

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

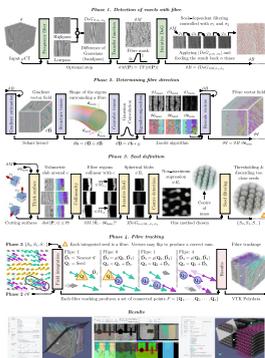
Adrià Julià i Juanola, Marc Ruiz i Altisent, Imma Boada i Oliveras*

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 xxx

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.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

* Corresponding author.

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

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

1. Introduction

Current μ CT (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 μ CT 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 μ CT 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 μ CT 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) μ CT 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 μ CT 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 μ CT 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 μ CT 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

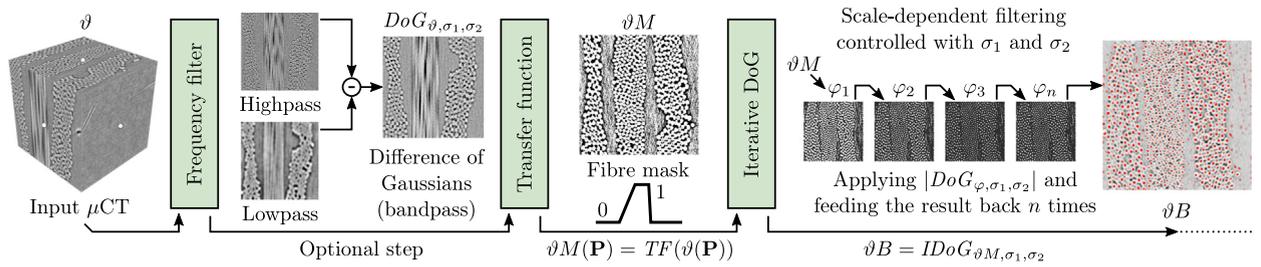


Fig. 1. Block diagram of the *Phase I* taking a μCT volume as the input. The ϑM segmentation is bandpass-filtered using the Iterative Difference of Gaussians filter in order to conform ϑB (overlaid red blobs). Images produced using the *cfnp_i17* dataset. (For interpretation of the references to colour 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 $\vec{\mathbf{t}}$. Some functions may be declared

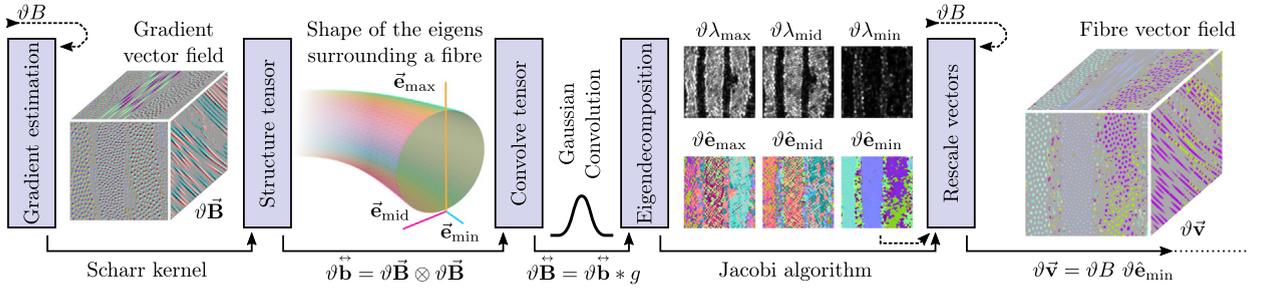


Fig. 2. Block diagram of the *Phase 2* using the ϑB from the previous *Phase 1* in order to produce the $\vartheta \vec{v}$ vector field representing the fibre directionality and magnitude. Images produced using the *cfip_i17* dataset mapping vector components into RGB colours. (For interpretation of 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 μ CT 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.

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

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.

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

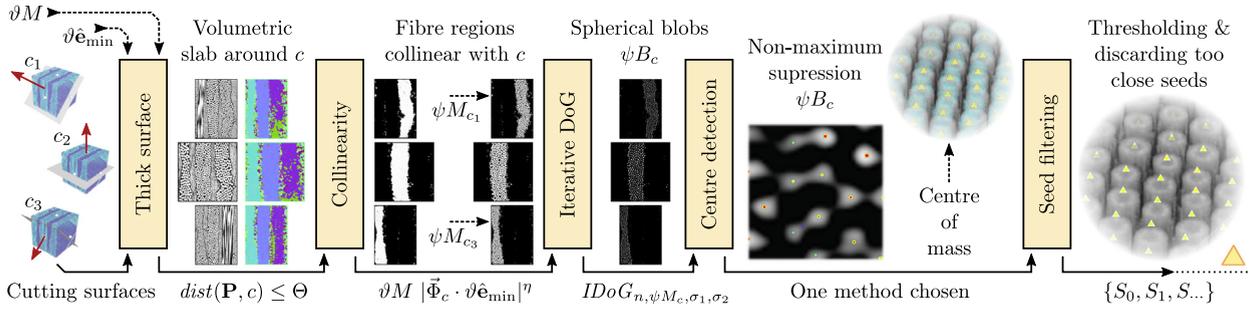


Fig. 3. Block diagram of the *Phase 3* where given a set of user defined cutting surfaces (three in this example), a set of seeds (yellow triangles) at the fibre centres is produced. Images produced using the *cfrp_i17* and Creveling et al. [7] (round) datasets. (For interpretation of the references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

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

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.

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

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)).

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

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

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.

$$G_{\varphi,\sigma} = \varphi * g_{\sigma} \quad (6a)$$

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

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.

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

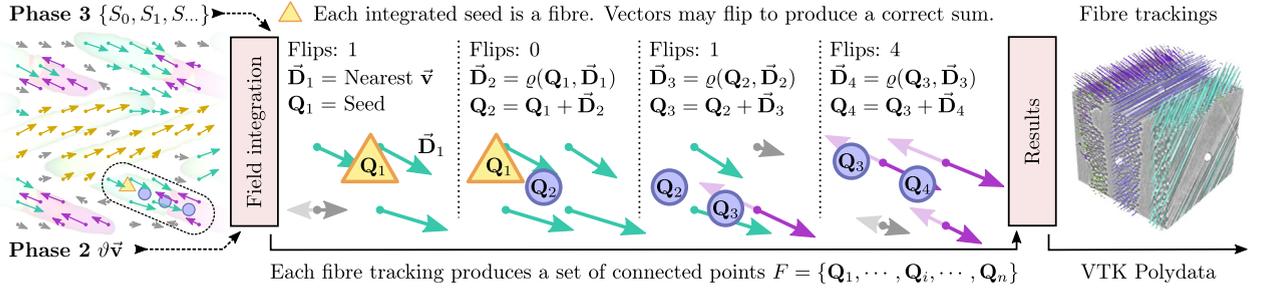


Fig. 4. Block diagram of the *Phase 4* where given the vector field $\vartheta \vec{v}$ from *Phase 2* (coloured arrows) tracks the fibres starting from *Phase 3* seeds (yellow triangles) by resolving each Q (blue points), flipping the vectors (semi-transparent arrows) if necessary. Images produced using 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.)

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

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 from the target diameter.

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

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 $\vartheta \vec{B}$ (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.

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

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

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

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

$$\vartheta \overleftrightarrow{B} = \vartheta \overleftrightarrow{b} * g_\sigma \quad (11a)$$

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

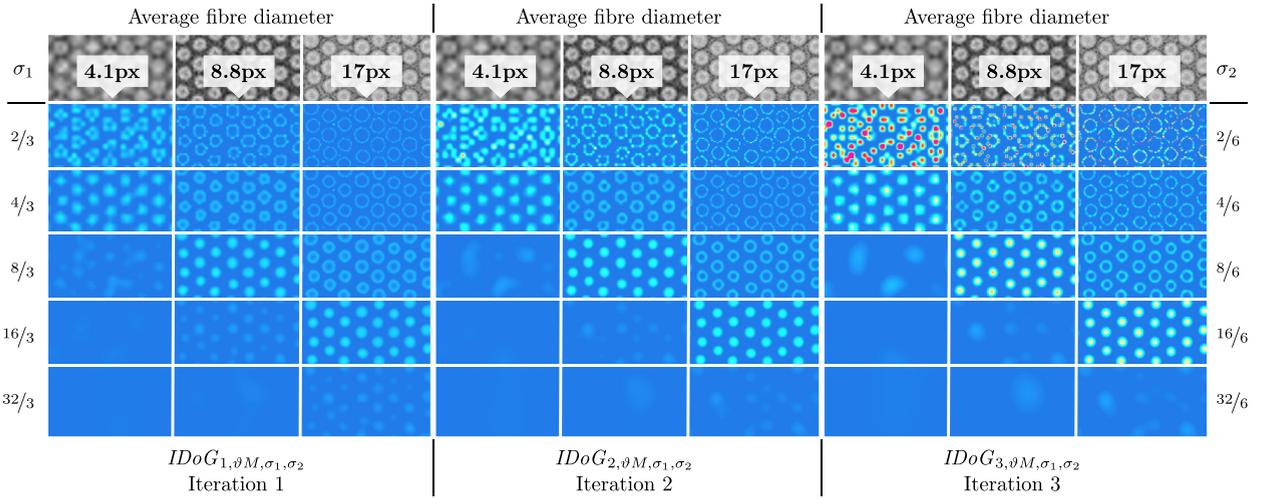


Fig. 5. Effect of altering σ_1 and σ_2 (y axis) of an IDoG filter (grouped by iterations) when the fibre diameter (x axis) varies. To allow a meaningful comparison, intensities are scaled within $[0, 10]$ range (blue to red). Images produced using Creveling et al. [7] datasets. (For interpretation of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

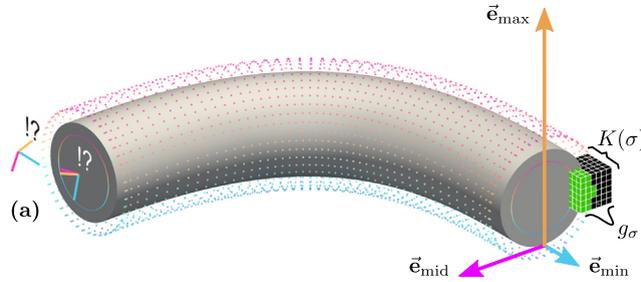


Fig. 6. An idealised fibre, with a tensor at the boundary and a cube representing the convolution kernel, alongside the effect on tensors at an unreachable $K(\sigma)$ distance (a).

Then, when each tensor voxel is eigendecomposed (diagonalised) using the Jacobi method [38,39] its λ_{\max} , λ_{mid} , λ_{\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 $C = \{C_1, \dots, C_m\}$ (Fig. 3). From each cutting surface c an undetermined amount of k seeds $S = \{S_1, \dots, S_k\}$ 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

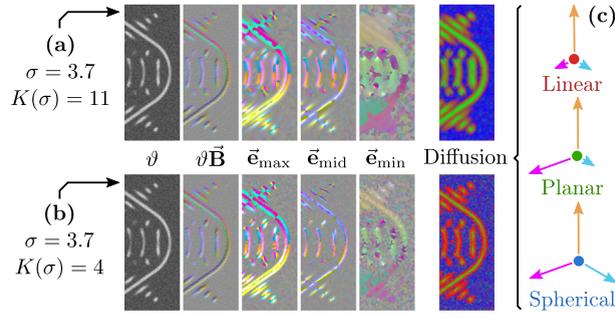


Fig. 7. The effect of varying the kernel size (a, b) on the tensor diffusion (c) mapped as follows: linear anisotropy (red), planar anisotropy (green), and isotropy (blue). Images produced using the *scaphandre* synthetic dataset. (For interpretation of the references to colour in this figure legend,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

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

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 .

$$\psi B_c(\mathbf{P}) = \text{IDoG}_{n, \psi M_c, \sigma_1, \sigma_2}(\mathbf{P}) \quad (14)$$

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 $\{\mathbf{S}_1, \dots, \mathbf{S}_k\}$ 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) $F = \{\mathbf{Q}_1, \dots, \mathbf{Q}_i, \dots, \mathbf{Q}_n\}$ 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.

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

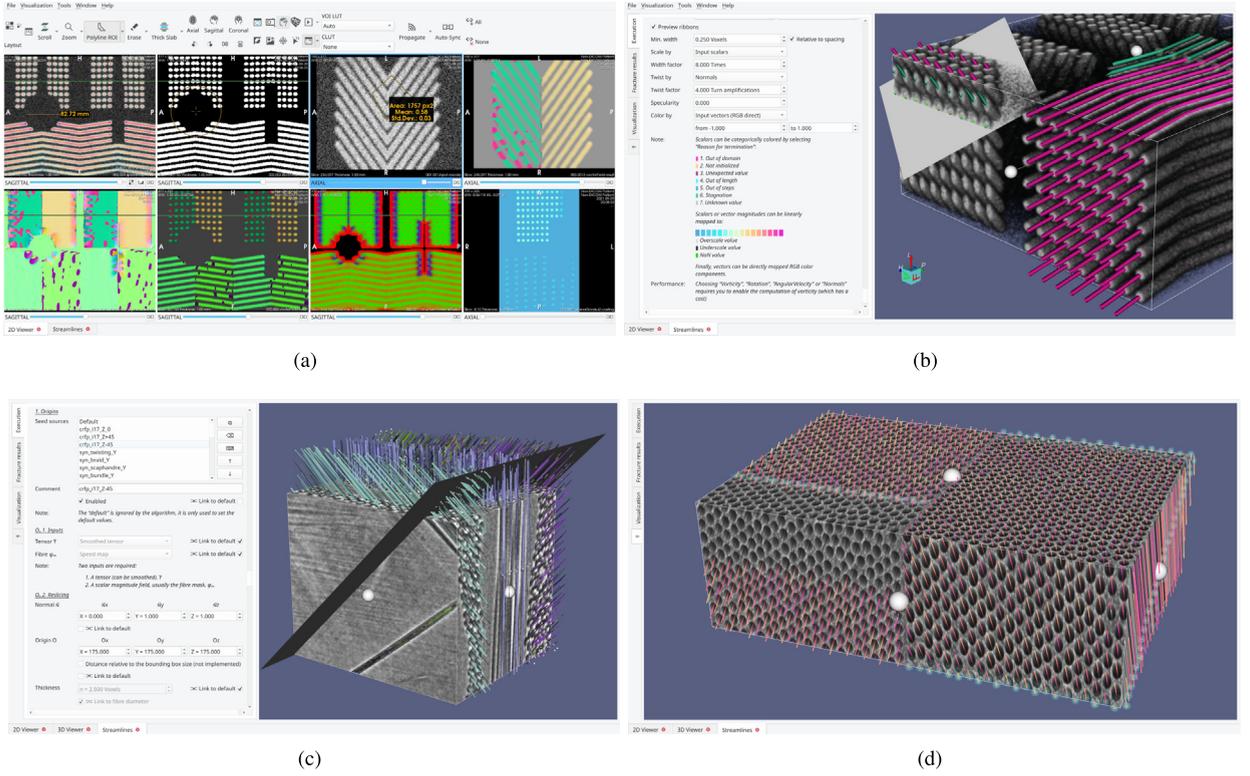


Fig. 8. The 2D interface (a) displaying some of the intermediate volumes of the *bundle* synthetic dataset alongside the 3D visualisations in (b) where one cutting plane is shown and the fibres are coloured according to their reason for termination. (c) The results produced using the real *cfrp_i17* 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.

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

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

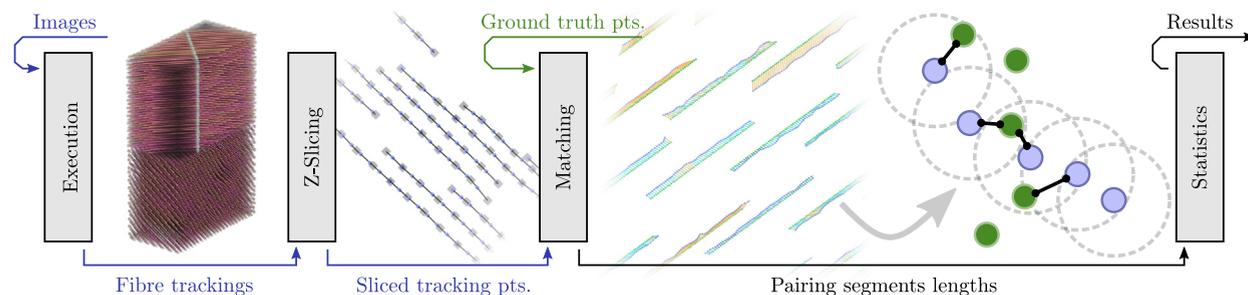


Fig. 9. Steps necessary to convert our fibre tracking results to a suitable Z-sliced point cloud comparable against the respective ground truth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

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 μ CT 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 *cfpr_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.

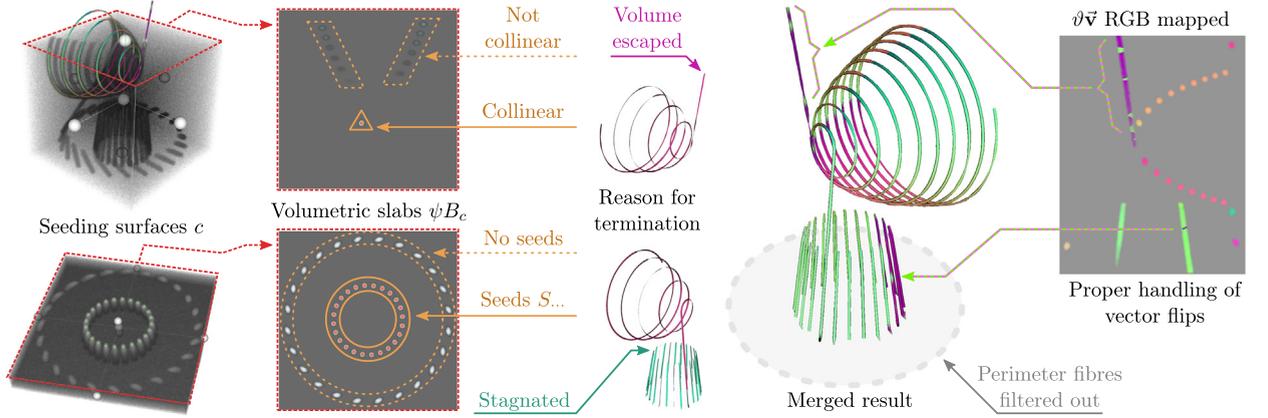


Fig. 10. Example of the method behaviour using the *twisting* synthetic dataset. Seeds (solid yellow) are only placed if they are collinear with the cutting surface (dashed red). Only the selected subset of fibres is tracked using the vector field $\vartheta\vec{v}$. (For interpretation of the references 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 $\vartheta\vec{v}$ 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 obviating it as some of the aforementioned approaches do. This processing is performed using Gaussian convolutions and Scharf [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.

$$\sigma_1 = \frac{\text{diameter}}{3} \quad \sigma_2 = \frac{\text{diameter}}{6} \quad (17)$$

A second phenomenon to consider is when the tensor field $\vartheta\vec{\mathbf{b}}$ components are convolved with their neighbours in order to produce $\vartheta\vec{\mathbf{B}}$. 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 $\hat{\mathbf{e}}_{\min}$ 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

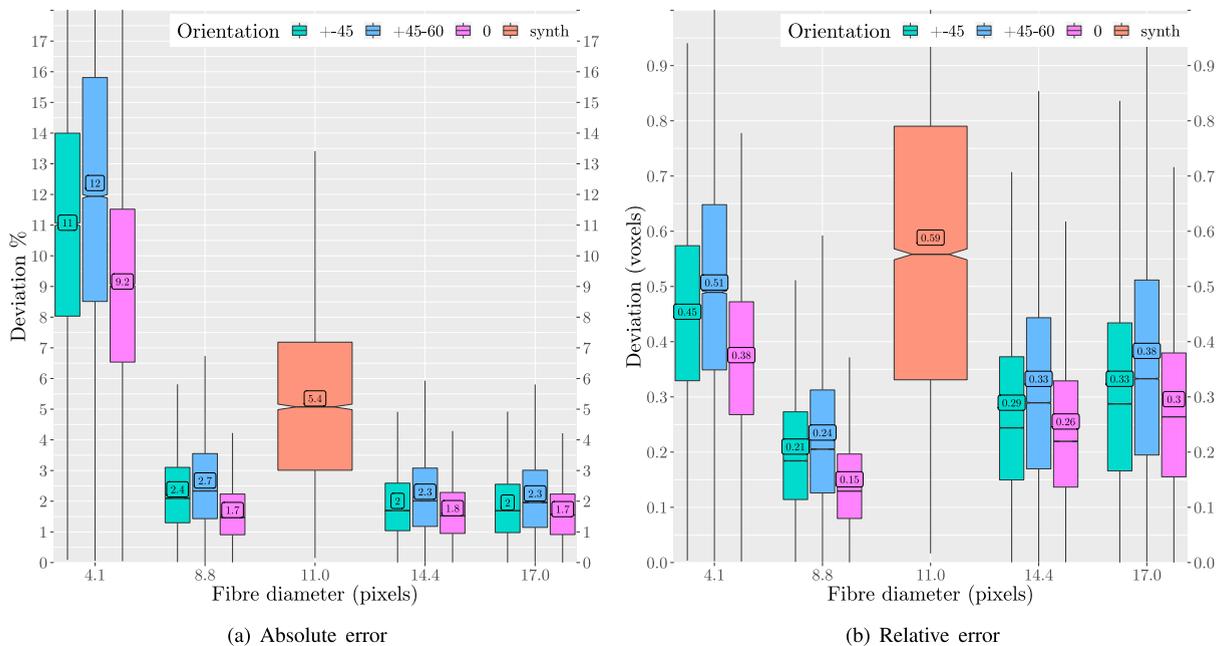


Fig. 11. Absolute (a) and relative (b) error of our approach measured using our synthetic (orange) and Creveling et al. [7] testing datasets.

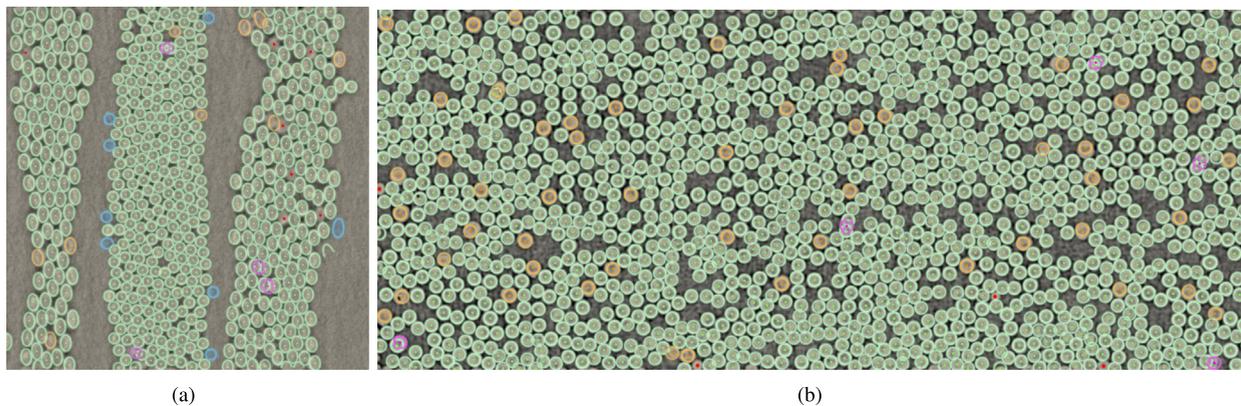


Fig. 12. Verification of the *cfrp_i17* (a) and *leuven_glass* (b) real datasets against our algorithm output (green) by manually identifying misplacements (yellow), duplicates (magenta), phantoms (blue), and misses (red). (For interpretation of the references to colour in this figure legend, the reader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

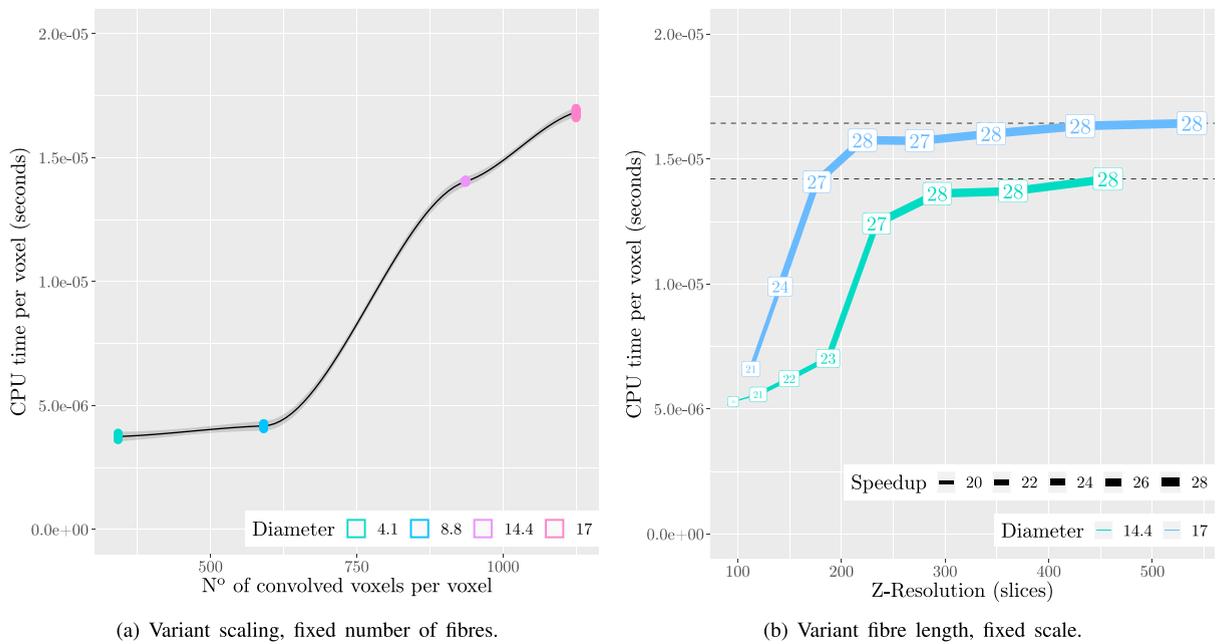


Fig. 13. Growth of the cost when some of the input volume characteristics are held invariant. Creveling et al. [7] datasets were used.

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 $350 \times 350 \times 350$ *cfp_i17* dataset is extrapolated to an extent of $2500 \times 2500 \times 1500$, 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

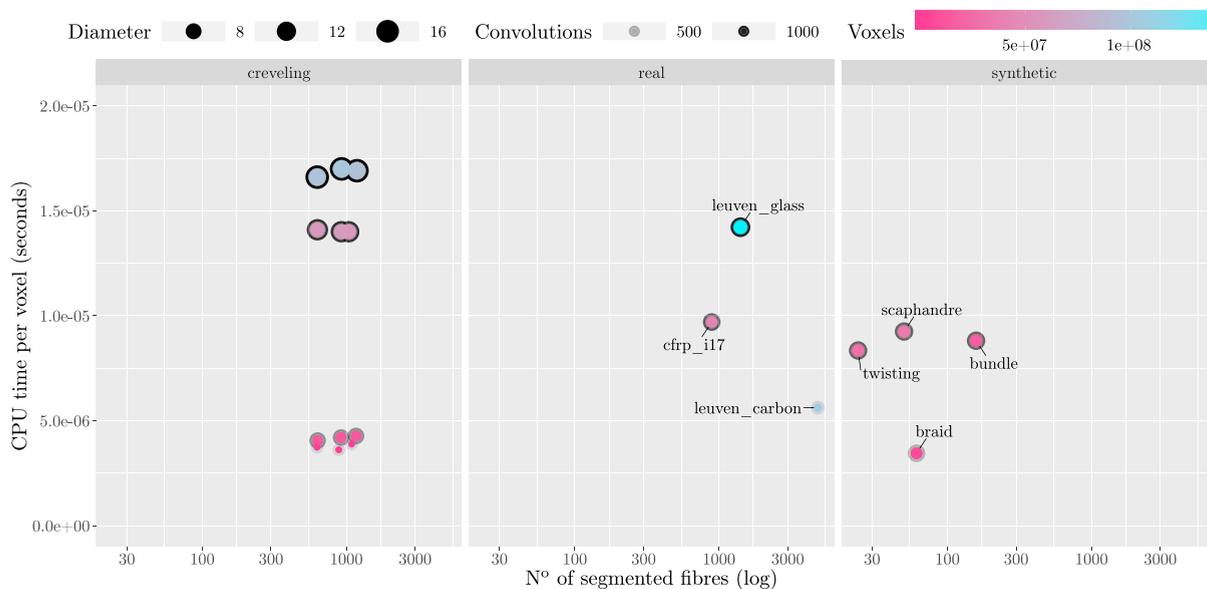


Fig. 14. Summary of the cost growth using all cases under test in relation to their main characteristics. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

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.

For 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 $\vartheta \vec{v}$ 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 μ CT 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.

2a5c58c78c64e26e592866564161c2a7fb27c0342824fcd5c86fa38b059c67 (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 e1deeac8118949ecf630043731b282c1b6fd21700fb75aafa26c7392d3df878 (SHA256) 81fc93def85f981611e4acc9ccdfad23c8143c7b (SHA1)

fhd_real.mkv

e02746bae484368515a8e5e0fa0b8ebe28ee7fe421f8dd1a4319ed7fe01dd05e (SHA256) d082dc00a135c118f783efb2184d0874bcc2521a (SHA1)

fhd_curved.mkv

ece543e87ba796f610998ad26d80cd118717cea975f50dc27c1845d5cfba3da3 (SHA256) c54e29492aa9b76dfb5984e64d7d0f6c6ec0c6dc1 (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](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](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. Kehler, 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](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](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>.