

METHODOLOGY

Open Access



Segmentation and characterization of macerated fibers and vessels using deep learning

Saqib Qamar^{1,2,3}, Abu Imran Baba⁴, Stéphane Verger^{2,4,5*} and Magnus Andersson^{1,2,6*}

Abstract

Purpose Wood comprises different cell types, such as fibers, tracheids and vessels, defining its properties. Studying cells' shape, size, and arrangement in microscopy images is crucial for understanding wood characteristics. Typically, this involves macerating (soaking) samples in a solution to separate cells, then spreading them on slides for imaging with a microscope that covers a wide area, capturing thousands of cells. However, these cells often cluster and overlap in images, making the segmentation difficult and time-consuming using standard image-processing methods.

Results In this work, we developed an automatic deep learning segmentation approach that utilizes the one-stage YOLOv8 model for fast and accurate segmentation and characterization of macerated fiber and vessel form aspen trees in microscopy images. The model can analyze 32,640 x 25,920 pixels images and demonstrate effective cell detection and segmentation, achieving a mAP_{0.5–0.95} of 78 %. To assess the model's robustness, we examined fibers from a genetically modified tree line known for longer fibers. The outcomes were comparable to previous manual measurements. Additionally, we created a user-friendly web application for image analysis and provided the code for use on Google Colab.

Conclusion By leveraging YOLOv8's advances, this work provides a deep learning solution to enable efficient quantification and analysis of wood cells suitable for practical applications.

Keywords Instance segmentation, YOLO, Wood, Fibers, Optical microscopy

Introduction

Wood fibers, including tracheids and other fiber types, are essential components of wood, providing mechanical strength to withstand external mechanical stresses such as wind. Moreover, they are crucial for supporting tree growth in height [16]. These fibers have a high economic importance as they are the basic constituent of most wood-derived products. Wood fibers are extracted by pulping, which separates these cells into independent fiber cells. Once separated into individual fibers, they can be reassembled into paper and are increasingly utilized in applications such as bio-composites, smart papers, and new packaging materials [14]. The term “wood fiber” generally refers to the cell type called tracheids in softwood (Gymnosperm) and fibers in hardwood (angiosperm)

*Correspondence:

Stéphane Verger
stephane.verger@umu.se

Magnus Andersson
magnus.andersson@umu.se

¹ Department of Physics, Umeå University, Umeå 90187, Sweden

² Integrated Science Lab, Department of Physics, Umeå University, Umeå 90187, Sweden

³ Department of Intelligent System, Robotics, Perception, and Learning (RPL), KTH Royal Institute of Technology, Stockholm 100 44, Sweden

⁴ Department of Forest Genetics and Plant Physiology, Umeå Plant Science Centre (UPSC), Swedish University of Agricultural Sciences, Umeå 90183, Sweden

⁵ Department of Plant Physiology, Umeå Plant Science Centre (UPSC), Umeå University, Umeå 90187, Sweden

⁶ Umeå Centre for Microbial Research (UCMR), Umeå University, Umeå 90187, Sweden



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

[47]. In hardwood, fibers coexist with vessels and ray cells, which collectively determine wood characteristics. Notably, hardwood fiber cells elongate intrusively at their tips. Fiber cells are derived from the same stem cell progenitors as vessels, which do not elongate intrusively [39]. The ratio of fiber and vessel length can indicate the degree of intrusive growth (elongation), an essential factor in determining wood quality [14]. However, despite its importance, little is known about the development and elongation of these fibers or the genetic factors that regulate their length [29]. With this knowledge, fiber properties could be improved to produce better and stronger wood fiber-based products in more sustainable production systems [43]. As the demand for renewable fiber-based products continues to grow, there is a need to develop robust, high-throughput methods for studying fiber length and characteristics.

A classical method to study wood fiber length consists of macerating (soaking) wood samples in a solution that separates individual cells [39]. Cells suspended in liquid solutions can be transferred onto microscopy slides for examination. Once prepared, these slides can be observed using a standard light microscope. The images captured can then be analyzed using image processing software like ImageJ, which allows for the manual measurement of individual fiber lengths [38]. This task is, however, highly time-consuming and prone to user bias and errors during the manual measurement step. Overall, this highly limits the throughput of this type of analysis. Alternatively, it is possible to use so-called fiber analyzers. These machines allow the high-throughput image acquisition of fibers floating in a constantly stirred solution, generating high-speed and unbiased measurements [11]. However, the resolution of data acquisition is often limited, and it can be difficult to differentiate individual fibers from clumps of fibers that have not been properly detached from each other. While this may not be a limitation for industrial applications, it can become limiting when it comes to accurately studying the biology of fibers and their individual length, width, and other morphometric descriptors [27]. This type of equipment can also be very costly and is not widely available to many research labs contrary to light microscopes, even those equipped with a motorized stage.

A more widely accessible solution for fundamental research application is thus to make use of commonly available light microscopes equipped with motorized stages. Those are commonly available in biology research labs and most university core microscopy facilities. The motorized stage can be used to automatically acquire very large fields of view with high resolution, usually within 1–2 min per slide. Typically, wood macerates mounted between slides and cover slips allow the capture

of hundreds to thousands of fibers and vessels per slide. This high-throughput image acquisition can also, in principle, be coupled with new image processing technologies to automate the time-consuming task of identifying and measuring fiber cells in those high-resolution images. The use of automated microscopy slide scanning and automated image processing can thus largely alleviate the drawback of the light microscopy-based approach compared to fiber analyzers while also delivering much higher quality images for detailed fiber characterization.

While most recent light microscopes are equipped with motorized stages and the capacity to automatically generate large stitched images, the remaining limiting step is the image processing for the detection (segmentation) and shape analysis of the fibers. Several studies have developed image processing methods to analyze fibers from wood section images [5, 7, 22, 31]. However, most of these methods are only adapted for images from cross-sections of wood samples. These samples are typically prepared with cross-section in a transverse orientation. Although this method provides crucial information about certain aspects of cell arrangement and wood density, it does not directly yield information about the length of fiber cells. Automatic segmentation of such images is also less challenging for classical image segmentation algorithms, as individual cells in the images do not overlap. As such, the image can be divided into regions (e.g., individual cells and background) where each pixel is only assigned to one cell or region. This is typically readily achievable with the classical watershed segmentation algorithm [23] and more recently with deep learning based segmentation tools such as Cellpose [42] for more challenging samples. However, 2D images obtained from wood macerates contain many fibers that frequently overlap but which are still fully visible thanks to the translucency of the fibers imaged with the light microscope (see Fig. 1). Thus, many pixels in the image can belong to more than one cell of interest. This situation is challenging for most existing segmentation algorithms, including for deep learning techniques. Those are generally able to segment objects that are overlapping in nature, i.e., a human in front of a car, but can only output segmented masks that delineate the visible contours of each object and, thus, in this case, a full mask of the human in front and, an incomplete mask of the partially hidden car. With translucent objects, as is the case for light microscopy images of fibers, it would be, in principle, possible to obtain full masks for each of the overlapping objects since each object is fully visible in the image despite the overlaps. However, such an approach remains very little developed and applied [4, 8].

Deep learning, a specialized subset of machine learning, has achieved significant success in fields such as

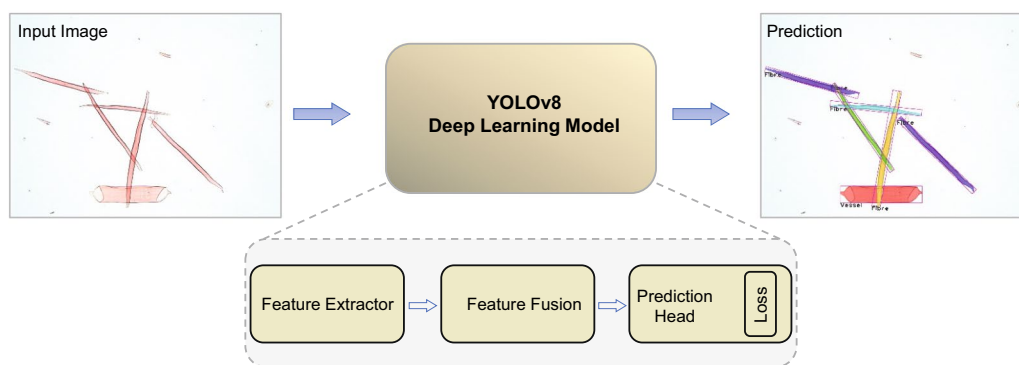


Fig. 1 Schematic of the YOLOv8 architecture for fiber and vessel segmentation. The model contains a Feature Extractor for feature extraction, Feature Fusion for feature aggregation, Prediction Head for predicting the objects' bounding boxes, classes, and masks. The loss component is used to optimize the model performance. An input image is passed through the network, which performs classification, detection, and segmentation jointly. This enables the delineation of individual cells even when overlapping, as shown in the prediction output

computer vision and image processing, particularly in tasks like image segmentation [34, 45]. For image segmentation, convolutional neural networks (CNNs), a type of deep neural network, are well-suited. CNNs process input images through multiple filters, thereby learning autonomously features from the images without the need for humans to design feature extractors manually. CNNs are extensively used for image segmentation applications [2, 28], including for segmenting cells in microscope images [13, 30, 36]. Given that cells in wood macerate images tend to overlap, it is essential that a neural network architecture supports the detection of individual objects, rather than merely dividing the image into regions. Instance segmentation, which identifies each cell individually and assigns pixels to the correct cell instance, is the most suitable approach in this regard. Popular existing instance segmentation methods such as Faster R-CNN [35], Mask R-CNN [18], and RetinaNet [25] have been successfully applied to cellular image segmentation [19, 21, 44]. These methods use deep convolutional neural networks (CNNs) like VGG [40] and ResNet [17] to extract features from the input images. However, these two-stage methods are slow for inference. One-stage methods such as Single Shot MultiBox Detector (SSD) [26] and You Only Look Once (YOLO) series streamline object detection by simultaneously predicting object locations and class probabilities in a single pass. One-pass approaches significantly improve inference speed over multi-stage methods. Among one-stage models, YOLO offers a good balance of speed and accuracy, achieving speeds suitable for real-time applications and that, for some cases, is more accurate than two-stage methods [9]. Most importantly, YOLOv8 can also be retrained

with a setting that allows it to deal with overlapping translucent objects and thus generate a full mask of such overlapping objects. Considering our need for a versatile model that is accessible to a broad audience and capable of analyzing high-resolution stitched images of macerated fibers containing many overlapping cells, the YOLO algorithm stands out as an ideal choice for addressing this challenge.

In this paper, we develop a model based on YOLOv8 [20] for segmenting and classifying fibers and vessels in 2D microscopy images, see Fig. 1. We compiled a dataset of 3850 wood macerate images using 1300 microscopy images. We annotated 9 617 fibers and 519 vessels, which, after augmentation resulted in 28 358 fibers and 1502 vessels used to train the model. We demonstrate that the enhanced model achieves fast inference speeds and high accuracy in detecting and classifying individual cells when processing large images (32,640 x 25,920 pixels). We also developed a browser interface for easy model access and image analysis. This interface, accessible after local installation, enables users to upload images via drag and drop. The system then provides measurements of the fibers and vessels, including their length, width, and area, as well as the full set of segmented masks corresponding to each segmented fiber and vessel, which can be used for further detailed morphometric analysis of fiber and vessel shapes. Additionally, users can access the training and prediction code run on Google Colab [3] in the GitHub repository [32]. Lastly, to validate the effectiveness of our new methods, we applied them to a well-known poplar transgenic line, which is recognized for having distinctly longer wood fibers than its wild-type relative. This approach aimed to corroborate previous manual measurements of fiber length differences in these lines.

Image capturing method

Sample preparation

To create training and test datasets, we collected stem samples from several 3-month-old hybrid aspens grown in a greenhouse, specifically from *Populus tremula* L. × *tremuloides* Michx.

To obtain data for comparative analysis of fiber length, 3 transgenic trees overexpressing the *GA20ox1* gene Arabidopsis gibberellin 20-oxidase (Ara GA20ox1 Line 1A) [11] and 3 trees of the T89 clone (control) were sampled.

To perform the maceration, stem segments between internodes above 10 cm from the soil were used. The bark was removed from the stem and 2–3 mm of the exposed vascular cambium region was trimmed, exposing the wood. From this surface, longitudinal match-like segments of 15 mm length and approximately 1.5 × 1.5 mm cross section were prepared from the newly exposed wood surface. Maceration was performed on these match-like samples by immersing them in maceration solution (30% Hydrogen peroxide: glacial acetic acid, 2:1 v/v) and heating at 90°C with periodic shaking for 5 h, as previously described in [39]. The macerated solution was then sedimented by low-speed centrifugation (1000 rpm) and washed a few times with water.

For safranin staining, a few drops from the obtained macerated solution were stained with the Safranin solution (1%) that stains lignified tissues in xylem cells. Similarly, for toluidine blue staining, a few drops from the macerated solution were stained with the (0.5% W/V) Toluidine blue stain.

Imaging

The samples were mounted between a slide and coverslip and imaged using a Leica DMI8 inverted microscope in brightfield mode with transmitted white light (Leica Microsystems, Germany). The microscope is equipped with a 10X objective lens and DFC7000T color camera mounted on a 0.70X C-mount adapter. Single images were acquired in RGB color mode with a resolution of 1920 × 1440 pixels and a pixel size of 0.65 × 0.65 μm. Tile images made of 19 × 19 (361) individual images were acquired with the Navigator function of the microscope using a 10 % overlap. Tile images were merged with the LAS X software (Leica Microsystems, Germany).

For training of the model, single tile images (1920 × 1440 pixel) of macerated fibers and vessels stained with safranin were obtained from the trimmed stem samples of wild-type (T89) trees. For quantification of the wildtype and the over-expression line, four 361-tile stitched images (containing fibers and vessels stained with safranin) were obtained from the trimmed stem

samples of each of these three biological replicates for both the control and over-expression lines.

To test the robustness of the model to different staining or imaging modes we also acquired several images of fibers and vessels either stained with Toluidine blue, non stained and non stained samples acquired in grayscale mode. We provide all the raw image data at Zenodo [33].

Statistics

We used a t-test statistical method to validate whether the trained model outputs consistent results across different image groups. Specifically, we tested for scale invariance by running the model on different-sized crops from the same large images. We also compared model predictions on an overexpression line GA20ox 1A and wildtype T89 as control. The t-statistic is calculated as $t = \frac{\bar{X}_1 - \bar{X}_2}{\text{SED}}$, where \bar{X}_1 and \bar{X}_2 are the sample means and SED is the standard error of the difference between the means. We employed Scipy's `ttest_ind()` to automatically compare two independent data samples, assessing significant mean differences based on t-test assumptions.

Computational method

Images annotation and dataset preparation

To identify fibers and vessels in microscope images, we created a dataset of over 1300 individual images showing these structures in different shapes and sizes. The images are 1920 × 1440 pixels in size. We carefully outlined the fibers and vessels in each image using the VGG Image Annotator [10]. This created a ground truth or guide to the actual fibers and vessels in the images. The outlines were saved as JSON files, which store the coordinates of the polygons drawn around each object, which can be seen in Fig. S1. We resized the large images into smaller 1024 × 1024 pixel images to avoid running out of memory on our graphics card. We also used data augmentation techniques to increase the number of training images. For example, by applying transformations like rotating, scaling, and flipping the original images, we created more variety in the dataset. This also helps the machine learning model learn robust features that apply to new images. From 1300 original images, we generated 3850 augmented images with around 29 861 annotated objects.

To train the YOLO model, we randomly split the dataset into training (85%) and validation (15%) sets. The training data is used to update the model's parameters. The validation data is used to evaluate the model during training but not to update parameters. This split helps prevent overfitting and ensures the model generalizes well.

Deep learning approach

We used the YOLOv8-seg deep learning method, which is a variant of YOLOv8 architecture designed explicitly for instance segmentation tasks [20]. During the model training process, we utilized the YOLOv8 model pre-trained on the COCO val2017 dataset and set `overlap_mask = False` in training parameters to deal with overlapping masks as a starting point. The architecture of the YOLOv8 algorithm consists of four main components: feature extractor, feature fusion, prediction head, and loss function. The components are shown in Fig. 1 and with more details in Fig. S2. Here, we look into the design concepts of each architecture module.

The feature extractor is the first part of the model and is responsible for extracting features at different stages from the input image. The output features from the different stages have different spatial resolutions. The earlier stages of the feature extractor network extract low-level features such as edges and corners. The later stages extract high-level features such as object shapes and parts. The feature extractor down-samples the input image because it extracts features at later stages.

The feature fusion module combines the output features from different stages of the feature extractor network to form a unified representation of the image. Deep neural networks capture increasingly detailed features as the network becomes deeper, which improves object prediction. However, as the network depth increases, the object localization accuracy for detecting small objects decreases owing to excessive convolution operations, resulting in the loss of important information. To address this tradeoff, the feature fusion module incorporates a multi-scale fusion of features using architectures such as a Feature Pyramid Network (FPN) and Path Aggregation Network (PAN). The feature fusion module performs different operations to extract higher-level features from the input features and consolidate the outputs from various stages of the feature extractor into a single representation. This unified representation enhances the object detection and segmentation performance.

The prediction-head module transforms the encoded image features into usable predictions for object detection and segmentation. This makes the final predictions based on the consolidated representation from the feature-fusion module. The head module combines features from earlier modules and leverages them to predict bounding boxes, classes, and masks for the objects in the image. By dividing the prediction into specialized branches, it efficiently performed classification, localization, and masking in a coordinated manner. The prediction head is the final component that outputs the actual detections and segmentations after processing using the complete YOLOv8 architecture.

The loss function in YOLOv8 measures how accurately the model detects and segments objects. It compares the model predictions with the ground-truth labels. The loss function is used to train the model to improve its performance. YOLOv8 has separate branches for classification, bounding-box regression, and masking. For classification and masking, the cross-entropy loss was used to minimize errors. Bounding box detection uses two losses: distributed focal loss (DFL) and CIOU Loss. These consider the aspect ratio between the predicted and ground-truth boxes. The overall loss is the sum of the losses from different branches. A lower loss indicates that the model is more accurate at detecting, classifying, and segmenting objects. Loss guides model training to improve these areas.

Measuring the morphology of objects

To identify and localize objects of interest within images, we use the YOLOv8 model to predict bounding boxes and contours around each instance. To enable more precise morphological measurement, we integrate functions from the OpenCV computer vision library to further analyze the object contours [6].

We use the contours to generate the mask of each predicted object by the model. We then calculate the length of the skeleton path of the binary mask. For this, we apply a thinning operation that is used to reduce the thickness of objects in an image to a single-pixel-wide skeleton. The goal of thinning is to preserve the essential structural and topological characteristics of the original objects while significantly reducing the amount of data. To apply the thinning process, we use the OpenCV library function `cv2.ximgproc.thinning()`. We then analyze the skeletons in the thinned image, extract the long path skeleton, and calculate its length using `np.sum(mask)` from the NumPy library. This centerline approach allows us to find the actual length of the object.

To calculate the width of the fiber, the code first computes the distance transform of the binary mask drawing. The distance transform assigns each pixel a value that corresponds to the distance from that pixel to the nearest background pixel. This is done using the `cv2.distanceTransform()` function. The maximum value in the distance transform represents the thickness of the thickest part of the fiber. This value is obtained using the `np.max()` function from the NumPy library. Since the distance transform gives the distance from the center of the fiber to the edge, the actual thickness of the fiber is twice this value.

In addition to length and width, we used OpenCV's `cv2.contourArea()` method to accurately measure the area enclosed within the detected contours. A schematic diagram showing the concept is shown in Fig. 2, where

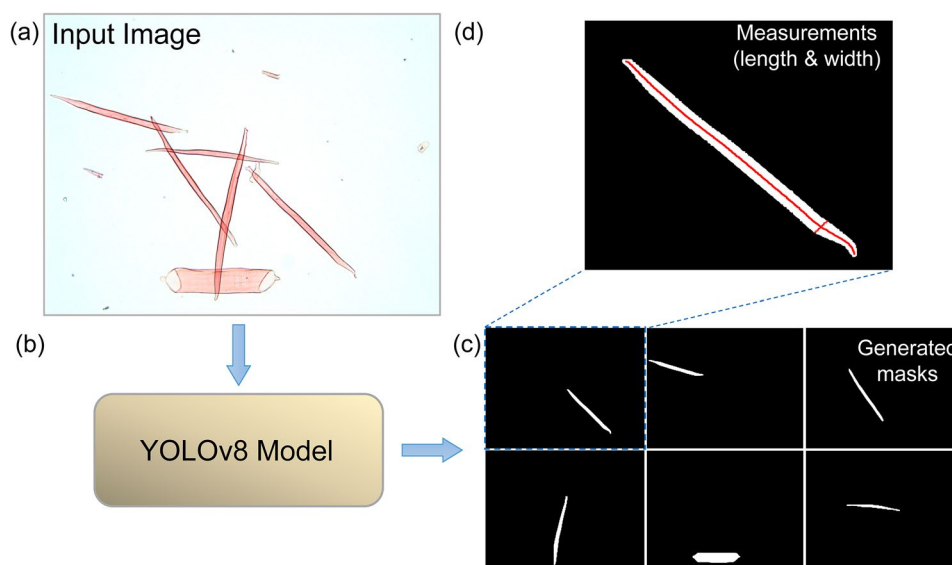


Fig. 2 Example of a segmented object and the corresponding image analysis that automatically measures the object's morphological traits using our YOLOv8 model. **a** displays the original microscopy input image used for object detection, which is sent to the YOLOv8 model **b** that outputs detected objects. **c** shows the individual mask generated for each individual cell in the image and also demonstrates that full masks are obtained from translucent overlapping objects. **d** illustrates the measurements of length and width for each detected object, providing quantitative data extracted from the image

the input image is passed to the model and the model gives the output. From the output, we generate masks for each instance and then use OpenCV functions to measure the length and width of each instance. The output from the model provides pixel values from the analysis, which need to be converted into micrometers.

By utilizing the polygon points in the contours, we perform post-processing to exclude segmented objects that touch the boundary edge of the image. This post-processing approach ensures that only fully segmented objects within the image are included.

Model retraining and parameter adjustment

YOLOv8 model retraining and hyperparameter tuning with wood macerate dataset

During the training process of all models, we utilized the pre-trained YOLOv8 models on the COCO val2017 dataset as a starting point. To train the YOLO model, we use 3350 images for training and 500 images for validation. The training data is used to update the model's parameters whereas the validation data is used to evaluate the model during training but not to update parameters. We used YOLOv8-m, l, and x in our experiment. YOLOv8-m uses a medium-sized feature extractor and more feature fusion levels, YOLOv8-l utilizes a larger feature extractor and more feature fusion levels compared to YOLOv8-m, and YOLOv8-x uses an even larger feature extractor and more feature fusion compared to YOLOv8-l. We find the convergence level and best optimizer for these models

during training. Based on experimental data from Ultralytics, we observed that YOLOv5 training required 300 epochs, while YOLOv8 training increased the number of epochs to 600. Initially, we set the number of epochs to 600 and incorporated a patience value of 50. This means that if no noticeable improvement occurred after waiting for 50 epochs, the training would terminate early. However, during the training of YOLOv8m, we found that the model reached its best performance at epoch 355 and training stopped early at epoch 405.

We chose hyperparameters for model training as suggested in reference [15]. The selection of an appropriate optimization algorithm is crucial when training YOLOv8-seg. The optimizer determines how the model parameters are updated during training to minimize the loss function. For small custom datasets, the Adam (Adaptive Moment Estimation) optimizer is recommended, while the SGD (Stochastic Gradient Descent) optimizer tends to perform better on larger datasets. Consequently, we trained YOLOv8-seg models separately using the Adam and SGD optimizers. The results of comparing the effects of these two optimizers on model training are presented in Table 1.

We opted for the Adam optimizer with a weight decay of 5×10^{-4} and an initial learning rate of 1×10^{-3} . Furthermore, we set the input image size to 1024 and trained the different models using a TITAN V100 16GB with a batch size of 8. The models were trained on Python 3.8 and PyTorch 1.10.0.

Evaluation metrics for the models

To evaluate the performance of the models, we used four metrics: precision, recall, F1-score, and mean average precision (mAP). These metrics are commonly used for object detection and segmentation tasks, and they are calculated based on the number of true positive (TP), true negative (TN), false positive (FP), and false negative (FN) predictions made by the model [41].

Precision quantifies how many of the predicted positive instances are actually correct. Recall quantifies how many of the actual positive instances are correctly identified by the model. F1-score combines both precision and recall giving a single value that represents the algorithm’s overall accuracy. A higher F1-score indicates a better algorithm performance in achieving both high precision and high recall simultaneously.

Mean Average Precision (mAP) is a widely adopted evaluation metric for assessing the performance of object detection algorithms across multiple classes. In this paper, we considered mAP 50 and mAP 50–95, in which, mAP 50 calculates the average precision for all classes at an IoU threshold of 0.5 while mAP 50–95 computes the average precision for all classes over a range of IoU thresholds from 0.5 to 0.95, with a step size of 0.05. This variation of mAP offers a more comprehensive evaluation by considering a wider range of IoU thresholds.

Table 1 The table compares YOLOv8m-seg, YOLOv8l-seg, and YOLOv8x-seg models trained on our generated dataset using both SGD and Adam optimizers

Model	Optimizer	Best Epoch	mAP 50	mAP 50-95	Speed GPU V100 (ms)
YOLOv8m-seg	SGD	405	0.93	0.73	8.3
	Adam	413	0.94	0.75	9.0
YOLOv8l-seg	SGD	421	0.93	0.74	17.4
	Adam	487	0.94	0.76	17.2
YOLOv8x-seg	SGD	553	0.93	0.76	24.5
	Adam	543	0.94	0.78	22.2

Table 2 The table shows the evaluation metrics for different YOLOv8 models when detecting fibers and vessels

Model	Precision	Recall	F1-Score	mAP 50	mAP 50-95	Weight
YOLOv8m-seg	0.97	0.91	0.94	0.95	0.75	27.2
YOLOv8l-seg	0.96	0.91	0.93	0.94	0.76	45.9
YOLOv8x-seg	0.97	0.90	0.93	0.94	0.78	70.1

Result and discussion

Model selection

In this work, we used the pre-trained YOLOv8 to detect and segment fibers and vessels. We trained with our custom dataset using three YOLO variants: m, l, and x. The quantitative results for precision, recall, F1 score, mAP@0.5, and mAP@0.5–0.95 values of the three YOLOv8 models in fiber and vessel segmentation are presented in Table 2. Among these models, YOLOv8m-seg demonstrated the highest precision (0.97), recall (0.91), and F1 score (0.94), while performing relatively lower in mAP: 0.5–0.95 (0.75). YOLOv8l-seg exhibited slightly lower precision (0.96) and recall (0.91) compared to YOLOv8m-seg, but had a comparable F1 score (0.93) and a marginally better mAP: 0.5–0.95 (0.76). YOLOv8x-seg matched YOLOv8m-seg in precision (0.97), had a slightly lower recall (0.90), and an F1 score (0.93) equal to YOLOv8l-seg, but it outperformed both in mAP: 0.5–0.95 with the highest value of 0.78. Furthermore, the YOLOv8x-seg model had the highest weight (70.1), compared to YOLOv8m-seg (27.2) and YOLOv8l-seg (45.9).

Based on these findings, the YOLOv8x-seg model was identified as the optimal choice for fiber and vessel detection and segmentation, exhibiting superior performance across the evaluation metric of mAP@0.5–0.95. Consequently, the YOLOv8x-seg model was selected for further analysis, specifically in estimating fiber and vessel length, width, and area. This selection ensures a consistent and focused evaluation of the model’s practical application, aligning with the study’s objectives.

Model performance

We thoroughly evaluated model performance for the tasks of detecting and segmenting fiber and vessel objects in images. We examined precision-recall curves to understand the trade-off between precision and recall for detection. For segmentation, we focused on how precisely the model delineates the boundaries and segments of the detected objects. Additionally, we analyzed F1-confidence curves to understand the relationship between F1 scores and model confidence levels for detection and segmentation. Examining the F1-confidence curves provided insights into how precision and recall were balanced across varying confidence thresholds.

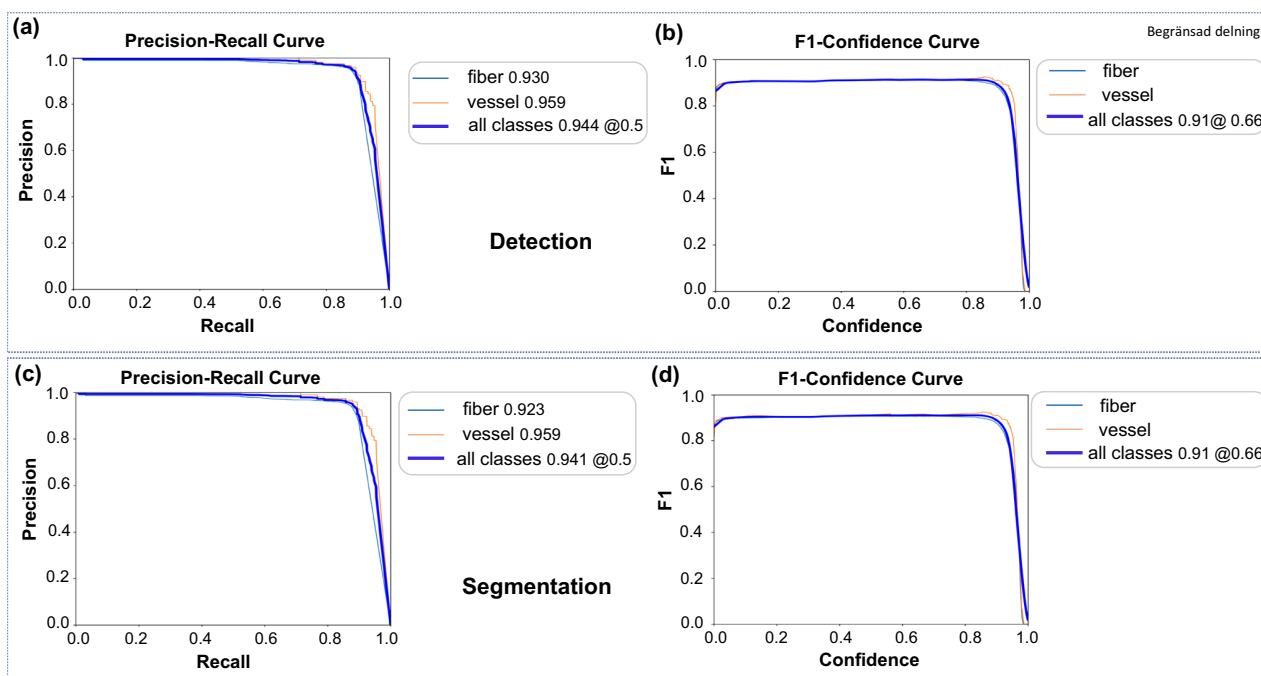


Fig. 3 These plots showcase the precision-recall (a, c) and F1-confidence (b, d) curves used to evaluate the performance of YOLOv8x in detecting (a, b) and segmenting (c, d) fibers and vessels. The model demonstrates strong mAP and F1-score across thresholds, confirming its effectiveness in object detection and segmentation tasks

Evaluating the model's performance step-by-step is important for assessing its effectiveness.

Figure 3 depicts the behavior of the selected YOLOv8x-seg model used for object detection and segmentation of fiber and vessel objects. The precision-recall curve in Fig. 3a represents the trade-off between precision and recall for the detection task. At a threshold of 0.5, the mean average precision (mAP) values were 0.930 for fiber and 0.959 for vessel detection. The overall mAP of 0.944 for all classes combined indicated the model's overall performance in object detection.

Figure 3b shows the F1-confidence curve, which illustrates the relationship between the F1 score and the model's confidence. At a confidence threshold of 0.66, the F1-score was 0.91 for both fiber and vessel classes. This score represents the balance between precision and recall. A higher F1-score indicates better model performance in terms of both precision and recall. These findings provide valuable insights into the model's performance and help assess its suitability for detecting fiber and vessel objects in images.

Moving on to the segmentation task, the model achieved values of 0.923 and 0.959 at a threshold of 0.5 for fiber and vessel segmentation, respectively. The overall mAP of 0.94 for all combined classes at the same threshold is shown in (Fig. 3c), highlighting the model's effectiveness in segmenting objects. Figure 3d displays

the F1-confidence curve for the segmentation task, where the model attained an F1-score of 0.91 at a confidence threshold of 0.66. This score represents the trade-off between precision and recall for both fiber and vessel classes.

We conclude that for fiber and vessel detection and segmentation, the YOLOv8x-seg model performed best compared to other models in terms of mAP. The YOLOv8m-seg model had the lowest mAP value specifically for fiber detection and segmentation.

Qualitative results

Figure 4 shows fiber and vessel detection and segmentation examples achieved using the YOLOv8x-seg model. These examples highlight the model's ability to accurately identify and outline fiber and vessel structures in the images. The visual results obtained from the model not only demonstrate its potential within the scope of our study but also provide valuable insight into its practical application for fiber and vessel analysis.

Because the model was primarily trained on RGB color images of safranin stained samples, we next aimed to test whether it was also able to detect and segment fibers in images obtained with different staining protocols and color acquisition parameters. We acquired several images with an imperfect white balance (Fig. 4a), Toluidine blue stained samples, unstained samples, and unstained

samples acquired in grayscale mode, all other images were acquired in RGB color mode (Fig. 4b). Despite these variations, our model consistently succeeded in detecting and segmenting fiber and vessel structures. This highlights the effectiveness of the model, as it still can perform well. Such robustness is particularly valuable in fiber and vessel segmentation, where images with diverse backgrounds are encountered.

Based on the findings of this study, we conclude that our retrained version of the YOLOv8x-seg model effectively detects and segments fibers and vessels in microscopy images. Despite achieving high accuracy in fiber detection and segmentation, our YOLOv8 model still made some mistakes by generating false positives and false negatives in certain cases. We highlight examples of this in Fig. 5 (a), where the model failed to identify several fibers and fiber segments, for instance, highlighted by the yellow ellipse. Similarly, in Fig. 5(b), the ellipse highlights one of the regions where YOLOv8 failed to identify the fiber and vessel. This failure to detect the fiber could be attributed to limitations in the training of the model but is also likely a consequence of the very high density and crowding of the fibers in this example image. This issue can also be easily alleviated during sample preparation by increasing the dilution of the sample before mounting it between the slide and coverslip and imaging. Nevertheless, to tackle this issue from the computational side, Adar et al. [12] suggested that training the model with a larger dataset containing more input features can greatly enhance its ability to generalize and perform well on new and unseen data. A larger dataset would enable the model to capture the subtle differences in fiber and vessel

structures in various background images. Additionally, a larger dataset can help to mitigate the risk of overfitting, where the model becomes too specialized to the training dataset and performs poorly on new samples.

Quantifying fibers and vessels in microscopy images

In order to complete the task of fiber and vessel detection and segmentation, we implemented standard measurements to describe the shape of the objects based on the model output. Specifically, we evaluated how well our model performed with images of different sizes. We separated the images into two categories based on size. The first type measured 1920 x 1440 pixels, similar to those used to train the model. The second type was large tile images of 33,384 x 25,112 pixels stitched from 361 small 1920 x 1440 images. The model first identified the locations of fibers and vessels in the image. Then, it assesses the morphological traits (e.g., length, width and area) of the detected objects using the method described in subsection 3.3. The total detected items in two large images are 1818 fibers and 62 vessels. Note that we implemented a post-processing step that removes fibers that are touching the borders in the image to avoid a bias of shorter fibers due to partially segmented objects.

Figures 6 and 7 show box plots summarizing the measurements of the detected fibers and vessels. The fibers' length ranged from 254–364 μm (average 310 μm), the width ranged from 21–25 μm (average 23 μm), and the area ranged from 3 283–5 145 μm^2 (average 4 347 μm^2). Whereas for most vessels, the length ranged from 218–275 μm (average 247 μm), the width ranged from 44–55 μm (average 51 μm), and the area ranged from 6 536–9

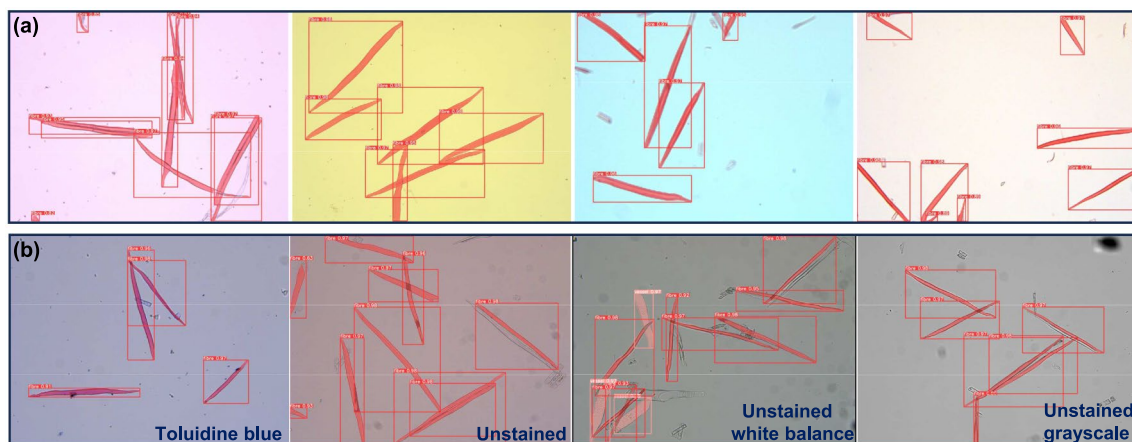


Fig. 4 The images depict YOLOv8 effectively detecting and segmenting images obtained with various staining protocols and color acquisition parameters. Panels in (a) depict only Safranin stained fibers with diverse backgrounds typically associated with imperfect white balance and (b) shows from left to right, Toluidine blue stained samples, unstained, unstained with properly adjusted white balance, and unstained acquired in grayscale mode. Note that in these representations, the overlaid mask of fibers and vessels is displayed as red and should not be confused with red staining

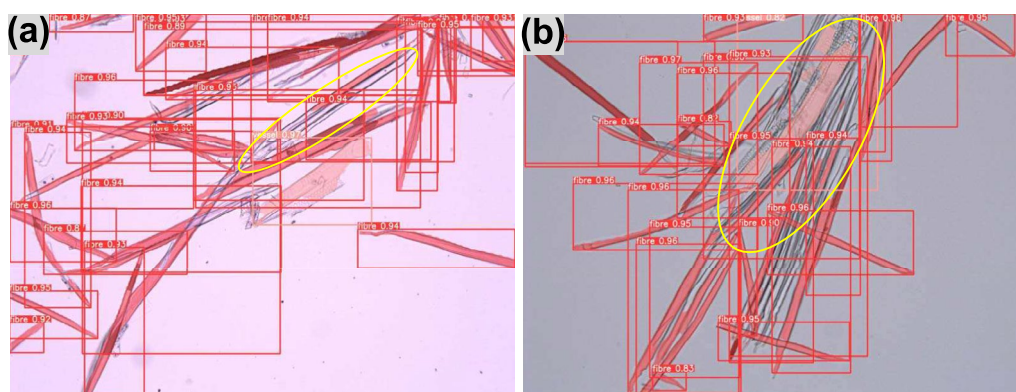


Fig. 5 The detection of fibers and vessels encountered some challenges in densely packed regions, as shown in these examples where several fibers and vessels were not properly detected or segmented

494 μm^2 (average 8 142 μm^2). In the supplementary materials, scattered and histogram plots offer a detailed visualization of the distribution and relationship between the length, width, and area as depicted in Fig. S3.

Further, testing on different image sizes allowed us to check if the model performance stayed consistent regardless of size changes. The results show that the model's effectiveness is independent of image dimensions. This scale-invariance makes the model more robust and can analyze images of varying sizes which are encountered in real applications. This property is crucial for real-world problems, as different image dimensions will inevitably be encountered in practice. To evaluate scale invariance and object detection and segmentation abilities, including overlapped objects, 20 small images of size 1920 x 1440 pixels were randomly selected, containing 115 fibers and 17 vessels. Measurements were compared to those from the large 33,384 x 25,112 pixel image.

We found the model extracted consistent average dimensions for fibers across these two image sizes. This includes an average length of 376 μm , a width of 23 μm , and an area of 10 689 μm^2 for the fibers. For fiber length, width, and area, the p-values were 0.954, 0.963, and 0.945, respectively, indicating no statistically significant difference between small and large image measurements.

However, we noted a variability when comparing the vessel measurements with a lower average length of 269 microns and an average width of 57 microns. The vessel area averages were closer between image sizes at 17 755 μm^2 on small images and 13 533 μm^2 on large images. The p-values for vessel length, width, and area were 0.275, 0.142, and 0.454, respectively. These results illustrate the scale invariance performance of the model across different image dimensions. In the supplementary material, we also provide the analysis of 20 small-sized images (1920 x 1440), see Fig. S4 and two mid-sized images consisting of 99 tiles (8275 x 7250), see Fig. S5, using box plots, scatter plots, and histogram plots to assess the robustness of the model.

In conclusion, the model's ability to accurately extract metrics for fibers and vessels including those that overlap, at various scales, underscores its proficiency in managing overlapping objects and scale invariance. Such capabilities are crucial for the robust detection and segmentation of high-resolution images, ensuring consistent performance regardless of image size. This model, therefore, becomes a valuable tool for research groups aiming to quantify metrics from extensive datasets.

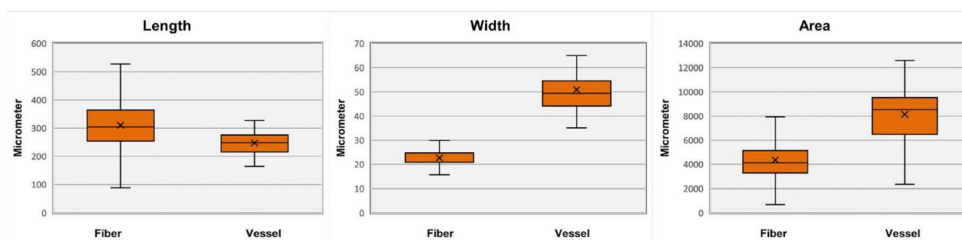


Fig. 6 Boxplots depict fibers' and vessels' length, width, and area measurements from two large images of 33,384 x 25,112 pixels. The box plots show the 1st and 3rd quartiles (box limits), mean values (cross sign), and median values (solid line)

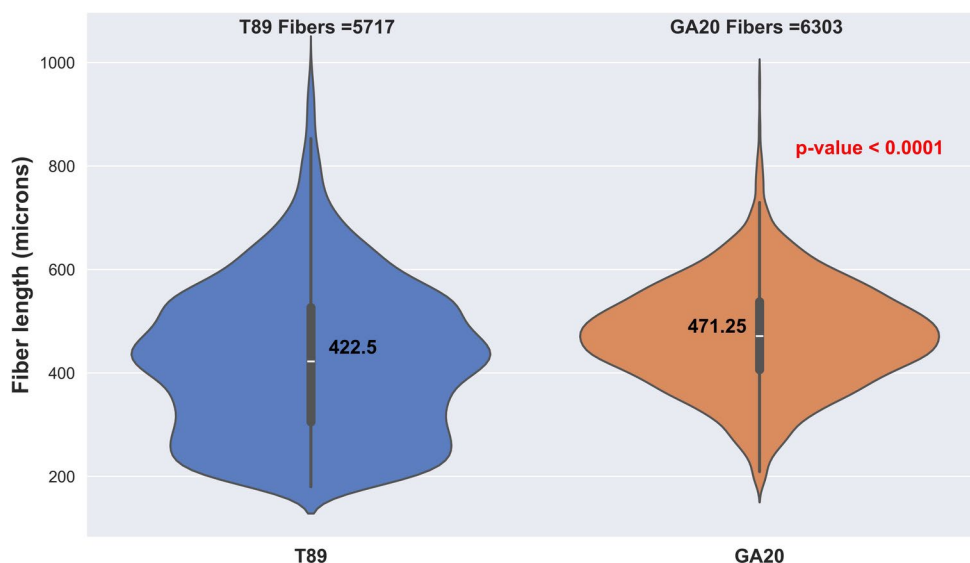


Fig. 7 Violin plots that depict fibers' length measurements from large images of 33,384 x 25,112 pixels from T89 and GA20ox line. The envelope shows the distribution, the thick center lines represent the 1st and 3rd quartiles (black line limits), and median values (center yellow dots)

Comparison of GA20ox 1A line and wildtype T89

To further evaluate our approach for research applications, we tested it on a new dataset consisting of 12 large images of samples taken from T89 wildtype trees and 12 large images from the transgenic line overexpressing the *GA20ox1* gene. This transgenic line was previously reported to show approximately 10 % increase in fiber length compared to the wildtype T89 [11]. In total, the model identified 5717 fibers in the 12 T89 images with an average length of 422.5 μm and for the GA20 line it identified 6303 fibers with an average of 471.25 μm , see Fig. 7. A student t-test indicated that the samples were highly significantly different, as evidenced by a p-value less than 0.0001. These results show an approximately 12 % length increase for the GA20ox line compared to the wildtype, in line with previously reported results [11]. Overall, these results suggest that our new high-throughput microscopy-based AI-assisted fiber characterization method works accurately to quantify differences between wild-type and mutant lines. As such it will be particularly useful in the future to study much larger mutant and natural variation tree collections.

GUI application and GitHub resources

We have developed a desktop application using Python's Flask framework, designed to simplify the process of uploading and analyzing images with our automated algorithm. The application's workflow is shown in Fig. S6. Users can select an image from a gallery or upload one via drag-and-drop. Note, that the image should be in a RGB format. Once an image is uploaded, the application

predicts the presence of fibers and vessels. It also quantifies each detected object, with results available in a downloadable data file. Note that segmented objects that touch the edge of the image and may thus not be fully segmentable, are automatically excluded from the analysis at this postprocessing step. This file includes details like object type, length, width, and area. It also generates a folder containing the full set of segmented masks corresponding to each segmented fiber and vessel. Such binary masks can then be further analysed to quantify other shape descriptors of interest using, for example, OpenCV [6, 46], and ImageJ/Fiji [1, 24, 37]. When installed on a local computer with an i9 10th Gen processor and 32GB of RAM, the application processes a 1024 x 1024 image in an average of 140 milliseconds.

Detailed instructions for installation of the developed code and instructions for implementation are found in the supporting information. Additionally, users can find instructions on how to retrain and test the model using Google colab in the GitHub resources [32]. The 'Train_custom_data.ipynb' file can be used to retrain the model; simply follow the instructions provided in the file. The 'prediction_file.ipynb' file is used to test the model on your custom dataset.

Conclusion

This paper introduces a deep learning solution using YOLOv8 to automatically analyze and quantify wood fibers and vessels in challenging microscope images, offering high-throughput capabilities. To achieve this, we trained multiple YOLOv8 models on diverse wood image

datasets and evaluated their performance in detecting and segmenting fibers and vessels. The most robust model was chosen. The model can consistently and reliably extract essential cell metrics across different image scales, such as length, width, and area. The model's consistent metric extraction underscores its strong practical applicability. We also created a web application pipeline that is useful in practical situations. Users can then upload images for automatic cell counting and shape quantification. Thus, we conclude that this study introduces an innovative high-throughput method for analyzing wood cells in densely populated 2D microscopy images, even when cells are partially obscured.

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s13007-024-01244-w>.

Additional file 1.

Author contributions

Saqib Qamar: Data curation, Software, Methodology, Validation, Formal analysis. Abu Imran Baba: Investigation, Validation. Stéphane Verger: Conceptualization, Resources, Supervision, Funding acquisition. Magnus Andersson: Conceptualization, Resources, Supervision, Funding acquisition, Project administration. All authors contributed to Writing, review and editing.

Funding

Open access funding provided by Umeå University. The project was funded by Kempeförelserna (JK-2129.3) and the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at Umeå, partially funded by Vetenskapsrådet (2022-06725). The authors acknowledge the facilities and technical assistance of the Umeå Plant Science Centre Microscopy facility and the plant growth facility. This work was also supported by grants from the Wallenberg Foundation (KAW 2016.0341 and KAW 2016.0352), VINNOVA (2016-00504), and the Novo Nordisk Foundation (NNF21OC0067282) to S.V. We also thank Bio4Energy for supporting this work.

Availability of data

All the code in this project was developed using Python and different public libraries as defined in the supporting information. The Google Colab project is found here [3] and the code can be downloaded at a public GitHub repository here [32]. All raw image data are available at Zenodo [33].

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

The authors declare that they have no Conflict of interest.

Competing interests

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

Received: 25 January 2024 Accepted: 22 July 2024

Published online: 14 August 2024

References

1. Abràmoff MD, Magalhães PJ, Ram SJ. Image processing with ImageJ. *Biophoton Int.* 2004;11(7):36–42.
2. Badrinarayanan V, Kendall A, Cipolla R. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans Pattern Anal Mach Intell.* 2017;39(12):2481–95.
3. Bisong E. Google Colaboratory. Berkeley: Apress; 2019. p. 59–64.
4. Böhm A, Ücker A, Jäger T, et al. Isoo dl: Instance segmentation of overlapping biological objects using deep learning. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), IEEE. 2018;1225–1229.
5. Boztoprak H, Ergun M. Determination of vessel and fibers in hardwoods. *Gaziosmanpasa J Sci Res.* 2017;6(2):87–96.
6. Bradski G. The OpenCV Library. *Dr Dobb's Journal of Software Tools*; 2000.
7. Brunel G, Borianne P, Subsol G, et al. Automatic identification and characterization of radial files in light microscopy images of wood. *Ann Bot.* 2014;114(4):829–40.
8. Chen L, Wu Y, Merhof D. Instance segmentation of dense and overlapping objects via layering. 2022. arXiv preprint [arXiv:2210.03551](https://arxiv.org/abs/2210.03551)
9. Diwan T, Anirudh G, Tembhrune JV. Object detection using yolo: challenges, architectural successors, datasets and applications. *Multimedia Tools Appl.* 2023;82(6):9243–75.
10. Dutta A, Zisserman A. The via annotation software for images, audio and video. In: Proceedings of the 27th ACM international conference on multimedia. 2019;2276–2279.
11. Eriksson M, Israelsson M, Olsson O, et al. Increased gibberellin biosynthesis in transgenic trees promotes growth, biomass production and xylem fiber length. *Nat Biotechnol.* 2000;18(7):784–8.
12. Frid M, Klang E, Amitai M, et al. Synthetic data augmentation using gan for improved liver lesion classification. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). 2018;289–293.
13. Fu H, Xu Y, Lin S, et al. Deepvessel: Retinal vessel segmentation via deep learning and conditional random field. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19, Springer. 2016;32–139.
14. Gholampour A, Ozbakkaloglu T. A review of natural fiber composites: properties, modification and processing techniques, characterization, applications. *J Mater Sci.* 2020;55(3):829–92.
15. Glenn J. YOLOv5 release v6. 1; 2022.
16. Gorshkova T, Brutch N, Chabbert B, et al. Plant fiber formation: state of the art, recent and expected progress, and open questions. *Crit Rev Plant Sci.* 2012;31(3):201–28.
17. He K, Zhang X, Ren S, et al. Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2016;770–778.
18. He K, Gkioxari G, Dollár P, et al. Mask R-CNN. In: Proceedings of the IEEE international conference on computer vision. 2017; 2961–2969.
19. Hollandi R, Szkalitsy A, Toth T, et al. A deep learning framework for nucleus segmentation using image style transfer. *Biorxiv.* 2019;580605.
20. Jocher G, Chaurasia A, Qiu J. Yolo by ultralytics. Github; 2023.
21. Johnson J. Adapting mask-rcnn for automatic nucleus segmentation. 2018. arXiv preprint [arXiv:1805.00500](https://arxiv.org/abs/1805.00500)
22. Kennel P, Subsol G, Guérout M, et al. Automatic identification of cell files in light microscopic images of conifer wood. In: 2010 2nd international conference on image processing theory, tools and applications, IEEE. 2010;98–103.
23. Kornilov AS, Safonov IV. An overview of watershed algorithm implementations in open source libraries. *J Imaging.* 2018;4(10):123.
24. Legland D, Arganda-Carreras I, Andrey P. Morpholibj: integrated library and plugins for mathematical morphology with imagej. *Bioinformatics.* 2016;32(22):3532–4.
25. Lin T, Dollár P, Girshick R, et al. Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2017;2117–2125.
26. Liu W, Anguelov D, Erhan D, et al. Ssd: Single shot multibox detector. In: Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14, Springer. 2016;21–37.

27. Lobo J, See EYS, Biggs M, et al. An insight into morphometric descriptors of cell shape that pertain to regenerative medicine. *J Tissue Eng Regen Med.* 2016;10(7):539–53.
28. Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2015;3431–3440.
29. Majda M, Kozlova L, Banasiak A, et al. Elongation of wood fibers combines features of diffuse and tip growth. *New Phytol.* 2021;232(2):673–91.
30. Minaee S, Boykov Y, Porikli F, et al. Image segmentation using deep learning: a survey. *IEEE Trans Pattern Anal Mach Intell.* 2021;44(7):3523–42.
31. Pan S, Kudo M. Recognition of wood porosity based on direction insensitive feature sets. *Trans Mach Learn Data Min.* 2012;5(1):45–62.
32. Qamar S. fiberseg. 2023. <https://github.com/sqbqamar/fiberseg>
33. Qamar S, Baba AI, Verger S, et al. Fiber and vessel dataset for segmentation and characterization. 2024. <https://doi.org/10.5281/zenodo.10913446>.
34. Razzak MI, Naz S, Zaib A. Deep learning for medical image processing: Overview, challenges and the future. *Classification in BioApps: automation of decision making.* 2018;323–350.
35. Ren S, He K, Girshick R, et al. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28. 2015.
36. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. Springer. 2015;234–241.
37. Schindelin J, Arganda-Carreras I, Frise E, et al. Fiji: an open-source platform for biological-image analysis. *Nat Methods.* 2012;9(7):676–82.
38. Schneider C, Rasband W, Eliceiri KW. Nih image to imagej: 25 years of image analysis. *Nat Methods.* 2012;9(7):671–5.
39. Siedlecka A, Wiklund S, Péronne MA, et al. Pectin methyl esterase inhibits intrusive and symplastic cell growth in developing wood cells of populus. *Plant Physiol.* 2008;146(2):554.
40. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. 2020. arXiv p. 1409.1556.
41. Sokolova M, Lapalme G. A systematic analysis of performance measures for classification tasks. *Info Proces Manag.* 2009;45(4):427–37.
42. Stringer C, Wang T, Michaelos M, et al. Cellpose: a generalist algorithm for cellular segmentation. *Nat Methods.* 2021;18(1):100–6.
43. Thumm A, Dickson AR. The influence of fibre length and damage on the mechanical performance of polypropylene/wood pulp composites. *Compos A Appl Sci Manuf.* 2013;46:45–52.
44. Tsai H, Gajda J, Sloan TF, et al. Usiigaci: instance-aware cell tracking in stain-free phase contrast microscopy enabled by machine learning. *SoftwareX.* 2019;9:230–7.
45. Voulodimos A, Doulamis N, Doulamis A, et al. Deep learning for computer vision: A brief review. *Computational intelligence and neuroscience.* 2018;2018.
46. Van der Walt S, Schönberger JL, Nunez-Iglesias J, et al. scikit-image: image processing in python. *PeerJ.* 2014;2: e453.
47. Wilson K, White DJB, et al. The anatomy of wood: its diversity and variability. *Stobart & Son Ltd;*1986.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.