J Mater Sci (2020) 55:16273–16289

Computation & theory

Automated segmentation of computed tomography images of fiber-reinforced composites by deep learning

Aly Badran[1,*], David Marshall[1], Zacharie Legault[2], Ruslana Makovetsky[3], Benjamin Provencher[3], Nicolas Piché[3], and Mike Marsh[3]

[1] *Department of Aerospace Engineering and Sciences, University of Colorado, Boulder, 3775 Discovery Dr. Boulder, Boulder, CO 80303, USA*
[2] *Department of Computer and Software Engineering, Polytechnique Montreal, Montreal, Canada*
[3] *Object Research Systems, Montreal, Canada*

Received: 17 April 2020
Accepted: 21 August 2020
Published online:
8 September 2020

ABSTRACT

A deep learning procedure has been examined for automatic segmentation of 3D tomography images from fiber-reinforced ceramic composites consisting of fibers and matrix of the same material (SiC), and thus identical image intensities. The analysis uses a neural network to distinguish phases from shape and edge information rather than intensity differences. It was used successfully to segment phases in a unidirectional composite that also had a coating with similar image intensity. It was also used to segment matrix cracks generated during in situ tensile loading of the composite and thereby demonstrate the influence of nonuniform fiber distribution on the nature of matrix cracking. By avoiding the need for manual segmentation of thousands of image slices, the procedure overcomes a major impediment to the extraction of quantitative information from such images. The analysis was performed using recently developed software that provides a general framework for executing both training and inference.

Handling Editor: Avinash Dongare.

Address correspondence to E-mail: aly.badran@colorado.edu

https://doi.org/10.1007/s10853-020-05148-7

GRAPHIC ABSTRACT

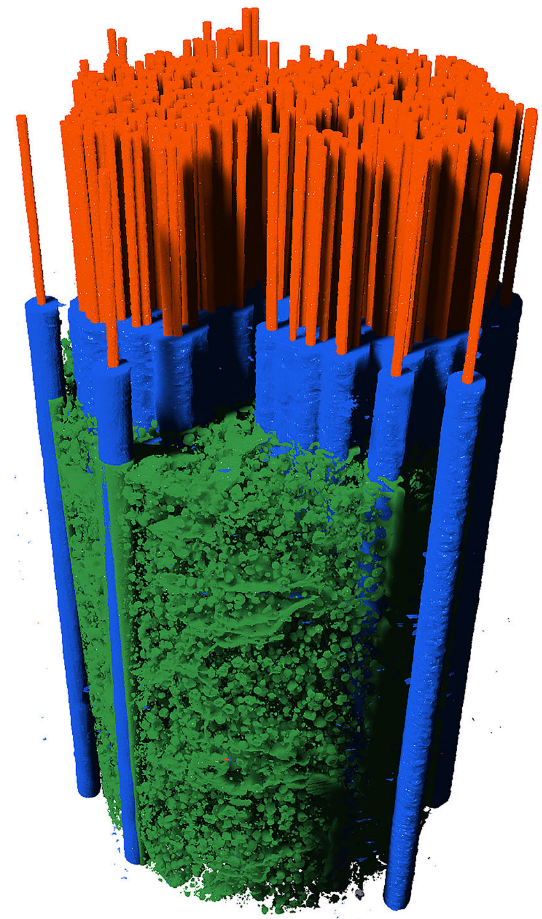

Introduction

X-ray micro-CT imaging has become a valuable tool for analysis of structural materials, both for visualizing complex 3D microstructures and for imaging internal defects and damage introduced during manufacture or service [1–10].

In the case of fiber-reinforced composites (polymer or ceramic matrices), useful mechanical properties are achieved by arranging high strength fibers in sometimes complex 3D arrangements and surrounding them with a matrix [11]. The performance and reliability of these composites are dependent on the internal fiber architecture and the nature of damage caused by external loads. During the past 10 years, synchrotron micro-CT imaging has been used to measure internal fiber architectures, including statistical deviations from the ideal (or intended) architecture, and provide the input needed to build numerical models for life prediction in both ceramic and polymer matrix composites [1, 4–7]. It has also been used with in situ mechanical loading and/or heating to detect and image the initiation and development of internal damage [2–6, 12], a critical step in guiding and validating life prediction and processing models.

In most studies, the identification and representation of image features of interest has been done manually. The large amount of data involved makes this a daunting and time-consuming task, which is a major impediment to the extraction of quantitative information from the images (a single experiment typically produces multiple 3D images, each consisting of more than 2000 slices). Automation of the process would enable more widespread and efficient use of such data.

The first step in automating the process involves image segmentation, the labeling of image pixels according to their constituent materials. Having all of the pixels labeled permits use of standard image processing routines to compute quantitative parameters, such as fiber volume fractions and fiber orientations, as well as visualize the spatial distribution of internal cracks and voids. Substantial progress has been made in developing and evaluating the accuracy of various automated algorithms for segmentation. The approaches range from simple thresholding to more sophisticated techniques involving k-means clustering, statistical region merging, and parallel Markov random fields [13, 14], which are able to avoid some of the errors invariably introduced by noise and reconstruction artifacts in real images. In a recent study, Czabaj et al. [15] used a synthetic template-matching technique combined with multi-target tracking to determine accurate locations of individual fibers in a graphite/epoxy composite with multidirectional reinforcing fibers, with very few errors. However, all of these approaches depend to some extent on intensity differences between the phases being separated. Consequently, they become less effective in composite materials with low contrast between fibers and matrix, or more complex microstructures (especially with multiple phases in the matrix) and intricate internal pores.

Ceramic composites consisting of SiC fibers in a SiC matrix have been developed for high-temperature applications in turbine engines, aerospace propulsion and nuclear power generation. These composites present an extreme challenge for automated image segmentation, since the fibers and matrix show no gray-level contrast. An example is shown in Fig. 1. This composite was made by taking a single straight tow of SiC fibers ($\sim$ 500 filaments), coating the fibers with a thin layer of BN ($\sim$ 500 nm thickness), then forming the matrix by chemical vapor infiltration of SiC. The BN layer is visible in Fig. 1a as a dark ring around each fiber. An environmental barrier coating (EBC) of alumina and silica was then applied to the exterior of the rod-shaped composite using a slurry deposition process. The CVI process used to form the SiC matrix invariably traps voids among the fibers, which are visible in Fig. 1a. Several regions are also visible where the EBC had infiltrated into voids that were connected to the surface of the composite. Despite the fact that there is no difference in gray levels between the fibers, matrix, and EBC, these three phases are easily discriminated visually in the section normal to the fibers (Fig. 1a). However, in the section parallel to the fibers (Fig. 1b), the fibers and matrix cannot be distinguished.

In this paper, we examine the application of an image segmentation method based on deep learning that is capable of segmenting images such as Fig. 1 automatically. This deep learning method relies on artificial neural networks to perceive the shape and edge information cues that enable visual discrimination in images such as Fig. 1a. Several other studies have explored application of deep learning to image segmentation in other fields. Haberl et al. [16] built a cloud-based tool (CDeep3M), which they benchmarked with images of biomedical systems obtained with several techniques. Sinchuk et al. evaluated both variational[1] and deep-learning-based segmentation approaches using images with low resolution and low (but nonzero) contrast from a carbon-epoxy composite [17]. The deep learning method used here, based on the software Dragonfly from ORS,[2] which runs locally on a workstation, is generalized, flexible, and straightforward to apply. It is also robust, even with images of low or zero contrast between phases on interest, as in Fig. 1a.

Methods

Material and imaging

The image slices in Fig. 1 are from an X-ray CT image of a rod-shaped SiC–SiC composite test specimen, fabricated as described above,[3] with diameter of approximately 2 mm. A series of CT images were also obtained from a second specimen of the same

[1] The variational method for segmentation at the fiber tow scale begins with a prior geometric model of the weave topology that is iteratively matched to the μCT image via an optimization process [47].

[2] Object Research Systems, Montreal, Canada (free of charge for non-commercial use) [23].

[3] The test specimen was supplied by Prof. G Morscher: further details of the fabrication method are given in [48].

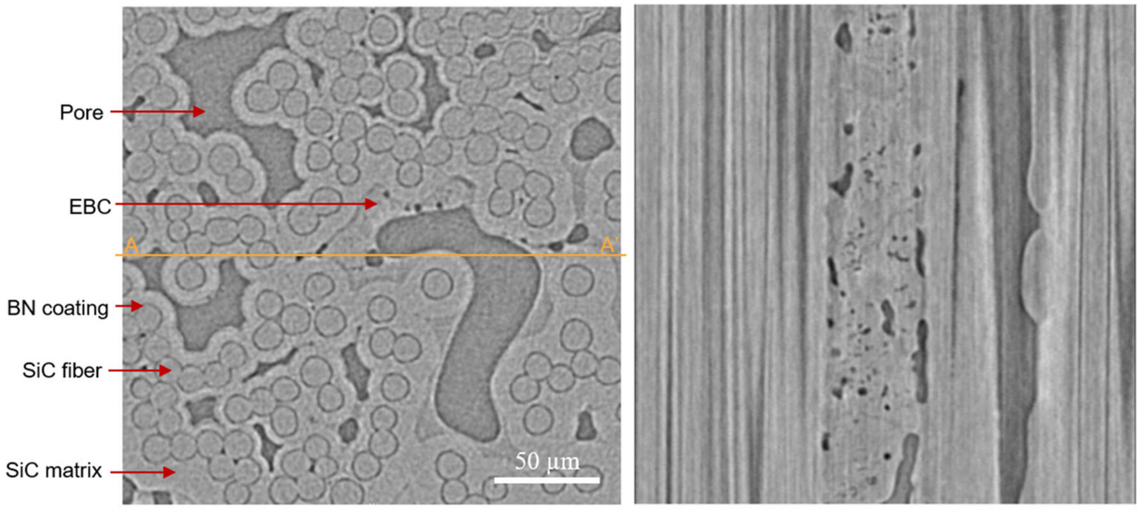

(a) (b)

Figure 1 **a** Transverse and **b** Longitudinal sections of CT image of SiC–SiC composite. The fibers and matrix can be identified visually through shape and surrounding recognition in the transverse section but not in the longitudinal section. Line A–A′ in **a** indicates location of longitudinal section in **b**.

composite, while the specimen was subjected to tensile loads parallel to the fibers (in the vertical direction), using an in situ test rig [2, 18] mounted on a synchrotron micro-tomography beamline (BL 8.3.2 at the Advanced Light Source at Lawrence Berkeley National Laboratory). For these images, a parallel white-light beam was used, with an exposure time of 40 ms for each of 1025 radiographs collected during a scan time of 1.5 min. Baseline images were also collected from both specimens beforehand with the specimen mounted outside the test rig, thereby allowing the scintillator to be positioned closer to the test specimen to minimize near-field diffraction effects. The baseline images were obtained using a monochromatic beam (17 keV), with an exposure time of 500 ms for each of 1025 radiographs, collected over 15 min. For each scan, a set of 1025 radiographs were collected and converted to a reconstructed 3D CT image using inverse Radon transforms implemented in Xi-cam, a software platform that builds on previous tomography toolkits TomoPy and gridrec [19]. Each reconstructed image in the baseline scan consisted of 2160 slices with a voxel size of 0.62 μm and total image height ∼ 1.7 mm, while the images obtained with the specimen mounted in the test rig, had voxel size of 0.58 μm and total image height ∼ 1.3 mm. The composite was scanned at 12 overlapping vertical positions (tiles) that can be stitched together to assemble an image stack of a 15 mm vertical test section of the composite. The reconstructed images were reduced from 32-bit to 8-bit for faster subsequent processing. No filters or any other image post-processing steps were conducted on the reconstructed images.

Deep learning strategy for image segmentation

Neural networks have been used recently to solve demanding image segmentation problems from other scientific domains as in medical imaging, face detection, and autonomous driving [20, 21]. When the network models are structured with multiple successive artificial neurons, the method is generally referred to as deep learning. Convolutional neural networks (CNNs) are a special case of deep learning where one or more layers of the network perform convolution operations; the specific convolution kernels are not programmed, but are learned from the input images by the deep-learning engine to extract the relevant features of an image that become useful discriminators in segmenting materials in complex images [22]. Further discussion of image texture decoding and discrimination, made possible by CNNs, is given in "Appendix A".

Operationally, CNNs can behave as image transform engines that take an input image and return a more useful output, such as a segmented image. The CNN architecture can be thought of as a formula of linear weights applied to the image pixel intensities, often combined through multiple network layers in a nonlinear fashion. The coefficients encoded in the neural network itself are learned from training data that couples example input images with example output images [22]. The iterative process of learning the weights that can reliably transform input images into output images is termed training; and it is the most computationally demanding phase of the deep-learning cycle. The segmented training data can be a selection of image slices that have been manually segmented to identify a material composition (or label) for each pixel. The trained model can then automatically segment the remaining unsegmented image slices. This process of using the trained model to transform the remaining unseen image slices is termed inference, and it is less computer intensive than the training phase.

The convolutional neural network model used in this work was designed using an online tool-sharing community repository, Infinite Toolbox, and the associated Dragonfly 4.1 software [23], which provides an interface where users can use manual segmentation tools, connect parameters, and execute network training (see "Appendix B"). Whereas the design of the network model itself and the tuning of various network parameters requires some expertise in neural network architectures, the application (training and inference) does not. Network parameters that need to be assigned include: (1) a "patch size" (in the training stage the images are split into a set of smaller 2D square patches that capture the features of interest in the image [24]); (2) a "stride-to-input ratio," which defines the positions of the neighboring patches; (3) a "batch-size" which defines the number of patches evaluated in each batch prior to updating the coefficients of the network model; (4) the number of epochs (an epoch is one training iteration, involving a pass over all batches of the training set); and (5) selection of a loss function (or cost function), to evaluate how far the output of the model deviates from the target output [25, 26] and an optimization algorithm to find optimal weights for the coefficients of the CNN [27, 28]. In the present work, the number of training slices and values of network parameters were selected after making trial runs with several combinations of parameters judged to be reasonable based on previous experience, until the quality of inference on unseen data was adequate (98% accuracy) with reasonable training time.

Results

CMC microstructure segmentation

Training data and parameters

For the training data, 17 of the 2160 transverse image slices from a baseline CT scan of the composite were selected and segmented manually to identify each pixel with one of four material phases: fibers, matrix, pores, and environmental barrier coating (EBC), in addition to empty space surrounding the composite. The image slices were oriented normal to the axis of the rod-shaped composite test specimen. In this orientation, the visual distinction between fibers and matrix is straightforward because of the presence of the thin dark coating of BN on each fiber (the BN itself, which was of sub-pixel thickness, was not segmented as a separate phase). The time taken for manual segmentation was approximately 1–3 h per slice.

The segmented image slices were used to train a neural network architecture FCDenseNet [29, 30], selected from the Dragonfly toolbox. The data from the image slices were divided into patches, with 80% of the patches being used for the training itself and the remaining 20% of the patches being retained for validation, to provide an unbiased evaluation of how accurately the trained model is capable of segmenting the images correctly. The trained model achieved 85% accuracy after 200 epochs and 98% accuracy after 300 epochs, (accuracy defined as the percentage of pixels with the material phase correctly identified). The computer time required for training depends on network complexity, the number of training images, the number of epochs required to achieve a satisfactory accuracy for each training image, and the performance of the hardware being used. In this work, each epoch required 10 min using an Nvidia Quadro P4000 graphics processing unit (GPU). This translates to 50 computer hours of training.

Segmentation results

After the training stage, the model was applied to segment the remaining (previously unseen) 2143 transverse image slices from the same CT scan, thereby producing a fully segmented image stack. The model succeeded in segmenting all four phases at a rate of 29 image slices per minute (approximately

75 min for the complete CT image). The accuracy of the segmentation output was assessed visually by comparing randomly selected segmented slices with the corresponding original images.

Typical results are shown in Fig. 2, which compares equivalent transverse and longitudinal slices from the raw and segmented CT images. Comparison of the transverse slices, where the different phases in the raw slice can be distinguished visually, indicates that the accuracy of the segmentation appears qualitatively to be high. Comparison of the longitudinal slices indicates that the fibers, matrix, and voids are clearly distinguished in the longitudinal slice from the segmented CT image, whereas in the equivalent slice from the raw CT image none of these phases can be distinguished visually.

Although the accuracy of training the model on the target images reached 98%, this does not guarantee error-free segmentation when running the automated segmentation on unseen images. Occasional isolated regions where the phase identification was incorrect were found in the segmented transverse image slices. An obvious example is shown in Fig. 3a, where an area of EBC was mistakenly labeled as matrix.

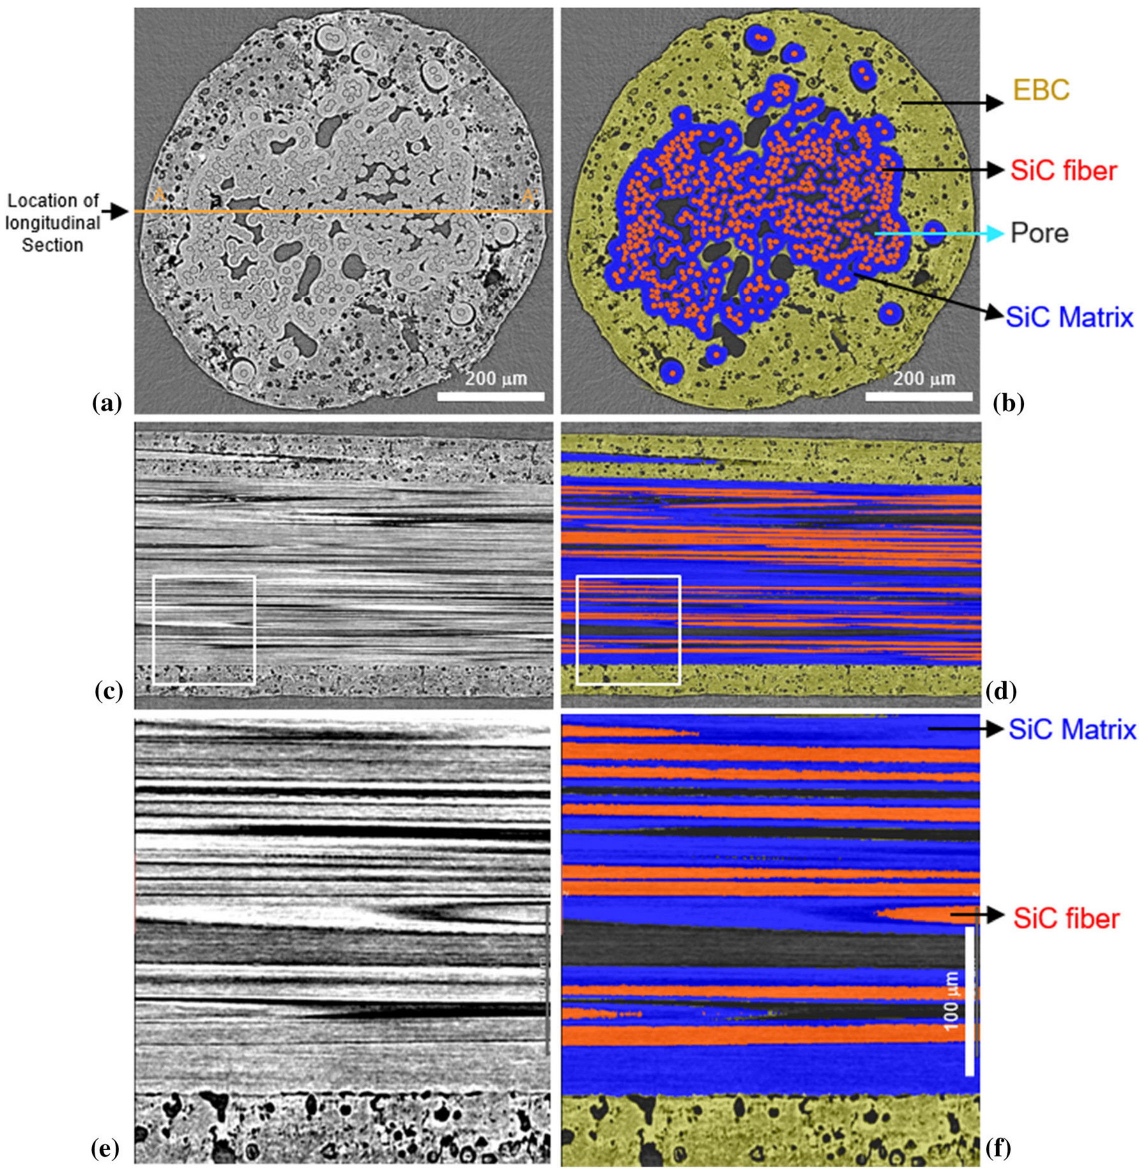

Figure 2 Transverse and Longitudinal sections of CT image of unidirectional CMC before (**a**, **c**, **d**) and after (**b**, **d**, **f**) deep-learning segmentation into the four phases indicated. **e**, **f** are enlargements of area indicated in **c**, **d**.

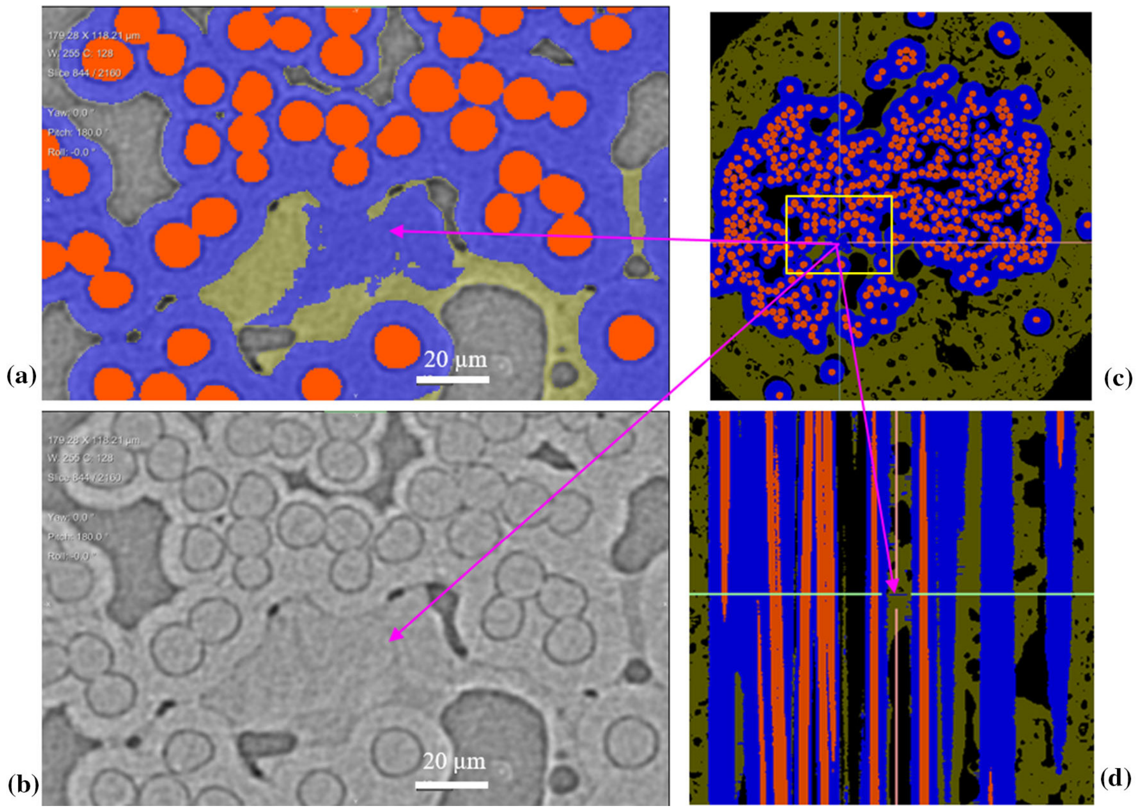

Figure 3 Example of incorrect phase identification in automated segmentation.

Viewing this area in a longitudinal section (Fig. 3d) reveals that the error is confined to a single transverse slice (the error appears on the longitudinal section as horizontal row of pixels that differ in phase from the pixels above and below). This type of error can be readily identified; several other examples are visible in Fig. 3d. Other less obvious isolated clusters of mislabeled pixels were also found in a small fraction of the slices.

Post-segmentation processing and quantification

Once the CT image is segmented by the deep-learning model, various post-segmentation image processing routines can be used to render the information in forms that help in visualizing the 3D arrangements of the different phases of the composite and to extract quantitative information on the microstructural characteristics. An example is shown in Fig. 4, where surface meshes of each of the segmented phases (fibers, pores, matrix, and EBC) are displayed both as separate entities and in combination. Such surface renderings allow simple visual inspection of microstructural characteristics such as the fiber architecture and porosity distribution. The volume fractions of each of the phases are readily calculated; the results are given in Fig. 4. The supplementaryVideo#1.mpg provides further 3D visualization of the surface meshes in Fig. 4.

Two distinct populations of pores exist in the composite: those in the environmental barrier coating and those within the CVI SiC matrix (Fig. 5). The matrix pores are typical of composites with matrices formed by CVI, in which regions between groups of fibers become closed off as the matrix builds up, thereby blocking further access by the CVI reactant gases. The restricted access of infiltrating gas as the matrix builds up is also responsible for the thickness of the layer of SiC deposited on the fibers being larger near the exterior of the composite than in the central regions (Fig. 5). The matrix pores are all extremely elongated parallel to the fibers, with many being interconnected and most being completely sealed by the matrix. However, some of the pores that appear to be closed in the transverse image slices over most of the length of the scanned section of the composite were open at one end to the exterior surface of the composite after the CVI infiltration was completed.

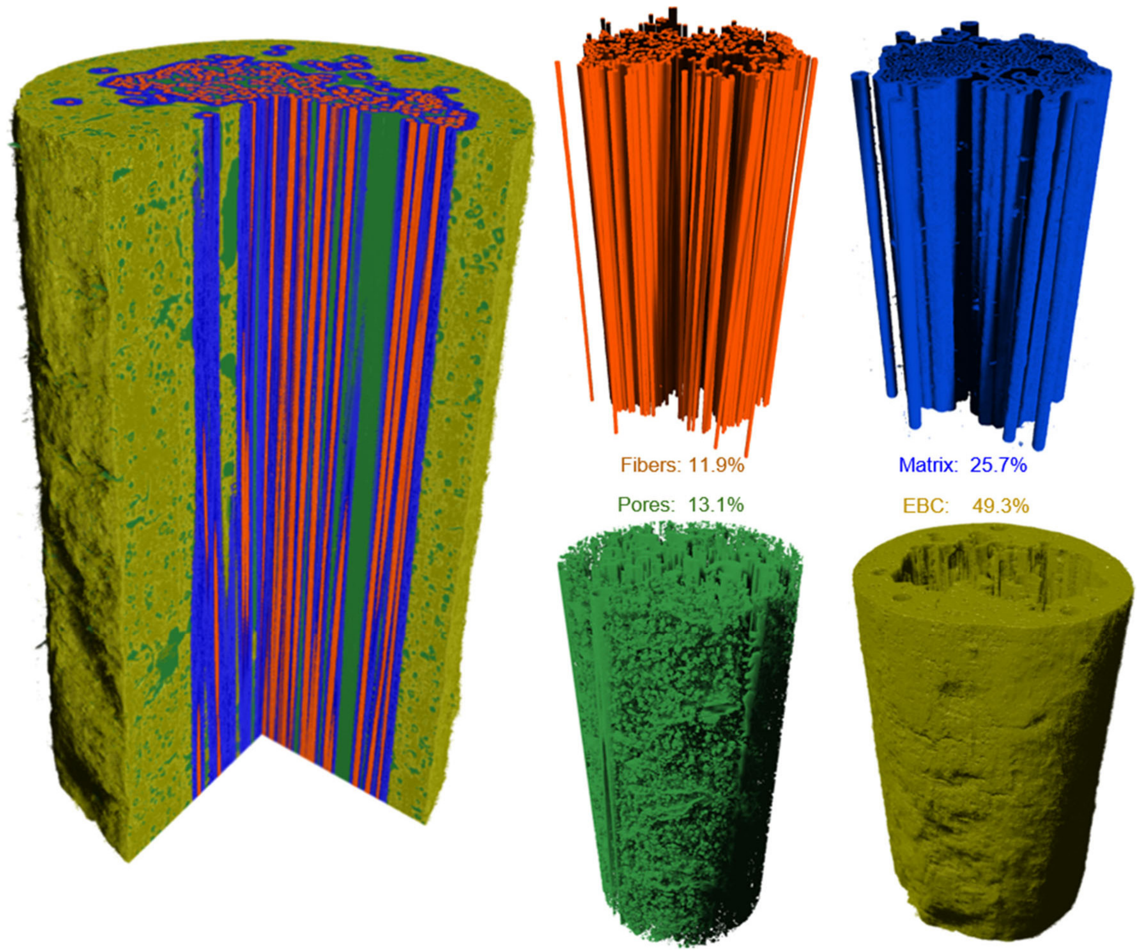

Figure 4 CT image of SiC–SiC composite after deep learning segmentation. Volume fractions of the different phases can be accurately estimated from the segmented volumes.

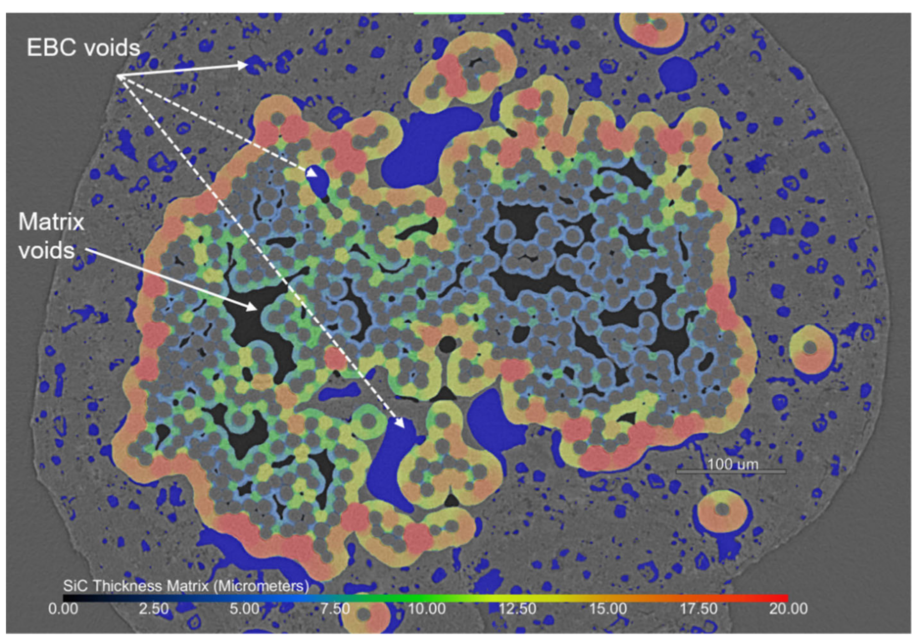

Figure 5 Transverse section color coded to indicate thickness of CVI SiC matrix material around fibers. Matrix pores are black; EBC pores are blue. Matrix pores are closed to the outside. EBC pores that appear to be surrounded by fibers and matrix (dashed arrows) are open to the surrounding EBC at locations above or below this image slice.

This allowed some of the EBC slurry to enter those pores.

The pores in the EBC are mostly not connected to each other or to the external surface. Two types of pores are evident: pores that are roughly equiaxed and contained completely within the EBC; and larger pores between the EBC and the adjacent surface of the SiC/SiC composite, which resulted from shrinkage-induced debonding between regions of the EBC and the CVI SiC matrix during sintering of the EBC (Fig. 5).

Segmentation of matrix cracks

When an axial tensile force above a critical value was applied to the composite test specimen using the in situ test rig, multiple cracks formed in the matrix and EBC, with crack surfaces nearly perpendicular to the fiber direction [1, 2, 18, 31, 32]. The cracks could be seen in longitudinal slices of CT images (Fig. 6a). Because of the presence of the thin BN coatings on the fibers, the cracks did not propagate through the fibers. Instead, debonding occurred along the BN coating between the fibers and matrix, while the matrix cracks bypassed the fibers. As the force was further increased, fiber fracture occurred at distributed locations and the opening displacements of the matrix cracks increased. Being able to track the evolution and spatial distribution of such cracks is key to understanding the relation between mechanical properties of the composite and microstructural features such as the spatial distribution of fibers.

If the crack opening is sufficiently large, the cracks can be segmented simply based on pixel intensity contrast between the matrix and the empty space between the crack surfaces. However, for crack opening displacements smaller than about 10 μm (as in Fig. 6a), the image consists of overlapping bright/dark phase-contrast fringes from the edges of the crack [18], which precludes a simple intensity-based segmentation. These fringes exist even when the crack opening is much smaller than the resolution of the image. In addition, the images are affected by intact fibers bridging the crack. The opening displacements of the crack surfaces are dependent on the applied load and microstructural properties of the composite, including frictional sliding forces between the fibers and matrix [33–35]. In ceramic composites of interest for high-temperature applications, the crack openings in the initial stages of

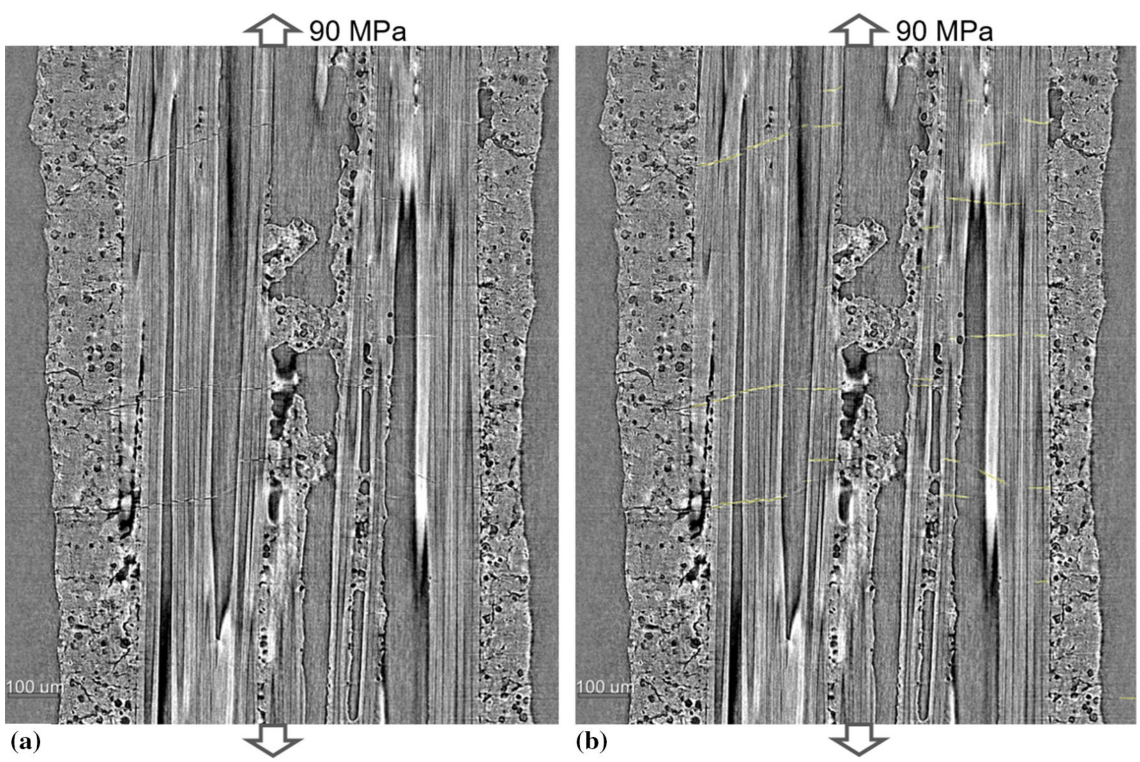

Figure 6 Longitudinal section with cracks before (**a**) and after (**b**) deep-learning segmentation. Segmented cracks labeled in yellow.

damage evolution fall within the range where simple intensity-based segmentation is not possible.

The results of applying the deep-learning procedure to detect matrix cracks in longitudinal image slices are shown in Figs. 6 and 7 (cracks found in the EBC surrounding the composite were not labeled as they were mostly generated during processing owing to shrinkage). The training was done using a set of 112 image slices from a CT scan obtained while a tensile stress of 90 MPa was applied to the test specimen. Manual segmentation of cracks was far less time consuming than manual segmentation of the material components in the transverse slices, taking approximately 5 min to segment all of the cracks in each longitudinal slice. The training stage was also much faster. Training of a U-net [36] CNN for 88 epochs, achieved an accuracy of 99% in approximately 7 h of computer time (20% of the manually segmented trained data were used for validation). The model then took 30 min to segment the cracks in an entire stack of longitudinal slices. A typical image slice before and after segmentation is shown in Fig. 6. A 3D view of the cracks from the full stack is shown in Fig. 7. Further 3D renderings of the cracks are given in supplementaryVideo#2.mpg.

In unidirectional ceramic composites with uniform distributions of fibers and matrix, the matrix cracks extend through the entire cross section of the composite and form in a near-periodic array, while the fibers remain intact, forming bridges between the crack surfaces [1, 2, 18, 31, 32]. However, in the composite specimens used in this study, the fiber distribution is very nonuniform and the CVI matrix, which was formed by deposition of SiC on the fibers, is correspondingly nonuniform (Fig. 6). As a result, there are many regions in a given transverse cross section where there are isolated islands consisting of groups of fibers or single fibers surrounded by matrix. When loaded in the axial direction, these islands act as independent composites, with no correlation of the axial positions of matrix cracks between different islands. Therefore, the 3D view of the cracks from the full image stack (Fig. 7) shows many cracks of limited lateral extent, rather than a few cracks that extend through the entire cross section, as seen in other similar composites with more uniform fiber distributions [1–3].

Discussion

Accuracy of microstructural segmentation

The accuracy of the deep-learning segmentation method employed here is dependent on three sources of error. One is in the quality of the original CT image data, as affected by the resolution of the image, near-field diffraction effects, the presence of noise, and the possible presence of artifacts. In the present study, this source of error does not appear to be a limitation in the segmentation of fibers, matrix, EBC coating,

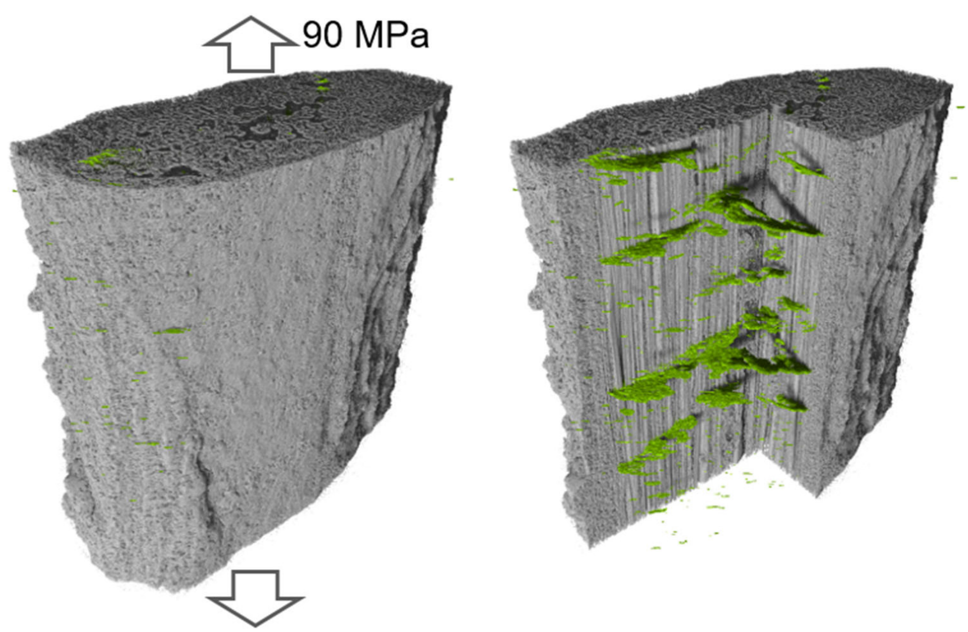

Figure 7 3D renderings of the composite in Fig. 6, with cracks highlighted in green.

and voids. However, we did not attempt to segment the thin BN coating on the surfaces of the fibers, which has thickness smaller than the voxel size.

A second source of error arises from the manual segmentation process that must be used to provide the training images. This is the most tedious step of the process, and there will inevitably be some inadvertent errors in labeling among the large number of pixels (2560 $\times$ 2560) in each image slice, as well as some judgment errors for pixels located at the boundaries of two phases. The first of these might be expected to occur at single isolated pixels and have minimal effect on the final trained model. The second might be expected to affect only the exact location of the boundaries within a distance determined by the pixel size or the image resolution.

The third source of error is mislabeling of pixels at the inference stage, owing to limitations in the accuracy of the coefficients of the network model. The accuracy of the model is dependent on some of the parameters set in the training stage (number of training images, number of iterations patch size and locations, batch size). Up to a certain limit, the error rate generally decreases as the number of training images and iterations increases, so there is a trade-off of accuracy of the network model with computer time needed for training. However, excessive iterations can lead to overfitting of the weights in the model and poor performance in segmentation of unseen images. The error functions were evaluated for both the training patches and the patches set aside for validation after each epoch to ensure that both errors trended down in successive epochs, thus ensuring that overfitting of the model was avoided (overfitting of the model being signaled by the validation error increasing). In the present study, an error rate of 2% was achieved with 300 epochs, using training parameters judged to be reasonable based on previous experience with other applications. A more systematic study to optimize the choice of training parameters and quantify the trade-off of computational time and error rate is under way.

The most obvious errors in the segmented images are of the type shown in Fig. 3, where an area of EBC was mistakenly labeled as matrix. In many (but not all) examples examined, the error was confined to a single transverse slice, as in Fig. 3. Since we have prior knowledge that there are no thin platelets of any phase in these composites, these patches of the segmentation can be dismissed as erroneous. Such errors could be detected and corrected automatically through use of a search routine that searches the segmented transverse image slices for platelets of pixels labeled with a different phase than the pixels above and below. A more general approach that would likely avoid such errors, albeit with higher computational cost, would be to use recently developed 3D CNN models [37, 38], and capable of processing multiple consecutive image slices simultaneously.

With the error rate achieved in this study, the volume fractions of each of the phases could be computed from the segmented image stack with high accuracy. Moreover, the spatial distributions of the phases were displayed with very few noticeable errors. Thus, the segmented images would be suitable for various post-processing routines for extracting other quantitative information, such as fiber tracking to determine whether the fibers within a fiber tow are intertwined relative to one another. This can be done by extracting the center lines of the segmented fibers and inspecting the spatial coordinate variance in X–Y, evaluated over the Z-axis to determine the straightness of the fiber, as demonstrated in the template matching technique by Czabaj et al. [15].

Matrix cracks

Observations of the longitudinal slices after being segmented for matrix cracks by the CNN model in the inference stage (see supplementaryVideo#3.mpg) show that the model succeeded in labeling all of the cracks that could be identified by visual inspection. However, there were some errors where additional crack labels were incorrectly assigned in voids and in air surrounding the composite. These errors occurred where there were image artifacts associated with phase contrast effects. Examples are visible as horizontal streaks outside the lower right side of the composite in Fig. 6, between the outside edge of the EBC and the edge of the micrograph. These errors can be readily removed by a post-segmentation algorithm that removes all the labeled cracks that fall within the segmented phases of voids or outside air. The mislabeling in the voids or outside air are found only in isolated patches and are not considered to affect the visualization of crack planes. Mislabels were eliminated in Fig. 7 and the supplementaryVideo#2.mpg.

The accuracy of crack segmentation is dependent on the same sources of errors mentioned in the previous section. Whereas the quality of the original CT image data, as affected by the resolution of the image, near-field diffraction effects, the presence of noise, and the presence of artifacts, was not a limiting factor for microstructural segmentation, it is likely the dominant factor in identification and segmentation of cracks that form during the initial stages of damage evolution. In the manual segmentation stage of training, some of these cracks are challenging to identify visually and distinguish from other image features, making them more susceptible to judgment error.

Generalization of deep learning image segmentation

The image segmentation challenge addressed in this study, involving low (or zero) contrast X-ray images from SiC–SiC composites, is not unique to these materials. For example, low-contrast images are typical of biological materials. The deep-learning solution described here, which was successful in segmenting phases that showed identical image intensities, is not restricted to any particular material system. It provides a general framework, described in more detail in "Appendix B," for executing both the training and inference stages of the deep learning cycle. In all applications of deep learning for image segmentation [39], this involves the following steps: manually labeled ground truth images are first paired with unprocessed images; then those images are presented to a deep-learning model that can accept them as training data; then, the trained model is used to process never-before-seen images. The general framework is not restricted to use of any particular neural network: it includes the capability to select from a list of commonly available neural networks and provides tools for preparing the ground truth images, for presenting those data to the deep learning model, and for analyzing the unseen images. Most importantly, it provides a workflow to iterate and tune simple deep-learning models that can be used without having to perform any programming or needing any expertise in the arcane details associated with the standard deep-learning frameworks. On the other hand, it also provides options that give expert users and programmers the flexibility to tune advanced parameters and derive and implement new deep learning models of varying network architectures.

Since previous deep-learning image segmentation studies have all relied directly on programming for their deep-learning solutions to solve specific applications, the availability of this tool can be expected to expand the use of deep-learning methods for segmentation beyond those with expertise in deep-learning language libraries.

Conclusion

A deep-learning image segmentation procedure was used successfully to segment 3D X-ray CT images from a fiber-reinforced ceramic composite, in which the fibers, matrix and an environmental barrier coating showed identical image intensities. This was achieved with reasonable run times on a local workstation, thus demonstrating a practical solution to the problem of dealing with the large amount of data involved in such images and the resulting barrier to extraction of quantitative information.

In transverse slices of the segmented image, the identification of the various phases was consistent with visual identification from the original image, enabled by shape and edge information. In longitudinal image slices, the phases could not be distinguished visually in the original image, whereas in the segmented image they were clearly distinguished.

The deep-learning method was also effective for automated segmentation of matrix cracks produced by in situ loading of the composite. The segmented images were amenable to post-segmentation image processing to render the distributions of phases and cracks in forms that help in visualizing their distributions in 3D space and to extract quantitative information on the microstructural characteristics. The results demonstrated that a nonuniform distribution of fibers in a unidirectional composite has a large effect on the nature and distribution of matrix cracks that form as the initial damage in tensile loading parallel to the fibers.

Acknowledgements

This research was supported by the US National Science Foundation PIRE program, Grant number

1743701, led by Prof. G Singh at Kansas State University. Beam time at the Berkeley Advanced Light Source was provided under an ALS Approved Program led by M. Czabaj. We thank P. Creveling, M. Czabaj, G. Morscher and D. Parkinson for assistance with in situ experiments at the Berkeley ALS.

Compliance with ethical standards

Conflict of interest Authors affiliated with Object Research Systems (ORS) developed Dragonfly, the software package that was used in this work. The software is licensed commercially, at a cost to most industry licensees, but at no cost for non-commercial users. University of Colorado workers have no interests to declare.

Appendix A: Deep-learning image segmentation by CNNs

Semantic image segmentation—the labeling of pixels in an image according to the object they constitute—is a deep-learning method that was first applied to scientific imaging with the description of the U-Net architecture in 2015 [37], although non-scientific applications predate that work [40]. These network models are built as CNNs, in which image data are subdivided into patches and fed through a network of neurons, which are arranged in sequential layers. Each neuron in a given layer receives input from neurons in the previous layer, transforms the input signal, and then passes the result to a set of neurons in the next layer. The signals are integrated in successive layers of the network, where ultimately higher order neurons may have remarkable discriminative value by selectively amplifying various signals from previous layers. CNNs employ convolution operations as their first layer(s) of neurons. The coefficients in the convolution kernels are seeded randomly initially and are refined iteratively in the learning phase, where the output of the CNN is compared with the manually segmented image (training). The learned weights of different neurons confer the extreme selectivity which gives the networks their power. Similar to biological neural networks, these models are able to interpret texture that is observed in the image and use that to help distinguish hallmarks of one visual object from another. The early neurons are able to encode texture, but not because they are programmed to recognize specific patterns, edges, gradients, or other primitive image descriptors. Rather, the coefficients of the kernels in those early-stage convolutional filters are learned through the reinforcement process of network training. Consequently, if a convolutional kernel conveys a meaningful signal that can provide discriminating value in the network, it will be preserved and up-weighted. CNNs are also used in object detection and other deep learning enabled computer vision techniques. Further discussion of CNNs can be found in Refs. [41, 42].

Appendix B: Implementation of deep learning

To make the deep-learning method easy to apply to a broad class of image segmentation problems, a general framework for executing both the training and inference stages of the deep-learning cycle was developed. Both the training and inference tools rely on software libraries TensorFlow and Keras available from Google (Mountain View, California USA) [27]. These libraries have been integrated into a desktop software platform for image manipulation and analysis named Dragonfly (developed and licensed by Object Research Systems, Montreal, Canada), available at no cost under non-commercial licensing terms. Further details on the software integration and instructions on how to download Dragonfly, the software described in this paper, can be found on the main Object Research Systems website [23]. To download the trained deep-learning models, training images and raw data used for this work along with the CMC image data set refer to the Materials Data Facility repository [43]. Other deep learning models can be found at the Object Research Systems online tool-sharing community repository (Infinite Toolbox).

It is an important goal to make the deep learning solution accessible by non-experts, but provide high-value flexibility for experts that want to use the same system. There are no major runtime parameters associated with the inference stage of deep learning, so this goal of serving both classes of users is accomplished in the interface of how users configure and train their models.

When setting up model training, the basic panel exposes standard parameters: patch size, stride ratio, and batch size, optimization function, and loss function. The training for the microstructural segmentation in "CMC microstructure segmentation" section was done using the neural network architecture FCDenseNet [29, 30], with a patch size of 64 × 64 pixels, stride ratio of 0.5 and batch size 16. The loss function "categorical-crossentropy" and optimization algorithm "Adam" were used as the default functions of the Keras Library. For the segmentation of matrix cracks, the U-net [36] CNN was used with patch size of 128 pixels, stride to input ratio 1, batch size 32, loss function "categorical crossentropy," and optimization function "Adadelta." Initially, a U-Net model was trained for the microstructural segmentation. However, the inference results had many errors, so a model with deeper neural net architecture, FCDenseNet, was deployed (with the cost of a much longer time taken for training). The patch size was set to 64 × 64 since the features of the microstructure, such as individual fibers, spanned small areas in the image. For the crack segmentation, a U-net model proved adequate. The cracks spanned larger horizontal pixel areas in the longitudinal images, thus allowing use of larger patch size (128 × 128 pixels). The default optimization function in the Keras library was used for both models.

To support greater control, an optional advanced parameters panel that exposes options for additional logging (including support for TensorBoard), fine tuning parameters for the optimization function, conditions for early termination of training, and conditions for learning rate reduction. These parameters are understood to be beyond the scope of non-experts, but match many of the controls experts would have if they were directly programming their own solution with lower level tools. This platform also integrates methods for data augmentation to reduce the raw training data that must be manually prepared, and data set aside for validation during the course of training. Models that were previously trained can be accessed for further iterations of training.

In the inference stage, the deep-learning model behaves as a simple image transform engine, which can take a single gray-scale image and return a single segmented image. From the user perspective, the trained deep-learning model behaves as a simple image filter. To simplify user interaction, the interface allows the selection from a library of trained models that can be applied like any standard image filter. A preview mode allows viewing of the output from any of the trained models applied to any single image or any sub-area of an image. This permits rapid assessment of whether any of those models is suitable for application at hand.

Appendix C: Validation of deep learning

The general application of the deep-learning solution described here was validated by applying it to a set of serial-section transmission electron micrographs of *Drosophila melanogaster* neurons that has been used as an image segmentation challenge [44, 45]. These images have many structural features (plasma membranes) that are difficult to discriminate with standard algorithms. The challenge micrographs include a stack of 30 raw images that have not been processed and a stack of 20 manually segmented images to be used as training data for machine learning technique development. A recent report documented the successful segmentation of these micrographs with a CNN called FusionNet [46]. The FusionNet architecture was reproduced here and implemented as a model in the new Dragonfly toolkit. Following training, a set of 10 validation slices were used to assess the inference quality. Visual inspection confirmed that proper segmentation of the plasma membrane was achieved.

Electronic supplementary material: The online version of this article (https://doi.org/10.1007/s10853-020-05148-7) contains supplementary material, which is available to authorized users.

References

[1] Bale HA, Blacklock M, Begley MR et al (2012) Characterizing three-dimensional textile ceramic composites using synchrotron X-ray micro-computed-tomography. J Am Ceram Soc 95:392–402. https://doi.org/10.1111/j.1551-2916.2011.04802.x

[2] Bale HA, Haboub A, Macdowell AA et al (2013) Real-time quantitative imaging of failure events in materials under load at temperatures above 1,600 C. Nat Mater 12:40–46. https://doi.org/10.1038/nmat3497

[3] Chateau C, Gélébart L, Bornert M et al (2011) In situ X-ray microtomography characterization of damage in SiCf/SiC minicomposites. Compos Sci Technol 71:916–924. https://doi.org/10.1016/j.compscitech.2011.02.008

[4] Wright P, Fu Z, Sinclair I, Spearing SM (2008) Ultra high resolution computed tomography of damage in notched carbon fiber-epoxy composites. J Compos Mater 42:1993–2002. https://doi.org/10.1177/0021998308092211
[5] Moffat AJ, Wright P, Buffière JY et al (2008) Micromechanisms of damage in 0 splits in a [90/0] s composite material using synchrotron radiation computed tomography. Scr Mater 59:1043–1046. https://doi.org/10.1016/j.scriptamat.2008.07.034
[6] Mazars V, Caty O, Couégnat G et al (2017) Damage investigation and modeling of 3D woven ceramic matrix composites from X-ray tomography in situ tensile tests. Acta Mater 140:130–139. https://doi.org/10.1016/j.actamat.2017.08.034
[7] Cox BN, Bale HA, Begley MR et al (2014) Stochastic virtual tests for high-temperature ceramic matrix composites. Annu Rev Mater Res 44:479–529. https://doi.org/10.1146/annurev-matsci-122013-025024
[8] Saucedo-Mora L, Zou C, Lowe T, Marrow TJ (2017) Three-dimensional measurement and cohesive element modelling of deformation and damage in a 2.5-dimensional woven ceramic matrix composite. Fatigue Fract Eng Mater Struct 40:683–695. https://doi.org/10.1111/ffe.12537
[9] Barnard HS, MacDowell AA, Parkinson DY et al (2017) Synchrotron X-ray micro-tomography at the advanced light source: developments in high-temperature in situ mechanical testing. J Phys: Conf Ser 849:012043. https://doi.org/10.1088/1742-6596/849/1/012043
[10] Larson NM, Cuellar C, Zok FW (2019) X-ray computed tomography of microstructure evolution during matrix impregnation and curing in unidirectional fiber beds. Compos Part A Appl Sci Manuf 117:243–259. https://doi.org/10.1016/j.compositesa.2018.11.021
[11] Marshall DB, Cox BN (2008) Integral textile ceramic structures. Annu Rev Mater Res 38:425–443. https://doi.org/10.1146/annurev.matsci.38.060407.130214
[12] Saucedo-Mora L, Lowe T, Zhao S et al (2016) In situ observation of mechanical damage within a SiC–SiC ceramic matrix composite. J Nucl Mater 481:13–23. https://doi.org/10.1016/j.jnucmat.2016.09.007
[13] Perciano T, Ushizima DM, Krishnan H et al (2017) Insight into 3D micro-CT data: exploring segmentation algorithms through performance metrics. J Synchrotron Radiat 24:1065–1077. https://doi.org/10.1107/S1600577517010955
[14] Straumit I, Lomov SV, Wevers M (2015) Quantification of the internal structure and automatic generation of voxel models of textile composites from X-ray computed tomography data. Compos Part A Appl Sci Manuf 69:150–158. https://doi.org/10.1016/j.compositesa.2014.11.016
[15] Czabaj MW, Riccio ML, Whitacre WW (2014) Three-dimensional imaging and numerical reconstruction of graphite/epoxy composite microstructure based on ultra-high resolution X-ray computed tomography. In: Proceedings of the American Society for Composites—29th technical conference ASC 2014; 16th US-Japan conference on composite materials ASTM-D30 Meeting, vol 105, pp 174–182
[16] Haberl MG, Churas C, Tindall L et al (2018) CDeep3M—Plug-and-Play cloud-based deep learning for image segmentation. Nat Methods 15:677–680. https://doi.org/10.1038/s41592-018-0106-z
[17] Sinchuk Y, Kibleur P, Aelterman J et al (2020) Variational and deep learning segmentation of very-low-contrast X-ray computed tomography images of carbon/epoxy woven composites. Materials (Basel) 13:936. https://doi.org/10.3390/ma13040936
[18] Haboub A, Bale HA, Nasiatka JR et al (2014) Tensile testing of materials at high temperatures above 1700 C with in situ synchrotron X-ray micro-tomography. Rev Sci Instrum 85:1–13. https://doi.org/10.1063/1.4892437
[19] Pandolfi RJ, Allan DB, Arenholz E et al (2018) Xi-cam: a versatile interface for data visualization and analysis. J Synchrotron Radiat 25:1261–1270. https://doi.org/10.1107/S1600577518005787
[20] Greenspan H, Van Ginneken B, Summers RM (2016) Guest editorial deep learning in medical imaging: overview and future promise of an exciting new technique. IEEE Trans Med Imaging 35:1153–1159
[21] Farfade SS, Saberian M, Li LJ (2015) Multi-view face detection using Deep convolutional neural networks. In: ICMR 2015—proceedings of the 2015 ACM international conference on multimedia retrieval. Association for Computing Machinery, Inc, New York, New York, USA, pp 643–650
[22] Badrinarayanan V, Kendall A, Cipolla R (2017) SegNet: a deep convolutional encoder-decoder architecture for image segmentation. IEEE Trans Pattern Anal Mach Intell 39:2481–2495. https://doi.org/10.1109/TPAMI.2016.2644615
[23] Dragonfly|3D visualization and analysis solutions for scientific and industrial data|ORS. https://www.theobjects.com/dragonfly/index.html. Accessed 21 Mar 2020
[24] Hamwood J, Alonso-Caneiro D, Read SA et al (2018) Effect of patch size and network architecture on a convolutional neural network approach for automatic segmentation of OCT retinal layers. Biomed Opt Express 9:3049. https://doi.org/10.1364/boe.9.003049
[25] Sudre CH, Li W, Vercauteren T et al (2017) Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Cardoso MJ, Arbel T, Carneiro G

et al (eds) Deep learning in medical image analysis and multimodal learning for clinical decision support. Springer, Cham, pp 240–248
[26] Losses - Keras Documentation. https://keras.io/losses/. Accessed 21 Mar 2020
[27] Geron A (2019) Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: concepts, tools, and techniques to build intelligent systems, 2nd edn. O'Rilly, Sebastopol
[28] Optimizers - Keras Documentation. https://keras.io/optimizers/. Accessed 21 Mar 2020
[29] Milletari F, Navab N, Ahmadi SA (2016) V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: Proceedings—2016 4th international conference on 3D vision, 3DV 2016. Institute of Electrical and Electronics Engineers Inc., pp 565–571
[30] Jegou S, Drozdzal M, Vazquez D et al (2017) The one hundred layers tiramisu: fully convolutional densenets for semantic segmentation. In: IEEE computer society conference on computer vision and pattern recognition workshops 2017 July, pp 1175–1183. https://doi.org/10.1109/CVPRW.2017.156
[31] Arhatari BD, Zonneveldt M, Thornton J, Abbey B (2017) Local structural damage evaluation of a C/C-SiC ceramic matrix composite. Microsc Microanal 23:518–526. https://doi.org/10.1017/S1431927617000459
[32] Spearing SM, Zok FW, Evans AG (1994) Stress corrosion cracking in a unidirectional ceramic-matrix composite. J Am Ceram Soc 77:562–570. https://doi.org/10.1111/j.1151-2916.1994.tb07030.x
[33] Marshall DB (1984) An indentation method for measuring matrix-fiber frictional stresses in ceramic composites. J Am Ceram Soc 67:C259–C260. https://doi.org/10.1111/j.1151-2916.1984.tb19690.x
[34] Marshall DB, Cox BN, Evans AG (1985) The mechanics of matrix cracking in brittle-matrix fiber composites. Acta Metall 33:2013–2021. https://doi.org/10.1016/0001-6160(85)90124-5
[35] Marshall DB, Evans AG (1985) Failure mechanisms in ceramic-fiber/ceramic-matrix composites. J Am Ceram Soc 68:225–231. https://doi.org/10.1111/j.1151-2916.1985.tb15313.x
[36] Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention, pp 234–241
[37] Kamnitsas K, Ledig C, Newcombe VFJ et al (2017) Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med Image Anal 36:61–78. https://doi.org/10.1016/j.media.2016.10.004
[38] Prasoon A, Petersen K, Igel C et al (2013) Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: Mori K, Sakuma I, Sato Y, Barillot C, Navab N (eds) Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). Springer, Berlin, pp 246–253
[39] Ronneberger O, Fischer P, Brox T (2015) U-net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells W, Frangi A (eds) Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). Springer, Berlin, pp 234–241
[40] Garcia-Garcia A, Orts-Escolano S, Oprea S et al (2017) A review on deep learning techniques applied to semantic segmentation. arXiv:1704.06857
[41] Aloysius N, Geetha M (2018) A review on deep convolutional neural networks. In: Proceedings of the 2017 IEEE international conference on communication and signal processing, ICCSP 2017. Institute of Electrical and Electronics Engineers Inc., pp 588–592
[42] Dettmers T (2015) Deep learning in a nutshell: core concepts|NVIDIA Developer Blog. In: Nvidia Dev. https://developer.nvidia.com/blog/deep-learning-nutshell-core-concepts/ . Accessed 30 Jun 2020
[43] Badran A, Marshall DB, Legault Z et al (2020) XCT dataset and deep learning models for automated segmentation of computed tomography images of fiber-reinforced composites. Mater Data Facil Open. https://doi.org/10.18126/SAIM-CV6C
[44] Arganda-Carreras I, Turaga SC, Berger DR et al (2015) Crowdsourcing the creation of image segmentation algorithms for connectomics. Front Neuroanat 9:1–13. https://doi.org/10.3389/fnana.2015.00142
[45] Cardona A, Saalfeld S, Preibisch S et al (2010) An integrated micro-and macroarchitectural analysis of the drosophila brain by computer-assisted serial section electron microscopy. PLoS Biol 8:1000502. https://doi.org/10.1371/journal.pbio.1000502
[46] Cheng D, Meng G, Xiang S, Pan C (2017) FusionNet: edge aware deep convolutional networks for semantic segmentation of remote sensing harbor images. IEEE J Sel Top Appl Earth Obs Remote Sens 10:5769–5783. https://doi.org/10.1109/JSTARS.2017.2747599
[47] Bénézech J, Couégnat G (2019) Variational segmentation of textile composite preforms from X-ray computed tomography. Compos Struct 230:111496. https://doi.org/10.1016/j.compstruct.2019.111496
[48] Zhou J, Almansour AS, Chase GG, Morscher GN (2017) Enhanced oxidation resistance of SiC/SiC minicomposites

via slurry infiltration of oxide layers. J Eur Ceram Soc 37:3241–3253. https://doi.org/10.1016/j.jeurceramsoc.2017.03.065