UC Davis

Title

A workflow for segmenting soil and plant X-ray computed tomography images with deep learning in Google's Colaboratory

Permalink

Authors

Rippner, Devin A
Raja, Pranav V
Earles, J Mason
et al.

Check for updates

A workflow for segmenting soil and plant X-ray computed tomography images with deep learning in Google's Colaboratory

Devin A. Rippner[1]*, Pranav V. Raja[2], J. Mason Earles[2,3], Mina Momayyezi[3], Alexander Buchko[4], Fiona V. Duong[5], Elizabeth J. Forrestel[3], Dilworth Y. Parkinson[6], Kenneth A. Shackel[7], Jeffrey L. Neyhart[8] and Andrew J. McElrone[3,9]

[1]Horticultural Crops Production and Genetic Improvement Research Unit-United States Department of Agriculture-Agricultural Research Service, Prosser, WA, United States, [2]Department of Biological and Agricultural Engineering, University of California, Davis, Davis, CA, United States, [3]Department of Viticulture and Enology, University of California, Davis, Davis, CA, United States, [4]Department of Computer Science, California Polytechnic and State University, San Luis Obispo, CA, United States, [5]Department of Integrative Biology, San Francisco State University, San Francisco, CA, United States, [6]Advanced Light Source, Lawrence Berkeley National Laboratory, Berkeley, CA, United States, [7]Department of Plant Sciences, University of California, Davis, Davis, CA, United States, [8]Genetic Improvement for Fruits and Vegetables Laboratory, United States Department of Agriculture-Agricultural Research Service, Chatsworth, NJ, United States, [9]Crops Pathology and Genetics Research Unit, United States Department of Agriculture-Agricultural Research Service, Davis, CA, United States

X-ray micro-computed tomography (X-ray μCT) has enabled the characterization of the properties and processes that take place in plants and soils at the micron scale. Despite the widespread use of this advanced technique, major limitations in both hardware and software limit the speed and accuracy of image processing and data analysis. Recent advances in machine learning, specifically the application of convolutional neural networks to image analysis, have enabled rapid and accurate segmentation of image data. Yet, challenges remain in applying convolutional neural networks to the analysis of environmentally and agriculturally relevant images. Specifically, there is a disconnect between the computer scientists and engineers, who build these AI/ML tools, and the potential end users in agricultural research, who may be unsure of how to apply these tools in their work. Additionally, the computing resources required for training and applying deep learning models are unique, more common to computer gaming systems or graphics design work, than to traditional computational systems. To navigate these challenges, we developed a modular workflow for applying convolutional neural networks to X-ray μCT images, using low-cost resources in Google's Colaboratory web application. Here we present the results of the workflow, illustrating how parameters can be optimized to achieve best results

using example scans from walnut leaves, almond flower buds, and a soil aggregate. We expect that this framework will accelerate the adoption and use of emerging deep learning techniques within the plant and soil sciences.

Introduction

Researchers have long been interested in analyzing the in-situ physical, chemical, and biological properties and processes that take place in plants and soils. To accomplish this, researchers have widely adopted the use of X-ray micro-computed tomography (X-ray μCT) for 3D analysis of flower buds, seeds, leaves, stems, roots, and soils (Petrovic et al., 1982; Crestana et al., 1985, 1986; Anderson et al., 1990; Brodersen et al., 2010; Hapca et al., 2011; Mooney et al., 2012; Helliwell et al., 2013; Cuneo et al., 2020; Théroux-Rancourt et al., 2020; Xiao et al., 2021; Duncan et al., 2022). In plants, researchers have used X-ray μCT to visualize the internal structures of leaves, allowing for the quantification of CO_2 diffusion through the leaf based on path length tortuosity from the stomata to the mesophyll (Mathers et al., 2018; Théroux-Rancourt et al., 2021). Other applications of X-ray μCT in plants include the visualization of embolism formation and repair in plant xylem tissue, allowing for the development of new models to better understand drought stress recovery, along with non-destructive quantification of carbohydrates in plant stems (Brodersen et al., 2010; Torres-Ruiz et al., 2015; Earles et al., 2018). In soils, X-ray μCT was used to visualize soil porosity, soil aggregate distribution, and plant root growth (Tracy et al., 2010; Mooney et al., 2012; Helliwell et al., 2013; Mairhofer et al., 2013; Ahmed et al., 2016; Yudina and Kuzyakov, 2019; Gerth et al., 2021; Keyes et al., 2022). Despite the wide use of advanced imaging techniques like X-ray μCT and imaging more generally in agricultural research, major limitations in both hardware and software hinder the speed and accuracy of image processing and data analysis.

Historically X-ray μCT data collection was extremely time consuming, and resource intensive as individual scans can exceed 50 Gb in size. Data acquisition rates were limited by the ability of X-ray detectors to transfer data to computers, limited hard drive storage capacity once the data was transferred, and intensive hardware requirements that limited the size of files that could be analyzed at any given time. Many of these constraints have been removed as detector hardware has improved, hard drive storage transfer speed and space has increased, and computing hardware has advanced. Now a major limiting step to the widespread use of X-ray μCT in the agricultural sciences is data analysis; while data can be acquired in hours to seconds, the laborious task of hand segmenting images can lead to analysis times of weeks to years for large data sets (Théroux-Rancourt et al., 2020).

Recent advances in machine learning, specifically the application of convolutional neural networks to image analysis, have enabled rapid and accurate segmentation of image data (Long et al., 2015; Ronneberger et al., 2015; Chen et al., 2018; Smith et al., 2020; Raja et al., 2021; von Chamier et al., 2021). Such applications have met with great success in medical imaging analysis, outperforming radiologists for early cancer diagnosis in X-ray μCT images (Lotter et al., 2021). However, challenges remain to applying convolutional neural networks to the analysis of agriculturally relevant X-ray μCT images. Specifically, training accurate models for image segmentation requires the production of hand annotated training datasets, which is time consuming and requires specialized expertise to properly annotate training image data (Kamilaris and Prenafeta-Boldú, 2018). Further, the computing resources required for training and applying deep learning models are unique, more common to computer gaming systems or graphics design work, rather than traditional computational systems (Gao et al., 2020; Ofori et al., 2022).

To navigate these challenges, we developed a modular workflow for image annotation and segmentation using open-source tools to empower scientists that use X-ray μCT in their work. Specifically, image annotation is done in ImageJ; while this does not prevent the need for experts to annotate images, it does allow experts to annotate their images without using proprietary software. The semantic segmentation of X-ray μCT image data is accomplished using Google's Colab to run PyTorch implementations of a Fully Convolutional Network (FCN) with a ResNet-101 backbone (He et al., 2015; Long et al., 2015; Ronneberger et al., 2015; Chen et al., 2017; Paszke et al., 2019). The FCN architecture, while older, allows for model development on variable size images due to the exclusion of fully connected layers in the FCN architecture (Long et al., 2015; Ronneberger et al., 2015). In addition to X-ray μCT datasets, the workflow is flexible enough to work on virtually any image dataset, as long as corresponding annotated images

are available for model training. Additionally, by developing and deploying the code in Google's Colaboratory, users have access to free or low-cost GPU resources that might otherwise be cost prohibitive to access (Rippner et al., 2022b). If users have access to better hardware than is available through Colaboratory, the notebooks can be run locally to utilize advanced hardware. Additional code is also available to use this workflow on high performance computing systems using batch scheduling (Rippner et al., 2022b). This method for analyzing X-ray μCT data allows users to rapidly extract information on important biological, chemical, and physical processes that occur in plants and soils from complex datasets without the need to learn to code extensively or invest in expensive computational hardware.

Materials and methods

In the following section we will describe the parameters under which our CT data was collected, how the CT image data was annotated for model training, and the parameters used to train the various models. The actual workflow and corresponding training video can be found on Github (Rippner et al., 2022b). For reproducibility purposes, the data sets used for training the models featured in this paper can be found a repository hosted by the United States Department of Agriculture, National Agricultural Library (Rippner et al., 2022a).

Computed tomography data acquisition

Six individual leaf sections (3 mm × 7 mm) from 6 unique accessions of English walnuts (*Juglans regia)* and an air dried soil aggregate collected from the top 15 cm of a Yolo silt loam (Fine-silty, mixed, superactive, non-acid, thermic Mollic Xerofluvents) at the UC Davis Russel Ranch Sustainable Agricultural Facility were scanned at 23 keV using the 10 × objective lens with a pixel resolution of 650 nanometers on the X-ray μCT beamline (8.3.2) at the Advanced Light Source (ALS) in Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, United States. Additionally, an almond flower bud (*Prunis dulcis*) was scanned using a 4 × lens with a pixel resolution of 1.72 μm on the same beamline. Raw tomographic image data was reconstructed using the TomoPy tomographic image reconstruction engine (Gürsoy et al., 2014). Reconstructions were converted to 8-bit tif or png format using ImageJ or the PIL package in Python before further processing (**Figure 1**; Schindelin et al., 2012; Kemenade et al., 2022).
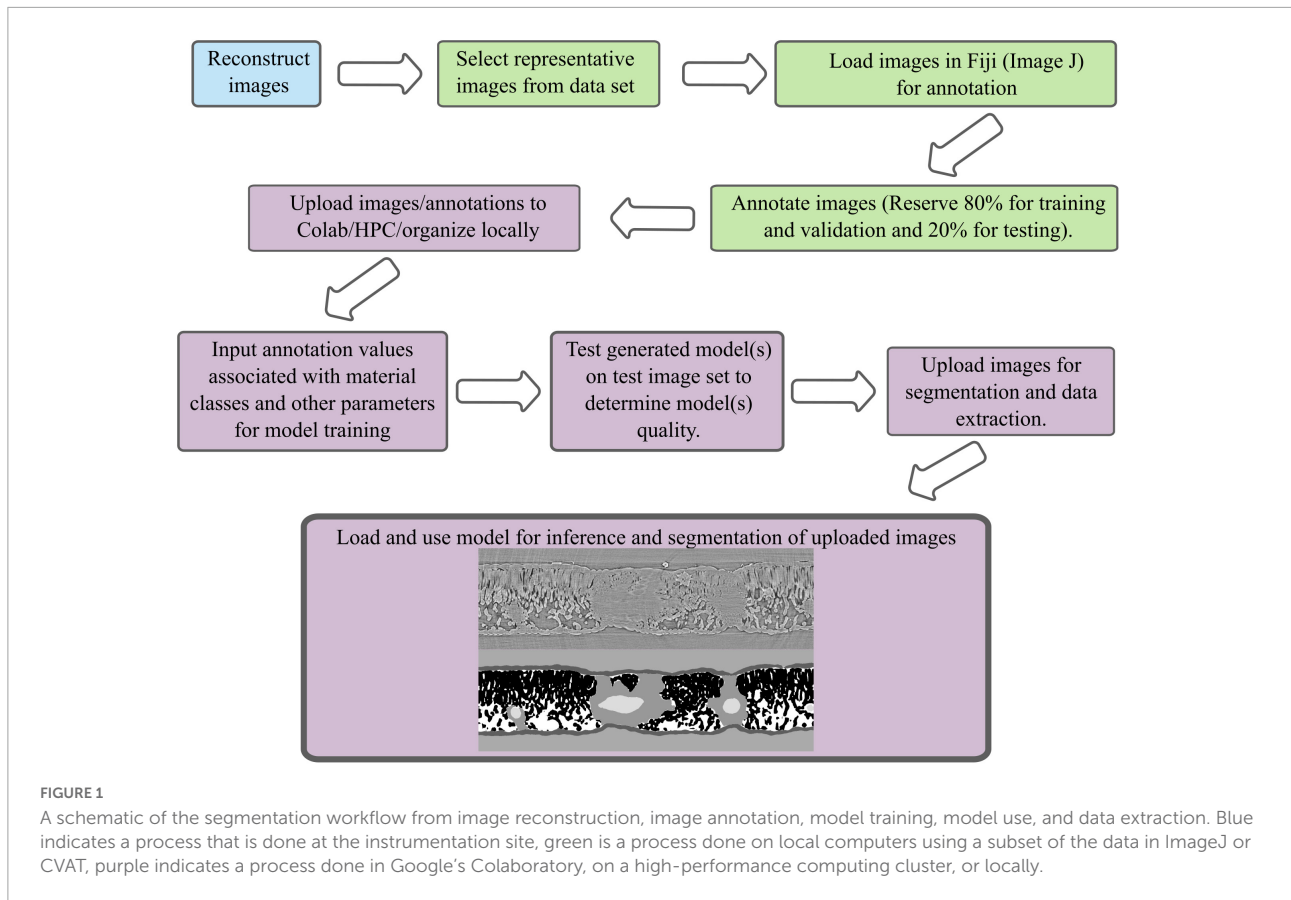
Image annotation

Leaf images were annotated in ImageJ following Théroux-Rancourt et al. (2020) (**Figure 1**). Flower bud and soil

aggregate images were annotated using Intel's Computer Vision Annotation Tool (CVAT) and ImageJ (**Figure 1**; Schindelin et al., 2012). Both CVAT and ImageJ are free to use and open source. To annotate the flower bud and soil aggregate, images were imported into CVAT. The exterior border of the bud (i.e., bud scales) and flower were annotated in CVAT and exported as masks. Similarly, the exterior of the soil aggregate and particulate organic matter identified by eye were annotated in CVAT and exported as masks. To annotate air spaces in both the bud and soil aggregate, images were imported into ImageJ. A Gaussian blur was applied to the image to decrease noise and then the air space was segmented using thresholding. After applying the threshold, the selected air space region was converted to a binary image with white representing the air space and black representing everything else. This binary image was overlaid upon the original image and the air space within the flower bud and aggregate was selected using the "free hand" tool. Air space outside of the region of interest for both image sets was eliminated. The quality of the air space annotation was then visually inspected for accuracy against the underlying original image; incomplete annotations were corrected using the brush or pencil tool to paint missing air space white and incorrectly identified air space black. Once the annotation was satisfactorily corrected, the binary image of the air space was saved. Finally, the annotations of the bud and flower or aggregate and organic matter were opened in ImageJ and the associated air space mask was overlaid on top of them forming a three-layer mask suitable for training the FCN.

Training general *Juglans* leaf segmentation model

Images and associated annotations from 6 walnut leaf scans were uploaded to Google Drive (**Figure 1**). Using Google's Colaboratory resources, a PyTorch implementation of a FCN with a ResNet-101 backbone was used to train 10 models using 5 image/annotation pairs from 1, 2, 3, 4, and 5 leaves (5, 10, 15, 20, and 25 images/annotation pairs, respectively) (**Figure 1**, **Supplementary Tables 1, 2** and **Supplementary Figures 1, 2**; He et al., 2015; Paszke et al., 2019). Models pre-trained on the COCO train 2017 dataset were imported and the original classifier was substituted for a new classifier based on 6 potential pixel classes: background, epidermis, mesophyll tissue, air space, bundle sheath extension tissue, or vein tissue. The pre-trained model weights were modified using an Adam optimizer for stochastic optimization with the learning rate set to 0.001 and a binary cross-entropy loss function (Kingma and Ba, 2017). To help avoid overfitting the training data, the data was augmented using Albumentations package in Python to flip and rotate a subset of the images during model training (Buslaev et al., 2020). Half of the image/annotation pairs were used for training and half were used for validation of the model during training. The batch size was set at 1 for

FIGURE 1

A schematic of the segmentation workflow from image reconstruction, image annotation, model training, model use, and data extraction. Blue indicates a process that is done at the instrumentation site, green is a process done on local computers using a subset of the data in ImageJ or CVAT, purple indicates a process done in Google's Colaboratory, on a high-performance computing cluster, or locally.

training due to graphics processing unit (GPU) constraints in Colaboratory. A mixture of NVIDIA T4, P100, V100 GPUs were used for training depending on the allocation assigned by Google Cloud Services. Such GPU's are available when using the free version of Google's Colaboratory, or the low cost ($9.99/month) subscription based Colaboratory Pro (Mountain View, CA, United States). A benefit of limiting the batch size to 1 was the ability to train on variably sized images.

The accuracy, precision, recall, and F1 score of these models were calculated after testing on 5 images from the 6th leaf that was not involved in training or validation of the generated models in any way. In our work accuracy, precision, recall, and f1 scores are defined as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

$$Precision = \frac{(TP + 1E - 9)}{(TP + FP + 1E - 9)} \tag{2}$$

$$Recall = \frac{(TP + 1E - 9)}{(TP + FN + 1E - 9)} \tag{3}$$

$$F1 = \frac{(TP + 1E - 9)}{(TP + 1E - 9) + \frac{1}{2}(FP + FN)} \tag{4}$$

Where TP = true positive prediction on a pixelwise basis, FP = false positive prediction on a pixelwise basis, TN = true negative prediction on a pixelwise basis, and FN = false negative prediction on a pixelwise basis. A correction factor of 1E-9 was included in Equations 2–4 to prevent *Not a Number* errors in python when the denominator of the equations is 0 due to the lack of TP, FP, or FN values when no prediction is made for a non-existent material class in a particular image (Powers, 2020).

The evaluation results for each of the 10 models generated after training on 1, 2, 3, 4, and 5 leaves were compiled using the Panda's library in Python and visualized using the Seaborn library (**Supplementary Table 1**; Oliphant, 2007; McKinney, 2011; Waskom, 2021).

The number of training epochs (i.e., iterative learning passes through the complete data set) to train a satisfactory model was also evaluated using the 25 image/annotation pairs taken from 5 annotated leaves. This number of training images was found to give the best results with a fixed number of epochs for model training. Ten models were generated after training for 10, 25, 50, 100, and 200 epochs. The accuracy, precision, recall, and F1 scores for these models were calculated after evaluating the same 5 images from the 6th leaf that was not used for model training and validation. Binary image outputs for each material type generated for the leaves were stacked and rendered in

3-dimensions using ORS Dragonfly (Object Research Systems, Montréal, Canada).

Training models for segmenting flower buds and soil aggregates

A mixture of NVIDA T4, P100, V100, and A100 GPUs were used for training models for segmenting an almond flower bud and a soil aggregate. Such GPU's are available when using the free version of Google's Colaboratory, the low cost ($9.99/month) subscription based Colaboratory Pro, or the higher cost ($49.99) subscription based Colaboratory Pro + (Mountain View, CA, United States). Training and validation images for both the almond flower bud and the soil aggregate

had to be downscaled to 50% size in the x and y dimensions to fit on the T4, P100, and V100 video cards due to VRAM limitations (16 Gb VRAM) (**Supplementary Table 1**). When using the A100 GPU (40 Gb VRAM) available through Colaboratory Pro +, images used for model training and validation from the flower bud and soil aggregate were only downscaled to 85% in the x and y dimensions, representing a large gain in image data for training and validation (**Supplementary Table 1**). Again, models pre-trained on the COCO train 2017 dataset were imported and the original classifier was substituted for a new classifier based on 4 potential pixel classes. For the soil aggregate, these were background, mineral solids, pore space and particulate organic matter; for the almond bud there were background, bud scales, leaf tissues, and air space. The pre-trained model weights were modified using an Adam optimizer
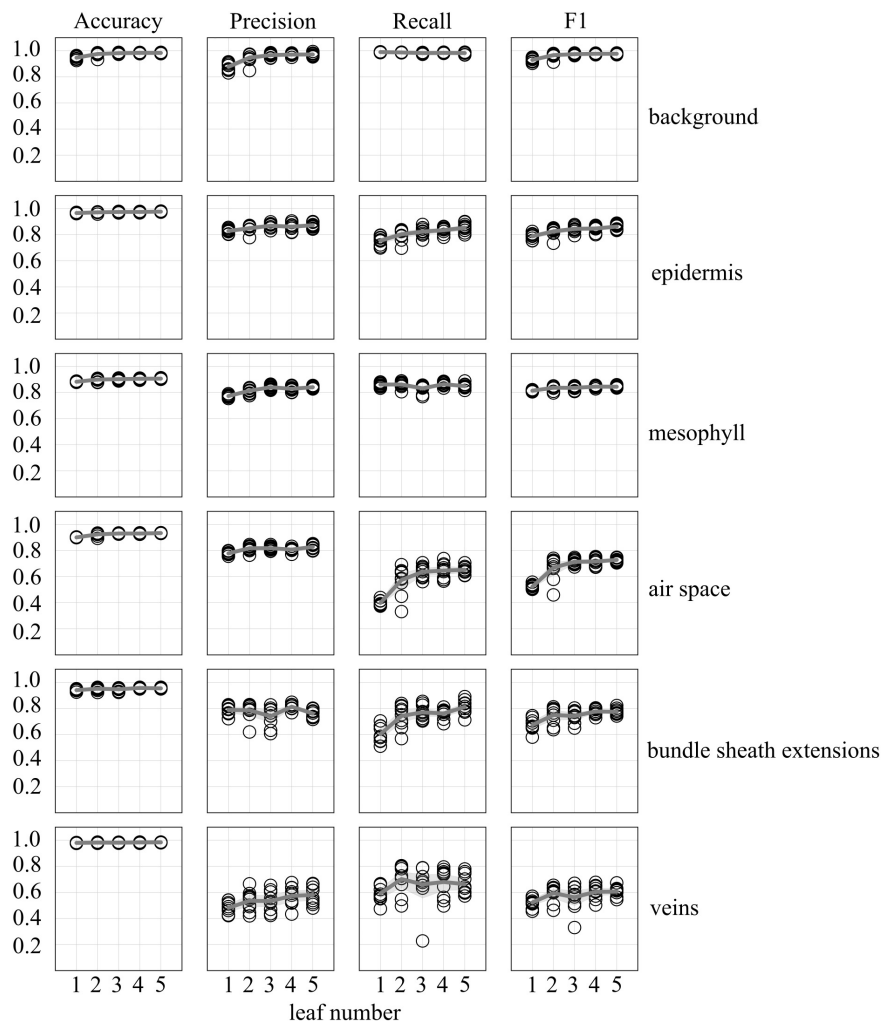
FIGURE 2
Accuracy, precision, recall, and F1 scores as a function of uniquely annotated leaf number (5 annotated images per leaf) used to predict tissue classes in X-ray μCT images from an independent leaf on which the models were not trained or validated. Circles represent unique predictions from 10 uniquely generated models per leaf number; dark gray lines represent the mean value of the 10 models while thick light gray lines represent the 95% confidence interval of the values.

for stochastic optimization with the learning rate set to 0.001 and a binary cross-entropy loss function (Kingma and Ba, 2017). To help avoid overfitting the training data, the data was augmented using Albumentations package in Python to flip and rotate a subset of the images during model training (Buslaev et al., 2020). Models were trained for 200 epochs as model loss for these data was previously found to plateau between 100 and 200 epochs. Model accuracy, precision, recall and F1 scores were calculated as above after testing on 5 independently annotated images from the same flower bud and soil aggregate that were not used for model training or validation. Binary image outputs for each

material type generated for the almond bud and soil aggregate were stacked and rendered in 3d using ORS Dragonfly (Object Research Systems, Montreal, Canada).

Results

Accuracy, precision, recall, and F1 scores for a general model to identify and segment specific walnut leaf tissues on a pixel-wise basis plateaued after training and validation on at least 3 annotated leaves. For epidermis and mesophyll tissues,

FIGURE 3
Visual representation of the model outputs for a single walnut leaf image; top image is a X-ray CT scan taken from a leaf that was not used for training or validation of the applied models; next is the hand annotated image of the scan followed by the outputs from the best performing model trained on 1, 3, and 5 leaves, respectively. For the walnut leaf segmentations the background is light gray, the epidermis is dark gray, the mesophyll is black, the air space is white, the bundle sheath extensions are middle gray, and the veins are lightest gray.

prediction F1 scores were > 80% while prediction F1 scores were generally > 75% for bundle sheath extensions and > 70% for airspaces. The lowest prediction F1 scores were achieved for veins tissues (∼60%) and the highest for the background class (>95%) (**Figures 2**, **3**). Precision scores were generally higher than recall scores except in the case of vein tissue identification, where the models tended to over-predict the occurrence of the vein tissue class. F1 score variability across all prediction classes decreased as the number of leaves used during training and validation increased. Model F1 score variability was highest for the bundle sheath extension and vein tissue classes, likely due to colocation of the two tissue classes. When evaluating generalized walnut leaf model performance with increasing training epochs, model F1 scores plateaued after 50 epochs. Model F1 score variability was consistent after 50 epochs, with

no improvement after additional training time (**Figures 4**, **5**). This was particularly true for the vein tissue class which took the most training for consistent identification.

Model performance for the segmentation of an almond flower bud and a soil aggregate was hindered by the downscaling necessary to fit the training and validation data on video cards with 16 Gb of VRAM. When downscaled to 0.5 size in the x and y dimensions (25% size), the best model F1 scores were 99.8, 99.1, 92.4, and 71.5% for the background, bud scale, flower, and air space classes, respectively. At the same scaling for the soil aggregate, the best model F1 scores were 99.4, 83.7, 55.6, and 71.3% for the background, solid, pore, and organic matter classes, respectively. With access to the A100 GPU with 40 Gb of VRAM, training images were only downscaled to 0.85 size in the x and y dimensions (72% size). This yielded
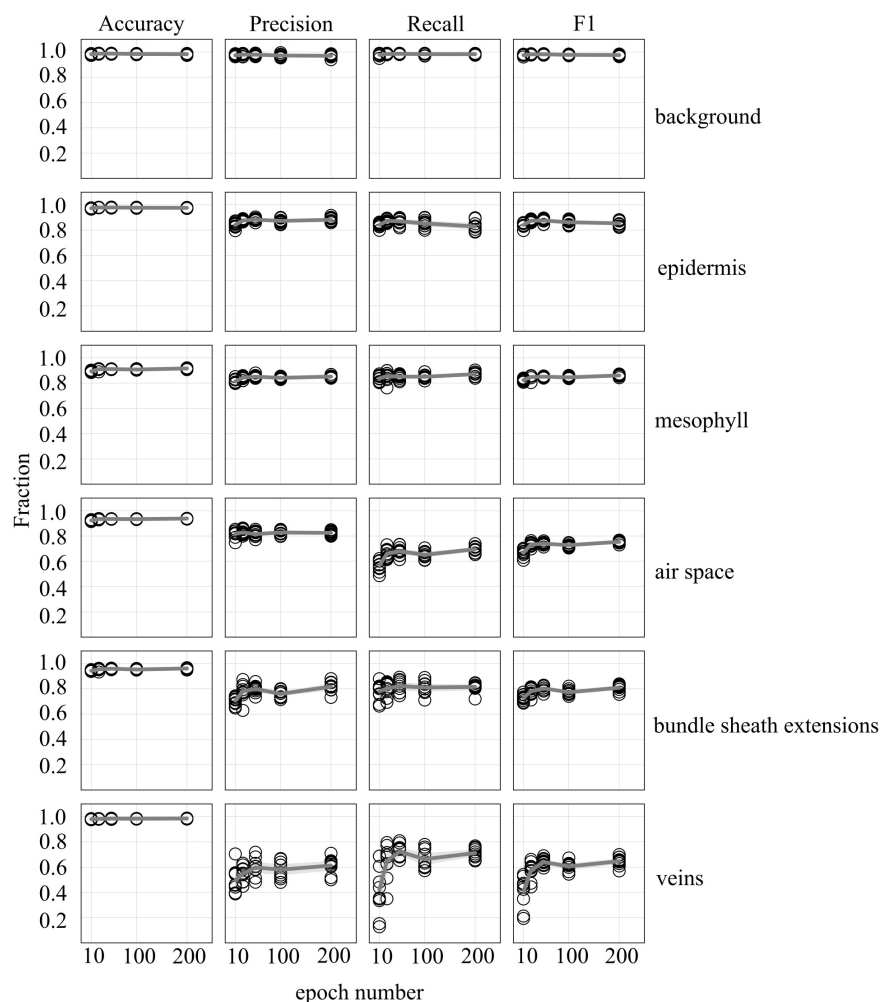
FIGURE 4
Accuracy, precision, recall, and F1 scores as a function increasing epoch number for models trained on annotated images from 5 leaves (5 annotated images per leaf) used to predict tissue classes from X-ray μCT images from an independent leaf on which the models were not trained or validated. Circles represent unique predictions from 10 uniquely generated models per epoch number; dark gray lines represent the mean value of the 10 models for each tissue type while thick light gray lines represent the 95% confidence interval of the values.
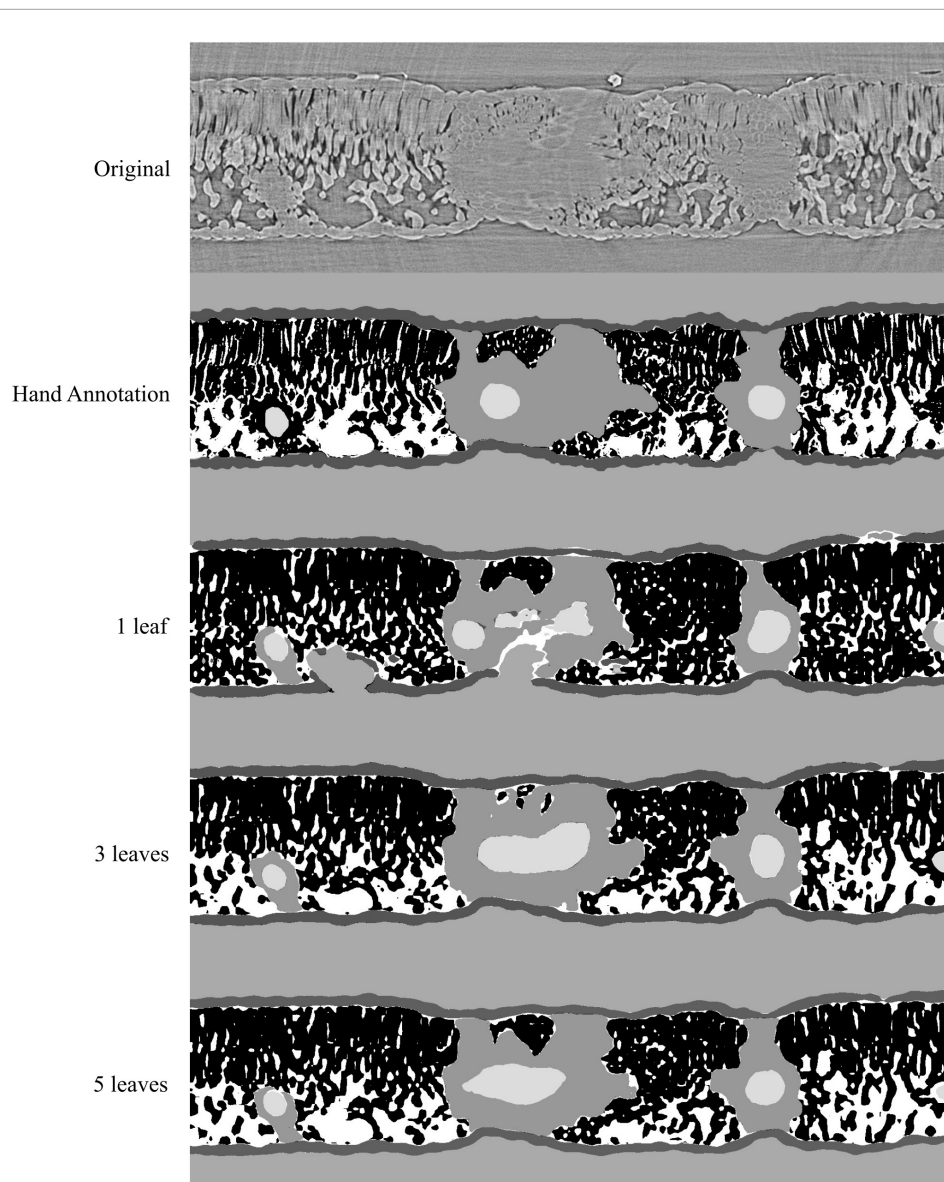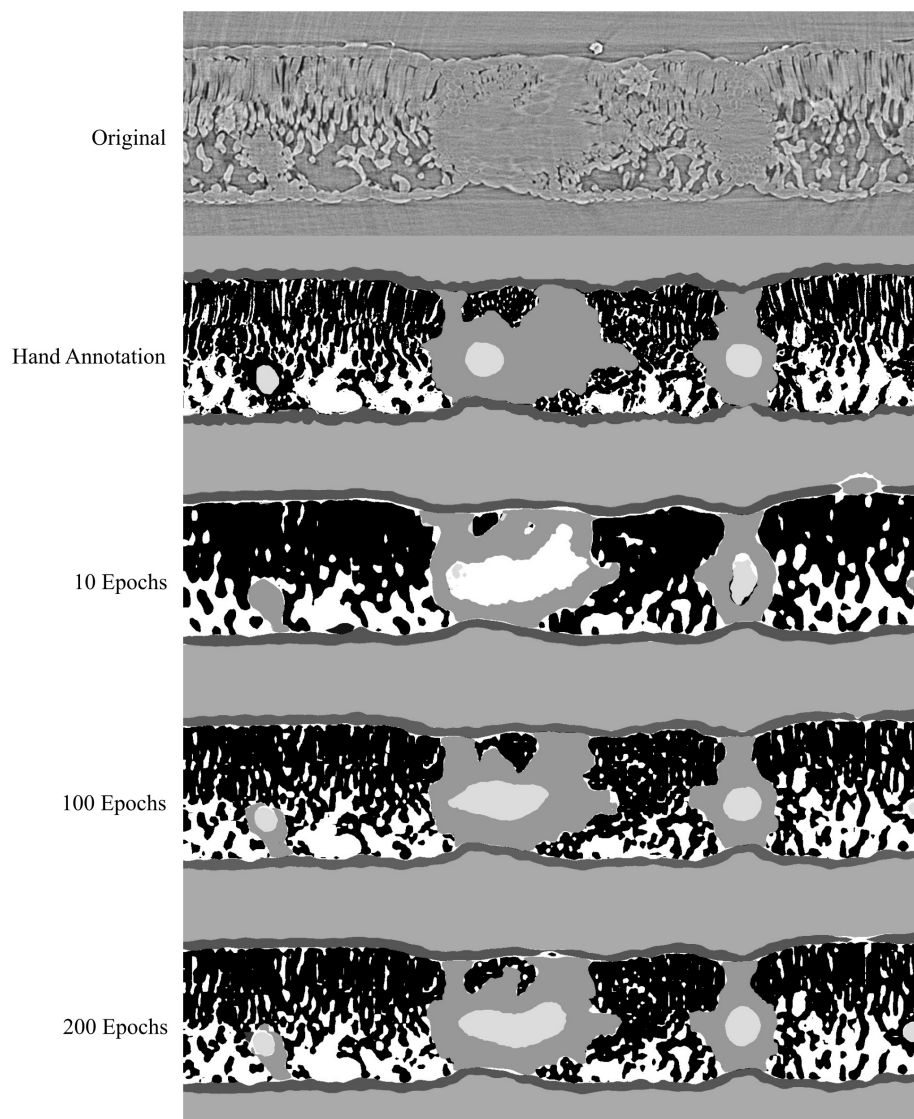
FIGURE 5

Visual representation of the model outputs for a single walnut leaf image cross-section; top image is a X-ray CT scan taken from a leaf that was not used for training or validation of the applied models; next is the hand annotated image of the scan followed by the outputs from the best performing model trained for 10, 100, and 200 epochs, respectively. For the walnut leaf segmentation, the background is light gray, the epidermis is dark gray, the mesophyll is black, the air space is white, the bundle sheath extensions are middle gray, and the veins are lightest gray.

best model F1 scores of 99.9, 99.4, 94.5, and 74.4% for the background, bud scale, flower, and air space classes, respectively. At the same scaling for the soil aggregate the best model F1 scores were 99.4, 91.3, 76.7, and 74.9% for the background, solid, pore, and organic matter classes, respectively. While the improvements in model performance with increased scaling were modest for the flower bud, they were large for the soil aggregate, likely due to the intricacy of the aggregate that was lost with downscaling (**Figure 6**).

Model outputs are binary 2-dimensional data from which information like material area, perimeter or other morphological traits can be extracted using downstream image analysis functions. These data are saved as sequences of arrays which allows for the 3d visualization of the segmented materials (**Figure 7**). Additionally, 3d data can be extracted from the array sequences using Python libraries like NumPy or using other programming languages like R or MATLAB.

Discussion

Generalized models for image segmentation are typically generated after training on hundreds to millions of images (Belthangady and Royer, 2019; O'Mahony et al., 2020;

Shahinfar et al., 2020; Khened et al., 2021). Due to the limited availability of leaf, bud, or soil x-ray CT scans, such a training image library simply doesn't exist (Moen et al., 2019). However, we found that we can generate accurate models using 5 annotated slices from at least 3 unique leaf scans. This discrepancy is likely the result of the consistent image collection settings (resolution) on the same imaging platform (Beamline 8.3.2) which simplifies the learning process (Fei et al., 2021; Silwal et al., 2021). Our results are comparable to those previously achieved by Théroux-Rancourt et al. (2020) on X-ray CT images of plant leaves. That method, based on random forest classification, requires hand annotating 6 images from every single scan and is designed for extracting data from leaf X-ray CT images exclusively. Specifically, the precision and recall scores for the background (>95%), mesophyll tissue (>80%), epidermis tissue (>80%), and bundle sheath extension (>75%) classes were equal to those achieved by Théroux-Rancourt et al. (2020). The current approach had lower precision and recall scores (>75%) for air space identification compared to Théroux-Rancourt et al. (2020) (>90%), but higher precision and recall scores for vein tissue identification (>60% vs. < 55%, respectively). Despite similarities in the quality of results, the current method decreases segmentation time from hours to minutes compared to Théroux-Rancourt et al. (2020) and can be applied to any X-ray CT image data set.

Our training batch size was limited to 1 by a combination of factors including variably sized training images for the leaf scans and hardware limitations for the almond bud and soil aggregate scans. Typically, batch size selection is constrained by a combination of hardware and the number of images used for training; the smaller the batch size, the longer training takes (Smith et al., 2017). This presents a significant barrier when training using millions of images but is not an issue when only tens of training images are available.

Epoch selection is an important component of maximizing model accuracy, precision, recall, and F1 scores. Training models for too few epochs leads to substandard model performance while over training with too many epochs wastes time and can lead to overfitting (Li et al., 2019; Baldeon Calisto and Lai-Yuen, 2020; Pan et al., 2020). Typically models trained with small batch sizes require more training epochs for satisfactory performance compared to models trained on large batch sizes (Smith, 2018). This can greatly increase training times if many images are required for training to achieve satisfactory model performance (von Chamier et al., 2021).

Beyond batch size and epoch number, image size plays a significant role in model performance. Large images take up more GPU memory than smaller images (Sabottke and Spieler, 2020). GPU memory can be conserved by decreasing batch size, but once batch size has been reduced to 1, the only option is to downscale images for training. However, this results in a significant loss of information in the images, hindering model performance as fine details are lost (Sabottke and Spieler, 2020).
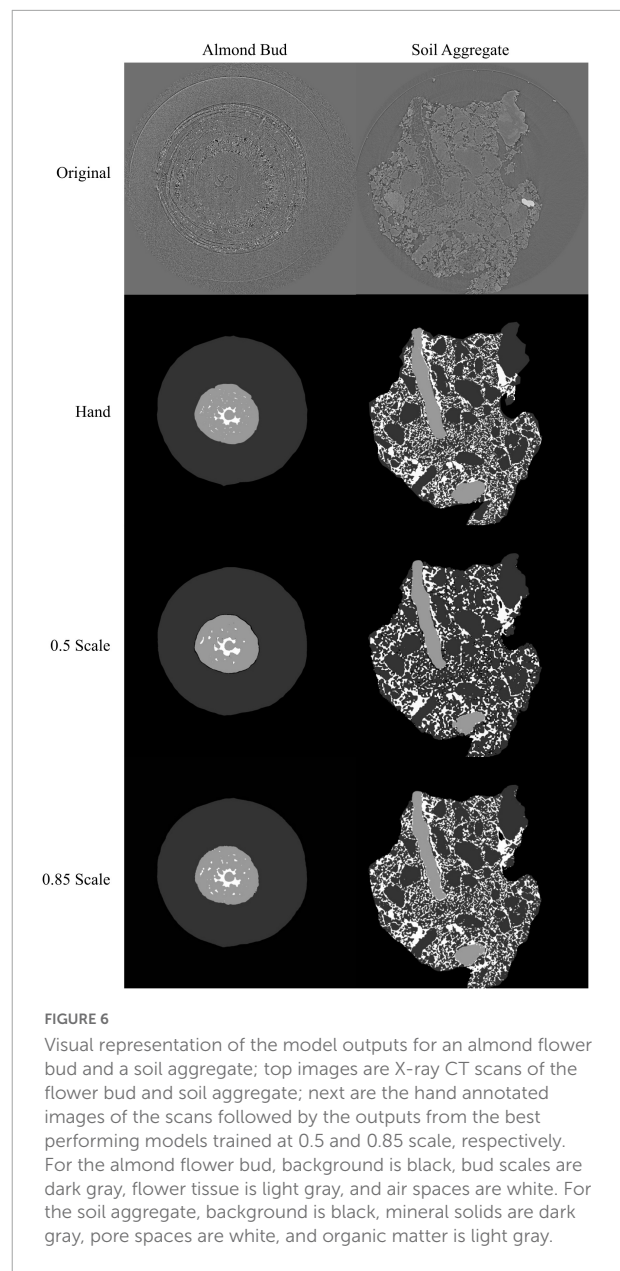
FIGURE 6

Visual representation of the model outputs for an almond flower bud and a soil aggregate; top images are X-ray CT scans of the flower bud and soil aggregate; next are the hand annotated images of the scans followed by the outputs from the best performing models trained at 0.5 and 0.85 scale, respectively. For the almond flower bud, background is black, bud scales are dark gray, flower tissue is light gray, and air spaces are white. For the soil aggregate, background is black, mineral solids are dark gray, pore spaces are white, and organic matter is light gray.

This was particularly problematic with the soil aggregate scans which contained fine pore spaces which were lost when the images were downscaled to 0.5 in the x and y dimensions to fit on GPUs with 16 Gb of VRAM. Downscaling to this degree results in a loss of 75% of the image information. Only when GPUs with more VRAM were used could images be downscaled less, resulting in improved performance for models trained on these higher resolution images.

By stacking the sequences of data arrays produced by the model, novel information can be gained about processes that occur over 3 dimensions. Taking this approach, Théroux-Rancourt et al. (2017) previously showed that mesophyll surface area exposed to intercellular air space is underestimated when
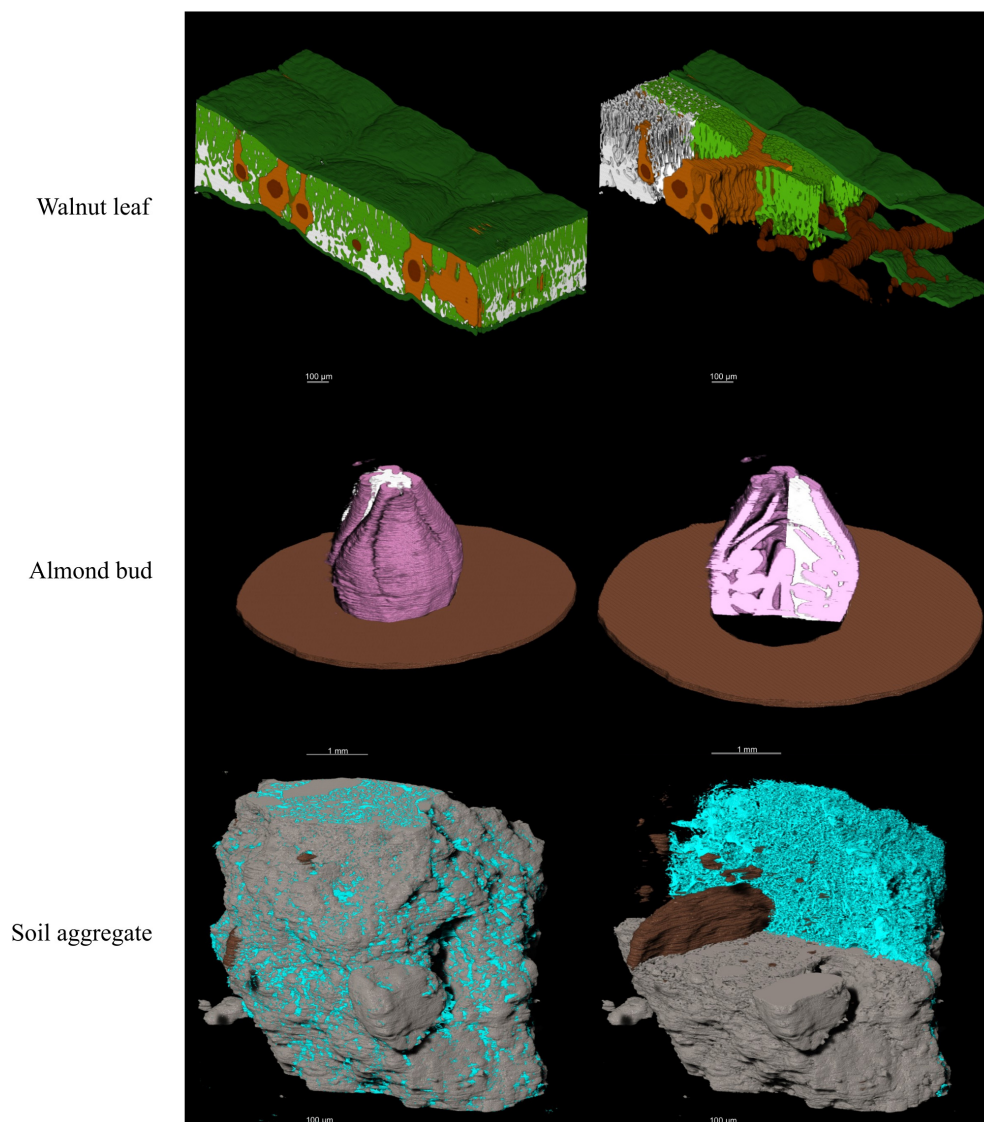
FIGURE 7
3D visual representation of the stacked model outputs for a walnut leaf, an almond flower bud, and a soil aggregate; top images are X-ray CT scans of the bud and soil aggregate; next are the hand. In the walnut leaf images, the background is black, the epidermis is dark green, the mesophyll is light green, the air space is white, the bundle sheath extensions burnt orange, and the veins are brown. In the almond bud images, background is black, bud scales are brown, flower tissue is pink, and air spaces are white. For the soil aggregate, background is black, mineral solids are dark gray, pore spaces are blue, and organic matter is brown.

using 2D rather than 3D approaches. Similarly, Trueba et al. (2021) found that the 3D organization of leaf tissues had a direct impact on plant water use and carbon uptake. In soils, it well understood that pore tortuosity plays a key role in understanding processes like water infiltration and O_2/CO_2 diffusion. As Baveye et al. (2018) highlighted in their review, the rapid segmentation of soil X-ray μCT data has long been a major hurdle to understanding these processes. Our workflow simplifies and accelerates this process, enabling researchers to rapidly extract information from their X-ray μCT data. Our approach is similar to those developed by Smith et al. (2020)

and von Chamier et al. (2021), but is more flexible as it works with variably sized images and allows for multi-label semantic segmentation.

Conclusion

With the current work, we present a workflow for using open-source software generate models to segment X-ray μCT images. These models can be specific to an image set (segmenting a single soil aggregate or almond bud) or be

generalized for a specific use case such as segmenting leaf scans. We demonstrated that a limited number of annotated images can achieve satisfactory results without excessively long training time. The workflow can be run locally, in Google's Colaboratory Notebook, or adapted for use on high performance computing platforms. By using GPU resources, the rate of segmentation can be dramatically increased, taking less than 0.02 s per image. This allows users to segment scans in minutes, a significant speed gain compared to other methods with similar precision and recall (often > 90%) across a variety of sample scans (Arganda-Carreras et al., 2017; Théroux-Rancourt et al., 2020). This will allow researchers to gain novel insights into the role that 3d architecture of soil and plant samples plays in a variety of important processes.

Data availability statement

The datasets presented in this study can be found on the National Agricultural Library Ag Data Commons website https://doi.org/10.15482/USDA.ADC/1524793.

Author contributions

AM, DR, JE, PR, EF, and DP contributed to the conception and design of the study. MM, FD, and KS annotated images. PR, DR, JE, JN, and AB wrote code for image segmentation and data extraction. DR wrote the first draft of the manuscript. MM helped write the "Materials and Methods" section of the manuscript. All authors contributed to the manuscript revision, read, and approved the submitted version.

Funding

Acknowledgments

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fpls.2022.893140/full#supplementary-material

SUPPLEMENTARY DATA SHEET 1
Precision, Recall, Accuracy, and F1 results for producing **Figures 2**, **3**.

SUPPLEMENTARY DATA SHEET 2
Precision, Recall, Accuracy, and F1 results for producing **Figures 4**, **5**.

References

Ahmed, S., Klassen, T. N., Keyes, S., Daly, M., Jones, D. L., Mavrogordato, M., et al. (2016). Imaging the interaction of roots and phosphate fertiliser granules using 4D X-ray tomography. *Plant Soil* 401, 125–134. doi: 10.1007/s11104-015-2425-5

Anderson, S. H., Peyton, R. L., and Gantzer, C. J. (1990). Evaluation of constructed and natural soil macropores using X-ray computed tomography. *Geoderma* 46, 13–29. doi: 10.1016/0016-7061(90)90004-S

Arganda-Carreras, I., Kaynig, V., Rueden, C., Eliceiri, K. W., Schindelin, J., Cardona, A., et al. (2017). Trainable Weka Segmentation: A machine learning tool for microscopy pixel classification. *Bioinformatics* 33, 2424–2426. doi: 10.1093/bioinformatics/btx180

Baldeon Calisto, M., and Lai-Yuen, S. K. (2020). AdaEn-Net: An ensemble of adaptive 2D–3D Fully convolutional networks for medical image segmentation. *Neural. Netw.* 126, 76–94. doi: 10.1016/j.neunet.2020.03.007

Baveye, P. C., Otten, W., Kravchenko, A., Balseiro-Romero, M., Beckers, É., Chalhoub, M., et al. (2018). Emergent properties of microbial activity in heterogeneous soil microenvironments: different research approaches are slowly converging, yet major challenges remain. *Front. Microbiol.* 9:1929. doi: 10.3389/fmicb.2018.01929

Belthangady, C., and Royer, L. A. (2019). Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction. *Nat. Methods* 16, 1215–1225. doi: 10.1038/s41592-019-0458-z

Brodersen, C. R., McElrone, A. J., Choat, B., Matthews, M. A., and Shackel, K. A. (2010). The dynamics of embolism repair in xylem: In vivo visualizations using high-resolution computed tomography. *Plant Physiol.* 154, 1088–1095. doi: 10.1104/pp.110.162396

Buslaev, A., Iglovikov, V. I., Khvedchenya, E., Parinov, A., Druzhinin, M., and Kalinin, A. A. (2020). Albumentations: Fast and flexible image augmentations. *Information* 11:125. doi: 10.3390/info11020125

Chen, L.-C., Papandreou, G., Schroff, F., and Adam, H. (2017). Rethinking atrous convolution for semantic image segmentation. *arXiv* [Preprint]. doi: 10.48550/arXiv.1706.05587

Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H. (2018). "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, 801–818. doi: 10.1007/978-3-030-01234-2_49

Crestana, S., Cesareo, R., and Mascarenhas, S. (1986). Using a computed tomography miniscanner in soil science. *Soil Sci.* 142:56. doi: 10.1097/00010694-198607000-00008

Crestana, S., Mascarenhas, S., and Pozzi-Mucelli, R. S. (1985). Static and dynamic three-dimensional studies of water in soil using computed tomographic scanning1. *Soil Sci.* 140, 326–332.

Cuneo, I. F., Barrios-Masias, F., Knipfer, T., Uretsky, J., Reyes, C., Lenain, P., et al. (2020). Differences in grapevine rootstock sensitivity and recovery from drought are linked to fine root cortical lacunae and root tip function. *New Phytol.* 229, 272–283. doi: 10.1111/nph.16542

Duncan, K. E., Czymmek, K. J., Jiang, N., Thies, A. C., and Topp, C. N. (2022). X-ray microscopy enables multiscale high-resolution 3D imaging of plant cells, tissues, and organs. *Plant Physiol.* 188, 831–845. doi: 10.1093/plphys/kiab405

Earles, J. M., Knipfer, T., Tixier, A., Orozco, J., Reyes, C., Zwieniecki, M. A., et al. (2018). In vivo quantification of plant starch reserves at micrometer resolution using X-ray microCT imaging and machine learning. *New Phytol.* 218, 1260–1269. doi: 10.1111/nph.15068

Fei, Z., Olenskyj, A. G., Bailey, B. N., and Earles, M. (2021). "Enlisting 3D crop models and GANs for more data efficient and generalizable fruit detection," in *Presented at the Proceedings of the IEEE/CVF International Conference on Computer Vision*, (Montreal, BC), 1269–1277.

Gao, Y., Liu, Y., Zhang, H., Li, Z., Zhu, Y., Lin, H., et al. (2020). "Estimating GPU memory consumption of deep learning models," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. Association for Computing Machinery*, (New York, NY), 1342–1352.

Gerth, S., Claußen, J., Eggert, A., Wörlein, N., Waininger, M., Wittenberg, T., et al. (2021). Semiautomated 3D root segmentation and evaluation based on X-ray CT imagery. *Plant Phenom.* 2021, 8747930. doi: 10.34133/2021/8747930

Gürsoy, D., De Carlo, F., Xiao, X., and Jacobsen, C. (2014). TomoPy: A framework for the analysis of synchrotron tomographic data. *J. Synchr. Radiat.* 21, 1188–1193. doi: 10.1107/S1600577514013939

Hapca, S. M., Wang, Z. X., Otten, W., Wilson, C., and Baveye, P. C. (2011). Automated statistical method to align 2D chemical maps with 3D X-ray computed micro-tomographic images of soils. *Geoderma* 164, 146–154. doi: 10.1016/j.geoderma.2011.05.018

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. *arXiv* [Priprint]. doi: 10.48550/arXiv.1512.03385

Helliwell, J. R., Sturrock, C. J., Grayling, K. M., Tracy, S. R., Flavel, R. J., and Young, I. M. (2013). Applications of X-ray computed tomography for examining biophysical interactions and structural development in soil systems: A review. *Eur. J. Soil Sci.* 64, 279–297.

Kamilaris, A., and Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Comput. Electr. Agric.* 147, 70–90. doi: 10.1016/j.compag.2018.02.016

Kemenade, H. V., Murray, A., Clark, A. W., Karpinsky, A., Baranoviè, O., Gohlke, C., et al. (2022). *python-pillow/Pillow: 9.0.1. Zenodo.* Available online at: https://doi.org/10.5281/zenodo.6788304

Keyes, S., van Veelen, A., McKay Fletcher, D., Scotson, C., Koebernick, N., Petroselli, C., et al. (2022). Multimodal correlative imaging and modelling of phosphorus uptake from soil by hyphae of mycorrhizal fungi. *New Phytol.* doi: 10.1111/nph.17980 [Epub ahead of print].

Khened, M., Kori, A., Rajkumar, H., Krishnamurthi, G., and Srinivasan, B. (2021). A generalized deep learning framework for whole-slide image segmentation and analysis. *Sci. Rep.* 11:11579. doi: 10.1038/s41598-021-90444-8

Kingma, D. P., and Ba, J. (2017). Adam: A Method for Stochastic Optimization. *arXiv* [Priprint]. doi: 10.48550/arXiv.1412.6980

Li, H., Li, J., Guan, X., Liang, B., Lai, Y., and Luo, X. (2019). "Research on overfitting of deep learning, in: 2019 15th International Conference on Computational Intelligence and Security (CIS)," in *Presented at the 2019 15th International Conference on Computational Intelligence and Security (CIS)*, 78–81. doi: 10.1109/CIS.2019.00025

Long, J., Shelhamer, E., and Darrell, T. (2015). "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (Boston, MA), 3431–3440.

Lotter, W., Diab, A. R., Haslam, B., Kim, J. G., Grisot, G., Wu, E., et al. (2021). Robust breast cancer detection in mammography and digital breast tomosynthesis using an annotation-efficient deep learning approach. *Nat. Med.* 27, 244–249. doi: 10.1038/s41591-020-01174-9

Mairhofer, S., Zappala, S., Tracy, S., Sturrock, C., Bennett, M. J., Mooney, S. J., et al. (2013). Recovering complete plant root system architectures from soil via X-ray μ-computed tomography. *Plant Methods* 9, 1–7. doi: 10.1186/1746-4811-9-8

Mathers, A. W., Hepworth, C., Baillie, A. L., Sloan, J., Jones, H., Lundgren, M., et al. (2018). Investigating the microstructure of plant leaves in 3D with lab-based X-ray computed tomography. *Plant Methods* 14, 1–12. doi: 10.1186/s13007-018-0367-7

McKinney, W. (2011). pandas: A foundational Python library for data analysis and statistics. *Python High Performance Sci. Comput.* 14, 1–9.

Moen, E., Bannon, D., Kudo, T., Graf, W., Covert, M., and Van Valen, D. (2019). Deep learning for cellular image analysis. *Nat. Methods* 16, 1233–1246. doi: 10.1038/s41592-019-0403-1

Mooney, S. J., Pridmore, T. P., Helliwell, J., and Bennett, M. J. (2012). Developing X-ray computed tomography to non-invasively image 3-D root systems architecture in soil. *Plant Soil* 352, 1–22.

O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Hernandez, G. V., Krpalkova, L., et al. (2020). "Deep learning vs. traditional computer vision," in *Advances in computer vision, advances in intelligent systems and computing*, eds K. Arai and S. Kapoor (Cham: Springer International Publishing), 128–144. doi: 10.1007/978-3-030-17795-9_10

Ofori, M., El-Gayar, O., O'Brien, A., and Noteboom, C. (2022). "A deep learning model compression and ensemble approach for weed detection," in *Proceedings of the 55th hawaii international conference on system sciences*, (Maui, HI: ScholarSpace).

Oliphant, T. E. (2007). Python for scientific computing. *Comput. Sci. Eng.* 9, 10–20.

Pan, Z., Emaru, T., Ravankar, A., and Kobayashi, Y. (2020). Applying semantic segmentation to autonomous cars in the snowy environment. *arXiv* [Preprint]. doi: 10.48550/arXiv.2007.12869

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *arXiv* [Preprint]. doi: 10.48550/arXiv.1912.01703

Petrovic, A. M., Siebert, J. E., and Rieke, P. E. (1982). Soil bulk density analysis in three dimensions by computed tomographic scanning. *Soil Sci. Soc. Am. J.* 46, 445–450.

Powers, D. M. (2020). Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation. *arXiv* [Preprint]. doi: 10.48550/arXiv.2010.16061

Raja, P., Olenskyj, A., Kamangir, H., and Earles, M. (2021). simultaneously predicting multiple plant traits from multiple sensors via deformable CNN regression. *arXiv* [Preprint]. doi: 10.48550/arXiv.2112.03205

Rippner, D., Pranav, R., Earles, J. M., Buchko, A., and Neyhart, J. (2022b). *Welcome to the workflow associated with the manuscript: A workflow for segmenting soil and plant X-ray CT images with deep learning in Google's Colaboratory.* Beltsville, MA: Ag Data Commons.

Rippner, D., Momayyezi, M., Shackel, K., Raja, P., Buchko, A., Duong, F., et al. (2022a). *X-ray CT data with semantic annotations for the paper "A workflow for segmenting soil and plant X-ray CT images with deep learning in Google's Colaboratory."* Washington, DC: Agricultural Research Service. doi: 10.15482/USDA.ADC/1524793

Ronneberger, O., Fischer, P., and Brox, T. (2015). "U-net: Convolutional networks for biomedical image segmentation," in *Proceeding of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, (Berlin: Springer), 234–241.

Sabottke, C. F., and Spieler, B. M. (2020). The Effect of Image Resolution on Deep Learning in Radiography. *Radiol. Artif. Intel.* 2:e190015. doi: 10.1148/ryai.2019190015

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., et al. (2012). Fiji: An open-source platform for biological-image analysis. *Nat. Methods* 9, 676–682. doi: 10.1038/nmeth.2019

Shahinfar, S., Meek, P., and Falzon, G. (2020). How many images do I need?" Understanding how sample size per class affects deep learning model performance metrics for balanced designs in autonomous wildlife monitoring. *Ecol. Inform.* 57:101085. doi: 10.1016/j.ecoinf.2020.101085

Silwal, A., Parhar, T., Yandun, F., and Kantor, G. (2021). A Robust Illumination-Invariant Camera System for Agricultural Applications. *arXiv* [Preprint]. doi: 10.48550/arXiv.2101.02190

Smith, A. G., Han, E., Petersen, J., Olsen, N. A. F., Giese, C., Athmann, M., et al. (2020). RootPainter: Deep learning segmentation of biological images with corrective annotation. *Bioarxiv* [Preprint]. doi: 10.1101/2020.04.16.044461

Smith, L. N. (2018). A disciplined approach to neural network hyper-parameters: Part 1 – learning rate, batch size, momentum, and weight decay. *arXiv* [Preprint]. doi: 10.48550/arXiv.1803.09820

Smith, S. L., Kindermans, P.-J., Ying, C., and Le, Q. V. (2017). Don't decay the learning rate, increase the batch size. *arXiv* [Preprint]. doi: 10.48550/arXiv.1711.00489

Théroux-Rancourt, G., Earles, J. M., Gilbert, M. E., Zwieniecki, M. A., Boyce, C. K., McElrone, A. J., et al. (2017). The bias of a two-dimensional view: comparing two-dimensional and three-dimensional mesophyll surface area estimates using noninvasive imaging. *New Phytol.* 215, 1609–1622. doi: 10.1111/nph.14687

Théroux-Rancourt, G., Jenkins, M. R., Brodersen, C. R., McElrone, A., Forrestel, E. J., and Earles, J. M. (2020). Digitally deconstructing leaves in 3D using X-ray microcomputed tomography and machine learning. *Appl. Plant Sci.* 8:e11380. doi: 10.1002/aps3.11380

Théroux-Rancourt, G., Roddy, A. B., Earles, J. M., Gilbert, M. E., Zwieniecki, M. A., Boyce, C. K., et al. (2021). Maximum CO_2 diffusion inside leaves is limited by the scaling of cell size and genome size. *Proc. Royal Soc. B* 288:20203145. doi: 10.1098/rspb.2020.3145

Torres-Ruiz, J. M., Jansen, S., Choat, B., McElrone, A. J., Cochard, H., Brodribb, T. J., et al. (2015). Direct X-ray microtomography observation confirms the induction of embolism upon xylem cutting under tension. *Plant Physiol.* 167, 40–43. doi: 10.1104/pp.114.249706

Tracy, S. R., Roberts, J. A., Black, C. R., McNeill, A., Davidson, R., and Mooney, S. J. (2010). The X-factor: Visualizing undisturbed root architecture in soils using X-ray computed tomography. *J. Exp. Bot.* 61, 311–313. doi: 10.1093/jxb/erp386

von Chamier, L., Laine, R. F., Jukkala, J., Spahn, C., Krentzel, D., and Nehme, E. (2021). Democratising deep learning for microscopy with ZeroCostDL4Mic. *Nat. Commun.* 12:2276. doi: 10.1038/s41467-021-22518-0

Waskom, M. L. (2021). Seaborn: Statistical data visualization. *J. Open Source Softw.* 6, 3021.

Xiao, Z., Stait-Gardner, T., Willis, S. A., Price, W. S., Moroni, F. J., Pagay, V., et al. (2021). 3D visualisation of voids in grapevine flowers and berries using X-ray micro computed tomography. *Austr. J. Grape Wine Res.* 27, 141–148. doi: 10.1111/ajgw.12480

Yudina, A., and Kuzyakov, Y. (2019). Saving the face of soil aggregates. *Glob. Change Biol.* 25, 3574–3577. doi: 10.1111/gcb.14779