open-source DICOM viewers are compared.

Informed consent

All anonymous subjects depicted in this paper signed an informed consent form allowing the use of the images for research purposes.

Authors' contributions

All authors have contributed equally.

Declarations of interest

None declared.

Acknowledgements

We want to thank all developers who contributed to Starviewer as well as others which have contributed to the project in some manner. We also want to thank Institut de Diagnòstic per la Imatge (IDI), our partner since 2004. This work has been funded in part by grants from the Spanish Government (Nr. TIN2016-75866-C3-3-R) and from the Catalan Government (Nr. 2017-SGR-1101).

Appendix A. Workflow to create a Starviewer extension

Focusing on final users, the main advantage of Starviewer is the possibility to create custom extensions for specific workflows. Currently, Starviewer only supports extensions as static libraries, thus one needs to be able to build Starviewer from source. Once the build environment is set up, creating an extension is very simple as illustrated in the following example.

Extensions must be placed under the `src/extensions` directory, and there are three possible directories to choose from: (i) `main`, intended for official stable extensions; (ii) `contrib`, intended for stable third-party contributed extensions; and (iii) `playground`, intended for unstable or experimental extensions. Thus, the example extension is a nice fit for `playground`.

The workflow to create an extension can be summarized in four steps: (i) create the extension subproject; (ii) create the bare minimum classes, including the UI; (iii) generate the translation files and add them to a resource file; and (iv) implement the actual functionality for the extension. Now we will extend on each of these steps.

Note that in this section we will include only a few code snippets which have been simplified for the sake of readability, omitting include guards, include directives, namespace declarations, and safety checks, and occasionally merging header and source files. Full code listings can be found in the appendix and in the supporting material.

A.1 Extension subproject

The first step is creating a subproject for the extension. We have to create a directory called `example` under `src/extensions/playground`. Inside this directory we must create a `qmake` project file named `example.pro` with this content:

```
EXTENSION_DIR = $$PWD
include(../../basicconfextensions.pri)
```

Finally, we have to add a reference to this new subproject in parent projects. Specifically, we must add `example` to the `SUBDIRS` variable in `src/extensions/playground/playground.pro` and the `PLAYGROUND_EXTENSIONS` variable in `src/extensions.pri`.

A.2 Minimum code

The bare minimum code required for an extension consists of the main extension class and a special class that is used by the core to initialize the extension.

All new files from now on will be placed inside the `example` directory.

By convention, all UI classes start with a 'Q' and all the code is placed in the `udg` namespace.

First, we create an UI file named `qexampleextensionbase.ui` which, for now, is just a `QWidget` containing a `QLabel` with the text "Example". We name the main widget "`QexampleExtensionBase`", and may be created using the Qt Designer graphical editor or the integrated one inside Qt Creator.

Then we create the main extension class named `QExampleExtension` which inherits `QWidget` publicly and `Ui::QExampleExtensionBase` privately. The latter is generated by `uic`, the Qt UI compiler, from the previously created `.ui` file. This class has just a constructor that calls `setupUi`(**this**).

Now, we need to create a special class named `ExampleExtensionMediator`. It is a subclass of `ExtensionMediator` and must implement two methods to initialize the extension and to give it a unique ID. In the header file, we also need to create a static instance of `InstallExtension` to make the extension known to Starviewer. This is all done in the following snippet:

```
class ExampleExtensionMediator : public ExtensionMediator {
    Q_OBJECT
public:
    explicit ExampleExtensionMediator(QObject *parent = nullptr) {}
    bool initializeExtension(QWidget *extension,
                             const ExtensionContext &extensionContext) override {
        QExampleExtension *exampleExtension;
        if (!(exampleExtension = qobject_cast<QExampleExtension*>(extension))) {
            return false;
        }
        return true;
    }
    DisplayableID getExtensionID() const override {
        return DisplayableID("ExampleExtension", tr("Example"));
    }
};
static InstallExtension<QExampleExtension, ExampleExtensionMediator>
```

Finally, we add both classes and the UI file to the project file `example.pro`.

A.3 Translations

Translation files must be generated and added to the project, even if they are empty. To generate these files we have to execute `lupdate example.pro` from the `example` directory; this will create a few `.ts` files. Then, we create a resource file `example.qrc` that includes the compiled translations (which will be created during the build from the `.ts` files, with the same names but with the `.qm` extension) under the prefix `"/extensions/ExampleExtension"`. The last part of the prefix must match the ID given in the extension mediator. Finally, we add the resource file to the `RESOURCES` variable in `example.pro`.

After this step, we can build Starviewer and try the new extension, as illustrated at Fig. A.1(a). It can be launched by selecting *Example* in the *Visualization* menu after opening a study.

A.4 Adding functionality

Having a minimal extension working, the next step is adding actual functionality to it. In this example we will add a 2D and a 3D viewer. Viewers can be added in the graphical editor as `QWidget`s and then promoted to `Q2DViewerWidget` and `Q3DViewer`, respectively.

Then we add a `setPatient` method to `QExampleExtension` with the following code:

```
void QExampleExtension::setPatient(Patient *patient) {
    m_2DViewer->setInputAsynchronously(patient->getVolumesList().first());
    m_3DViewer->setInput(patient->getVolumesList().first());
}
```

Finally we call this method from `ExampleExtensionMediator`, getting the patient from the `ExtensionContext`:

```
exampleExtension->setPatient(extensionContext.getPatient());
```

If we build and launch the extension in this moment, we can already see both viewers showing the first series of the opened study, as seen in Fig. A.1(b). Both viewers allow to independently change the series using the right click menu. The 2D viewer allows changing slices with a slider, but the

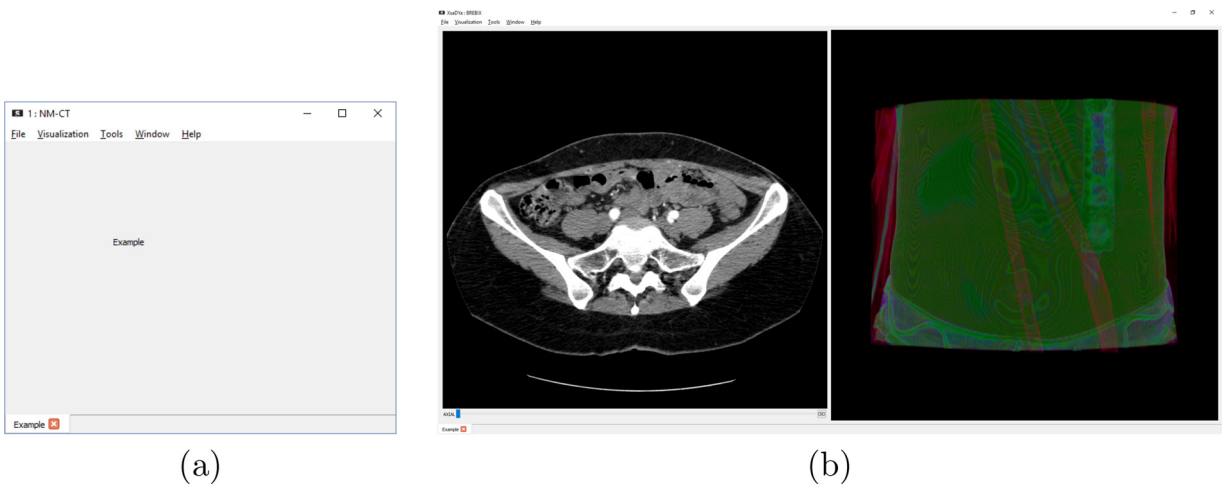

(a) (b)

Fig. A.1. Screenshots of the Example extension (a) containing just a label and (b) having a 2D and a 3D viewer.

3D viewer does not support any interactions yet.

To add interactions to the viewers, we need to register and activate several tools in each one. This is done with the help of a `ToolManager`. Since each viewer will have a different set of tools, we will need separated `ToolManager` instances for each one.

We add the new code to the extension's constructor:

```
QExampleExtension::QExampleExtension(QWidget *parent) : QWidget(parent) {
    setupUi(this);

    ToolManager *toolManager2D = new ToolManager(this);
    toolManager2D->registerTool("ZoomTool");            // left button
    toolManager2D->registerTool("TranslateTool");       // middle button
    toolManager2D->registerTool("WindowLevelTool");     // right button
    toolManager2D->registerTool("SlicingKeyboardTool"); // keyboard
    toolManager2D->registerTool("SlicingWheelTool");    // wheel
    toolManager2D->setupRegisteredTools(m_2DViewer->getViewer());
    toolManager2D->triggerTools({"ZoomTool", "TranslateTool", "WindowLevelTool",
                                 "SlicingKeyboardTool", "SlicingWheelTool"});

    ToolManager *toolManager3D = new ToolManager(this);
    toolManager3D->registerTool("ZoomTool");        // left button
    toolManager3D->registerTool("TranslateTool");   // middle button
    toolManager3D->registerTool("Rotate3DTool");    // right button
    toolManager3D->setupRegisteredTools(m_3DViewer);
    toolManager3D->triggerTools({"ZoomTool", "TranslateTool", "Rotate3DTool"});
}
```

If we rebuild and launch the extension again, we can try the new features. In the 2D viewer we can zoom while holding the left button, pan while holding the middle button, adjust the window level or VOI LUT while holding the right button, and change slices with the keyboard arrows and the mouse wheel.

In the 3D viewer we can zoom while holding the left button, pan while holding the middle button, and rotate the image while holding the right button.

Appendix B. Score calculation

The motivation of this appendix is to explain with detail how the actual scores are calculated following the reduced but complete example of Fig. 2.

Note that the actual spreadsheet has many more (hidden) intermediate columns that propagate values up and down in order to accomplish the tree-like structure calculations.

First of all we have to define some concepts: (i) *nodes*, which are rows, and must have a weight; (ii) *leaf nodes*, which are rows without children nodes, are the only directly evaluated items for each analysed software, and their scores always correspond with the evaluation; and (iii) *container nodes*, which have children nodes, are never directly evaluated, and their score is the weighted sum of their direct children scores.

The columns in bold are the ones to be considered user editable. Almost all numeric columns have values ranging from 0 to 1; being an exception the *weights* editable column, that may take any positive number. Then, for each sibling, all weights are summed and normalized to a 0 to 1 range that is placed in the *relative norm. w.* column. The sum of *relative norm. w.* across all siblings is exactly 1. The reason behind this methodology is to make the weight setting process more easy and natural.

Container nodes scores are computed recursively, starting from leaf nodes and ascending up to the root node, according to the following process: (i) each *rel. score* of each analysed software is multiplied by the *relative norm. w.*; (ii) each *rel. score* of each direct child is summed; and (iii) the result of the summation is the *rel. score* that the *container node* takes.

References

[1] National Electrical Manufacturers Association (NEMA), Digital Imaging and Communications in Medicine (DICOM) [Internet], (2018) (cited 08.01.18), https://www.dicomstandard.org/current/.

[2] D.H.K. Huang, PACS and Imaging Informatics: Basic Principles and Applications, 2nd ed., Wiley-Blackwell, 2010.

[3] G. Valeri, F.A. Mazza, S. Maggi, D. Aramini, L. La Riccia, G. Mazzoni, A. Giovagnoni, Open source software in a practical approach for post processing of radiologic images, La Radiol. Med. 120 (3) (2015) 309–323, https://doi.org/10.1007/s11547-014-0437-5.

[4] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, J. Buatti, S. Aylward, J.V. Miller, S. Pieper, R. Kikinis, 3D Slicer as an image computing platform for the quantitative imaging network, Magn. Reson. Imaging 30 (9) (2012) 1323–1341, https://doi.org/10.1016/j.mri.2012.05.001.

[5] 3D Slicer community, 3DSlicer [Internet], (2018) (cited 18.01.18), https://www.slicer.org.

[6] Inria, medInria [Internet], (2018) (cited 16.01.18), http://med.inria.fr.

[7] I. Wolf, M. Vetter, I. Wegner, T. Böttger, M. Nolden, M. Schöbinger, M. Hastenteufel, T. Kunert, H.-P. Meinzer, The medical imaging interaction toolkit, Med. Image Anal. 9 (6) (2005) 594–604, https://doi.org/10.1016/j.media.2005.04.005.

[8] German Cancer Research Center, Division of Medical Image Computing, The Medical Imaging Interaction Toolkit (MITK), [Internet], 2018 (cited 18.01.18), http://mitk.org/wiki/MITK.

[9] Kitware Inc, VolView [Internet], (2018) (cited 16.01.18), https://www.kitware.com/volview/.

[10] IRCAD, VR-Render [Internet], (2014) (cited 22.08.14 (archived), https://web.archive.org/web/20140822230414/http://www.ircad.fr:80/recherche/rd/rd.php.

[11] O. Ratib, A. Rosset, Open-source software in medical imaging: development of OsiriX, Int. J. Comput. Assist. Radiol. Surg. 1 (4) (2006) 187–196, https://doi.org/10.1007/s11548-006-0056-2.

[12] Pixmeo SARL, OsiriX [Internet], (2018) (cited 18.01.18), https://www.osirix-viewer.com/osirix/overview.
[13] D. Haak, C.-E. Page, T.M. Deserno, A survey of DICOM viewer software to integrate clinical research and medical imaging, J. Digit. Imaging 29 (2) (2016) 206–215, https://doi.org/10.1007/s10278-015-9833-1.
[14] M. Mcauliffe, F. Lalonde, D. McGarry, W. Gandler, K. Csaky, B. Trus, Medical image processing, analysis & visualization in clinical research, Proceedings of the 14th IEEE Symposium on Computer-Based Medical Systems, vol. 14 (2001) 381–386.
[15] National Institutes of Health, MIPAV (Medical Image Processing, Analysis and Visualization), [Internet], (2018) (cited 08.01.18), https://mipav.cit.nih.gov/.
[16] dcm4che, Weasis [Internet], (2018) (cited 08.01.18), https://nroduit.github.io/en/basics/architecture/.
[17] The Horos Project, Horos [Internet], (2018) (cited 26.07.18), https://horosproject.org/.
[18] T. Mettler, P. Rohner, Situational Maturity Models as Instrumental Artifacts for Organizational Design, (2009 01), https://doi.org/10.1145/1555619.1555649 http://www.alexandria.unisg.ch/Publikationen/67758.
[19] A. Rocha, Evolution of information systems and technologies maturity in healthcare, Int. J. Healthc. Inf. Syst. Inform. 6 (2011) 28–36, https://doi.org/10.4018/jhisi.2011040103.
[20] R. van de Wetering, R. Batenburg, A pacs maturity model: A systematic meta-analytic review on maturation and evolvability of pacs in the hospital enterprise, Int. J. Med. Inform. 78 (2008) 127–140, https://doi.org/10.1016/j.ijmedinf.2008.06.010.
[21] J. Carvalho, A. Rocha, A. Abreu, Maturity models of healthcare information systems and technologies: a literature review, J. Med. Syst. 40 (2016) 10, https://doi.org/10.1007/s10916-016-0486-5.
[22] U. e Laila, A. Zahoor, K. Mehboob, S. Natha, Comparison of open source maturity models, Proc. Comput. Sci. 111 (2017) 348–354, https://doi.org/10.1016/j.procs.2017.06.033 the 8th International Conference on Advances in Information Technology. http://www.sciencedirect.com/science/article/pii/S1877050917312061.
[23] M.A. Babar, L. Zhu, R. Jeffery, A framework for classifying and comparing software architecture evaluation, Proceedings Australian Software Engineering Conference (ASWEC). (2004) (2004) 309–318.
[24] M.-P. Gagnon, M. Desmartis, M. Labrecque, J. Car, C. Pagliari, P. Pluye, P. Fremont, J. Gagnon, N. Tremblay, F. Légaré, Systematic review of factors influencing the adoption of information and communication technologies by healthcare professionals, J. Med. Syst. 36 (2012) 241–277, https://doi.org/10.1007/s10916-010-9473-4.
[25] P. Nagy, Open source in imaging informatics, J. Digit. Imaging 20 (1) (2007) 1–10, https://doi.org/10.1007/s10278-007-9056-1.
[26] G.L. Presti, M. Carbone, D. Ciriaci, D. Aramini, M. Ferrari, V. Ferrari, Assessment of DICOM viewers capable of loading patient-specific 3D models obtained by different segmentation platforms in the operating room, J. Digit. Imaging 28 (5) (2015) 518–527, https://doi.org/10.1007/s10278-015-9786-4.
[27] F. Kamseu, N. Habra, Adoption of Open Source Software: Is it the Matter of Quality? (2004).
[28] Raphaël Semeteys, Qualification and Selection of Opensource Software [Internet], (2018) (cited 18.12.19), https://www.qsos.org/.
[29] V. del Bianco, L. Lavazza, S. Morasca, D. Taibi, D. Tosi, The Qualispo Approach to Oss Product Quality Evaluation, (2010), https://doi.org/10.1145/1833272.1833277.
[30] Qualipso, Qualification and Selection of Opensource Software [Internet], (2013) (cited 13.16.13), https://web.archive.org/web/20130613043131/http://www.qualipso.eu/.
[31] W.J. Sung, J.H. Kim, S.Y. Rhew, A quality model for open source software selection, Sixth International Conference on Advanced Language Processing and Web Information Technology (ALPIT 2007) (2007) 515–519, https://doi.org/10.1109/ALPIT.2007.81.
[32] L. Aversano, M. Tortorella, Evaluating the Quality of Free/Open Source Systems: A Case Study vol. 73, (2010), pp. 119–134, https://doi.org/10.1007/978-3-642-19802-1_9.
[33] L. Aversano, M. Tortorella, Quality evaluation of floss projects: application to erp systems, Inf. Softw. Technol. 55 (2013) 1260–1276, https://doi.org/10.1016/j.infsof.2013.01.007.
[34] The Qt Company, Qt Project [Internet], (2017) urlhttps://www.qt.io/https://www.qt.io/ (cited 19.11.17).
[35] Kitware Inc, The Visualization ToolKit [Internet], (2017) (cited 19.11.17), https://www.vtk.org.
[36] Kitware Inc, The Insight Segmentation and Registration ToolKit [Internet], (2017) (cited 19.11.17), https://www.itk.org.
[37] OFFIS e.V, DCMTK ToolKit [Internet], (2017) (cited 19.11.17), http://dcmtk.org/.
[38] M. Malaterre, Grassroots DICOM [Internet], (2018) (cited 03.01.18), https://gdcm.sourceforge.net/.
[39] KDE e.V, ThreadWeaver [Internet], (2018) (cited 03.01.18), https://download.kde.org/Attic/frameworks/5.3.0/threadweaver-5.3.0.tar.xz.
[40] Google, Breakpad [Internet], (2018) (cited 03.01.18), https://chromium.googlesource.com/breakpad/breakpad/.
[41] S. Graf, B. List, An evaluation of open source E-learning platforms stressing adaptation issues, Proceedings of the Fifth IEEE International Conference on Advanced Learning Technologies, ICALT'05, IEEE Computer Society, Washington, DC, USA, 2005, pp. 163–165, https://doi.org/10.1109/ICALT.2005.54.
[42] N. Cavus, The evaluation of learning management systems using an artificial intelligence fuzzy logic algorithm, Adv. Eng. Softw. 41 (2) (2010) 248–254, https://doi.org/10.1016/j.advengsoft.2009.07.009.
[43] A. Crabb, I Do Imaging [Internet], (2018) (cited 31.07.18), https://idoimaging.com/.
[44] AlternativeTo, AlternativeTo – Crowdsourced Software Recommendations [Internet], (2018) (cited 31.07.18), https://alternativeto.net/.
[45] GitHub, GitHub [Internet], (2018) (cited 31.07.18), https://github.com/.
[46] Atlassian Pty Ltd, Bitbucket [Internet], (2018) (cited 31.07.18), https://bitbucket.org/.
[47] GitLab, GitLab [Internet], (2018) (cited 31.07.18), https://gitlab.com/.
[48] Slashdot Media, SourceForge – Download, Develop and Publish Free Open Source Software [Internet], (2018) (cited 31.07.18), https://sourceforge.net/.
[49] Aliza Medical Imaging, Aliza Medical Imaging & DICOM Viewer [Internet], (2018) (cited 26.07.18), http://www.aliza-dicom-viewer.com.
[50] Amide, AMIDE: Amide's a Medical Imaging Data Examiner [Internet], (2012) (cited 26.07.18), http://amide.sourceforge.net/.
[51] G. Wollny, Ginkgo CADx [Internet], (2018) (cited 08.01.18), http://ginkgo-cadx.com/en/.
[52] Centro de Tecnologia da Informação Renato Archer, InVesalius [Internet], (2018) (cited 26.07.18), https://www.cti.gov.br/en/invesalius.
[53] S. Min, Z. Wang, N. Liu, An evaluation of html5 and webgl for medical imaging applications, J. Healthc. Eng. 2018 (2018) 1–11, https://doi.org/10.1155/2018/1592821.
[54] R. Yuan, M. Luo, Z. Sun, S. Shi, P. Xiao, Q. Xie, Rayplus: a web-based platform for medical image processing, J. Digit. Imaging 30 (2016 11), https://doi.org/10.1007/s10278-016-9920-y.

3

Article 2
An efficient and uniformly behaving streamline-based μCT fibre tracking algorithm using volume-wise structure tensor and signal processing techniques

Adrià Julià i Juanola; Marc Ruiz i Altisent; Imma Boada i Oliveras, **An efficient and uniformly behaving streamline-based μCT fibre tracking algorithm using volume-wise structure tensor and signal processing techniques**, Computer Methods in Applied Mechanics and Engineering [0045-7825], *vol. 394 (2022): 114898*, 10.1016/j.cma.2022.114898, (accepted in February 2022)

SCImago Journal Rank (SJR): 2.180 (2021), Q1 Computational Mechanics, H-Index 205

Journal Impact Factor (JIF): 6.756 (2020), 2/108 Q1 Mathematics, Interdisc. Appl.

Journal Citation Indicator (JCI): 2.33 (2020), 1/156 Q1 Mechanics

Available online at www.sciencedirect.com
ScienceDirect
Comput. Methods Appl. Mech. Engrg. 394 (2022) 114898
Computer methods in applied mechanics and engineering
www.elsevier.com/locate/cma

An efficient and uniformly behaving streamline-based μCT fibre tracking algorithm using volume-wise structure tensor and signal processing techniques

Adrià Julià i Juanola, Marc Ruiz i Altisent, Imma Boada i Oliveras*

Graphics and Imaging Laboratory (GILAB), Universitat de Girona, Edifici P-IV, 17001 Girona, Catalonia, Spain

Received 26 January 2022; received in revised form 22 February 2022; accepted 22 February 2022
Available online xxxx

Graphical Abstract

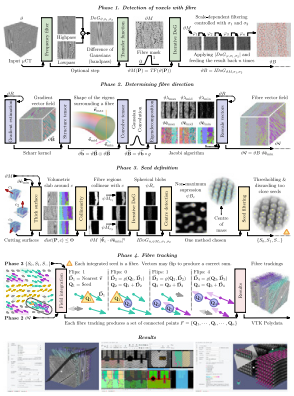

Abstract

A method for reconstructing polygonal paths of fibres in reinforced composites imaged using micro-computed tomography is formally described, implemented and tested. The algorithm has been crafted to be explicable, require no training data and behave uniformly in all axes or orientations. It consists of four phases: (1) segmenting fibre regions using a scale-dependent Iterative Difference of Gaussians approach, (2) extracting directionality using the structure tensor minimum eigenvector, (3) automatically placing the seeds near a set of user-defined restricting surfaces, and (4) tracking fibres using a streamline-based integration method. The algorithm cost grows in relation to the target fibre diameter and is proportional to the number of voxels in the input volume. Its behaviour, ability to process very curved fibres, and error have been assessed using both synthetic and real datasets. The C++ implementation is performant and parallelizable, and produces helpful visualisations to gain insight of the intermediate and final results.

* Corresponding author.
E-mail addresses: adria.julia@udg.edu (A. Julià), marc.ruiz@udg.edu (M. Ruiz), imma.boada@udg.edu (I. Boada).

https://doi.org/10.1016/j.cma.2022.114898

Keywords: Composites; Fibre tracking; Micro-computed tomography; Structure tensor; Image processing; Streamlines

1. Introduction

Current μCT (micro-computed tomography) techniques are able to non-destructively capture precise and detailed information of materials inner composition, and represent it as a stack of tomographic images that define a three-dimensional volume [1,2]. The capabilities of this technology have been exploited in the analysis and evaluation of materials in different applications such as control of manufacturing processes, tensile and compression loading studies, or evaluation of materials resistance, among others. In these applications, two key processes to be considered are: (i) identification and reconstruction of individual fibres from the μCT volumetric acquisitions; and (ii) graphical volume representations alongside the reconstructed data to better evaluate, extract and contextualise information. Both processes are challenging and have become important focuses of research in the last years.

Centring on the identification and reconstruction of fibres, although visible in μCT images, their correct identification and quantification is a hard process that requires a specific post-processing. This is generally based on image processing techniques commonly used in medical science to investigate connectivity and diffusion within the human brain [3–5]. Particularly, reconstruction of fibres in composite materials is usually performed by first identifying their centres and then connecting those with the ones of adjacent images to obtain the fibre tracts throughout the 3D image [6]. Unfortunately, both processes need to face different issues caused by the nature of μCT volume data such as: limited amounts of pixels representing fibre cross-sections; densely packed fibres with fuzzy boundaries; and varying sizes, orientations, and curvatures of individual fibres [7]. Proposed strategies include techniques to overcome these limitations, being in most of the cases designed to support a specific type of material and imaging technique [8,9]. For instance, Schöttl et al. [10] presented and compared several segmentation approaches based on layer thickness, fibre orientation angle, and degree of fibre isotropy to separate each unidirectional tape to analyse the layers individually. They also introduced mapping approaches to transfer local fibre orientation of each tape layer onto a discrete surface. Those approaches were applied to carbon fibre-reinforced polyamide 6 (PA6-CF) samples. Sencu et al. [11] presented an algorithm that uses a global overlapping stack filtering step followed by a local fibre tracking step, both steps based on the Bayesian inference theory, to efficiently locate the fibre centrelines from carbon fibre reinforced polymer (CFRP) μCT images; the centrelines are later used to generate micro-scale finite element models (FEM).

There are also approaches based on template matching techniques which require the provision of dictionaries, i.e. image excerpts of the relevant features to detect. The *InSegt* [12] tool detects single fibres in μCT images using a dictionary-based probabilistic method to determine local fibre misalignments. Similarly, Emerson et al. [13] characterised fibre orientations for unidirectional fibre reinforced polymers, proposing a segmentation method to accurately extract individual fibres from low contrast μCT scans taking the template matching approach, thus requiring a predefined training dictionary; later, the detected centres are connected to produce the polygonal fibre tracking. Likewise, but focusing more on the method efficiency, Amjad et al. [14] proposed an algorithm which also requires training and complements the approach with image correlation techniques on each 2D slice as the algorithm proceeds to detect the fibre centres; their method has been tested with uniform and non-uniform volume spacings. Although being accurate in estimating fibre centres (and thus fibre trajectories), these methods tie the user-provided dictionaries to a specific imaging configuration and sample features such as the fibre diameter, material, matrix, etc. While based on template matching, Creveling et al. [7] proposed a method that synthetically generates fibre-templates from general specifications such as expected fibre diameters and directions. This approach offloads the user from the laborious task of composing a dictionary. The centres detected at each 2D slice are then connected and the individual fibres tracked using a Kalman-filter estimator. Their method was applied to segment fibres in multi-directional fibre reinforced composites. Additionally, a methodology and the datasets necessary for evaluating the algorithm are produced from synthetically generated μCT slices with respect to a ground truth.

So far, the majority of fibre methods will, among other aspects, require the fibre to: (i) be straight or not have significant curvatures, (ii) not be coplanar with the acquired 2D slice, or (iii) exist within certain global orientations. The paradigm of considering a volume as a set of separated slices, is the root cause of those shortcomings. Instead, when algorithms observe the volume as a single 3D image, the slicing axis (the Z) stops being a special case. For

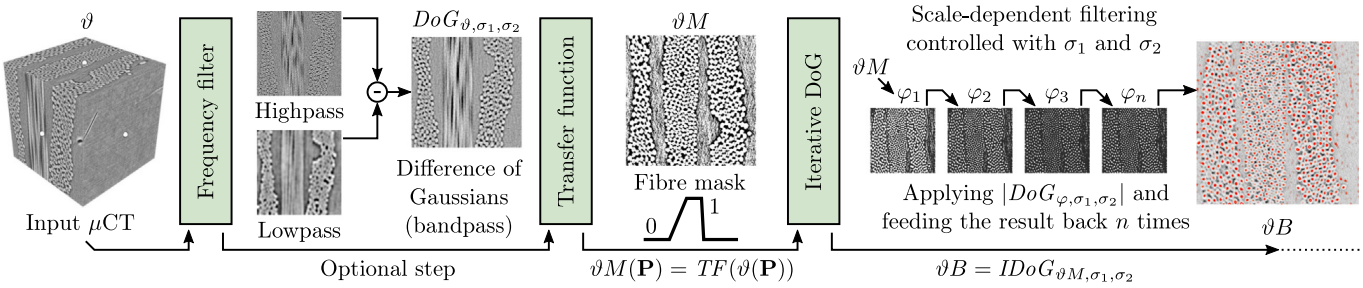

Fig. 1. Block diagram of the *Phase 1* taking a μCT volume as the input. The ϑM segmentation is bandpass-filtered using the Iterative Difference of Gaussians filter in order to conform ϑB (overlaid red blobs). Images produced using the *cfrp_i*17 dataset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

instance, Gaiselmann et al. [15] take this approach and connect the polygonal tracks of curved fibres whose medial axis has been detected using a skeletoning approach applied on the whole volume.

Visualisation of acquired and reconstructed information is also of special interest since graphical representations can enhance data interpretation. Several authors have developed tools to provide these representations. Fritz et al. [16] presented a visual approach to explore and quantify features of interest based on transfer functions in the parameter spaces of specific application scenarios. The applicability and effectiveness of the approach were illustrated using two specific scenarios: the first focused on steel fibre reinforced sprayed concrete, and the second on the designation of the microstructure of ductile cast irons with respect to the contained graphite. Bhattacharya et al. [17] introduced *MetaTracts*, a tool to interactively explore and visualise fibre bundles and weaving patterns from μCT scans of carbon fibre reinforced polymers. Their tool was able to generate volumetric representations and surface models from the extracted fibre bundles. Heinzl and Stappen [18] identified the major tasks regarding visual computing and analysis in materials sciences alongside simulation and testing techniques. They also reviewed input data characteristics, direct and derived outputs, visualisation techniques and metaphors used, and interactions of employed workflows. Weissenböck et al. [19] developed *PorosityAnalyzer*, a novel tool for analysing and evaluating pore segmentation pipelines in fibre-reinforced polymers. Recently, Fröhler et al. [20] developed *FIAKER*, a tool implementing methods for analysing and comparing the results of fibre reconstruction algorithms. An extension of the paper [21] analysed and compared different techniques capable of reconstructing curved fibres. From a commercial standpoint, tools such as, *GeoDict* [22], *VGStudio Max* [23], and *Avizo Software* [24] which employs cylinder fitting techniques [25,26], provide functionalities to visualise and extract relevant information from μCT data.

Taking into account the importance of fibre reconstructions and visualisations from the material science perspective, the aim of this paper is three-fold: (i) present a technique to efficiently and effectively reconstruct the fibres in a μCT dataset by taking an explainable 3D-based approach instead of the common 2D-based (slice-by-slice) strategies used by most of the state-of-the-art methods; (ii) integrate it in a framework that offers different functionalities to explore the input volume as well as the intermediate data produced throughout the method in a user-friendly manner; and (iii) evaluate the performance of the proposal with real and synthetic datasets.

2. Material and methods

2.1. Fibre reconstruction algorithm

The proposed fibre reconstruction algorithm requires as inputs a μCT volume and a set of user-defined surfaces which are then fed into a four-phase process resulting in a set of polygonal lines representing the fibres together with some measurements used in the different graphical representations that give insight into the results.

Before describing the method, some mathematical notation remarks shall be clarified. First, the ϑ function represents a volume whereas ψ is used when it has been constrained/associated to a volume cutting surface c. Additionally, a φ function may be used when referring to an unspecified or intermediate volume. Function names may be suffixed in order to group or express the provenance of intermediary functions such as ϑM or ψB. Bold fonts emphasise which variables or functions return points $\mathbf{P}$, vectors $\vec{\mathbf{v}}$, or tensors $\overleftrightarrow{\mathbf{t}}$. Some functions may be declared

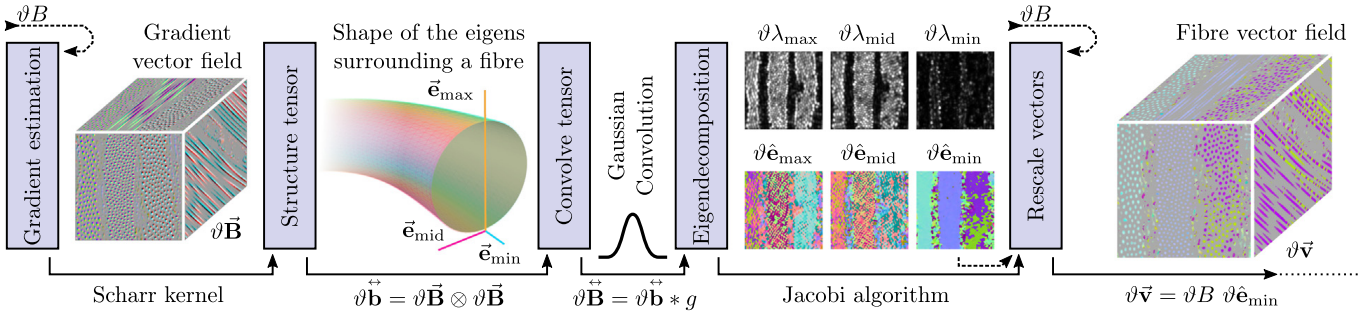

Fig. 2. Block diagram of the *Phase 2* using the ϑB from the previous *Phase 1* in order to produce the $\vartheta\vec{\mathbf{v}}$ vector field representing the fibre directionality and magnitude. Images produced using the *cfrp_i17* dataset mapping vector components into RGB colours. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

with a list of parameters on its subindex; when used a one-to-one assignment between the given values and the declared parameters is made taking priority over any externally provided default value. Finally, the convolution operator is represented by $*$ and the tensor product by $\otimes$.

2.1.1. Phase 1. Detection of voxels with fibre

This phase (see Fig. 1) takes as input a set of r_z μCT image slices of $r_x \times r_y$ pixels in size, which are stored as a voxel model, i.e. a three-dimensional array of $r_x \times r_y \times r_z$ positions with an associated intensity value $\vartheta(\mathbf{P})$, where $\mathbf{P}$ is a $\mathbb{R}^3$ point with (P_x, P_y, P_z) coordinates, and ϑ is a discrete version of a twice continuously differentiable function in order to be integrable, differentiable and interpolatable to its values at non-integers [27,28]. To ease the comprehension, it is assumed that intensity values are normalised to the [0, 1] range and floating point arithmetic is used, although the implemented algorithm allows the use of integer data types which improve the computation performance and optimise memory usage if numeric overflows and precision are kept under control.

To identify the voxels containing fibres, segmentation techniques can be applied. However, under the scope of materials analysis, aspects such as noise, blurred boundaries or multiple fibres lying on a common voxel make segmentations challenging and thus special techniques are required [8,9,13,15]. Among them, thresholding techniques are a common approach to separate the different regions in a volume according to their voxel intensities [29,30]. In our proposal, a non-binary threshold defined like a transfer function (TF) is applied with the intent to preserve as much information as possible by having a fuzzy boundary, a smooth transition, between the *inside* (i.e. region with fibre) and *outside* classifications. Otherwise small intensity variations produced by noise become determinant when a voxel is inevitably assigned in one of the two categories. Those random artefacts increase the gradient magnitude artificially thus negatively impacting the subsequent *Phase 2* where fibre directionality is determined.

$$TF(x) = \begin{cases} 0 & \text{if } x \in (-\infty, a) \\ \frac{x-a}{b-a} & \text{if } x \in [a, b] \\ 1 & \text{if } x \in (b, c) \\ \frac{-x+d}{d-c} & \text{if } x \in [c, d] \\ 0 & \text{if } x \in (d, +\infty) \end{cases} \tag{1}$$

Our one-dimensional transfer function $TF(x)$ (Eq. (1)) has four user-adjustable parameters $a < b < c < d$ which define the thresholds and transition smoothness. The $(-\infty, a)$, $(d, +\infty)$ intervals represent the *outside* (black) values clamped to 0; the (b, c) interval represents the *inside* (white) values clamped to 1; and the $[a, b]$, $[c, d]$ ones define a linear interpolation between the [0, 1] and [1, 0] clamping values respectively. When each voxel represented by ϑ is mapped through the TF, a new ϑM function (Eq. (2)) representing the probability of each voxel to be considered fibre is obtained.

$$\vartheta M(\mathbf{P}) = TF(\vartheta(\mathbf{P})) \tag{2}$$

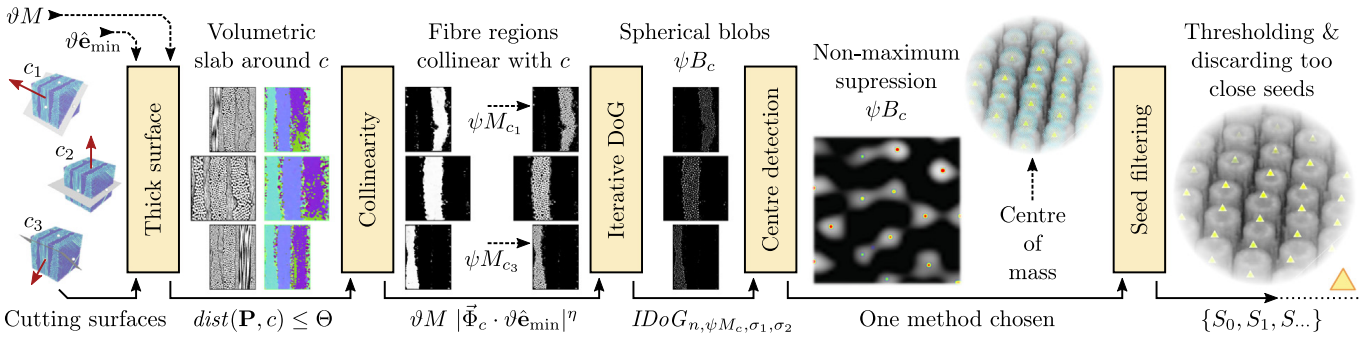

Fig. 3. Block diagram of the *Phase 3* where given a set of user defined cutting surfaces (three in this example), a set of seeds (yellow triangles) at the fibre centres is produced. Images produced using the *cfrp_i17* and Creveling et al. [7] (round) datasets. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The user can interactively evaluate and visualise ϑM; if the result is not satisfactory enough, optionally, a low-pass, band-pass or high-pass frequency filter can be applied on the input ϑ to, among other things, filter high frequency noise in exchange for an improved dynamic range. Then, the *TF* is reapplied, readjusting its parameters if necessary.

$$g_\sigma(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \mathrm{e}^{-\frac{x^2}{2\sigma^2}} \tag{3}$$

These filtering techniques take an unspecified input volume $\varphi(\mathbf{P})$ and convolve it using the Gaussian distribution $g_\sigma(x)$ (Eq. (3)) in order to define a new $G_{\varphi,\sigma}(\mathbf{P})$ (Eq. (6)) function representing a blurry volume. Additionally, when two of them are subtracted (Eq. (4)), the Difference of Gaussians band-pass filter $DoG_{\varphi,\sigma_1,\sigma_2}$ is obtained.

$$DoG_{\varphi,\sigma_1,\sigma_2}(\mathbf{P}) = G_{\varphi,\sigma_1}(\mathbf{P}) - G_{\varphi,\sigma_2}(\mathbf{P}) \tag{4}$$

These convolutions are approximated using finite kernels whose size is linearly bound to the σ parameter, but, in our implementation, the size is governed by a four parameter $\{q, s, h_x, h_y\}$ user-configurable equation $KFactor(\sigma)$ (see Eq. (5a)), whose application and ceiling to the final kernel size takes place at $K(\sigma)$ (Eq. (5b)).

$$KFactor(\sigma) = q\ \mathrm{e}^{-\frac{\ln\left(\frac{s+q}{s+h_y}\right)\sigma}{h_x}} + s, \quad \in \mathbb{R}^+ \tag{5a}$$

$$K(\sigma) = \lceil KFactor(\sigma)\ \sigma \rceil, \quad \in \mathbb{N}^+ \tag{5b}$$

The $KFactor(\sigma)$ function exhibits an exponential behaviour (q) for smaller σ, a linear one (s) for bigger σ, and a control point (h_x, h_y) smoothly separating the two behaviours. In this way, one can improve the quality of smaller kernels by proportionally considering more discrete neighbours without impacting the performance of bigger convolutions.

$$G_{\varphi,\sigma} = \varphi * g_\sigma \tag{6a}$$

$$G_{\varphi,\sigma}(\mathbf{P}) \approx \int_{-K(\sigma)}^{K(\sigma)} \int_{-K(\sigma)}^{K(\sigma)} \int_{-K(\sigma)}^{K(\sigma)} g_\sigma(\|\mathbf{U}\|)\ \varphi(\mathbf{P}-\mathbf{U})\ dU_x\, dU_y\, dU_z \tag{6b}$$

While ϑM segments the regions with fibre, it still has noise, sharp boundaries and almost no gradual gradient towards the fibre medial axis. These limitations are overcome by applying the Iterative Difference of Gaussians (IDoG) filter (Eq. (7)). Starting from an initial φ volume function, the filter proceeds by: (i) applying $DoG_{\varphi,\sigma_1,\sigma_2}$ (see Eq. (4)) on the current iteration φ function; (ii) discarding the negative values (absolute value); and (iii) repeating the process (feeding the result back) n times until the trivial case $IDoG_0$ is reached and thus $\varphi(\mathbf{P})$ returned.

$$IDoG_{n,\varphi,\sigma_1,\sigma_2}(\mathbf{P}) = IDoG_{n-1,\varphi',\sigma_1,\sigma_2}(\mathbf{P}), \quad \varphi' = |DoG_{\varphi,\sigma_1,\sigma_2}| \tag{7a}$$

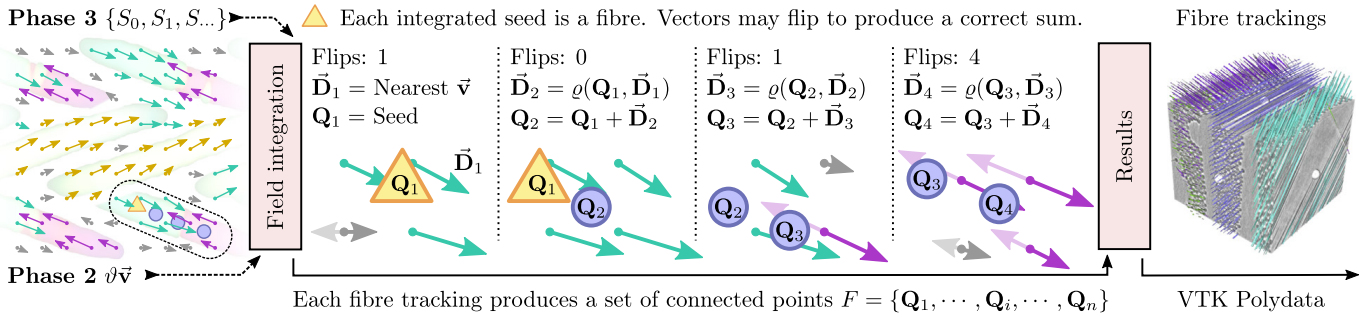

Fig. 4. Block diagram of the *Phase 4* where given the vector field $\vartheta\vec{\mathbf{v}}$ from *Phase 2* (coloured arrows) tracks the fibres starting from *Phase 3* seeds (yellow triangles) by resolving each $\mathbf{Q}$ (blue points), flipping the vectors (semi-transparent arrows) if necessary. Images produced using the *cfrp_i17* dataset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$IDoG_{0,\varphi,\sigma_1,\sigma_2}(\mathbf{P}) = \varphi(\mathbf{P}) \tag{7b}$$

Experimentally, it has been observed that from 1 to 3 iterations suffice to filter and shape the fibres of interest. Eq. (8) applies the described technique over ϑM, to produce a new ϑB. This type of filtering is commonly used in medical imaging to enhance vessel structures [31] as those are scale dependent and the blobs (i.e. fibres) outside the target size have their magnitude reduced, which (in addition to the filtering benefits) may serve to indicate the deviation form the target diameter.

$$\vartheta B(\mathbf{P}) = IDoG_{n,\vartheta M,\sigma_1,\sigma_2}(\mathbf{P}) \tag{8}$$

2.1.2. Phase 2. Determining fibre direction

The goal of this phase is to determine the local fibre directionality at any position in the volume. While for 2D images the gradient vector field rotated by 90° could fit this purpose, it is unsuitable for higher dimensions as it would only restrict the directionality within the hyperplane. For this reason, the more advanced structure tensor [32,33] technique commonly used in material analysis [27,27,34–36] is required.

As illustrated in Fig. 2, this step begins by taking the ϑB as the input in order to produce a gradient vector field $\vartheta\vec{\mathbf{B}}$ (Eq. (9)) estimated using a $5 \times 5 \times 5$ 3D Scharr kernel [37] which is optimised to achieve a high level of precision and rotational invariance.

$$\vartheta\vec{\mathbf{B}}(\mathbf{P}) = \left[\frac{\partial\,\vartheta B}{\partial x}(\mathbf{P}),\ \frac{\partial\,\vartheta B}{\partial y}(\mathbf{P}),\ \frac{\partial\,\vartheta B}{\partial z}(\mathbf{P})\right] \tag{9}$$

From this $\vartheta\vec{\mathbf{B}}$, a symmetric rank-2 structure tensor field $\vartheta\overleftrightarrow{\mathbf{B}}$ expresses the gradient distribution around a finite Gaussian convolution window of $K(\sigma)$ voxels as described by Eqs. (10) and (11).

$$\vartheta\overleftrightarrow{\mathbf{b}}(\mathbf{P}) = \vartheta\vec{\mathbf{B}}(\mathbf{P}) \otimes \vartheta\vec{\mathbf{B}}(\mathbf{P}) \tag{10a}$$

$$\vartheta\overleftrightarrow{\mathbf{b}}(\mathbf{P}) = \begin{bmatrix} \frac{\partial\,\vartheta B}{\partial x}(\mathbf{P})^2 & \cdots & sym \\ \frac{\partial\,\vartheta B}{\partial y}(\mathbf{P})\frac{\partial\,\vartheta B}{\partial x}(\mathbf{P}) & \frac{\partial\,\vartheta B}{\partial y}(\mathbf{P})^2 & \vdots \\ \frac{\partial\,\vartheta B}{\partial z}(\mathbf{P})\frac{\partial\,\vartheta B}{\partial x}(\mathbf{P}) & \frac{\partial\,\vartheta B}{\partial z}(\mathbf{P})\frac{\partial\,\vartheta B}{\partial y}(\mathbf{P}) & \frac{\partial\,\vartheta B}{\partial z}(\mathbf{P})^2 \end{bmatrix} \tag{10b}$$

$$\vartheta\overleftrightarrow{\mathbf{B}} = \vartheta\overleftrightarrow{\mathbf{b}} * g_\sigma \tag{11a}$$

$$\vartheta\overleftrightarrow{\mathbf{B}}(\mathbf{P}) \approx \int_{-K(\sigma)}^{K(\sigma)}\int_{-K(\sigma)}^{K(\sigma)}\int_{-K(\sigma)}^{K(\sigma)} g_\sigma(\|\mathbf{U}\|)\ \vartheta\overleftrightarrow{\mathbf{b}}(\mathbf{P}-\mathbf{U})\ dU_x\, dU_y\, dU_z \tag{11b}$$

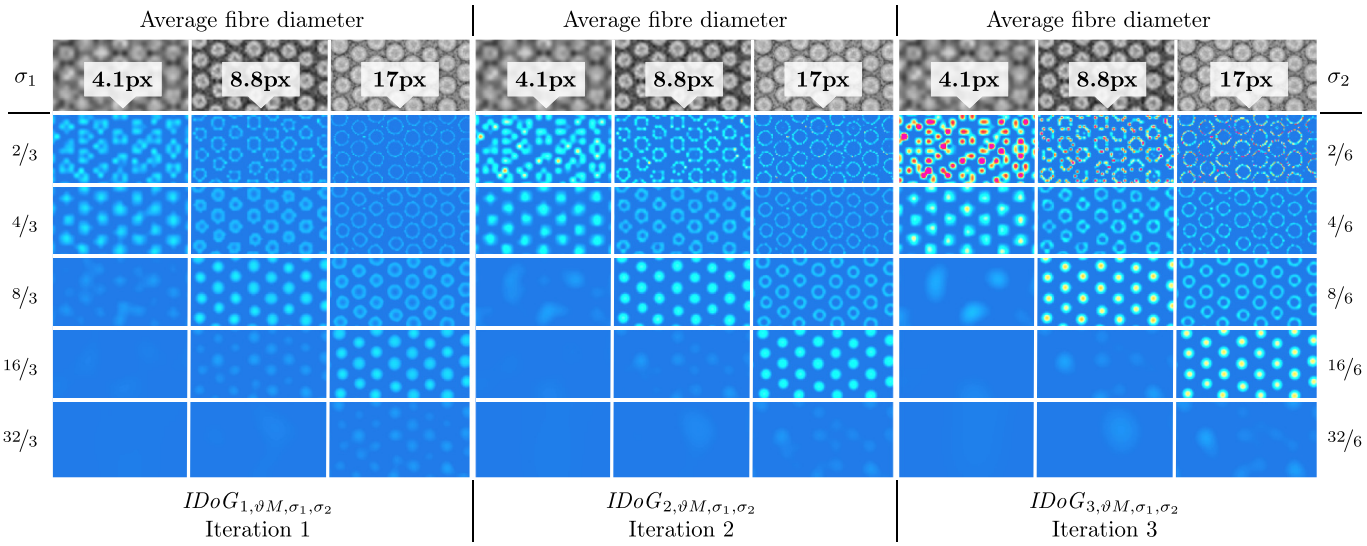

Fig. 5. Effect of altering σ_1 and σ_2 (y axis) of an IDoG filter (grouped by iterations) when the fibre diameter (x axis) varies. To allow a meaningful comparison, intensities are scaled within $[0, 10]$ range (blue to red). Images produced using Creveling et al. [7] datasets. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

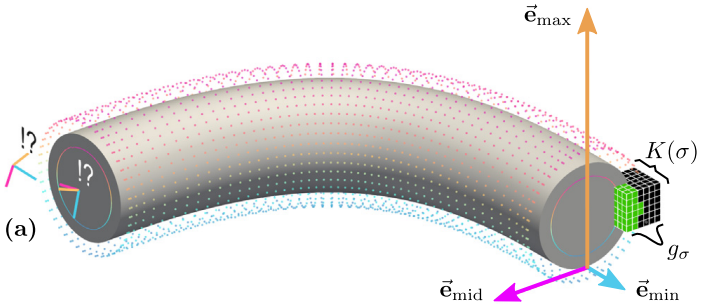

Fig. 6. An idealised fibre, with a tensor at the boundary and a cube representing the convolution kernel, alongside the effect on tensors at an unreachable $K(\sigma)$ distance (a).

Then, when each tensor voxel is eigendecomposed (diagonalised) using the Jacobi method [38,39] its $\lambda_{\max}$, λ_{mid}, $\lambda_{\min}$ eigenvalues, and $\hat{\mathbf{e}}_{\max}$, $\hat{\mathbf{e}}_{\text{mid}}$, $\hat{\mathbf{e}}_{\min}$ unitary eigenvectors can be extracted and the respective $\vartheta\lambda_{\max}(\mathbf{P})$, $\vartheta\lambda_{\text{mid}}(\mathbf{P})$, $\vartheta\lambda_{\min}(\mathbf{P})$, $\vartheta\hat{\mathbf{e}}_{\max}(\mathbf{P})$, $\vartheta\hat{\mathbf{e}}_{\text{mid}}(\mathbf{P})$, and $\vartheta\hat{\mathbf{e}}_{\min}(\mathbf{P})$ functions produced. From these, other supplementary volumes expressing morphological features of the tensor [40,41] can be generated in order to examine aspects such as the tensor diffusion (Fig. 7) among others (Fig. 8(a)).

The relevant fibre orientation information for the posterior steps is in the $\vartheta\hat{\mathbf{e}}_{\min}(\mathbf{P})$ unitary vector field function which corresponds to the direction of least gradient (collinear with the medial axis) as illustrated in Fig. 6.

$$\vartheta\vec{\mathbf{v}}(\mathbf{P}) = \vartheta B(\mathbf{P})\ \vartheta\hat{\mathbf{e}}_{\min}(\mathbf{P}) \tag{12}$$

This phase ends scaling the unitary vectors by ϑB in order to produce the final vector field $\vartheta\vec{\mathbf{v}}$ where the magnitude increases as a fibre medial axis is approached (Eq. (12)).

2.1.3. Phase 3. Seed definition

A common need of fibre tracking algorithms is the provision of seeding points from which each tracking emanates. In our case, seeds are automatically placed near a set of m user-defined cutting surfaces $C = \{C_1, \ldots, C_m\}$ (Fig. 3). From each cutting surface c an undetermined amount of k seeds $S = \{\mathbf{S}_1, \ldots, \mathbf{S}_k\}$ will be computed as follows. Each cutting surface defines a ψM_c (Eq. (13)) volume function where only the locations $\mathbf{P}$ around a distance Θ from the surface have a non-zero return value; thus producing a thin volume with short cylindrical fibres when Θ is set to the target fibre diameter. Moreover, ψM_c is able to suppress nearby regions whose local

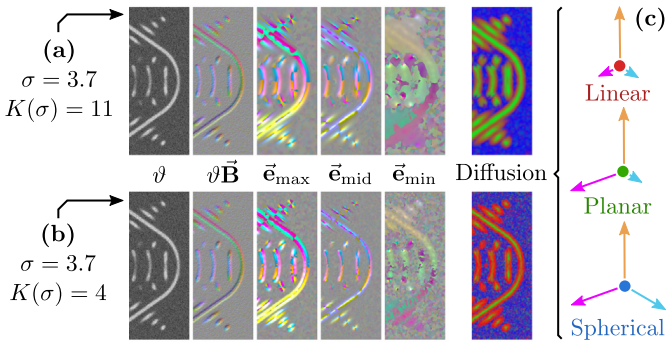

Fig. 7. The effect of varying the kernel size (a, b) on the tensor diffusion (c) mapped as follows: linear anisotropy (red), planar anisotropy (green), and isotropy (blue). Images produced using the *scaphandre* synthetic dataset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

fibre directionality $\vartheta\hat{\mathbf{e}}_{\min}$ is not collinear with the closest surface normal. This filtering is performed using the dot product and exponentiating the result by a factor η to provide a continuous user-adjustable tolerance parameter. The closest normal from a $\mathbf{P}$ location with respect to an arbitrary c surface is resolved by a $\vec{\Phi}_c(\mathbf{P})$ function, which will produce an invariant result if c is flat.

$$\psi M_c(\mathbf{P}) = \begin{cases} \vartheta M(\mathbf{P}) |\vec{\Phi}_c(\mathbf{P}) \cdot \vartheta\hat{\mathbf{e}}_{\min}(\mathbf{P})|^{\eta} & dist(\mathbf{P},\ c) \leq \Theta \\ 0 & dist(\mathbf{P},\ c) > \Theta \end{cases} \tag{13}$$

However, a similar issue as in *Phase 1* arises: the cylindrical fibre sections in ψM_c are not smooth nor spherical. For this reason, a ψB_c function (Eq. (14)) turns those sections into blobs with a rising intensity towards its centres by applying the IDoG filter (Eq. (7)) over ψM_c.

$$\psi B_c(\mathbf{P}) = IDoG_{n,\psi M_c,\sigma_1,\sigma_2}(\mathbf{P}) \tag{14}$$

This smoothness leads to a gradient peak at the blob centre which is detectable using a non-maximum suppression step, or, alternatively, if the blobs are clearly separated from each other, by computing the centre of masses using a region-growing method which further improves the precision to sub-voxel locations. Finally, for each c surface the corresponding set of seeds $\{\mathbf{S}_1, \ldots, \mathbf{S}_k\}$ is made from the $\mathbb{R}^3$ coordinates detected at the centres which: (i) have a scalar value above a certain user-defined threshold; and (ii) are separated by a minimum distance usually linked to the target fibre diameter in order to discard some occasional close peaks duplicates.

For the sake of simplicity, our current implementation assumes planar surfaces which are interactively placed and oriented by its normal vector and origin. Briefly summarised, the ψM_c (then morphed into ψB_c) requires the underlying volume ϑM to be affine transformed (resliced) in order to have its Z axis match the surface normal. The different miniatures in Fig. 3 are examples of these transformations; and although complex to implement, they are performant as only a small and thin portion of the whole dataset has to be considered.

2.1.4. Phase 4. Fibre tracking

Starting from every seed $\mathbf{s}$, this last phase creates fibre tracks by integrating the vector field $\vartheta\vec{\mathbf{v}}$ numerically, but employing a slightly modified interpolation methodology. Formally, for each seed $\mathbf{s}$ obtained from the previous *Phase 2*, an associated path of points (i.e. fibre tracking) $F = \{\mathbf{Q}_1, \ldots, \mathbf{Q}_i, \ldots, \mathbf{Q}_n\}$ conforms a list of $n-1$ segments. Every $\mathbf{Q}_i$ is computed like a particle travelling on the steady vector field $\vartheta\vec{\mathbf{v}}$ using a numeric method such as RK2, RK4, RK4.5 [42,43] and evolving as follows: (i) from a $\mathbf{Q}_i$ location, the local field directionality is evaluated; (ii) a small advancement in that particular direction is made; (iii) a new location $\mathbf{Q}_{i+1}$ is found; and (iv) the process is systematically repeated to integrate a full trajectory.

Unfortunately, as illustrated in Fig. 4, $\vartheta\vec{\mathbf{v}}$ cannot be straightforwardly integrated since at some orientations, the interpolated sum of collinear and randomly signed vectors will net in a non-collinear sum escaping the real fibre trajectory (see turquoise and purple arrows). This issue emerges on RGB mapped vector volumes $\vartheta\hat{\mathbf{e}}_{\min}$ or $\vartheta\vec{\mathbf{v}}$ in Figs. 2, 10, and 7 as alternating colour patches.

$$flip(\vec{\mathbf{v}}, \vec{\mathbf{D}}) = \begin{cases} \vec{\mathbf{v}}, & \text{if } \vec{\mathbf{v}} \cdot \vec{\mathbf{D}} \geq 0 \\ \vec{\mathbf{v}}(-1), & \text{if } \vec{\mathbf{v}} \cdot \vec{\mathbf{D}} < 0 \end{cases} \tag{15}$$

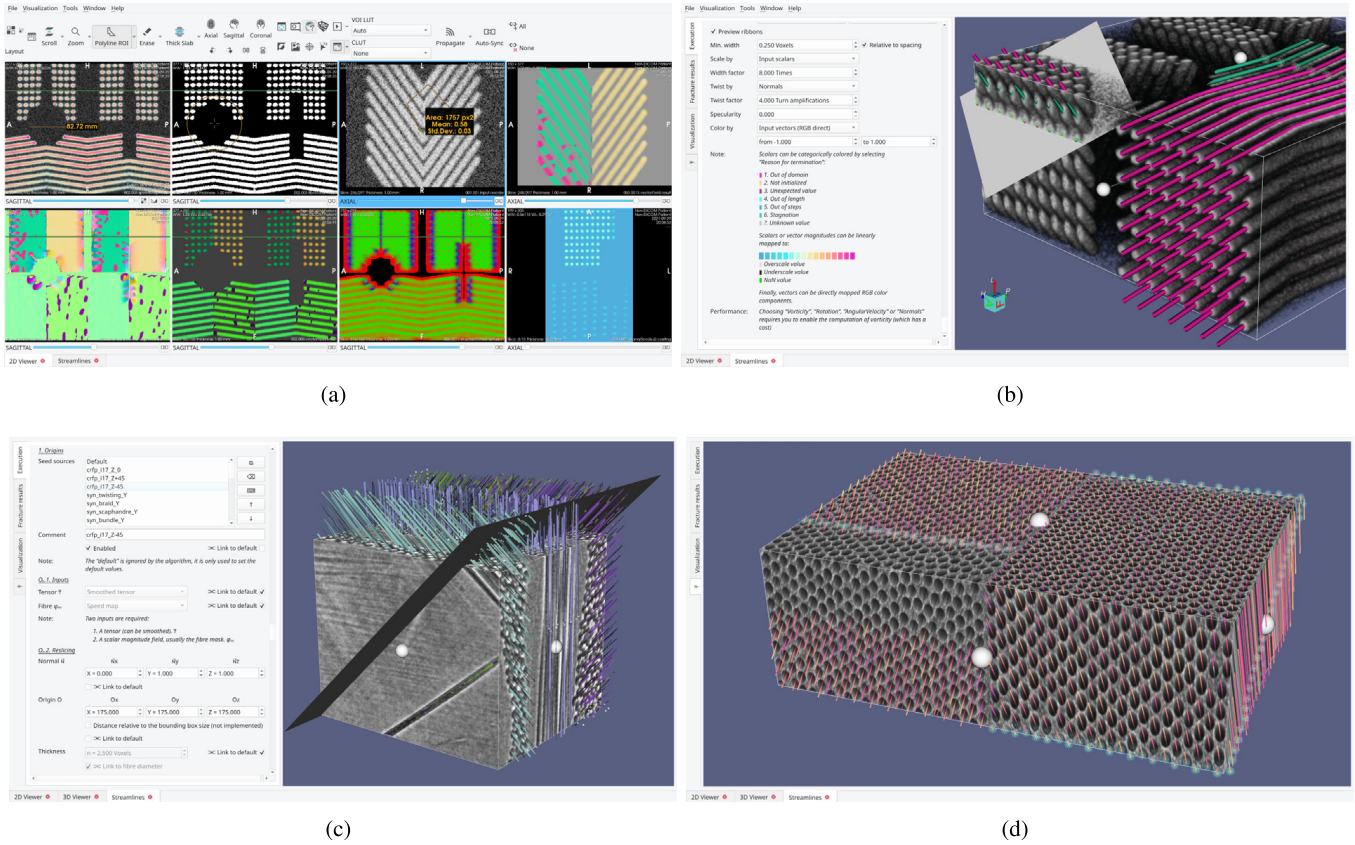

(a) (b)

(c) (d)

Fig. 8. The 2D interface (a) displaying some of the intermediate volumes of the *bundle* synthetic dataset alongside the 3D visualisations in (b) where one cutting plane is shown and the fibres are coloured according to their reason for termination. (c) The results produced using the real *cfrp_i17* dataset. (d) The centre-of-mass point clouds and seeds using the 17px Creveling et al. [7] dataset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

To tackle this ambiguity, a reference vector $\vec{\mathbf{D}}$ taken from the previous step is fed into a ϱ function (Eq. (16)) which interpolates the eight discrete samples around a $\mathbf{P}$ location as given by $\vartheta\vec{\mathbf{v}}(\mathbf{P})$ with $\mathbf{P} \in \mathbb{N}^3$. Before adding them all, these eight vectors may be flipped according to the $flip(\vec{\mathbf{v}}, \vec{\mathbf{D}})$ function (Eq. (15)) in order to cohere with the reference vector $\vec{\mathbf{D}}$ sign. At each iteration, $\vec{\mathbf{D}}$ is updated; and as it can be safely assumed that fibres will never take a very sharp U-turn in a single step distance, then critical cases such as curls (see Fig. 10) are seamlessly handled as they traverse those problematic regions. Initially, without any previous step to refer, $\vec{\mathbf{D}}$ takes the vector value from the nearest discrete sample of $\vartheta\vec{\mathbf{v}}$ around the seed.

$$\varrho(\mathbf{P}, \vec{\mathbf{D}}) = \sum_{U_x=\lfloor P_x \rfloor}^{\lceil P_x \rceil} \sum_{U_y=\lfloor P_y \rfloor}^{\lceil P_y \rceil} \sum_{U_z=\lfloor P_z \rfloor}^{\lceil P_z \rceil} weight(\mathbf{P}-\mathbf{U})\, flip(\vartheta\vec{\mathbf{v}}(\mathbf{U}), \vec{\mathbf{D}}), \quad \mathbf{U} \in \mathbb{N}^3 \tag{16}$$

Additionally, for each $\mathbf{Q}_i$ point, extra attributes such as field magnitude, rotation, vorticity, etc. can be associated; and, for each fibre as a whole, attributes such as the reason for termination can be specified. The most common ones are: (i) the volume domain has been escaped, (ii) a maximum path length is reached, (iii) a maximum number of steps is achieved, or (iv) stagnation occurs because the field magnitude is below a certain threshold. Figs. 8(b), and 10 illustrate some of those circumstances.

2.2. Visualisation framework

The proposed approach has been integrated in a visualisation framework implemented in C++ in the Starviewer platform [44] with the support from several open-source libraries, the main ones being Qt [45] and VTK [46]. As

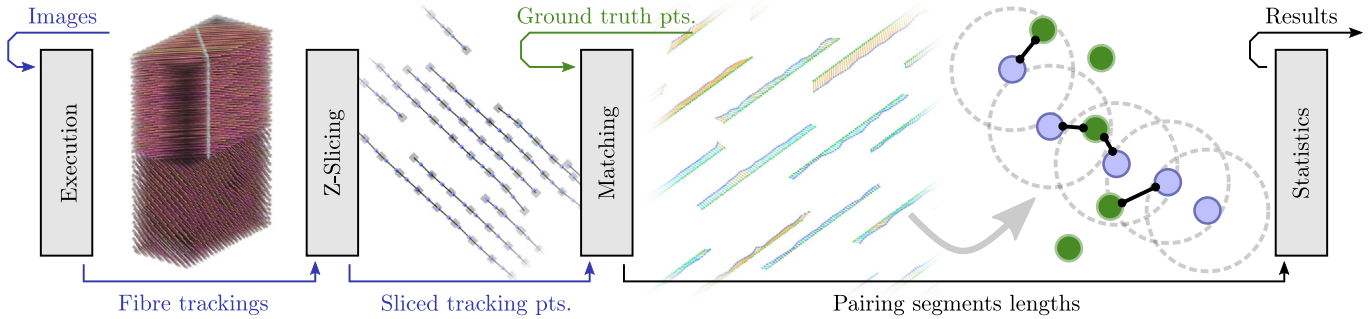

Fig. 9. Steps necessary to convert our fibre tracking results to a suitable Z-sliced point cloud comparable against the respective ground truth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

illustrated in Fig. 8 and Supplementary Material 5, the user interface can be divided in two parts: the configuration side panel and the visualisation area. The first, although complex, contains several aids such as integrated help tips and automatically computed parameters derived from a more general one (i.e. "link to" checkboxes). Briefly summarised, the essential configurations are: (i) parameters of the fibre transfer function, (ii) target fibre diameter, (iii) *IDoG* iterations, and (iv) placement of seeding surfaces.

The other part, the visualisation area, provides graphical representations of the information obtained from the input model at different steps of the process. In this way, the user can interactively explore the fibres with respect to a rendered volume, supervise the seed placement, explore the intermediate volumes, and visualise the fibre tracking attributes among other tasks. Such interactions provide insight of the input model and ultimately lead to a better definition of the algorithm parameters.

2.3. Experiments and testing datasets

The proposed method has been tested using several synthetic and real datasets. To evaluate the error, the synthetic models created by Creveling et al. [7] have been used as they provide a total of 12 cases along its ground truth data consisting of: three straight fibre orientations on the Z axis and four scaling levels to produce the different diameters (4.1px, 8.8px, 14.4px and 17.0px). While the comparison and error measurement could be performed using tools such as FIAKER [20], we opted for the ad hoc approach in Fig. 9 which converts our 3D fibre results (grey points) to a regular Z-sliced point cloud (blue points) which can be seamlessly matched (connecting segments) against the closest points in the Creveling et al. [7] ground truth coordinates (green points). Note that due to the filtering nature of our seeding methodology, some fibres on the ground truth data shall be discarded if they are at an unrealistic distance away in relation to the fibre diameter (grey dashed lines). The actual C++ implementation of the comparison method and results accompanied by the 3D VTK data necessary to visually inspect the paired points is accessible at Supplementary Material (1).

Four synthetic datasets named *twisting*, *scaphandre*, *bundle* and *braid* exhibiting different combinations of curvatures with severe amounts of noise have been created in order to assess the method behaviour under these complex scenarios. Alongside their corresponding ground truth, they are made available in Supplementary Material 4.

Finally, three real μCT models alongside their execution results are provided in Supplementary Material 3. Unfortunately as there is no ground truth to compare against, a more subjective validation methodology as described in Section 3.2 had to be employed. The [47] datasets offer a sharp fibreglass case (*leuven_glass*) alongside a noisy carbon one (*leuven_carbon*). The third *cfrp_i17* dataset, acquired with the European Synchrotron Radiation Facility's ID19 Beamline (ESRF, Grenoble) at 26 keV, with a pixel size of 650 nm, contains three main fibre orientations, and a more challenging set of features and defects.

3. Results and discussion

In this section the different aspects of the proposed method have been evaluated by: (i) assessing the behaviour on edge-cases; (ii) quantitatively determining the error against a ground truth; and (iii) analysing the key parameters affecting the computation cost.

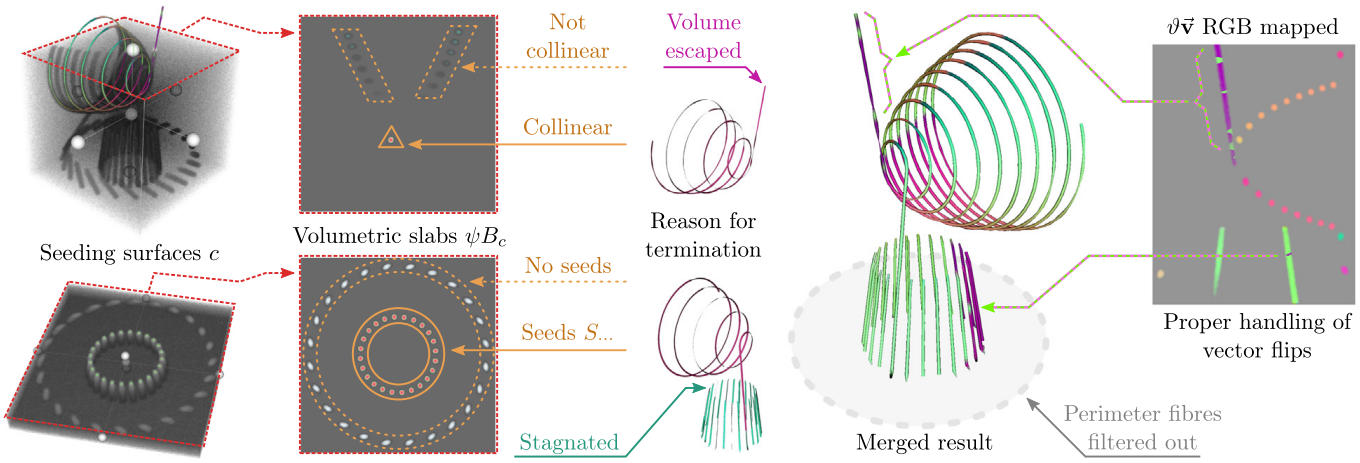

Fig. 10. Example of the method behaviour using the *twisting* synthetic dataset. Seeds (solid yellow) are only placed if they are collinear with the cutting surface (dashed red). Only the selected subset of fibres is tracked using the vector field $\vartheta\vec{\mathbf{v}}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.1. Method behaviour

The first aspect to note is that our algorithm, unlike two-dimensional or slice-dependent approaches [7,11–14], considers the input as a whole volume where all dimensions and orientations shall be treated equally. To properly assess our disregard with respect to a slicing axis (usually Z), *twisting*, *scaphandre*, *bundle* and *braid* synthetic datasets were employed. Putting aside the assumption that physical fibres do not merge nor bifurcate, Fig. 10 illustrates the method's ability to reconstruct the path of two closely whirling fibres. While these may be an exaggerated example, slight curvatures can be observed in the real *cfrp_i17* dataset, an aspect not present in the perfectly aligned Creveling et al. synthetic datasets. However, in order to achieve this, the fibre tracking (*Phase 4*) has to be fed with a quality vector field $\vartheta\vec{\mathbf{v}}$ also respecting those principles. In order to achieve this, the volume is observed as a valuable signal to process using analogue-like filtering principles implemented over digital means with the aim to maximise the signal-to-noise ratio. These, in addition to behaving in predictable ways, when performed 3D-wise, consider the information between neighbouring slices instead of obliviating it as some of the aforementioned approaches do. This processing is performed using Gaussian convolutions and Scharr [37] gradient estimators in order to behave in all directions uniformly and achieve rotational invariance. At the same time, another side benefit of Gaussian filters is their selectivity effect towards a particular fibre diameter (i.e. frequency) as illustrated for the *IDoG* case in Fig. 5. The values of each σ can be bound to particular target size (Eq. (17)), and the tolerance adjusted according to the number of iterations.

$$\sigma_1 = \frac{diameter}{3} \qquad \sigma_2 = \frac{diameter}{6} \tag{17}$$

A second phenomenon to consider is when the tensor field $\vartheta\overleftrightarrow{\mathbf{b}}$ components are convolved with their neighbours in order to produce $\vartheta\overleftrightarrow{\mathbf{B}}$. In this particular case, the kernel dimensions as dictated by $K(\sigma)$ can make the otherwise low significance samples relevant when determining the orientation of the unitary $\hat{\mathbf{e}}_{\text{min}}$ eigenvector field. This effect is discernible on the diffusion maps in Fig. 7 where the planar isotropy (green) increases with the kernel size even though the Gaussian distribution σ is kept invariant.

A third aspect to note is that while most cited methods aim to reconstruct all the fibres in a volume; we opted for enabling the user to restrict the tracking by three main criteria: (i) cutting surface of origin, a user orientable plane on the current implementation; (ii) amount of collinearity with respect to this surface; and (iii) the minimum magnitude to sustain through all the fibre before stagnating. These, in addition to other indirect parameters governing the image processing filters (e.g. fibre diameter, iterations), can be strategically combined to selectively reconstruct only certain bundles of interest as illustrated in Fig. 10. In this way, a user familiarised with those fundamentals can properly fine tune the relevant parameters in order to reconstruct datasets with varying circumstances, imaging

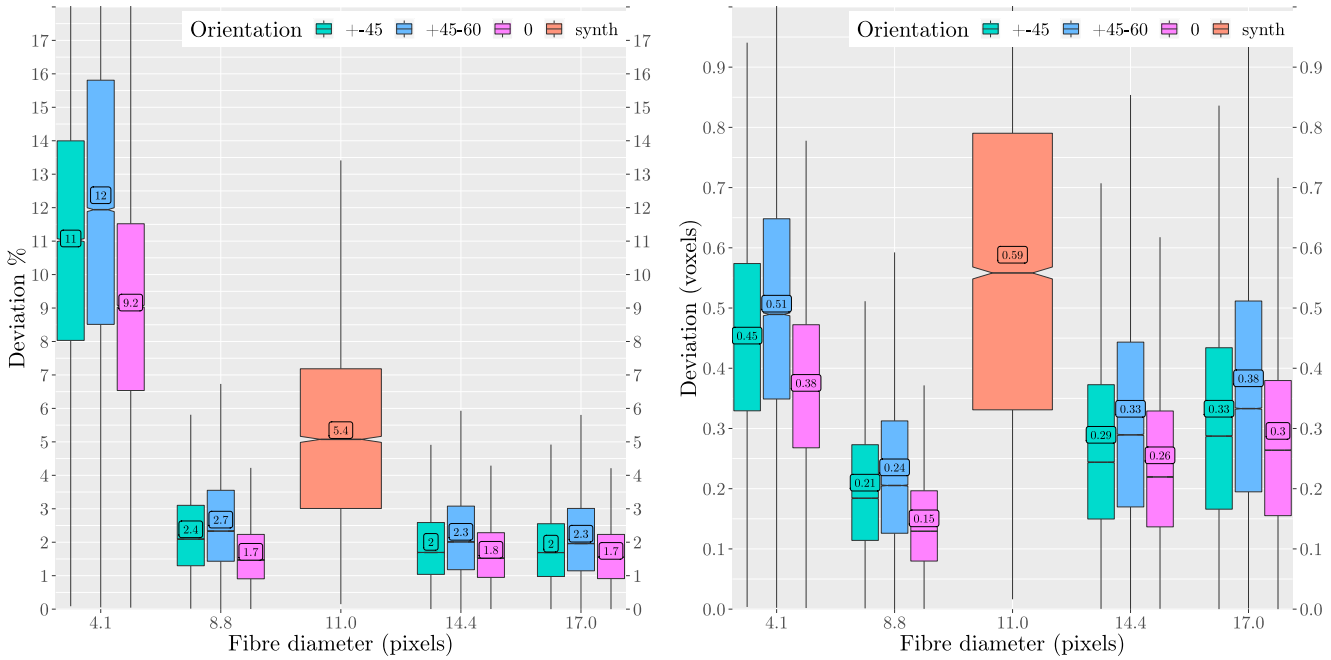

(a) Absolute error (b) Relative error

Fig. 11. Absolute (a) and relative (b) error of our approach measured using our synthetic (orange) and Creveling et al. [7] testing datasets.

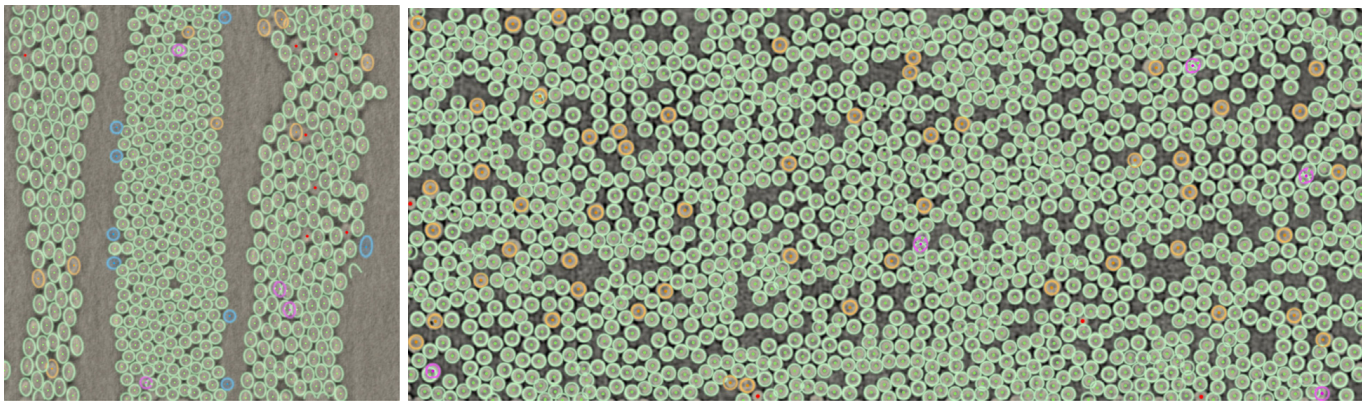

(a) (b)

Fig. 12. Verification of the *cfrp_i17* (a) and *leuven_glass* (b) real datasets against our algorithm output (green) by manually identifying misplacements (yellow), duplicates (magenta), phantoms (blue), and misses (red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

techniques, and difficult acquisitions (i.e. low contrast, noisy, etc.), which are attainable if the relevant fibre signal does not mix with under-sampling artefacts (e.g. ring artefacts, non-linearities) at the same frequencies.

We consider all those aspects beneficial compared to more opaque approaches requiring training data. An explainable behaviour is mainly provided by the use of: (i) common signal processing techniques, (ii) constrained seed placement at peak intensities, and (iii) the use of fluid dynamics principles (streamlines) to trace the fibre paths.

Table 1
Absolute (voxels) and relative (% fibre diameter) errors.

Case	⌀	Median	Mean	Median	Mean
±45	4.1	0.45 vx.	0.45 vx.	11.03%	11.07%
45–60	4.1	0.49 vx.	0.50 vx.	11.93%	12.37%
0	4.1	0.37 vx.	0.38 vx.	9.05%	9.16%
±45	8.8	0.18 vx.	0.21 vx.	2.09%	2.38%
45–60	8.8	0.21 vx.	0.24 vx.	2.33%	2.67%
0	8.8	0.13 vx.	0.15 vx.	1.47%	1.71%
synth	11.0	0.56 vx.	0.59 vx.	5.07%	5.35%
±45	14.4	0.24 vx.	0.29 vx.	1.69%	2.01%
45–60	14.4	0.29 vx.	0.33 vx.	2.00%	2.31%
0	14.4	0.22 vx.	0.26 vx.	1.52%	1.77%
±45	17.0	0.29 vx.	0.33 vx.	1.69%	1.95%
45–60	17.0	0.33 vx.	0.38 vx.	1.95%	2.25%
0	17.0	0.26 vx.	0.30 vx.	1.55%	1.74%

3.2. Error

To evaluate the error of the proposed approach, the methodology described in Section 2.3 was applied in all 12 cases of Creveling et al. [7] and our four synthetic datasets in Supplementary Material 4. Fig. 11 summarises the average displacement error (labels) using boxplots grouped by the average fibre diameter (x axis) and orientation case (colour). Fig. 11(a) is produced by measuring the Euclidean length of the small line segments between the ground truth and fibre tracking point clouds obtained with our approach as shown in Fig. 9. The complementary Fig. 11(b) is crafted to match Figure 9 of Creveling et al. [7] paper and shows the relative error in relation to the fibre diameter. In those cases, except for the 4.1px one, we have a very useable error of about 2%, which is in par with [7] results although it can be observed that the absolute error is slightly bigger in the 14.4px and 17px cases at a steady absolute value of about 0.3 voxels. However, on our four synthetic 11px cases, a greater error is expected as the ground truth from which the fibres are generated arise from a fuzzy point cloud with at least an error of 1px. While we cannot attribute an exact reason for some of those fluctuations, at those high levels of accuracy, clearly below the pixel size, they will in practice be limited by external factors such as the imaging technique, quality, uncertainty in the measurement, etc. In addition, the small variations in the parameters and settings used may have also contributed to this phenomenon although they were only linearly scaled to match the average fibre diameter stated for each dataset. For more details, data, visualisations and code refer to Table 1 and Supplementary Material (1).

In order to further validate the results when no ground truth is available: a manual classification of the defects as illustrated in Fig. 12 was performed on the *cfrp_i17* and *leuven_glass* real datasets. Although not flawless, most instances were of slightly misplaced fibres (yellow); phantoms (blue), detectable when abrupt turns occur; and duplicates (magenta) which could be suppressed by eliminating correlated paths at physically impossible distances. Given these post-filtering abilities, the sensitivity to start a tracking could be increased in order to reduce the amount of undetected fibres (red). Supplementary Material 3 contains more overlaid slices, including the *leuven_carbon* case which had to be dismissed as the image quality was not good enough to make a comparable judgement, especially for misplacements or duplicates.

3.3. Cost

In order to evaluate the cost a set of tests were performed and contrasted with the theoretical expectations. In that regard, the first aspect to consider is that whole volume convolutions can be characterised by the number of neighbouring voxels visited for each input voxel. For this reason, the voxel-wise growth is expected to correlate with the convolution size and not the overall volume size. Note that Gaussian convolutions are performed

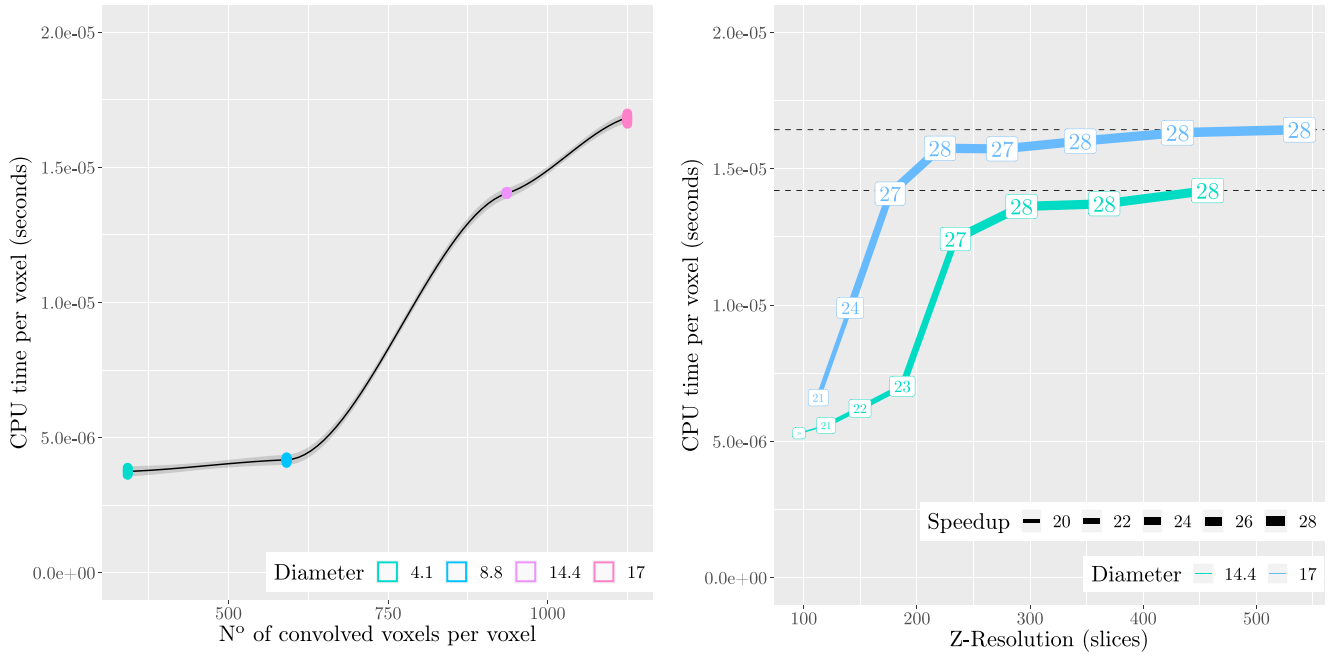

(a) Variant scaling, fixed number of fibres.

(b) Variant fibre length, fixed scale.

Fig. 13. Growth of the cost when some of the input volume characteristics are held invariant. Creveling et al. [7] datasets were used.

separately for each dimension thus reducing the overall cost significantly. To experimentally test this, and ease the comprehensibility among all the variables impacting it, experimental results are analysed in relation to the *CPU time* dedicated for each voxel (y axis in all plots) because it is independent of the volume size and number of threads on a particular testing system (in this case a 32-thread AMD ThreadRipper 1950X CPU). For each dataset under test, the following data was automatically collected: (i) metadata regarding the input size alongside the key parameters to estimate convolution sizes; (ii) *wall time* or elapsed period as experienced by the user; and (iii) *CPU time* which sums the individual CPU utilisation times across all threads. Fig. 14 summarises and groups all tests in relation to some of the most important variables: (i) number of fibres tracked (logarithmic x axis), (ii) volume size in voxels (fill colour scale), (iii) diameter (dot size), and (iv) convolution operations required on each voxel (border opacity). From these, it can be experimentally observed that the major factor penalising the performance is the fibre diameter and not the overall volume size, confirming the theoretical expectation.

These facts are better explained by Fig. 13, when Creveling et al. [7] datasets (which have an invariant number of fibres) are employed to restrict some of the variables. In Fig. 13(a), the scale is varied, and it can be observed how the number of convolutions required (x axis) alongside the cost (y axis) increases with respect to the fibre diameter (colour) with little variability between the different orientation cases. On the other hand in Fig. 13(b), where the scaling is fixed and the overall volume size (x axis) is altered by cropping in the Z direction (collinear with the fibres) for the 14.4px and 17px cases (colour), a very important fact arises: as the overall volume size grows, the cost per voxel is kept constant, thus achieving a very desirable linear behaviour.

Under the current testing environment, if the 16 s processing time of the 350 × 350 × 350 *cfrp_i17* dataset is extrapolated to an extent of 2500 × 2500 × 1500, provided enough memory, a very reasonable time of 1 h is projected. Fig. 13(b) labels show the level of parallelism or speedup achieved (limited to a theoretical maximum of 32 in our current setup), and is computed by dividing the *CPU time* and *wall time*. However, unlike the computations performed on a volume per-voxel basis, the stream tracing workload can only be atomised up to a seed level, as each tracking must be performed sequentially from its initial conditions. Estimating this cost is challenging as the number of steps required will vary depending on the user-defined restrictions. If no limits are imposed on the length or number of segments, the computation may never halt. Although the current tracking implementation is

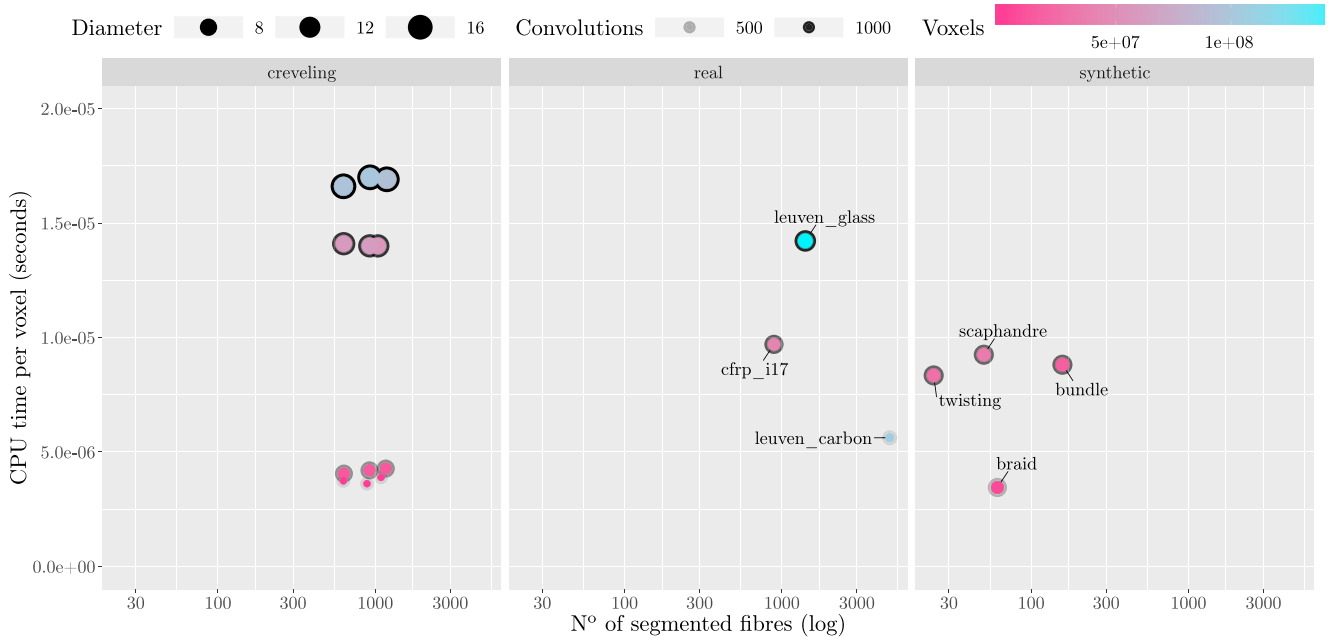

Fig. 14. Summary of the cost growth using all cases under test in relation to their main characteristics. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

single-threaded, in a practical scenario the attainable speed-up is very high as the number of seeds will naturally outnumber the available CPUs; and although memory access to the whole dataset is formally required, the accessions will follow a locality principle. This pitfall may have negatively impacted the experimental speed-up shown on the labels of Fig. 13(b) and will be addressed in our future work.

Fore more details refer to Supplementary Material 2 where the data and extra plots are provided alongside a coloured spreadsheet with the calculations performed to estimate the theoretical number of convolutions required for each voxel.

3.4. Shortcomings

The presented results have shown good performance and error figures along with robustness when addressing complex use cases. Although the implementation is feature complete, improvements can be made. For example, the ability to define a curved surface, or refinement on the automations in order to simplify the overwhelming amount of parameters and details exposed by default on the user interface. Also, the manually defined planes could be automatically suggested according to the $\vartheta\vec{\mathbf{v}}$ orientation histogram [28]. These, alongside other restrictions, could be implemented in order to mandate the traversal of different surfaces, sharing certain sub-paths or correlations with nearby fibres. While many of the intermediate and final results provide insight on the analysed dataset, some could be extended or improved. For instance the fibre size attribute associated at each point $\mathbf{Q}$ in a tracking is unable to indicate if it deviates positively or negatively from the estimated diameter. Those values are taken from the underlying ϑB volume, which is conformed to respond to a given target diameter. To overcome this limitation a fourth dimension could be added to the volume by making use of scale-space techniques [48,49] which, in addition to respecting our design principles and plausibly providing a better diameter indicator, could be suitable to properly address complex cases with different fibre diameters or filtering out undesirable features while preserving the efficiency by upholding separable convolutions.

4. Conclusions

In this paper, we proposed a new method for the reconstruction and visualisation of fibres from μCT images that differs from common state-of-the art methods that process volumes in a two dimensional image-by-image fashion. Our four-step approach considers the volume as a whole where streamlines and three-dimensional image processing techniques (mainly Gaussian convolutions) are applied to achieve a uniform behaviour in all axes. First, it determines the probability of each voxel to belong to a fibre using a transfer function with four user-adjustable parameters. The obtained mask is then improved by an Iterative Difference of Gaussians bandpass filter producing a gradual intensity rise from fibre boundaries to its medial axes. Second, using the structure tensor technique, a vector field representing the directionality of the fibre at each voxel is obtained. Third, given a set of user-defined cutting surfaces, seeds are automatically placed at the nearby fibre centres collinear with the surface normal. Fourth, a streamline-based tracking approach reconstructs the polygonal fibre paths starting from the seeds and integrating through the directionality vector field. The method has been implemented into a framework that provides, on the one hand, the different interfaces necessary to adjust the numerous algorithm parameters and, on the other hand, the different visualisations to interpret the results interactively for each step.

The proposed approach has been tested on both synthetic and real datasets. From the tests it has been observed that, in addition to capitalising on signal-to-noise improvements, common limitations of slice-by-slice based paradigms such as coplanar fibres on the acquisition plane are seamlessly avoided. In addition, despite facing an increased complexity arising from the uniform treatment of the third dimension of the volume, the algorithm performance is good, parallelizable, with a cost growing linearly with respect to the total number of voxels, and attaining a measured error below 3%. Moreover, key parameters such as the target fibre diameter or the user-defined seeding surfaces can be strategically combined to selectively process certain bundles.

As a future work, we will focus on reducing required user interaction with the aim to automate and assist some of tasks without degrading the algorithm explainability doctrine. In addition, we aim to further extend the framework with different visualisation strategies to enhance data interpretation.

CRediT authorship contribution statement

Adrià Julià i Juanola: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Marc Ruiz i Altisent:** Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Imma Boada i Oliveras:** Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

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

Acknowledgments

The authors want to thank the support from AMADE-UdG *(Analysis and Advanced Materials for Structural Design)* research group.

Funding sources

This work has been financially supported by grants from the Spanish Government *(Ministerio de Ciencia, Innovación y Universidades)* PID2019-106426RB-C31.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.cma.2022.114898.

(1) **error.tar.xz** Contains (among other things) the execution results, extra graphs and statistics used to perform the comparison of our execution results against our synthetic, and Creveling et al. [7] datasets ground truth. The C++ implementation of the comparison methodology is also included.

`1efe7a14eb69b9c5972b4efa6a71bb7beb8a2833ee9eb2e7d0679648a330be76 (SHA256)` `be02b93daaafd241b4aca12aa72a53ab6440e11b (SHA1)`

(2) **cost.tar** Benchmarks, results, plots and scripts used to characterise the cost of our method.

`c5c0bb509044be0ede9a93d0f091c04600c0a3be7caafe24668482d242bc6435 (SHA256)` `09bb5fa97b9b7e4773b12c0d29232b985ae95315 (SHA1)`

(3) **real.tar.xz** Contains the execution results and manual validation of our method using the real *leuven_carbon*, *leuven_glass* [47], and *cfrp_i17* datasets.

`a2cd54a731f606198d02d6a0739a08369fb743da9f92313e582768a47a097c10 (SHA256)` `60269a3917347c8d15bad3aa01b0ce39efc34448 (SHA1)`

(4) **synthetic.tar.xz** Contains the *twisting*, *scaphandre*, *bundle* and *braid* synthetic datasets, alongside their ground truth and our execution results.

`2a5c58c78c64e26e592866564161c2a7fbfb27c0342824fcd5c86fa38b059c67 (SHA256)` `0b9dabb209bb11e0f9e02f38141b5015936cd825 (SHA1)`

(5) Videos and screenshots of the implemented solution. The `.mkv` files are encoded in VP9 and contain a subtitle track with comments. We recommend visualising them with VLC media player.

screencaptures.tar `ed1eeacf8118949ecf630043731b282c1b6fd21700fb75aafa26c7392d3fd878 (SHA256)` `81fc93def85f981611e4acc9ccdfad23c8143c7b (SHA1)`

fhd_real.mkv

`e02746bae484368515a8e5e0fa0b8ebe28ee7fe421f8dd1a4319ed7fe01dd05e (SHA256)` `d082dc00a135c118f783efb2184d0874bcc2521a (SHA1)`

fhd_curved.mkv

`ece543e87ba796f610998ad26d80cd118717cea975f50dc27c1845d5cfba3da3 (SHA256)` `c54e29492aa9b76dfb5984e64d7d0fc6ec0c6dc1 (SHA1)`

fhd_synthetic.mkv

`4957bf7e4168052a2f42cd4538bc34ad0c12d14d4cc91f9736a73ad56b3516b5 (SHA256)` `0bc7a9577ff67e1c19ce2f3e9fe1a51b14e004d2 (SHA1)`

References

[1] J. Kruth, M. Bartscher, S. Carmignato, R. Schmitt, L. De Chiffre, A. Weckenmann, Computed tomography for dimensional metrology, CIRP Ann. 60 (2) (2011) 821–842, http://dx.doi.org/10.1016/j.cirp.2011.05.006, URL https://www.sciencedirect.com/science/article/pii/S0007850611002083.

[2] S. Garcea, Y. Wang, P. Withers, X-ray computed tomography of polymer composites, Compos. Sci. Technol. 156 (2018) 305–319, http://dx.doi.org/10.1016/j.compscitech.2017.10.023, URL https://www.sciencedirect.com/science/article/pii/S0266353817312460.

[3] S. Mori, B. Crain, V. Chacko, P. van zijl, Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging, Ann. Neurol. 45 (1999) 265–269, http://dx.doi.org/10.1002/1531-8249(199902)45:23.0.CO;2-3.

[4] S. Mori, P.C.M. van Zijl, Fiber tracking: principles and strategies – a technical review, NMR Biomed. 15 (7–8) (2002) 468–480, http://dx.doi.org/10.1002/nbm.781, arXiv:https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/pdf/10.1002/nbm.781, URL https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/abs/10.1002/nbm.781.

[5] P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, A. Aldroubi, In vivo fiber tractography using DT-MRI data, Magn. Reson. Med. 44 (4) (2000) 625–632, http://dx.doi.org/10.1002/1522-2594(200010)44:4<625::AID-MRM17>3.0.CO;2-O, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/1522-2594%28200010%2944%3A4%3C625%3A%3AAID-MRM17%3E3.0.CO%3B2-O, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/1522-2594%28200010%2944%3A4%3C625%3A%3AAID-MRM17%3E3.0.CO%3B2-O.

[6] A. Bhattacharya, C. Heinzl, A. Amirkhanov, J. Kastner, R. Wenger, MetaTracts - A method for robust extraction and visualization of carbon fiber bundles in fiber reinforced composites, in: 2015 IEEE Pacific Visualization Symposium, PacificVis, 2015, pp. 191–198, http://dx.doi.org/10.1109/PACIFICVIS.2015.7156377.

[7] P.J. Creveling, W.W. Whitacre, M.W. Czabaj, A fiber-segmentation algorithm for composites imaged using X-ray microtomography: Development and validation, Composites A 126 (2019) 105606, http://dx.doi.org/10.1016/j.compositesa.2019.105606, URL https://www.sciencedirect.com/science/article/pii/S1359835X19303550.

[8] R. Blanc, C. Germain, J.D. Costa, P. Baylou, M. Cataldi, Fiber orientation measurements in composite materials, Composites A 37 (2) (2006) 197–206, CompTest 2004, https://doi.org/10.1016/j.compositesa.2005.04.021, URL https://www.sciencedirect.com/science/article/pii/S1359835X05002472.

[9] M. Kronenberger, K. Schladitz, B. Hamann, H. Hagen, Fiber segmentation in crack regions of steel fiber reinforced concrete using principal curvature, Image Anal. Stereol. 37 (2) (2018) 127–137, http://dx.doi.org/10.5566/ias.1914, URL https://www.ias-iss.org/ojs/IAS/article/view/1914.

[10] L. Schöttl, D. Dörr, P. Pinter, K.A. Weidenmann, P. Elsner, L. Kärger, A novel approach for segmenting and mapping of local fiber orientation of continuous fiber-reinforced composite laminates based on volumetric images, NDT & E Int. 110 (2020) 102194, http://dx.doi.org/10.1016/j.ndteint.2019.102194, URL https://www.sciencedirect.com/science/article/pii/S0963869519303743.

[11] R. Sencu, Z. Yang, Y. Wang, P. Withers, C. Rau, A. Parson, C. Soutis, Generation of micro-scale finite element models from synchrotron X-ray CT images for multidirectional carbon fibre reinforced composites, Composites A 91 (2016) 85–95, http://dx.doi.org/10.1016/j.compositesa.2016.09.010, URL https://www.sciencedirect.com/science/article/pii/S1359835X16303049.
[12] V.A. Dahl, M.J. Emerson, C.H. Trinderup, A.B. Dahl, Content-based propagation of user markings for interactive segmentation of patterned images, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW, 2020, pp. 4280–4288, http://dx.doi.org/10.1109/CVPRW50498.2020.00505.
[13] M.J. Emerson, K.M. Jespersen, A.B. Dahl, K. Conradsen, L.P. Mikkelsen, Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials, Composites A 97 (2017) 83–92, http://dx.doi.org/10.1016/j.compositesa.2016.12.028, URL https://www.sciencedirect.com/science/article/pii/S1359835X16304560.
[14] K. Amjad, W. Christian, K. Dvurecenska, M. Chapman, M. Uchic, C. Przybyla, E. Patterson, Computationally efficient method of tracking fibres in composite materials using digital image correlation, Composites A 129 (2020) 105683, http://dx.doi.org/10.1016/j.compositesa.2019.105683, URL https://www.sciencedirect.com/science/article/pii/S1359835X19304324.
[15] G. Gaiselmann, I. Manke, W. Lehnert, V. Schmidt, Extraction of curved fibers from 3D data, Image Anal. Stereol. 32 (1) (2013) 57–63, http://dx.doi.org/10.5566/ias.v32.p57-63, URL https://www.ias-iss.org/ojs/IAS/article/view/979.
[16] L. Fritz, M. Hadwiger, G. Geier, G. Pittino, M.E. Groller, A visual approach to efficient analysis and quantification of ductile iron and reinforced sprayed concrete, IEEE Trans. Vis. Comput. Graphics 15 (6) (2009) 1343–1350, http://dx.doi.org/10.1109/TVCG.2009.115.
[17] A. Bhattacharya, J. Weissenböck, R. Wenger, A. Amirkhanov, J. Kastner, C. Heinzl, Interactive exploration and visualization using MetaTracts extracted from carbon fiber reinforced composites, IEEE Trans. Vis. Comput. Graphics 23 (8) (2017) 1988–2002, http://dx.doi.org/10.1109/TVCG.2016.2582158.
[18] C. Heinzl, S. Stappen, Star: Visual computing in materials science, Comput. Graph. Forum 36 (3) (2017) 647–666, http://dx.doi.org/10.1111/cgf.13214, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13214, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13214.
[19] J. Weissenböck, A. Amirkhanov, E. Gröller, J. Kastner, C. Heinzl, PorosityAnalyzer: Visual analysis and evaluation of segmentation pipelines to determine the porosity in fiber-reinforced polymers, 2016, pp. 101–110, http://dx.doi.org/10.1109/VAST.2016.7883516.
[20] B. Fröhler, T. Elberfeld, T. Möller, H. Hege, J. Weissenböck, J. De Beenhouwer, J. Sijbers, J. Kastner, C. Heinzl, A visual tool for the analysis of algorithms for tomographic fiber reconstruction in materials science, Comput. Graph. Forum 38 (3) (2019) 273–283, http://dx.doi.org/10.1111/cgf.13688, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13688, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13688.
[21] B. Fröhler, T. Elberfeld, T. Möller, H.-C. Hege, J.D. Beenhouwer, J. Sijbers, J. Kastner, C. Heinzl, Analysis and comparison of algorithms for the tomographic reconstruction of curved fibres, Nondestruct. Test. Eval. 35 (3) (2020) 328–341, http://dx.doi.org/10.1080/10589759.2020.1774583.
[22] Math2Market GmbH, GeoDict, 2021, [Internet] URL https://www.geodict.com/Solutions/aboutGD.php.
[23] Volume Graphics, VGSTUDIO MAX, 2021, [Internet] URL https://www.volumegraphics.com/en/products/vgsm.html.
[24] Thermo Fisher Scientific Inc., Avizo software, 2021, [Internet] URL https://www.thermofisher.com/es/es/home/electron-microscopy/products/software-em-3d-vis/avizo-software.html.
[25] F. de Pascalis, M. Nacucchi, Relationship between the anisotropy tensor calculated through global and object measurements in high-resolution X-ray tomography on cellular and composite materials, J. Microsc. 273 (1) (2019) 65–80, http://dx.doi.org/10.1111/jmi.12762, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/jmi.12762, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jmi.12762.
[26] M. Mehdikhani, C. Breite, Y. Swolfs, M. Wevers, S.V. Lomov, L. Gorbatikh, Combining digital image correlation with X-ray computed tomography for characterization of fiber orientation in unidirectional composites, Composites A 142 (2021) 106234, http://dx.doi.org/10.1016/j.compositesa.2020.106234, URL https://www.sciencedirect.com/science/article/pii/S1359835X2030470X.
[27] M. Krause, J. Hausherr, B. Burgeth, C. Herrmann, W. Krenkel, Determination of the fibre orientation in composites using the structure tensor and local X-ray transform, J. Mater. Sci. 45 (2010) 888–896, http://dx.doi.org/10.1007/s10853-009-4016-4.
[28] I. Straumit, S.V. Lomov, M. Wevers, Quantification of the internal structure and automatic generation of voxel models of textile composites from X-ray computed tomography data, Composites A 69 (2015) 150–158, http://dx.doi.org/10.1016/j.compositesa.2014.11.016, URL https://www.sciencedirect.com/science/article/pii/S1359835X14003625.
[29] N. Otsu, A threshold selection method from gray-level histograms, IEEE Trans. Syst. Man Cybern. 9 (1) (1979) 62–66, http://dx.doi.org/10.1109/TSMC.1979.4310076.
[30] D.-Y. Huang, C.-H. Wang, Optimal multi-level thresholding using a two-stage otsu optimization approach, Pattern Recognit. Lett. 30 (3) (2009) 275–284, http://dx.doi.org/10.1016/j.patrec.2008.10.003, URL https://www.sciencedirect.com/science/article/pii/S0167865508002985.
[31] A.F. Frangi, W.J. Niessen, K.L. Vincken, M.A. Viergever, Multiscale vessel enhancement filtering, in: W.M. Wells, A. Colchester, S. Delp (Eds.), Medical Image Computing and Computer-Assisted Intervention — MICCAI'98, Springer Berlin Heidelberg, Berlin, Heidelberg, 1998, pp. 130–137.
[32] G.M.P.V. Kempen, N. van den Brink, L.J. van Vliet, M.V. Ginkel, P.W. Verbeek, H. Blonk, The application of a local dimensionality estimator to the analysis of 3-D microscopic network structures, in: SCIA'99, Proceedings of the 11th Scandinavian Conference on Image Analysis (Kangerlussuaq 1999), 1999, pp. 447–455.
[33] K. Robb, O. Wirjadi, K. Schladitz, Fiber orientation estimation from 3D image data: Practical algorithms, visualization, and interpretation, in: 7th International Conference on Hybrid Intelligent Systems, HIS 2007, 2007, pp. 320–325, http://dx.doi.org/10.1109/HIS.2007.26.
[34] M.D. Budde, J.A. Frank, Examining brain microstructure using structure tensor analysis of histological sections, Neuroimage 63 (1) (2012) 1–10, http://dx.doi.org/10.1016/j.neuroimage.2012.06.042, URL https://www.sciencedirect.com/science/article/pii/S105381191200657X.

[35] P. Pinter, S. Dietrich, B. Bertram, L. Kehrer, P. Elsner, K. Weidenmann, Comparison and error estimation of 3D fibre orientation analysis of computed tomography image data for fibre reinforced composites, NDT & E Int. 95 (2018) 26–35, http://dx.doi.org/10.1016/j.ndteint.2018.01.001, URL https://www.sciencedirect.com/science/article/pii/S0963869517303821.

[36] R. Karamov, L.M. Martulli, M. Kerschbaum, I. Sergeichev, Y. Swolfs, S.V. Lomov, Micro-CT based structure tensor analysis of fibre orientation in random fibre composites versus high-fidelity fibre identification methods, Compos. Struct. 235 (2020) 111818, http://dx.doi.org/10.1016/j.compstruct.2019.111818, URL https://www.sciencedirect.com/science/article/pii/S0263822319322986.

[37] H. Scharr, Optimale Operatoren in der Digitalen Bildverarbeitung (Ph.D. thesis), Naturwissenschaftlich-Mathematischen Gesamtfakultät, 2000, pp. 145–156, http://dx.doi.org/10.11588/heidok.00000962, URL http://www.ub.uni-heidelberg.de/archiv/962.

[38] G.H. Golub, H.A. van der Vorst, Eigenvalue computation in the 20th century, J. Comput. Appl. Math. 123 (1) (2000) 35–65, Numerical Analysis 2000. Vol. III: Linear Algebra, https://doi.org/10.1016/S0377-0427(00)00413-1, URL https://www.sciencedirect.com/science/article/pii/S0377042700004131.

[39] C. Jacobi, Über ein leichtes Verfahren die in der Theorie der Säcularstörungen vorkommenden Gleichungen numerisch aufzulösen*, 1846 (30) (1846) 51–94, http://dx.doi.org/10.1515/crll.1846.30.51.

[40] C.-F. Westin, S. Maier, H. Mamata, A. Nabavi, F. Jolesz, R. Kikinis, Processing and visualization for diffusion tensor MRI, Med. Image Anal. 6 (2) (2002) 93–108, http://dx.doi.org/10.1016/S1361-8415(02)00053-1, URL https://www.sciencedirect.com/science/article/pii/S1361841502000531.

[41] S. Mori, Introduction to Diffusion Tensor Imaging, Elsevier Science B.V., Amsterdam, 2007, http://dx.doi.org/10.1016/B978-044452828-5/50021-1, URL https://www.sciencedirect.com/book/9780444528285,

[42] W. Press, B. Flannery, S. Teukolsky, W. Vetterling, Numerical Recipes in FORTRAN 77: Volume 1, Volume 1 of Fortran Numerical Recipes: The Art of Scientific Computing, Cambridge University Press, 1992, URL https://books.google.es/books?id=LRO0zAEACAAJ.

[43] J.R. Cash, A.H. Karp, A variable order Runge-Kutta method for initial value problems with rapidly varying right-hand sides, ACM Trans. Math. Software 16 (3) (1990) 201–222, http://dx.doi.org/10.1145/79505.79507.

[44] M. Ruiz, A. Julià, I. Boada, Starviewer and its comparison with other open-source DICOM viewers using a novel hierarchical evaluation framework, Int. J. Med. Inform. 137 (2020) 104098, http://dx.doi.org/10.1016/j.ijmedinf.2020.104098, URL https://www.sciencedirect.com/science/article/pii/S1386505619301108.

[45] Qt Group Oyj, Qt libraries, 2021, [Internet] URL https://www.qt.io.

[46] Kitware Inc., The visualization ToolKit, 2021, [Internet] URL https://www.vtk.org.

[47] M. Mehdikhani, C. Breite, Y. Swolfs, M. Wevers, S.V. Lomov, L. Gorbatikh, A dataset of micro-scale tomograms of unidirectional glass fiber/epoxy and carbon fiber/epoxy composites acquired via synchrotron computed tomography during in-situ tensile loading, Data in Brief 34 (2021) 106672, http://dx.doi.org/10.1016/j.dib.2020.106672, URL https://www.sciencedirect.com/science/article/pii/S2352340920315511.

[48] T. Lindeberg, Scale-space for discrete signals, IEEE Trans. Pattern Anal. Mach. Intell. 12 (3) (1990) 234–254, http://dx.doi.org/10.1109/34.49051.

[49] T. Lindeberg, Scale-Space Theory in Computer Vision, first ed., in: The Springer International Series in Engineering and Computer Science, Royal Institute of Technology, Stockholm, Sweden, Springer, 1994, p. XII, 424, http://dx.doi.org/10.1007/978-1-4757-6465-9.

4

Article 3
A frequency-limited waviness and curvature measurement algorithm for noisy and irregularly sampled composite fibre trackings

Adrià Julià i Juanola; Marc Ruiz i Altisent; Narcís Coll i Arnau; Imma Boada i Oliveras **A frequency-limited waviness and curvature measurement algorithm for noisy and irregularly sampled composite fibre trackings**, Computer Methods in Applied Mechanics and Engineering [0045-7825], (submitted in June 2022)

SCImago Journal Rank (SJR): 2.180 (2021), Q1 Computational Mechanics, H-Index 205

Journal Impact Factor (JIF): 6.756 (2020), 2/108 Q1 Mathematics, Interdisc. Appl.

Journal Citation Indicator (JCI): 2.33 (2020), 1/156 Q1 Mechanics

A frequency-limited waviness and curvature measurement algorithm for noisy and irregularly sampled composite fibre trackings

Adrià Julià i Juanola[a], Marc Ruiz i Altisent[a], Narcís Coll i Arnau[a], Imma Boada i Oliveras[a]

[a]*Graphics and Imaging Laboratory (GILAB)*
Universitat de Girona, Edifici P-IV
17001 Girona, Catalonia

Abstract

Given a set of irregularly sampled 3D polygonal curves representing composite fibres within a micro-computed tomography volume, a new approach based on the Frenet-Serret formulas is proposed to measure the point curvature and waviness along a polyline even when its oscillations are not coplanar. However, a direct computation of the measures would lead to ill-formed results depending on variant externalities across acquisitions such as noise, sampling, resolution, fractality, etc. Consequentially, we also propose a decoupling mechanism employing a low-pass Gaussian frequency filter to gradually discard features smaller than a certain user-specified σ wavelength referenced in actual space units. This proposal has been tested, characterized and visualized using both real and synthetic datasets contemplating complex waveform features to asses the filter selectivity and convergence across varying sampling frequencies (i.e. polyline resolution). The implementation (in the form of a C++ VTK filter), alongside an extra amount of supplementary materials encompassing the execution results, synthetic datasets and its generative code is provided.

Keywords: Composites; Waviness; Fractal coastline paradox; Streamlines; Micro-computed tomography

Graphical abstract

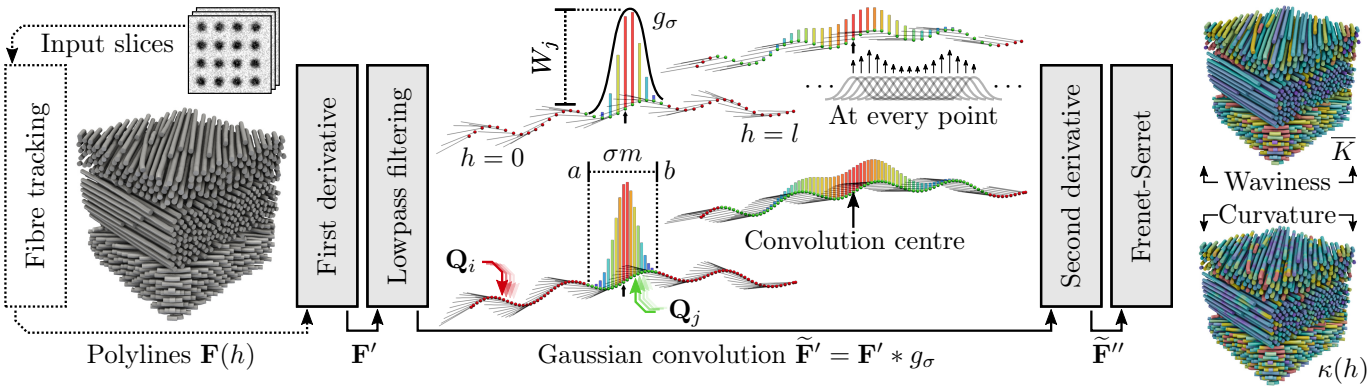

⋆Corresponding author
Email addresses: adria.julia@udg.edu (Adrià Julià i Juanola), marc.ruiz@udg.edu (Marc Ruiz i Altisent), narcis.coll@udg.edu (Narcís Coll i Arnau), imma.boada@udg.edu (Imma Boada i Oliveras⋆)

1. Introduction

In the last years, 3D X-ray micro-computed tomography (μCT) has become an extensively used technique in materials science as is able to non-destructively capture detailed and precise information of materials composition as a stack of X-ray CT images defining a three-dimensional volume [1, 2]. Depending on the use case, in order to explore and obtain information from these volumetric datasets, specific methods have to be designed; a design that can be challenging depending on the application, specially, when pure 3D-based approaches want to be considered. In this paper, our interest has been focused on the quantification of fibre waviness: a geometrical defect in fibre reinforced polymer composites with direct implications on mechanical properties as exposed in reviews such as [3, 4].

Generally, fibre waviness evaluation methods consider a μCT input volume, and analyse it with a two step process, where first, an approximation of the fibres is obtained and then measurements over the reconstructed fibre orientations are performed. They define (or implicitly assume) a set of cutting planes over the μCT volume, ideally, perpendicularly oriented to the principal fibre directions. Then, some evaluate the fibre sections which will be circular in case of normally aligned fibres and elliptical otherwise [5, 6]. This is usually complemented by the more generalist approach of detecting the centres and then knitting or correlating them between slices in order to obtain their in-plane (the deviation of fibre from the normal fibre direction in a plane of the lamina) and out-plane (bending of a single or multiple fibre separated from the laminate fibrous layer) deviations [7, 8, 9, 10, 11, 12, 13, 14]. There are also alternative methods such as [15, 16] capable of reconstructing fibres. Other approaches quantify the waviness directly at every voxel avoiding the complexities of producing polygonal fibre trackings. These voxel-based methods, generally, after a fibre segmentation process, analyse the neighbourhood around each voxel using kernels and applying the tensor or similar strategies in order to obtain the directions of least image diffusion (collinear with the fibres) [17, 18, 19, 20, 21, 22, 16]. Others transform the volume using Hough, Fourier, Radon, or similar approaches to detect the fibre directions and distributions [23, 24, 25, 26]. A key advantage of these voxel-based methods is their ability to perform measurements on challenging or low resolution datasets where individual fibres are not distinguishable. Unfortunately, they usually lack the ability to treat fibres as separate entities.

In any case, once the fibre orientation can be determined, in-plane and out-plane misalignment angles can be obtained as $\alpha = \arctan\left(\frac{u}{l}\right)$ and $\beta = \arctan\left(\frac{v}{l}\right)$, respectively, where l is the distance between two consecutive slices (or samples), with u and v corresponding to the two out-plane and in-plane local displacement vectors [27, 7, 23, 25, 9, 28]. From these, measures such as the severity factor [29], waviness, or wavelength ratios [30, 28, 31, 32, 33, 34, 3, 35, 4] can be computed. These approaches have extensively been used to evaluate the effect of misalignments to mechanical properties. They have been applied not only on μCT data but also on other imaging modalities such optical microscopy [23, 25, 12]. Despite its advantages, most approaches reduce the 3D fibre information to 2D counterparts by emphasizing on the angles or sinusoid shapes. However, as performed by [8] there is the possibility to consider a more fundamental measurement for 3D curves such as the Frenet-Serret [36, 37] κ curvature formula. This assumes that each polygonal fibre of length l is a continuous function where at any arc-position location h a $\mathbf{P}_h$ point can be obtained by a $\mathbf{F}(h)$ interpolatable function. Building upon this continuous curvature definition a first proposal of this paper is a new measurement to summarize the whole fibre into a single scalar which can serve as an indicator to determine the misalignments relevance.

Unfortunately, varying aspects such as the scale of imaged fibres, noise, resolution, sampling, etc. will have an impact on the measurements if these are directly computed over the input data. When one considers the fractal nature [38] of the noisy imaged fibres, then, a more finely sampled polygonal curve will inevitably consider more roughness and thus overestimate the measure. As this dependence on acquisition parameters is troublesome, a second proposal of this paper is a method to decouple it by smoothing the fibre trackings using a low-pass Gaussian convolution and binding its standard deviation σ parameter to the underlying dataset space units, i.e. nm, µm, etc. Note that most of the discussed methods, including the voxel-based ones do not take into account the aforementioned issue. Our approach overcomes these limitations by smoothing and interpolating respect physical space units, to achieve a decoupling from acquisition and/or reconstruction parameters.

The proposed measure and method have been tested using real and synthetic datasets to properly asses the behaviour and filtering capabilities. These synthetic cases were produced from exact definitions of curves swept using Gaussian splatters in order to generate slices mimicking real acquisition while controlling the morphological details from a signal processing point of view.

2. Material and methods

2.1. A new measure for waviness evaluation

As mentioned, the waviness of a 2D fibre $\mathbf{F}$ has been usually estimated by the maximum misalignment angle θ_{max} defined by:

$$\theta_{max} = \max_x \measuredangle(\mathbf{F}'(x), \overrightarrow{i}),$$

where $\overrightarrow{i}$ represents the vector $(1, 0)$. However, in real composite materials, fibre-axis are ideally assumed to have a sinusoidal shape. That is, fibre-axes are modelled by

$$\mathbf{F}(x) = \left(x, A\sin\left(\frac{2\pi x}{\lambda}\right)\right) \quad x \in [0, \lambda],$$

where A and λ correspond the amplitude and wavelength, respectively. In this case, the waviness severity is measured by the ratio A/λ. Hsiao and Daniel in [31] made the following observation $\theta_{max} = \arctan(2\pi A/\lambda)$, and concluded that in unidirectional composites the major Young's modulus is severely degraded as the ratio A/λ increases. From this work, it can be derived that any measure of a fibre also increasing with respect to A/λ or θ_{max} (for instance), can be employed to assess the waviness severity. However, this θ_{max} parameter may not be sufficient to describe a fibre and its behaviour, sometimes, can be misleading.

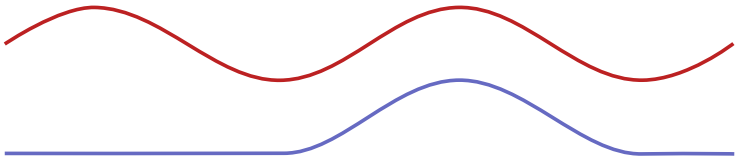

Figure 1: Two fibres of similar scale, shape and length producing the same θ_{max} angle despite being different.

For instance, in Figure 1 the axes of two different fibres with the same θ_{max} are illustrated, however the blue one is less wavy than the red. Consequently, it seems necessary to make use of more geometric parameters to obtain a real description of the fibre characteristics. For this reason, in order to obtain a new measure of the waviness we propose using the total curvature K along the fibre as in Equation 1, and then study how it is related with the θ_{max} parameter.

$$K = \int_{\mathbf{F}} \kappa \, d\mathbf{F} \tag{1}$$

First, let's study how the angle θ_{max} behaves under the assumption of a fibre shaped as a circular arc. In Figure 2a one can observe that

$$\theta_{max} = \frac{l}{2R} = \frac{l}{2}\kappa,$$

with $\kappa = 1/R$ being the curvature of this circular arc. If this arc were to be composed by a finite number of sub-arcs of lengths l_i and central angles θ_i (see Figure 2b), θ_{max} could be computed as follows:

$$\theta_{max} = \frac{1}{2}\sum_i \theta_i = \frac{1}{2}\sum_i l_i\kappa = \frac{1}{2}\int_0^l \kappa \, dh = \frac{K}{2}$$

Therefore, it seems that the total curvature is related with the fibre waviness severity. To further check this idea if we assume a fibre modelled by a sinusoid, the following must hold:

$$\mathbf{F}(x) = \left(x, A\sin\left(\frac{2\pi x}{\lambda}\right)\right), \qquad \mathbf{F}'(x) = \left(1, A\cos\left(\frac{2\pi x}{\lambda}\right)\frac{2\pi}{\lambda}\right), \qquad \mathbf{F}''(x) = \left(0, -A\sin\left(\frac{2\pi x}{\lambda}\right)\left(\frac{2\pi}{\lambda}\right)^2\right).$$

Then,

$$K = \int_{\mathbf{F}} \kappa d\mathbf{F} = \int_0^{\lambda} \frac{|\det(\mathbf{F}'(x), \mathbf{F}''(x))|}{\|\mathbf{F}'(x)\|^3} \|\mathbf{F}'(x)\| dt = 2 \int_0^{\lambda/2} \frac{A \sin\left(\frac{2\pi x}{\lambda}\right)\left(\frac{2\pi}{\lambda}\right)^2}{1 + A^2 \cos^2\left(\frac{2\pi x}{\lambda}\right)\left(\frac{2\pi}{\lambda}\right)^2} dx = 4 \arctan\left(\frac{2\pi A}{\lambda}\right) = 4\theta_{max}\,.$$

So, if the fibre was made up of n periods of λ wavelength, we would have $K = 4n\theta_{max}$. In this way, the blue fibre in Figure 1 would have waviness of $4\theta_{max}$, while the red one $8\theta_{max}$. For this reason, it can be concluded, that the total curvature is a better descriptor of the waviness in fibres.

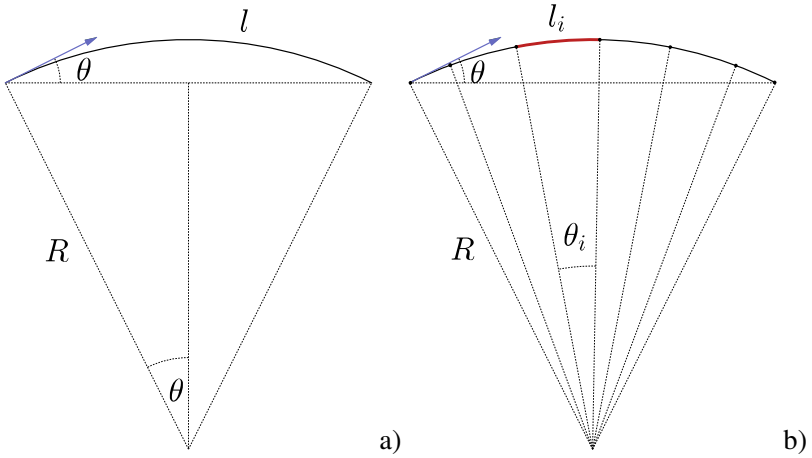

Figure 2: Maximum misalignment angle of a circular fibre arc.

However, note that the fibres depicted in Figure 1 have similar length. So, in order to produce a more intuitive scalar for the whole fibre indicating the waviness independently of the fibre length; we propose to consider the waviness as the average curvature $\overline{K}$ along the fibre, as in Equation 2 where l corresponds to the fibre length.

$$\overline{K} = \frac{K}{l} = \frac{1}{l} \int_{\mathbf{F}} \kappa d\mathbf{F} \tag{2}$$

2.2. Algorithm

Given a set of polygonal trackings representing real fibres we want to obtain a continuous curvature indicator κ along the fibre, and a single K waviness scalar summarizing the whole fibre itself. To compute these, as illustrated in Figure 3 (dashed box), a preliminary step to reconstruct the fibres to polygonal paths is required. Although different methods can be applied, in our case we used the Julià et al. [16] method. This considers the input volume as a whole, employs an uniformly behaving 3D reconstruction algorithm, and is capable of seamlessly tracking complex curves using variable step sizes.

Irrespective to the reconstruction algorithm employed, the method assumes the input fibres as a set of 3D curve functions, each independently represented by a chain of irregularly sampled points $\{\mathbf{Q}_0, \cdots, \mathbf{Q}_i, \cdots, \mathbf{Q}_n\}$ conforming a path of $n-1$ linked segments up to length of l space units mapped by a function $\mathbf{F}$ which given a parametric coordinate $h \in [0, l]$ produces an interpolated $\mathbf{F}(h)$ point between the neighbouring $\mathbf{Q}_i$ samples. In what follows, h_i will denote the arc-position parameter snapped to a discrete sample, thus satisfying $\mathbf{F}(h_i) = \mathbf{Q}_i$. With these definitions, the curvature at any position $\mathbf{F}(h)$ in a fibre can be defined by the κ parameter extracted using the Frenet-Serret formulas [36, 37] as in Equation 3. Moreover, the average curvature $\overline{K}$ along the fibre, that is, the new measure of the waviness proposed, can be computed as in Equation 4.

$$\kappa(h) = \frac{\|\mathbf{F}'(h) \times \mathbf{F}''(h)\|}{\|\mathbf{F}'(h)\|^3} \tag{3}$$

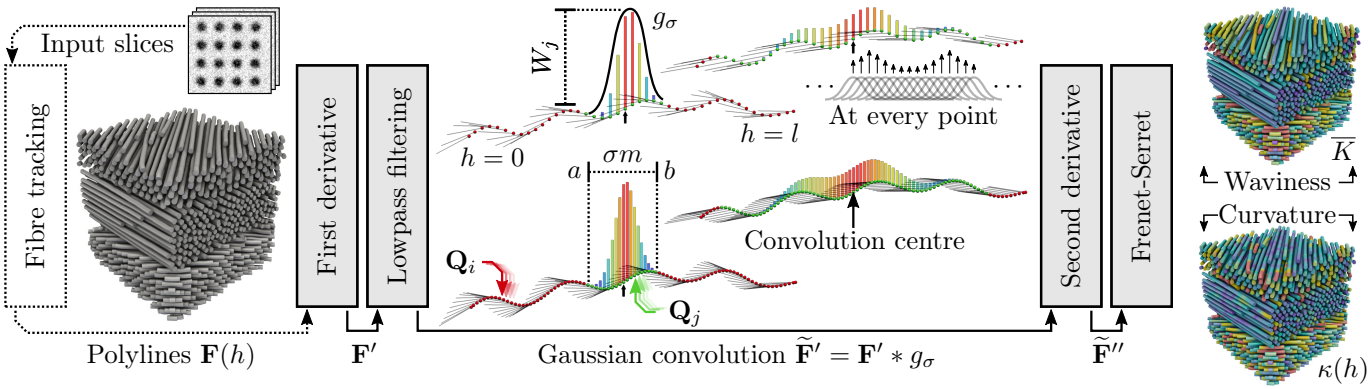

Figure 3: Method pipeline starting from the input slices (top left) up to the final measurements (right) with a detailed overview of the convolution state at two different σ and step sizes as seen from a single point (dark arrow) which is convolved according to the Gaussian distribution g_σ weights (coloured bars) until a σm distance (green dots) is reached. Fine black lines show the orientation of the first derivative. The curvature is computed for each fibre position (bottom right) or can be integrated into a single waviness scalar for the whole fibre (top right).

$$\overline{K} = \frac{1}{l}\int_0^l \kappa(h)\|\mathbf{F}'(h)\|\,dh \tag{4}$$

Note that these measures, if applied directly, will overestimate the results depending on undesired factors such as the input volume noise level, sampling, etc. For this reason, as illustrated in Figure 3 (filled grey boxes) the following four steps are applied: (i) the first derivative is computed; (ii) curvatures below a certain size are attenuated; (iii) the second derivative is computed; and (iv) the aforementioned curvature and waviness formulas are calculated for every tracked fibre. In the next sections we describe each one in more detail.

2.2.1. *First derivative (tangent vector)*

The algorithm begins by computing at each discrete $\mathbf{Q}_i$ sample an approximation $\mathbf{F}'(h_i)$ of the first derivative respect the arc distance h in space units. When three samples are available (i.e. $0 < i < n$), the first derivative at h_i is computed using the second order Lagrange interpolating polynomial using the pairs $(h_{i-1}, \mathbf{Q}_{i-1})$, $(h_i, \mathbf{Q}_i)$ and $(h_{i+1}, \mathbf{Q}_{i+1})$. Otherwise, at $\mathbf{Q}_0$ and $\mathbf{Q}_n$ extremes, the approximation is similarly computed using two samples and a first order interpolating Lagrange polynomial. Note that the results of this step, despite being valid vary upon undesired external factors such as the volume resolution, fibre scale, sampling frequency (i.e. step size), etc.

2.2.2. *Low-pass filtering*

As ignoring the aforementioned externalities will lead to ill-formed measurements, the first derivative is convolved (Equation 6) using a Gaussian distribution (Equation 5) whose σ parameter smoothly adjusts (in space units) the scale below which perturbations and thus high frequency information shall be attenuated, and thus decoupled from the underlying sampling and/or acquisition parameters.

$$g_\sigma(x) = \frac{1}{\sigma\sqrt{2\pi}}\mathrm{e}^{-\frac{x^2}{2\sigma^2}} \tag{5}$$

$$\widetilde{\mathbf{F}}'(h) = \mathbf{F}' * g_\sigma = \int_{-\infty}^{+\infty} \mathbf{F}'(s)\, g_\sigma(s-h)\,ds \tag{6}$$

This smoothed derivative $\widetilde{\mathbf{F}}'$ needs to be approximated at each discrete sample $\mathbf{Q}_i$. To do this, the integration interval (i.e. convolution window size) is restricted to have at most a reach of σm space units, with the m parameter being a user-adjustable setting trading off performance versus accuracy. In parametric coordinates, this $[a, b]$ interval is determined by:

$$a = \max\left(0,\, h_i - \frac{\sigma m}{2}\right), \qquad b = \min\left(l,\, h_i + \frac{\sigma m}{2}\right).$$

Then, for each discrete sample $\mathbf{Q}_j$ whose parameter h_j lies in $[a, b]$, a weight W_j is computed by:

$$W_j = \int_{\frac{h_{j-1}+h_j}{2}}^{\frac{h_j+h_{j+1}}{2}} g_\sigma(s - h_i)\, ds\,.$$

Finally, $\widetilde{\mathbf{F}}'(h_i)$ is approximated as follows in Equation 7a. In addition, one can refer to Equation 7b for a more generalized definition.

$$\widetilde{\mathbf{F}}'(h_i) \simeq \frac{\sum_j \mathbf{F}'(h_j) W_j}{\sum_j W_j} \tag{7a}$$

$$\widetilde{\mathbf{F}}'(h_i) \simeq \frac{\int_a^b \mathbf{F}'(s - h_i) g_\sigma(s - h_i)\, ds}{\int_a^b g_\sigma(s - h_i)\, ds}\,, \quad ds = \text{variable step size} \tag{7b}$$

Figure 3 illustrates how the reach (green dots) and weights (bars) of the convolutions become independent from the underlying sampling. In addition the smoothing effect in tangent vectors (thin lines) when σ is increased can be observed by visually comparing between the four examples.

2.2.3. Second derivative (normal vector)

In this third step, and with the undesired high frequency information removed, the method proceeds to calculate the second derivative $\widetilde{\mathbf{F}}''(h_i)$ from the smoothed first derivative $\widetilde{\mathbf{F}}'(h_i)$ results by applying the same method explained in Section 2.2.1. Note that the normal vector produced will have an unstable direction when the fibre travels straight as illustrated by the ribbons normals in Figure 6 (right) and Figure 9 (bottom row).

2.2.4. Computing the curvature and waviness

In this last step, as our implementation employs the VTK (Visualization ToolKit) and each fibre is represented by a polyline, data values can be associated at two levels: (i) for every $\mathbf{Q}_i$ sample (point data), (ii) or to the whole line $\mathbf{F}$ (cell data). Consequently, the tangent $\widetilde{\mathbf{F}}'(h_i)$ and normal $\widetilde{\mathbf{F}}''(h_i)$ vectors are stored as point data, alongside the curvature measurement $\kappa(h)$ defined by Equation 3 and computed as expressed by Equation 8.

$$\kappa(h_i) = \frac{\|\widetilde{\mathbf{F}}'(h_i) \times \widetilde{\mathbf{F}}''(h_i)\|}{\|\widetilde{\mathbf{F}}'(h_i)\|^3} \tag{8}$$

Finally, this point curvature can be integrated using the trapezoidal rule into a scalar summarizing the average waviness $\overline{K}$ (see Equation 9) and stored as cell data information (i.e. single value assigned to the whole fibre).

$$\overline{K} \simeq \frac{1}{h_n} \sum_{i=0}^{n-1} \frac{\kappa(h_i)\|\widetilde{\mathbf{F}}'(h_i)\| + \kappa(h_{i+1})\|\widetilde{\mathbf{F}}'(h_{i+1})\|}{2} (h_{i+1} - h_i) \tag{9}$$

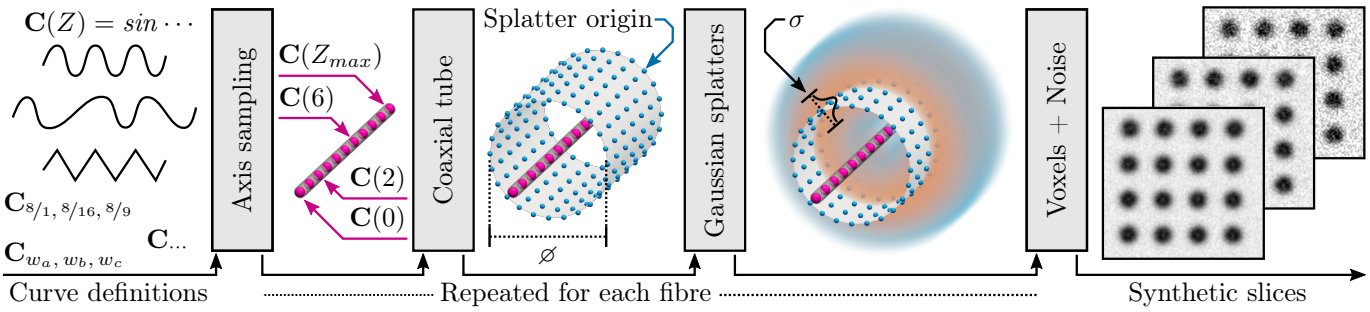

Figure 4: Each curve definition is sampled at discrete points (magenta) which are then radiated outwards in order to produce a coaxial tube where at each of its surface vertex (blue) a Gaussian blur is splattered, and then sampled into the resulting synthetic slices (right).

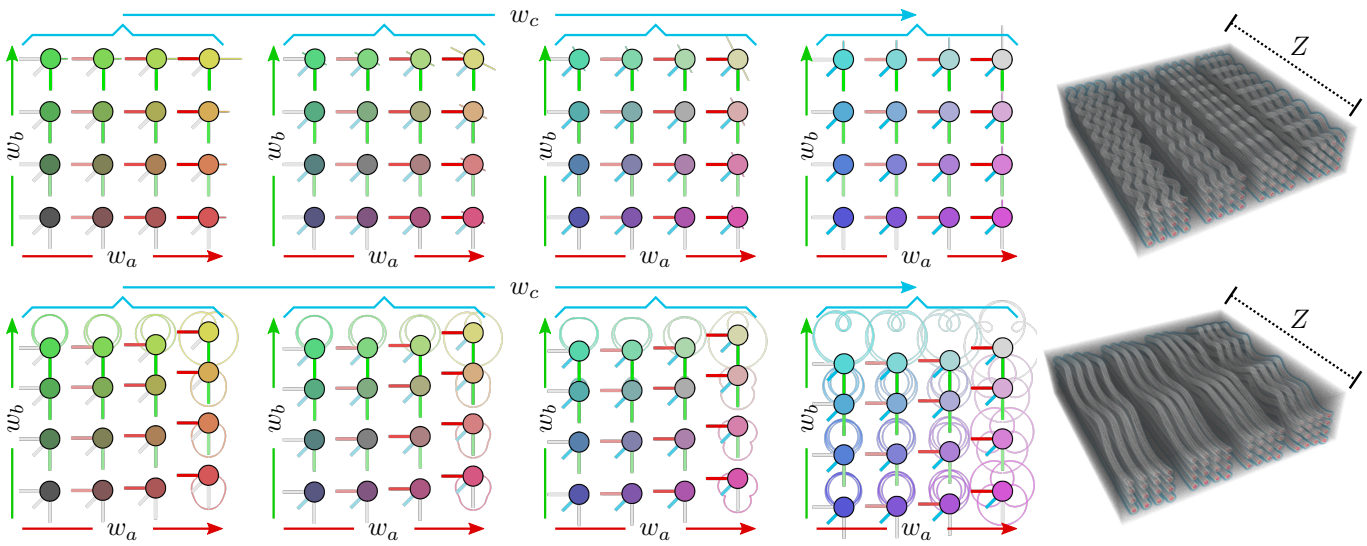

Figure 5: Arrangement as seen from the X-Y plane of the three weighting variables: w_a (red, columns), w_b (green, rows), and w_c (blue, blocks); with circles colour-mapping them in RGB respectively; combined with an overlay and a thumbnail of the produced curves. The *polarized* (top) and *harmonics* (bottom) cases are illustrated.

2.3. *Experiments and testing datasets*

The proposed method has been tested using the real *cfrp_i17* dataset with three straight fibre orientations and acquired with the European Synchrotron Radiation Facility's ID19 Beamline (ESRF, Grenoble) at 26 keV, and pixel size of 650nm. It has been cropped in two different $250 \times 250 \times 250$ and $500 \times 500 \times 500$ sub-volumes named *cfrp250* and *cfrp500* respectively and made available in Supplementary Material 1. However, in order to properly assessing the invariance and filtering capabilities of the proposed algorithms, four synthetic datasets named *harmonics*, *hifreq*, *mixed*, and *polarized*, have been generated in order to produce a set of complex fibre paths by morphing and combining characteristics (frequencies) at different scales.

2.3.1. *Synthetic fibre generation*

As illustrated by Figure 4, synthetic datasets are generated using a three-phase process which begins by: (i) discretely mapping (in our case along the Z slicing axis) a medial axis curve $\mathbf{C}$ in order to obtain a finite set of high resolution points (pink dots); (ii) sweeping a coaxial tube around it to produce a new surface (white); and (iii) splatting Gaussian blurs of σ size at each surface vertex (blue dots). These steps are repeated for each input $\mathbf{C}$ such that a diffuse cloud of intensities is then quantized into a finite resolution volume where random white noise is mixed before producing the image slices. Depending on the relationship between the σ and the tube diameter $\varnothing$, the central intensity decay will be more or less accentuated; a useful feature mimicking beam-hardening artifacts [1] (Section 6.1) .

2.3.2. *Synthetic cases*

With the ability to produce synthetic fibres from an exact $\mathbf{C}$ function, the next step is defining and arranging several of them in a volume. Most of our synthetic datasets are organized in 4 different blocks of 4×4 fibres, where each position corresponds to a unique combination of the three w_a, w_b, and w_c weighting variables in order to obtain 64 unique (but similar) $\mathbf{C}_{w_a, w_b, w_c}$ fibres exhibiting different amplitudes for each tunable feature. Every fibre function is sampled along the Z-axis, and shifted on the X-Y plane as illustrated in *Figure* 5 in order to achieve the aforementioned arrangement. These volumes have a size of $512 \times 128 \times 512$ voxels and a dynamic range of 8 bits.

While all synthetic datasets can be expressed as a mathematical equation, the simpler *harmonics* case is expressed in Equation 10 where a fibre combines three circular oscillations: (i) a lower frequency one, turning once from start to end; (ii) a mid frequency revolving twice; and (iii) a higher one oscillating four times. Each is respectively adjusted according to the w_a, w_b, and w_c weights and arranged in rows, columns and blocks as disposed in Figure 5 obtaining a total of

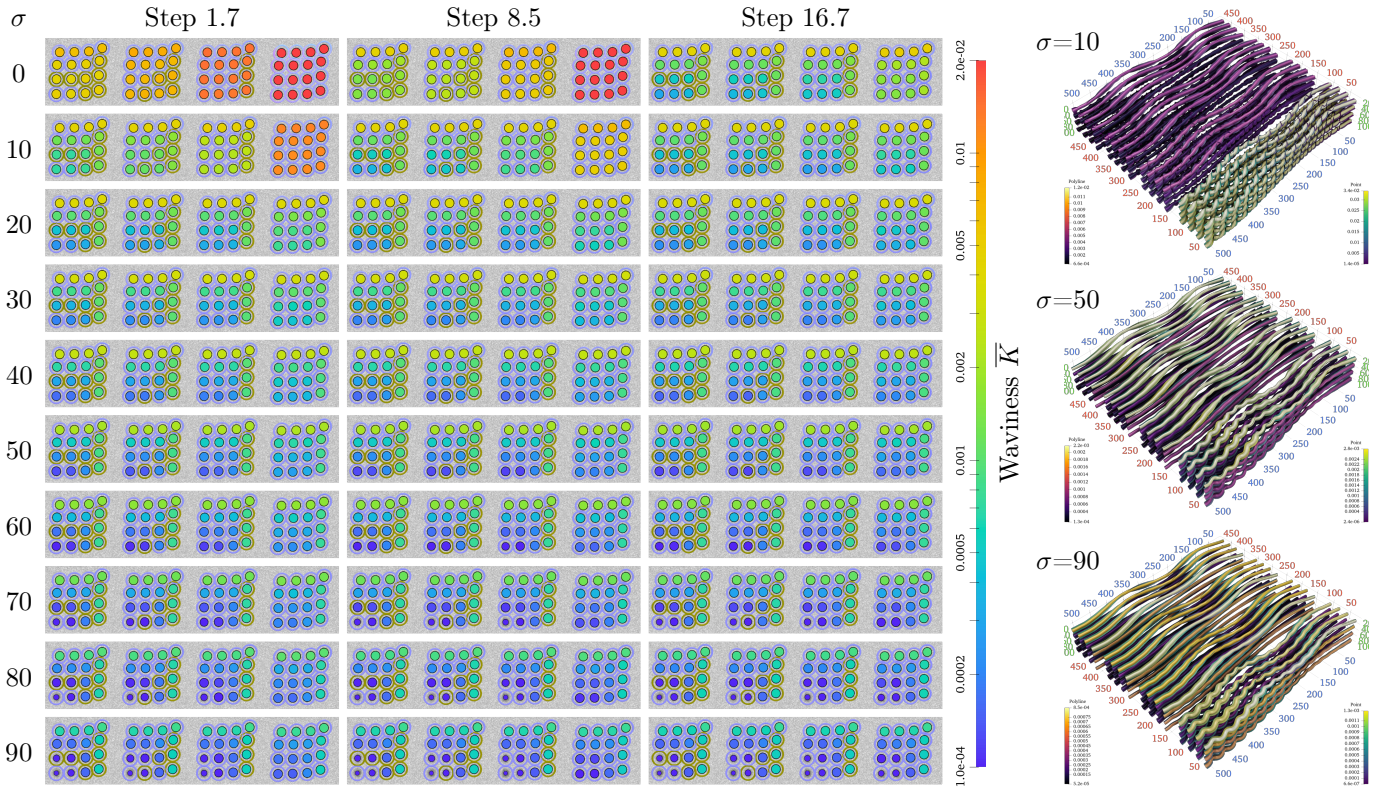

Figure 6: On the right, a 3D visualization of the waviness (tube colour), normal vectors (ribbon orientation), and point curvature (ribbon colour) at three different σ values. On the left side, a matrix of slices comparing the waviness at different σ (rows) and sampling step sizes(columns). The waviness of each fibre (circle colour) is represented using the same scale to ease the comparison. Perimeter lines around circles RGB map the tangent vectors components.

$4 \times 4 \times 4 = 64$ fibres as different 4 values for each weight are combined. In this way, one can easily analyse and isolate each variable; for instance, the effects of high frequencies are perceived observing the differences between blocks.

$$\begin{aligned} \mathbf{C}_{w_a, w_b, w_c}(z) = (0,\, 0,\, z) \quad &+ \quad (\sin(2\pi z/Z_{max}),\, \cos(2\pi Z/Z_{max}),\, 0)\, w_a \\ &+ \quad (\sin(4\pi z/Z_{max}),\, \cos(4\pi Z/Z_{max}),\, 0)\, w_b \\ &+ \quad (\sin(8\pi z/Z_{max}),\, \cos(8\pi Z/Z_{max}),\, 0)\, w_c \end{aligned} \quad , \quad \begin{cases} w_{a,b,c} & \in \{8/1,\, 8/4,\, 8/9,\, 8/16\} \\ Z_{max} & = 512 \\ z & \in [0,\, Z_{max}] \end{cases} \tag{10}$$

The *polarized* dataset is characterized by its oscillations in a single plane. It employs the same methodology for arranging the w_a, w_b, and w_c weights as in Figure 5 in order to: (i) encode a lower frequency turn from start to end in the rows w_a; (ii) a high frequency (8 turns) in the columns w_b; and (iii) a Z-axis rotation of the oscillating plane between blocks w_c.

The *hifreq* dataset reuses the same low w_a and mid w_b frequency components of the *harmonics* case, but regarding the very high frequency w_c feature varied for each block, is a more complex combination of circular and linear oscillations whose amplitude gradually decays to zero as one of the fibre extremes is reached. This incremental waviness shall be observable in the local curvature.

Finally, the *mixed* dataset combines and exaggerates the features of the aforementioned cases in order to produce complex shapes at different scales.

3. Results and discussion

The proposed measurement algorithm has been implemented in C++ as the VTK [39] filter presented in Supplementary Material 2, which has been integrated in the Starviewer [40] platform as illustrated in Figure 10. In Supplemen-

tary Material 1 the results, slices, figures, extra visualization, plots, and the code to produce the synthetic cases alongside its ground truth is made available to ensure the reproducibility.

3.1. Scale selectivity

A key distinctive feature of the proposed algorithm is its ability to suppress fibre variations below a certain size employing a low-pass Gaussian filter. Its behaviour is better understood when a **F** fibre curve is observed in the frequency space (i.e. Fourier-transformed) as a signal where high frequencies shall be eliminated. With the advantage of avoiding computationally intensive transformations; the Gaussian convolution employed behaves as a low-pass filter suppressing fibres below a user-specified σ wavelength in space units. These capabilities are demonstrated in Figure 6 using the *hifreq* dataset containing an increasing amount finer variations at each block from left to right (refer to Section 2.3.2 and Figure 5). When the σ parameter is smaller than the wavelength (first rows) the waviness measurements are distributed across blocks and the filter has little effect. But, as σ increases, as expected, these higher frequencies are attenuated, the waviness is decreased, and its relative distribution shifts from the blocks to the top rows (which have a mid-frequency). Then (in this example) at $\sigma > 30$, the lowest frequencies (columns) become more dominant and distinguishable.

If one observes the remaining comparative figures in Supplementary Material 1 the following key features of each dataset are fulfilled: (i) in *harmonics* a gradual and smooth transition from blocks to the upper-right corners where more curvature is present; (ii) in *polarized*, as expected the blocks present the same curvature, although with slight asymmetries in the intermediately rotated middle blocks; (iii) in *mixed* the desired selectivity is visually appreciated; and (iv) in the *crfp* cases, the outliers can be properly detected by their curvatures, and its main slightly curved regions can be identified at big σ values (see first row in Figure 9).

3.2. Definition of the signal and noise

From a signal processing perspective, the highest frequencies will not describe the actual shape of physical fibres, but instead manifest varying acquisition parameters such as voxel size, scale of the fibres, resolution, dynamic range, noise characteristics of the imaging device, etc. Hence, choosing a proper σ parameter to decouple these from our measurements is crucial to obtain comparable and non-overestimated measures across acquisitions. However, as illustrated in Figure 6 (rows) and Figure 3 (tangent lines) employing a σ well above the acquisition noise is also useful to discard physically present features. So, the σ parameter (arbitrarily set by the user) defines boundary between what shall be considered noise (higher frequencies), and signal (lower frequencies). As it specifies a wavelength in space units, it can be employed as an invariant reference value across different acquisitions. For these reasons, values below the sampling size and/or fibre diameter shall be discouraged.

3.3. Sampling

While the Gaussian convolution method is able to cut high frequencies, these must exist in the first place. In practice, fibres will never be infinitely sampled, consequently, they will be band-limited due to the Whittaker–Nyquist–Shannon sampling theorem [41, 42, 43]. This is an observable fact in Figure 6 where the smaller σ in coarser cases do not have any effect until its value is greater than the sampling spacing. However, once exceeded, the measurements converge irrespective of the sampling as illustrated by Figure 8. In the employed Julià et al. [16] reconstruction method, steps sizes are user adjustable and may vary within a user-defined range. But, most reconstruction algorithms discussed in Section 1 implicitly set their sampling frequency by design choices such as snapping to integer voxel locations, reconstructing in a slice-by-slice approach, etc. In any case, all reconstruction methods, in addition to the ones performing measures voxel-wise, will have their highest frequencies capped according to the input volume resolution. However, relying on the sampling alone or other non-space referenced smoothing techniques (e.g. moving average of discrete samples) undermines the meaning of the measurement across different dataset acquisitions leading to a disparity in the results as observable in Figure 7 when our method is set at $\sigma = 0$ (i.e. filtering disabled). For this reason, when one considers the fractal nature [38] of the noisy imaged fibres, then, a more finely sampled polygonal curve will inevitably consider more roughness and thus overestimate the measure.

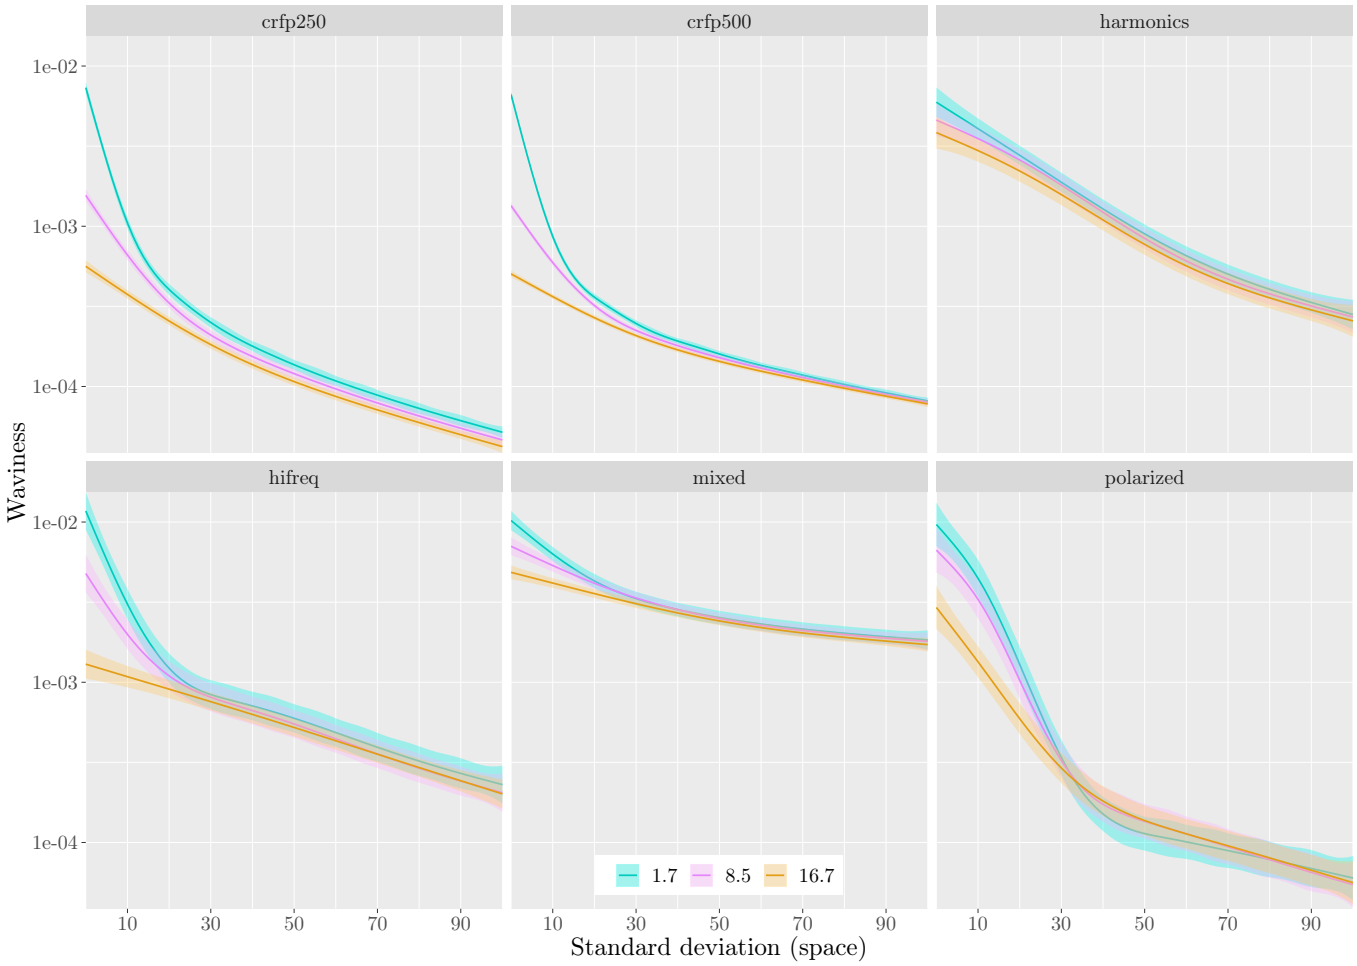

Figure 7: For every dataset (boxes), a comparison of the mean waviness results (Log10 Y-axis) for each combination σ (X-axis) and step size (colours) values. Confidence intervals (semi-transparent colours) up to the 99th percentile.

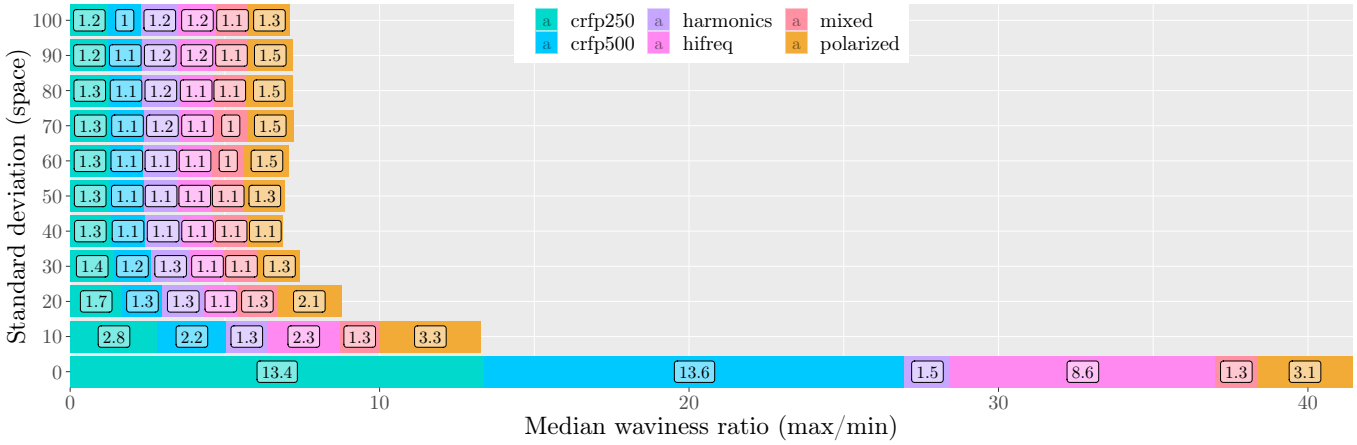

Figure 8: For each case (colours), ratio (X-axis) between the maximum and minimum median values at different σ (Y-axis) of the whole volume waviness (labels, bar size) grouped by each step size tested. In practice, it is the ratio between the 1.7 and 16.7 steps sizes. Values approaching 1 indicate convergence of the measurement irrespective the sampling.

3.4. Numeric stability and edge behaviour

One of the first aspects to consider is the finite lengths of the fibres; this approach was tackled by normalizing the vectors of the smoothed first derivatives in order to take in account the otherwise missing weight at the limits. Due to the nature of the curvature formula, its results are numerically stable even when a fibre travels straight (i.e. it also approaches zero), however, as future work we plan to further extend our measurements and provide more advanced visualizations. Also, relying on a pure Frenet-Serret frame is troublesome due to its abrupt changes in the normal vectors direction (see last row in Figure 9). We plan to build upon slightly altered approaches such as [44] to mitigate these aforementioned issues, or by employing rotation minimizing frames such as [45].

3.5. Datasets and testing

The usage of datasets whose ground truths curves were analytically defined in terms of frequencies has been crucial to assess whether the theoretical expectations matched the actual measurements of the algorithm. Although these curves alone could have been directly fed into the measurement algorithm for testing purposes, this would have ignored the effect of white pixel noise, whose maximum frequency is in turn band-limited by the volume resolution. This noise affects the Julià et al. [16] tracking algorithm and it is manifested as small variations when a finer step size is employed (an expected behaviour). Despite the availability of datasets such as [6] with richer features better mimicking aspects such as beam hardening, we opted for a more conservative approach using Gaussian splatters in order to ensure an isotropic noise behaviour as the complex shapes are sampled into finite resolution slices. In addition to the volume and ground truths; in Supplementary Material 1 one will encounter the generative pipeline implemented in Paraview [46] to ensure the reproducibility by anybody and ease the creation of new datasets.

As illustrated in Figure 9 visualizations, the method has been tested using a real dataset, and while the results and behaviour are coherent with synthetic cases, as future work we plan to extend the testing by combining specimens with varied fibre diameters and altering acquisition parameters affecting resolution, contrast, noise levels, etc.

3.6. Cost

The algorithm cost grows linearly respect the overall number of samples within a volume, but the amount of operations required at each one will depend on the amount of discrete neighbours convolved in a σm arc distance. While similar to discrete convolution kernels (e.g. Gaussian blur of an image), our kernel size cannot be precisely known nor allocated beforehand. As illustrated in Figure 3, the active voxel (upwards black arrow) must access its discrete neighbours (green points) around a $\frac{\sigma m}{2}$ arc distance in space units. As by definition the spacing is irregular (even within a polyline), a reasonable way to quantify the cost is applying Equation 11 using estimative averages of the fibre step sizes $\overline{ds}$, samples $\overline{n}$ and length $\overline{h}$.

$$\text{Operations per arc unit length} \simeq \frac{\overline{h}\sigma m}{\overline{n}\,\overline{ds}} \tag{11}$$

While this linear behaviour is comparable to the more straightforward approaches discussed in Section 1, in practice, larger convolutions in our method will severely impact the runtime. However, with larger σ values respect the step size, one may voluntarily use coarser step sizes in order to reduce them, and thus indirectly pre-filter the highest frequencies as discussed in Section 3.3. Finally, the implemented VTK filter in Supplementary Material 2, although usable, does not support multithreading yet, and as future work we plan to parallelize it at a fibre level, i.e. by polydata cells.

4. Conclusions

In this paper we have presented a method to assess the point curvature and overall waviness of fibres represented as polygonal curves reconstructed from μCT slices. While the required input is independent respect the underlying tracking methodology, we used an inherently 3D one with the ability to produce irregularly spaced samples in order to ensure a high level of compatibility should any other state-of-the-art tracking method be used in-place.

Our Frenet-Serret based measurements ensure an inherent 3D behaviour which has been proven numerically stable for the measurement of fibre misalignments. Additionally it has the advantage of being a decoupled measure from the world-orientation of the dataset; in this way, there is no need to perform fine realignments to the specimens across acquisitions

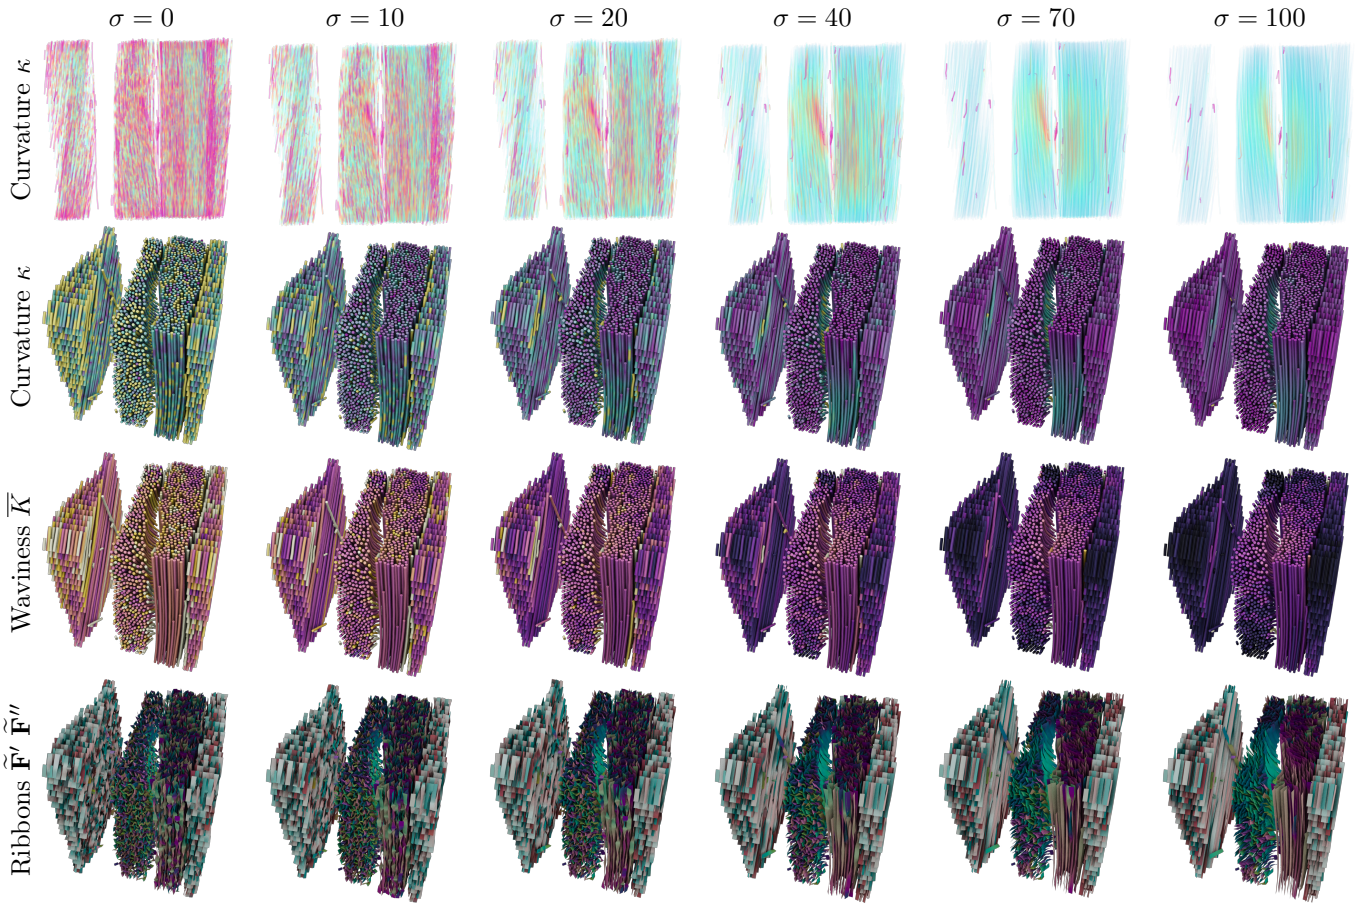

Figure 9: Results and visualizations for the *crfp500* real dataset. Each column in the matrix corresponds to a different σ value. Colour scales are not rescaled within a row. The first row maps the curvature κ using semi-transparent lines; the second one also maps κ employing 3D tubes around the reconstructed curves; the third one maps the waviness $\overline{K}$ as a single colour for the whole fibre; and the last one employs ribbon surfaces to express the second derivative $\widetilde{\mathbf{F}}''$ orientation combined with an RGB colour mapping of the first derivative $\widetilde{\mathbf{F}}'$.

thus avoiding the introduction of further error. Also, we have demonstrated the relation of our proposal with common waviness measurements approaches which assume the fibres as simpler waveforms coplanar to a plane. This ability to address more complex cases with a more intuitive or expected behaviour makes our results differ numerically from other measurements; however its mechanical implications can be deemed equivalent.

A key concept introduced in this paper is the importance of setting a limit to the smaller curvatures, and considering only the ones above a certain size. In our case, this scale selectivity is attained using a Gaussian low-pass frequency filter with an user-adjustable sigma parameter corresponding to a wavelength in space units. However, most state-of-the-art measurements in composite fibres ignore this fact, rely only on the fibre resolution (i.e. sampling frequency) to limit the over-estimative effect of high frequencies, or may employ more advanced techniques, such as moving averages or polynomial interpolations not bound to actual space units making the results variant respect the acquisition parameters. Solving these issues and attaining coherent measurements across datasets can be challenging, especially with varying factors such as noise, resolution, sampling, etc. Our method attains convergence of results for fibres sampled at different resolutions when a sufficiently large sigma above the sampling frequency is employed. This adjustable low-pass filter, while a valuable tool to suppress the aforementioned fractality effects of high frequencies, can also be used to suppress smaller but undesired curvatures already present in the physical fibres deemed irrelevant for a particular case of study.

While the behaviour is explainable by well established signal processing fundamentals, these have been experimentally tested using one real and several synthetic datasets. For this reason, we also introduced a methodology to generate synthetic volumes from a given ground truth defined as a composition of several oscillations. Although these synthetic

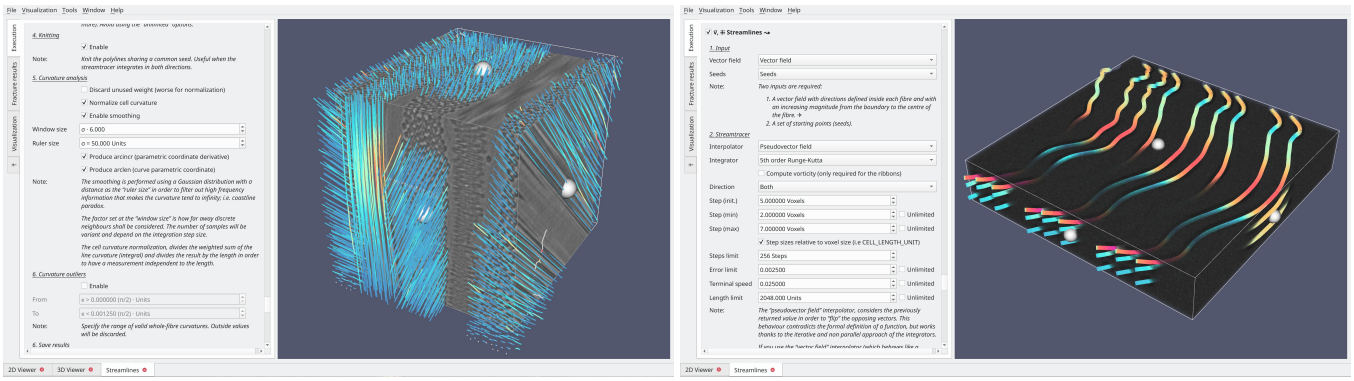

Figure 10: The measurement method has been integrated in Starviewer [40, 16]. All algorithm parameters are exposed to the end user (left sidepanels) who can later interactively explore and validate the results using fused volume renders of the underlying datasets. The *crfp500* (left) and *mixed* (right) cases are illustrated.

slices do not reflect the richness of defects present in real acquisitions, their uniform behaviour and wave-based definition have been valuable tools to validate the theoretical expectations with actual results.

Our future work will be focused on further improving the algorithm performance and visualizations while combining it with other measurements employing rotation minimizing frames; performing the analysis at multiple scales; or expressing the curvature as an spectrum by binning the spectrum of spatial frequencies composing the fibre shapes.

Acknowledgements

The authors want to thank the support from AMADE-UdG *(Analysis and Advanced Materials for Structural Design)* research group.

Declarations of interest

The authors declare no conflict of interest.

Funding sources

This work has been financially supported by grants from the Spanish Government *(Ministerio de Ciencia, Innovación y Universidades)* `PID2019-106426RB-C31`.

Supplementary material

1. **results.tar.xz** Contains the code to generate synthetic volumes, the actual datasets, results for all tested combinations, code, figures, and extra visualizations.

   `c5c3d4b37f103988e52a1a4d0e831286af4fe16e6512b2f765930e43dd513e8d (SHA256)` `ef36a0d68e734a5cc1e84e439cf7984582f002bf (SHA1)`

2. **vtkfilter.tar.xz** C++ code of the measurement algorithm implemented in the form of a VTK polydata filter. It is released under a BSD license compatible with VTK.

   `096081b24182d6ef747d7a75079d196162cacec85d8548b91cd6a4db81b24649 (SHA256)` `07597a780286c2336eb6a41d27af1ee0b1b43503 (SHA1)`

3. **video.mkv** Video of the implemented solution. The `.mkv` files are encoded in VP9 and contain a subtitle track with comments. We recommend visualizing them with VLC media player.

   `d26304f991087a1a32bd63602b3da26857371f70231e54f64641eda60c756b58 (SHA256)` `b65a5960f37d478ca723bff7ecff823ddbcf22a2 (SHA1)`

References

[1] J. Kruth, M. Bartscher, S. Carmignato, R. Schmitt, L. De Chiffre, A. Weckenmann, Computed tomography for dimensional metrology, CIRP Annals 60 (2) (2011) 821–842 (2011). doi:10.1016/j.cirp.2011.05.006.
URL https://www.sciencedirect.com/science/article/pii/S0007850611002083

[2] S. Garcea, Y. Wang, P. Withers, X-ray computed tomography of polymer composites, Composites Science and Technology 156 (2018) 305–319 (2018). doi:10.1016/j.compscitech.2017.10.023.
URL https://www.sciencedirect.com/science/article/pii/S0266353817312460

[3] P. Kulkarni, K. D. Mali, S. Singh, An overview of the formation of fibre waviness and its effect on the mechanical performance of fibre reinforced polymer composites, Composites Part A: Applied Science and Manufacturing 137 (2020) 106013 (2020). doi:10.1016/j.compositesa.2020.106013.
URL https://www.sciencedirect.com/science/article/pii/S1359835X20302529

[4] M. Alves, C. Cimini Junior, S. Ha, Fiber waviness and its effect on the mechanical performance of fiber reinforced polymer composites: An enhanced review, Composites Part A: Applied Science and Manufacturing 149 (2021) 106526 (2021). doi:10.1016/j.compositesa.2021.106526.
URL https://www.sciencedirect.com/science/article/pii/S1359835X21002487

[5] N. C. Davidson, A. R. Clarke, G. Archenhold, Large-area, high-resolution image analysis of composite materials, Journal of Microscopy 185 (2) (1997) 233–242 (1997). arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1365-2818.1997.1560712.x, doi:10.1046/j.1365-2818.1997.1560712.x.
URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2818.1997.1560712.x

[6] P. J. Creveling, W. W. Whitacre, M. W. Czabaj, A fiber-segmentation algorithm for composites imaged using x-ray microtomography: Development and validation, Composites Part A: Applied Science and Manufacturing 126 (2019) 105606 (2019). doi:10.1016/j.compositesa.2019.105606.
URL https://www.sciencedirect.com/science/article/pii/S1359835X19303550

[7] A. Clarke, G. Archenhold, N. Davidson, A novel technique for determining the 3d spatial distribution of glass fibres in polymer composites, Composites Science and Technology 55 (1) (1995) 75–91 (1995). doi:10.1016/0266-3538(95)00087-9.
URL https://www.sciencedirect.com/science/article/pii/0266353895000879

[8] A. Clarke, C. Eberhardt, The representation of reinforcing fibres in composites as 3d space curves, Composites Science and Technology 59 (8) (1999) 1227–1237 (1999). doi:10.1016/S0266-3538(98)00159-6.
URL https://www.sciencedirect.com/science/article/pii/S0266353898001596

[9] 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 Manufacturing 40 (2) (2009) 152–163 (2009). doi:10.1016/j.compositesa.2008.10.014.
URL https://www.sciencedirect.com/science/article/pii/S1359835X08002662

[10] R. Sencu, Z. Yang, Y. Wang, P. Withers, C. Rau, A. Parson, C. Soutis, 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 (2016). doi:10.1016/j.compositesa.2016.09.010.
URL https://www.sciencedirect.com/science/article/pii/S1359835X16303049

[11] M. J. Emerson, K. M. Jespersen, A. B. Dahl, K. Conradsen, L. P. Mikkelsen, Individual fibre segmentation from 3d x-ray computed tomography for characterising the fibre orientation in unidirectional composite materials, Composites Part A: Applied Science and Manufacturing 97 (2017) 83–92 (2017). doi:10.1016/j.compositesa.2016.12.028.
URL https://www.sciencedirect.com/science/article/pii/S1359835X16304560

[12] D. Wilhelmsson, L. Asp, A high resolution method for characterisation of fibre misalignment angles in composites, Composites Science and Technology 165 (2018) 214–221 (2018). doi:10.1016/j.compscitech.2018.07.002.
URL https://www.sciencedirect.com/science/article/pii/S0266353818309400

[13] K. Amjad, W. Christian, K. Dvurecenska, M. Chapman, M. Uchic, C. Przybyla, E. Patterson, Computationally efficient method of tracking fibres in composite materials using digital image correlation, Composites Part A: Applied Science and Manufacturing 129 (2020) 105683 (2020). doi:10.1016/j.compositesa.2019.105683.
URL https://www.sciencedirect.com/science/article/pii/S1359835X19304324

[14] M. Mehdikhani, C. Breite, Y. Swolfs, M. Wevers, S. V. Lomov, L. Gorbatikh, Combining digital image correlation with x-ray computed tomography for characterization of fiber orientation in unidirectional composites, Composites Part A: Applied Science and Manufacturing 142 (2021) 106234 (2021). doi:10.1016/j.compositesa.2020.106234.
URL https://www.sciencedirect.com/science/article/pii/S1359835X2030470X

[15] C. N. Eberhardt, A. R. Clarke, Automated reconstruction of curvilinear fibres from 3d datasets acquired by x-ray microtomography, Journal of Microscopy 206 (1) (2002) 41–53 (2002). arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1365-2818.2002.01009.x, doi:10.1046/j.1365-2818.2002.01009.x.
URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2818.2002.01009.x

[16] A. Julià i Juanola, M. Ruiz i Altisent, I. Boada i Oliveras, An efficient and uniformly behaving streamline-based μct fibre tracking algorithm using volume-wise structure tensor and signal processing techniques, Computer Methods in Applied Mechanics and Engineering 394 (2022) 114898 (2022). doi:10.1016/j.cma.2022.114898.
URL https://www.sciencedirect.com/science/article/pii/S0045782522001359

[17] Y. H. Lee, S. W. Lee, J. R. Youn, K. Chung, T. J. Kang, Characterization of fiber orientation in short fiber reinforced composites with an image processing technique, Materials Research Innovations 6 (2) (2002) 65–72 (2002). arXiv:https://doi.org/10.1007/s10019-002-0180-8, doi:10.1007/s10019-002-0180-8.
URL https://doi.org/10.1007/s10019-002-0180-8

[18] K. Robb, O. Wirjadi, K. Schladitz, Fiber orientation estimation from 3d image data: Practical algorithms, visualization, and interpretation, in: 7th International Conference on Hybrid Intelligent Systems (HIS 2007), 2007, p. 320–325 (Sep. 2007). doi:10.1109/HIS.2007.26.

[19] M. Krause, J. Hausherr, B. Burgeth, C. Herrmann, W. Krenkel, Determination of the fibre orientation in composites using the structure tensor and local x-ray transform, Journal of Materials Science 45 (2010) 888–896 (02 2010). doi:10.1007/s10853-009-4016-4.

[20] I. Straumit, S. V. Lomov, M. Wevers, Quantification of the internal structure and automatic generation of voxel models of textile composites from x-ray computed tomography data, Composites Part A: Applied Science and Manufacturing 69 (2015) 150–158 (2015). doi:10.1016/j.compositesa.2014.11.016.
URL https://www.sciencedirect.com/science/article/pii/S1359835X14003625

[21] P. Pinter, S. Dietrich, B. Bertram, L. Kehrer, P. Elsner, K. Weidenmann, Comparison and error estimation of 3d fibre orientation analysis of computed tomography image data for fibre reinforced composites, NDT & E International 95 (2018) 26–35 (2018). doi:10.1016/j.ndteint.2018.01.001.
URL https://www.sciencedirect.com/science/article/pii/S0963869517303821

[22] R. Karamov, L. M. Martulli, M. Kerschbaum, I. Sergeichev, Y. Swolfs, S. V. Lomov, Micro-ct based structure tensor analysis of fibre orientation in random fibre composites versus high-fidelity fibre identification methods, Composite Structures 235 (2020) 111818 (2020). doi:10.1016/j.compstruct.2019.111818.
URL https://www.sciencedirect.com/science/article/pii/S0263822319322986

[23] C. Creighton, M. Sutcliffe, T. Clyne, A multiple field image analysis procedure for characterisation of fibre alignment in composites, Composites Part A: Applied Science and Manufacturing 32 (2) (2001) 221–229 (2001). doi:10.1016/S1359-835X(00)00115-9.
URL https://www.sciencedirect.com/science/article/pii/S1359835X00001159

[24] R. Blanc, C. Germain, J. D. costa, P. Baylou, M. Cataldi, Fiber orientation measurements in composite materials, Composites Part A: Applied Science and Manufacturing 37 (2) (2006) 197–206, compTest 2004 (2006). doi:10.1016/j.compositesa.2005.04.021.
URL https://www.sciencedirect.com/science/article/pii/S1359835X05002472

[25] K. Kratmann, M. Sutcliffe, L. Lilleheden, R. Pyrz, O. Thomsen, A novel image analysis procedure for measuring fibre misalignment in unidirectional fibre composites, Composites Science and Technology 69 (2) (2009) 228–238 (2009). doi:10.1016/j.compscitech.2008.10.020.
URL https://www.sciencedirect.com/science/article/pii/S0266353808004132

[26] T. Elberfeld, J. De Beenhouwer, A. J. den Dekker, C. Heinzl, J. Sijbers, Parametric reconstruction of glass fiber-reinforced polymer composites from x-ray projection data-a simulation study, Journal of nondestructive evaluation 37 (3) (2018) 62 (2018). doi:10.1007/s10921-018-0514-0.
URL https://europepmc.org/articles/PMC6314276

[27] S. Yurgartis, Measurement of small angle fiber misalignments in continuous fiber composites, Composites Science and Technology 30 (4) (1987) 279–293 (1987). doi:10.1016/0266-3538(87)90016-9.
URL https://www.sciencedirect.com/science/article/pii/0266353887900169

[28] M. Sutcliffe, Modelling the effect of size on compressive strength of fibre composites with random waviness, Composites Science and Technology 88 (2013) 142–150 (2013). doi:10.1016/j.compscitech.2013.09.002.
URL https://www.sciencedirect.com/science/article/pii/S0266353813003552

[29] J. Wang, K. Potter, K. Hazra, M. Wisnom, Experimental fabrication and characterization of out-of-plane fiber waviness in continuous fiber-reinforced composites, Journal of Composite Materials 46 (17) (2012) 2041–2053 (2012). arXiv:https://doi.org/10.1177/0021998311429877, doi:10.1177/0021998311429877.
URL https://doi.org/10.1177/0021998311429877

[30] D. Liu, N. Fleck, M. Sutcliffe, Compressive strength of fibre composites with random fibre waviness, Journal of the Mechanics and Physics of Solids 52 (7) (2004) 1481–1505 (2004). doi:10.1016/j.jmps.2004.01.005.
URL https://www.sciencedirect.com/science/article/pii/S0022509604000110

[31] H. Hsiao, I. Daniel, Effect of fiber waviness on stiffness and strength reduction of unidirectional composites under compressive loading, Composites Science and Technology 56 (5) (1996) 581–593 (1996). doi:10.1016/0266-3538(96)00045-0.
URL https://www.sciencedirect.com/science/article/pii/0266353896000450

[32] L. Bloom, J. Wang, K. Potter, Damage progression and defect sensitivity: An experimental study of representative wrinkles in tension, Composites Part B: Engineering 45 (1) (2013) 449–458 (2013). doi:10.1016/j.compositesb.2012.05.021.
URL https://www.sciencedirect.com/science/article/pii/S1359836812003472

[33] K. Potter, B. Khan, M. Wisnom, T. Bell, J. Stevens, Variability, fibre waviness and misalignment in the determination of the properties of composite materials and structures, Composites Part A: Applied Science and Manufacturing 39 (9) (2008) 1343–1354 (2008). doi:10.1016/j.compositesa.2008.04.016.
URL https://www.sciencedirect.com/science/article/pii/S1359835X08001176

[34] K. Mizukami, Y. Mizutani, K. Kimura, A. Sato, A. Todoroki, Y. Suzuki, Y. Nakamura, Visualization and size estimation of fiber waviness in multidirectional cfrp laminates using eddy current imaging, Composites Part A:

Applied Science and Manufacturing 90 (2016) 261–270 (2016). doi:10.1016/j.compositesa.2016.07.008.
URL https://www.sciencedirect.com/science/article/pii/S1359835X16302287

[35] M. Thor, U. Mandel, M. Nagler, F. Maier, J. Tauchner, M. Sause, R. Hinterhoelzl, Numerical and experimental investigation of out-of-plane fiber waviness on the mechanical properties of composite materials, International Journal of Material Forming 14 (01 2021). doi:10.1007/s12289-020-01540-5.

[36] J.-A. Serret, Sur quelques formules relatives à la théorie des courbes à double courbure., Journal de Mathématiques Pures et Appliquées (1851) 193–207 (1851).
URL http://eudml.org/doc/235002

[37] F. Frenet, Sur les courbes à double courbure., Journal de Mathématiques Pures et Appliquées (1852) 437–447 (1852).
URL http://eudml.org/doc/233946

[38] B. Mandelbrot, How long is the coast of britain? statistical self-similarity and fractional dimension, Science 156 (3775) (1967) 636–638 (1967). arXiv:https://www.science.org/doi/pdf/10.1126/science.156.3775.636, doi:10.1126/science.156.3775.636.
URL https://www.science.org/doi/abs/10.1126/science.156.3775.636

[39] Kitware Inc., The visualization toolkit, [Internet] (2022) [cited 5/5/2022].
URL https://www.vtk.org

[40] M. Ruiz i Altisent, A. Julià i Juanola, I. Boada i Oliveras, Starviewer and its comparison with other open-source dicom viewers using a novel hierarchical evaluation framework, International Journal of Medical Informatics 137 (2020) 104098 (2020). doi:10.1016/j.ijmedinf.2020.104098.
URL https://www.sciencedirect.com/science/article/pii/S1386505619301108

[41] E. T. Whittaker, Xviii.—on the functions which are represented by the expansions of the interpolation-theory, Proceedings of the Royal Society of Edinburgh 35 (1915) 181–194 (1915). doi:10.1017/S0370164600017806.

[42] H. Nyquist, Certain topics in telegraph transmission theory, Transactions of the American Institute of Electrical Engineers 47 (2) (1928) 617–644 (April 1928). doi:10.1109/T-AIEE.1928.5055024.

[43] C. Shannon, Communication in the presence of noise, Proceedings of the IRE 37 (1) (1949) 10–21 (Jan 1949). doi:10.1109/JRPROC.1949.232969.

[44] J. Velut, A spline-driven image slicer, The VTK Journal (09 2011). doi:10.54294/tih9wh.
URL http://hdl.handle.net/10380/3318

[45] W. Wang, B. Jüttler, D. Zheng, Y. Liu, Computation of rotation minimizing frames, ACM Trans. Graph. 27 (1) (mar 2008). doi:10.1145/1330511.1330513.
URL https://doi.org/10.1145/1330511.1330513

[46] Kitware Inc., Paraview, [Internet] (2022) [cited 5/5/2022].
URL https://www.paraview.org

5 Results and discussion

The purpose of this thesis has been the application of common medical imaging techniques for the analysis and visualization of composite materials using non-destructive methods such as μCT. While these datasets are technically similar to medical CT, the morphology of imaged structures is not. Thus, in order to aid an expert engineer or researcher in composite materials, specific methods had to be developed in order to fulfill the measurement and quantification needs. Notwithstanding, several medical visualization features could be identified and repurposed in the field of materials analysis such as: 2D viewers, measurement tools, reference planes, reslicing, projections, ROI (region of interest), volume rendering, tensor maps, among others (see Figure 5.1). In this chapter, in addition to discussing in detail the developed algorithms, we also provide a more interrelated insight of Chapter 2, Chapter 3 and Chapter 4 respectively in the following orderly manner: (i) analysing the evaluation methodology and features breakdown; (ii) the viewer itself, its medical use, the potential repurpose of certain DTI characteristics, among others; and (iii) the novel fibre reconstruction and measurement algorithms we propose.

(a) Comparison at different tensor sizes.

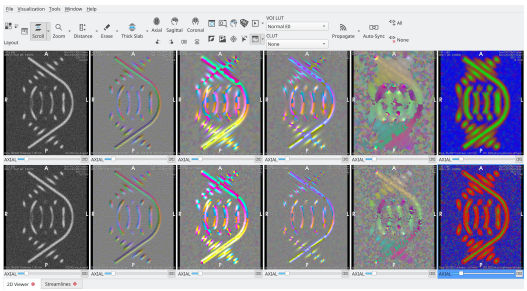

(b) More tensor maps loaded in 2D viewers.

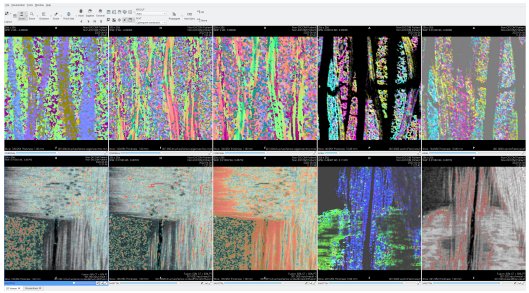

(c) Other intermediate and fused visualizations.

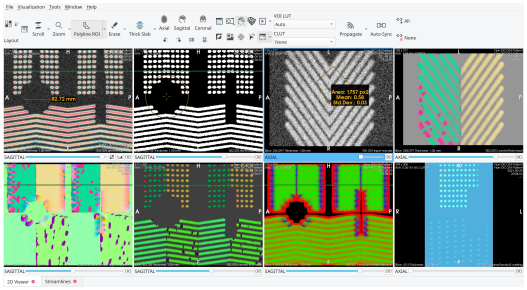

(d) The developed materials extension.

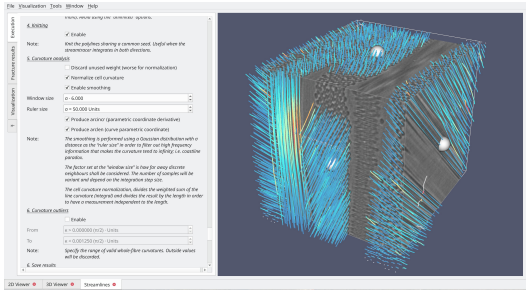

Figure 5.1.: Visualization of volumetric intermediate results (i.e. virtual series) employing the medical 2D viewer (a,b,c), alongside the implemented materials analysis workflow in the Starviewer platform (d).

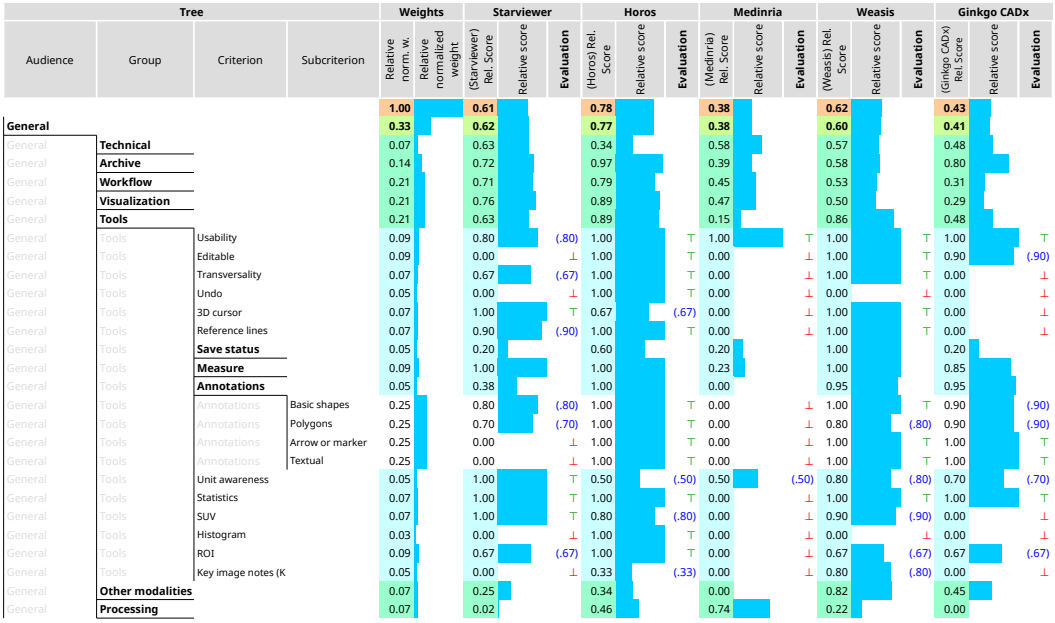

Figure 5.2.: Semi-expanded hierarchy of the evaluation methodology applied in Chapter 2 for the 5 different viewers analysed.

5.1 Evaluation framework

In Chapter 2 we began introducing the most relevant features expected in medical imaging software with special emphasis on the clinical practice (where a more mature usage is observed) in order to posteriorly extend the use case to the non-destructive analysis of composite materials imaged with the μCT technique. Some of these key features had been described in more detail through all the Section 1.2 in order to put in context the relevance and potential of this interdisciplinary use case.

While considering the compendium of evaluation criteria in the state-of-the-art, a new and exhaustive collection has been presented and elaborated. In order to reduce the bias, the proposed criteria have been broken down at three hierarchical levels: group, criterion, and subcriterion (see Figure 5.2). Using these and focusing on open-source software, the following candidates were evaluated: Starviewer, Horos, MedInria, Weasis, and Ginkgo CADx. Although many more were considered, they have been dismissed for reasons such as being unmaintained, or deemed unsuitable in a productive diagnostic workflow. The chosen ones were installed in a testing environment (medical workstation with diagnostic-grade displays) in order to evaluate all criteria. Out of the 220 items, only the innermost nodes (i.e. leafs) were assigned a floating point score ranging from 0 to 1 considering the descriptive definition of each criterion. Most of them could be evaluated in a binary fashion, however some were assigned an intermediate mark when all aspects could not be fulfilled. In addition to this fine granularity as a measure for minimizing the subjectivity and biasing issues, this was also complemented by the self-obligation of providing an explanation/justification when a full mark could not be assigned.

Despite having a concrete evaluation for each combination of item and viewer; we also wanted to produce final and intermediate scores for each viewer and level in the hierarchy. One naive way to perform such task is averaging scores contained within each group. While

simple, this approach was discarded because otherwise the relevance would instead depend on the amount of items and not on the actual importance for the final-user. For this reason, each item in the hierarchy had a weight assigned within its group in order to express its relevance. Although this assignment was also performed in a bottom-up approach, this also constitutes a potential source for bias when the real needs of the end-users are disregarded. Putting aside the interdisciplinary composite material analysis, the spotlight was set on evaluating the importance of each criterion focusing on the daily diagnostic needs in the medical sector. In this mature environment, and through talks with expert radiologists, we could establish two clearly distinguished user groups: (i) the ones requiring the manipulation of 3D datasets (MRI, CT, PET, etc.); and (ii) the ones exclusively focused on the treatment of 2D images (radiography, ultrasound, mammography, etc.). In between these two, there is also a third group of users performing a more balanced use of both groups features due to their medical speciality. For this reason the assignment of weights was split in three different groups (3D based, 2D based, and General) in order to reflect the real needs of each audience.

5.2 The Starviewer platform

By observing the comparison results in Fig. 6 of Chapter 2, it can be concluded that Starviewer is capable of fulfilling most daily clinical practice needs. In addition to demonstrating its capabilities, Chapter 2 also introduces the platform itself, its modules, structure, and the possibilities for developing specific extensions (see Fig. 1 of Chapter 2). Apart from its core/production-ready features, throughout its history several extensions have been developed to address specific medical workflows. In that regard, and while being a different field, the needs of an expert engineer analysing composite materials acquired through the μCT technique match almost all core-features identified in the 3D-based audience.

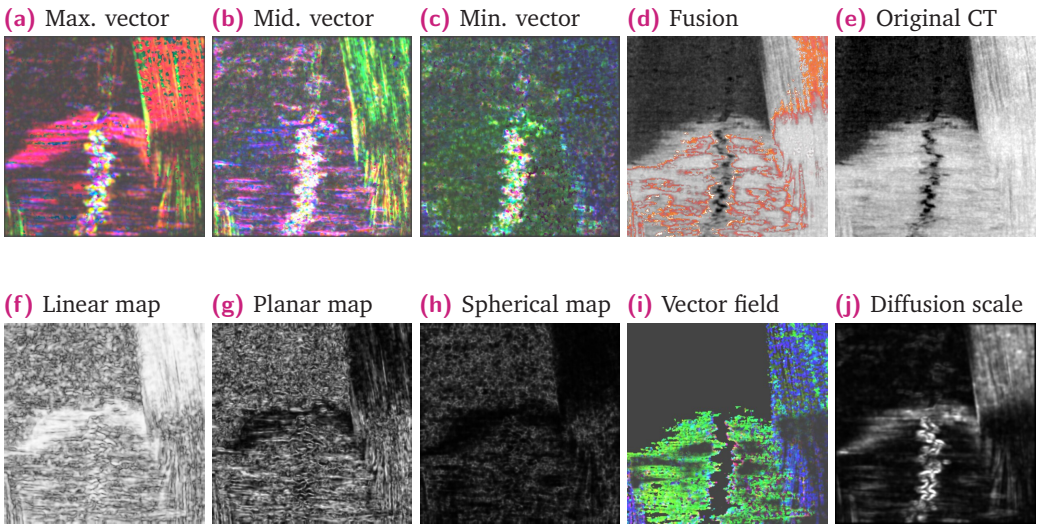

Figure 5.3.: Examples of the intermediate volumes produced on-the-fly which are loadable as regular DICOM series in the 2D viewer in order to gain insight of the underlying tensor field. These images are generated from a fibre glass CT where individual fibres can not identified (more similar to the DTI scenario).

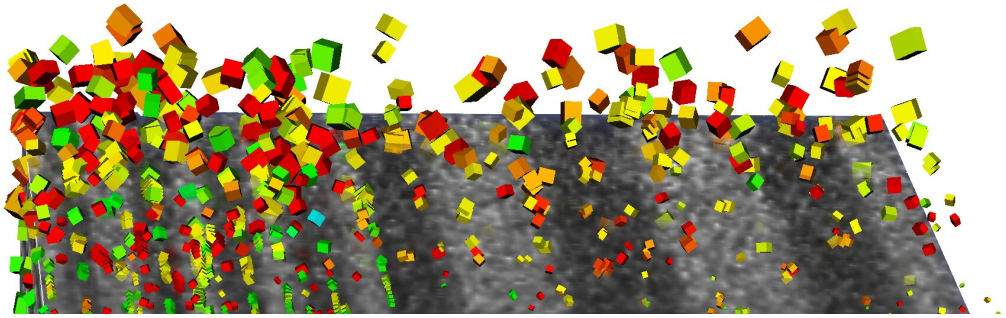

Figure 5.4.: Oriented glyphs corresponding to the underlying structure tensor field arising from the image gradient of a synthetic dataset where noise is progressively increased along the Z-axis direction.

Unlike the simpler extension implemented as an example in the Annex A of Chapter 2, the explainability objective requires exposing to the user the intermediate results of the specialized processing pipeline (see Figure 5.5 and Section 5.3). This has resulted in the development of a specific material analysis extension (see Figure 5.1d) which produces extra/virtual volumetric datasets loadable as additional DICOM series in order to seamlessly repurpose some of the basic exploration features fulfilling the medical diagnostic needs. One example is the production of several scalar and vectorial maps understandably exposing the behaviour of the underlying structure tensor volume.

Among others, as illustrated by Figure 5.1a, Figure 5.1b, Figure 5.1c, and Figure 5.3, one can interactively visualize several intermediate volumes such as: (i) the tensor *diffusion scale* (Figure 5.3j), or its morphology in terms of *linear anisotropy* (Figure 5.3f), *planar anisotropy* (Figure 5.3g) or *isotropy* (Figure 5.3h) [66]; (ii) the eigenvector fields scaled according to its corresponding eigenvalues (Figure 5.3a, 5.3b, 5.3c); and (iii) the fibre orientation vector field (Figure 5.3i), scaled only for the zones where fibre has been detected (Figure 5.3d). All these series can also be interactively overlaid on top of the original μCT acquisition (Figure 5.3e) by repurposing the fusion feature commonly used for PET acquisitions (Figure 5.3d, and Fig. 3b of Chapter 2). Also, tensors can be visualized as oriented glyphs as shown in Figure 5.4 where the effects of a z-increasing noise floor alters the image gradient magnitude and hence the scale of the tensor (see Figure 5.3j). Note how all these intermediate visualizations for composite materials are based on well established techniques in the medical practice for the analysis of DTI datasets [66].

Should the user interface enabling the volume exploration capabilities be implemented from scratch instead of relying on the features provisioned by Starviewer, the development time and overall quality of the whole would have diminished. While most of the method implementation code interfaces directly with VTK, the Starviewer platform also offloads the need to reimplement the input/output features related to the loading of datasets. Starviewer, besides reading from a local filesystem, can also perform a network retrieval via external PACS server. Although positive overall, the development experience has highlighted the need for Starviewer to better support mixed visualizations of non-volumetric data types such as polydata, point clouds, etc. in the 2D-based viewers. As future work, some of the specific visualizations implemented exclusively for the analysis of materials will be ported in a more generic form to the platform itself in order to further ease the development of new extensions.

5.3 Analysis of composite materials

Although Starviewer covered the most important visualization topics, some advanced aspects such as the orientation or curvature measurements required the identification of individual fibre paths. This has lead to the development of a specific extension to cover the topics in Chapter 3 and Chapter 4. Together, and as illustrated by Figure 5.5 they conform an image processing pipeline which can be subdivided in the following phases: (i) *segmentation*, where the relevant regions with fibre are selectively detected (Figure 5.5a); (ii) *directionality*, where each region in space is assigned a vector representing the fibre orientation (Figure 5.5b); (iii) *seeding*, where the fibre centres normally crossing a certain user-defined surface are detected (Figure 5.5c); (iv) *fibre tracking*, which starts from the obtained seeds and integrates over the fibre directionality vector field (Figure 5.5d); and (v) *measurement* of the curvature and waviness above certain user-defined scale using the reconstructed trackings (Figure 5.5e). In this section we also discuss the explainability, performance, and testing of the methods performed with both real and synthetic datasets. Some of these synthetic datasets had to be generated using the novel methodology proposed in Section 2.3.1 of Chapter 4 which can uniformly produce the complex morphological shapes required for ensuring a proper validation of the pipeline and its measurements.

5.3.1 Fibre reconstruction

While being more faithful to the reality, volumetric datasets are computationally challenging due to their large size. Realising that our purpose is the analysis of fibre paths and their directionality, summarizing and concentrating the most relevant information for that purpose is beneficial. For this reason, representing each one as a polygonal fibre path aids the development of more advanced measurement methods. Unlike image/volume processing algorithms operating around a neighbouring region or kernel, polygonal representations allow for a separate and isolated processing of each entity (i.e. fibre). Hence, are not affected by the inability to discern voxels belonging to different fibres within a kernel.

The proposed fibre reconstruction algorithm, unlike two-dimensional or slice-dependent approaches [10, 14, 2, 11, 59], considers the input as a whole and consequently, all dimensions and orientations are treated uniformly. In turn, when performing 3D-based operations (e.g. convolutions, projections, etc.), the amount of noise is reduced as the information correlation between neighbouring slices emerges instead of being obliviated as the aforementioned slice-by-slice approaches inevitably do. At each step in the pipeline, the data is observed as a valuable signal to preserve and process using analogue-like filtering principles implemented over digital means. Images are not rescaled, transfer functions are preferred over binary segmentations, filters respond smoothly, premature discretization is avoided, etc. Notable examples of this practice are illustrated in Figure 5.5a for: (i) the optional low-pass, band-pass or high-pass prefiltering; (ii) the piecewise fibre transfer function delimited by four user-adjustable regions; and (iii) the iterative Difference of Gaussians filter (IDoG) which selectively amplifies the features (i.e. fibres) at a certain scale using a band-pass filter with two σ_1 and σ_2 parameters. As with any filter, the output energy (i.e. voxel intensity) is expected to decay the narrower its selectivity is.

Their application leads to a volume with many blurry and elongated blobs where each one represents a fibre with an intensity decay from its centre to its perimeter. While having a

(a) Detection of the voxels with fibre (segmentation) matching a certain user-defined target scale.

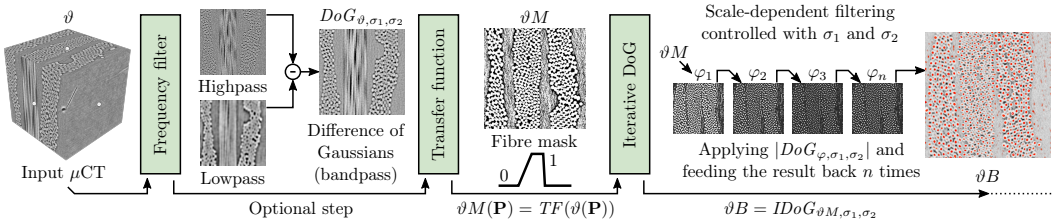

(b) Determination of the fibre directionality using the structure tensor over the volume image gradient.

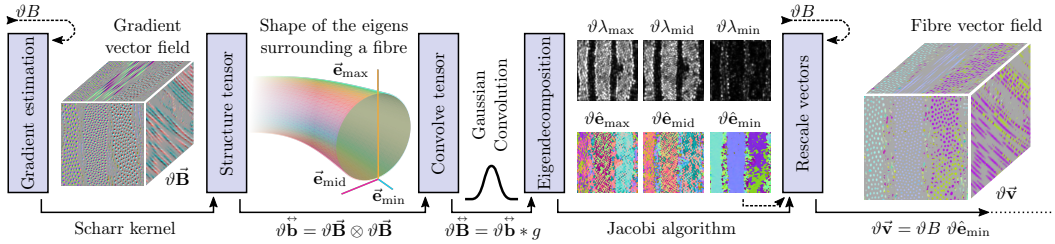

(c) Finding the individual fibre centres piercing through a set of user-defined planes.

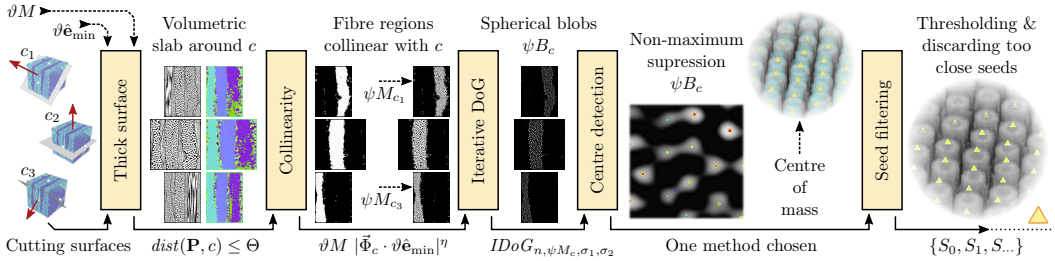

(d) Tracing streamlines over the fibre directionality vector field from the previously detected centres.

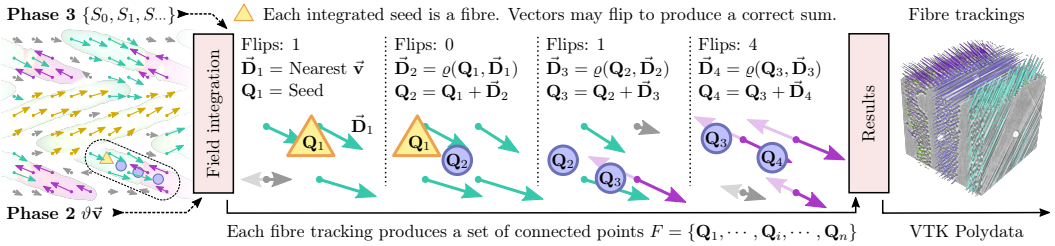

(e) Measuring the curvature of the polygonal reconstructions only above a certain user-defined scale.

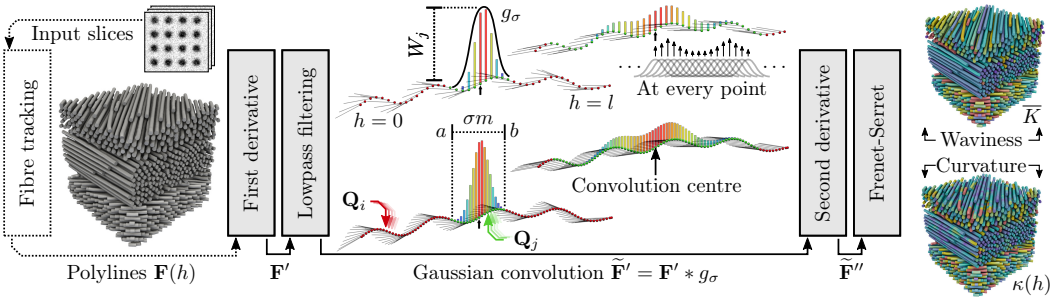

Figure 5.5.: Overview of the whole image processing pipeline starting from the acquired slices (i.e. volume) to the final measurements as defined by Chapter 3 and Chapter 4 respectively.

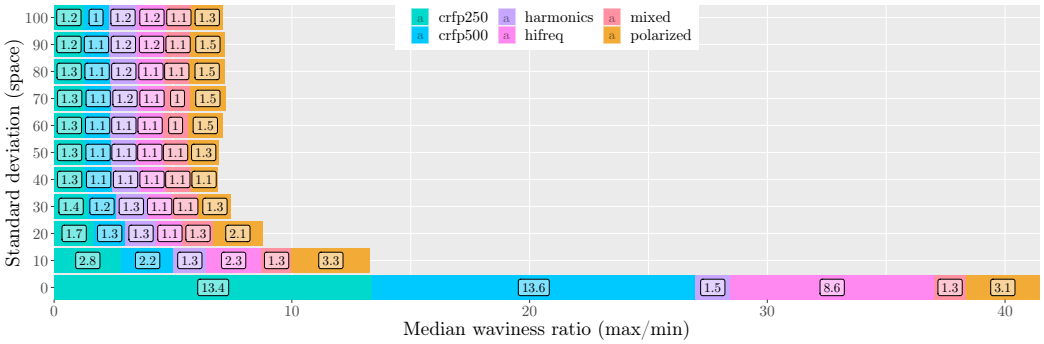

Figure 5.6.: Ratio between the maximum and minimum median values at different σ sizes. Once information is properly sampled, ratios approaching 1 indicate convergence independently of the step size. Refer to the legend in Figure 8 of Chapter 4 for more information.

proper segmentation we are required to obtain its directionality, hence the computation of the image gradient is required. Unfortunately this new volumetric field contains vectors pointing from the perimeter towards a fibre medial axis. In a 2D scenario/approach, the orthogonal direction would correspond to the actual fibre directionality, however in the inherently 3D method we pursue, the vector field collinear with the fibres must be obtained by: (i) computing the structure tensor of the gradient; (ii) eigendecomposing it; and (iii) picking the minimum eigenvector field for the directionality measurement. Note that its quality and precision will be subject to numeric gradient estimator method. Common techniques such as the Sobel kernel have leaded to unsuitable results; the tracked fibres deviated from the desired path at certain orientations. For this reason, it is of utmost importance the employment of precise estimators such as the Scharr [56] kernel which has been especially crafted to have an uniform behaviour and thus achieve rotational invariance in real-world conditions.

The implemented pipeline exposes an extremely long list of user-adjustable parameters; withal the developed extension is capable of automatically setting most of them by precomputing a fail-safe ratio respect a more general one (e.g. fibre diameter). While these defaults can be overridden at any time, in practice, the most essential ones to specify are: (i) the transfer function intercepts, (ii) the target diameter, (iii) the amount of IDoG iterations, and (iv) the placement of seeding surfaces. In this way, a user familiarized with the underlying fundamentals (e.g. signal processing, filtering, segmentation) can properly fine tune the relevant ones for challenging datasets with varying factors, imaging techniques or different acquisitions circumstances (e.g. low contrast, noise, artefacts).

Note that while most of the cited methods aim to reconstruct all the fibres in a volume, we opted for enabling the user to strategically restrict the tracking by adjusting three main aspects: (i) the *cutting surface* where the seeds (i.e. fibre centres) are automatically placed at the intersection region; (ii) the amount of *collinearity* or tolerance respect the surface normal and the fibre vector field directionality; and (iii) the minimum field magnitude to sustain through all the fibre tracking before stagnating. These, in addition to other indirect parameters governing the image processing filters (e.g. fibre diameter, iterations), can be strategically combined to selectively reconstruct only certain bundles of interest as illustrated in Fig. 10 of Chapter 3.

As it can be observed, the whole process relies heavily on image processing techniques and defers the polygonal conversion to the very last step. The tracking inevitably simplifies and thus trims part of the original dataset information; nonetheless it is very suited for the pursued goal. Its integration principles employed in DTI tractograpies (hyperstreamlines), or more commonly in the visualization and analysis of flow simulations as illustrated in Figure 5.7. However, in our specific case, the fibre orientation field produced in the second phase (see Figure 5.5b) is in fact a pseudo-vector field. While vectors are collinear with fibres, at some regions the vector sign (i.e. the tip) is randomly flipped. As there is no global way to tackle this ambiguity, the Runge-Kutta implementation has been altered to interpolate the underlying field considering the previous result in the path in order to flip the signs if necessary and thus produce a meaningful vector sum. Without this, the sum at these opposing regions becomes zero and the streamline is prematurely stagnated or experiences an abrupt change in direction.

5.3.2 Curvature and waviness measurements

In Section 1.2.6 we distinguished between two measurement approaches: (i) the volume-wise ones which employ image processing techniques to obtain a direct measurement; and (ii) the polygonal-based ones which irrespective of its reconstruction means, measure a curve function (i.e. polygonal path). While the method proposed in Chapter 4 is of the latter type, as described in Section 5.2 the fibre reconstruction process itself leads to several intermediate visualizations emerging from its image processing pipeline (see Figure 5.3). The vector field produced by the *phase 2* (see Figure 5.5b) which is later employed in the tracking phase (see Figure 5.5d) can be repurposed to perform voxel-wise measurements. For this reason, whether individual fibres are discernible or not, this field can be analysed by transforming the 3D vectors into the two angles required to produce a 2D orientation histogram. Also, they can be directly integrated for the whole volume in order to obtain statistics. It is worth noting that in addition this raw vectorial field emerges from the fibre segmentation and filtering process. Hence, only the relevant features (i.e. fibres) matching a

(a) Surface warped respect the velocity.

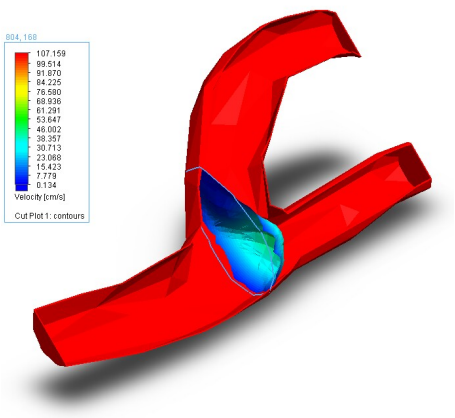

(b) Streamlines at the carotid bifurcation.

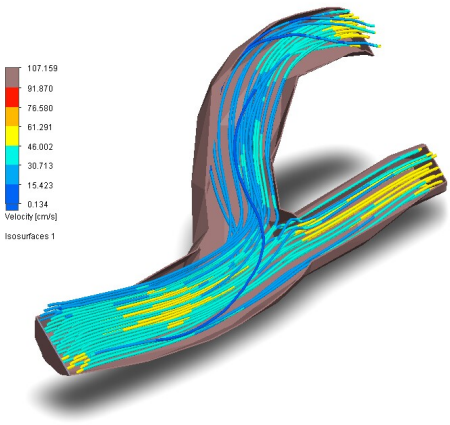

Figure 5.7.: Examples of blood flows simulations performed from segmented vessels acquired with the angiographic techniques as shown in Figure 1.5.

certain scale are considered due to the band-pass capabilities of the IDoG filter as illustrated in Fig. 5 of Chapter 3. In turn, the valuable maps produced by the tensor eigendecomposition as shown in Figure 5.3 could be repurposed and combined to filter out certain artefacts or feed-back to enhance the segmentation selectivity.

Putting aside these future work proposals, we proceed to discuss the more advanced measurement method in Chapter 4 and its advantages respect pure voxel-wise approaches. First of all, the proposed measure does not depend on the reconstruction in Chapter 3; both papers are independent from each other. Its required input is a set of curve functions represented by a sorted list of $\mathbb{R}^3$ points (i.e. discrete samples) consecutively connected by rectilinear segments such that the paths of imaged fibres are faithfully reproduced. Unlike more restrictive definitions, the position of samples is independent from the underlying voxels and in turn these can also be irregularly sampled (i.e. spaced). These flexible input requirements ensure that any other reconstruction methodology can be seamlessly used in place of the one proposed in Chapter 3.

Focusing on the measurement definition itself, it consists in the application of the Frenet-Serret curvature formulas [60, 16] in order to: (i) obtain a punctual *curvature* at any position in the curve (i.e. point data); and (ii) produce a summarizing *waviness* scalar expressing the amount of bending irrespective of its length (cell data). The proposed measurement definitions are inherently 3D behaving, numerically stable (in their intended purpose), are isolated from their neighbouring counterparts, and can properly quantify spiralling fibres beyond a single revolution unlike pure image-processing measurements or other polygonal-based methods which measure from a global world-coordinates perspective.

Unlike the volumetric image signal in the image processing operations, these 3D curves are somewhat similar to a 1D waveform (e.g. audio or radio signal). As the discussed signal processing fundamentals still apply, and a curve will never be infinitely sampled,

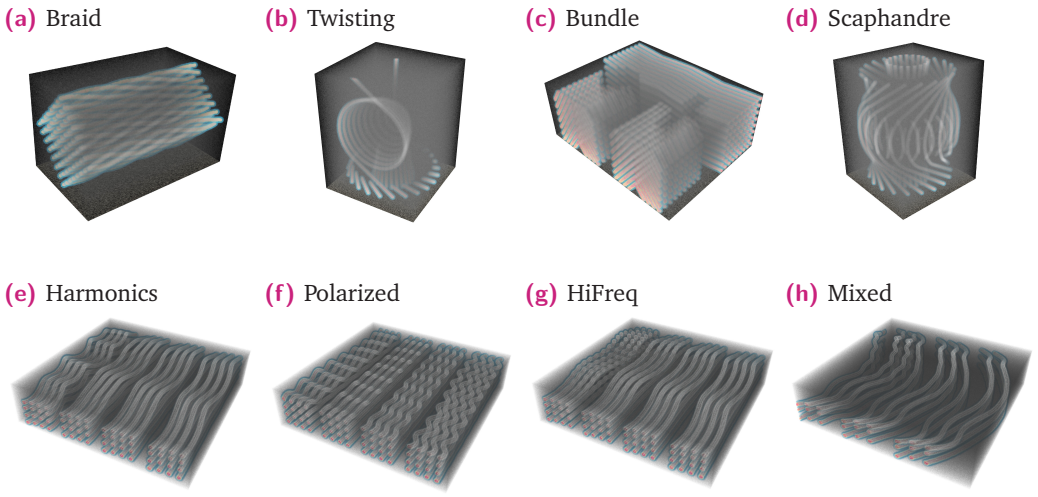

Figure 5.8.: Volume renders of the synthetic datasets generated using the proposed Gaussian-splatter method from a diverse set of intertwined ground truth curves. The actual slices and the generative code is made available in the supporting material of Chapter 3 (top row), and Chapter 4 (bottom row).

the signal (i.e. the curve) will be band-limited due to the Whittaker–Nyquist–Shannon sampling theorem [68, 47, 61]. The high frequency components will not describe the actual shape of physical fibres; their amount varies due to acquisition parameters such as voxel size, scale of the fibres, resolution, dynamic range, noise characteristics of the imaging device, etc. Depending on the sampling frequency, the amount of noise will vary and so will the measured curvature. This undesired phenomenon is problematic as the spectral energy induced by these random high-frequencies produces a fractal-behaving signal which must be limited should any proper measurements be made [39]. The key proposal of our measurement method is the addition of a decoupling mechanism to mitigate these effects. Using a low-pass Gaussian frequency filter the features smaller than a certain user-specified σ wavelength referenced in actual space units can be attenuated. Also, the convolution process considers the irregularity of samples and its reach is also referenced in space units. In addition, it can be repurposed to not only discard noise but also dismiss real variations in the fibres not deemed relevant for a particular study. As illustrated by Figure 5.6 the method converges irrespective of the underlying step size.

So far all state-of-the-art measurements we analysed do not perform such considerations. Some may inadvertently limit these effects by the employment of smoothing polynomials or by indirectly filtering the noise with coarser samplings. For this reason, our proposal enables a proper and comparable measurement across different acquisitions.

5.3.3 Testing

The proposed methods have been tested with real and synthetic datasets. To evaluate the error of the fibre reconstruction proposal in Chapter 3 we employed the synthetic models created by Creveling et al. [10] in order to compare with his error figures in the most transparent and unbiased way. Our method results show that an error below the pixel-size is attained as shown in Figure 5.9b. While we cannot attribute an exact reason for some of the

(a) Manual verification of a real dataset.

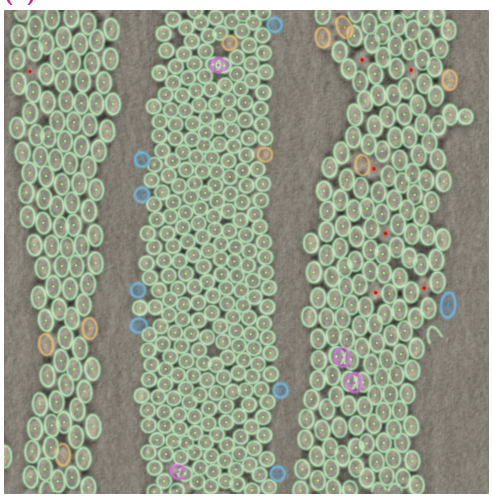

(b) Error respect a synthetic ground truth.

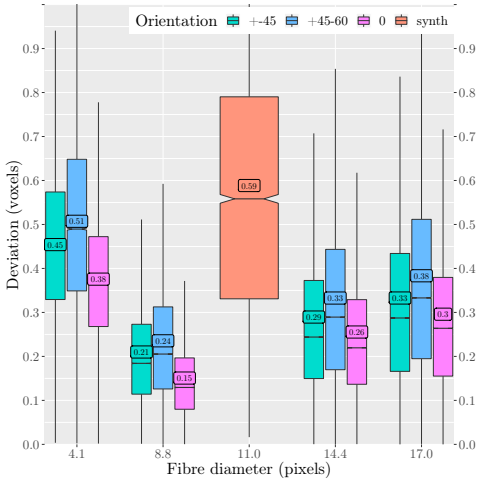

Figure 5.9.: Error assessment of Chapter 3 corresponding to Fig. 11 (left), and Fig. 12a (right).

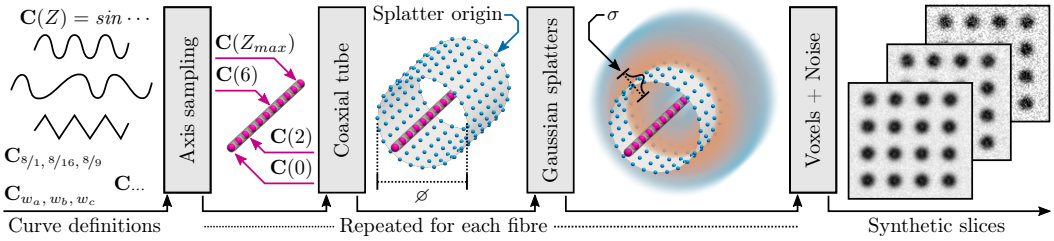

Figure 5.10.: Proposed pipeline for the production of synthetic datasets from a known ground truth. Refer to the legend in Figure 4 of Chapter 4 for more information.

fluctuations, at such high levels of accuracy, external factors such as the imaging technique, quality and uncertainty in the measurement, among others, will have to be considered.

Complementing the aforementioned third-party datasets, in Chapter 4 a novel approach to produce synthetic cases is presented. The proposed method, illustrated in Figure 5.10, generates a tube around the medial axis from which several Gaussian blurs can be splattered. The combined field is then discretized into a finite resolution volume. Due to its 3D and uniform behaviour, the complex curves required to characterize the proposed methods could be produced. In the case of Chapter 3, the purpose of *twisting*, *scaphandre*, *bundle*, and *braid* datasets (see Figure 5.8, top row) was to assess the uniform behaviour of the fibre tracking method and evaluate its abilities to reconstruct any kind of fibre paths as illustrated in Fig. 10 of Chapter 3. This proves that our method can overcome some of the limitations posed by several state-of-the-art algorithms. In the case of Chapter 4, the *harmonics*, *hifreq*, *mixed*, and *polarized* datasets were used (see Figure 5.8, bottom row). Their corresponding ground truths were mathematically described as wave equations with different combinations of wavelengths and amplitudes. These shapes, combined with random pixel-noise were crucial in order to validate the proper behaviour and filtering capabilities of the proposal as illustrated in Figure 6 of Chapter 4.

Complementing the synthetic cases, real acquisitions with individually identifiable fibres were also used. As they lack a reference ground truth, in the case of Chapter 3, a manual verification process was performed as shown in Figure 5.9a. These datasets were also used in Chapter 4, where one can visually check that the slightly curving fibres were properly identified as illustrated in Figure 9 (Chapter 4). Complementing these, several acquisitions with regular CT (see Figure 1.11a, 1.11b, 1.11c) were also used in this thesis to demonstrate the capabilities of intermediate results and visualizations when individual fibres are not discernible (see Figure 5.3 and Figure 5.1b).

5.3.4 Performance and implementation

The proposed approach has been developed in C++ and integrated in the Starviewer platform in the form of a specialized materials analysis extension. At the same time, the implementation is supported by several open-source libraries, the main one being VTK [28].

On the one hand, the specific fibre reconstruction in Chapter 3 was assessed by timing the execution and then contrasting whether the results matched the theoretical expectations. In that regard, the first aspect to consider is that the cost of convolutions is characterized

by the amount of neighbours visited for each input voxel, hence its growth is proportional to the target fibre diameter and not the overall volume size. This expectation has been corroborated by the experiments in Chapter 3.

On the other hand, the stream tracing workloads can only be atomized up to a seed level as each tracking must be performed sequentially from its initial conditions. This paradigm also applies for the measurement algorithm in Chapter 4 although some optimizations could be implemented by parallelizing the convolution for each sample. In any case, as the amount of fibres will usually exceed the core-count, future work will prioritize tackling the parallelism at a cell level (i.e. for each fibre path). While in this type of data structures the size of convolution is not known beforehand, the measurement algorithm cost can be deemed linear respect the overall number of samples, and the cost growth is attributable to the amount of convolving operations required, which will depend on the amount of discrete neighbours reachable in a σm arc distance.

In any case, the overall performance of the whole can be considered good and efficient. Most of the proposed pipeline is partitionable in terms of memory and processing; and the computationally image-processing operations have been successfully parallelized. The performance figures were assessed in a GNU/Linux workstation with a ThreadRipper 1950X whose memory is arranged in two separate NUMA nodes (non-uniform memory access). Under the current testing environment, if the $16\,\mathrm{s}$ processing time of the $350 \times 350 \times 350$ *cfrp_i17* dataset is extrapolated to an extent of $2500 \times 2500 \times 1500$, provided enough memory, a very reasonable processing time of $1\,\mathrm{h}$ is projected.

5.3.5 Reproducibility

All findings and experiments stated in the articles are accompanied with an extensive set of supporting material files with the results, visualizations, code, extra figures and videos, among others. Although not all code for the Starviewer extension has been made publicly available, the algebraic descriptions in Chapter 2 and Chapter 3 are a complete description of the actual implementation.

6 Conclusions

The purpose of this thesis has been the research and development of methods, algorithms and tools for the analysis of composite materials. These require specific techniques in order to gain insight of their constituent materials and their behaviour. For this reason, the non-destructive μCT technique is an established way to perform such detailed analysis. Due to its technical similarity with medical CT, in Chapter 2 we analysed the established techniques in this mature sector, and then we proposed specialized techniques for the reconstruction of fibres in Chapter 3 and the analysis of its curvature in Chapter 4.

In that regard, DICOM viewers are essential in many clinical processes, either to diagnose, plan operations, or follow up the evolution of pathologies. Due to the uprising reliance of digital-based acquisitions, the key role in modern medical imaging workflows is crucial in the daily medical practice. For this reason, and focusing on the open-source DICOM viewers, we presented an adaptable evaluation framework that considers different audiences and criteria weighted according to user needs. In addition we introduce Starviewer, an open-source multi-platform DICOM viewer that supports the main needs of medical image diagnosis and also serves as a platform to develop custom extensions to cover specific workflows, including interdisciplinary ones such as the analysis of materials. However, in order to tackle the potential bias of self-evaluating Starviewer respect others, several measures have been taken in order to achieve an unambiguous and fact-based evaluation. From this exercise, it has been observed that the needs of a 3D-based audience fits many of the needs for the analysis of composite materials. For this reason, the specialised measurement and reconstruction techniques had to be implemented and integrated into a specific extension containing the whole workflow. Also, in its implementation, we considered established medical modalities such as DTI in order to repurpose some of its principles for the analysis and reconstruction of fibre orientations and the handling of tensorial volumes.

Hence, in order to process the μCT acquisitions with individually identifiable fibres, in Chapter 3 we successfully implemented a performant fibre reconstruction algorithm which fulfils the explainability objectives by: (i) applying known signal processing techniques based on physical principles; (ii) constraining the seed placement at peak intensities controllable by user-defined surfaces; and (iii) employing fluid dynamics principles (streamlines) to trace the actual fibre paths. These, in turn, are complemented by a very extensive set of user-adjustable parameters controlling all aspects of the pipeline. In addition, as the execution evolves the user can seamlessly interact with its intermediate results as if they were regular medical imaging series. All these aspects make this whole process supervisable.

Also, some of the intermediate results themselves can be repurposed to gain insight in the μCT acquisitions where individual fibres are not identifiable. By studying its structure tensor maps and/or its orientations at a voxel-level one can easily extract valuable information at a much lower acquisition cost provided the use-case permits it. In any case the proposed four-step reconstruction algorithm considers the input volume as a whole where streamlines and three-dimensional image processing techniques (mainly Gaussian convolutions) are

applied in order to achieve a uniform behaviour in all three axes. Considering the quality and precision of the reconstruction, we consider all those aspects beneficial compared to more opaque approaches requiring training datasets, dictionaries, slice-by-slice methods, etc. Performance-wise, the more voxels, the bigger the processing time, however its growth is linked to the size of the convolutions, which is in turn governed by the target fibre diameter and the kernel precision factor. For this reason, tackling large volumes is attainable even in modest hardware provided enough memory and time.

Building upon this polygonal reconstruction technique, in Chapter 4 we propose a waviness and curvature measurement. Several state-of-the-art articles have found a decrease in the mechanical properties when fibres are misaligned. While some propose specific techniques to quantify these defects from μCT datasets, they present some fundamental limitations regarding their reconstruction and numeric quantification. For this reason we proposed a compatible measurement employing the inherently 3D and uniformly behaving curvature definition of the Frenet-Serret formulas combined with a band-pass filter in order to limit the over-estimating noise effects that attenuate the curvatures below a certain user-defined wavelength (i.e. size) in space units. These finer variations will always be present and they infer a fractal-like shape to the imaged fibres; considering this issue is crucial in order to perform meaningful and coherent measurements across acquisitions and specimens. In our case, this scale selectivity is attained using a Gaussian low-pass frequency filter with an user-adjustable sigma parameter corresponding to a cut-off wavelength defined in space units. With it, our method attains convergence of results for fibres sampled at different resolutions, varying amounts of noise, etc., when a sufficiently large sigma wavelength above the sampling band-limit is employed. This cut-off frequency can also be used to discard actual curvatures present in the physical fibres when these are not relevant in a certain application.

6.1 Future work

Both in the fibre reconstruction algorithm and the curvature measurement method, it has been observed that the scale-space principles [36, 37] can be applied. For the curvature measurement, the cut-off wavelength could be probed at different sizes and then visualized as an spectrum. This has been already indirectly performed as the sigma was gradually increased through all the testing in Figure 6 of Chapter 4. Regarding the fibre reconstruction algorithm, this multi-scale analysis could be implemented as a fourth-dimension for the IDoG volume in order to track fibres with varying diameters. For instance, such line of work could be repurposed for medical purposes as a way to reconstruct blood vessels from angiographic images.

In addition, some of the implemented features will be ported as a more generalized abstraction to the Starviewer platform itself in order to improve some of its drawbacks such as the handling and fusion of polydata datasets in 2D viewers, enable more advanced visualizations for the DTI acquisitions and interactive tractographies, among others. Finally, regarding the specific materials extension, the workflow could be further automated to ease the adjustment of parameters for a novice user.

Bibliography

[1] M.P. Alves, C.A. Cimini Junior, and S.K. Ha. „Fiber waviness and its effect on the mechanical performance of fiber reinforced polymer composites: An enhanced review“. In: *Composites Part A: Applied Science and Manufacturing* 149 (2021), p. 106526. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2021.106526 (cit. on p. 8).

[2] K. Amjad, W.J.R. Christian, K. Dvurecenska, et al. „Computationally efficient method of tracking fibres in composite materials using digital image correlation“. In: *Composites Part A: Applied Science and Manufacturing* 129 (2020), p. 105683. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2019.105683 (cit. on pp. 8, 73).

[3] Peter J. Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. „In vivo fiber tractography using DT-MRI data“. In: *Magnetic Resonance in Medicine* 44.4 (2000), pp. 625–632. DOI: 10.1002/1522-2594(200010)44:4<625::AID-MRM17>3.0.CO;2-O (cit. on p. 9).

[4] R. Blanc, Ch. Germain, J.P. Da costa, P. Baylou, and M. Cataldi. „Fiber orientation measurements in composite materials“. In: *Composites Part A: Applied Science and Manufacturing* 37.2 (2006). CompTest 2004, pp. 197–206. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2005.04.021 (cit. on p. 8).

[5] L.D. Bloom, J. Wang, and K.D. Potter. „Damage progression and defect sensitivity: An experimental study of representative wrinkles in tension“. In: *Composites Part B: Engineering* 45.1 (2013), pp. 449–458. ISSN: 1359-8368. DOI: 10.1016/j.compositesb.2012.05.021 (cit. on p. 8).

[6] Joao Carvalho, Alvaro Rocha, and Antonio Abreu. „Maturity Models of Healthcare Information Systems and Technologies: a Literature Review“. In: *Journal of Medical Systems* 40 (Apr. 2016), p. 10. DOI: 10.1007/s10916-016-0486-5 (cit. on p. 5).

[7] Nadire Cavus. „The Evaluation of Learning Management Systems Using an Artificial Intelligence Fuzzy Logic Algorithm“. In: *Adv. Eng. Softw.* 41.2 (Feb. 2010), pp. 248–254. ISSN: 0965-9978. DOI: 10.1016/j.advengsoft.2009.07.009 (cit. on p. 5).

[8] A.R. Clarke, G. Archenhold, and N.C. Davidson. „A novel technique for determining the 3D spatial distribution of glass fibres in polymer composites“. In: *Composites Science and Technology* 55.1 (1995), pp. 75–91. ISSN: 0266-3538. DOI: 10.1016/0266-3538(95)00087-9 (cit. on p. 8).

[9] Ashley Clarke and Colin Eberhardt. „The representation of reinforcing fibres in composites as 3D space curves“. In: *Composites Science and Technology* 59.8 (1999), pp. 1227–1237. ISSN: 0266-3538. DOI: 10.1016/S0266-3538(98)00159-6 (cit. on p. 8).

[10] Peter J. Creveling, William W. Whitacre, and Michael W. Czabaj. „A fiber-segmentation algorithm for composites imaged using X-ray microtomography: Development and validation“. In: *Composites Part A: Applied Science and Manufacturing* 126 (2019), p. 105606. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2019.105606 (cit. on pp. 8, 73, 78).

[11] Vedrana A. Dahl, Monica J. Emerson, Camilla H. Trinderup, and Anders B. Dahl. „Content-based Propagation of User Markings for Interactive Segmentation of Patterned Images“. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2020, pp. 4280–4288. DOI: 10.1109/CVPRW50498.2020.00505 (cit. on p. 73).

[12] C. N. Eberhardt and A. R. Clarke. „Automated reconstruction of curvilinear fibres from 3D datasets acquired by X-ray microtomography“. In: *Journal of Microscopy* 206.1 (2002), pp. 41–53. DOI: 10.1046/j.1365-2818.2002.01009.x (cit. on p. 8).

[13] Tim Elberfeld, Jan De Beenhouwer, Arnold J den Dekker, Christoph Heinzl, and Jan Sijbers. „Parametric Reconstruction of Glass Fiber-reinforced Polymer Composites from X-ray Projection Data-A Simulation Study“. In: *Journal of nondestructive evaluation* 37.3 (2018), p. 62. ISSN: 0195-9298. DOI: 10.1007/s10921-018-0514-0 (cit. on p. 8).

[14] Monica J. Emerson, Kristine M. Jespersen, Anders B. Dahl, Knut Conradsen, and Lars P. Mikkelsen. „Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials“. In: *Composites Part A: Applied Science and Manufacturing* 97 (2017), pp. 83–92. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2016.12.028 (cit. on pp. 8, 73).

[15] R.P. Feynman, R.B. Leighton, and M. Sands. *The Feynman Lectures on Physics, Vol. I: The New Millennium Edition: Mainly Mechanics, Radiation, and Heat*. Vol. 1. The Feynman Lectures on Physics. Basic Books, 2011. Chap. 36. ISBN: 9780465024933. URL: https://www.feynmanlectures.caltech.edu/I_36.html (cit. on p. 1).

[16] F. Frenet. „Sur les courbes à double courbure.“ fre. In: *Journal de Mathématiques Pures et Appliquées* (1852), pp. 437–447. ISSN: 0021-7874. URL: http://eudml.org/doc/233946 (cit. on p. 77).

[17] Bernhard Fröhler, Tim Elberfeld, Torsten Möller, et al. „Analysis and comparison of algorithms for the tomographic reconstruction of curved fibres“. In: *Nondestructive Testing and Evaluation* 35.3 (2020), pp. 328–341. DOI: 10.1080/10589759.2020.1774583 (cit. on p. 8).

[18] Marie-Pierre Gagnon, Marie Desmartis, Michel Labrecque, et al. „Systematic Review of Factors Influencing the Adoption of Information and Communication Technologies by Healthcare Professionals“. In: *Journal of medical systems* 36 (Feb. 2012), pp. 241–77. DOI: 10.1007/s10916-010-9473-4 (cit. on p. 5).

[19] S.C. Garcea, Y. Wang, and P.J. Withers. „X-ray computed tomography of polymer composites“. In: *Composites Science and Technology* 156 (2018), pp. 305–319. ISSN: 0266-3538. DOI: 10.1016/j.compscitech.2017.10.023 (cit. on p. 8).

[20] Sabine Graf and Beate List. „An Evaluation of Open Source E-Learning Platforms Stressing Adaptation Issues“. In: *Proceedings of the Fifth IEEE International Conference on Advanced Learning Technologies*. ICALT '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 163–165. ISBN: 0-7695-2338-2. DOI: 10.1109/ICALT.2005.54 (cit. on p. 5).

[21] Daniel Haak, Charles-E. Page, and Thomas M. Deserno. „A Survey of DICOM Viewer Software to Integrate Clinical Research and Medical Imaging“. In: *Journal of Digital Imaging* 29.2 (2016), pp. 206–215. ISSN: 1618-727X. DOI: 10.1007/s10278-015-9833-1 (cit. on p. 5).

[22] C. Heinzl and S. Stappen. „STAR: Visual Computing in Materials Science“. In: *Computer Graphics Forum* 36.3 (2017), pp. 647–666. DOI: 10.1111/cgf.13214 (cit. on p. 7).

[23] H.M. Hsiao and I.M. Daniel. „Effect of fiber waviness on stiffness and strength reduction of unidirectional composites under compressive loading“. In: *Composites Science and Technology* 56.5 (1996), pp. 581–593. ISSN: 0266-3538. DOI: 10.1016/0266-3538(96)00045-0 (cit. on p. 8).

[24] Dr. H. K. Huang. *PACS and Imaging Informatics: Basic Principles and Applications*. 2nd edition. Wiley-Blackwell. ISBN: 978-0-470-37372-9 (cit. on p. 5).

[25] ISO Central Secretary. *ISO 25178 - Geometrical product specifications (GPS) - Surface texture: Areal*. Standard. (all parts of the standard referenced). Geneva, CH: International Organization for Standardization, 2016 (cit. on p. 1).

[26] N. Jeppesen, L.P. Mikkelsen, A.B. Dahl, A.N. Christensen, and V.A. Dahl. „Quantifying effects of manufacturing methods on fiber orientation in unidirectional composites using structure tensor analysis“. In: *Composites Part A: Applied Science and Manufacturing* 149 (2021), p. 106541. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2021.106541 (cit. on p. 8).

[27] Radmir Karamov, Luca M. Martulli, Martin Kerschbaum, et al. „Micro-CT based structure tensor analysis of fibre orientation in random fibre composites versus high-fidelity fibre identification methods“. In: *Composite Structures* 235 (2020), p. 111818. ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2019.111818 (cit. on p. 8).

[28] Kitware Inc. *The Visualization ToolKit*. [Internet]. 2022. URL: https://www.vtk.org (cit. on pp. 11, 79).

[29] K.K. Kratmann, M.P.F. Sutcliffe, L.T. Lilleheden, R. Pyrz, and O.T. Thomsen. „A novel image analysis procedure for measuring fibre misalignment in unidirectional fibre composites“. In: *Composites Science and Technology* 69.2 (2009), pp. 228–238. ISSN: 0266-3538. DOI: 10.1016/j.compscitech.2008.10.020 (cit. on p. 8).

[30] M. Krause, J. Hausherr, Bernhard Burgeth, C. Herrmann, and Walter Krenkel. „Determination of the fibre orientation in composites using the structure tensor and local X-ray transform“. In: *Journal of Materials Science* 45 (Feb. 2010), pp. 888–896. DOI: 10.1007/s10853-009-4016-4 (cit. on p. 8).

[31] J.P. Kruth, M. Bartscher, S. Carmignato, et al. „Computed tomography for dimensional metrology“. In: *CIRP Annals* 60.2 (2011), pp. 821–842. ISSN: 0007-8506. DOI: 10.1016/j.cirp.2011.05.006 (cit. on p. 8).

[32] Pritam Kulkarni, Kiran D. Mali, and Sandeep Singh. „An overview of the formation of fibre waviness and its effect on the mechanical performance of fibre reinforced polymer composites“. In: *Composites Part A: Applied Science and Manufacturing* 137 (2020), p. 106013. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2020.106013 (cit. on p. 8).

[33] Denis Le Bihan and E. Breton. „Imagerie de diffusion in-vivo par résonance magnétique nucléaire“. In: *Comptes-Rendus de l'Académie des Sciences* 93.5 (Dec. 1985), pp. 27–34. URL: https://hal.archives-ouvertes.fr/hal-00350090 (cit. on p. 9).

[34] Yong Hoon Lee, Seok Won Lee, Jae Ryoun Youn, Kwansoo Chung, and Tae Jin Kang. „Characterization of fiber orientation in short fiber reinforced composites with an image processing technique“. In: *Materials Research Innovations* 6.2 (2002), pp. 65–72. DOI: 10.1007/s10019-002-0180-8 (cit. on p. 8).

[35] M. Levoy. „Display of surfaces from volume data“. In: *IEEE Computer Graphics and Applications* 8.3 (1988), pp. 29–37. ISSN: 1558-1756. DOI: 10.1109/38.511 (cit. on p. 7).

[36] T. Lindeberg. „Scale-space for discrete signals“. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12.3 (1990), pp. 234–254. ISSN: 1939-3539. DOI: 10.1109/34.49051 (cit. on p. 82).

[37] Tony Lindeberg. 1st ed. The Springer International Series in Engineering and Computer Science. Springer, 1994, pp. XII, 424. ISBN: 978-1-4757-6465-9. DOI: 10.1007/978-1-4757-6465-9 (cit. on p. 82).

[38] William E. Lorensen and Harvey E. Cline. „Marching Cubes: A High Resolution 3D Surface Construction Algorithm“. In: *SIGGRAPH Comput. Graph.* 21.4 (1987), pp. 163–169. ISSN: 0097-8930. DOI: 10.1145/37402.37422 (cit. on pp. 6, 8).

[39] Benoit Mandelbrot. „How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension“. In: *Science* 156.3775 (1967), pp. 636–638. DOI: 10.1126/science.156.3775.636 (cit. on pp. 4, 78).

[40] Mahoor Mehdikhani, Christian Breite, Yentl Swolfs, et al. „Combining digital image correlation with X-ray computed tomography for characterization of fiber orientation in unidirectional composites“. In: *Composites Part A: Applied Science and Manufacturing* 142 (2021), p. 106234. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2020.106234 (cit. on p. 8).

[41] Klaus-Dietmar Merboldt, Wolfgang Hanicke, and Jens Frahm. „Self-diffusion NMR imaging using stimulated echoes“. In: *Journal of Magnetic Resonance (1969)* 64.3 (1985), pp. 479–486. ISSN: 0022-2364. DOI: 10.1016/0022-2364(85)90111-8 (cit. on p. 9).

[42] Koichi Mizukami, Yoshihiro Mizutani, Kenshi Kimura, et al. „Visualization and size estimation of fiber waviness in multidirectional CFRP laminates using eddy current imaging“. In: *Composites Part A: Applied Science and Manufacturing* 90 (2016), pp. 261–270. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2016.07.008 (cit. on p. 9).

[43] Susumu Mori and Peter C. M. van Zijl. „Fiber tracking: principles and strategies – a technical review“. In: *NMR in Biomedicine* 15.7-8 (2002), pp. 468–480. DOI: 10.1002/nbm.781 (cit. on p. 9).

[44] Paul Nagy. „Open Source in Imaging Informatics“. In: *Journal of Digital Imaging* 20.1 (2007), pp. 1–10. ISSN: 1618-727X. DOI: 10.1007/s10278-007-9056-1 (cit. on p. 5).

[45] National Electrical Manufacturers Association (NEMA). *Digital Imaging and Communications in Medicine (DICOM)*. [Internet]. 2018. URL: https://www.dicomstandard.org/current/ (cit. on p. 5).

[46] J. W. Nelson, T. W. Riddle, and D. S. Cairns. „Effects of defects in composite wind turbine blades – Part 1: Characterization and mechanical testing“. In: *Wind Energy Science* 2.2 (2017), pp. 641–652. DOI: 10.5194/wes-2-641-2017 (cit. on p. 8).

[47] H. Nyquist. „Certain Topics in Telegraph Transmission Theory“. In: *Transactions of the American Institute of Electrical Engineers* 47.2 (1928), pp. 617–644. ISSN: 2330-9431. DOI: 10.1109/T-AIEE.1928.5055024 (cit. on pp. 4, 78).

[48] William H. Oldendorf. „The quest for an image of brain“. In: *Neurology* 28.6 (1978), pp. 517–517. ISSN: 0028-3878. DOI: 10.1212/WNL.28.6.517 (cit. on p. 2).

[49] P. Pinter, S. Dietrich, B. Bertram, et al. „Comparison and error estimation of 3D fibre orientation analysis of computed tomography image data for fibre reinforced composites“. In: *NDT & E International* 95 (2018), pp. 26–35. ISSN: 0963-8695. DOI: 10.1016/j.ndteint.2018.01.001 (cit. on p. 8).

[50] Giuseppe Lo Presti, Marina Carbone, Damiano Ciriaci, et al. „Assessment of DICOM Viewers Capable of Loading Patient-specific 3D Models Obtained by Different Segmentation Platforms in the Operating Room“. In: *Journal of Digital Imaging* 28.5 (2015), pp. 518–527. ISSN: 1618-727X. DOI: 10.1007/s10278-015-9786-4 (cit. on p. 5).

[51] Johann Radon. „On the determination of functions from their integral values along certain manifolds“. In: *IEEE Transactions on Medical Imaging* 5.4 (1986). English translation, pp. 170–176. ISSN: 1558-254X. DOI: 10.1109/TMI.1986.4307775. URL: https://people.csail.mit.edu/bkph/courses/papers/Exact_Conebeam/Radon_Deutsch_1917.pdf (cit. on p. 3).

[52] Guillermo Requena, Georg Fiedler, Bernhard Seiser, et al. „3D-Quantification of the distribution of continuous fibres in unidirectionally reinforced composites“. In: *Composites Part A: Applied Science and Manufacturing* 40.2 (2009), pp. 152–163. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2008.10.014 (cit. on p. 8).

[53] Caroline Richmond. „Sir Godfrey Hounsfield“. In: *BMJ* 329.7467 (2004), p. 687. ISSN: 0959-8138. DOI: 10.1136/bmj.329.7467.687 (cit. on p. 2).

[54] Katharina Robb, Oliver Wirjadi, and Katja Schladitz. „Fiber Orientation Estimation from 3D Image Data: Practical Algorithms, Visualization, and Interpretation“. In: *7th International Conference on Hybrid Intelligent Systems (HIS 2007)*. 2007, pp. 320–325. DOI: 10.1109/HIS.2007.26 (cit. on p. 8).

[55] Alvaro Rocha. „Evolution of Information Systems and Technologies Maturity in Healthcare“. In: *International journal of healthcare information systems and informatics* 6 (Apr. 2011), pp. 28–36. DOI: 10.4018/jhisi.2011040103 (cit. on p. 5).

[56] Hanno Scharr. „Optimale Operatoren in der Digitalen Bildverarbeitung“. PhD thesis. Heidelberg University, Jan. 2000. Chap. Anhang B, Optimale Parameter einiger Filter, pp. 145–156. DOI: 10.11588/heidok.00000962. URL: http://www.ub.uni-heidelberg.de/archiv/962 (cit. on p. 75).

[57] Ludwig Schöttl, Dominik Dörr, Pascal Pinter, et al. „A novel approach for segmenting and mapping of local fiber orientation of continuous fiber-reinforced composite laminates based on volumetric images“. In: *NDT & E International* 110 (2020), p. 102194. ISSN: 0963-8695. DOI: 10.1016/j.ndteint.2019.102194 (cit. on p. 8).

[58] W. Schroeder, K. Martin, B. Lorensen, and Inc Kitware. *The Visualization Toolkit: An Object-oriented Approach to 3D Graphics*. Kitware, 2006. ISBN: 9781930934191. URL: https://gitlab.kitware.com/vtk/textbook/raw/master/VTKBook/VTKTextBook.pdf (cit. on p. 9).

[59] R.M. Sencu, Z. Yang, Y.C. Wang, et al. „Generation of micro-scale finite element models from synchrotron X-ray CT images for multidirectional carbon fibre reinforced composites“. In: *Composites Part A: Applied Science and Manufacturing* 91 (2016), pp. 85–95. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2016.09.010 (cit. on pp. 8, 73).

[60] J.-A. Serret. „Sur quelques formules relatives à la théorie des courbes à double courbure.“ fre. In: *Journal de Mathématiques Pures et Appliquées* (1851), pp. 193–207. URL: http://eudml.org/doc/235002 (cit. on p. 77).

[61] C.E. Shannon. „Communication in the Presence of Noise“. In: *Proceedings of the IRE* 37.1 (1949), pp. 10–21. ISSN: 2162-6634. DOI: 10.1109/JRPROC.1949.232969 (cit. on pp. 4, 78).

[62] Ilya Straumit, Stepan V. Lomov, and Martine Wevers. „Quantification of the internal structure and automatic generation of voxel models of textile composites from X-ray computed tomography data“. In: *Composites Part A: Applied Science and Manufacturing* 69 (2015), pp. 150–158. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2014.11.016 (cit. on p. 8).

[63] D G Taylor and M C Bushell. „The spatial mapping of translational diffusion coefficients by the NMR imaging technique“. In: *Physics in Medicine and Biology* 30.4 (1985), pp. 345–349. DOI: 10.1088/0031-9155/30/4/009 (cit. on p. 9).

[64] Michael Thor, Ulrich Mandel, M. Nagler, et al. „Numerical and experimental investigation of out-of-plane fiber waviness on the mechanical properties of composite materials“. In: *International Journal of Material Forming* 14 (Jan. 2021). DOI: 10.1007/s12289-020-01540-5 (cit. on p. 8).

[65] Gianluca Valeri, Francesco Antonino Mazza, Stefania Maggi, et al. „Open source software in a practical approach for post processing of radiologic images“. In: *La radiologia medica* 120.3 (2015), pp. 309–323. ISSN: 1826-6983. DOI: 10.1007/s11547-014-0437-5 (cit. on p. 5).

[66] C.-F. Westin, S.E. Maier, H. Mamata, et al. „Processing and visualization for diffusion tensor MRI“. In: *Medical Image Analysis* 6.2 (2002), pp. 93–108. ISSN: 1361-8415. DOI: 10.1016/S1361-8415(02)00053-1 (cit. on pp. 9, 72).

[67] Rogier van de Wetering and Ronald Batenburg. „A PACS maturity model: A systematic meta-analytic review on maturation and evolvability of PACS in the hospital enterprise“. In: *International journal of medical informatics* 78 (Sept. 2008), pp. 127–40. DOI: 10.1016/j.ijmedinf.2008.06.010 (cit. on p. 5).

[68] E. T. Whittaker. „XVIII.—On the Functions which are represented by the Expansions of the Interpolation-Theory“. In: *Proceedings of the Royal Society of Edinburgh* 35 (1915), pp. 181–194. DOI: 10.1017/S0370164600017806 (cit. on pp. 4, 78).

[69] D. Wilhelmsson and L.E. Asp. „A high resolution method for characterisation of fibre misalignment angles in composites“. In: *Composites Science and Technology* 165 (2018), pp. 214–221. ISSN: 0266-3538. DOI: 10.1016/j.compscitech.2018.07.002 (cit. on p. 8).

[70] D. Wilhelmsson, R. Gutkin, F. Edgren, and L.E. Asp. „An experimental study of fibre waviness and its effects on compressive properties of unidirectional NCF composites“. In: *Composites Part A: Applied Science and Manufacturing* 107 (2018), pp. 665–674. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2018.02.013 (cit. on p. 8).

Colophon

This document and the whole research has been made exclusively with Free/Libre and Open Source Software.