Segmentation Method Comparison for Residual Fiber Length Measurement Across Tiled Microscopy Images

Authors: Chris O'Brien[1,2], Jay Pike[1,2], Luis Martinez[2], Chad Duty[1,2], Seokpum Kim[2], Kris Villez[1,2]

Affiliations: [1]University of Tennessee – Knoxville, Knoxville, TN, [2]Oak Ridge National Laboratory, Oak Ridge, TN

Corresponding author: Chris O'Brien, cobrien8@vols.utk.edu

Address: 414 Dougherty Engineering Bldg. Knoxville, TN 37996-2210

Abstract

Fiber length distribution (FLD), in part, governs mechanical properties in discontinuous fiber composites, yet manual measurement methods limit the high-throughput characterization needed for materials design optimization. This study compares deep learning segmentation approaches for automated FLD measurement in large-field microscopy, evaluating how method choice affects the microstructural descriptors used in structure-property-processing relationships. A critical challenge is that high-resolution microscopy images (10,000 x 10,000 pixels) must be tiled for deep learning analysis, fragmenting fibers at boundaries. We demonstrate that segmentation method proves crucial for measurement accuracy. For example, instance segmentation with Slicing Aided Hyper Inference (SAHI) preserves individual fiber integrity across tiles while semantic segmentation prioritizes speed. Comparing against manual measurement of extracted carbon fibers, YOLOv11-SAHI matched manual ground truth (238 μm weighted mean) with 40x speedup (4.5 vs 167 minutes per image). U-Net provides rapid quantification although it is at the cost of reduced accuracy due only reliably measuring stand-alone fibers. Our comparative analysis reveals that instance segmentation with SAHI better preserves length measurements while semantic segmentation prioritizes speed, providing empirical guidance for method selection. The characterization provides essential inputs for mechanical property prediction models and inverse design workflows, accelerating composite materials development cycles.

Keywords: residual fiber length**,** deep learning, image segmentation, composites, additive manufacturing**,** extrusion

Highlights

- Large-format AM fiber characterization enabled by tiled image processing
- Comparison of DL segmentation effects on fiber distribution metrics
- Instance segmentation maintains accuracy; semantic offers speed
- Foundation for open-source, standardized fiber measurement across material systems

1. Introduction

1.1 Additive Manufacturing

Polymer-based additive manufacturing (AM) consists of the layer-wise deposition of material to construct three-dimensional objects. Guided by computer-aided designs, the manufacturing technique aids in the construction of parts that consist of complex geometries not possible by

traditional manufacturing methods. AM has been rapidly adopted for prototyping, innovation, and development of parts for the automotive, aerospace, and energy sectors [1], [2], [3], [4]. Large-format additive manufacturing (LFAM) is of significant interest in these fields due to its ability to offer rapid prototyping and design flexibility without requiring expensive molds and tooling.

1.2 Residual Fiber Length in Composite Material Design

Fiber length distribution (FLD) is a critical microstructural descriptor that governs mechanical properties of discontinuous fiber-reinforced polymers (FRPs). FRPs are made from a polymer matrix combined with fibers, such as carbon, glass, or natural fibers. These composite materials enable significant weight reduction without compromising performance by leveraging the low density of polymers with the high strength and stiffness of the fiber reinforcement [2]. Load transfer between the matrix and reinforcement depends on fiber length relative to the critical fiber length, typically ~6 millimeters for carbon fiber. Fibers above this threshold can reach their ultimate strength, while shorter fibers provide diminished reinforcement [5]. Since manufacturing processes produce a distribution of fiber length, rather than scalar values, characterizing the complete FLD becomes essential for accurate property prediction [6].

Processing-induced fiber breakage presents a fundamental challenge in composite manufacturing. Value-added fibers play a large role in the minimization of final part distortion while also increasing strength [7]. In screw-extrusion processes (Figure 1), such as LFAM and injection molding, previous work has shown that shear-inducing processing parameters (i.e. screw speed, temperature, screw geometry, nozzle diameter) directly relate to the reduction of fiber length [8], [9], [10], [11]. Previous work demonstrates that polymer mixing can significantly reduce fiber length, shortening 6 mm chopped fibers (aspect ratio ~833) to aspect ratios as low as 1-20 during extrusion [12]. This fiber damage creates a complex relationship between manufacturing parameters and final microstructure [13]. For example, increase in extrusion speed can lead to reduced fiber length resulting in decreased elastic modulus and tensile strength [14]. Furthermore, changes in fiber length, orientation, and dispersion can directly lead to anisotropic properties of the final part [15]. Understanding and controlling residual fiber length, the final length distribution after processing, is therefore essential for composites optimization. The trade-off between manufacturability and performance makes accurate FLD measurement critical for establishing process-structure-property relationships and enabling data-driven design approaches in composite development for LFAM.

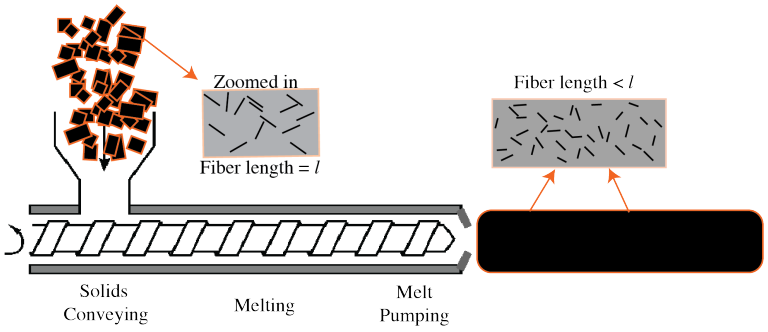

Figure 1: Extrusion process results in a reduction of fiber length. Presented is a simplified depiction of the material extrusion process where pellets contain fibers with length *l*. After being processed, the extrudate contains fibers of length < *l* due to, in part, shear forces and thermal degradation.

1.3 Measurement of Residual Fiber Length

The methods of measurement of fiber length can be labor-intensive, costly, and lack consistency. The imaging phase of the process requires multiple steps, making the entire process difficult to scale and time-consuming [14]. A small sample may contain millions of fibers, yet typically only 2000 fibers are used to represent a part. In practice, this only represents a narrow section of the part's spatial domain, limiting a comprehensive understanding of fiber properties across the entire structure. Current techniques predominantly rely on manual analysis, utilizing software like ImageJ for human identification and measurement of fibers. As a result, the measurement of 2000 fibers can take 2-4 hours to complete and becomes highly uncertain. Additionally, the geometric characteristics of fiber pose different measurement challenges. Straight, non-overlapping, and short fibers (< 1mm) are straightforward due to the small profile and low variance. Analyses become much more difficult the denser the fibers are clustered and the longer the fibers. This is in part because longer fibers have a much higher likelihood of overlapping and curving.

In the past, there have been several different attempts to automate the fiber measurement process [16], [17], [18], [19]. Commercial systems like FASEP combine specialized optical hardware with proprietary image processing software, offering reproducible results but with associated cost and accessibility constraints. Open-source alternatives such as CT-FIRE provide transparent methodologies and employ curvelet transforms but are optimized for biological collagen fibers rather than composite materials. Deep learning has recently emerged as a promising approach, with recent work demonstrating successful fiber segmentation and microstructure characterization [20], [21]. A method tailored specifically to the needs of the composites space could prove to be valuable. Currently, a major limiting factor in open-source measurement processes is the inability to capture fibers in dense, large images.

Large-field composites microscopy generates images of sizes that can exceed 10,000 x 10,000 pixels. This necessitates tiling strategies to make input image processable by deep learning methods. However, as a result fiber can be split between tiles resulting in a single fiber being broken up and causing the same fiber to be counted multiple times or missed entirely. The impact of segmentation approach on maintaining fiber integrity across tiles has not been systematically evaluated for materials research.

Beyond 2D optical methods, advanced X-ray computed tomography (CT) approaches enable three-dimensional fiber characterization. Previous work has demonstrated 3D fiber segmentation using X-ray μCT, providing complete spatial information about fiber orientation and length within composite volumes [22]. Additional work has been conducted for fiber segmentation using deep learning [23], [24]. While X-ray CT methods offer superior 3D information, they require specialized equipment, longer acquisition times, and extensive computational resources for reconstruction and analysis. In contrast, 2D optical microscopy of matrix-removed samples remains the common practice due to its accessibility and standardized procedures according to ASTM D3171 [25].

1.4 Deep Learning Application to Additive Manufacturing

Deep learning (DL) is a subset of artificial intelligence where algorithms modeled loosely after the human brain learn from the data provided to them. DL models expressly use neural networks, which excel at identifying complex patterns and making predictions or decisions based on their "learning." A key advantage of these models is their ability to extract and learn features directly

from data, eliminating the need for feature engineering [26]. Convolutional neural networks (CNNs) have revolutionized image-based tasks such as image classification [27], [28], semantic segmentation [29], [30], object detection, [31], and instance segmentation [32]. Within these networks, convolutions are used to generate feature maps, which are representations that refine important aspects of an input image, such as edges. These convolutional operations can be stacked to capture the complex spatial patterns contained within images.

Image-based CNN methods can be categorized into two broad categories: semantic segmentation and instance segmentation. The goal of a general segmentation task is the classification of pixels into pre-defined groups of object classes. Semantic segmentation is a method for labeling each pixel in an image with a class but not creating a distinction between individual objects (Figure 2, middle panel). Alternatively, instance segmentation is a method for classifying each pixel into a category but also distinguishing between different instances of the same object classes (Figure 2, right panel).

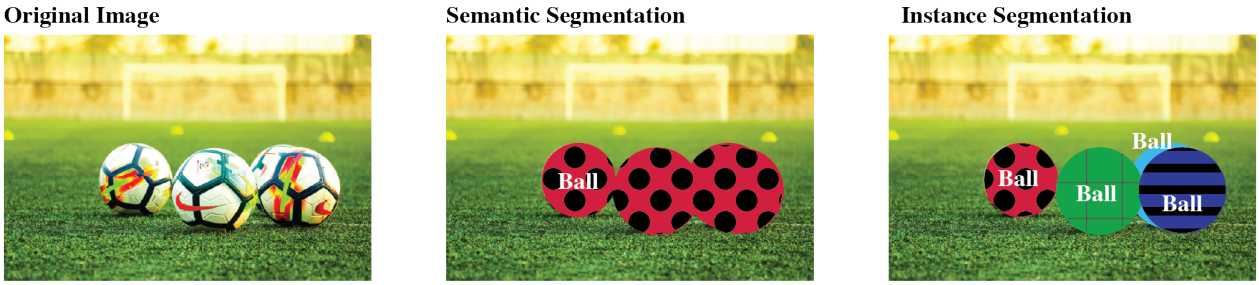

Figure 2: Comparison of segmentation methods. Semantic segmentation allows for the identification of the object in a class, but instance segmentation can differentiate between objects of the same class.

The application of DL for machine vision and image-based tasks in AM is rapidly increasing, with use cases focused on predicted stress-strain behavior, property prediction, process optimization, and in-situ monitoring for reliable resultant products [33], [34], [35]. The implementation of DL has been more prevalent in metal AM. A factor contributing to this is the vast investment from the aerospace, automotive, and biomedical industries into powder-bed AM technologies [36].
A recent 2023 review highlighted several ongoing advancements in applying machine learning to polymer-AM [33]. The execution of these methods is complicated due to the sensitivity and variability of polymer materials, complex material-parameter interactions, and limited training data [33]. Despite notable obstacles, imaged-based learnings are rapidly being published in this field. Image-based methods have demonstrated reliable prediction and correction of over/under material extrusion [37] and defect detection [38]. Non-imaged-based applications for polymer AM include defect detection [39], [40], [41], delamination [42], [43], and microstructure prediction [44]. Furthermore, non-image-based applications continue to grow rapidly such as for predicting mechanical behaviors of biopolymers [45] and print parameter optimization [46].

The segmentation methods presented in the paper are applicable to residual fiber analysis from any discontinuous composite manufacturing process. We demonstrate their application in AM contexts. Accurate automated measurement is critical for rapid quality-control feedback and for quantifying process-induced fiber breakage during print process optimization.

1.5 Challenges for Deep Learning and Large-Field Microscopy

DL applications in materials science and additive manufacturing face unique data challenges that distinguish them from traditional computer vision domains. A 2025 review of DL for AM identified managing limited, imbalanced, and noisy data as critical challenges, recommending deep generative models and few-shot learning as promising solutions [47]. These challenges are particularly an issue in microscopy-based characterization where data acquisition costs are prohibitive, and sample preparation is destructive. These conditions necessitate tailored strategies for data availability, as reflected in field-wide dataset sizes. The precedent for such strategies was, in part, established by the seminal U-Net paper, which achieved breakthrough biomedical segmentation performance using only 30 training images [48]. This work demonstrated that limited microscopy datasets can be sufficient when data augmentation is properly leveraged.

A literature review found that 67% of materials science ML datasets contain fewer than 1000 samples, with 57% containing fewer than 500 samples [49]. Recent successes demonstrate effective strategies for limited data. Expert-level segmentation has been achieved by using only 5-8 images per microstructural feature through few-shot learning [50]. Similarly, Mask R-CNN has been used for microstructure classification with limited training data, validating these approaches can succeed in materials science despite data constraints [51],[52].

In fiber analysis specifically, Kurkin et al. demonstrated that 37 real images combined with synthetic augmentation could achieve quality results, validating that augmentation strategies can effectively extend small datasets for fiber segmentation tasks [20]. Building on these successes, important questions for practical implementation include how different segmentation methods affect the resulting distribution measurements and how to maintain fiber integrity across tile boundaries in large-field microscopy.

Recent advances in large-field microscopy segmentation have addressed fundamental challenges in automated fiber characterization. Domain-specific pre-training has proven to be key [53], while few-shot learning has achieved expert-level segmentation with minimal data [54]. Synthetic data generation through physics-informed models has been able to address aforementioned data scarcity in materials science [54]. Tiled-based inference methods enable processing of large images despite GPU memory constraints [55], though maintaining fiber integrity across boundaries of complex scenes remains challenging [56], [57].

Methods that balance computational constraints (image size) with measurement precision while minimizing training data represent critical steps needed for automated characterization of FLD. The objectives of this study are to 1) Compare how different deep learning segmentation approaches (instance vs. semantic) affect fiber length distribution measurements, providing empirical guidance for method selection based on specific characterization needs, 2) Demonstrate the feasibility of automated fiber analysis in large-field tiled microscopy, exploring challenges of maintaining fiber integrity across tile boundaries. 3) Enable rapid process-structure-property mapping for composite material design by achieving measurement speeds suitable for high-throughput characterization.

2. Material and Methods

2.1 Fiber Isolation, Imaging, and Manual Measurement

All fiber isolation and imaging techniques below are referenced from Rhodes et al. [58]. Briefly, this process can be summarized in four parts (Figure 3). First, circular composite coupons that have a diameter at least twice the length of the longest possible fiber are extracted from a final part. Next, the extracted coupon undergoes constrained removal of the matrix material via burn-off or acid digestion. The fibers are then isolated and down-sampled. Care is taken to try to obtain fibers only from the center of the coupon to avoid any that may have been broken during coupon extraction. Finally, the fibers are then dispersed and imaged via optical microscopy.

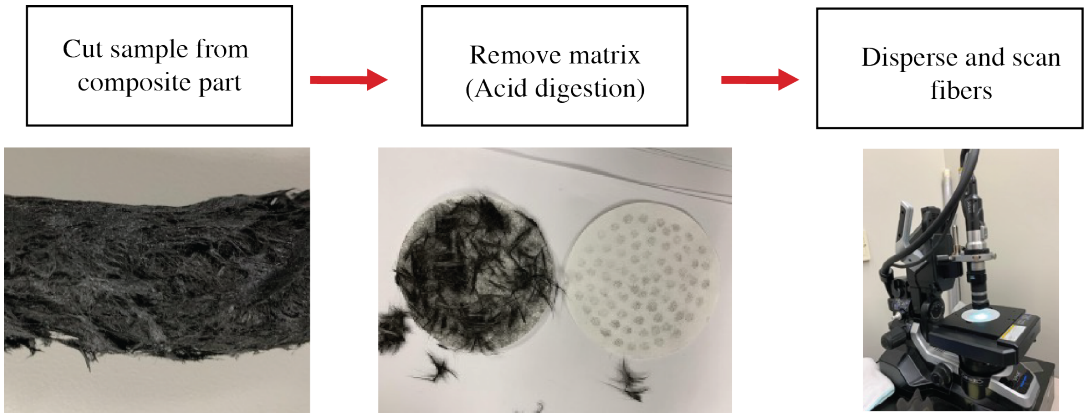

Figure 3: Fiber sample acquisition and imaging process. The material extrudate was obtained and the fibers were isolated by constrained removal of matrix material. These fibers were subsequently dispersed and imaged with optical microscopy. Figure adapted from [58].

Micrographs were obtained using VHX-5000 digital microscope (Keyence Corporation of America, Itasca, IL) at 50x magnification and a polarized lens to enhance contrast and visibility of the fibers. The microscope's automated stitching function was used to capture the entire slide area, generating full-field images of approximately 10,000 x 10,000 pixels at 0.225 pixels / μm. Images were saved in lossless PNG format to preserve detail for segmentation.

The result is a high-resolution image of the extracted fibers. To serve as the ground truth for the remainder of the work, these fibers were manually identified using Image-J. To assess measurement consistency and quality, two independent operators performed manual measurements on the test image. Inter-operator comparison revealed excellent agreement with mean difference of 3.19 μm (1.83%) and length-weighted mean difference of 3.33 μm (1.38%), demonstrating that fiber identification and measurement is reproducible across operators (Supplementary Table S1). A zoomed-in subsection of an image can be seen in Figure 4, A and B. It is important to note that burn-off and acid digestion methods, while standard for fiber length measurement, inherently destroy the polymer matrix and original fiber architecture. This process eliminates spatial orientation information for additional mechanical property prediction.

2.2 Model Training

To train the deep learning models, two custom datasets were generated. One for the instance segmentation method and another for the semantic method. The dataset consisted of the same images.

2.2.1 Training Set Development

First, a dataset was created for the semantic image segmentation process. Images were manually segmented using GIMP (v2.10.34) and represented pixel-wise identification of fibers. The dataset consisted of 29, 1000 x 1000-pixel images (Figure 4A: left). The result is an image tile coupled with a binary mask resulting in 1044 tiles of size 160 x 160 pixels. A separate dataset, consisting of the same images, was generated for the instance segmentation process images (Figure 4A). Images were tiled and bounding boxes and segmentations were manually annotated with CVAT. The result was ~2000 annotated 160 x 160-pixel images representing ~9000 unique fiber instances. The images are coupled with a JSON file describing the annotated information i.e. segmentation and bounding box locations.

All manual annotations were performed by a single experienced operator to ensure consistency across the training dataset. Clear guidelines were established for fiber boundary identification, with ambiguous cases (e.g., touching fibers, partial fibers at image edges) handled according to predetermined rules. While using a single annotator eliminates inter-annotator variability, we validated annotation quality through comparison with independent manual measurements (Supplementary Table S1), which showed excellent inter-operator agreement.

2.2.2 Data Augmentation

Since hand segmentation is labor intensive, few-shot learning is used to break each image into additional training samples, or tiles [59]. This method is powerful for up-sampling manually annotated images to generate a larger dataset. This method of data augmentation is accomplished by modifying samples at random with horizontal and vertical flips, color modification (brightness, contrast, hue, and saturation changes), and Gaussian blurring in hopes of training a more robust model that is ready for image inconsistencies contained throughout the dataset.

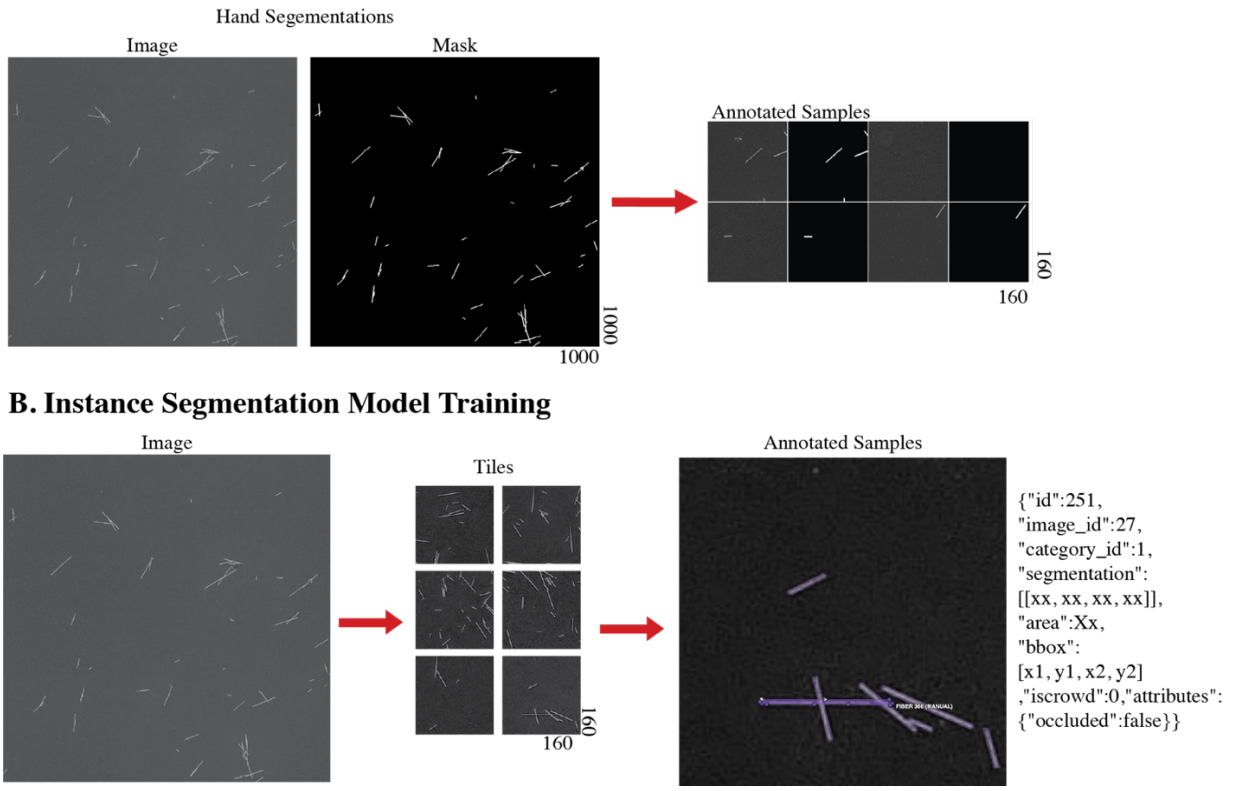

Figure 4: CNN model training for both segmentation techniques. **A.** For semantic segmentation, hand segmentations training examples were tiled and augmented to upsampled the data and generate ~1,000,000 training tiles (160 x 160 pixels). **B.** Individual instances of fibers were identified manually via CVAT resulting in annotated samples with pixel coordinates for each fiber segmentation and bounding box.

2.2.3 Software and Hardware

All code was written in Python 3 and utilized TensorFlow (U-Net, Mask R-CNN) or PyTorch (YOLOv11) as well as other various packages such as OpenCV, and NumPy. Different environments were used to manage software versions for each model through Anaconda (v23.7.4). Hardware configurations were as follows: **OS**: Ubuntu (22.04.3 LTS), **CPU**: Intel Core i9 13900K, **GPU**: 2x Nvidia RTX4090 24 GB, and **Memory**: 96 GB DDR5.

3. Theory

Initially, traditional computer vision methods were evaluated as baseline approaches. While thresholding can segment well-dispersed fibers with good contrast, the methods fail catastrophically in suboptimal images or when fibers overlap (Figure 5). Furthermore, the inability to adapt to varying image quality without manual intervention (retuning) limits automating processes and throughput. These fundamental limitations motivated the adoption of deep learning approaches that maintain fiber identity through overlapping regions and generalize across imaging conditions.

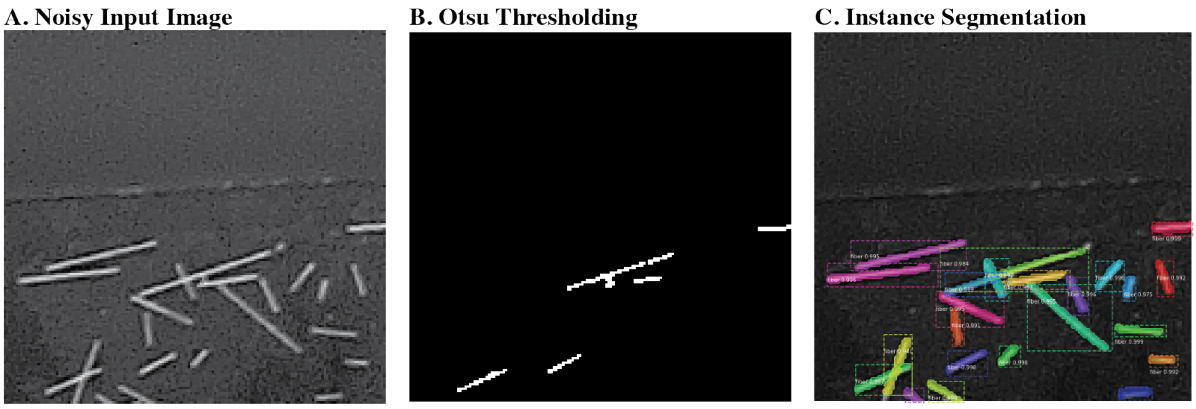

Figure 5: Comparison of traditional segmentation methods versus instance segmentation for fiber detection. **A.** Original microscopy image showing carbon fibers with noise, varying contrast and overlapping regions. **B.** Otsu thresholding highlighting missed and merged fibers. **C.** Instance segmentation output with individual fibers identified by unique colors.

3.1 Deep Learning Models

For this work, three deep learning architectures were chosen that represent fundamentally different approaches to the fiber segmentation challenge. U-Net provides semantic segmentation for maximum throughput, Mask R-CNN offers precise instance delineation for overlapping fibers, and YOLOv11 delivers efficient instance detection without anchor box constraints. Each architecture addresses specific aspects of the large-field microscopy challenge.

3.1.1 U-Net

Since first published in 2015, U-Net has proven to be a premier semantic image segmentation model notably due to the network requiring few images for training. The U-Net structure can be viewed as a U or a cone and can be broken into three important aspects. The first of which is a contracting path where the model funnels downwards using max pooling layers to capture context. As the path contracts, the feature information increases enabling the network to learn increasingly complex features of the image. The network then inversely and symmetrically expands through up-convolution layers to provide precise localization [48]. Recent applications have successfully used U-Net for materials microstructure segmentation and downstream classification [60]. For fiber length measurement, U-Net's pixel-wise classification means tiles are processed independently without tracking individual fibers. This enables rapid processing but cannot distinguish between touching fibers, making it suitable for quality control screening where speed outweighs individual fiber resolution.

3.1.2 Mask R-CNN

Mask R-CNN [61], an instance segmentation model, aims to extend the Faster R-CNN [31] architecture to include segmentations, or masks, for identified objects. Mask R-CNN adopts the same two-stage procedure as the Faster R-CNN. The first stage of Mask R-CNN, the Region Proposal Network (RPN) works to identify and return candidate bounding boxes and is identical to the first stage of Faster R-CNN. Mask R-CNN differs in the second stage, in parallel to predicting both the class label and bounding box, it also outputs a mask for each Region of Interest (RoI). The goal of Mask R-CNN is not necessarily speed, but rather, to be methodical and accurate for precise results in scenes that are highly complex with overlapping objects. The instance-level segmentation is particularly valuable for fiber analysis as it maintains individual fiber identity even

in dense regions. When combined with SAHI, Mask R-CNN can track the same fiber across multiple tiles, preventing double-counting at boundaries which is critical for accurate length distribution measurement.

3.1.3 YOLOv11

In comparison to Mask R-CNN, You Only Look Once (YOLO) is an instance segmentation model designed for speed and real-time applications [62]. The network's goal is the unification of steps for simultaneous segmentation and classification of objects within an image [63]. The architecture can be broken into three parts. The backbone is responsible for feature extraction and is designed to efficiently capture the spatial hierarchies within the input image. Next, there is the neck which further distills the features extracted by the backbone. Finally, the information feeds into a head where the identified features are processed, and the outputs are presented. YOLOv11's anchor-free detection eliminates hyperparameter tuning for fiber aspect ratios, making it robust to the wide range of fiber lengths (50-1000 μm) in our samples. The single-pass architecture balances the speed of semantic segmentation with the fiber preservation of instance methods. The recent release of YOLOv11 aimed to increase object detection accuracy when compared to the former YOLOv8.

3.1.4 Differences between Mask R-CNN and YOLOv11

A major difference between the two architectures is that Mask R-CNN relies on its Region Proposal Network for an anchor-based approach to object detection. The network architecture is constructed in such a manner that for each anchor point, the K number of predefined candidate bounding boxes are evaluated as RoIs (Figure 6, left). For the candidates, scale and aspect ratio are predefined complicating the evaluation of objects of vastly different sizes (i.e., small, long, curved fibers). Despite Mask R-CNN's accuracy, it is inherently slow for complex or dense scenes that consist of overlapping objects [64]. Primarily, this can be attributed to the model's generation of many RoIs. Each anchor has a set number of K candidate bounding boxes that must be processed individually to produce pixel-precise segmentations.

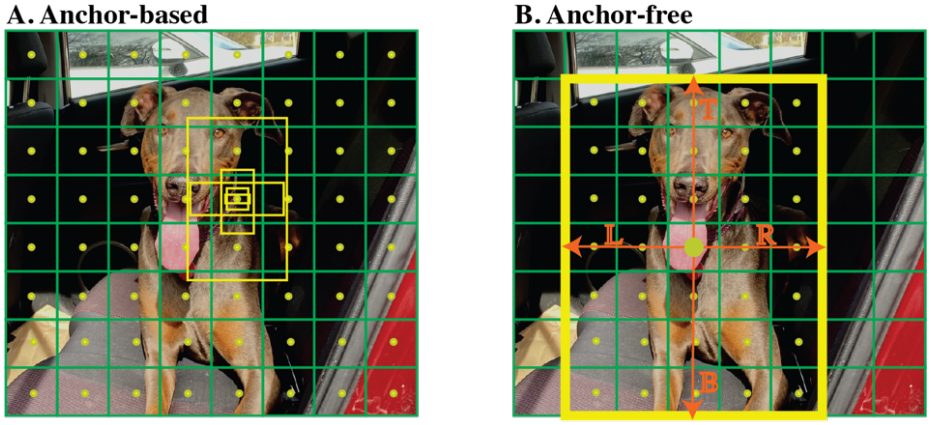

Figure 6: Differing schemes for object detection. Mask R-CNN is an anchor-based method that relies on hyperparameters defining anchor size and aspect ratio. YOLO is an anchor-free allowing for the direct localization of objects.

YOLOv11 is computationally efficient due to its single-pass architecture, which allows the processing of the whole image at one time [62]. In that single pass, the model can quickly identify RoIs through its utilization of anchor-free detection. This method allows for the direct prediction

of the location of objects within an image through the identification of important characteristics in the feature maps. This technique improves efficiency by eliminating complicated computations related to anchor boxes and removing all their hyper-parameters (scale, aspect ratio, etc.) [65]. In return, this allows the production of fewer candidate regions (Figure 6, right). However, it could in turn struggle with small objects since they are represented by fewer pixels, leading to weaker feature representations and subsequent loss of detail in the feature maps.

3.2 Data Processing

Input images can be large and on the order of magnitude of ~100,000,000 pixels. It is essential to break these images up into smaller, more digestible chunks for computational efficiency. A subsequent challenge of this is the recreation of the original image and the maintaining the results.

3.2.1 Image Tiling

Figure 7B & C shows the semantic segmentation dataset being broken into 160 x 160-pixel tiles. The purpose of this is for memory constraints, computational efficiency, and **preservation of local features**. Memory constraints and computational efficiency go hand in hand. Processing large images in a single pass can result in hardware limitations. Smaller chunks can be processed quicker while ensuring that fine-grain details of small objects are preserved rather than being lost in the large input image.

3.2.2 Slicing Aided Hyper Inference

Although tiling is a necessity for the identification of small objects in large images, it does create potential issues when an instance is split across tiles (Figure 7B & C). This is largely an issue for instance segmentation and not semantic segmentation. For the latter, when an image is restitched together the segmentation masks are just merged (Figure 7D). SAHI is designed to aid in the merging of split instances, particularly for small objects [55]. For our application, fibers spanning tile boundaries may be fragmented (Figure 7C), counted multiple times, or missed entirely.

The method is structured to merge overlapping bounding box prediction results back to their original dimensions using Non-Maximum Merging (NMM) visually shown in Figure 7C & E. In this process, bounding boxes with Intersection over Smaller (IoS) ratios exceeding a predetermined threshold are identified and merged. Differing from other bounding box suppression methods that discard bounding boxes not used, NMM retains and combines these overlapping boxes. The end goal is to merge multiple bounding boxes, that represent a single instance, into one (Figure 7E).

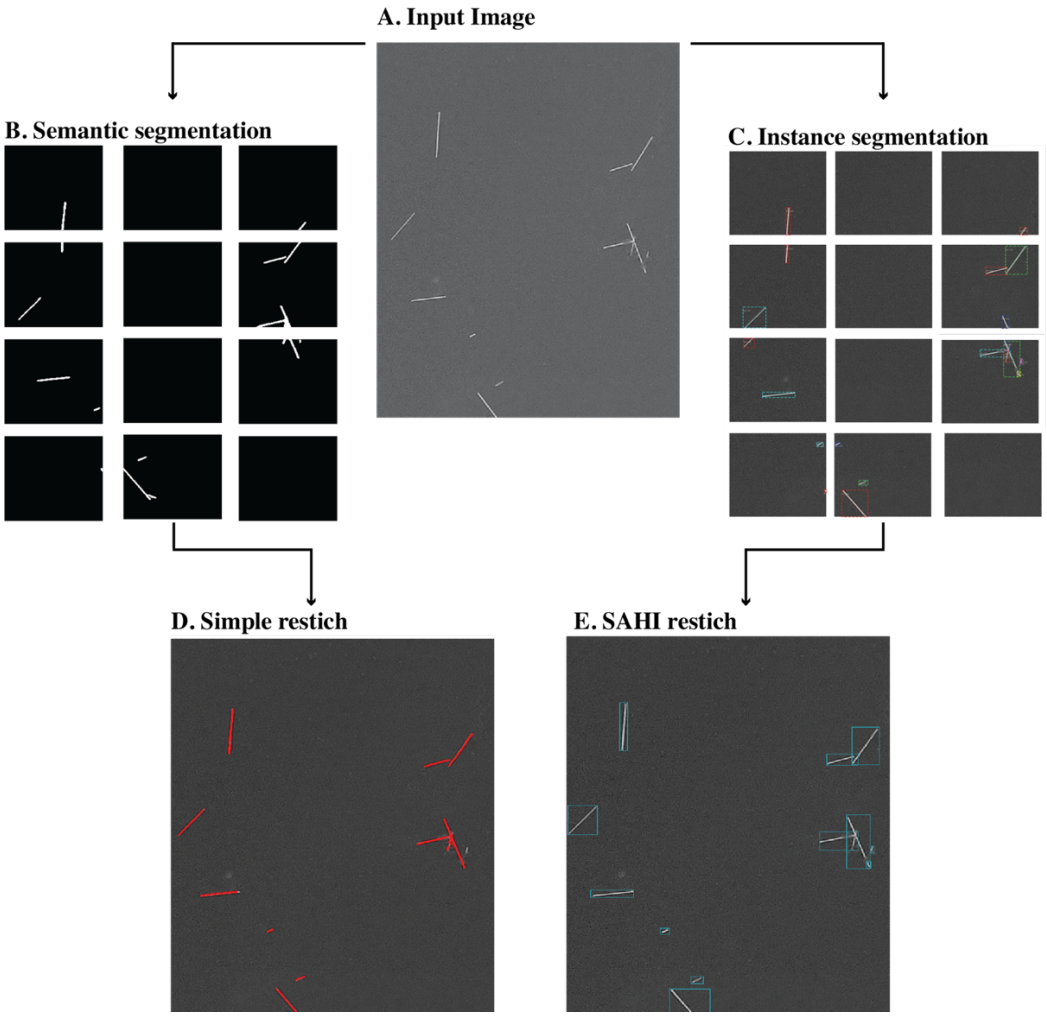

Figure 7: Image processing differences between the semantic and instance segmentation models along with the visualization of the handling of splitting of fiber instances and subsequent merging for each model. An example image is split, and each part is then used for object detection. The smaller tiles are then merged and measured.

3.3 Fiber Length Estimation

Once individual objects (fibers) are identified, they must be measured for the estimation of fiber length. This is accomplished through determining the width (W) and height (H) of the fiber and using the Pythagorean Theorem. Where c is the micron-to-pixel ratio of the image.

$$L = c * \sqrt{H^2 + W^2} \tag{1}$$

3.4 Evaluation Metrics

Assessment of fiber length distribution and CNN segmentation performance is essential before moving forward. Length measurement metrics provide different perspectives on fiber distribution, while the Wasserstein distance quantifies differences between distributions. For CNN output evaluation, DSC is used to measure overlap accuracy between prediction and ground truth.

3.4.1 Length Measurement

3.4.1.1 Arithmetic and Length-Weighted Mean

Both metrics provide a different understanding of the resulting fiber length distribution. Arithmetic mean (L_n) helps understand the presence and impact of the shorter fibers within the distribution. Alternatively, length-weighted mean (L_w) gives more importance to longer fibers [58].

$$L_n = \sum N_i l_i / \sum N_i \tag{2}$$

$$L_w = \sum N_i l_i^2 / \sum N_i l_i \tag{3}$$

3.4.1.2 Wasserstein Distance

Sometimes referred to as the Earth Movers Distance (**EMD**), the metric can be used for the understanding of the amount of 'work' it takes to transform one distribution into another. This is equivalent to finding the area between two different 1-D distributions. The lower the value the less work that is needed. With the assumption that the distributions X and Y are cumulative distribution functions, the metric is represented by the following equation:

$$EMD = \int_{-\infty}^{+\infty} |X - Y| \tag{4}$$

3.4.1.3 Kullback-Leibler Divergence

Kullback-Leibler (KL) divergence quantifies the proximity of two probability distributions or how much the modeled distribution (q) represents the true probability distribution (p).

$$D_{KL}(p \parallel q) = \sum_i p_i \log(p_i / q_1) \tag{5}$$

3.4.1.4 Jensen Shannon Divergence

Jensen Shannon (JS) Divergence represents how similar two distributions are. It is a symmetrized and smoothed version of KL Divergence**.**

$$JSD(p \parallel q) = /2D_{KL}(p \parallel m) + 1/2D_{KL}(q \parallel m) \text{ where } m = 1/2(p + q) \tag{6}$$

3.4.2 CNN Output

3.4.2.1 Sørensen–Dice coefficient

Often referred to as the Dice Similarity Coefficient (DSC), its purpose is to gauge the similarity between two segmentations.

$$DSC = 2 * |X \cap Y| / |X| + |Y| \quad (7)$$

$|X \cap Y|$ represents the common pixels with both images and$|X|$ and $|Y|$ are the counts of the total number of segmented pixels within the image. In the case of this paper, the metric is used to compare the ground truth (hand segmentations) to the CNN outputted results. A DSC of 1 indicates perfect similarity or overlap, whereas a DSC of 0 represents no overlap at all.

4. Results and Discussion

Initial results generated a major question. The simple aspect of the work is measuring single, stand-alone fibers. However, are these fibers representative of the imaged fiber population? Figure 8 presents the initial results from each model.

4.1 Importance of Capturing Overlapping Fibers

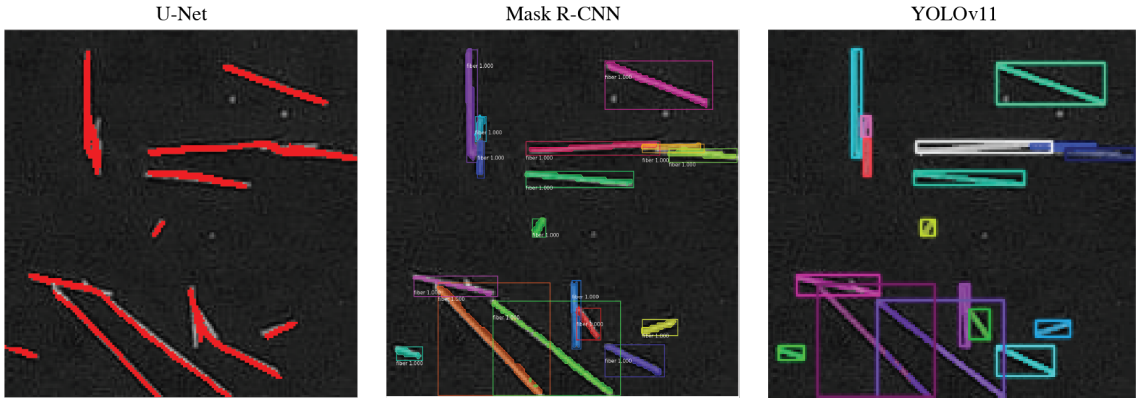

Figure 8: Visual comparison of segmentation models. U-Net simply classifies pixels to belong to either a fiber or background. Where both Mask R-CNN and YOLOv11 work to identify individual instances of the fibers allowing the identification of closely grouped fibers.

From the sample U-Net segmentation (Figure 8, left), it is evident that when fibers are not close to other fibers, the segmentations visually seem reliable. As soon as the fibers become close to each other or overlap, they merge. Initial approaches simply classified identified objects as either a single fiber or overlapping. Subsequently, only the single fibers were measured, and the rest were ignored. Figure 8 provides a rationale that this is a poor approach. The two distributions represent the manual measurement of fibers that were identified as stand-alone, single fibers, or the measurement of all fibers that were classified as overlapping. The distributions notably differ, with the single fibers being much shorter than those in the other distribution. It can be assumed that there is a higher probability that a longer fiber is to overlap other imaged fibers. The inclusion of the overlapping fibers results in a nearly 40% increase in the total number of fibers measured. The two instance segmentation techniques Mask R-CNN (Figure 8, middle) and YOLOv11 (Figure 8, right) highlight their ability to differentiate between fibers even in complex or overlapping orientations.

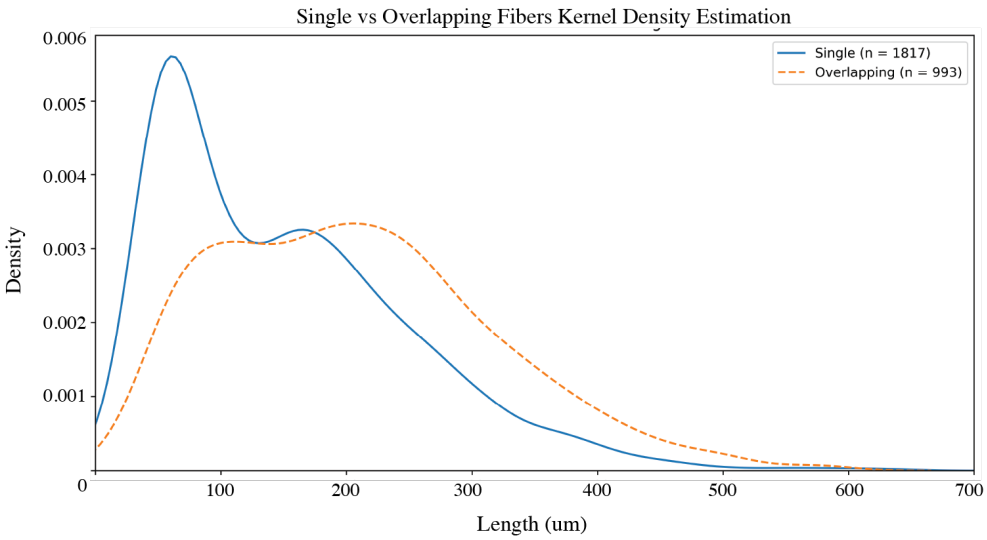

Figure 9: Hand measurement of single and overlapping fibers. Objects identified in the input image were to be either a single fiber or a cluster of fibers. Subsequently, each fiber in both classes was hand measured.

Not only does the measurement of single fibers miss the full picture of the residual fiber length of the part, but initial measurements also using U-Net led to consistent over-measurement of the fiber length for singular fibers. Figure 9 comprises two distinct plots that present comparisons between U-Net-based and manual measurement techniques. The left plot is a scatter plot of the measurement of the same fibers by each method. Although a strong linear relationship is presented, each U-Net measurement is on average 1.5 pixels or 7 μm longer than the hand measurement. This overestimation likely reflects different handling of ambiguous boundary pixels. At the fiber edges where intensity gradually transitions, these faint pixels still have a higher probability of being fiber than background in the model's classification, where human annotators can conservatively exclude such regions. The initial results showed that fiber length could be reasonably estimated for single fibers but highlighted the necessity for capturing all fibers contained in an image and a more precise method.

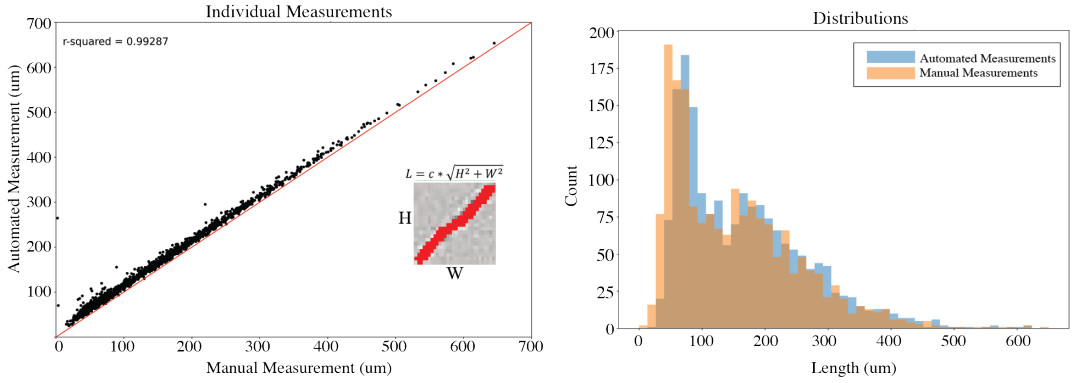

Figure 10: Manual and automated measurements of single fibers. The measurements agree and are consistent. However, automated measurements are consistently larger by 1.5 pixels on average.

4.2 Determination of the Quality of CNN Outputs

Before using outputs for the estimation of fiber length, we evaluated segmentation performance using the Dice Similarity Coefficient (DSC) on a large-field test image. The mean DSC for U-Net,

Mask R-CNN, and YOLOv11 were 0.8926, 0.8324, and 0.7461, respectively. When considering the results, it is important to discuss that the fibers being identified in the images are less than 30 pixels in total. The denominator of the DSC calculation is the total number of pixels identified with both the ground truth and CNN output. If the CNN misclassifies even a few pixels that can make a notable impact on the output. While U-Net achieved the highest pixel-level accuracy, its semantic segmentation approach cannot separate touching fibers, fundamentally limiting its utility for fiber measurement.

Additional evaluation metrics were assessed using a held-out validation set of annotated tiles nor used during training (Supplementary Table S3). YOLOv11 demonstrated balanced precision-recall (0.82/0.82) while Mask R-CNN showed high recall (0.91) but low precision (0.68), possibly indicating boundary overestimation. These findings highlight a fundamental difference between pixel-level accuracy vs measurement distribution accuracy. YOLOv11 achieves the lowest DSC yet produces the most accurate fiber length distribution, revealing that segmentation metrics do not accurately predict measurement quality. DSC weights all pixels equally, but length measurements require only accurate boundaries. While missing 14% of fibers (primarily small ones) reduces DSC, these omissions minimally affect the weighted mean where longer fibers dominate quadratically (Table 1).

Method	N	Mean ± SD (μm)	95% CI	Weighted Mean (μm)	EMD (μm)	CV (%)	Time (s)
Manual	2810	170.9 ± 107.0	(167.0, 174.0)	237.9	-	62.6	10,020
U-Net	2139	155.9 ± 103.6	(151.5, 160.3)	224.7	16.34	66.5	3.92
Mask R-CNN	3008	174.9 ± 110.4	(170.9, 178.8)	244.6	5.78	63.1	448
YOLOv11	2409	164.8 ± 110.1	(160.4, 169.2)	238.4	6.77	66.8	268

Table 1: Summary of the Measurement Results. Fiber length is compared against the manual measurements via evaluation time, mean fiber length, and mean weighted fiber length. Note, 5 μm is equivalent to 0.75 pixels. EMD: Earth Mover's Distance (Wasserstein Distance), CV%: Coefficient of Variation.

4.3 Measurement Comparison

All models were compared using a test image of size 10,286 x 10,092 pixels, which was completely held out from training and validation. Visually, all presented distributions in Figure 11B are of the same shape. However, each automated measurement technique resulted in subtle differences. The mean and weighted fiber length average (μm) for each model are presented in Table 1. The Mask R-CNN and YOLO estimates behave the most like the manual measurements based on these summary statistics.

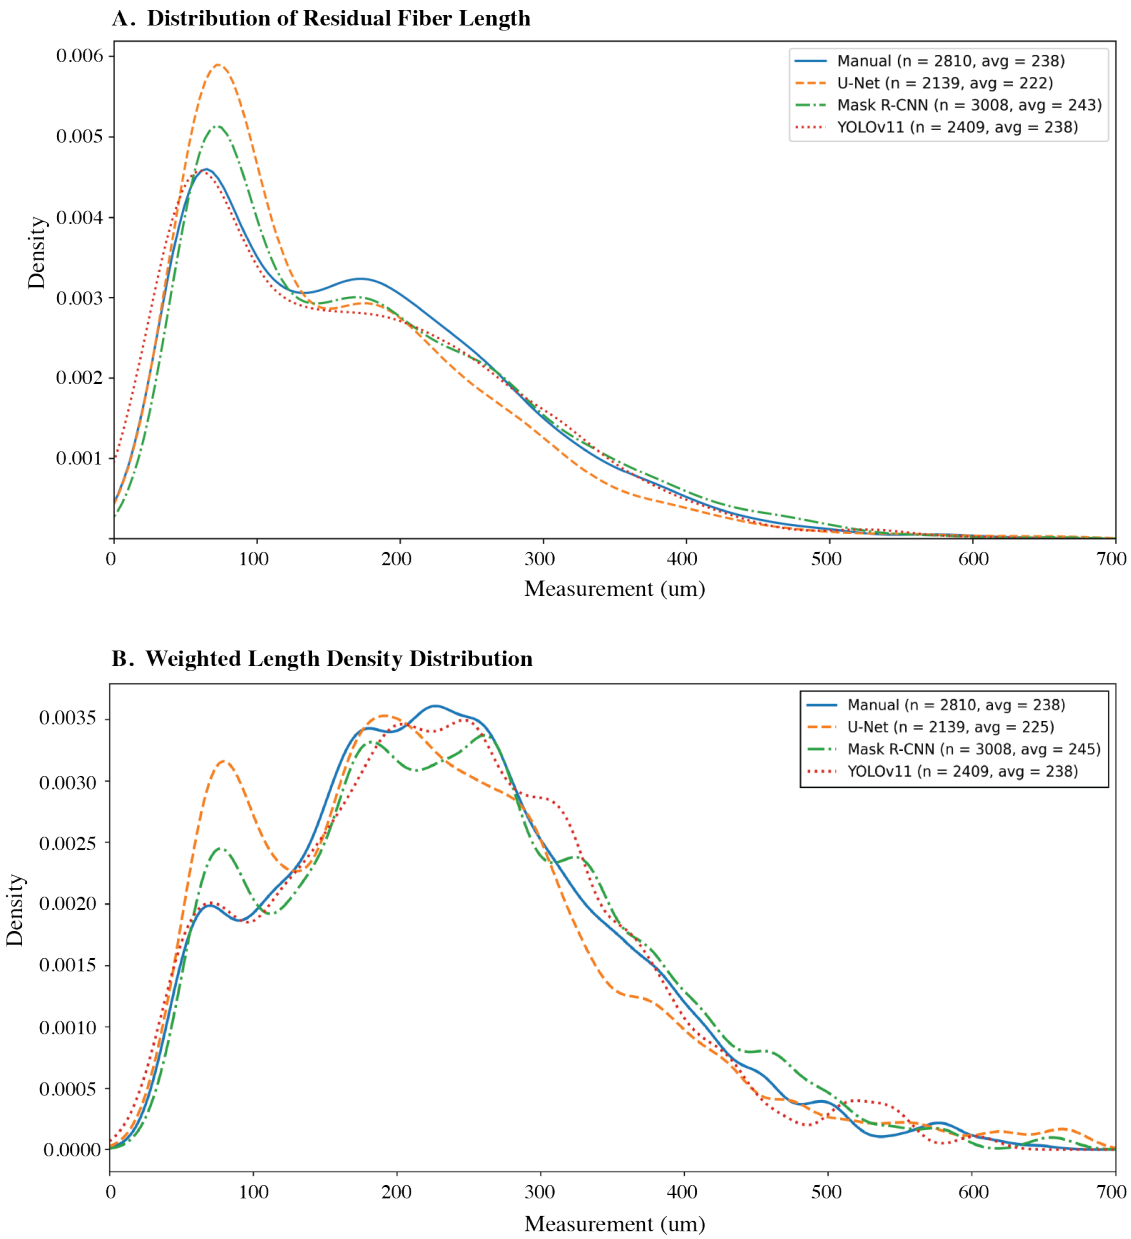

Figure 11: Comparison of fiber length distributions from manual measurements and automated methods. A. Kernel density estimation shows Mask R-CNN over-detects short fibers, all models underrepresent 150-250 μm, and U-Net missed long fibers. B. Length-weighted kernel densities highlight model-specific biases.

Upon inspection of the residual fiber length distributions in Figure 11A, it is evident that measurements obtained from the U-Net segmentations have far fewer values greater than 200 μm as compared to all other measurement methods. This can be attributed to the fact that only fibers classified as stand-alone were measured reliably. Mask R-CNN was able to capture the short fibers, but the longer fibers as well. The length-weighted distribution, where each fiber contributes proportionally to its length, provides insight into how different measurement techniques affect the assessment of total fiber content. The analysis of length-weighted residual fiber, Figure 11B, shows that Mask R-CNN demonstrates enhanced sensitivity to short fibers, evident from the elevated density peak below 100 μm. While this captures a population potentially missed by other methods,

it may indicate false positive detections at small length scales. The model's ability to detect both short and long fibers results in a broader distribution, though with a notable depression between 175 - 300 μm. Within this range and below 100 μm, YOLOv11 shows great agreement with manual measurement. Table 1 shows how the measurements differ between models. Interestingly, when compared to manual measurement, Mask R-CNN mean fiber length is most like the manual measurement. However, YOLO's weighted fiber length is nearly identical to the manual measurements. Further quantitative results are needed to evaluate the outputs.

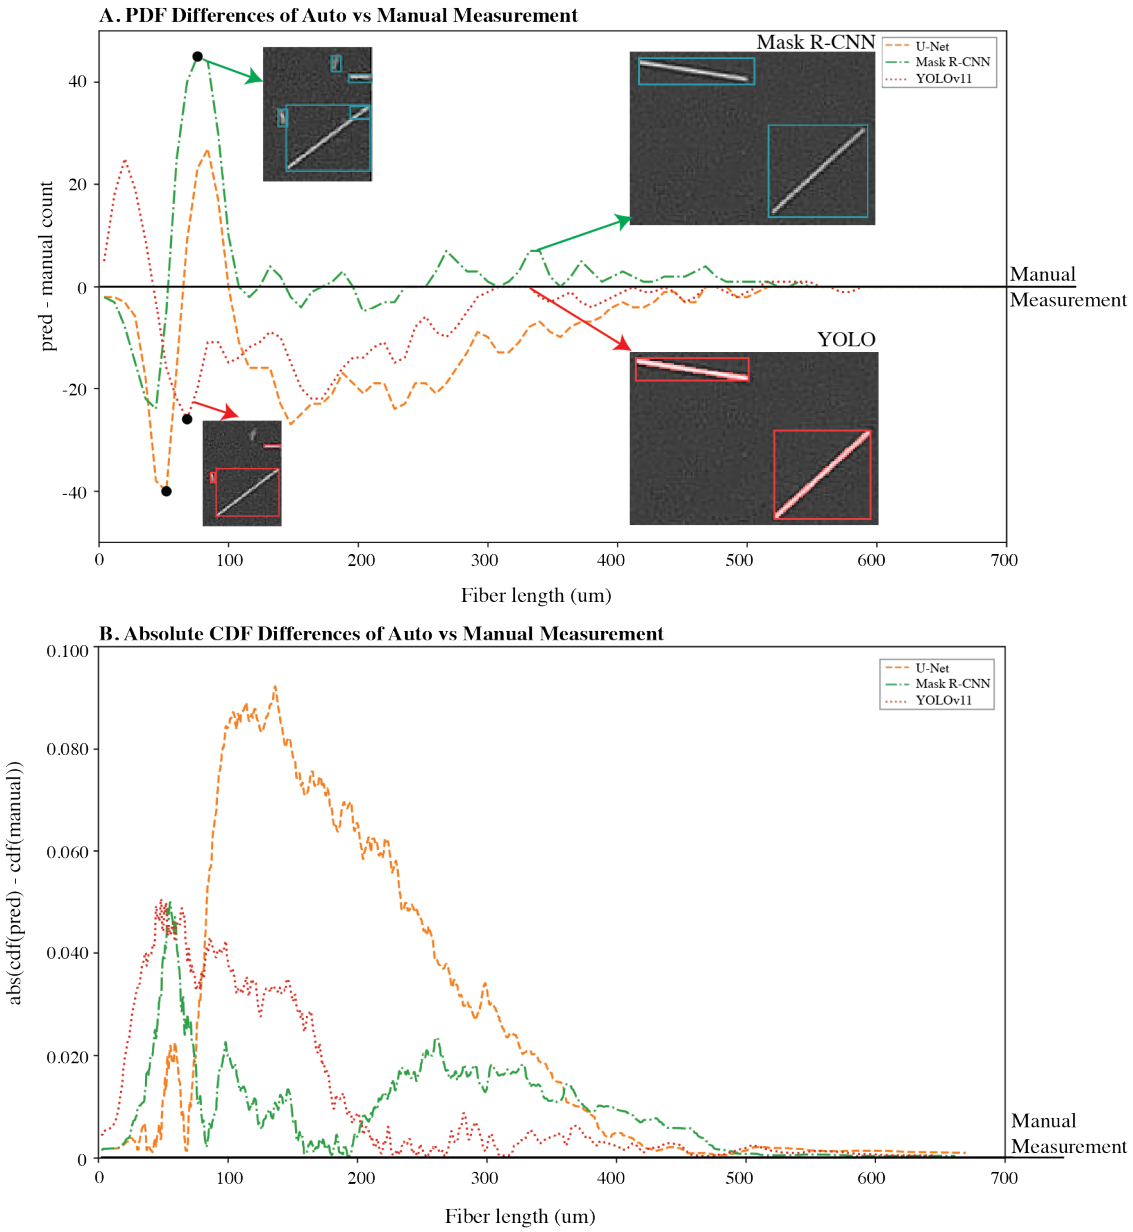

Figure 12: PDF and CDF difference plots comparing each automated method to manual measurements. The horizontal black line at y = 0 denotes manual baseline. A. PDF differences showing over- and under- predictions across the length ranges, with insets illustrating representative examples for YOLO and Mask R-CNN segmentations. B. Absolute CDF differences quantifying divergence from manual distributions.

Each model was compared against manual measurements using distributional metrics and measurement time. As shown by the previous discussion, U-Net exhibited the largest deviation from ground truth, showing strong negative bias, compressed variance, and the highest CDF difference across the distribution. In contrast, Mask R-CNN and YOLO performed similarly well, each maintaining low divergence from the manual distribution and closely matching the weighted fiber length average. While Mask R-CNN achieved the best overall accuracy, YOLO provided comparable results with significantly faster runtime, suggesting an attractive trade-off between speed and precision. All quantitative numerical findings are summarized in Table 2.

A kernel density estimate (KDE) for the three CNN-measured FLD's subtracted from the manual measurement's KDE further reveals characteristics of CNN model performances at different fiber lengths as shown in Figure 12A. This plot aims to compare the three auto-measured FLD's and manually measured FLD without complicated histogram overlays. At first glance, this plot very simply reveals from the large alternating peaks from all three models at fiber lengths < 100 μm that shorter fibers are much more prone to over or under-representation. The most straightforward explanation would be that fibers < 100 μm are simply being missed by the models or that the model is measuring fibers that are not present. The potential for these false negatives and false positives was analyzed qualitatively by observing the output images from each of the three models. Upon inspection of the output images, all 3 models were found to occasionally miss fiber entirely, resulting in false negatives. This occurrence was noticed to be inversely proportional to fiber length, where longer fibers were not observed to have missed measurement, but small to medium-length fibers were missed at higher frequencies. False positives on the other hand were less likely as they were rarely observed with Mask R-CNN or YOLO; however, U-Net would frequently misattribute pixels to small fibers that the human annotators ignored during measurement.

Upon further qualitative inspection of the images after processing, other effects were also discovered that would contribute to a mismatch between the auto-measured and manual FLD's. As shown in the fiber images in Figure 12A (top right), the Mask R-CNN was found to occasionally overestimate the bounding box dimensions for fibers at larger lengths. In contrast, YOLO was rarely found to misapply the bounding boxes to the fibers, which leads to more accurate length estimates. This mismatch in accuracy at higher lengths is supported in Figure 12A where it can be observed that at fiber lengths past 300 μm, Mask R-CNN's difference is consistently farther from the manual measurement line than YOLO and never dips below the manual prediction line. This hypothesis is also supported by noting that Mask R-CNN's highest measurements exceed that of YOLO's as shown in Fig 12A. Additionally, restitching errors were also found to occur with both Mask R-CNN and YOLO where a single fiber was measured multiple times in shorter segments (Figure 12A, top left). This finding could help to explain the overrepresentation of shorter fibers present for each of the three CNN models; however, more analysis is required to more precisely understand the effect of this issue. Future work could also include a larger parameter search with SAHI as this impacts restitching results.

These effects render direct comparison between FLD's for drawing strong conclusions on performance at certain fiber lengths difficult, but other statistics allow for interpretability. The CDF difference plot (Figure 12B) reaffirmed that U-Net diverged substantially from manual measurements, with a maximum CDF difference between in the 100–150 μm fiber length range. This indicates a persistent accumulation error in mid-to-long fibers. Meanwhile, both Mask and

YOLO demonstrated significantly better alignment with the manual CDF, with nearly identical max differences (~0.05) and better performance across long-tail regions. Mask showed tighter cumulative agreement throughout measurements <200 μm, while YOLO appeared to better represent long fiber behavior beyond 200 μm. The intersection at 200 μm shows that YOLOv11's lower fiber count (due to missed detections) is balanced by its high accuracy on the fibers it does capture, while Mask R-CNN is catching more fibers but at the cost of slight over-predictions in that length range. This is further highlighted in Table 1.

Model	Evaluation Time (s)	Variance Diff	Max CDF Diff	Mean Abs CDF Diff	EMD	KL Divergence	JS Divergence
U-Net	3.92	-720.7	0.0936	0.0468	16.344	0.1729	0.1410
Mask R	448	737.5	0.0531	0.0136	5.7895	0.1183	0.1174
YOLOv11	268	677.3	0.05412	0.0192	6.828	0.1513	0.1364

Table 2: Overview of quantitative analyses of fiber measurement distributions. CDF: Cumulative Distribution Function, EMD: Earth Mover's Distance (Wasserstein Distance), KL Divergence: Kullback-Leibler Divergence, Jensen-Shannon Divergence.

U-Net underpredicted fiber lengths significantly. Its low variance suggests a poor representation of fiber length variability compared to manual measurements, which is unsurprising due to capturing only stand-alone fibers. Both Mask R-CNN and YOLO produced much lower bias, with Mask R-CNN slightly overpredicting the mean and YOLO slightly underpredicting. In terms of distribution shape, Mask R-CNN had the lowest mean absolute CDF difference, JS divergence, KL divergence.

Evaluation time is also a key consideration for practical implementation. Each model evaluated the 100,000,000-pixel test image 50 times, and the average time is presented. U-Net demonstrated the fastest processing time at ~4 seconds per image, compared to 268s for YOLO and 448s for Mask R-CNN (Table 2). Given that YOLO's accuracy metrics **closely** match those of Mask R-CNN, it may offer a more scalable solution when high throughput is needed. However, although Mask R-CNN is slower it remains the top performer where time is less critical.

4.4 Implications for Materials Design

Accurate FLD measurement is fundamental to establish structure-property-processing relationships in discontinuous fiber composites. Efficient processes to quickly characterize fiber length distributions could speed up rapid prototyping and part quantification efforts. Current methods, requiring >2 hours per 2000 fibers, create a non-negligible bottleneck in this workflow. While commercial systems (FASEP) achieve <1% measurement error, they require specialized hardware. The rapid automated methods in the presented work achieve comparable accuracy (0.2% weighted mean error for YOLOv11) using standard microscopy images and completing measurements rapidly.

More fundamentally, automated FLD measurements enable a paradigm shift from trial-and-error approaches to data-driven design methodologies. Established composite theories including the Halpin-Tsai equations and modified rule of mixtures require FLDs to predict modulus and strength. Recent advances have demonstrated that CNN-based microstructural measurement methods can directly inform materials design and development [34], [35], [66]. The automated FLD methods described provide high-throughput microstructural data needed for inverse design workflows and

process optimization in discontinuous fiber composites. This enables statistical models that could connect extrusion conditions to FLDs and resulting final part properties.

5. Conclusions and Perspectives

This study has demonstrated the efficacy of deep learning-based techniques in automating the measurement of fiber length distribution in discontinuous fiber composites. Advanced image processing algorithms enabled the reduction of the time associated with manual measurement methods while simultaneously increasing the consistency and repeatability of the results. Of the models compared, the instance segmentation methods excelled in capturing an accurate representation of the residual fiber length due to whole domain understanding. This can be attributed to the ability to capture overlapping instances, allowing for a more comprehensive representation of the part due to capturing ~40 percent more fibers. Given the improved accuracy, efficiency, and scalability of these deep learning models, demonstrated in our comparative study, the results suggest that this approach has the potential to become a promising alternative method for fiber length measurement, reducing reliance on labor-intensive manual techniques or costs associated with automated methods.

However, important limitations constrain current applicability. Our validation is limited to carbon fibers with models trained on 29 images under consistent imaging conditions. Generalization to other fiber types (glass, natural), requires consideration of their distinct optical properties. For example, glass fibers present unique challenges such as their translucency creating inconsistent contrast depending on backing material. Current measurement approaches assume straight fibers of consistent width, limiting accuracy for curved fibers. Different preparation methods (burn off vs. acid digestion) may leave varying residues, potentially degrading model performance. Additionally, variations in microscopy settings (magnification changes, filters, illumination) or differing image modalities would require additional model fine-tuning. Ongoing work is expanding material system applicability through new architectures [67] and addresses curved fiber measurement through skeleton-based algorithms that track centerlines.

Future deployments will include systematic validation through cross-material testing, synthetic data generation for training diversity, and transfer learning studies to quantify adaptation requirements for new systems. Furthermore, the burn-off process fundamentally limits our method to length distribution measurement, destroying the orientation information essential for complete mechanical property prediction. Yet, fiber length distribution remains a critical quality control parameter that directly impacts strength, toughness, and processability.

Despite these limitations, our methodological contributions, demonstrating how segmentation choice affects distribution measurements and implementing SAHI for tile boundary handling, provide insights that extend beyond our specific validation case. The inverse DSC-measurement accuracy relationship demonstrates that materials characterization requires measurement-specific metrics, not computer vision benchmarks. The consistent performance differences between instance and semantic segmentation (speed versus distribution fidelity) appear fundamental rather than dataset-specific, though this requires validation across multiple material systems. As a result, the methods provide a promising alternative to established residual fiber length measurement protocols. A summary and future directions for the work are provided below:

- No automated measurement methods are perfect - these techniques are just estimations of the physical information represented by the images.

- **Computational efficiency:** U-Net 2500x faster, YOLOv11 40x, and Mask R-CNN 22x compared to manual measurements

- **Tile boundary handling:** Implementation of SAHI proved critical for maintaining tile fiber integrity at boundary conditions to avoid fibers being split and measured as multiple short fragments

- **Measurement accuracy:** YOLOv11 based segmentation resulted in optimal weighted mean accuracy (238.4 vs 237.9, 0.2% error), critical for property prediction, while Mask R-CNN showed best distribution matching (EMD = 5.78)

- **Method selection guidance:** U-Net provides rapid screening of dispersed fibers; YOLOv11 optimal for property critical applications, Mask R-CNN best for whole distribution analysis

- **Practical impact:** Deep learning-based methods reduce processing time from hours to minutes. These automated methods enable quality control and processing understanding without specialized measurement hardware

- Potential future applications include:
    - Non-destructive evaluation of final parts using imaging methods such as micro-CT
    - 3D segmentation methods for advanced defect detection
    - In-situ process monitoring of the SMC process, aiding in downstream fiber dispersion and orientation control

6. References

[1] C. E. Duty *et al.*, 'Structure and mechanical behavior of Big Area Additive Manufacturing (BAAM) materials', *Rapid Prototyp J*, vol. 23, no. 1, pp. 181–189, 2017, doi: 10.1108/RPJ-12-2015-0183/FULL/XML.

[2] D. K. Rajak *et al.*, 'Fiber-Reinforced Polymer Composites: Manufacturing, Properties, and Applications', *Polymers 2019, Vol. 11, Page 1667*, vol. 11, no. 10, p. 1667, Oct. 2019, doi: 10.3390/POLYM11101667.

[3] M. Y. Khalid *et al.*, '3D printing of active mechanical metamaterials: A critical review', *Mater Des*, vol. 246, p. 113305, Oct. 2024, doi: 10.1016/J.MATDES.2024.113305.

[4] A. Nazir *et al.*, 'Multi-material additive manufacturing: A systematic review of design, properties, applications, challenges, and 3D printing of materials and cellular metamaterials', *Mater Des*, vol. 226, p. 111661, Feb. 2023, doi: 10.1016/J.MATDES.2023.111661.

[5] H. Ning *et al.*, 'A review of Long fibre thermoplastic (LFT) composites', *Intl Mat Rev*, vol. 65, no. 3, pp. 164–188, Apr. 2020, doi: 10.1080/09506608.2019.1585004

[6] S. R. Doshi and J. -M Charrier, 'A simple illustration of structure-properties relationships for short fiber-reinforced thermoplastics', *Polym Compos*, vol. 10, no. 1, pp. 28–38, 1989, doi: 10.1002/pc.750100105.

[7] D. M. Sánchez *et al.*, 'Development of carbon fiber acrylonitrile styrene acrylate composite for large format additive manufacturing', *Mater Des*, vol. 191, p. 108577, Jun. 2020, doi: 10.1016/J.MATDES.2020.108577.

[8] H. J. Wolf, 'Screw plasticating of discontinuous fiber filled thermoplastic: Mechanisms and prevention of fiber attrition', *Polym Compos*, vol. 15, no. 5, pp. 375–383, Oct. 1994, doi: 10.1002/PC.750150508.

[9] K. Ramani *et al.*, 'Effect of screw design on fiber damage in extrusion compounding and composite properties', *Polym Compos*, vol. 16, no. 3, pp. 258–266, Jun. 1995, doi: 10.1002/PC.750160310.

[10] R. W. Gray IV *et al.*, 'Effects of processing conditions on short TLCP fiber reinforced FDM parts', *Rapid Prototyp J*, vol. 4, no. 1, pp. 14–25, 1998, doi: 10.1108/13552549810197514/FULL/XML.

[11] A. Inoue *et al.*, 'Effect of screw design on fiber breakage and dispersion in injection-molded long glass-fiber-reinforced polypropylene', *http://dx.doi.org/10.1177/0021998313514872*, vol. 49, no. 1, pp. 75–84, Dec. 2013, doi: 10.1177/0021998313514872.

[12] Y. Abderrafai *et al.*, 'Additive manufacturing of short carbon fiber-reinforced polyamide composites by fused filament fabrication: Formulation, manufacturing and characterization', *Mater Des*, vol. 214, p. 110358, Feb. 2022, doi: 10.1016/J.MATDES.2021.110358.

[13] Y. Hu *et al.*, 'Carbon fibre damage during 3D printing of polymer matrix laminates using the FDM process', *Mater Des*, vol. 205, p. 109679, Jul. 2021, doi: 10.1016/J.MATDES.2021.109679.

[14] A. P. Rhodes, 'Correlating Large-Format Additive Manufacturing Processing Correlating Large-Format Additive Manufacturing Processing Parameters to Fiber Length and the Mechanical Performance of Parameters to Fiber Length and the Mechanical Performance of Reinforced Polymer Composites Reinforced Polymer Composites', Accessed: Mar. 04, 2024. [Online]. Available: https://trace.tennessee.edu/utk_gradthes

[15] S. Y. Fu and B. Lauke, 'Effects of fiber length and fiber orientation distributions on the tensile strength of short-fiber-reinforced polymers', *Compos Sci Technol*, vol. 56, no. 10, pp. 1179–1190, Jan. 1996, doi: 10.1016/S0266-3538(96)00072-3.

[16] M. Frei and F. E. Kruis, 'FibeR-CNN: Expanding Mask R-CNN to improve image-based fiber analysis', *Powder Technol*, vol. 377, pp. 974–991, Jan. 2021, doi: 10.1016/J.POWTEC.2020.08.034.

[17] J. S. Bredfeldt *et al.*, 'Computational segmentation of collagen fibers from second-harmonic generation images of breast cancer', *J Biomed Opt*, vol. 19, no. 1, p. 016007, Jan. 2014, doi: 10.1117/1.JBO.19.1.016007.

[18] S. Goris *et al.*, 'A novel fiber length measurement technique for discontinuous fiber-reinforced composites: A comparative study with existing methods', *Polym Compos*, vol. 39, no. 11, pp. 4058–4070, Nov. 2018, doi: 10.1002/PC.24466.

[19] 'Fiber Length Analysis - FASEP Fiber Length Distribution'. Accessed: Dec. 10, 2023. [Online]. Available: http://www.fasep.biz/Service.html

[20] E. Kurkin *et al.*, 'Computer Vision Technology for Short Fiber Segmentation and Measurement in Scanning Electron Microscopy Images', *Technologies 2024, Vol. 12, Page 249*, vol. 12, no. 12, p. 249, Nov. 2024, doi: 10.3390/TECHNOLOGIES12120249.

[21] S. Medghalchi *et al.*, 'Automated segmentation of large image datasets using artificial intelligence for microstructure characterisation and damage analysis', *Mater Des*, vol. 243, p. 113031, Jul. 2024, doi: 10.1016/J.MATDES.2024.113031.

[22] M. J. Emerson *et al.*, 'Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials', *Compos Part A Appl Sci Manuf*, vol. 97, pp. 83–92, Jun. 2017, doi: 10.1016/J.COMPOSITESA.2016.12.028.

[23] A. Badran *et al.*, 'Validation of Deep Learning Segmentation of CT Images of Fiber-Reinforced Composites', *J Comp Sci, Vol. 6, Page 60*, vol. 6, no. 2, p. 60, Feb. 2022, doi: 10.3390/JCS6020060.

[24] A. Badran *et al.*, 'Automated segmentation of computed tomography images of fiber-reinforced composites by deep learning', *J Mater Sci*, vol. 55, no. 34, pp. 16273–16289, Dec. 2020, doi: 10.1007/S10853-020-05148-7/FIGURES/7.

[25] 'Test Methods for Constituent Content of Composite Materials', Mar. 2022, doi: 10.1520/D3171-22.

[26] J. Schmidhuber, 'Deep learning in neural networks: An overview', *Neural Networks*, vol. 61, pp. 85–117, 2015, doi: 10.1016/j.neunet.2014.09.003.

[27] K. Simonyan and A. Zisserman, 'Very Deep Convolutional Networks for Large-Scale Image Recognition', *ICLR 2015 - Conference Track Proceedings*, Sep. 2014, Accessed: Mar. 20, 2024. [Online]. Available: https://arxiv.org/abs/1409.1556v6

[28] K. He *et al.*, 'Deep Residual Learning for Image Recognition', *Proceedings of the IEEE CVPR*, vol. 2016-December, pp. 770–778, Dec. 2015, doi: 10.1109/CVPR.2016.90.

[29] E. Shelhamer *et al.*, 'Fully Convolutional Networks for Semantic Segmentation', *IEEE Trans Pattern Anal Mach Intell*, vol. 39, no. 4, pp. 640–651, Nov. 2014, doi: 10.1109/TPAMI.2016.2572683.

[30] T. He *et al.*, 'Knowledge Adaptation for Efficient Semantic Segmentation', *Proceedings of the IEEE CVPR*, vol. 2019-June, pp. 578–587, Mar. 2019, doi: 10.1109/CVPR.2019.00067.

[31] S. Ren *et al.*, 'Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks', *IEEE Trans Pattern Anal Mach Intell*, vol. 39, no. 6, pp. 1137–1149, Jun. 2015, doi: 10.1109/TPAMI.2016.2577031.

[32] W. Gu *et al.*, 'A review on 2D instance segmentation based on deep neural networks', *Image Vis Comput*, vol. 120, p. 104401, Apr. 2022, doi: 10.1016/J.IMAVIS.2022.104401.

[33] T. Nasrin *et al.*, 'Application of machine learning in polymer additive manufacturing: A review', *J Polym Sci*, 2023, doi: 10.1002/POL.20230649.

[34] Y. Tu *et al.*, 'Towards an instant structure-property prediction quality control tool for additive manufactured steel using a crystal plasticity trained deep learning surrogate', *Mater Des*, vol. 213, p. 110345, Jan. 2022, doi: 10.1016/J.MATDES.2021.110345.

[35] C. Yang *et al.*, 'Prediction of composite microstructure stress-strain curves using convolutional neural networks', *Mater Des*, vol. 189, p. 108509, Apr. 2020, doi: 10.1016/J.MATDES.2020.108509.

[36] I. C. Weber *et al.*, 'Process defects and in situ monitoring methods in metal powder bed fusion: a review', *Meas Sci Technol*, vol. 28, no. 4, p. 044005, Feb. 2017, doi: 10.1088/1361-6501/AA5C4F.

[37] G. D. Goh *et al.*, 'Anomaly Detection in Fused Filament Fabrication Using Machine Learning', *3D Print Addit Manuf*, vol. 10, no. 3, pp. 428–437, Jun. 2023, doi: 10.1089/3DP.2021.0231.

[38] L. Lu *et al.*, 'Deep learning-assisted real-time defect detection and closed-loop adjustment for additive manufacturing of continuous fiber-reinforced polymer composites', *Robot Comput Integr Manuf*, vol. 79, p. 102431, Feb. 2023, doi: 10.1016/J.RCIM.2022.102431.

[39] S. Meister *et al.*, 'Investigations on Explainable Artificial Intelligence methods for the deep learning classification of fibre layup defect in the automated composite manufacturing', *Compos B Eng*, vol. 224, p. 109160, Nov. 2021, doi: 10.1016/J.COMPOSITESB.2021.109160.

[40] H. S. Kim *et al.*, 'Non-destructive detection of thin micro-defects in glass reinforced polymer composites using a terahertz electro–magnetic wave based on a convolution neural network', *Compos B Eng*, vol. 257, p. 110694, May 2023, doi: 10.1016/J.COMPOSITESB.2023.110694.

[41] N. Saeed *et al.*, 'Automatic defects detection in CFRP thermograms, using convolutional neural networks and transfer learning', *Infrared Phys Technol*, vol. 102, p. 103048, Nov. 2019, doi: 10.1016/J.INFRARED.2019.103048.

[42] A. Khan *et al.*, 'Structural vibration-based classification and prediction of delamination in smart composite laminates using deep learning neural network', *Compos B Eng*, vol. 161, pp. 586–594, Mar. 2019, doi: 10.1016/J.COMPOSITESB.2018.12.118.

[43] D. Cristiani *et al.*, 'Strain-based delamination prediction in fatigue loaded CFRP coupon specimens by deep learning and static loading data', *Compos B Eng*, vol. 241, p. 110020, Jul. 2022, doi: 10.1016/J.COMPOSITESB.2022.110020.

[44] G. Ammasai Sengodan 'Prediction of two-phase composite microstructure properties through deep learning of reduced dimensional structure-response data', *Compos B Eng*, vol. 225, p. 109282, Nov. 2021, doi: 10.1016/J.COMPOSITESB.2021.109282.

[45] M. Laurenti *et al.*, 'Predicting the mechanical behavior in FDM printing of biopolymers through boosting artificial neural networks', *Mater Des*, vol. 257, p. 114475, Sep. 2025, doi: 10.1016/J.MATDES.2025.114475.

[46] J. H. S. Almeida and G. F. Gomes, 'Optimising 3D-printed carbon fibre composites using machine learning: Balancing strength and efficiency', *Mater Des*, vol. 256, p. 114325, Aug. 2025, doi: 10.1016/J.MATDES.2025.114325.

[47] A. I. Saimon *et al.*, 'Advancing additive manufacturing through deep learning: A comprehensive review of current progress and future challenges', *IISE Trans*, 2025, doi: 10.1080/24725854.2024.2443592.

[48] W. Weng and X. Zhu, 'U-Net: Convolutional Networks for Biomedical Image Segmentation', *IEEE Access*, vol. 9, pp. 16591–16603, May 2015, doi: 10.1109/ACCESS.2021.3053408.

[49] Y. Liu *et al.*, 'Data quantity governance for machine learning in materials science', *Natl Sci Rev*, vol. 10, no. 7, May 2023, doi: 10.1093/NSR/NWAD125.
[50] S. Akers *et al.*, 'Rapid and flexible segmentation of electron microscopy data using few-shot machine learning', *npj Computational Materials 2021 7:1*, vol. 7, no. 1, pp. 1–9, Nov. 2021, doi: 10.1038/s41524-021-00652-z.
[51] K. Wijesinghe *et al.*, 'Characterization of microscopic deformation of materials using deep learning algorithms', *Mater Des*, vol. 208, p. 109926, Oct. 2021, doi: 10.1016/J.MATDES.2021.109926.
[52] X. Wang *et al.*, 'Quantitative assessment of cement bridges and voids in cement-stabilized permeable base materials using a mask R-CNN-based CT image segmentation strategy', *Mater Des*, vol. 241, p. 112907, May 2024, doi: 10.1016/J.MATDES.2024.112907.
[53] J. Stuckner *et al.*, 'Microstructure segmentation with deep learning encoders pre-trained on a large microscopy dataset', *npj Computational Materials 2022 8:1*, vol. 8, no. 1, pp. 200-, Sep. 2022, doi: 10.1038/s41524-022-00878-5.
[54] R. Li *et al.*, 'Microscopy image reconstruction with physics-informed denoising diffusion probabilistic model', *Communications Engineering 2024 3:1*, vol. 3, no. 1, pp. 186-, Dec. 2024, doi: 10.1038/s44172-024-00331-z.
[55] F. C. Akyon *et al.*, 'SLICING AIDED HYPER INFERENCE AND FINE-TUNING FOR SMALL OBJECT DETECTION', *Proceedings - ICIP*, pp. 966–970, 2022, doi: 10.1109/ICIP46576.2022.9897990.
[56] M. Majurski and P. Bajcsy, 'Exact Tile-Based Segmentation Inference for Images Larger than GPU Memory', *J Res Natl Inst Stand Technol*, vol. 126, p. 126009, 2021, doi: 10.6028/JRES.126.009.
[57] E. Buglakova *et al.*, 'Tiling artifacts and trade-offs of feature normalization in the segmentation of large biological images', Mar. 2025, Accessed: Nov. 25, 2025. [Online]. Available: http://arxiv.org/abs/2503.19545
[58] A. Rhodes *et al.*, 'Analysis of Fiber Attrition and Mechanical Performance in Large-Format Additive Manufacturing of Long-Fiber Reinforced Polymer Composites'.
[59] J. Lagergren *et al.*, 'Few-Shot Learning Enables Population-Scale Analysis of Leaf Traits in Populus trichocarpa', *Plant Phenomics*, vol. 5, Jul. 2023, doi: 10.34133/PLANTPHENOMICS.0072
[60] P. Zhao *et al.*, 'A new method for classifying and segmenting material microstructure based on machine learning', *Mater Des*, vol. 227, p. 111775, Mar. 2023, doi: 10.1016/J.MATDES.2023.111775.
[61] K. He *et al.*, 'Mask R-CNN', *IEEE Trans Pattern Anal Mach Intell*, vol. 42, no. 2, pp. 386–397, Mar. 2017, doi: 10.1109/TPAMI.2018.2844175.
[62] J. Redmon *et al.*, 'You Only Look Once: Unified, Real-Time Object Detection', *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2016-December, pp. 779–788, Jun. 2015, doi: 10.1109/CVPR.2016.91.
[63] 'ultralytics/ultralytics: NEW - YOLOv8 🚀 in PyTorch > ONNX > OpenVINO > CoreML > TFLite'. Accessed: Dec. 11, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics
[64] S. Fang *et al.*, 'Improved Mask R-CNN Multi-Target Detection and Segmentation for Autonomous Driving in Complex Scenes', *Sensors (Basel)*, vol. 23, no. 8, Apr. 2023, doi: 10.3390/S23083853.
[65] Z. Tian *et al.*, 'FCOS: Fully Convolutional One-Stage Object Detection', *Proceedings of the IEEE ICCV*, vol. 2019-October, pp. 9626–9635, Apr. 2019, doi: 10.1109/ICCV.2019.00972.
[66] M. Li *et al.*, 'Machine learning and materials informatics approaches for predicting transverse mechanical properties of unidirectional CFRP composites with microvoids', *Mater Des*, vol. 224, p. 111340, Dec. 2022, doi: 10.1016/J.MATDES.2022.111340.
[67] N. Carion *et al.*, 'SAM 3: Segment Anything with Concepts', Nov. 2025, Accessed: Dec. 01, 2025. [Online]. Available: https://arxiv.org/abs/2511.16719v1

Funding

This research was sponsored by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Advanced Manufacturing Office, under contract DE-AC05-00OR22725 with UT-Battelle, LLC. This research was also supported by the National Science Foundation under Award No. 2055529.

Data Availability

Data will be made available upon reasonable request.

Acknowledgments

A special thanks to Andrew Rhodes for highlighting the need for this work. The authors also appreciate John Schepens, Roo Walker-Moir, and Nina Bhat for fiber length measurements.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used ChatGPT and Claude to help with editing for the improvement of sentence clarity and iterability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Descriptive Statistics	Annotator 1	Annotator 2
N	2810	2474
Mean ± SD (μm)	170.9 ± 107.0	174.1 ± 108.1
Weighted Mean (μm)	237.9	241.3
Median (μm)	156.38	162.15
Coeff. of Variation (%)	62.6	62.1
Mean difference	-3.19 (-1.83%)	----------
Length-weighted mean difference	-3.33 (-1.38%)	----------

Supplementary Table S1: Inter-operator agreement for manual fiber length measurement.

Parameter	U-Net	Mask R-CNN	YOLOv11
Architecture			
Backbone	Custom (Residual)	ResNet101	YOLOv11L
Pre-trained weights	No	Yes (COCO)	Yes
Input size (Training)	160 x 160	160 x 160	160 x 160
Training Configuration			
Training Epoch	500	500	1000
Batch Size	128	6	8
Learning Rate	0.001	0.003	0.001
Optimizer	Adam	SGD w/ momentum	AdamW
Early Stopping Patience	10 epochs	10 epochs	50 epochs
Loss Weights			
Segmentation/Mask	1.0	1.4	1.0
Classification	-	1.7	0.5
Bounding Box	-	3.0	-
RPN Class	-	1.0	-
BPN Bounding Box	-	2.75	-
SAHI Parameters			
Inference Tile Size	N/A	160 x 160	160 x 160
Tile Overlap Ratio	N/A	0.1	0.5
Post-process Method	N/A	Greedy NMM	NMM

Supplementary Table S2: Deep learning model training parameters and configurations.

Method	Type	DSC	IoU	Precision	Recall	F1
U-Net	Semantic	0.893	------	------	------	------
Mask R-CNN	Instance	0.746	0.631	0.675	0.910	0.768
YOLOv11	Instance	0.832	0.695	0.818	0.823	0.805

Supplementary Table S3: Segmentation metrics comparison. U-Net (semantic) evaluated on pixel classification only; Mask R-CNN and YOLOv11 (instance) evaluated on individual fiber detection and segmentation. All metrics calculated on validation images. DSC: Dice Similarity Coefficient, IoU: Intersection over Union.