RESEARCH

RootNav 2.0: Deep learning for automatic navigation of complex plant root architectures

Robail Yasrab [ID][1], Jonathan A. Atkinson [ID][2], Darren M. Wells [ID][2], Andrew P. French [ID][1,2], Tony P. Pridmore [ID][1] and Michael P. Pound [ID][1,*]

[1]School of Computer Science, University of Nottingham, Jubilee Campus, Wollaton Road, Nottingham NG8 1BB, UK and [2]School of Biosciences, Sutton Bonington Campus, University of Nottingham, Nottingham LE12 5RD, UK

*Correspondence address. Michael P. Pound, School of Computer Science, University of Nottingham, Jubilee Campus, Wollaton Road, Nottingham NG8 1BB, UK. Email: michael.pound@nottingham.ac.uk [ID] http://orcid.org/0000-0002-5016-1078

Abstract

Background: In recent years quantitative analysis of root growth has become increasingly important as a way to explore the influence of abiotic stress such as high temperature and drought on a plant's ability to take up water and nutrients. Segmentation and feature extraction of plant roots from images presents a significant computer vision challenge. Root images contain complicated structures, variations in size, background, occlusion, clutter and variation in lighting conditions. We present a new image analysis approach that provides fully automatic extraction of complex root system architectures from a range of plant species in varied imaging set-ups. Driven by modern deep-learning approaches, RootNav 2.0 replaces previously manual and semi-automatic feature extraction with an extremely deep multi-task convolutional neural network architecture. The network also locates seeds, first order and second order root tips to drive a search algorithm seeking optimal paths throughout the image, extracting accurate architectures without user interaction. **Results:** We develop and train a novel deep network architecture to explicitly combine local pixel information with global scene information in order to accurately segment small root features across high-resolution images. The proposed method was evaluated on images of wheat (*Triticum aestivum* L.) from a seedling assay. Compared with semi-automatic analysis via the original RootNav tool, the proposed method demonstrated comparable accuracy, with a 10-fold increase in speed. The network was able to adapt to different plant species via transfer learning, offering similar accuracy when transferred to an *Arabidopsis thaliana* plate assay. A final instance of transfer learning, to images of *Brassica napus* from a hydroponic assay, still demonstrated good accuracy despite many fewer training images. **Conclusions:** We present RootNav 2.0, a new approach to root image analysis driven by a deep neural network. The tool can be adapted to new image domains with a reduced number of images, and offers substantial speed improvements over semi-automatic and manual approaches. The tool outputs root architectures in the widely accepted RSML standard, for which numerous analysis packages exist (http://rootsystemml.github.io/), as well as segmentation masks compatible with other automated measurement tools. The tool will provide researchers with the ability to analyse root systems at large scales than ever before, at a time when large scale genomic studies have made this more important than ever.

Keywords: convolutional neural network (CNN); plant phenotyping; computer vision; encoder-decoder; root system

Figure 1: An example of the challenge root phenotyping presents for computer vision. a) A sample input image of a *Brassica napus* seedling grown on germination paper. This plant phenotype exhibits a single primary root and numerous lateral roots. b) Cluttered scenes make segmentation challenging. c) Complex occlusion and intersection makes extracting root topology difficult. d) Many small image features, such as root tips, occur in close proximity, making identification difficult.

Background

Plant phenotyping plays a key role in plant science research, underpinning large-scale genetic discovery and the breeding of more resilient traits [1]. This innovation makes a fundamental contribution to the push for global food security. In recent years quantitative analysis of root growth has become increasingly important as a way to explore the influence of abiotic stresses such as high temperate and drought on a plant's ability to take up water and nutrients [2]. Segmentation and feature extraction of plant roots from images presents a significant computer vision challenge. Root images contain complicated structures, variations in size, background, occlusion, clutter, and variation in lighting conditions. Fig. 1 shows an exemplar root image captured on germination paper. Even a straightforward imaging assay presents numerous challenges to a classic computer vision pipeline.

In recent years machine learning has driven advances throughout many computer vision domains [3]. Indeed, much of the recent progress in plant phenotyping has also been driven by new and so-called deep learning techniques, a branch of artificial intelligence, often centring around convolutional neural networks (CNNs) [4–6]. The sharp increase in the availability of performant techniques in image analysis has coincided with an increase in the availability of genomic information in plant biology, providing an opportunity for robust and high-throughput solutions. The scale of the data challenge seen within plant science means that now, all but the truly fully automatic approaches will quickly become bottlenecks that hinder progress [7].

Analysis of root system architectures

In this article, we focus on the analysis of root systems where improvements promise increases to water and nutrient use efficiency [8]. Historically, automated root phenotyping has proven challenging, owing partly to the concealed nature of roots in the soil, but also to the architectural complexity and variability of root systems between species, and even individuals. Progress has been made through a combination of innovative approaches and tools [9,10], and new imaging technologies such as X-ray and magnetic resonance imaging [11,12].

The prevailing methodologies in root image analysis can be broadly categorized on the basis of the level of automation they provide. Fully automated tools attempt to quantify the traits of a root system without human guidance, often through a process of image segmentation followed by post-processing. These are what might be termed "bottom-up" approaches, which perform successive filtering over images in order to best distinguish between the foreground root material, and the background. Tools such as DIRT, GiA Roots, IJ_Rhizo, and EZ-Rhizo [13–16] offer a familiar pipeline in which an image is first segmented into 2 classes, root system and background, before noise removal (such as image filters and morphology [17]) and skeletonization techniques [18] are used to clean the image. These tools then quantify the distribution of root mass within an image, providing summary statistics such as root system width, height, and more complex measures such as density. Some tools, e.g., EZ-Rhizo, will measure root width at each location, providing more detailed analysis of the distribution of roots of different sizes.

A limitation of automated systems such as these is that errors propagate from early processing stages through to measurement. Noisy images or unexpected phenotypes will lead to errors in thresholding, which are challenging to remove and may lead to incorrect measurement of the root system. For this reason, most automated tools have placed heavy focus on cruder organ-scale measurements such as the total width of the root system because these are most robust to small errors in image segmentation. Owing to the challenge of reliably segmenting and analysing root systems automatically, many tools place strict requirements on the type of image they will analyse. RhizoScan [19], for example, offers an automatic pipeline similar to the above, based on the OpenAlea platform [20], but supports only root systems grown on Petri plates.

Beyond the problem of low-level image analysis, by framing the problem as one of identifying root pixels at a low level, these tools struggle to extract high-level root architectural information. More detailed phenotypic traits such as the number of lateral roots are out of reach of many existing tools simply because disambiguating the category of a root within a system may prove impossible in the presence of noise, especially once growth is at a mature stage where roots begin to overlap. Semantically untangling such a root system requires a higher-level understanding of the image than pixel-based processing methods provide.

Manual root analysis tools such as ImageJ's polyline function [21] and DART [22] offer an entirely different approach. They place reliance on an expert human annotator to successfully identify the structure of the root system by asking the expert to label each root by hand. The advantage here is that if sufficiently well trained, an annotator could conceivably reconstruct an entire root system, using their advanced knowledge to clear up disambiguation in cluttered areas of the image. The obvious drawback to this approach is that this is an extremely time-consuming process. In practice, many experiments will therefore have to severely limit the number of measurements captured per image, such as by focusing on primary root length, to bring the time required into a reasonable range. Some tools, e.g., RootScape [23], have been designed with this in mind, requiring that a user highlight only 20 key landmarks on a root system. These landmarks are then used to explore phenotypic differences between genotypes via principal component analysis. In those instances where detailed analysis is required, the burden

on annotators is huge, and the cost of mistakes may be high. Outside of plant science, obtaining cheap and efficient annotation has become a widely researched topic in and of itself [24,25]. In plant science, noisy and low-cost annotation may not be acceptable, depending on the experimental requirements, and ultimately offers few benefits over the automated tools described above.

Alongside the development of manual and automated tools, a selection of widely used semi-automatic tools have been released. These approaches aim to bridge the gap between speed and accuracy, offering a compromise acceptable for many use cases. Tools such as RootReader [26] perform a similar automatic function to the tools above but provide the user with the ability to manipulate some of the output to correct mistakes. Most of the tools in this category are not bottom up, and instead model the root system in some way, guided by the user, in order to better understand the image on which they are run. Smartroot [10], a plugin for the popular ImageJ tool [21], operates by tracing along each root in a guided way, at each step searching for the optimal direction in which to travel based on the current orientation of the root at that point. Smartroot is semi-automatic, with initiation of roots and correction of errors often requiring human intervention. Nevertheless, with some user effort Smartroot can potentially be used to reconstruct full root system architectures. RootNav [9], a precursor to the work presented here, offers a point-to-point path search between labelled seed locations and root tips. Images are first segmented into background/foreground classes, before a user is required to label root tip and seed locations. Shortest path search is used to trace between key organ landmarks, resulting in a complete reconstruction of the root system. However, RootNav does not include a reliable method for *detecting* seeds and root tips (the user must perform this step), nor is the segmentation step robust to image noise. This means that significant user interaction is still required to guide the software, but as with Smartroot, the output is a full and architecturally correct root system architecture. Many tools that are able to output root system architectures have been adapted to provide output in the popular RSML format [10]. RSML is an XML-based standard for the sharing of root system architectures, including information on geometry, and relative position within the system. Numerous tools exist to read and write RSML files, allowing customized pipelines between tools, and the ability to decouple the image analysis from the ultimate measurement of traits, as well as view the final architecture labelling.

Deep learning for root systems

The prevailing methodology when working with images in deep learning is the CNN. CNNs improve upon traditional machine learning via their ability to learn not only solutions to problems but also the most effective way in which to transform data to make this goal easier. This representation learning provides CNNs with unparalleled discriminative power and has seen them quickly move into a dominant position within the field of computer vision [3]. A CNN is a layered structure that performs successive image-filtering operations that transform an image from a traditional RGB input into a new feature representation. This transformation is learned during training and provides the final layers of the CNN with the best possible view of those data from which to base decisions. The deeper into a CNN data flows, the more abstracted and powerful the representation becomes. While the initial layers may compute simple primitives such as edges and corners, deeper into the network fea-

ture maps may highlight groups of primitives. Deeper still, feature maps may contain complex arrangements of features representing real-world objects [5]. These features are learnt by the CNN training algorithms and are not hand-coded, meaning that with sufficient training data any number of different problems can be addressed. Within the biosciences, such networks have been used to perform a variety of tasks ranging from classification, assigning discrete labels to images and objects [27], through to regression problems; i.e., of directly predicting values [28]. For root systems, Pound et al. [29] used a deep classification network to scan an image for probable root tip locations in 32×32 pixel tiles. Despite promising results, the drawback of this approach is that using a small field of view, customarily called a "receptive field" within the machine learning literature, is computationally less efficient and may produce additional false-positive results where the small field of view is not sufficient to distinguish true roots from image noise. This system also only currently detects root tips, which means more complex traits involving other organs cannot be computed.

Image segmentation and feature localization

The measurement of complex phenotypic traits requires analysis at a finer scale than that of whole-root-system traits but sensitive to more than only a small selection of plant features such as just root tips. To address this, the research community has begun to move towards networks that output a richer array of information. Recent work has been based around newer CNN designs in what we term an encoder-decoder configuration, aimed at segmentation of images, or the location of key feature points. Traditional CNNs perform spatial downsampling such that by the end of the network, features spatially correspond to the entire image, i.e., they have lost location resolution. This is ideal for classification tasks, where a decision must be made on an image scale. This is not appropriate, however, for situations in which a 2D segmentation result is required. Encoder-decoders therefore upsample again from the feature space, back into a spatially high-resolution image (Fig. 2). This process can be thought of as combining a CNN with a second, reversed CNN that learns to produce images once again; these images might be trained to predict the locations of objects, or to segment pixels into background and foreground classes. Encoder-decoders are being used in plant science to, among other tasks, segment plant shoots [30,31], other plant organs [32], and fill gaps in rhizotron images of root systems [33]. Pound et al. [5] first introduced the concept of heat map regression to the plant phenotyping domain, in which a segmentation output is replaced by a heat map showing likely target locations. Our development in this article combines both of these approaches, simultaneously segmenting a root system and predicting the likely locations of root tips and seeds.

Automated root phenotyping

We present here a new tool for the automatic analysis of root systems that is designed to work across a wide variety of plants and imaging conditions. Our pipeline is driven by a deep encoder-decoder network, similar to that presented by Pound et al. [5] but adapted to handle higher-resolution images. The network is trained to simultaneously segment root material, classify root type, and locate key features from which root geometry can be derived. To our knowledge this is the first use of deep learning to perform multi-task segmentation and localization in plant phenotyping. The output of the network is refined using an

Figure 2: A simplified example illustrating the major components of a CNN in an encoder-decoder configuration. The encoder performs a combination of filtering operations including convolutional filters, spatial downsampling, and normalization. These layers convert the original image into a high-dimensional feature space but with very low spatial resolution. The decoding network performs similar layer operations but replaces downsampling with upsampling to return the feature representation back into a spatially high-resolution image.

A∗ shortest path algorithm to determine the most likely path of each root, connecting located second-order roots to appropriate first-order roots, and first-order roots back to the seed location. Full root geometry is extracted per plant and is robust to multiple plants and highly varied architectures. The tool outputs the standard RSML format [34], widely supported by the community, from which root system architecture (RSA) traits can be derived. The tool also outputs the underlying segmentation masks for first- and second-order roots, from which global traits can be derived. An overview of the tool can be seen in Fig. 3. The system first performs pixel-wise segmentation of the image and heat map regression to locate key features; it next extracts the root topology via a series of guided shortest-path searches before finally extracting the entire root architecture into a portable RSML format.

We first demonstrate the performance of the tool on a large wheat dataset grown on germination paper. We perform a quantitative comparison with traits measured using the original semi-automatic RootNav tool [9], hereby referred to as RootNav 1.0, in which an expert performed detailed manual intervention to ensure accuracy. We next demonstrate the ability of RootNav 2.0 to adapt to new image types with a much smaller training set. We retrained the network on 200 images of *Arabidopsis thaliana* grown on agar plates, in which up to 5 plants appear per image. We again compare quantitatively against human-labelled images generated using RootNav 1.0. Finally, we transfer learn once more using an even smaller, rapeseed dataset, comprising only 91 training images. Beyond accuracy measures, we have assessed our system's performance in terms of inference time and resource efficiency to provide a comparative analysis of user burden for root architecture analysis. The trained networks, tool, and all training datasets have been made publicly available.

Data description
Primary dataset

Our primary dataset is composed of images of wheat (*Triticum aestivum* L.) seedlings totalling 3,630 images of $1,900 \times 2,000$ pixel resolution. Images include those released by Pound et al. [5], plus additional images captured using the same methodology. Images were captured as per Atkinson et al. [35]; seeds were sieved to uniform size, sterilized, and pre-germinated before transfer to growth pouches in a controlled environment chamber (12-hour photoperiod: 20°C day, 15°C night, with a light intensity of 400 μmol m^{-2} s^{-1} photosynthetically active radiation). After 9 days (with plants at the 2-leaf stage), individual pouches were transferred to a copy stand for imaging using a Nikon D5100 DSLR camera controlled using NKRemote software (Breeze Systems Ltd, Camberley, UK). Ground truth annotations

for all plants were obtained using the original RootNav 1.0 software [9] and stored in RSML format [34]. Each annotation was provided by an expert user, and because we intended to use RootNav 1.0 as a quantitative baseline for accuracy, emphasis was placed on accuracy over speed during this process.

Ground truth images for network training and validation were generated from these RSML files by rendering appropriate segmentation masks and heat maps. The dataset was split into training and validation sets totalling 2,864 and 716 images, respectively. An additional 50 images were held back as a final testing set. More details on this methodology can be found in the Methods section. Example images can be found in Fig. 4a.

Transfer learning datasets

Our second dataset is composed of images of *Arabidopsis thaliana* grown on agar plates as detailed by Wilson et al. [36]. Images of individual plates were acquired using near-infrared imaging utilizing the system described by Wells et al. [37]. In this system, multiple seeds are sown on each plate, and thus, unlike the primary dataset, each image typically contained up to 5 plants (Fig. 4b). This dataset is considerably smaller, totalling 277 images, and is used as a demonstration of transfer learning with our approach despite limited annotated data. The dataset was split into training and validation sets of 200 and 27 images, respectively, and as with the primary dataset, 50 holdout test images were used for final quantitative evaluation.

Our final dataset is composed of images of rapeseed (*Brassica napus*) seedlings, grown in the same system as used in the primary dataset above. This dataset is small, containing only 120 images of individual plants. Our hypothesis was that despite the reduced size, transfer learning from a network trained on both the wheat (similar image background) and *Arabidopsis* (similar root system organization) datasets would lead to sufficient accuracy. The dataset was split into training and validation sets of 91 and 14 images, respectively. We used 15 holdout test images for the final quantitative evaluation. Example images for the 2 transfer learning datasets can be found in Fig. 4b and c.

Analyses

This section will present a comprehensive performance analysis of RootNav 2.0, including a quantitative evaluation of both the underlying segmentation approach and the root architecture extraction. We evaluate segmentation accuracy via 3 common metrics, mean average pixel classification accuracy (both global and class averages) and mean intersection over union (mIoU). We compare the segmentation performance of our approach against the well-known benchmark architectures VGG [38], FCN [39], SegNet [40], UNet [41], and DeepLab [42]. We
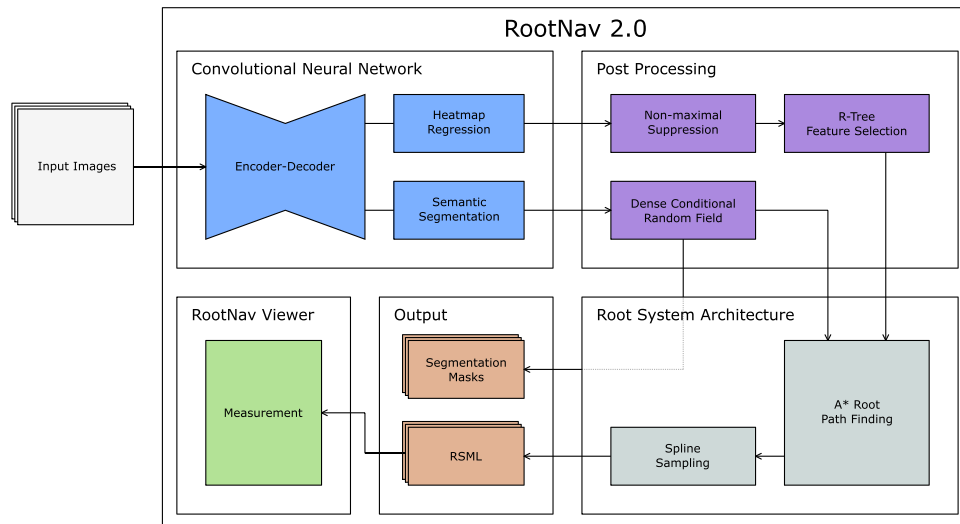
Figure 3: An overview of RootNav 2.0. The input enters a CNN that performs both segmentation of the root structure and localization of key points. These are post-processed to extract information for a path-finding algorithm. A∗ search then extracts likely paths taken by each root, generating an entire architecture for an arbitrary number of plants in an image. All roots are resampled as smooth splines, before all topology and geometry are output into an RSML file. Segmentation masks for first- and second-order roots are also saved.

Figure 4: Example images from each of the 3 datasets used during this work. (a) Wheat (*Triticum aestivum* L.). (b) *Arabidopsis* (*Arabidopsis thaliana*). (c) Rapeseed (*Brassica napus*). Scale bars are 50 mm long.

then evaluate the automatic reconstruction of root systems using a comparison of common root phenotypic traits such as the dimensions of the root system, and root counts. For ground truth, we use semi-automatic measurements obtained through expert annotation using RootNav 1.0. Finally, we perform the same experiments to outline the accuracy on the 2 additional datasets, which contain fewer training images, to demonstrate the efficacy of transfer learning to new species and imaging modalities.

Table 1: Quantitative comparison: a quantitative analysis of trainable parameters and memory requirements of different benchmark architectures used during experiments

CNN model	Trainable parameters, input (3 × 256 × 256 pixels)	GPU memory requirements (bytes)
VGG-16 [38]	138,357,544	1,253,048,320
FCN [39]	134,815,994	1,766,850,560
SegNet [40]	29,572,256	1,603,272,704
UNet [41]	13,395,329	1,276,116,992
Stacked Hourglass [43]	6,720,132	6,309,281,792
LinkNet [44]	11,546,148	**533,725,184**
PSPNet [45]	65,589,332	1,934,622,720
DeepLab-V3 [42]	59,344,309	596,639,744
RootNav 2.0	**1,595,782**	892,338,176

The input size was set at a constant 3 × 256 × 256 pixel size for this comparison.

Root image segmentation

RootNav 2.0 is driven by a deep network that segments images of root systems into classes: background, first-order roots, and second-order roots. Crucial to the accuracy of any subsequent path-finding approach is a reliable segmentation. Segmenting whole-root images is important in order to provide sufficient context when distinguishing first- or second-order roots. Splitting images into efficient tiles reduces memory consumption but makes distinguishing root type problematic. With this in mind, we designed the network to be efficient by reducing the number of trainable parameters, intermediate feature sizes, and thus overall memory requirements. This allows larger 1,024 × 1,024 resolution input. Table 1 shows a comparison of the memory requirements and parameter sizes of commonly used segmentation networks, and our own architecture.

We trained each network on the wheat dataset as described in the Methods. To provide a fair comparison of each network, we allocated 2 Nvidia GPUs with >11 GB onboard memory each for training each network, then trained using consistent hyperparameters such as learning rates, and equal batch sizes. Image

Table 2: A quantitative comparison of the segmentation performance on root images of RootNav 2.0 against other commonly used CNN architectures

CNN architecture	Global average	Class average	mIoU
VGG [38]	37.32	38.5	31.11
FCN [39]	43.78	48.56	39.47
SegNet [40]	68.68	70.23	47.05
UNet [41]	70.47	69.65	51.88
DeepLab-V3 [42]	87.7	85.90	54.0
RootNav.2.0	**99.6**	**95.1**	**66.1**

Performance is measured using global average accuracy, class average accuracy, and mean intersection over union. The classes evaluated are background (no root), first-order, and second-order roots.

resolution was maximized for each network depending on its resource requirements. Accuracy was measured using 3 standard metrics: Global average accuracy, class average accuracy, and mIoU. Global average accuracy measures the performance of segmentation over all pixels in the validation set. High values indicate that the majority of pixels have been classified correctly. Because most pixels are background in root images, high values indicate few false-positive results but do not necessarily demonstrate good root segmentation. Class average accuracy measures the performance of each class separately, before computing a final average. High values here represent good performance across all classes. Finally, mIoU represents the percentage of overlap between each class and the ground truth. Higher values indicate predictions closer to that of the ground truth labelling.

Example image output from each network can be found in Fig. 5, with quantitative results for all tested networks across the validation set shown in Table 2. The larger networks contain more features, which while in some cases may improve performance of a deep network, here hinders the ability of each network to resolve finer detail because they cannot operate at 1-megapixel image resolution. The strong performance of RootNav 2.0 in this experiment can be attributed to its efficient use of features throughout the network, lower memory requirements, and thus larger 1-megapixel input sizes.

Extraction of root system architecture

After segmentation and feature localization, segmentation masks are converted into a weighted graph structure amenable to traversal with a shortest-path algorithm. RootNav 2.0 extracts a full root architecture by performing a series of heuristic searches across the image. First, shortest paths are found between all first-order root tips and the most appropriate seed location (defined as the seed first reached during a heuristic search). This generates a series of first-order roots, to which second-order root paths are found from all second-order root tips. The output of this process is a complete root architecture description, stored in RSML format, from which phenotypic traits can be derived. We compare the output of RootNav 2.0 against ground truth measurements captured using RootNav 1.0 in collaboration with an expert user. Quantitative traits were measured directly using the RSML output by both tools; results are presented in Fig. 6.

We chose a range of root traits that are both representative of the measurements commonly used in the root phenotyping literature but also ones that exercise various aspects of our partic-

ular approach. For example, we include traits that measure the accuracy of feature detection (e.g., total root count) and those that also measure the accuracy of the shortest-path approach (e.g., total root length). In Fig. 6 it can be seen that there is strong agreement between the results of RootNav 2.0 and the ground truth measurements. Measurements based on the extremities of the root system (maximum depth, maximum width, and convex hull area) produced values very close to those in the ground truth, with r^2 values >0.99. Traits that summarize the entire root system, such as centroid depth, provided r^2 values in the range 0.64–0.72.

The prediction of first- and second-order root counts achieved r^2 values of 0.641 and 0.724, respectively. For first-order roots, we observed that the majority of incorrect predictions were either 1 count higher or 1 count lower than the ground truth and that these confusions often occurred near the seed position, where a seminal root may visually appear similar to a second-order root that emerges near the seed, or vice versa. Other failures were produced by roots leaving the field of view of the camera, but that had been annotated by the expert, or where 2 root tips grew in very close proximity (within a few pixels). Second-order roots were typically much shorter and often in close proximity. Some missed root tips would be caused by non-maximal suppression, when the R-Tree data structure is used to remove possible duplicates. We also found that the contrast on second-order roots was lower because they were usually thinner, which might account for some missed tips in this class. Errors in the detection of root tips will also propagate errors into the total root length measurements because these roots will not be detected. For second-order roots, we found that most of the error in root length can be attributed to missed roots rather than errors in path finding. For primary roots, path finding was usually robust, except in cases where 2 roots grow side by side. RootNav 1.0 handled these errors by allowing a user to intervene and correct any mistakes; in RootNav 2.0 we wish the process to remain fully automatic, so we do not explicitly correct for this. However, the occurrence of this type of growth is in the minority, in our experience. Centroid depth is measured as the mean position of all roots and so is influenced by the detection and path finding of every root.

An understanding of where and how RootNav 2.0 may produce errors provides insight into these results. An accurate measurement of maximum depth depends on only 2 variables: the location of the seed, and the location of the first-order root tip that is lowest (in terms of y-position) in the image. The graph of maximum depth in Fig. 6 reflects the fact that these 2 features were successfully found in every case. Similarly, maximum width depends only on the left- and right-most roots, and convex hull only on the outermost roots throughout the architecture. A missed second-order root within a root system will not affect these traits, so these results are robust even where some roots have been missed. This tells us that for the majority of images, the locations of the seeds, lowest tips, and outermost roots are detected successfully and that these traits that measure the extremities of the root system are robust.

Transfer learning to new species and images

To demonstrate the adaptability of our approach to different species and imaging modalities, we retrained the network first on an *Arabidopsis* dataset, comprising ∼277 images of *A. thaliana* grown on agar plates. We then trained once more from the wheat dataset to the rapeseed dataset, comprising 120 images of *B. napus* on germination paper. In both cases we extracted RSML

Figure 5: Example image output from each trained network architecture. (a) An example hydroponic wheat image. (b) VGG [38]. (c) FCN [39]. (d) SegNet [40]. (e) UNet [41]. (f) DeepLab-V3 [42]. (g) RootNav 2.0.
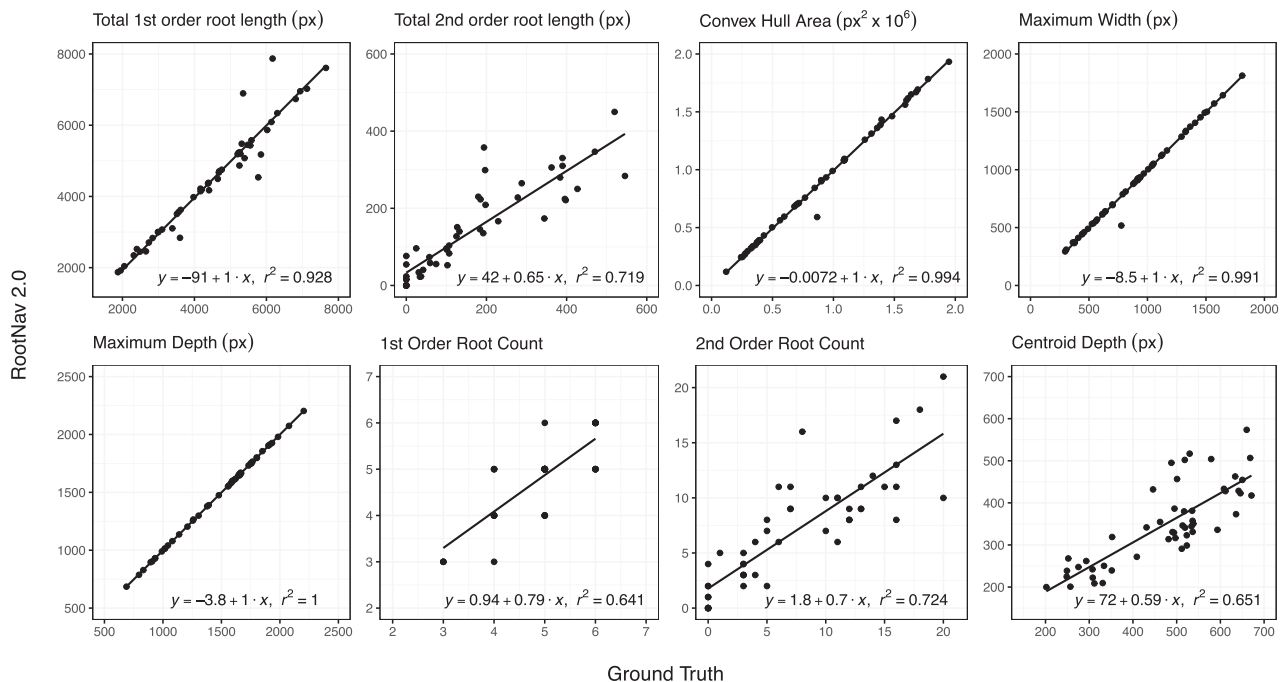
Figure 6: Numerical results showing a range of root system traits measured in RootNav 2.0 against ground truth measurements on the wheat test set. For each trait we also fit a linear regression model and report the r^2 value.

root descriptions and quantified these in the same way as the wheat dataset. We also trained both networks from randomly initialized weights, rather than transfer learning, and found that the datasets were too small to train effectively (Supplementary Figs 6 and 7).

Arabidopsis thaliana

The *Arabidopsis* dataset contains marked differences from the wheat data. The dataset contains many fewer images, which makes transfer learning essential to reduce overfitting. This species has a taproot structure that contains a single primary root, from which lateral roots emerge, rather than multiple first-order roots in the form of a primary and multiple seminal roots. This dataset is also imaged under infrared illumination and so contains no colour information, and a very different background arrangement consisting of a plastic plate containing semi-transparent growth medium instead of blue germination paper. Finally, each plate typically contains 5 plants rather than a single plant. We found that the network and heuristic searches adapted well to this new domain. We made minor modifications to the path-finding approach to support multiple plants, which are discussed further in the Methods. Quantita-

tive results are shown in Fig. 7, with full results found in Supplementary Fig. 1 and example image output in Supplementary Fig. 4.

Despite the smaller number of images available for training, the results show a good performance after transfer learning to the new data. Overfitting on the smaller dataset led to noisier segmentation and feature detection on some instances of the test data, examples of which can be found in Supplementary Fig. 8. Not every plant was successfully detected; a missing or additional primary root tip or seed location would mean that the number of plants was under- or overestimated. In 60% of images examined, the tool correctly identified the same number of plants as were marked in the ground truth. In 6% of the images, a single plant was missed, usually owing to the plant being extremely underdeveloped but having been annotated by the user anyway. We found only a single instance in 1 image that contained a well-established plant that had not been identified by our network. Overcounting of plants was more common, with 20% of images identifying an additional plant and 14% identifying more beyond this. In the majority of cases we found that these errors were caused by unusual angles in the leaves and germinated seeds at the top of the plant, producing false-positive detections. This is something that would likely be

Figure 7: Numerical results showing a sample of root system traits measured in RootNav 2.0 against ground truth measurements on the *Arabidopsis* test set. Each image contains up to 5 plants, and results are presented per plant. For each trait we also fit a linear regression model and report the r^2 value.

Figure 8: Numerical results showing a sample of root system traits measured in RootNav 2.0 against ground truth measurements on the *Brassica napus* test set. For each trait we also fit a linear regression model and report the r^2 value.

corrected with additional training data; remember, we are using a very small amount of training data for this image class, versus the wheat images. Where duplicate plants were found, they were often extremely close to, or even above, an existing plant location. These duplicates could be removed easily via

post-processing; this is something we do not address in this article.

Of the plants that were successfully identified, the traits captured by the tool offer a good agreement with the ground truth. As with the wheat dataset, measures of the extremities of the root system such as maximum depth performed with the highest accuracy, but we also found that total first-order root length was very close to the ground truth in the majority of cases. Errors here usually indicated a second primary root incorrectly detected alongside an existing one, a feature that we do not yet remove in post-processing as with duplicate plants, although this would be possible. The detection of second-order roots was also highly correlated with the ground truth measurements, and the total length of all second-order roots (measured per plant) correlated with the ground truth with an r^2 of 0.91.

Brassica napus
This dataset uses the imaging format of blue germination paper with single plants (like the wheat dataset) but contains the same species root structure as *Arabidopsis* (a single taproot from which all other roots derive). This dataset contains the fewest images, with only 90 images used for training. We use this small dataset as a demonstration of the efficacy of transfer learning, but we also note that training over a slightly larger dataset in practice would be worthwhile for improving robustness. Results can be found in Fig. 8, and in full in Supplementary Fig. 2. As with the *Arabidopsis* dataset, the low number of training images produces some overfitting that results in noisier output than the wheat dataset on some of the test data. Examples can be seen in Supplementary Fig. 9.

It can be seen that the correlation between RootNav 2.0 and ground truth ranges from r^2 of 0.539 (first-order root length) to 0.941 (convex hull area). Because this is such a small dataset, the test set contains only 15 images, meaning that there is inevitably more noise in the results than in the previous experiments. Nevertheless, the results are promising, particularly given the tiny size of this dataset compared with typical standards for deep learning. As with both previous datasets, convex hull and other extremity-based measures provided the most reliable results. Accuracy of the total length and root count metrics had a lower r^2 than the other datasets, caused we believe by the smaller training set, meaning that the approach is slightly less robust to noise. We found that in a few images the longest lateral tips were incorrectly classified as first order, causing erroneously high measures of first-order root length. We believe that this occurred where these laterals are mistaken for wheat seminal roots, on which the network was originally trained. This occurred on the minority of test images and is a problem that we are confident would be resolved with more training data. It would be possible to use prior domain knowledge, e.g., the knowledge that rapeseed has a single taproot, to clean the output during post-processing; as with the *Arabidopsis* dataset, we did not perform any post-processing of this kind in this work.

After segmentation, the structure of these root systems was quite amenable to traversal using a shortest-path approach. In many cases the longest roots grow close together, which causes errors where a search may travel along the same path as another root. We found that this did not substantially increase the error in total root length because many of these roots grew in close proximity and were of similar length. Nevertheless, dealing with root overlap in an efficient and automatic way is a topic worth exploring in future work.
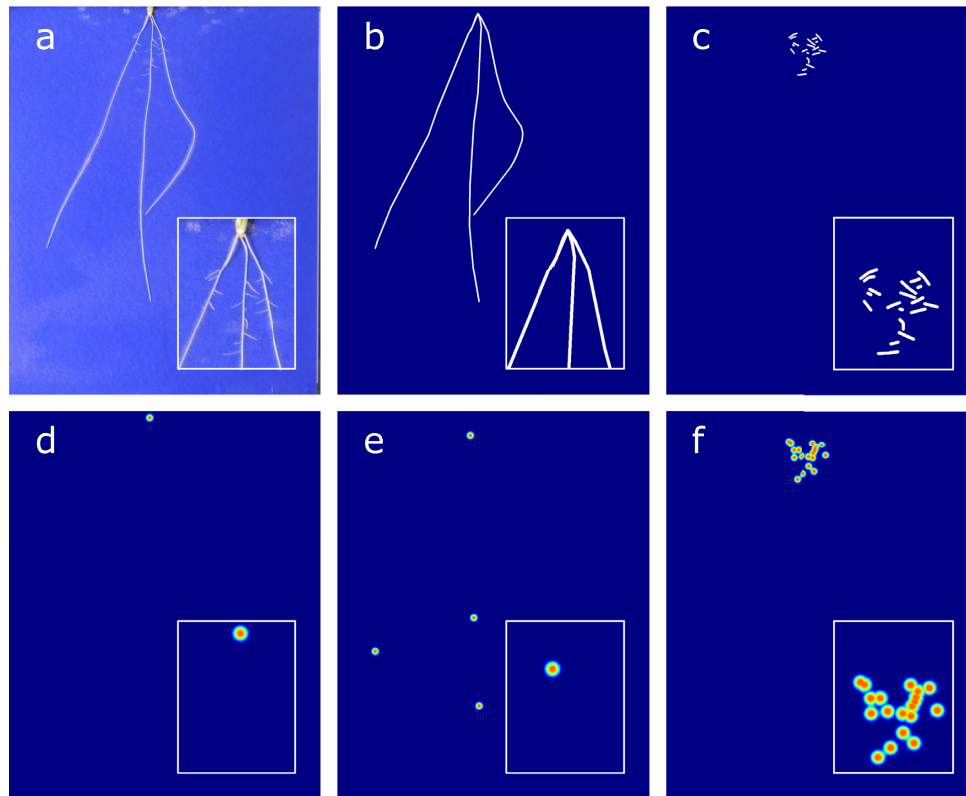
Figure 9: Sample training image and ground truth, zoomed sections and colour added for clarity. (a) An example wheat image. (b) Segmentation mask for first-order roots. (c) Segmentation mask for second-order roots. (d–f) Heat maps for seed, first-, and second-order tip locations.

Performance analyses

We measured the time taken for both tools to complete the full pipeline, from image to RSML output. We timed RootNav 1.0 by annotating random images from each test set from scratch; the total time taken to annotate 10 images of wheat and *Arabidopsis* and 5 of rapeseed was recorded, and averages computed. The annotation was performed by an expert user who ha many years of familiarity with the tool. For RootNav 2.0, we processed each test set and then calculated the average inference time per image. For all tests each image was annotated in sufficient detail to measure the traits seen in Figs 6–8. Results can be found in Table 3.

In both systems a more complex root architecture typically leads to a longer analysis time. In RootNav 1.0 this is due to the human input required; with RootNav 2.0 the path finding takes longer if there are more lateral roots or roots are longer. On some images such as those in the rapeseed dataset, RootNav 1.0 would take less time to process each image if only course traits such as first-order root length were required because significant user time is taken in annotating and correcting second-order root positions. We anticipate that RootNav 2.0 will be used to measure as many traits as possible automatically and so present here a like-for-like comparison where both tools are used to measure the same features of each root architecture. In this comparison, RootNav 2.0 offers a substantial speed advantage over the original tool on each dataset. It should also be noted that the time presented here for RootNav 1.0 requires that the user engage with the software continuously. Because RootNav 2.0 is fully automatic, the human time cost is essentially zero because images could be batch-processed overnight. This test was also run on a single CPU and GPU, where additional computa-

tional resource would linearly scale the speed of the system. If performance were a serious consideration, a dedicated parallel hardware set-up could streamline RootNav 2.0 performance considerably.

Discussion

In this article we have introduced RootNav 2.0, a state-of-the-art, fully automated root-phenotyping tool. It is powered by a deep CNN in an encoder-decoder configuration, designed to perform segmentation efficiently in high-resolution images. The network segments root from background and can distinguish first- and second-order roots. This deep learned root segmentation provides a strong foundation upon which users can derive common architectural traits, such as those based on RSA skeletonization. We have adapted the network, however, to simultaneously predict the location of key root architectural features: the seed location, and first- and second-order root tips. This knowledge then drives a heuristic search that reconstructs the entire root system. This topology is represented as spline curves, and output in RSML format.

A quantitative analysis of RootNav 2.0 shows that it offers comparable accuracy against the original RootNav on large training sets. Over a range of standard trait measurements the new tool produced highly correlated results against the ground truth, with r^2 values ranging from 0.64 to 1. Performance on traits representing the bounds of the root system yielded among the highest r^2 values. On smaller datasets, we have demonstrated that transfer learning produces accurate results despite many fewer training examples. This adaptability is a key advantage for those within the research community looking to use Root-

Table 3: Performance comparison of RootNav 1.0 against RootNav 2.0

Dataset	Average processing time (s)	
	RootNav 1.0	RootNav 2.0
Wheat	68.8	8.2
Arabidposis	109.8	7.5
Brassica	132.4	14.0

The time to process a random sample of images from each test set was measured, and an average time per image calculated.

Nav 2.0; those who use different growth conditions, image capture approaches, or require the analysis of different species can adapt one of our existing trained models with a minimum of effort using transfer learning.

While the accuracy of the fully automatic approach here does not yet match a human-annotated approach across all traits, we believe that this system still offers a benefit over all existing approaches for forward genetic screens on root system architecture, such as quantitative trait locus (QTL) analysis or genome-wide association studies (GWAS). RootNav 2.0 is substantially more convenient to use than previous semi-automated tools, with no human interaction required at any point during the pipeline. The entire process requires <15 seconds processing time per image. The output can be analysed using the RootNav viewer tool or any compatible RSML analysis pipeline. Training the original network took a few days on suitable hardware, with transfer learning to a new dataset typically taking about half a day. The speed of the system allows many more images to be processed, images that may represent additional species or more replicates. Previous work has shown that larger sample sizes have a positive effect on the performance of both QTL analysis and GWAS [46,47] and that good results can be obtained on even simple automatic measures compared with existing semi-automatic approaches [29]. We believe that RootNav 2.0 will prove to be a key milestone in root phenotyping, further encouraging the uptake of machine learning in addressing these important challenges.

In future work, we will continue to adapt this approach to new and varied datasets, maximizing the potential for use in the research community. We will also explore the use of more robust heuristic searches, combined with appropriate segmentation output from the network, to address the challenge of crossing and intersecting root systems. We will also continue to explore deep network developments at the core of the tool, with a view to closing the gap between automatic and semi-automatic approaches on the most challenging traits. To encourage community uptake and engagement, we will also develop mechanisms to ease the sharing of network models, and indeed the retraining process required to adapt them to specific scenarios.

Potential implications

We believe that RootNav 2.0 offers a substantial increase in accuracy over bottom-up approaches to root image analysis. It also offers an increase in throughput over existing semi-automatic tools. Importantly, results on the *Arabidopisis* dataset suggest that the approach will be applicable to images obtained with other phenotyping systems such as rhizotrons. With continued community support, RootNav 2.0 has the potential to be the first true species—and platform—agnostic analysis tool in the plant sciences. This will provide researchers with the ability to analyse root systems at larger scales than ever before, at a time when large-scale genomic studies have made this more important than ever.

Methods

Training, validation, and test image preparation

For each image we obtained ground truth annotations using the original RootNav 1.0 software. This software is semi-automatic and allows users to manually intervene to correct errors in either segmentation or RSA extraction. We used these data as ground truth, rather than to evaluate the accuracy of RootNav 1.0, and as such annotators were instructed to spend sufficient time on each image to correct all mistakes that they could identify. This semi-automatic process often requires a large amount of human interaction and is time consuming, but the approach has provided very reliable ground truth annotations. All ground truth was stored in RSML format.

RSML data for each image were converted into a series of segmentation masks and feature heat maps for use in training. Segmentation masks were created separately for both first- and second-order roots by rendering them as polylines over a blank image. RootNav 1.0 does not measure diameter information for root systems, but the seedlines are sufUbbficiently young that root diameter is quite consistent across species and images. We rendered each root with a width of 8 pixels. For heat map output, the seed location and first- and second-order root tip locations were rendered as in Pound et al. [5], as separate images of blurred Gaussian points of standard deviation 1.0 pixels. The result of these processes is that for each input image there were 5 associated output images, 2 segmentation masks for first- and second-order roots, and 3 heat maps for seed position and first- and second-order root tips (Fig. 9).

At this point, we have constructed 3 suitable training sets of images based on manual annotations. The next task is to construct a suitable encoder-decoder architecture capable of segmenting these images and locating root features.

CNN Design

Input and output resolution

We used the PyTorch [48] framework to develop the network, training, and validation code that drives our segmentation approach. The network is based around an encoder-decoder architecture (Fig. 2) but has been adapted to handle the higher-resolution images seen in the datasets. Encoder-decoder CNNs are memory intensive, particularly at points towards the start and end of the network where the spatial resolution is high. Each layer calculates many features, each of which exists as an image stored in memory. Over many layers, the computational cost becomes prohibitive. Previous work, such as that by Pound et al. [5], used small input and output sizes of 256 × 256 pixels. Other commonly used networks such as VGG-FCN [39] and U-

Table 4: The proposed CNN's layers

Layer architecture		Dimensions	Feature maps
Input (RGB)		$1{,}024 \times 1{,}024$	3
Convolution (7×7)		512×512	64
Residual block 1		512×512	128
Maximum pooling		256×256	128
Residual block 2		256×256	128
Maximum pooling		128×128	128
Residual plock 3		128×128	256
Hourglass (Block = 1)		128×128	256
Transposed convolution		256×256	256
ResNet block 4		256×256	256
Transposed convolution		512×512	256
ResNet block 5		512×512	128
Convolution (1×1)		512×512	128
Network split			
Convolution (1×1)	Convolution (1×1)	512×512	3
Softmax	Convolution (1×1)	512×512	3

Batch normalization layers and ReLU activation functions are used between layers and within residual blocks. The hourglass used is equivalent to a single stack of the type used by Pound et al. [5].

Net [41] use similar input sizes. Root images pose a challenge in this situation because roots may be only a few pixels in diameter but exist as part of a large, connected architecture covering many megapixels. Shrinking the image to a convenient size will make processing simpler but also badly degrade the quality of these small features. In scenarios such as this, where shrinking the input this far may represent a significant loss in quality, it is common to tile the input into small cropped sections and run the network repeatedly. This is the approach taken by Pound et al. [5], in which wheat images are tiled, processed, and then reconstructed. The drawback of tiling images is that each tile is then considered in isolation, removing vital context on its position in the wider image. In the root datasets, for example, first- and second-order roots often appear identical when not viewed as part of a larger architecture.

In this work, we limit the size of input images to $1{,}024 \times 1{,}024$ pixels. For the wheat and rapeseed datasets, this necessitated downsampling of the input and output images, but only by a moderate amount, in which fine root detail is preserved. For the *Arabidopsis* plate images, no downsampling was required because they were already of a suitable size. Upon completion of the deep learning, images are returned to native resolution to ensure that the output measurement scale is preserved.

Network architecture

Our complete network operates on input images of $1{,}024 \times 1{,}024$ pixels and outputs segmentation masks and regression maps of 512×512 pixels. A diagrammatic overview of the network can be found in Fig. 10, with a description of the layers in Table 4. The core of the network is an hourglass architecture similar to those used by Pound et al. [5] and Newell et al. [43], but here we use a restricted number of features throughout and do not use stacked structure (repeated encoder-decoders after one another). These alterations to the network allow it to successfully process the 1-megapixel input size without reaching the limit of available memory. We also perform additional downsampling and upsampling at either end of the network. Initial strided convolutional layers with large filter sizes of 7×7 are used to extract fea-

tures and downsample the image size, before interleaved residual blocks and maximum-pooling operations are used to further reduce the spatial size of the input to 128×128 pixels. The hourglass architecture performs the primary encoder-decoder role, with downsampling performed using maximum pooling and upsampling performed using bilinear interpolation. The output of the hourglass is a set of 128×128 pixel feature maps, after which learned deconvolutional filters and residual blocks are used to return to a 512×512 pixel spatial resolution. Finally, 2 paths are used to separately predict segmentation masks and feature heat maps. Each branch comprises 1×1 convolutional layers for prediction, with the segmentation output also passed through a sigmoid output, as required by the binary cross entropy loss function.

Loss functions

The output of our network is divided into 2 paths with different objectives. The first outputs segmentation masks containing locations for first- and second-order roots. Each of these is a 2D binary output and is trained using a cross-entropy loss. It is common in root images that the number of background pixels heavily outweigh the foreground. Calculating a loss over an unbalanced dataset such as this is likely to cause a bias towards background pixels, causing error and undersegmentation of the foreground. We apply a class-balancing approach to the standard cross-entropy loss, based on median frequency balancing [49]. Weights are assigned to each class inversely proportional to the median frequency in which that class appears throughout the entire training set. This reduces the weight for classes that appear more often, in this case background, and increases the weight of foreground classes such as first-order roots. The accuracy of classes with higher weights is prioritized during training. The proposed loss L_1 is given in Equation (1).

$$L_1 = \frac{\alpha_c}{N} \sum_{n=1}^{N} \sum_{x=1}^{W} \sum_{y=1}^{H} \left[g_{xy}^n \log \left(\hat{g}_{xy}^n \right) + \left(1 - g_{xy}^n \right) \log \left(1 - \hat{g}_{xy}^n \right) \right],$$

(1)

where for N features, $\hat{g}_{xy}^n$ is the predicted class output at location (x, y) and g_{xy}^n is the ground truth prediction at that location. The weight of each class is scaled by its frequency relative to the median frequency of all classes by α_c, given by

$$\alpha_c = \frac{\text{median\_freq}}{\text{freq}(c)},$$

(2)

where $\text{freq}(c)$ is the frequency of occurrences of pixels of class c divided by the number of pixels in any image containing that class and median_freq is the median of these frequencies over all classes.

The second path is responsible for predicting key feature locations on the root system, specifically, the seed location, first-order root tips, and second-order root tips. The output is three 2D outputs, trained using a mean squared error loss, predicting likely locations for root features, represented by 2D Gaussians centred at each feature location. The proposed loss L_2 is depicted in Equation (3).

$$L_2 = \frac{\alpha_c}{N} \sum_{n=1}^{N} \sum_{x=1}^{W} \sum_{y=1}^{H} \left\| p_{xy}^n - \hat{p}_{xy}^n \right\|^2,$$

(3)

Figure 10: The network used is an extended encoder-decoder architecture. The input is filtered and downsampled into an efficient size, before an hourglass network [43] performs the remaining encoding and some decoding. Finally, a series of learned deconvolutional layers upsample back to 50% of the original input size. The network is split into 2 fully convolutional branches at the output, which are responsible for learning the segmentation and heat map regression outputs separately.

where for each of N features, $\hat{p}_{xy}^{n}$ is the predicted feature likelihood at pixel (x, y) and p_{xy}^{n} is the expected ground truth at the same location.

The final loss $L = L_1 + L_2$ trains the network end to end, balancing the objectives of both paths. We found that additional scaling factors applied to the loss of either path were not necessary for accurate training.

Training

Beginning with the wheat dataset, the network was trained end to end from scratch using the rmsprop optimizer. The initial learning rate was set to $1e^{-4}$ and reduced by a factor of 10 after 50,000 iterations. The network was trained using a batch size of 6 for 500,000 iterations, although we found that performance plateaued after approximately 400,000–450,000 iterations. During training, we selected the best-performing model from the validation set.

We applied random augmentation to the training set to reduce potential overfitting. We added random horizontal flipping with a 50% probability during training, as well as random rotation in the range $[-30°, 30°]$. We experimented with random cropping as in Pound et al. [5] but found that cropping often caused the removal of parts of the root system, sacrificing context crucial, e.g., in distinguishing first- and second-order roots. For this reason we did not use random cropping during these experiments.

Transfer learning

Transfer learning is the process of training on new data by beginning with an existing trained network's parameters, rather than randomly initialized weights. We began by training the wheat network to completion. This is a large dataset, with more than sufficient images to train a network reliably from scratch. As noted in the above section, successful training simply means acceptable performance on the validation images. We began with the existing wheat network and retrained on the *A. thaliana* dataset. This is a smaller dataset, but the use of pre-trained weights allows a network to make use of any useful image filters learned during the initial training. We experimented with training from scratch on the smaller dataset but found that we were unable to train a network that performed reliably on the validation set (Supplementary Figs 6 and 7). The same process of training and validation was used to complete training on the new dataset, except that we limited training duration to 120,000 iterations. Normally minor modifications of the un-

derlying deep network would be required to support single-channel near-infrared images rather than RGB. However, to ensure compatibility between models, in particular to simplify transfer learning between images that may move between RGB and single channel, we chose to fix the network and use 3 input channels in all cases. For the *Arabidopsis* dataset, we duplicate the grayscale channel into 3 RGB channels prior to use. The computational overhead of this only affects the weights in the first layer of a very deep network and is marginal. The result is that the *Arabidopsis* model does not make use of the additional channel information in the first layer but remains structurally identical to the other 2 models.

Finally, we repeated transfer learning from the wheat dataset to the *B. napus* data, which have the fewest images of the datasets we use. We explored training from scratch, as well as using pre-trained weights from either the wheat or arabidopsis datasets, and found that the wheat network offered the most reliable starting point, a fact we attribute to the similar background and foreground colours, and scales for both datasets. As above, we trained for 120,000 iterations and selected the model with highest validation performance.

Post-processing

Dense CRF

Each segmentation mask was passed to a dense conditional random field (CRF) to improve smoothness and maximize agreement between similar neighbouring pixels. We found that this approach had a subtle but helpful effect on the separation between roots growing in close proximity, and the smoothness of the boundaries of segmented roots. We used the dense CRF proposed by Krähenbühl and Koltun [50], in which each pair of random variables (pixels) are connected by an edge [51], weighted using a Gaussian pairwise potential. The effect was a smoothing of conflicting regions of pixels where the image was cluttered; but where segmentation was already successful, the approach had no notable negative effect on the results.

Feature localization and non-maximal suppression

The heat map regression output contains the probable locations of the seed, as well as first- and second-order root tips. These are represented as 2D Gaussian distributions, the centre of which lies on likely feature locations. We obtain a discrete location for each feature via a non-maximal suppression (NMS) approach [52], which suppresses all predicted pixels except those that are greater than their surrounding neighbours.

For a Gaussian-based distribution, this has the effect of locating the centre of the distribution. Our implementation of NMS uses pixel-wise search, immediately discounting any pixel output below a pre-defined threshold (in our case, 0.7), for speed. Each pixel is then compared with its neighbour pixels in a 3×3 window, where the central pixel "c" is non-maximal; if another pixel of greater or equal intensity is discovered in its neighbourhood, the algorithm skips to the next pixel in the scan line [53].

For the majority of root tips in isolation, NMS will successfully return a single location for each true root tip. This relies on the heat map regression layers of the network returning well-formed Gaussian distributions in all instances, which while likely, may not occur in the presence of image clutter, confusing root hairs, or multiple tips in close proximity. To avoid 2 positions being returned for a single underlying root feature, we identify and suppress neighbouring features. We use an R-Tree data structure to efficiently query for neighbours within close proximity. When NMS returns a new position on the image, the R-Tree is searched for nearby features that have already been added and prevents locations from being added twice. In our experiments, we considered a new position a duplicate if it fell within 8 pixels radius of an existing feature, which is derived from the scale of the roots in our datasets.

Root architecture reconstruction

After successful pixel-wise segmentation of the complete root system architecture and extraction of the tips and seed locations, we are now able to reconstruct the whole root skeleton. This procedure is similar to the original RootNav 1.0 tool [9], except it is now driven by more accurate class-aware segmentation, rather than error-prone root likelihood estimations. Root-Nav 2.0 can place more reliance on the accuracy of the segmentation and make use of each segmentation map separately to ensure that roots are not traversed over the wrong material, e.g., that first-order roots prioritize image locations of that class. We establish an 8-way connected graph structure throughout the image, where the weights travelling to neighbouring pixels are calculated as a function of their class, the path we are trying to find, and the distance between them. Each segmentation mask is converted into a distance map of values [0, 1] indicating the distance from any background pixel. We then convert this distance into a weighting that prioritizes paths along root centres; the maximum weight we assign to any root pixel is 0.1, for pixels near the root edge. The weight decreases towards the centre of the root, to a minimum of 0.01. Because the graph includes diagonal connections, these are weighted by an additional cost of $\sqrt{2}$ to account for the longer distance. Finally, any pixel that does not belong to the specific class being traversed, e.g., first-order root only, is assigned a weight of 10.0, representing a much stronger penalty for traversing these pixels. Unlike RootNav 1.0, we use separate graphs and searches for first- and second-order roots. A value of 10.0 was chosen simply as a very large increase in weight when compared to the minimum cost for any segmented root material. Different weight values are effective, as long as they are large enough relative to root material to avoid the shortest path taking shortcuts across background pixels where this is unnecessary. In practice, these weights are only traversed if there is a gap in the segmentation for true root material.

A∗ search [54] is a path-finding algorithm that in our implementation seeks to find a path of minimal cost between locations on a root system. It is an extension of Dijkstra's shortest-


Figure 11: Example output from RootNav 2.0. (a) Input image. (b) Colour-coded segmentation mask. (c, d) Binary segmentation masks for first- and second-order roots. (e) A sample of the RSML file representing the entire architecture.

path algorithm [55], and along with distance travelled also considers a heuristic measure of the remaining distance to the goal. Pixels are explored based on the lowest cost first, in order to minimize the function

$$f(p) = g(p) + h(p), \qquad (4)$$

where $g(p)$ is the sum of all weights to p, and $h(p)$ is the remaining distance, which we calculate as the Manhattan distance, or L^1-norm.

In the case of RSA traversal, minimal cost paths between key features such as first-order root tips and seeds represent reconstructed roots. A∗ searches are initialized from all first-order root tips, travelling along segmented roots until they reach any seed point. Upon reaching a seed location, the entire path is recorded as a first-order root. Once all first-order roots have been traversed, a new series of searches are begun from second-order root tip locations, ending at any encountered first-order root. The second-order root searches use Dijkstra's algorithm by not including a remaining distance heuristic $h(p)$, which would be inefficient to calculate over many possible goal locations. The output of each search is a list of pixel coordinates representing the individual roots within the RSA.

The *Arabidopsis* dataset contains multiple plants per image, which can be handled through minor modifications to the path-finding approach. Where the deep network detects multiple seed locations, these are used to initialize multiple candidate plants. We use Dijkstra's algorithm rather than A∗ as with second-order root searches with multiple candidate goal locations. First-order root searches proceed as above but now terminate upon encountering the first seed, which in a Dijkstra search represents the shortest possible path between that tip and any of the seeds. This root path is then assigned to the corresponding plant, and the process is repeated for each first-order tip location. Once the process is complete, any remaining plants whose seed locations did not result in a connection with a first-order root are removed.

Spline fitting

The use of a distance map that prioritizes the centre lines of roots generally acts to smooth the paths found throughout the root system. This may not be the case where there is noise in the segmentation output, or roots cross, and the distance map is less reliable. We smooth each root path using a spline curve representation. Control points are sampled at equal spacing along each path, before the path is resampled using cubic splines. Each spline includes a tension parameter that we set at a constant of 0.5 for our experiments. Both the spline and a polyline representation are output into the final RSML file to ensure maximum compatibility with other tools.

RSML and output

The RSA reconstruction approach in RootNav 2.0 does not perform phenotypic measurements itself; rather it extracts a root topology along with segmented images from which traits can be derived. The entire root system for each plant in an image is exported using the RSML format [34], providing a standard and interoperable format. RSML is an XML document specifically designed to store 2D and 3D root architectures. It also stores metadata and plant properties and is compatible with numerous analysis tools. Some existing plant-phenotyping tools offer RSML import support, meaning that they may also load root systems created automatically using RootNav 2.0. Our approach also outputs first- and second-order root segmentation images, representing an alternative source of quantitative data. Many tools such as Ez-Rhizo operate on such images, but the segmentation masks generated here contain very little image noise, making them more amenable to further automated analysis. An example output can be seen in Fig. 11. For this work we performed quantification entirely using the RSML output. Phenotypic measurements were calculated from each RSML file using the existing Root-Nav Viewer tool, which has been extended and updated for this publication.

Availability of Source Code and Requirements

Project name: RootNav 2.0
Project home page: https://github.com/robail-yasrab/RootNav-2.0
SciCrunch RRID:SCR_015584
biotoolsID: biotools:Rootnav_2.0
Operating system(s): Platform independent
Programming language: Python, C# (RootNav Viewer)
Other requirements: Python 3.6, PyTorch 1.0.1
License: GNU GPL
Any restrictions to use by non-academics: None

Availability of Supporting Data and Materials

Datasets will be made available at https://plantimages.nottingham.ac.uk. Snapshots of our code and other supporting data are available in the GigaScience repository, GigaDB [56].

Additional Files

Additional Figure S1: Extended plots showing quantitative results on the *Arabidopsis thaliana* dataset.
Additional Figure S2: Extended plots showing quantitative results on the Brassical napus dataset.
Additional Figure S3: Example output of RootNav 2.0 on the wheat (*Triticum aestivum* L.) dataset.
Additional Figure S4: Example output of RootNav 2.0 on the *Arabidopsis thaliana* dataset.
Additional Figure S5: Example output of RootNav 2.0 on the *Brassica napus* dataset.
Additional Figure S6: Example output on the *Arabidopsis* dataset showing the results of transfer learning vs. the results when the network is trained from scratch (randomly initialised weights).
Additional Figure S7: Example output on the *Brassica napus* dataset showing the results of transfer learning vs. the results when the network is trained from scratch (randomly initialised weights).
Additional Figure S8: Examples of incorrect segmentation and root system extraction on the *Arabidopsis* dataset.
Additional Figure S9: Examples of incorrect segmentation and root system extraction on the *Brassica napus* dataset.

Abbreviations

CNN: convolutional neural network; CRF: conditional random field; DART: Data Analysis of Root Tracings; DIRT: Digital Imaging of Root Traits; DSLR: digital single-lens reflex; FCN: fully convolutional network; GPU: graphical processing unit; GWAS: genome-wide association study; mIoU: mean Intersection over Union; NMS: non-maximal suppression; QTL: quantitative trait locus; RSA: root system architecture; RSML: Root System Markup Language;

Competing Interests

The authors declare that they have no competing interests.

Funding

Authors' Contributions

M.P.P., A.P.F. and T.P.P. designed the project; R.Y. implemented the software and trained the deep networks under the supervision of M.P.P. with support from A.P.F. and T.P.P.; M.P.P. wrote additional code for release; J.A.A. and D.M.W. designed the protocols for and carried out the collection of the datasets. R.Y. and M.P.P. wrote the manuscript. All authors contributed to and approved the final manuscript.

Acknowledgements

References

1. Pieruschka R, Schurr U. Plant phenotyping: past, present, and future. Plant Phenomics 2019, doi:10.34133/2019/7507131.

2. Lynch JP. Steep, cheap and deep: an ideotype to optimize water and N acquisition by maize root systems. Ann Bot 2013;**112**(2):347–57.

3. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015;**521**(7553):436.

4. Singh A, Ganapathysubramanian B, Singh AK, et al. Machine learning for high-throughput stress phenotyping in plants. Trends Plant Sci 2016;**21**(2):110–24.

5. Pound MP, Atkinson JA, Wells DM, et al. Deep learning for multi-task plant phenotyping. In: 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), Venice, Italy. IEEE; 2017:2055–63, doi:10.1109/ICCVW.2017.241.

6. Ubbens JR, Stavness I. Deep plant phenomics: a deep learning platform for complex plant phenotyping tasks. Front Plant Sci 2017;**8**:1190.

7. Minervini M, Scharr H, Tsaftaris SA. Image analysis: the new bottleneck in plant phenotyping [applications corner]. IEEE Signal Proc Mag 2015;**32**(4):126–31.

8. Atkinson JA, Pound MP, Bennett MJ, et al. Uncovering the hidden half of plants using new advances in root phenotyping. Curr Opin Biotech 2019;**55**:1–8.

9. Pound MP, French AP, Atkinson JA, et al. RootNav: navigating images of complex root architectures. Plant Physiol 2013;**162**(4):1802–14.

10. Lobet G, Pagès L, Draye X. A novel image-analysis toolbox enabling quantitative analysis of root system architecture. Plant Physiol 2011;**157**(1):29–39.

11. Mairhofer S, Zappala S, Tracy S, et al. Recovering complete plant root system architectures from soil via X-ray μ-computed tomography. Plant Methods 2013;**9**(1):8.

12. Schulz H, Postma JA, van Dusschoten D, et al. Plant root system analysis from MRI images. In: Csurka G, Kraus M, Laramee RS , et al., eds. Computer Vision, Imaging and Computer Graphics: Theory and Application. Berlin and Heidelberg: Springer; 2013:411–25.

13. Das A, Schneider H, Burridge J, et al. Digital imaging of root traits (DIRT): a high-throughput computing and collaboration platform for field-based root phenomics. Plant Methods 2015;**11**(1):51.

14. Galkovskyi T, Mileyko Y, Bucksch A, et al. GiA Roots: software for the high throughput analysis of plant root system architecture. BMC Plant Biol 2012;**12**(1):116.

15. Pierret A, Gonkhamdee S, Jourdan C, et al. IJ_Rhizo: an open-source software to measure scanned images of root samples. Plant Soil 2013;**373**(1–2):531–9.

16. Armengaud P, Zambaux K, Hills A, et al. EZ-Rhizo: integrated software for the fast and accurate measurement of root system architecture. Plant J 2009;**57**(5): 945–56.

17. Haralick RM, Sternberg SR, Zhuang X. Image analysis using mathematical morphology. IEEE Trans Pattern Anal Mach Intell 1987;**PAMI-9**(4):532–50.

18. Lobregt S, Verbeek PW, Groen FC. Three-dimensional skeletonization: principle and algorithm. IEEE Trans Pattern Anal Mach Intell 1980;**PAMI-2**(1):75–7.

19. Berzin I, Cohen B, Mills D, et al. RHIZOSCAN: a semiautomatic image processing system for characterization of the morphology and secondary metabolite concentration in hairy root cultures. Biotechnol Bioeng 2000;**70**(1): 17–24.

20. Pradal C, Dufour-Kowalski S, Boudon F, et al. OpenAlea: a visual programming and component-based software platform for plant modelling. Funct Plant Biol 2008;**35**(10):751–60.

21. Abràmoff MD, Magalhães PJ, Ram SJ. Image processing with ImageJ. Biophotonics Int 2004;**11**(7):36–42.

22. Le Bot J, Serra V, Fabre J, et al. DART: a software to analyse root system architecture and development from captured images. Plant Soil 2010;**326**(1–2):261–73.

23. Ristova D, Rosas U, Krouk G, et al. RootScape: a landmark-based system for rapid screening of root architecture in *Arabidopsis*. Plant Physiol 2013;**161**(3):1086–96.

24. Settles B, Craven M, Friedland L. Active learning with real annotation costs. In: Proceedings of the NIPS Workshop on Cost-Sensitive Learning, Vancouver, BC, Canada. 2008:1–10.

25. Rolnick D, Veit A, Belongie S, et al. Deep learning is robust to massive label noise. arXiv 2017:1705.10694.

26. Clark RT, Famoso AN, Zhao K, et al. High-throughput two-dimensional root system phenotyping platform facilitates genetic analysis of root growth and development. Plant Cell Environ 2013;**36**(2):454–66.

27. Mohanty SP, Hughes DP, Salathé M. Using deep learning for image-based plant disease detection. Front Plant Sci 2016;**7**:1419.

28. Itzhaky Y, Farjon G, Khoroshevsky F, et al. Leaf counting: multiple scale regression and detection using deep CNNs. In: Proceedings of the British Machine Vision Conference, Newcastle, UK, 2018. BMVA Press; 2018:328.

29. Pound MP, Atkinson JA, Townsend AJ, et al. Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. Gigascience 2017;**6**(10):gix083.

30. Aich S, Stavness I. Leaf counting with deep convolutional and deconvolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision; 2017:2080–9.

31. Keller K, Kirchgessner N, Khanna R, et al. Soybean leaf coverage estimation with machine learning and thresholding algorithms for field phenotyping. In: Proceedings of the British Machine Vision Conference, Newcastle, UK, 2018. BMVA Press; 2018:32.

32. Atanbori J, Chen F, French AP, et al. Towards low-cost image-based plant phenotyping using reduced-parameter CNN. In: Proceedings of the British Machine Vision Conference, Newcastle, UK, 2018. BMVA Press; 2018:326

33. Chen H, Giuffrida MV, Tsaftaris SA, et al. Root gap correction with a deep inpainting model. In: Proceedings of the British Machine Vision Conference, Newcastle, UK, 2018. BMVA Press; 2018:325.

34. Lobet G, Pound MP, Diener J, et al. Root System Markup Language: toward a unified root architecture description language. Plant Physiol 2015;**167**(3):617–27.

35. Atkinson JA, Wingen LU, Griffiths M, et al. Phenotyping pipeline reveals major seedling root growth QTL in hexaploid wheat. J Exp Bot 2015;**66**(8):2283–92.

36. Wilson MH, Holman TJ, Sørensen I, et al. Multi-omics analysis identifies genes mediating the extension of cell walls in the *Arabidopsis thaliana* root elongation zone. Front Cell Dev Biol 2015;**3**:10.

37. Wells DM, French AP, Naeem A, et al. Recovering the dynamics of root growth and development using novel image acquisition and analysis methods. Philos Trans R Soc Lond B Biol Sci 2012;**367**(1595):1517–24.

38. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. arXiv 2014: 1409.1556.

39. Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2015:3431–40.

40. Badrinarayanan V, Kendall A, Cipolla R. Segnet: a deep convolutional encoder-decoder architecture for image segmentation. arXiv 2015:1511.00561.

41. Ronneberger O, Fischer P, Brox T. U-net: convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer; 2015: 234–41.

42. Chen LC, Zhu Y, Papandreou G, et al. Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Ferrari V, Hebert M, Sminchisescu C , et al., eds. Proceedings of the European Conference on Computer Vision (ECCV). Cham, Springer; 2018:801–18.

43. Newell A, Yang K, Deng J. Stacked hourglass networks for human pose estimation. In: Liebe B, Mata J, Sebe N , et al., eds. European Conference on Computer Vision. Cham: Springer; 2016:483–99.

44. Chaurasia A, Culurciello E. Linknet: exploiting encoder representations for efficient semantic segmentation. In: 2017 IEEE Visual Communications and Image Processing (VCIP). IEEE; 2017:1–4.

45. Zhao H, Shi J, Qi X, et al. Pyramid scene parsing network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017:2881–90.

46. Li H, Hearne S, Bänziger M, et al. Statistical properties of QTL linkage mapping in biparental genetic populations. Heredity 2010;**105**(3):257.

47. Hong EP, Park JW. Sample size and statistical power calculation in genetic association studies. Genomics Iinform 2012;**10**(2):117.

48. Paszke A, Gross S, Chintala S, et al. Automatic differentiation in pytorch. In: NIPS-W. 2017.

49. Eigen D, Fergus R. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile. IEEE; 2015:2650–8, doi:10.1109/ICCV.2015.304.

50. Krähenbühl P, Koltun V. Efficient inference in fully connected CRFs with Gaussian edge potentials. In: Proceedings of the 24th International Conference on Neural Information Processing Systems, Grenada, Spain. 2011:109–117.

51. Desmaison A, Bunel R, Kohli P, et al. Efficient continuous relaxations for dense CRF. In: European Conference on Computer Vision. Springer; 2016:818–33.

52. Lindeberg T. Edge detection and ridge detection with automatic scale selection. Int J Comput Vision 1998;**30**(2):117–56.

53. Pham TQ. Non-maximum suppression using fewer than two comparisons per pixel. In: International Conference on Advanced Concepts for Intelligent Vision Systems. Springer; 2010:438–51.

54. Doran JE, Michie D. Experiments with the Graph Traverser program. Proc R Soc Lond A Math Phys Sci 1966;**294**(1437):235–59.

55. Dijkstra EW. A note on two problems in connexion with graphs. Numer Math 1959;**1**(1):269–71.

56. Yasrab R, Atkinson JA, Wells DM, et al. Supporting data for "RootNav 2.0: Deep learning for automatic navigation of complex plant root architectures. GigaScience Database 2019; http://dx.doi.org/10.5524/100651.