Image Anal Stereol 2026; 45:277-290 doi: 10.5566/ias.4096
Original Research Paper

SINGLE FIBER SEGMENTATION IN 3D IMAGES OF HIGHLY POROUS WOOD FIBER INSULATION MATS

AISWARYA NAIR[1,2], KATJA SCHLADITZ ✉[,1], CLAUDIA REDENBACH[2] AND MARKUS KRONENBERGER[1]

[1]Fraunhofer Institute for Industrial Mathematics, Kaiserslautern, [2]RPTU University Kaiserslautern-Landau
e-mail: aiswarya.nair.manikantan.geetha@itwm-extern.fraunhofer.de, katja.schladitz@itwm.fraunhofer.de, claudia.redenbach@rptu.de, markus.kronenberger@itwm.fraunhofer.de
(Received July 20, 2026; revised October 13, 2026; accepted October 16, 2026)

ABSTRACT

The thermal conductivity of highly porous wood fiber insulation mats is hard to predict but known to be strongly determined by the geometry and spatial arrangement of the fibers. Characterization of individual fibers is essential for numerically simulating and ultimately optimizing the insulation mats by altering their microstructure. The latter requires fitting stochastic geometric models to the mats' microstructure and this in turn number-weighted distributions of fiber characteristics such as length, thickness, orientation, and number of fiber-fiber contacts. Therefore, instance segmentation of wood fibers in 3D images obtained by computed tomography is needed. The wood fibers' features render this task unsolvable for established classical approaches: Besides all the challenges coming with the elongated shape, they are hollow, yet the lumen as well as the fiber walls are frequently damaged by the mat production process resulting in high variability of the fibers' cross-sectional sizes and shapes. Moreover, the mats additionally include fiber bundles and polymer fibers. For machine learning, there is not enough data. We therefore suggest a new method detecting fiber junctions by an orientation coherence map derived from the 3D structure tensor. Removal of these junctions yields the individual fiber objects, whose original shape is restored in a post-processing step. Our method proves to be applicable even to stitched 3D images with a size of $4\,000^2 \times 2\,100$ voxels. Moreover, comparing various fiber features derived from the fully resolved image and a coarsened version, we verify that most geometric characteristics are robust with respect to downsampling the images by a factor of three. The full resolution is however needed for measuring the size of fiber-fiber contact regions.

Keywords: fiber-fiber contact, instance segmentation, structure tensor, synchrotron radiation, X-ray computed tomography.

INTRODUCTION

The microstructure of fibers plays a crucial role in determining the mechanical, biological, and functional properties of many natural and man-made materials, such as fiber reinforced composites, paper products, and wood fiber based materials. Understanding the microstructure of such fibrous materials is essential in predicting their material properties by simulations. The microstructure can be imaged non-destructively by X-ray computed tomography (CT). The resulting volumetric (3D) datasets contain valuable information on the spatial arrangement and morphology of the fibers within the material. However, this information is fully revealed only if the respective microstructure components can be extracted (segmented) from the gray-value image.

A common first step in CT data analysis is therefore phase or semantic segmentation, where each voxel is assigned to a material class, for example, fibers or background. This step is sufficient for measurements of global properties such as fiber volume fraction, mean fiber orientation, and porosity. However, if the materials properties are not only to be predicted but the final goal is to improve them by optimizing the microstructure as sketched in Fig. 1, then a stochastic geometry model has to be fit to the observed microstructure.

Here, we consider highly porous ($\approx 95\%$) wood fiber mats for thermal insulation, manufactured from softwood. Compared to man-made fibers such as glass fibers, wood fibers exhibit substantially more complex geometries. They have highly irregular, non-circular cross-sections as a result of their biological structure. Intact wood fibers possess a lumen – a hollow interior channel running along the fiber axis – that causes the cross-section to appear as an annular ring rather than a solid disk. As the production process of the insulation mats involves milling and cooking, fiber walls are frequently torn and the lumen is irregular or narrowed (Fig. 2).

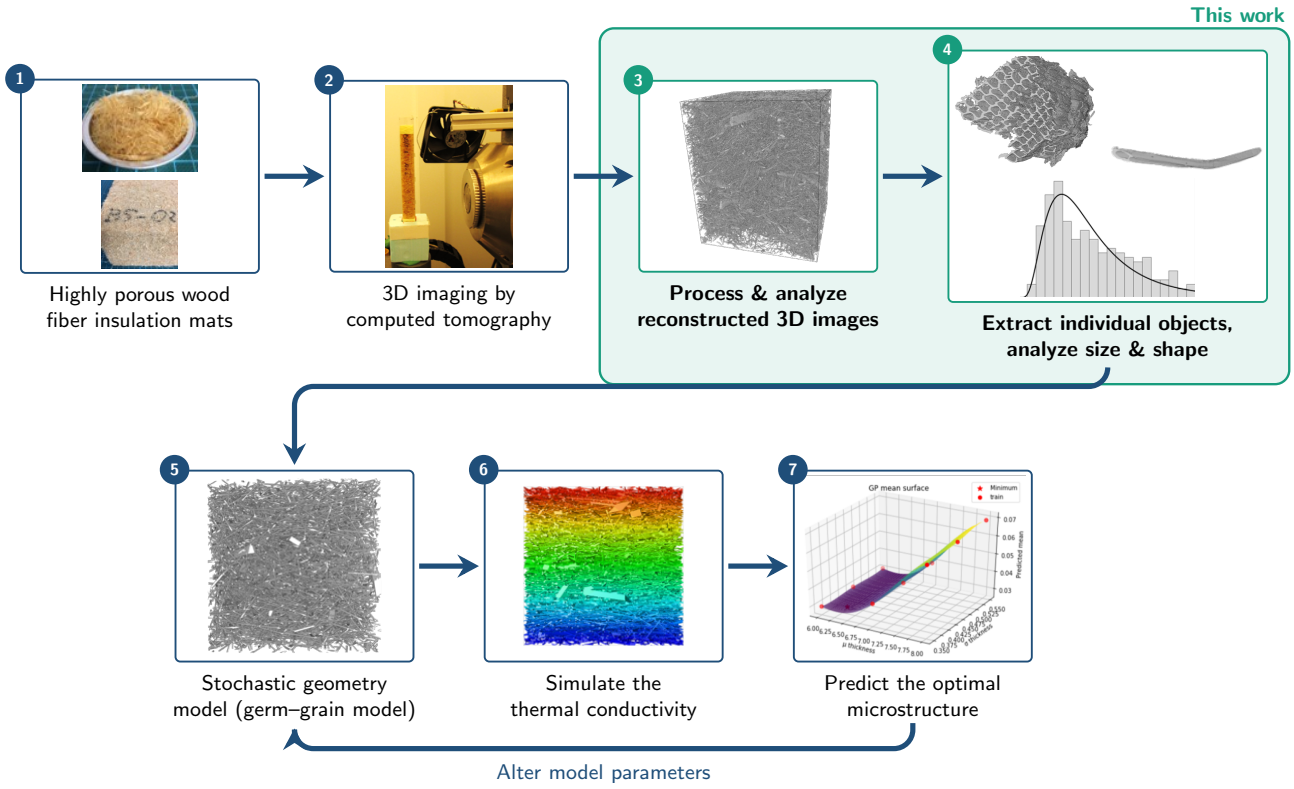


Fig. 1. *Scheme of the micro-structural optimization we are aiming at. True optimization (5-7) is enabled by fitting a stochastic geometry model (4-5) to the observed micro-structure. The fitting step in turn requires number-weighted empirical size and shape distributions (4) which necessitates segmentation of individual objects (3-4). That is, instance segmentation.*

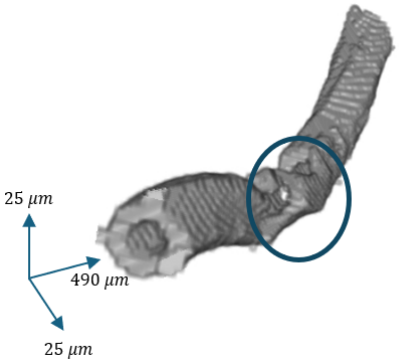


Fig. 2. *3D rendering of a single fiber with hollow part (lumen), extracted from a synchrotron radiation X-ray sub-micro CT image with voxel size 0.65 µm. The circle marks a local collapse of the fiber wall. Such deformations introduce irregular surface geometry and ambiguous cross-sectional boundaries. Annotated numbers are the approximate dimensions: length (straight extent along the main axis)* $\approx 490\,\mu m$*, width* $\approx$ *thickness* $\approx 25\,\mu m$.

For instance segmentation of approximately spherical or blob-like objects, the classical morphological approach is often effective: The local maxima of a distance transform mark object centers, and touching particles are separated by a watershed transform (Vincent and Soille, 1991; Beucher and Meyer, 1993). This approach is particularly suitable for compact, convex, similarly sized particles with one dominant distance maximum per object, but known to over-segment elongated structures like fibers. The tubular geometry produces multiple markers, causing single fibers to be incorrectly split into several segments. Locally broken fiber walls worsen this effect.

Many earlier 3D fiber segmentation methods relied on centerline-based approaches, classically grounded in the medial axis transformation (Blum, 1967) and topology-preserving thinning algorithms (Lee *et al.*, 1994; Palágyi and Kuba, 1999). By reducing fibers to their centerline, these methods enable analysis of orientation, length, and curvature. They have been mainly applied to fibrous materials in which fibers are solid and nearly cylindrical with smooth and well-defined boundaries. For example, Lux (2013) paired skeleton segments using curvature and radius continuity constraints to recover individual fibers. These centerline tracking methods have also been used to quantify fiber distributions and fiber arrangements

in tomographic data (Requena *et al.*, 2009; Latil *et al.*, 2011; Gaiselmann *et al.*, 2013; Zhou *et al.*, 2016; Sencu *et al.*, 2016). Further extensions refine the centerline detection through template-based cross-correlation (Zauner *et al.*, 2015; Czabaj *et al.*, 2014; Creveling *et al.*, 2019), geometry-based voting (Pinter *et al.*, 2016), and machine learning-based extraction (Emerson *et al.*, 2017). However, the basic assumption for such methods is violated for hollow fibers, such as wood fibers. The medial axis of a hollow cross-section no longer collapses to a single point but instead forms a ring-like structure, producing ambiguous or degenerate skeletons. Likewise, slice-based centroid detection becomes unreliable when the fiber wall is incomplete or irregular, as is typical for wood fibers. Therefore, an instance segmentation approach explicitly designed for the structural complexity of wood fibers is required.

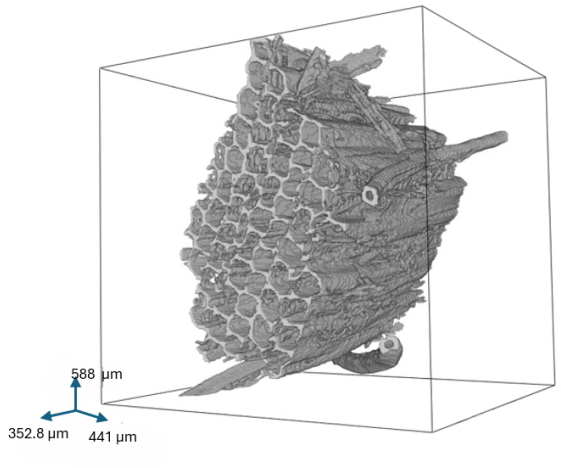


Fig. 3. *3D rendering of wood structure obtained from an X-ray micro CT image, illustrating a prominent fiber bundle with multiple lumens, varying wall thickness, and wall discontinuities. Individual fibers of varying orientation adhering to the bundle surface emphasize the structural complexity and geometric irregularity of the wood fiber system. Annotated numbers are the dimensions of the cuboid.*

To address these challenges, several researchers have explored segmentation of wood fibers using the lumen as a marker. Wernersson *et al.* (2009) first closed the pores of the fibers with a curvature-based filter, and then separated individual fibers from the segmented lumen using a graph-cut formulation. Sharma *et al.* (2015) tracked the lumen in 3D by connected component labelling across slices and reconstructed fiber walls outward from each tracked lumen using a distance transform. Such lumen based methods depend on the lumen being closed and continuous whereas collapsed or crushed lumen merge with the background and cannot be used for segmentation. Sharma *et al.* (2015) segmented more than 40% of the fiber volume in paper handsheets, a substantial improvement over earlier region-growing approaches, but the remaining fibers could not be recovered. In wood fiber insulation mats, production induces frequent wall collapses resulting in locally narrowed, irregular, or absent lumen (Fig. 2). Additionally, we face the challenge of fiber bundles (Fig. 3) featuring multiple lumens within a single connected structure, rendering it impossible to uniquely associate lumen regions with individual microstructure components. Thus, segmentation algorithms based on lumen as markers are unsuitable.

Deep learning approaches face fundamental practical obstacles when applied to materials such as wood fiber insulation mats: they require large volumes of annotated 3D training data. For hollow, irregularly shaped wood fibers this must be produced by manual annotation, since synthetic training data cannot yet faithfully reproduce the lumen geometry and wall discontinuities encountered in real samples. Furthermore, networks trained on one fiber type or specific imaging conditions do not readily generalize to other conditions without further re-training, and the lack of publicly available annotated datasets for wood fiber CT volumes means that each new study must bear the full cost of ground-truth generation. Kibleur *et al.* (2022) demonstrated that even for identifying only the wood fiber bundles from fiberboards, a dedicated 3D U-Net trained on manually annotated real CT data was required to achieve reliable results.

An alternative strategy for fiber separation is to identify and subsequently remove the contact regions between fibers. Our approach is motivated by Viguié *et al.*'s (2013) analysis of the local orientation variation, which detects fiber-fiber contacts by thresholding the local angular deviation C_p of the principal direction field derived from a directed-distance-based orientation tensor. While effective for fibers with complex cross-sections, the method has some limitations. First, the underlying orientation tensor is computed from directed Euclidean distance transforms along a discrete set of predefined directions. Exploiting the strong anisotropy of their material, they sampled 52 directions predominantly in the fiber plane. In the insulation mats however, a significant fraction of the fibers is oriented out of plane. Wirjadi *et al.* (2016) compared four local orientation estimation algorithms and showed that derivative-based approaches, specifically the structure tensor and the Hessian matrix, achieve errors of only a few degrees, compared to methods discretely sampling the orientation space, while also being faster. Second, the method of Viguié *et al.* (2013) is validated

predominantly on rather small image volumes (up to approximately 36 million or $500 \times 500 \times 146$ voxels), requiring approximately ten hours of computation in MATLAB, making it impractical for the large synchrotron CT volumes acquired in the present study (35.2 billion or $4\,000 \times 4\,000 \times 2\,100$ voxels).

Therefore, a discretization-free estimate of the local orientation is required. The structure tensor provides continuous, discretization-free directional estimates directly from the image gray values with computational efficiency. It has been widely used in image analysis for detecting orientation and intersection structures. Early formulations were proposed by Bigün and Granlund (1987) for orientation detection and by Harris and Stephens (1988) for corner detection. Later work by Weickert (1999) established the structure tensor as a fundamental tool for analyzing local orientation in images. The eigenvalues and eigenvectors of the structure tensor describe the distribution of intensity gradients within a local neighborhood. Edge-like structures exhibit one dominant eigenvalue, whereas corner or junction regions exhibit multiple significant eigenvalues due to multi-directional variation (Köthe, 2003). Here, we exploit the eigenvalues for detecting locally ambiguous contact regions through the local coherence of the orientation field, rather than relying solely on angular changes in the principal fiber direction.

MATERIAL AND METHODS

We investigate highly porous wood fiber insulation mats from STEICO's "Flex" product line. The raw density of this product is approximately 50 kg/m^3 (Andrä *et al.*, 2023). Two processes are commonly used to manufacture wood fiber insulation materials: the wet process, in which soaked fibers are formed into panels and bound by the wood's own natural lignin, and the dry process, in which dried fibers are blended with a small amount of binder before panel formation (STEICO SE, 2025a). In both processes, fresh coniferous wood is broken down by steam and mechanical treatment (STEICO SE, 2025a). The wood fibers produced by thermo-mechanical refining are highly heterogeneous in geometry: they vary in length from a few hundred micrometers to several millimeters, in cross-sectional shape from nearly circular to highly flattened or collapsed, and also in wall thickness.

According to the manufacturer STEICO SE (2025b), the product composition comprises wood fibers, polyolefin fibers, and ammonium sulphate as a flame retardant. The polyolefin fibers serve as thermally activated binder fibers: during the heating step in production, their outer shells soften and fuse at fiber-to-fiber contact points while the core stays solid, creating a mechanically stable three-dimensional network. In the CT data acquired for this study, the binder fibers appear as solid structures with smooth, approximately circular cross-sections with nearly constant diameter (Fig. 4), distinguishing them geometrically from the hollow, irregularly shaped wood fibers – a distinction that is exploited during image analysis to separate the two fiber populations. These binder fibers are called bico (bi-component) fibers. The resulting mat is a low-density, highly porous yet flexible three-dimensional network.



Fig. 4. *Three-dimensional rendering of a binder fiber extracted from the CT volume. Annotated numbers are the approximate dimensions (length* $\approx 445\,\mu m$*, diameter* $\approx 16\,\mu m$*) of the bico fiber.*

Due to the high porosity of the material, cutting small specimens for CT measurements without destroying the microstructure is extremely difficult. Therefore, a dedicated freeze-cutting method was developed (Andrä *et al.*, 2023). The specimens were first cooled in liquid nitrogen and then immersed in water. The prior cooling prevents the fibers from swelling on contact with water. This freeze-immerse cycle was repeated several times until a solid block of ice encased the fiber network. The frozen block was then cut on a band saw to the required dimensions (5 mm short edge for synchrotron CT, 15 mm short edge for laboratory CT). Immediately after cutting, the specimens were placed in plastic containers to stabilize the structure (Fig. 5). The ice was removed by sublimation in a vacuum chamber while avoiding liquid-phase melting, which would distort the microstructure through capillary forces.

From the cuboidal samples, regions of interest were imaged at beamline ID19 of the European Synchrotron Radiation Facility (ESRF) in Grenoble, France, at a photon energy of 19.5 keV, obtained from the u17.6 undulator (27 mm gap). Projections were acquired by propagation-based phase-contrast imaging using a pco.edge camera and a 10× microscope objective lens (propagation distance: 40 mm; effective pixel size: 0.65 µm; 10 µm GGG:Eu scintillator). Paganin *et al.*'s 2002 single-distance phase-retrieval

approach was applied to the projections, the extended field of view was obtained by stitching projections weighted according to the position of the rotation axis (Haberthür *et al.*, 2010). The scans were obtained from the top, bottom, and center of one of the cuboidal samples representing the full mat thickness as shown in Fig. 5.

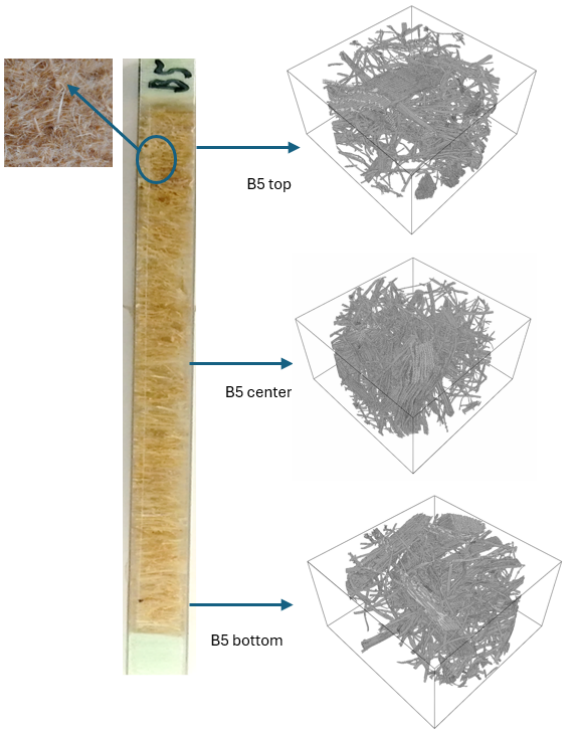


Fig. 5. *Cuboidal specimen (5 mm short edge length) extracted for scanning. On the left, a close-up photograph of the wood fiber insulation mat prior to specimen extraction. On the right of the sample there are three-dimensional renderings of the binarized CT volumes for the top, center, and bottom scans.*

The reconstructed volumes measure approximately $4\,000 \times 4\,000 \times 2\,100$ voxels, corresponding to an edge length of $2.6 \times 2.6 \times 1.3$ mm^3. The voxel size of 0.65 µm was chosen to resolve the thin fiber walls of the hollow wood fibers. At the chosen resolution, the wall thickness spans approximately 3-30 voxels, so that even the thinnest walls are sampled across by at least three voxels. This is essential for preserving connectivity. Otherwise, the partial volume effect reduces gray values along the wall, causing them to drop locally below the segmentation threshold. The large field of view – compared to the voxel size – was required as many fibers exceed 2 mm in length (Andrä *et al.*, 2023). A wider field of view helps to capture more fibers that are not cut by the boundaries. The sample was scanned at the top, center, and bottom to analyze the change in microstructure and alignment of wood fibers at varying positions over the thickness of the mat. In the present work, only the top sub-volume is analyzed. A systematic comparison of microstructure and fiber alignment across the three positions is presented as part of the microstructure optimization workflow (Keilmann *et al.*, 2026).

METHODOLOGY

This subsection outlines the algorithm designed to extract individual fiber structures from the 3D scans of wood fiber insulation mats.

Binary mask for the solid material

The first step is to separate the main of two phases – solid material (wood fibers, wood chunks, and bico fibers) and air in the original gray value image. The solid material attenuates the monochromatic synchrotron radiation much more strongly than air. Therefore, they occupy distinct gray value ranges. The monochromatic beam prevents beam hardening artifacts, hence, each phase has almost the same gray value throughout. Therefore, a single global gray value threshold is sufficient.

Otsu's (1979) method picks the value that best separates the image's gray value histogram into two classes. Every voxel brighter than t^* is labeled solid, and every voxel as dark as t^* or darker is labeled air. This yields the binary mask $\mathscr{B}$, with $\mathscr{B}(x) = 1$ for solid and $\mathscr{B}(x) = 0$ for air as in Fig. 6.

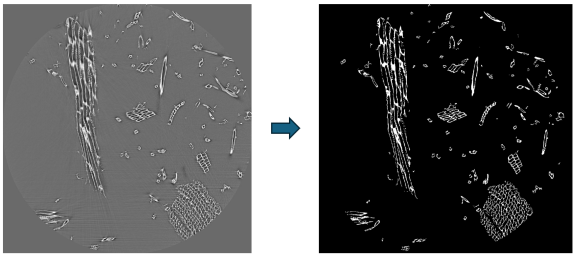


Fig. 6. *Original slice (left) and slice through the binary mask derived by global thresholding (right).*

Bico fiber removal

As mentioned earlier, the "Flex" insulation mats contain bi-component (bico) fibers whose outer shell melts and then glues the wood fibers together, while their core remains stiff and yields the particular flexible behavior of the mats. See Fig. 7 for the rendering of a small cutout of the network formed that way.

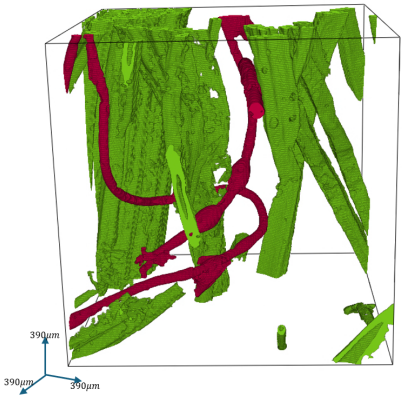


Fig. 7. *Rendering of (parts of) bico fibers (magenta) along with the wood fibers and chunks (green) attached to them.*

In the 3D images, when these bico fibers are in contact with or in proximity to wood fibers, the similarity in thickness of the fiber wall and the bico fibers leads to over-segmentation of the wood fiber walls, splitting individual fibers into multiple fragments. Beyond this segmentation artifact, the bico fibers constitute a distinct material property; retaining them would bias the number-weighted distributions of fiber length, wall thickness, and contact area that form the basis of the subsequent stochastic geometric modeling. For these reasons, the bico fibers are identified and removed before per-fiber analysis.

To separate the two fiber populations, the Frangi filter for multiscale structure enhancement based on the eigenvalues of the Hessian is applied to the gray value CT volume (Frangi *et al.*, 1998; Lindeberg, 1996). This filter was originally designed to enhance tubular structures (vessels). The Frangi index for tubular structures does not directly separate the two fiber populations. Both are locally elongated and have similar diameters. As a result, the corners of the wood fibers are also enhanced as tubular. Therefore, we rather exploit the wall enclosing the lumen being locally plate-like as discriminating feature. We apply the Frangi filter enhancing locally lamellar structure as previously used to detect cracks in tomographic data of C/SiC ceramics (Wirjadi *et al.*, 2014) and concrete (Barisin *et al.*, 2022) and closed facets in ceramic foams (Kampf *et al.*, 2015) to emphasize the thin walls of the hollow wood fibers.

The bico fibers, being solid and approximately cylindrical, are not enhanced. Finally, the detected wood fiber regions are masked with the binary mask $\mathscr{B}$ to obtain a binary mask $\hat{\mathscr{B}}$ corresponding to only the wood fibers, and $\hat{\mathscr{B}}$ subtracted from the binarized material mask $\mathscr{B}$, yields the bico fibers as the residual.

Structure Tensor Computation

To extract local orientation features from the 3D CT volumes, we compute the structure tensor at each voxel following the approach by Weickert (1999) for 2D gray-scale images. In particular, adjusting the *smoothing scale* σ and the *integration scale* ρ allows the algorithm to be sensitive to fine local variations while ensuring stability through spatial smoothing within a neighborhood of the current voxel. The first step is to compute the image gradients after convolving the original intensity image $I(x,y,z)$ with a Gaussian filter with standard deviation σ:

$$\nabla I_\sigma = \begin{pmatrix} I_x \\ I_y \\ I_z \end{pmatrix},$$

where $I_x = \frac{\partial I_\sigma}{\partial x}$, $I_y = \frac{\partial I_\sigma}{\partial y}$, $I_z = \frac{\partial I_\sigma}{\partial z}$ and $I_\sigma = G_\sigma * I$, and G_σ is a Gaussian kernel with a standard deviation of σ. This ensures that the edge detector ignores clusters of noise voxels smaller than σ (Weickert, 1999). Hence, choosing a small σ, we preserve the fiber walls while ignoring small noise voxel agglomerations.

The outer product of the gradient vector forms the unsmoothed structure tensor:

$$J_{\text{raw}} = \nabla I_\sigma (\nabla I_\sigma^T) = \begin{pmatrix} I_x^2 & I_x I_y & I_x I_z \\ I_x I_y & I_y^2 & I_y I_z \\ I_x I_z & I_y I_z & I_z^2 \end{pmatrix}$$

To integrate this information over a neighborhood, each component of the structure tensor is convolved with a Gaussian of standard deviation ρ:

$$J = G_\rho * J_{\text{raw}}$$

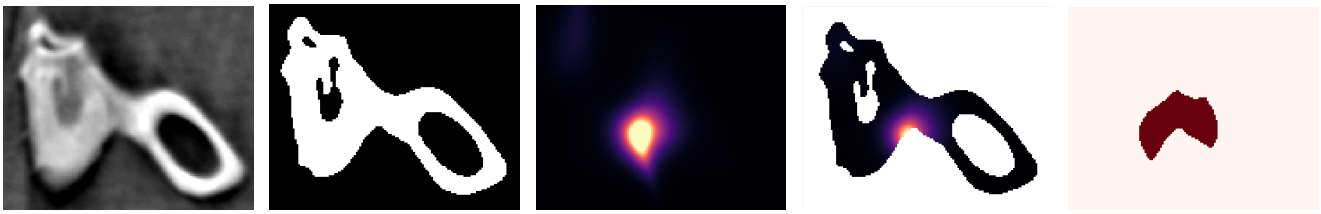
where the convolution is applied component-wise to each entry of the tensor.

This results in a symmetric positive semi-definite 3×3 matrix at each voxel – the *structure tensor* or *second moment matrix*. The eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ and the corresponding eigenvectors of this matrix represent the local orientation variation.

Orientation Coherence for Fiber Contact Detection

In 3D CT images of fiber mats, contacts between fibers appear as regions where local orientation information is strong in different directions (i.e., multiple dominant directions instead of a single one). This can be identified by observing the eigenvalues of the structure tensor. Depending on the local structure, we observe the following relations of the eigenvalues:

– **Inside the fiber walls and on background:** $\lambda_1 \approx \lambda_2 \approx \lambda_3 \approx 0$. Intensity remains homogeneous, the gradient is ≈ 0.



(a) *Raw grayscale* (b) *Binary mask* $\hat{\mathscr{B}}$ (c) *Coherence map* c_h *(magma-color scale)* (d) *Highlighted contact region* (e) *Removed junction voxels*

Fig. 8. *Processing pipeline for contact detection. We show 2D cross-sections of the analyzed cropped part of a 3D CT volume obtained from ESRF. The color-mapped coherence* $c_h = (\lambda_2\lambda_3/(\lambda_1^2+\varepsilon))(\lambda_1+\lambda_2+\lambda_3)$ *highlights fiber contact zones.*



(a) *Binary mask* $\hat{\mathscr{B}}$ (b) *Contact region* $\mathscr{J}$ (c) *Labeled image* $\mathscr{L}$ (d) *After recovery* $\hat{\mathscr{L}}$

Fig. 9. *Post-processing pipeline for contact recovery. Labels of connected components are shown using different colors.*

– **Fiber surface:** $\lambda_1 \gg \lambda_2 \approx \lambda_3 \approx 0$. One dominant direction indicates a locally planar surface.
– **Contacts or junctions:** $\lambda_1 \approx \lambda_2 \approx \lambda_3 \gg 0$. All directions exhibit similar variation, consistent with fiber intersections or background noise.

So we define a coherence term c_h to locate regions of high intensity change in different directions based on the eigenvalues of the structure tensor:

$$c_h = \frac{\lambda_2\lambda_3}{(\lambda_1^2+\varepsilon)}(\lambda_1+\lambda_2+\lambda_3),\ \lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0,\ \varepsilon \geq 0, \tag{1}$$

where ε is a small constant preventing numerical instability. Here, we use $\varepsilon = 10^{-7}$. The term $\lambda_2\lambda_3/(\lambda_1^2+\varepsilon)$ is close to 0 when one eigenvalue dominates, as on the fiber surface. It approaches 1 when all three eigenvalues are comparable, as at fiber intersections and in homogeneous regions (background, fiber interior). To distinguish fiber junctions from homogeneous regions, the ratio is multiplied by the trace $(\lambda_1+\lambda_2+\lambda_3)$. It is large only where the gradient is strong in all directions (fiber junctions) and small in homogeneous regions. Local maxima of the coherence term c_h thus indicate fiber junctions. The coherence term c_h for each voxel is stored in an image of the same size as the input data – the coherence map. Fiber contact regions are identified by thresholding this image:

$$\text{Contact voxel} \quad \rightarrow \quad c_h > \tau,$$

where τ is chosen locally according to the algorithm by Sauvola and Pietikäinen (2000). Fig. 8 shows an example. For details on parameter choice, see the Results section.

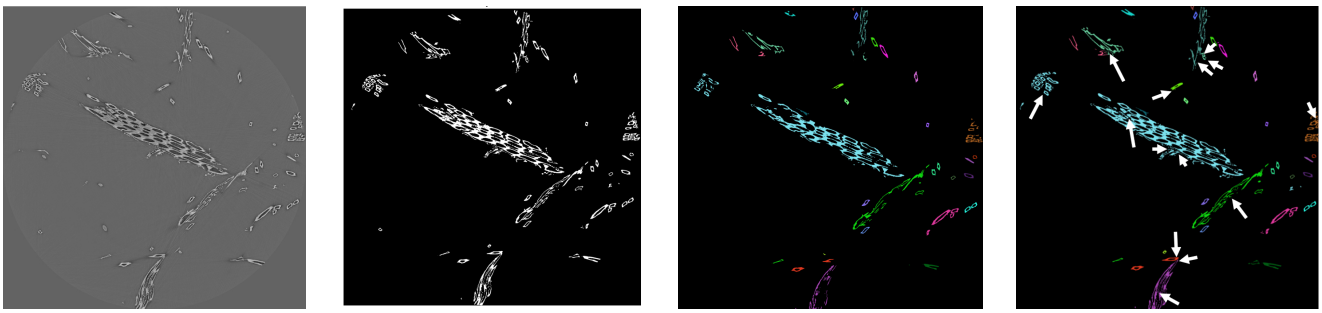
Recovering removed junctions

After fiber contacts are removed via thresholding, we obtain separated fibers, which we label using 26-connectivity. The contact regions are, however, lost at this point. Nevertheless, the ultimate goal of segmenting the CT image is to analyze the thermal conductivity, and the fiber-fiber heat conductivity occurs primarily through these fiber contacts (Andrä *et al.*, 2023). It is therefore crucial to recover the contact surface. To recover these lost voxels and assign them to appropriate fiber labels, a constrained label propagation is applied.

Let $\mathscr{L}$ be the labeled image after the junctions are removed with M connected components $\mathscr{L}_1, \mathscr{L}_2, \ldots, \mathscr{L}_M$, and $\mathscr{J}$ the binary image of the identified junctions. The removed junction voxels form 26-connected components $\mathscr{J}_1, \mathscr{J}_2, \ldots, \mathscr{J}_N$. For each "junction" $\mathscr{J}_i$, the set of neighboring fiber labels is determined by dilating $\mathscr{J}_i$ by a structuring element of radius r (default r=1) and intersecting with $\mathscr{L}$:

$$\mathscr{N}(\mathscr{J}_i) = \{\ell \in \{1,\ldots,M\} \mid (\mathscr{J}_i \oplus S_r) \cap \mathscr{L}_\ell \neq \emptyset\}$$

where $J_i \oplus S_r$ denotes the dilation of J_i by structuring element S_r. The junctions $\mathscr{J}_i$ are then



(a) *Grey-value CT reconstruction.* (b) *Binarized wood-fiber mask* $\hat{\mathscr{B}}$. (c) *Labeled fiber instances* $\mathscr{L}$. (d) *After junction recovery* $\hat{\mathscr{L}}$.

Fig. 10. *Instance segmentation across the processing stages on a representative slice of dimension* $1\,309 \times 1\,309$ *constituting an area of* $2.57 mm^2$. *(a) Gray-value CT reconstruction. (b) Binarized wood-fiber mask, in which contacting fibers form a connected network. (c) Labeled fiber instances after junction removal (176 objects over the full volume). (d) Labels after EDT-based junction recovery. White arrows point to particularly affected and now recovered foreground regions.*

treated based on the cardinality $|\mathscr{N}(J_i)|$:

- **Case 1:** $|\mathscr{N}(\mathscr{J}_i)| = 1$**.** The component $\mathscr{J}_i$ is adjacent to a single fiber label ℓ. All voxels in $\mathscr{J}_i$ are assigned to label ℓ, indicating a likely false positive detection by the coherence-based thresholding.
- **Case 2:** $|\mathscr{N}(\mathscr{J}_i)| \geq 2$**.** The component $\mathscr{J}_i$ is adjacent to two or more fiber labels, corresponding to a multi-fiber junction, and its voxels are split between the neighboring fibers by competitive growing. For each neighboring fiber $\mathscr{L}_\ell$, $\ell \in \mathscr{N}(\mathscr{J}_i)$, a binary image is formed in which $\mathscr{L}_\ell$ is set to one and all other voxel values to zero. Subsequently the Euclidean distance transform is computed on the background. This yields for each neighboring fiber, the map of the distance to this fiber. Each neighboring fiber is then grown into $\mathscr{J}_i$ one voxel layer at a time by dilation, restricted at each step to the still-unassigned voxels. In each step, all fibers are dilated, and every voxel of $\mathscr{J}_i$ is assigned to the fiber whose front reaches it first. Where two fronts reach a voxel in the same step, it is assigned to the nearer fiber according to the distance maps. The growth continues until no unassigned voxel remains.
- **Case 3:** $|\mathscr{N}(\mathscr{J}_i)| = 0$**.** The component $\mathscr{J}_i$ has no neighboring fiber labels in L. Such isolated components are likely to arise from image noise or small imaging artifacts that were incorrectly retained in the binary mask $\hat{\mathscr{B}}$ rather than representing true fiber material. Components satisfying this condition are therefore discarded and their voxels assigned to the background.

The recovered labels are merged back into $\mathscr{L}$, yielding the final instance segmentation $\hat{\mathscr{L}}$, in which every voxel of the binary mask $\hat{\mathscr{B}}$ carries a fiber label and all contact zone voxels are explicitly assigned to one of the adjacent fiber instances. See Fig. 9 for an illustration of the recovery of the junction removed in Fig. 8.

TWO-SCALE SEGMENTATION WORKFLOW

The pipeline was applied to the image of the top of the sample described in the Materials and Data section. The 3D image has approximately $4\,000 \times 4\,000 \times 2\,100$ voxels, totaling roughly 130 GB. Since tuning parameters for large volume data is computationally expensive and time-consuming, parameters were first established on a downsampled version of the data before being applied at full resolution. Data were downsampled by a factor of three, assigning each output voxel the mean gray value of the corresponding $3 \times 3 \times 3$ block of the input image (resulting in a voxel size of 1.95 µm).

Segmentation of downscaled image

The bico fibers were separated from the wood fibers as residuals of the locally flat wood-fiber walls emphasized by the Frangi flatness filter, as implemented in MAVIkit (2025) with scale range $\sigma \in [0.5, 7.0]$ (step 1.0) and parameters $\alpha = 0.5$, $\beta = 2.5$, $\gamma = 1.0$. The result is globally thresholded to obtain fiber walls, followed by a morphological closing with a cubic structuring element of size 4 voxels to recover thick fiber walls, and finally masked with the binary image. This yields the wood fibers, only. The removed bico component is shown in Fig. 11.

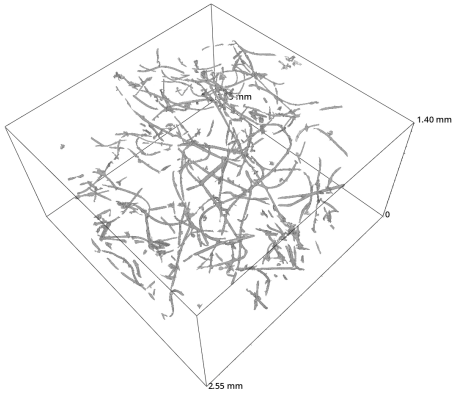
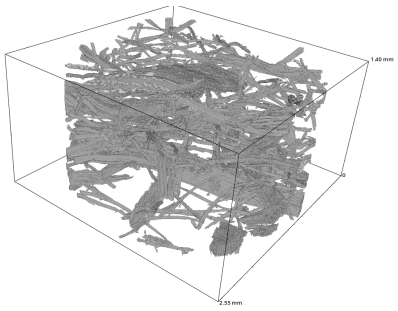


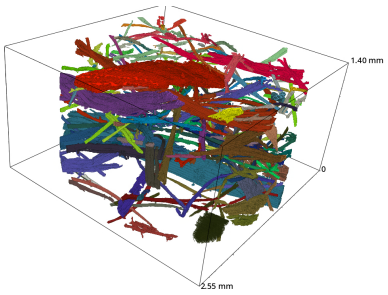
Fig. 11. *Bico component detected in the downsampled (*$1\,309 \times 1\,309 \times 720$*) voxel image.*

Following the bico removal from the binary image, the wood fibers are separated as described above. See Fig. 10 for slice views of the essential steps of the segmentation algorithm. The structure tensor was computed on the gray value CT volume using MAVIkit (2025). The pre-smoothing scale σ was chosen as 0.75, and the integration scale ρ was chosen as 4.5 with a convolution precision of 3, i.e., the kernel is truncated at 3 times the standard deviation. The characteristic fiber-wall thickness ranges from approximately 1 to 10 voxels at this resolution. Hence, the integration scale $\rho = 4.5$ was chosen as the center of this interval. The eigenvalue decomposition in each voxel yielded the ordered eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ at each voxel. Overall, the computation of the coherence map (1) on the downsampled image takes around one hour on a compute node with two AMD EPYC 7742 processors (128 cores in total) and 1 TiB of main memory.

Junctions were identified by locally adaptively thresholding the coherence map (1) with window size 5 and coefficient 0.2 using Sauvola's local thresholding. The flagged junction voxels were subtracted from the bico-removed binary fiber image, separating the connected network into individual fibers. These individual fiber objects were labeled using 26-connectivity (N26), assigning a distinct label to each fiber instance as in Fig. 10c. The lost voxels are recovered using distance based recovery as described above and shown in Fig. 10d. Fig. 12 additional provides 3D renderings of the wood component and the corresponding segmentation result.



(a) *Rendering of the wood component.*



(b) *Segmentation results including recovered junctions for the same image.*

Fig. 12. *Segmentation result for the downsampled sample of dimension* $1\,309 \times 1\,309 \times 720$*. Individual fibers are color-coded.*

Segmentation of full-resolution image

The tuned pipeline was subsequently applied to the full-resolution volume (0.65 µm voxel size). Since rerunning the Frangi filter is computationally expensive on the full-resolution image of size around 130 GB, the bico fiber mask obtained at 1.96 µm was upscaled using the nearest-neighbor method with an upsampling factor of 3 (Gonzalez and Woods, 2010). This was followed by a dilation with a spherical structuring element of radius 2 to compensate for interpolation and recover lost boundaries, and was masked with the binary image. The structure tensor was then computed on the full-resolution gray-scale image with $\sigma = 2$, convolution precision 3, and integration scale $\rho = 13$. The computation of the coherence map (1) of the full volume took about 31.5 hours. The rest of the steps were applied as described for the downsampled data.

RESULTS

After segmentation and voxel recovery, we compute the intrinsic volumes and related features (Ohser and Schladitz, 2009), and geodesic features and features of the minimal volume bounding cuboid for each image object both at full resolution and in the downsampled image using MAVI (2025). All these

features are implemented as described by Vecchio *et al.* (2012) with the minimal volume bounding box using the algorithm of Barequet and Har-Peled (2001) and the geodesic features based on the fast marching approach of Pêtrès *et al.* (2005). Here, we report the geodesic length. That is, the longest shortest discrete path connecting two points of the object and being completely contained inside the object.

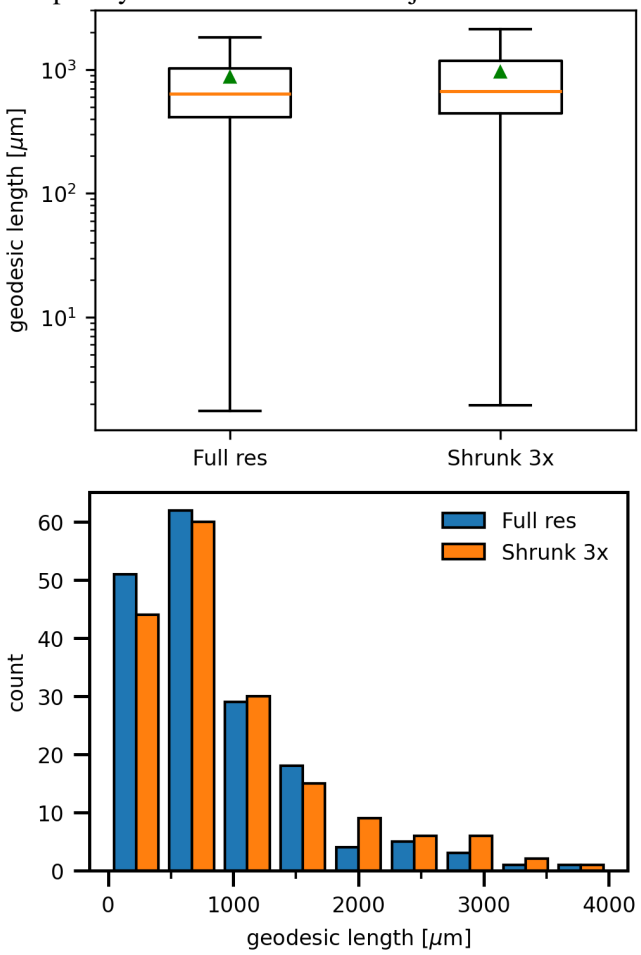


Fig. 13. *Comparison of geodesic fiber lengths derived from the full resolution image (0.65 µm) and the downsampled one (1.95 µm). Top: Boxplot plotted on a logarithmic axis. The box marks the middle 50 % of the fiber population. The orange line is the median and the green triangle is the mean. Bottom: Histograms of the geodesic lengths.*

Fig. 13 shows the geodesic length distributions for the full resolution and the downsampled one. The wide spread and logarithmic axis reflect the heterogeneous distribution of fiber length. The median geodesic length increases slightly under downsampling (640.9 µm to 666.6 µm, an increase by 4.0 %). This can be explained by fibers merging at lower resolution, which connects adjacent fibers into longer paths.

Fig. 14 compares volume, surface area, and the maximal Feret diameter between the full resolution and the downsampled image. All three features closely agree across both resolutions.

Fig. 15 shows the three aspect ratios of the minimal volume bounding cuboid. All three are preserved under downsampling. Table 1 summarizes the median of each feature at both resolutions as well as the p-values of two-sample Kolmogorov-Smirnov tests for equality of distributions.

Fiber-fiber contacts were compared too between resolutions. At full resolution, 121 fibers were in contact, corresponding to 113 fiber-fiber contacts. After downsampling, these rose to 125 fibers and 140 contacts, confirming that contacts are preserved rather than being lost under downsampling. Fig. 16 shows per fiber contact distributions at both resolutions. We observe that most fibers contact one or two neighbors at both resolutions, and the histograms overlap closely with a slight shift towards higher counts for the downsampled image. The number of contacts per fiber did not differ significantly between resolutions (permutation-based Kolmogorov-Smirnov test, $D = 0.12$, $p = 0.077$), whereas the total contact area per fiber was significantly lower in the downsampled volume ($D = 0.27$, $p < 10^{-3}$), with the median falling from 9.5×10^{-3} to 4.6×10^{-3} mm^2.

Therefore, we conclude that the downsampled SR-µCT image is adequate for separating individual fibers and wood chunks, analyzing microstructural object features, and counting the total fiber-fiber contacts. However, for the contact area and probably other geometric features of the contact regions, correct measurements can only be expected from the full resolution data.

DISCUSSION

We present a pipeline for 3D instance segmentation of wood fibers in highly porous insulation mats, imaged by SR-subµCT. The fibers are bound together by bico fibers and touch and cross throughout the network, so connected component labelling merges many fibers into a few large objects. Therefore, as a pre-processing step, we remove the bico fibers exploiting local shape and a Frangi filter to enhance it. Subsequently, we separate the fiber instances by removing the junctions using an orientation coherence map derived from the structure tensor eigenvalues. For the (sub-)sample considered in this paper, we observe the number of separate fiber instances to rise from 32 to 176. The final recovery step reassigns removed junction voxels to the closest fiber, so the separation does not come at the cost of losing fiber material.

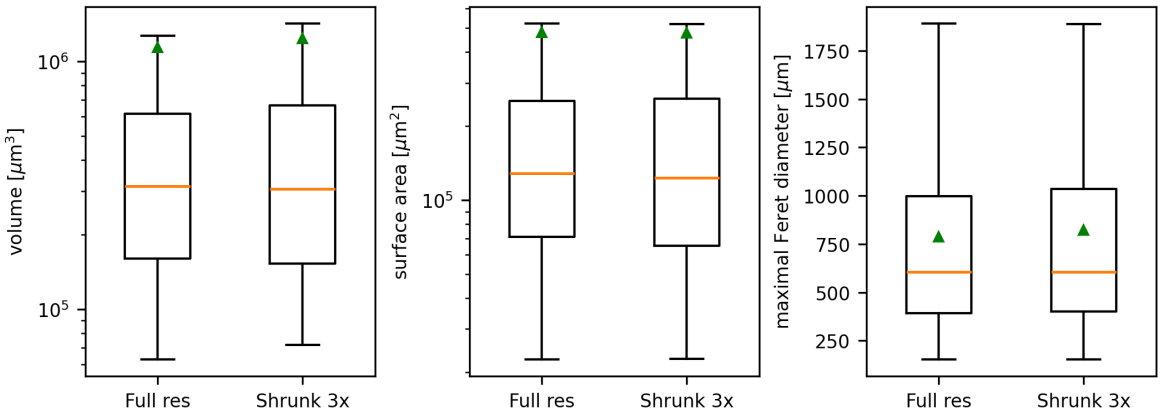


Fig. 14. *The morphological features stay almost unchanged when the resolution is reduced. Each subfigure compares the fiber values at full resolution (voxel size 0.65 µm) and after downsampling the image by factor 3 (voxel size 1.95 µm).*

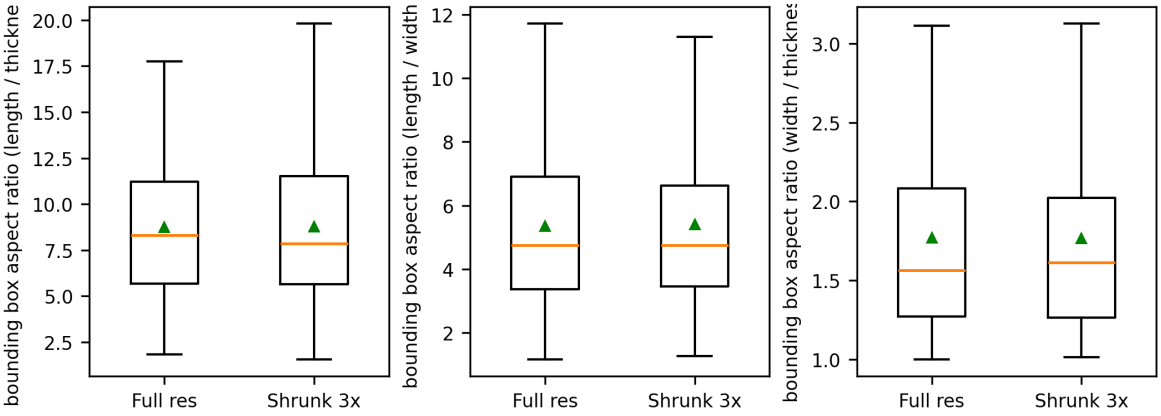
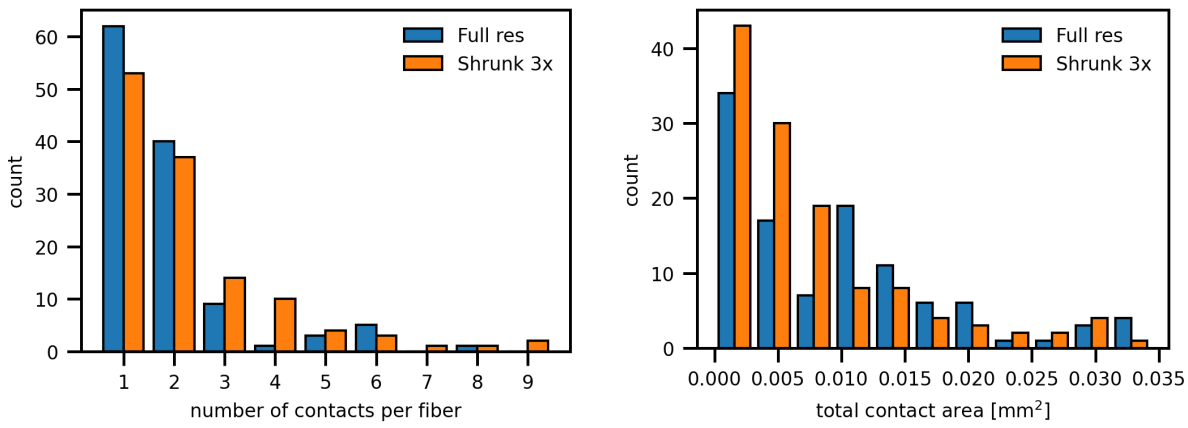


Fig. 15. *Bounding box aspect ratios stay almost unchanged when the resolution is reduced. Each subfigure compares the fiber values at full resolution (voxel size 0.65 µm) and after downsampling the image by factor 3 (voxel size 1.95 µm).*



(a) *Number of contacts per fiber* (b) *Fiber contact area*

Fig. 16. *Histogram of per fiber contact for full resolution (0.65 µm) and after 3 times downsampling (1.95 µm). (a) Number of neighbors in contact per fiber. (b) Total contact area per fiber.*

Table 1. *Median of each feature at full resolution (0.65 µm) and after downsampling (1.95 µm), with the relative difference and the two-sample Kolmogorov–Smirnov test p-value.*

Feature	Unit	Full resolution	Downsampled 3×	Δ (%)	p-value
Volume	mm^3	3.145×10^{-4}	3.128×10^{-4}	−0.5	0.81
Volume-equivalent spherical diameter	mm	8.437×10^{-2}	8.422×10^{-2}	−0.2	0.81
Surface area	mm^2	1.295×10^{-1}	1.237×10^{-1}	−4.4	0.99
Maximal Feret diameter	mm	6.058×10^{-1}	6.056×10^{-1}	−0.0	0.99
Bounding-box length	mm	6.037×10^{-1}	6.043×10^{-1}	+0.1	0.99
Bounding-box width	mm	1.177×10^{-1}	1.202×10^{-1}	+2.2	0.88
Bounding-box thickness	mm	7.326×10^{-2}	7.757×10^{-2}	+5.9	0.55
Geodesic length	mm	6.409×10^{-1}	6.666×10^{-1}	+4.0	0.47
Length / width	–	4.752	4.745	−0.1	0.99
Length / thickness	–	8.318	7.865	−5.4	0.88
Width / thickness	–	1.565	1.613	+3.1	0.72

This pipeline is tailor-made for wood fibers. Specifically, we can separate fibers that are hollow and thin-walled, with partially collapsed lumen and torn walls resulting from the manufacturing process. This breaks the assumptions behind the most commonly used fiber segmentation methods. From the structure tensor, we derive a local orientation coherence criterion for junction detection, that does not depend on the fiber being tubular, which is why it applies where these methods do not. Our method enables separation of the fibers in a $4\,000 \times 4\,000 \times 2\,100$ voxel image and subsequent analysis of their geometric characteristics as well as analysis of the fiber-fiber contact regions. Note that the separation thrives heavily on the strong edge contrast thanks to the phase retrieval and the low noise, both enabled by the use of synchrotron radiation, only.

Furthermore, we compared geometric characteristics of the image objects segmented at the full resolution (voxel size 0.65 µm) and after downsampling by a factor of 3 (voxel size 1.95 µm). Most fiber characteristics agree within a few percentage points (6 %). The largest changes occur in bounding-box thickness (+5.9%) and, due to that, for the length-to-thickness ratio (−5.4%). Contact behavior separates the two resolutions more clearly: the number of contacts per fiber does not differ significantly between resolutions, whereas the total contact area per fiber is significantly smaller after downsampling. Downsampling the SR-sub-µCT images is therefore adequate for object morphology and for counting contacts, but contact area should be quantified at full resolution. Thus, if size and shape of the contact regions are crucial for the target property like e. g. mechanical properties, sub-µ resolution is recommended.

A limitation of the junction-based separation is that it distinguishes fibers where they cross but not where they run parallel and in contact over an extended area. The coherence map does not reveal such configurations. However, for the wood fiber insulation mats, this is actually an asset: Parallel fibers essentially occur as bundles or wood chunks as featured in Fig. 3. That is, fibers are bonded along their length and form a single connected structure. Our pipeline segments these chunks correctly as single objects. A further limitation is that junction detection also depends on a threshold applied to the coherence map. This threshold is selected by Sauvola's locally adaptive method with manually tuned parameters. The separation is sensitive to this choice, and the step is not yet automated.

To summarize, we suggest a new separation method for fibers with very irregular cross-sectional shape and prove its usability for very large 3D image data. Based on the derived segmentation, we compared geometric characteristics of fibers and wood chunks in SR-sub-µCT images at the original resolution and after downsampling. According to our findings, downsampling by a factor of three by averaging is feasible except if one is explicitly interested in the geometry of the fiber-fiber contacts.

ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry for Research, Technology, and Space (BMFTR) under the focus programme "Mathematics for innovations", Grant Agreement No: 05M2022. We thank Michael Makas from STEICO SE for providing the material, Max Engelhardt and Sebastian Treml from Forschungsinstitut für Wärmeschutz e.V. München for sample preparation, Alexander Rack and Ludovic Broche for exceptional support during measurements at ID19 of ESRF, and Alex Keilmann from RPTU for help in data acquisition at ESRF and verifying the quality of segmentations.

DATA AVAILABILITY STATEMENT

The data will be available on demand and will be published in the future after being accordingly prepared and embedded in the sketched microstructure optimization.

REFERENCES

Andrä H, Dobrovolskij D, Engelhardt M, Godehardt M, Makas M, Mercier C, Rief S, Schladitz K, Staub S, Trawka K, Treml S (2023). Image-based microstructural simulation of thermal conductivity for highly porous wood fiber insulation boards: 3D imaging, microstructure modeling, and numerical simulations for insight into structure–property relation. Wood Sci Technol 57:5–31.

Barequet G, Har-Peled S (2001). Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. J Algorithms 38:91–109.

Barisin T, Jung C, Müsebeck F, Redenbach C, Schladitz K (2022). Methods for segmenting cracks in 3D images of concrete: A comparison based on semi-synthetic images. Pattern Recogn 129:108747.

Beucher S, Meyer F (1993). The morphological approach to segmentation: The watershed transformation. In: Dougherty ER, ed., Mathematical Morphology in Image Processing. New York: Marcel Dekker, 433–81.

Bigün J, Granlund GH (1987). Optimal orientation detection of linear symmetry. In: Proceedings of the First International Conference on Computer Vision. London: IEEE Computer Society Press, 433–8.

Blum H (1967). A transformation for extracting new descriptors of shape. In: Wathen-Dunn W, ed. Models for the Perception of Speech and Visual Form. Cambridge, MA: MIT Press, 362–80.

Creveling PJ, Whitacre WW, Czabaj MW (2019). A fiber-segmentation algorithm for composites imaged using X-ray microtomography: Development and validation. Compos Part A-Appl S 126:105606.

Czabaj MW, Riccio ML, Whitacre WW (2014). Numerical reconstruction of graphite/epoxy composite microstructure based on sub-micron resolution X-ray computed tomography. Compos Sci Technol 105:174–82.

Emerson MJ, Jespersen KM, Dahl AB, Conradsen K, Mikkelsen LP (2017). Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials. Compos Part A-Appl S 97:83–92.

Frangi AF, Niessen WJ, Vincken KL, Viergever MA (1998). Multiscale vessel enhancement filtering. In: Wells WM, Colchester A, Delp S, eds. Medical Image Computing and Computer-Assisted Intervention (MICCAI 1998), vol. 1496 of Lecture Notes in Computer Science. Berlin: Springer, 130–7.

Gaiselmann G, Manke I, Lehnert W, Schmidt V (2013). Extraction of curved fibers from 3D data. Image Anal Stereol 32:57–63.

Gonzalez RC, Woods RE (2010). Digital Image Processing. 3rd ed. Upper Saddle River, NJ: Pearson Prentice Hall.

Haberthür D, Hintermüller C, Marone F, Schittny JC, Stampanoni M (2010). Radiation dose optimized lateral expansion of the field of view in synchrotron radiation X-ray tomographic microscopy. J Synchrotron Radiat 17:590–9.

Harris C, Stephens M (1988). A combined corner and edge detector. In: Proceedings of the 4th Alvey Vision Conference. Manchester: Alvey Vision Club, 147–51.

Kampf J, Schlachter AL, Redenbach C, Liebscher A (2015). Segmentation, statistical analysis, and modelling of the wall system in ceramic foams. Mater Charact 99:38–46.

Keilmann A, Khatab M, Leithäuser C, Makas M, Nair A, Redenbach C, Rief S, Schladitz K, Totzeck C, Treml S (2026). Microstructure modelling and optimization of highly porous wood fiber insulation panels. In preparation.

Kibleur P, Boone MN, Van Acker J, Aelterman J (2022). Deep learning segmentation of wood fiber bundles in fiberboards. Compos Sci Technol 221:109287.

Köthe U (2003). Edge and junction detection with an improved structure tensor. In: Michaelis B, Krell G, eds. Pattern Recognition (DAGM 2003), vol. 2781 of Lecture Notes in Computer Science. Berlin: Springer, 25–32.

Latil P, Orgéas L, Geindreau C, Dumont PJJ, Rolland du Roscoat S (2011). Towards the 3D in situ characterisation of deformation micromechanisms within a compressed bundle of fibres. Compos Sci Technol 71:480–8.

Lee TC, Kashyap RL, Chu CN (1994). Building skeleton models via 3-D medial surface axis thinning algorithms. CVGIP-Graph Model Im 56:462–78.

Lindeberg T (1996). Edge detection and ridge detection with automatic scale selection. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 1996). San Francisco: IEEE Computer Society, 465–70.

Lux J (2013). Automatic segmentation and structural characterization of low density fiberboards. Image Anal Stereol 32:13–25.

MAVI (2025). MAVI – modular algorithms for volume images, version 2025, by fraunhofer itwm, 3d image analysis. https://www.mavi-3d.de.

MAVIkit (2025). MAVI toolbox for ToolIP, version 2025, by Fraunhofer ITWM, 3D Image Analysis. http://www.itwm.fraunhofer.de/toolip.

Ohser J, Schladitz K (2009). 3D Images of Materials Structures – Processing and Analysis. Weinheim: Wiley-VCH.

Otsu N (1979). A threshold selection method from gray-level histograms. IEEE T Syst Man Cyb 9:62–6.

Paganin D, Mayo SC, Gureyev TE, Miller PR, Wilkins SW (2002). Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object. J Microsc-Oxford 206:33–40.

Palágyi K, Kuba A (1999). A parallel 3D 12-subiteration thinning algorithm. Graph Model Im Proc 61:199–221.

Pêtrès C, Pailhas Y, Petillot Y, Lane D (2005). Underwater path planing using fast marching algorithms. In: Europe Oceans 2005, vol. 2. Brest: IEEE, 814–9.

Pinter P, Bertram B, Weidenmann KA (2016). A novel method for the determination of fibre length distributions from µCT-data. In: Proceedings of the 6th Conference on Industrial Computed Tomography (iCT 2016). Wels: NDT.net.

Requena G, Fiedler G, Seiser B, Degischer P, Di Michiel M, Buslaps T (2009). 3D-quantification of the distribution of continuous fibres in unidirectionally reinforced composites. Compos Part A-Appl S 40:152–63.

Sauvola J, Pietikäinen M (2000). Adaptive document image binarization. Pattern Recogn 33:225–36.

Sencu RM, Yang Z, Wang YC, Withers PJ, Rau C, Parson A, Soutis C (2016). Generation of micro-scale finite element models from synchrotron X-ray CT images for multidirectional carbon fibre reinforced composites. Compos Part A-Appl S 91:85–95.

Sharma Y, Phillion AB, Martinez DM (2015). Automated segmentation of wood fibres in micro-CT images of paper. J Microsc-Oxford 260:400–10.

STEICO SE (2025a).
STEICO insulation materials.
https://www.steico.com/uk/solutions/product-advantages/steico-insulation-materials.
Accessed: 2026-05-20.

STEICO SE (2025b).
Technical data sheet: STEICOflex – flexible thermal insulation from wood fibre.
https://www.steico.com/fileadmin/user_upload/Public_Docs/EUR_en/EUR_en_Technical_Data_Sheets/1._Product_information_wood_fibre_insulation/STEICO_technical-data-sheet_flex-038_EUR_en_i.pdf. Accessed: 2026-05-20.

Vecchio I, Schladitz K, Godehardt M, Heneka M (2012). 3D geometric characterization of particles applied to technical cleanliness. Image Anal Stereol 31:163–74.

Viguié J, Latil P, Orgéas L, Dumont PJJ, Rolland du Roscoat S, Bloch JF, Marulier C, Guiraud O (2013). Finding fibres and their contacts within 3D images of disordered fibrous media. Compos Sci Technol 89:202–10.

Vincent L, Soille P (1991). Watersheds in digital spaces: An efficient algorithm based on immersion simulations. IEEE T Pattern Anal 13:583–98.

Weickert J (1999). Coherence-enhancing diffusion filtering. Int J Comput Vision 31:111–27.

Wernersson ELG, Brun A, Luengo Hendriks CL (2009). Segmentation of wood fibres in 3D CT images using graph cuts. In: Foggia P, Sansone C, Vento M, eds. Image Analysis and Processing – ICIAP 2009, vol. 5716 of Lecture Notes in Computer Science. Berlin: Springer, 92–102.

Wirjadi O, Rack A, Liebscher A, Meinhardt J, Schladitz K, Shafei B, Steidl G (2014). Microstructural analysis of a C/SiC ceramic based on the segmentation of X-ray phase contrast tomographic data. Int J Mater Res 105:702–8.

Wirjadi O, Schladitz K, Easwaran P, Ohser J (2016). Estimating fibre direction distributions of reinforced composites from tomographic images. Image Anal Stereol 35:167–79.

Zauner H, Salaberger D, Heinzl C, Kastner J (2015). 3D image processing for single fibre characterization by means of XCT. In: Proceedings of the 14th International Congress for Stereology and Image Analysis. Liège: Acta Stereologica.

Zhou Y, Yu H, Simmons J, Przybyla CP, Wang S (2016). Large-scale fiber tracking through sparsely sampled image sequences of composite materials. IEEE T Image Process 25:4932–41.