

1

Root Anatomy based on Root Cross-Section Image Analysis with Deep Learning

Chaoxin Wang¹, Xukun Li¹, Doina Caragea^{1,*}, Raju Bheemanahalli² and S.V. Krishna Jagadish²

¹*Department of Computer Science, Kansas State University, Manhattan, KS, USA*

²*Department of Agronomy, Kansas State University, Manhattan, KS, USA*

Correspondence*:
Doina Caragea
dcaragea@ksu.edu

2 ABSTRACT

3 The aboveground plant efficiency has improved significantly in recent years, and the
4 improvement has led to a steady increase in global food production. The improvement of
5 belowground plant efficiency has the potential to further increase food production. However,
6 the belowground plant roots are harder to study, due to inherent challenges presented by
7 root phenotyping. Several tools for identifying root anatomical features in root cross-section
8 images have been proposed. However, the existing tools are not fully automated and require
9 significant human effort to produce accurate results. To address this limitation, we propose a fully
10 automated approach, called Deep Learning for Root Anatomy (DL-RootAnatomy), for identifying
11 anatomical traits in root cross-section images. Using the Faster Region-based Convolutional
12 Neural Network (Faster R-CNN), the DL-RootAnatomy models detect objects such as root, stele
13 and late metaxylem, and predict rectangular bounding boxes around such objects. Subsequently,
14 the bounding boxes are used to estimate the root diameter, stele diameter, and late metaxylem
15 number and average diameter. Experimental evaluation using standard object detection metrics,
16 such as intersection-over-union and mean average precision, has shown that our models can
17 accurately detect the root, stele and late metaxylem objects. Furthermore, the results have shown
18 that the measurements estimated based on predicted bounding boxes have very small root
19 mean square error when compared with the corresponding ground truth values, suggesting that
20 DL-RootAnatomy can be used to accurately detect anatomical features. Finally, a comparison
21 with existing approaches, which involve some degree of human interaction, has shown that
22 the proposed approach is more accurate than existing approaches on a subset of our data. A
23 webserver for performing root anatomy using our deep learning pre-trained models is available at
24 <https://rootanatomy.org>, together with a link to a GitHub repository that contains code that can
25 be used to re-train or fine-tune our network with other types of root-cross section images. The
26 labeled images used for training and evaluating our models are also available from the GitHub
27 repository.

28 **Keywords:** Image Analysis, Deep Learning, Object Detection, Faster R-CNN, Root Anatomy

1 INTRODUCTION

29 The crop scientific community has made significant strides in increasing global food production through
30 advances in genetics and management, with majority of the progress achieved by improving aboveground
31 plant efficiency (Araus et al., 2008; Khush, 2013; Bishopp and Lynch, 2015). The belowground plant roots,
32 which provide water and nutrients for plant growth, are relatively less investigated. This is primarily because
33 of the difficulty in accessing the roots, and the complexity of phenotyping root biology and function (Jung
34 and Mccouch, 2013; E. Schmidt and C.M. Gaudin, 2017). Hence, root potential has largely been untapped
35 in crop improvement programs (Jung and Mccouch, 2013; E. Schmidt and C.M. Gaudin, 2017). Over the
36 past decade, different root phenotyping approaches have been developed for studying root architecture, e.g.,
37 basket method for root angle (Uga et al., 2013), rhizotron method for tracking root branching, architecture
38 and growth dynamics (Bucksch et al., 2014), shovelomics, a.k.a., root crown phenotyping (Colombi et al.,
39 2015), among others. Recent advances in magnetic resonance imaging and X-ray computed tomography
40 detection systems have provided the opportunity to investigate root growth dynamics in intact plants at
41 high temporal frequency (Mooney et al., 2012; Schulz et al., 2013; Topp et al., 2013; van Dusschoten et al.,
42 2016; Pfeifer et al., 2015). However, each of these techniques comes with a range of inherent biases or
43 limitations (such as artificial plant growth conditions), with none of the techniques currently available
44 clearly standing out as a promising approach that could become a blanket fit (Durham Brooks et al., 2010;
45 Clark et al., 2011; Sozzani et al., 2014). The recent non-destructive technologies, such as X-ray computed
46 tomography, are extremely expensive, and thus beyond the reach of common crop improvement programs,
47 in addition to not having the bandwidth to capture large genetic diversity.

48 Machine learning approaches have been used successfully to address a wide variety of biological problems,
49 including problems relevant to crop sciences, such as genome annotation (Yip et al., 2013), predicting
50 gene functions (Rhee and Mutwil, 2014), discovery of genetic variation and genotyping (DePristo et al.,
51 2011), identification of genomic regions of interest (Topp et al., 2013; Heslot et al., 2012), high throughput
52 phenotyping based on aerial image analysis (Khan et al., 2018). Furthermore, applications of advanced deep
53 learning and image analysis techniques to challenging problems in crop analysis have led to state-of-the-art
54 results that outperform the results of traditional machine learning and image analysis techniques (Kamilaris
55 and Prenafeta-Boldú, 2018; Jones et al., 2017).

56 Most relevant to this work, machine learning, and more specifically, deep learning, are expanding the
57 ability to accurately predict a plant phenotype (Tardieu et al., 2017; Singh et al., 2016; Pound et al.,
58 2017a; Namin et al., 2017; Aich and Stavness, 2017; Dobrescu et al., 2017; Ubbens and Stavness, 2017;
59 Namin et al., 2017), enabling the researchers to capture a wide range of genetic diversity, a task which has
60 been hardly possible in the past, given the amount of time and effort involved in manual analysis (Singh
61 et al., 2016). Several recent studies have used use deep learning regression approaches for identifying
62 and quantifying aboveground plant traits such as the number of leaves in rosette plants based on high-
63 resolution RBF images (Aich and Stavness, 2017; Dobrescu et al., 2017; Ubbens and Stavness, 2017). Other
64 investigations have focused on identifying diseases (Mohanty et al., 2016) or phenotyping for stress/nutrient
65 deficiencies (Singh et al., 2016).

66 Furthermore, several prior studies have focused on data-driven approaches and tools for belowground
67 plant phenotyping, including identifying and quantifying root morphological parameters, such as changes
68 in root architecture, or branching and growth (Kuijken et al., 2015; Jiansan et al., 2017; Delory et al.,
69 2018; Pound et al., 2017b; Betegón-Putze et al., 2018; Reeb et al., 2018). Such approaches generally rely
70 on standard image analysis techniques as opposed to state-of-the-art machine learning approaches.

71 While the study of root morphological parameters is important in relation to the health and productivity
72 of crops, the study of root anatomical parameters, such as stele and the xylem vessels, is equally important.
73 Root anatomical parameters represent the conduits for transport of water and nutrients to the plant's
74 aboveground parts. Hence, they are significantly affected by different rhizosphere conditions, and in turn,
75 affect crop productivity (E. Schmidt and C.M. Gaudin, 2017).

76 Innovations in image acquisition technologies have made it possible to gather relatively large sets of root
77 cross-section images, enabling studies on root anatomy. Several approaches and tools for quantifying root
78 anatomical variation based on cross-section images have been proposed in recent years (Burton et al., 2012;
79 Chopin et al., 2015; Lartaud et al., 2015). However, the existing tools are only partially automated, as they
80 require user input and fine-tuning of the parameters for each specific image or for a batch of images. Fully
81 automated tools exist for the analysis of hypocotyl cross-sections (i.e., the region in between seed leafs and
82 roots) in the context of secondary growth (Hall et al., 2016; Sankar et al., 2014), but they are not directly
83 applicable to the analysis of root cross-section images. Thus, there is a pressing need for automated root
84 cross-section image analysis tools that can use to perform root anatomy at a low cost.

85 To address this limitation, we have taken advantage of recent advances in deep learning and image analysis,
86 and developed a state-of-the-art, fully-automated deep learning approach for identifying and quantifying
87 root anatomical parameters, indicative of the physiological and genetic responses of root anatomical
88 plasticity in field crops. Specifically, we have considered the following parameters: root diameter (RD),
89 stele diameter (SD), late metaxylem diameter (LMXD) and late metaxylem number (LMXN), which were
90 studied and found important in relation to water-deficit stress in our prior work (Kadam et al., 2015, 2017).
91 A graphical illustration of these parameters is shown in Figure 1.

92 Our proposed approach is based on the Faster R-CNN network (Ren et al., 2015), and can be used to
93 produce models that can detect objects of interest in a root cross-section image (i.e., root, stele and late
94 metaxylem), together with their corresponding bounding boxes. Subsequently, the bounding boxes can
95 be used to estimate anatomical parameters such as RD, SD, LMXD, LMXN. Once trained, our models
96 generalize well to unseen images, thus eliminating the need for the end-user to hand-draw a stele border or
97 manually choose the metaxylem cells, tasks that are time-consuming, and also prone to noise and errors.

98 To summarize, our contributions are as follows:

- 99 • We have proposed an approach based on Faster R-CNN to detect root, stele and late metaxylem objects,
100 and their corresponding bounding boxes, in root cross-section images.
- 101 • We have investigated the Faster R-CNN models with respect to the number of instances needed to
102 accurately detect the objects of interest, and their corresponding bounding boxes.
- 103 • We have evaluated the ability of the predicted bounding boxes to produce accurate estimates for
104 anatomical properties, and performed error analysis to identify sources of errors.
- 105 • We have compared the results of the proposed fully-automated DL-RootAnatomy approach with the
106 results obtained from existing approaches in terms of accuracy.

2 MATERIALS AND METHODS

107 While there are many anatomical traits that can be identified, and measured or counted (e.g., RootScan
108 outputs more than 20 anatomical parameters), as a proof-of-concept, we have focused on measuring the
109 root diameter (RD), stele diameter (SD), and late metaxylem diameter (LMXD), and counting the number
110 of late metaxylem inside the stele (LMXN), as motivated by Kadam et al. (2015; 2017), who showed the

111 importance of these traits in relation to water-deficit stress. The tasks that we target can be achieved with
112 modern object detection techniques, as described below.

113 2.1 Proposed Approach

114 We have used Faster R-CNN (Ren et al., 2015), a state-of-the-art network for object detection, to detect
115 objects of interest (i.e., root, stele, late metaxylem), and subsequently mark each object with a bounding
116 box. More precisely, we have trained a Faster R-CNN model to identify the root and stele within a
117 cross-section image, and another similar model to identify the late metaxylem within the stele region of
118 a cross-section. Given the bounding box of an object, identified by the Faster R-CNN model, we have
119 calculated its diameter by averaging the width and height of the bounding box. The count of late metaxylem
120 was obtained by counting the number of late metaxylem objects detected by the Faster R-CNN network.

121 The proposed Faster R-CNN model architecture used to detect the root and stele in a root cross-section
122 image is shown in Figure 2. Faster R-CNN has two main components. The first component consists of a
123 Region Proposal Network (RPN), which identifies Regions of Interest (which potentially contain objects of
124 interest), and also their location. The second component consists of a Fast R-CNN (Girshick, 2015), which
125 classifies the proposed regions (i.e., objects) into different classes (e.g. root and stele), and also refines the
126 location parameters to generate an accurate bounding box for each detected object. The two components
127 share the convolutional layers of VGG-16 (Simonyan and Zisserman, 2014), which is used as the backbone
128 of the Faster R-CNN model. More details on convolutional neural networks, VGG-16 and Faster R-CNN
129 approach used to detect objects and generate bounding boxes are provided below.

130 2.1.1 Convolutional Neural Networks and VGG-16

131 Convolutional Neural Networks (CNNs) (LeCun et al., 1989) are widely used in image analysis. While
132 originally designed for image classification, the features extracted by CNNs are informative for other image
133 analysis tasks, including object detection. A CNN consists of convolutional layers followed by non-linear
134 activations, pooling layers and fully connected layers, as seen in the example in Figure 3 (which shows
135 a specific CNN architecture called VGG-16 (Simonyan and Zisserman, 2014)). A convolutional layer
136 uses a sliding window approach to apply a set of filters (low-dimensional tensors) to the input image. The
137 convolution operation captures local dependencies in the original image, and it produces a feature map.
138 Different filters produce different feature maps, consisting of different features of the original image (e.g.,
139 edges, corners, etc.). A convolution layer is generally followed by a non-linear activation function, such
140 as the Rectified Linear Unit (i.e., ReLU), applied element-wise to generate a rectified feature map. The
141 ReLU activation replaces all the negative pixels in a feature map with zero values. A pooling layer is used
142 to reduce the dimensionality of the rectified feature map. Intuitively, the pooling operation retains the
143 most important information in a feature map by taking the maximum or the average pixel in each local
144 neighborhood of the feature map. As a consequence, the feature map becomes equivariant to scale and
145 translation (LeCun et al., 2015).

146 After a sequence of convolutional layers (together with non-linear activations) and pooling layers, a
147 CNN has one or more fully connected layers. In fully connected layers all neurons in the current layer
148 are connected to all neurons in the next layer. The first fully connected layer is connected to the last
149 downsized feature map. The fully connected layers are used to further reduce the dimensionality and to
150 capture non-linear dependencies between features (LeCun et al., 2015). The last fully connected layer uses
151 a softmax activation function, and has as many output neurons as the number of targeted classes.

152 There are several pre-trained CNN architectures available, including VGG-16 (Simonyan and Zisserman,
153 2014), shown in Figure 3. VGG-16 has been shown to give very good performance in the ImageNet
154 competition, where the network was trained on millions on images with 1000 categories (Simonyan and
155 Zisserman, 2014). Furthermore, VGG-19 was used with good results in the original Faster R-CNN study
156 (Ren et al., 2015), which motivated us to use it also in our model. As can be seen in Figure 3, VGG-
157 16 has 13 convolutional+ReLU layers, 5 pooling layers, and 3 fully connected layers. The dimensions
158 corresponding to each layer are also shown in Figure 3.

159 2.1.2 Region Proposal Network (RPN)

160 As mentioned above, the region proposal network identifies regions that could potentially contain objects
161 of interest, based on the last feature map of the pre-trained convolutional neural network that is part of the
162 model, in our case VGG-16 (Simonyan and Zisserman, 2014). More specifically, using a sliding window
163 approach, k regions are generated for each location in the feature map. These regions, are represented as
164 boxes called *anchors*. The anchors are all centered in the middle of their corresponding sliding window,
165 and differ in terms of scale and aspect ratio (Ren et al., 2015), to cover a wide variety of objects. The
166 region proposal network is trained to classify an anchor (represented as a lower-dimensional vector) as
167 containing an object of interest or not (i.e., it outputs an “objectness” score), and also to approximate
168 the four coordinates of the object (a.k.a., location parameters). The ground truth used to train the model
169 consists of bounding boxes provided by human annotators. If an anchor has high overlap with a ground
170 truth bounding box, then it is likely that the anchor box includes an object of interest, and it is labeled as
171 positive with respect to the *object* versus *no object* classification task. Similarly, if an anchor has small
172 overlap with a ground truth bounding box, it is labeled as negative. Anchors that don’t have high or small
173 overlap with a ground truth bounding box are not used to train the model. During training, the positive and
174 negative anchors are passed as input to two fully connected layers corresponding to the classification of
175 anchors as containing *object* or *no object*, and to the regression of location parameters (i.e., four bounding
176 box coordinates), respectively. Corresponding to the k anchors from a location, the RPN network outputs
177 $2k$ scores and $4k$ coordinates.

178 2.1.3 Fast R-CNN

179 Anchors for which the RPN network predicts high “objectness” scores are passed to the last two layers
180 (corresponding to object classification and location parameter refinement, respectively) of a network that
181 resembles the original Fast R-CNN network (Girshick, 2015), except for how the proposed regions are
182 generated. Specifically, in the original Fast R-CNN, the regions were generated from the original image
183 using an external region proposal method (e.g., selective search).

184 As opposed to the original Fast R-CNN (Girshick, 2015), in the Fast R-CNN component of the Faster R-
185 CNN model, the external region proposal method is replaced by an internal RPN, described in the previous
186 subsection, which is trained to identify regions of interest (Ren et al., 2015). Highly overlapping regions,
187 potentially corresponding to the same object, can be filtered using a non-maximum suppression (NMS)
188 threshold. A pooling layer is used to extract feature vectors of fixed length for the regions of the interest
189 proposed by RPN. Subsequently, the feature vectors are provided as input to two fully connected layers,
190 corresponding to the classification of the object detected and the regression of its location, respectively.

191 The object classification layer in Fast R-CNN uses the softmax activation, while the location regression
192 layer uses linear regression over the coordinates defining the location as a bounding box. All parameters of
193 the network are trained together using a multi-task loss (Girshick, 2015).

194 2.1.4 Faster R-CNN Training

195 The Fast R-CNN network and the region proposal network share several convolutional layers, specifically
196 the 13 convolutional layers of VGG-16. We have initialized the parameters of the 13 convolutional layers
197 using the VGG-16 network pre-trained on the ImageNet dataset. As many image features are highly
198 transferable between different datasets, this initialization based on VGG-16 allowed us to learn accurate
199 models from a relatively small number of root cross-section labeled images.

200 Given that the region proposal network and the Fast R-CNN network share 13 convolutional layers, they
201 are co-trained using an iterative process that alternates between fine-tuning the RPN and fine-tuning the
202 Fast R-CNN network (with fixed proposed regions produced by RPN) (Ren et al., 2015). All the model
203 parameters are updated using stochastic gradient descent (SGD).

204 In our preliminary experimentation, we found that it is difficult to accurately detect the late metaxylem at
205 the same time with the root and stele. Thus, we have trained a Faster R-CNN model to detect root and stele,
206 and another Faster R-CNN model to detect late metaxylem.

207 2.2 Existing Approaches

208 There are several approaches and tools for quantifying root anatomical variation based on cross-section
209 images (Burton et al., 2012; Chopin et al., 2015; Lartaud et al., 2015). Approaches in this category can
210 be roughly categorized as manual, semi-automated, and automated approaches. Manual analysis of root
211 images relies heavily on subjective assessments, and is suitable only for low throughput analysis. ImageJ
212 (Schneider et al., 2012) is an image analysis tool that has been extensively used to manually identify and
213 quantify root anatomical traits (Kadam et al., 2015; Yamauchi et al., 2013; Kadam et al., 2017), given that
214 it enables researchers to mark objects of interest and obtain their measurements. In particular, the ImageJ
215 software was used to acquire the original ground truth annotations for the images used in this study.

216 Semi-automated tools require user feedback to tune parameters for individual images in order to get
217 accurate results. *RootScan* (Burton et al., 2012) and *PHIV-RootCell* (Lartaud et al., 2015) are semi-
218 automated tools that identify and quantify anatomical root traits. *RootScan* was originally designed for
219 analyzing maize root cross-section images. The analysis of each image involves several steps. *RootScan*
220 starts by isolating the cross-section from the background using a global thresholding technique (Otsu,
221 1979). Subsequently, the stele is segmented based on the contrast between pixel intensities within and
222 outside the stele. Different cells within the stele (e.g., late metaxylem) are classified based on their area
223 according to background knowledge on root anatomy for a particular species. After each step, the user has
224 to “approve” the classification performed automatically or alternatively correct it, before moving to the
225 next step. The tool can be run on a set of images in batch mode, but the user still needs to provide input for
226 each step of the analysis for each image, as explained above. The output of *RootScan* consists of a table
227 with area measurements and counts of different anatomical traits.

228 The *PHIV-RootCell* tool for root anatomy is built using the ImageJ software (Schneider et al., 2012),
229 and provides options for selecting regions of interest (ROI) such as root, stele, xylem, and for measuring
230 properties of these regions. It was designed for analyzing rice root cross-section images. Similar to
231 *RootScan*, domain knowledge is used to identify ROIs. The *PHIV-RootCell* tool uploads and analyzes one
232 image at a time, and does not have an option for batch uploading or processing. Furthermore, it requires
233 user’s supervision at each segmentation and classification step (Lartaud et al., 2015). For example, it
234 requires the user to validate the root selection, stele selection, central metaxylem selection, among others.

235 As opposed to semi-automated tools that require user feedback, a fully automated approach should
236 involve “a single click” and should produce accurate results without any human intervention during the
237 testing and evaluation phases. However, human input and supervision in the form of background knowledge
238 or labeled training examples may be provided during the training phase. In this sense, *RootAnalyzer*
239 (Chopin et al., 2015) is an automated tool, which incorporates background knowledge about root anatomy.
240 The first step in *RootAnalyzer* is aimed at performing image segmentation to distinguish between root
241 pixels (corresponding to boundaries of individual root cells) and background pixels. To achieve this,
242 *RootAnalyzer* utilizes a local thresholding technique to analyze each pixel’s intensity by comparing it with
243 the mean pixel intensity in a small square neighborhood around that pixel (defined by a width parameter,
244 W). Subsequently, *RootAnalyzer* constructs a difference image, and classifies pixels as root or background
245 pixels based on a threshold, T , used on the difference image. The next step is focused on detecting root
246 cells and closing small leaks in cell boundaries, using an interpolation approach. Finally, cells are classified
247 in different categories, such as stele cells, cortex cells, epidermal cells, etc. based on size, shape, and
248 position. Two thresholds are used to classify cells as small or large: a threshold, A_s , for small cells, and a
249 threshold, A_l , for large cells. Furthermore, stele cells are classified based on an additional threshold, N ,
250 on the maximum distance from a cell to any of its nearest neighbor cells. The *RootAnalyzer* tool allows
251 for both single image processing and batch processing. Single image processing allows the user to adjust
252 and tune parameters, and also to interact with the tool at each stage of the segmentation and classification.
253 Batch processing requires the user to provide the parameters to be used with a specific batch of plant
254 images. Similar to *RootScan*, *RootAnalyzer* outputs a table of area measurements and counts for regions of
255 interest. This tool was designed for wheat and also tested on maize (Chopin et al., 2015).

256 2.3 Dataset

257 Twenty-five accessions of *Oryza* species were grown in plastic pots (25 cm in height; 26 and 20 cm
258 diameter at the top and bottom, respectively), filled with 6 kg of clay loam soil. Three replications per each
259 accession were maintained under well-watered conditions and roots were sampled 60 days after sowing,
260 to obtain fully mature roots. The roots were harvested and washed thoroughly and stored in 40% alcohol.
261 To obtain the cross-section images used in this study, root samples stored in 40% alcohol were hand
262 sectioned with a razor blade using a dissection microscope. Images of root sections were acquired with the
263 Axioplan 2 compound microscope (Zeiss, Germany) at 50x and 100x magnification. For each of the 25 rice
264 accessions included in the study, three biological replicate root samples from root-shoot junction and 6 cm
265 from the root tip were obtained. From each replicate, 2-3 images were taken at root-shoot junction, and 2-3
266 images at 6 cm from the tip of the root. Images may have two versions: a 50× magnification version, which
267 captures the whole root diameter (top image in Fig. 1), and a 100× magnification version, which captures
268 only the stele diameter (bottom image in Fig. 1). However, not all 50× images have a 100× correspondent.
269 Specifically, there are 388 images at 50× magnification, and 339 images at 100× magnification.

270 For each root image, we manually measured root anatomical parameters, such as root cross-section
271 diameter, stele diameter, late metaxylem average diameter and late metaxylem number, using the ImageJ
272 software (Schneider et al., 2012). The manual measurements and counts constitute our ground truth to
273 which we compared the measurements produced by our models. Statistics about the dataset, including the
274 minimum, maximum, average and standard deviation for root diameter, stele diameter, late metaxylem
275 average diameter and late metaxylem number, respectively, are presented in Table 1.

276 In addition to measuring root anatomical parameters, each 50× magnification image was also manually
277 labeled by independent annotators with bounding boxes that represent root, stele, and late metaxylem,

Table 1. Ground Truth Statistics: minimum (Min), maximum (Max), and average together with standard deviation (Avg \pm std) are shown for the ground truth measurements (RD, SD, LMXD and LMXN).

Statistics	RD	SD	LMXD	LMXN
Min	354	115	15	1
Max	1352	419	65	12
Avg \pm std	869 \pm 194	216 \pm 55	36 \pm 8	5.4 \pm 1.8

278 respectively, and each $100\times$ magnification image was labeled with boxes that represent the late metaxylem.
279 We used the LabelImg tool (Tzutalin, 2015) to perform the bounding box labeling. This tool produces
280 annotations in the Pascal Visual Object Classes (VOC) XML format (Everingham et al., 2015), a standard
281 format used for annotating images with rectangular bounding boxes corresponding to objects. The bounding
282 boxes in the $50\times$ and $100\times$ magnification images constitute the ground truth to which we compared the
283 bounding boxes of the objects detected by our models.

284 We should note that the $50\times$ magnification images contain all the anatomical features that we target
285 in this study, and are sufficient for training the proposed deep learning models. However, we also train
286 models from the $100\times$ magnification images, independently, to understand how much the identification of
287 the LMX objects and their measurements may be improved by using images with a higher resolution. In
288 general, any resolution can be used for training, as long as all the features that need to be identified are
289 contained in the image.

290 2.4 Experimental Setup

291 2.4.1 Training, Development and Test Datasets

292 We performed a set of experiments using 5-fold cross-validation. Specifically, we split the set of $50\times$
293 magnification images into five folds, based on accessions, such that each fold contained 5 accessions out
294 of the 25 accessions available. The exact number of $50\times$ magnification images (instances) in each fold
295 is shown in Table 2. For each fold, Table 2 also shows the number of corresponding $100\times$ magnification
296 images (instances) available (note that not every $50\times$ magnification image has a corresponding $100\times$
297 magnification image). In each 5-fold cross-validation experiment, four folds were used for training, and the
298 fifth fold was used for test. To tune hyper-parameters, we used one of the training folds as the development
299 dataset. The results reported were averaged over the 5 folds. The reason for splitting the set of images based
300 on accessions was to avoid using images from the same plant or the same replicate both in the training and
301 test datasets.

Table 2. Number of instances in each of the 5 folds used to perform cross-validation for the $50\times$ and $100\times$ magnification images, respectively. The total number of instances in the dataset is also shown.

Fold	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Total
Instances ($50\times$)	71	79	86	77	75	388
Instances ($100\times$)	62	60	80	69	68	339

302 2.4.2 Evaluation Metrics

303 We used three standard metrics in our evaluation, driven by preliminary observations. First, given that
304 there exist exactly one root and one stele in an image, we observed that these objects are always detected in
305 the $50\times$ magnification images. We used the Intersection-over-Union (IoU) metric to measure how well the
306 predicted bounding boxes overlap with the ground truth bounding boxes. Second, given that the number

307 of LMX objects varies between 1 and 12, and these objects are relatively small, the corresponding object
308 detection models are prone to both false positive and false negative mistakes. Thus, we used mean average
309 precision (mAP), a standard metric in object detection, to evaluate the ability of our models to accurately
310 identify the LMX objects. Both IoU and mAP metrics range between 0 and 1, and higher values are better.
311 Finally, we used the root mean square error (RMSE) metric to measure the ability of the proposed approach
312 to detect objects and corresponding bounding boxes that lead to root/stele/LMX diameter measurements
313 and LMX counts close to those available as ground truth. For RMSE, smaller values are better.

314 2.4.3 Hyper-parameter Tuning

315 Deep learning models, in general, and the Faster R-CNN models, in particular, have many tunable hyper-
316 parameters. We tuned several hyper-parameters shown by Zhang et al. (2016) to affect the performance
317 of the Faster R-CNN models, and used the values suggested by Ren et al. (2015) for the other hyper-
318 parameters. More specifically, we tuned the IoU threshold used in the RPN network to identify anchors
319 that could potentially include an object of interest (i.e., positive instances/anchors). Furthermore, we
320 tuned the non-maximum suppression (NMS) threshold which is used to filter region proposals produced
321 by the trained RPN network (specifically, if two proposals have IoU larger than the NMS threshold, the
322 two proposals will be considered to represent the same object). At last, we tuned the fraction of positive
323 instances in a mini-batch.

324 The specific values that we used to tune the IoU threshold were 0.4, 0.5 and 0.6, the valued used to tune
325 the NMS threshold were 0.6, 0.7 and 0.8, and the values used to tune the fraction of positive instances in
326 a mini-batch were 1:5 and 1:4. To observe the variation of performance with the tuned parameters, and
327 select the values that gave best performance, we trained a model corresponding to a particular combination
328 of parameters on three training folds, and evaluated the performance of the model on the development
329 fold. The performance of the models for root and stele detection was measured using the IoU metric (by
330 comparing the predicted bounded boxes with the ground truth bounded boxes), while the performance
331 of the models for LMX detection was measured using the mAP metric (by comparing the detected LMX
332 objects with the ground truth LMX objects) to ensure that the Faster R-CNN models can accurately detect
333 all the LMX objects.

334 Our tuning process revealed that the performance did not vary significantly with the parameters for our
335 object detection problem. However, the best combination of parameters for the root/stele models consisted
336 of the following values: 0.4 for the IoU threshold, 0.8 for the NMS threshold and 1:4 for the fraction of
337 positive anchors in a mini-batch. The best combination of parameters for the LMX models consisted of
338 the following values: 0.5 for the IoU threshold, 0.8 for the NMS threshold, and 1:4 for the fraction of
339 positive anchors in a mini-batch. We used these combinations of values for the root/stele and LMX models,
340 respectively, in the our experiments described in the next section.

3 RESULTS AND DISCUSSION

341 In this section, we present and discuss the results of our proposed approach, *DL-RootAnatomy*, and
342 also the results of a comparison between our approach and two related approaches, *RootAnalyzer* and
343 *RootScan*. Furthermore, we outline time requirements for DL-RootAnatomy and discuss the availability of
344 our approach as a tool.

345 **3.1 Variation of Performance with the Number of Training Instances**

346 As opposed to the existing tools for identifying anatomical parameters in root cross-section images, which
347 incorporate background knowledge about the root anatomy of a particular species and the types of images
348 used, our proposed deep learning approach is easily generalizable to various species and types of images,
349 given that a representative set of annotated images is provided as training data. Under the assumption that
350 data annotation is expensive and laborious, we aim to understand how many images are necessary for
351 good performance on roots from a particular species. Intuitively, the number of required images should be
352 relatively small, given that our model relies on a VGG-16 network pre-trained to detect a large number of
353 objects that are generally more complex than root, stele and late metaxylem.

354 To validate our intuition, we have performed an experiment where we varied the number of images used
355 for training, while keeping the number of test images fixed. Specifically, we used 5, 10, 25, 50, 75, 100,
356 150, 200, 250, and all available training images in that split, respectively, to train models for detecting the
357 root, stele and late metaxylem in an images. The 50 \times magnification images were used to train the models
358 for root/stele/LMX, while and 100 \times magnification images were used to train models for LMX. The trained
359 50 \times models were used to detect the root, stele, and LMX objects in the test images. Similarly, the trained
360 100 \times models were used to detect LMX objects in test images, with the goal of understanding the benefits
361 provided by higher resolution images.

362 The performance of the models was measured by comparing the predicted objects with the ground truth
363 objects. We used the IoU metric to evaluate the predicted bounded boxes for root/stele by comparison
364 with the corresponding ground truth bounding boxes. We used the mAP metric to measure the ability of
365 the models to accurately detect LMX objects. The variation of performance with the number of training
366 images is shown in Figure 4 for root/stele (Left plot) and LMX (Right plot). As can be seen, in the case
367 of the 50 \times models, the performance increases with the number of training images, but tends to stabilize
368 generally around 250 images. This confirms our intuition that only a small number of labeled images is
369 needed to learn accurate models for the problem at hand. Furthermore, the left plot in the figure shows
370 that the IoU values for both root and stele objects are around 0.95, when all the training images are used,
371 although the root bounding boxes are slightly better than the stele bounding boxes. Similarly, the LMX
372 objects are detected with high accuracy, as shown on the right plot of Figure 4, where the mAP values are
373 close to 0.9 consistently for models trained with smaller or larger number of 100 \times magnification images.
374 Similar performance is obtained with the models trained from all 50 \times magnification images. The plots for
375 both root/stele and LMX also show that generally the variance decreases with the size of the data. The slow
376 decrease in performance that is observed sometimes between two training set sizes can be explained by the
377 addition of some inconsistently labeled images present in the original dataset, as shown in Figure 5.

378 **3.2 Performance Evaluation Using RMSE**

379 The trained root/stele and LMX detection models were used to detect root/stele/LMX objects in the test
380 data. Subsequently, the detected objects were further used to calculate RD, SD, LMXD and LMXN. To
381 evaluate the models in terms of their ability to produce the right root/stele/LMX diameter and LMX number,
382 we have used the RMSE error computed by comparing the measurement/count estimates obtained from the
383 predicted bounded boxes with the ground truth measurements/counts. The RD and SD measurements were
384 evaluated based on models trained/tested with the 50 \times magnification images, while LMXD and LMXN
385 were evaluated based on models trained/tested with 50 \times and 100 \times magnification images, respectively.
386 Intuitively, the LMXD/LMXN results obtained with the models trained on the 100 \times magnification

Table 3. RMSE Results. The RMSE results for root diameter (RD), stele diameter (SD), late metaxylem diameter (LMXD) and late metaxylem number (LMXN) for 5 splits, together with the average and standard deviation over the 5 splits. The number of 50× magnification images used in these experiments is 388, while the number of 100× magnification images is 339. For each measurement, the magnification of the images that were used to train the model that produced that measurement (i.e., 50× or 100×) is also shown. Furthermore, for each split, the test fold corresponding to that split is shown.

Split (Test Fold)	RD	SD	LMXD	LMXD	LMXN	LMXN
	(50×)	(50×)	(50×)	(100×)	(50×)	(100×)
Split 1 (Fold 5)	62.77	21.93	3.67	2.45	0.81	1.37
Split 2 (Fold 4)	32.18	17.54	3.77	3.13	0.71	0.45
Split 3 (Fold 3)	61.19	21.96	3.53	3.22	0.91	0.83
Split 4 (Fold 2)	33.12	20.01	3.58	3.56	1.90	0.63
Split 5 (Fold 1)	43.67	20.94	2.43	1.61	0.74	0.25
Average	46.59	20.39	3.40	2.79	1.02	0.71
Standard deviation	14.77	1.81	0.55	0.77	0.50	0.43

387 images should be more accurate, as those images have higher resolution. The results of the experiments
388 corresponding to the five splits, together with their average and standard deviation, are shown in Table 3.

389 As can be seen from Table 3, the average RMSE error for RD over the 5 splits is 46.59 μm . Given that
390 root diameter for the images in our dataset varies between 354 μm and 1352 μm (see Table 1), this result
391 is very encouraging. Similarly, the average RMSE error for SD over the five splits is 20.39 μm , which is
392 low, given that the stele diameter varies between 115 μm and 419 μm . As opposed to root and stele, the
393 LMXD is significantly smaller, varying between 15 μm and 65 μm . However, the average RMSE error
394 is 3.40 μm for the model trained using the 50× magnification images, and decreases by almost 1 μm for
395 the model trained with the 100× magnification images (the exact value is 2.79 μm). In terms of LMXN,
396 the ground truth numbers vary between 1 and 12, with an average of 5.4 LMX objects per image. The
397 average RMSE error for LMXN is 1.02 for the models trained on the 50× magnification images, and down
398 to 0.71 for the models trained on the 100× magnification images. Thus, we can say that our models miss
399 roughly one LMX per image, when trained with the 50× magnification images, and less than that, when
400 trained with the 100× magnification images. We performed error analysis to understand if these results
401 might be useful in practice. Specifically, we analyzed images where our models made mistakes in terms
402 of LMXN, and observed that some of those images were annotated in an inconsistent way by the human
403 annotators, as can be seen in Figure 5. This observation is not surprising, as human annotators are prone
404 to mistakes and inconsistencies. As opposed to that, the automated models produced by our proposed
405 approach produce more consistent results, once they are well trained. At last, Table 3 shows that the RMSE
406 results obtained do not vary significantly with the split, as shown by the relatively small standard deviation.
407 Together, these results suggest that the proposed approach has the potential to replace the labor-intensive
408 manual annotations of root cross-section images.

409 3.3 Comparison with RootAnalyzer and RootScan

410 We aimed to compare DL-RootAnatomy with RootAnalyzer and RootScan tools on all 388 50×
411 magnification images in our dataset.

412 Given the batch processing capability of RootAnalyzer by comparison with the amount of user effort
413 involved by the RootScan, we started our comparative analysis with RootAnalyzer. As described in Section
414 2.2, RootAnalyzer has five parameters that need to be tuned: W , T , A_l , A_s and N . To understand the range
415 of parameters that our images require for good performance, we experimented with a variety of images

Table 4. Number of instances that could be analyzed with RootAnalyzer in each of the 5 folds, out of the total number of instances in each fold. These images were used in the comparison between our approach and existing tools. The total number of instances used in the comparison is also shown in the last column.

Fold	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Total
Instances	57	70	73	58	65	323
Total-fold	71	79	86	77	75	388

Table 5. Comparison between the RMSE results of our proposed approach (called DL-RootAnatomy), and the RMSE results of RootAnalyzer, RootScan (automated) and RootScan (adjusted). The number of images used in these experiments is 323. RootScan (adjusted) is seen as an estimate of the human error. DL-RootAnatomy was run on both 50 \times and 100 \times magnification images to detect LMX objects. The other tools were used with the 50 \times magnification images (as they do not work properly with 50 \times magnification images). The RMSE is calculated for: RD, SD, LMXD, and LMXN. The results are averaged over five splits. Corresponding to each average, standard deviation is also show.

Method	RD	SD	LMXD	LMXN
RootAnalyzer	208.44 \pm 22.40	172.21 \pm 20.65	32.89 \pm 10.62	4.01 \pm 0.54
RootScan (automated)	132.33 \pm 40.08	428.89 \pm 13.29	45.20 \pm 2.88	19.58 \pm 1.66
DL-RootAnatomy (50 \times)	43.67 \pm 16.80	20.51 \pm 1.84	3.58 \pm 0.57	1.13 \pm 0.43
DL-RootAnatomy (100 \times)	N/A	N/A	2.79 \pm 0.93	0.64 \pm 0.31
RootScan (adjusted) \approx Human Error	66.82 \pm 20.86	42.27 \pm 25.54	6.26 \pm 2.39	0.72 \pm 0.23

416 and parameters, and observed that RootAnalyzer freezes for some images, regardless of the parameters
417 used. Specifically, it freezes or produces degenerate results on images for which the root has a dark, solid
418 boundary, and the epidermal cells are not clearly identifiable, while it works as expected on images for
419 which the root has identifiable epidermal cells. Examples of images that can or cannot be analyzed by
420 RootAnalyzer are shown in Figure 6 (a) and (b), respectively. Out of the 388 images, we identified 323
421 with clear epidermal cells and used those in the comparison between our tool and other related tools. The
422 distribution of the 323 images analyzed with RootAnalyzer over the 5 folds is shown in Table 4.

423 As mentioned before, RootScan is a semi-automated tool, which requires human interaction/approval
424 at each stage. For example, after automatically detecting the root border (and similarly the stele border
425 and late metaxylem border), the tool presents the user with the opportunity to manually redraw or adjust
426 the border, if the automatically detected border does not look as expected. We ran RootScan with human
427 interaction to estimate the human error/bias in our dataset, under the assumption that without human
428 error, the adjusted borders should lead to minimal differences between the original ground truth manual
429 measurements and the RootScan measurements. We refer to this experiment as RootScan (adjusted) in
430 what follows. We also run RootScan in an automated fashion, where we approved the borders identified by
431 the tool, without any further adjustment. We refer to this experiment as RootScan (automated).

432 First, we ran RootAnalyzer on the 323 images in batch mode, with a set of overall good parameters
433 manually identified in our preliminary examination of these images. With just one click, all the images were
434 processed and the results were output in a csv file. We then fed the 323 images into the semi-automated
435 tool, RootScan, and performed experiments in either the automated or adjusted modes. We used the
436 measurements produced by each tool to calculate the RMSE error for each fold and averaged the error
437 over the 5 folds. Finally, we tested our DL-RootAnatomy approach on the 323 images using the 5-fold
438 cross-validation setup. For LMXD and LMXN, we used both models trained on the 50 \times magnification
439 images and models trained on the 100 \times magnification images. The results of the experiments described
440 are shown in Table 5.

441 As can be seen from Table 5, our proposed approach, DL-RootAnatomy, is the most accurate when
442 compared to the existing tools, specifically RootAnalyzer and RootScan. First, the DL-RootAnatomy
443 models have an average RMSE error of $43.67 \mu m$ for RD (slightly different from the error reported in
444 Table 3, as a smaller number of images were used in this experiment, as described above). The RootScan
445 (adjusted) average RMSE, which is assumed to approximate the human error in the ground truth, is 66.82
446 μm , slightly higher than the error obtained by the DL-RootAnatomy. Second, our approach gives the
447 smallest error also for SD, specifically, $20.51 \mu m$, followed again by RootScan (adjusted) with an error of
448 42.27 . Third, our approach has the smallest error for LMXD ($3.58 \mu m$ for the $50\times$, and $2.79 \mu m$ for the
449 $100\times$ images), with the RootScan (adjusted) as the second winner (with RMSE $6.26 \mu m$). Finally, in terms
450 of LMXN, our models trained with $50\times$ images are slightly worse than RootScan (adjusted) (1.13 versus
451 0.72), while the models trained with $100\times$ images are slightly better than RootScan (adjusted) (0.64 versus
452 0.72). The other two automated tools, RootAnalyzer and RootScan (automated), give significantly higher
453 errors overall, with RootAnalyzer being better between the two, but not good enough to be used for the
454 automated analysis of rice images.

455 Thus, based on this comparison, we claim that the existing tools do not generalize well on the rice root
456 images studied in this article. We identified several possible reasons:

- 457 (1) For a given tool, it is hard to find parameters that are universally good for all images in our dataset.
458 For example, for a given set of parameters, the segmentation result from the RootAnalyzer in Figure
459 7 shows that the parameters are appropriate for the left image (a) where the LMX are reasonably
460 well identified, but not appropriate for the right image (b) where no LMX are identified. As opposed
461 to that, our experiments have shown that the performance of our models does not vary much with
462 hyper-parameters. Once a model is properly trained, it performs accurately on a big variety of images.
- 463 (2) Plant samples for imaging are grown in different conditions, for example in hydroponic (water based
464 nutrient supply) or in soil, and root cross-section images are collected using different techniques
465 (e.g., hand sectioning or sectioning using tools like vibratomes). Plant growing or image acquisition
466 differences lead to differences in image's color, contrast and brightness. Figure 8 shows input images
467 for RootAnalyzer, RootScan, PHIV-RootCell, and DL-RootAnatomy, respectively. As opposed to
468 other tools, our approach is not very sensitive to the light conditions or to the structure of the
469 root cross-section images (including the epidermis thickness, epidermis transparency, and distorted
470 cross-sections), assuming the models are trained with a variety of root cross-section images.
- 471 (3) Each tool is designed with certain image characteristics in mind, and may not work on images that do
472 not exhibit those characteristics. As described above, RootAnalyzer assumes a clear cell boundary and
473 does not work for images that contain a solid boundary where the cells are not clearly identifiable. Our
474 models simply reflect the broad characteristics of the images that they are trained on, instead of being
475 built with some characteristics in mind. No specific background knowledge is provided, except for
476 what is inferred automatically from training images.
- 477 (4) Each tool is designed for a particular species, and incorporates background knowledge for that particular
478 species. As different species may have different root anatomy, a tool designed for a species may not
479 work for other species. For example, RootAnalyzer is designed to automatically analyze maize and
480 wheat root cross-section images, and "may work" for other species (Chopin et al., 2015). However,
481 our models can be easily adapted to other species, assuming some annotated training images from
482 those species are provided. No other background knowledge is required. Along the same lines, our
483 models can be easily adapted to images with a different resolutions, assuming those images include
484 the features of interest.

485 3.4 Tool Availability and Time Requirements

486 DL-RootAnatomy can be run from a terminal or as a web-based application, which is also mobile friendly.
487 The web-based application is available at <https://rootanatomy.org>, and links to a GitHub repository that
488 contains the source code and our pre-trained models and the ground truth data. The web-based application
489 is user-friendly and does not require any programming skills. It can be run with one of our sample images
490 displayed on the site, or with an image uploaded by the user.

491 In terms of time/image requirements, our experiments have shown that accurate models can be trained
492 from scratch with 150 to 250 images. The average time for labeling an image with LabelImg is
493 approximately 2 minutes. The average time for training a model on an EC2 p2-xlarge instance available
494 from Amazon Web Services (AWS) is approximately 10 hours, and does not require any human intervention
495 during that time. Once the model is trained, the average time to annotate a new image is less than one
496 second (using an EC2 p2-xlarge instance). If using our webserver (hosted on a local machine), the running
497 time for a annotating a new image is approximately 9 seconds, as this includes the time to setup the virtual
498 environment, the time to retrieve the input image from the server, the time to perform the annotation, and
499 the time to download the image to the user's browser. Given these time requirements, assuming that a
500 relatively large number of images need to be annotated for a biological study (on the order of thousands),
501 the human time can be potentially reduced from days or weeks (the time would take to manually annotate
502 all images) to hours (the time may take to manually label images for training) or minutes (the time for
503 automatically annotating images with our tool).

504 To gain insights into the time to "adapt" our models to other types of root cross-section images, we
505 identified 14 images that have been used to demonstrate RootAnalyzer and 10 images that have been
506 used to demonstrate PHIV-RootCell. Out-of-the-box, our trained models were not very accurate on these
507 images. However, we fine-tuned our models using 10 images from RootAnalyzers and 6 images from
508 PHIV-RootCell, and tested the new models on the remaining 4 images from RootAnalyzer and 4 images
509 from PHIV-RootCell. The results were impressive, showing that the models fine-tuned with such a small
510 number of images from RootAnalyzer and PHIV-RootCell learn to predict those types of images accurately,
511 in addition to our images, as can be seen at <https://rootanatomy.org>. Thus, the human time for labeling
512 images for training can be dramatically reduced to less than an hour, if one is fine-tuning our models as
513 opposed to training a model from scratch.

4 CONCLUSIONS

514 In this paper, we presented a fully automated approach for processing root cross-section images to extract
515 anatomical root features. The approach is based on Faster R-CNN networks, which were used to detect
516 root, stele and LMX objects, and to predict bounding boxes for each detected object. Subsequently, the
517 bounding boxes were used to obtain anatomical properties, specifically, root diameter, stele diameter, LMX
518 diameter and LMX number. Our Faster R-CNN models had VGG-16 as a backbone, to take advantage of
519 the extensive training of the VGG-16 network, and were fine-tuned on root cross-section images.

520 We evaluated the models in terms of their ability to detect the objects of interest, and also in terms of their
521 ability to lead to accurate measurements for RD, SD, LMXD and LMXN. The results of the evaluation
522 showed that our models produced accurate and consistent annotations, when trained on a relatively small
523 number of training images. For LMXD and LMXN, we trained models from both 50× magnification
524 images and 100× magnification images. Our results showed that the performance is slightly better for
525 the 100× magnification images, although this magnification is not a requirement for good performance.

526 Furthermore, a comparison with existing tools for analyzing root cross-section images showed that our
527 automated tool identifies anatomical features more accurately than those obtained with tools that require
528 manual adjustment. Overall, these results suggest that our tool, DL-RootAnatomy, can potentially be used
529 in practice to accelerate the speed at which root cross-section images are analyzed, and save significant
530 human efforts and costs.

531 The evaluation in this paper was done on rice images. However, an important observation was that our tool
532 can be easily adapted to other types of root cross-section images and also to other species, by fine-tuning
533 our existing models with a small number of labeled images from the species of interest. Similarly, additional
534 anatomical features can be extracted by fine-tuning our existing models with images labeled according
535 to the traits that are targeted. As part of future work, we plan to thoroughly study domain adaptation
536 approaches that allow the transfer of knowledge from our existing rice models to models for other plant
537 species (or for other traits), without labeling a large number of images from the other species of interest.

CONFLICT OF INTEREST STATEMENT

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

AUTHOR CONTRIBUTIONS

540 XL carried out the model design and implementation, with input from DC. XL, CW and DC carried out
541 the computational experiment design, with input from SVKJ. RB and SVKJ performed the biological
542 experiment design and collection of the data. XL and CW carried out all the computational experiments.
543 RB performed the labeling of the data according to RD, SD, LMXD and LMXN measurements. XL and
544 CW performed the bounding box labeling. SVKJ is the agronomy project leader with technical background
545 in root phenotyping. DC is the computational project leader, with background in machine learning and
546 deep learning. XL and CW drafted the first version of the manuscript, and DC and SVKJ contributed to the
547 preparation of the final version of the manuscript. RJ contributed biological knowledge to the manuscript
548 and provided feedback on the preliminary version. CW designed and developed the webserver. All authors
549 read and approved the final manuscript.

FUNDING

550 Contribution No. 19-072-J from Kansas Agriculture Experiment Station.

DATA AVAILABILITY STATEMENT

551 The image datasets used in this study can be found in a GitHub repository at [https://github.com/cwang16/Root-](https://github.com/cwang16/Root-Anatomy-Using-Faster-RCNN)
552 [Anatomy-Using-Faster-RCNN](https://github.com/cwang16/Root-Anatomy-Using-Faster-RCNN).

REFERENCES

- 553 Aich, S. and Stavness, I. (2017). Leaf counting with deep convolutional and deconvolutional networks.
554 *arXiv preprint arXiv:1708.07570*
- 555 Araus, J. L., Slafer, G. A., Royo, C., and Serret, M. D. (2008). Breeding for yield potential and stress
556 adaptation in cereals. *Critical Reviews in Plant Sciences* 27, 377–412. doi:10.1080/07352680802467736

- 557 Betegón-Putze, I., González, A., Sevillano, X., Blasco-Escámez, D., and Caño-Delgado, A. I. (2018).
558 Myroot: A novel method and software for the semi-automatic measurement of plant root length. *bioRxiv*
559 , 309773
- 560 Bishopp, A. and Lynch, J. (2015). The hidden half of crop yields 1, 15117
- 561 Bucksch, A., BurrIDGE, J., York, L. M., Das, A., Nord, E., Weitz, J. S., et al. (2014). Image-based
562 high-throughput field phenotyping of crop roots. *Plant Physiology* doi:10.1104/pp.114.243519
- 563 Burton, A. L., Williams, M., Lynch, J. P., and Brown, K. M. (2012). Rootscan: software for high-throughput
564 analysis of root anatomical traits. *Plant and Soil* 357, 189–203
- 565 Chopin, J., Laga, H., Huang, C. Y., Heuer, S., and Miklavcic, S. J. (2015). Rootanalyzer: a cross-section
566 image analysis tool for automated characterization of root cells and tissues. *PloS one* 10, e0137655
- 567 Clark, R. T., MacCurdy, R. B., Jung, J. K., Shaff, J. E., McCouch, S. R., Aneshansley, D. J., et al. (2011).
568 Three-dimensional root phenotyping with a novel imaging and software platform. *Plant Physiology* 156,
569 455–465. doi:10.1104/pp.110.169102
- 570 Colombi, T., Kirchgessner, N., Le Marié, C. A., York, L. M., Lynch, J. P., and Hund, A. (2015). Next
571 generation shovelomics: set up a tent and rest. *Plant and Soil* 388, 1–20
- 572 Delory, B. M., Li, M., Topp, C. N., and Lobet, G. (2018). archidart v3. 0: A new data analysis pipeline
573 allowing the topological analysis of plant root systems. *F1000Research* 7
- 574 DePristo, M. A., Banks, E., Poplin, R., Garimella, K. V., Maguire, J. R., Hartl, C., et al. (2011). A
575 framework for variation discovery and genotyping using next-generation dna sequencing data. *Nature*
576 *genetics* 43, 491
- 577 Dobrescu, A., Giuffrida, M. V., and Tsiftaris, S. A. (2017). Leveraging multiple datasets for deep leaf
578 counting. In *ICCV Workshops 2017*. 2072–2079
- 579 Durham Brooks, T. L., Miller, N. D., and Spalding, E. P. (2010). Plasticity of arabidopsis root gravitropism
580 throughout a multidimensional condition space quantified by automated image analysis. *Plant Physiology*
581 152, 206–216. doi:10.1104/pp.109.145292
- 582 E. Schmidt, J. and C.M. Gaudin, A. (2017). Toward an integrated root ideotype for irrigated systems 22
- 583 Everingham, M., Eslami, S. M. A., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. (2015).
584 The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*
585 111, 98–136
- 586 Girshick, R. (2015). Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*.
587 1440–1448
- 588 Hall, H. C., Fakhrzadeh, A., Luengo Hendriks, C. L., and Fischer, U. (2016). Precision automation of
589 cell type classification and sub-cellular fluorescence quantification from laser scanning confocal images.
590 *Frontiers in plant science* 7, 119
- 591 Heslot, N., Yang, H.-P., Sorrells, M. E., and Jannink, J.-L. (2012). Genomic selection in plant breeding: a
592 comparison of models. *Crop Science* 52, 146–160
- 593 Jiansan, Z., Peter, S., Gernot, B., and Boris, R. (2017). Root traits of european vicia faba cultivars-using
594 machine learning to explore adaptations to agroclimatic conditions. *Plant, Cell & Environment* 0.
595 doi:10.1111/pce.13062
- 596 Jones, W., Alasoo, K., Fishman, D., and Parts, L. (2017). Computational biology: deep learning. *Emerging*
597 *Topics in Life Sciences* 1, 257–274
- 598 Jung, J. and McCouch, S. (2013). Getting to the roots of it: Genetic and hormonal control of root architecture
599 4, 186

- 600 Kadam, N., Tamilselvan, A., Lawas, L. M. F., Quinones, C., Bahuguna, R., Thomson, M. J., et al. (2017).
601 Genetic control of plasticity in root morphology and anatomy of rice in response to water-deficit. *Plant*
602 *physiology*, pp-00500
- 603 Kadam, N., Yin, X., Bindraban, P., Struik, P., and Jagadish, K. (2015). Does morphological and anatomical
604 plasticity during the vegetative stage make wheat more tolerant of water-deficit stress than rice? *Plant*
605 *physiology*, pp-114
- 606 Kamilaris, A. and Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and*
607 *Electronics in Agriculture* 147, 70–90
- 608 Khan, Z., Rahimi-Eichi, V., Haefele, S., Garnett, T., and Miklavcic, S. J. (2018). Estimation of vegetation
609 indices for high-throughput phenotyping of wheat using aerial imaging. *Plant methods* 14, 20
- 610 Khush, G. S. (2013). Strategies for increasing the yield potential of cereals: case of rice as an example.
611 *Plant Breeding* 132, 433–436. doi:10.1111/pbr.1991
- 612 Kuijken, R. C., van Eeuwijk, F. A., Marcelis, L. F., and Bouwmeester, H. J. (2015). Root phenotyping:
613 from component trait in the lab to breeding. *Journal of experimental botany* 66, 5389–5401
- 614 Lartaud, M., Perin, C., Courtois, B., Thomas, E., Henry, S., Bettembourg, M., et al. (2015). Phiv-rootcell:
615 a supervised image analysis tool for rice root anatomical parameter quantification. *Frontiers in plant*
616 *science* 5, 790
- 617 LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature* 521, 436
- 618 LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., et al. (1989).
619 Backpropagation applied to handwritten zip code recognition. *Neural computation* 1, 541–551
- 620 Mohanty, S. P., Hughes, D. P., and Salathé, M. (2016). Using deep learning for image-based plant disease
621 detection. *Frontiers in plant science* 7, 1419
- 622 Mooney, S. J., Pridmore, T. P., Helliwell, J., and Bennett, M. J. (2012). Developing x-ray computed
623 tomography to non-invasively image 3-d root systems architecture in soil. *Plant and Soil* 352, 1–22.
624 doi:10.1007/s11104-011-1039-9
- 625 Namin, S. T., Esmailzadeh, M., Najafi, M., Brown, T. B., and Borevitz, J. O. (2017). Deep phenotyping:
626 Deep learning for temporal phenotype/genotype classification. *bioRxiv*, 134205
- 627 Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems,*
628 *man, and cybernetics* 9, 62–66
- 629 Pfeifer, J., Kirchgessner, N., Colombi, T., and Walter, A. (2015). Rapid phenotyping of crop root systems
630 in undisturbed field soils using x-ray computed tomography 11:14
- 631 Pound, M. P., Atkinson, J. A., Townsend, A. J., Wilson, M. H., Griffiths, M., Jackson, A. S., et al.
632 (2017a). Deep machine learning provides state-of-the-art performance in image-based plant phenotyping.
633 *GigaScience* 6, 1–10
- 634 Pound, M. P., Fozard, S., Torres, M. T., Forde, B. G., and French, A. P. (2017b). Autoroot: open-source
635 software employing a novel image analysis approach to support fully-automated plant phenotyping.
636 *Plant methods* 13, 12
- 637 Reeb, C., Kaandorp, J., Jansson, F., Puillandre, N., Dubuisson, J.-Y., Cornette, R., et al. (2018).
638 Quantification of complex modular architecture in plants. *New Phytologist* 218, 859–872
- 639 Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster R-CNN: Towards real-time object detection with
640 region proposal networks. In *Advances in neural information processing systems*. 91–99
- 641 Rhee, S. Y. and Mutwil, M. (2014). Towards revealing the functions of all genes in plants. *Trends in plant*
642 *science* 19, 212–221
- 643 Sankar, M., Nieminen, K., Ragni, L., Xenarios, I., and Hardtke, C. S. (2014). Automated quantitative
644 histology reveals vascular morphodynamics during arabidopsis hypocotyl secondary growth. *Elife* 3

- 645 Schneider, C. A., Rasband, W. S., and Eliceiri, K. W. (2012). Nih image to imagej: 25 years of image
646 analysis. *Nature methods* 9, 671
- 647 Schulz, H., Postma, J. A., van Dusschoten, D., Scharr, H., and Behnke, S. (2013). Plant root system analysis
648 from mri images. In *Computer Vision, Imaging and Computer Graphics. Theory and Application*, eds.
649 G. Csurka, M. Kraus, R. S. Laramée, P. Richard, and J. Braz (Berlin, Heidelberg: Springer Berlin
650 Heidelberg), 411–425
- 651 Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image
652 recognition. *arXiv preprint arXiv:1409.1556*
- 653 Singh, A., Ganapathysubramanian, B., Singh, A. K., and Sarkar, S. (2016). Machine learning for
654 high-throughput stress phenotyping in plants. *Trends in Plant Science* 21, 110 – 124. doi:<https://doi.org/10.1016/j.tplants.2015.10.015>
- 655
- 656 Sozzani, R., Busch, W., Spalding, E. P., and Benfey, P. N. (2014). Advanced imaging techniques
657 for the study of plant growth and development. *Trends in Plant Science* 19, 304 – 310. doi:<https://doi.org/10.1016/j.tplants.2013.12.003>
- 658
- 659 Tardieu, F., Cabrera-Bosquet, L., Pridmore, T., and Bennett, M. (2017). Plant phenomics, from sensors to
660 knowledge. *Current Biology* 27, R770–R783
- 661 Topp, C. N., Iyer-Pascuzzi, A. S., Anderson, J. T., Lee, C.-R., Zurek, P. R., Symonova, O., et al.
662 (2013). 3d phenotyping and quantitative trait locus mapping identify core regions of the rice genome
663 controlling root architecture. *Proceedings of the National Academy of Sciences* 110, E1695–E1704.
664 doi:[10.1073/pnas.1304354110](https://doi.org/10.1073/pnas.1304354110)
- 665 [Dataset] Tzutalin (2015). Labelimg
- 666 Ubbens, J. R. and Stavness, I. (2017). Deep plant phenomics: a deep learning platform for complex plant
667 phenotyping tasks. *Frontiers in plant science* 8, 1190
- 668 Uga, Y., Sugimoto, K., Ogawa, S., Rane, J., Ishitani, M., Hara, N., et al. (2013). Control of root system
669 architecture by deeper rooting 1 increases rice yield under drought conditions 45
- 670 van Dusschoten, D., Metzner, R., Kochs, J., Postma, J. A., Pflugfelder, D., Buehler, J., et al. (2016).
671 Quantitative 3d analysis of plant roots growing in soil using magnetic resonance imaging. *Plant*
672 *Physiology* doi:[10.1104/pp.15.01388](https://doi.org/10.1104/pp.15.01388)
- 673 Yamauchi, T., Watanabe, K., Fukazawa, A., Mori, H., Abe, F., Kawaguchi, K., et al. (2013). Ethylene and
674 reactive oxygen species are involved in root aerenchyma formation and adaptation of wheat seedlings to
675 oxygen-deficient conditions. *Journal of experimental botany* 65, 261–273
- 676 Yip, K. Y., Cheng, C., and Gerstein, M. (2013). Machine learning and genome annotation: a match meant
677 to be? *Genome biology* 14, 205
- 678 Zhang, L., Lin, L., Liang, X., and He, K. (2016). Is faster r-cnn doing well for pedestrian detection? In
679 *European Conference on Computer Vision* (Springer), 443–457

FIGURE CAPTIONS

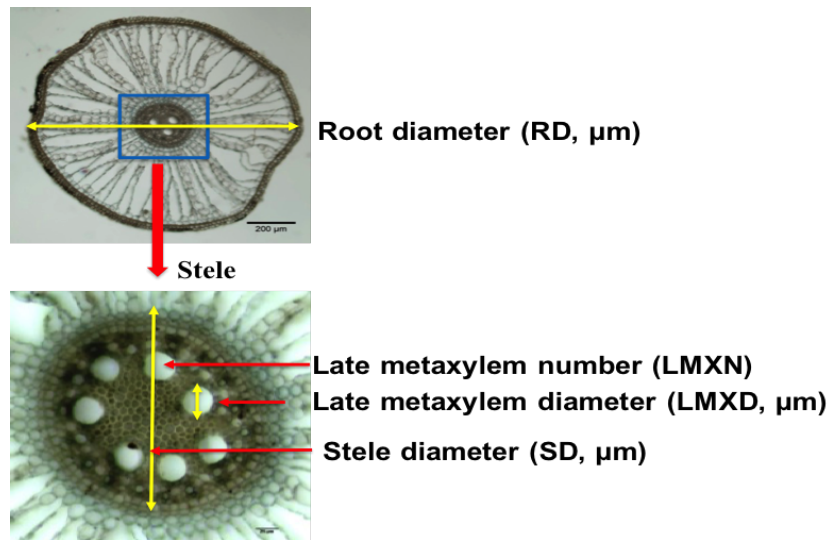


Figure 1. Root anatomical traits. (Top) Root cross-section with highlighted *root diameter* and *stele*. Image taken at 50x magnification. (Bottom) Enlarged stele with highlighted *stele diameter*, and *late metaxylem diameter*. The *late metaxylem number* is also a trait of interest. Image taken at 100x magnification.

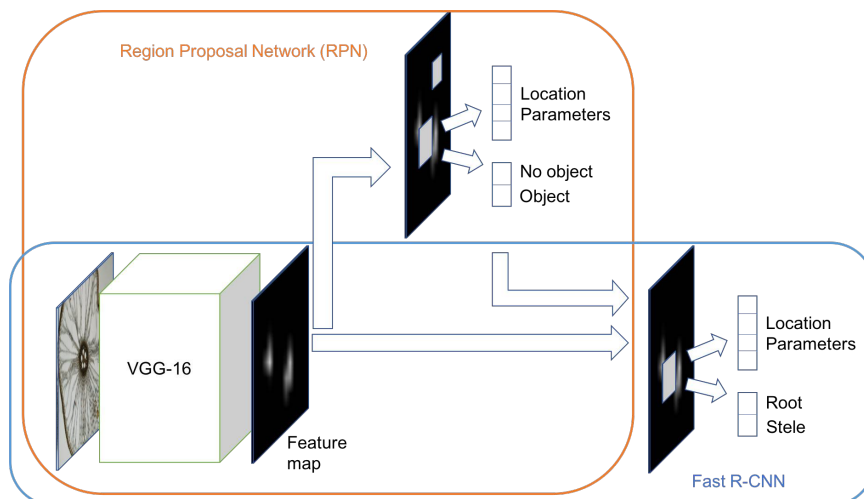


Figure 2. Proposed Faster R-CNN model architecture, which has two main components: 1) a region proposal network (RPN), which identifies regions that may contain objects of interest and their approximate location; and 2) a Fast R-CNN network, which classifies objects as root or stele, and refines their location, defined using bounding boxes. The two components share the convolutional layers of the pre-trained VGG-16 Simonyan and Zisserman (2014).

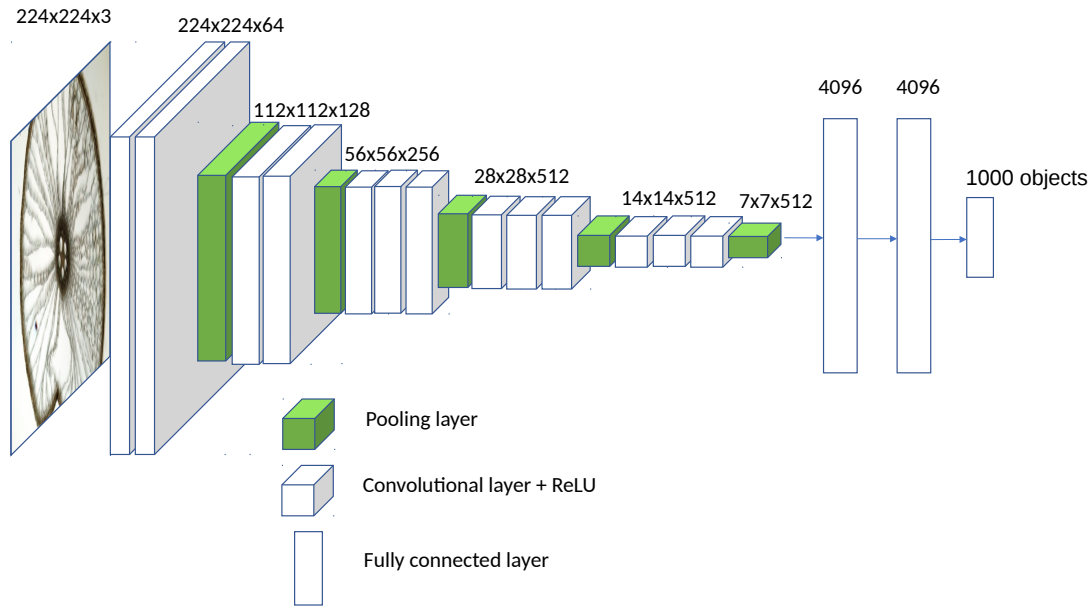


Figure 3. VGG-16. The original VGG-16 architecture consists of 13 convolution+ReLU layers, five pooling layers, and three fully connected layers. A convolution+ReLU layer produces a feature map, while a pooling layer reduces the dimensionality of the feature map. The last fully connected layer uses a softmax activation function to predict one of the 1000 categories. The dimensions corresponding to each layer are also shown.




Figure 4. Variation of performance with the number of training images for root/stele detection model (Left plot), and for the LMX detection model (Right plot), respectively. We used 50x magnification images to detect root and stele objects, and both 50x and 100x magnification images to detect LMX. The performance of the root/stele detection model was measured using the IoU metric (which shows how accurately the predicted bounding boxes match the ground truth), while the performance of the LMX detection model was measured using the mAP metric (which shows how accurately LMX objects were detected). The plots show average values over 5 splits together with standard deviation.




Figure 5. Examples of inconsistent annotations in our dataset. Specifically, image (a) was labeled as having LMXN=4 (the smaller LMX was included in the count), while image (b) was labeled as having LMXN=3 (the smaller LMX was not included in the count although it has size comparable with the smaller LMX counted in (a)). Our approach consistently identified 4 LMX objects in both (a) and (b) images. Similarly, image (c) was incorrectly labeled as having LMXN=4, while the similar image in (d) was properly labeled as having LMXN=3. Our approach correctly identified 3 LMX objects in both (c) and (d) images. Finally, images (e) and (f) show a larger number of LMX which have variable size, but it is not very clear which LMX were counted and which were not counted to get the 12 and 11 counts, respectively. Our approach identified 7 LMX objects in image (e) and 10 LMX objects in image (f).



Figure 6. Examples of root boundaries: (a) boundary with clear and identifiable epidermal cell; (b) dark solid boundary with unclear or unidentifiable epidermal cell.




Figure 7. RootAnalyzer Annotations: With the same set of parameters, in the left image the root border (red), stele border (yellow), endodermis (green) and late-metaxylem (purple) are detected reasonably well, while in the right image, only half of the stele border is detected. Given that the tool fails to properly detect the stele border, it also fails to detect the late metaxylem.




Figure 8. Sample input images used by different tools. (a) Sample input image for RootAnalyzer, which shows a clear difference between background pixels and cell border pixels. (b) Sample input image for RootScan. (c) Sample input image for PHIV-RootCell, which works with root cross-section images that contain a central metaxylem (marked with a white rectangle). (d) Sample root cross-section image from the dataset used in this study.