Artificial intelligence enables the identification and quantification of arbuscular mycorrhizal fungi in plant roots

Edouard Evangelisti[1*], Carl Turner[2], Alice McDowell[3], Liron Shenhav[1], Temur Yunusov[1], Aleksandr Gavrin[1], Emily K. Servante[4], Clément Quan[1], Sebastian Schornack[1*].

[1] The Sainsbury Laboratory, Cambridge University (SLCU). Cambridge, UK.

[2] Department of Applied Mathematics and Theoretical Physics (DAMTP), University of Cambridge. Cambridge, UK.

[3] Present address: Department of Biochemistry, University of Cambridge. Cambridge, UK.

[4] Department of Plant Sciences, University of Cambridge. Cambridge, UK.

* Correspondence: edouard.evangelisti@slcu.cam.ac.uk, sebastian.schornack@slcu.cam.ac.uk

Summary

Soil fungi establish mutualistic interactions with the roots of most vascular land plants. Arbuscular mycorrhizal (AM) fungi are among the most extensively characterised mycobionts to date. Current approaches to quantifying the extent of root colonisation and the relative abundance of intraradical hyphal structures in mutant roots rely on staining and human scoring involving simple, yet repetitive tasks prone to variations between experimenters.

We developed the software AMFinder which allows for automatic computer vision-based identification and quantification of AM fungal colonisation and intraradical hyphal structures on ink-stained root images using convolutional neural networks.

AMFinder delivered high-confidence predictions on image datasets of colonised roots of *Medicago truncatula*, *Lotus japonicus*, *Oryza sativa* and *Nicotiana benthamiana* obtained via flatbed scanning or digital microscopy enabling reproducible and transparent data analysis. A streamlined protocol for sample preparation and imaging allowed us to quantify dynamic increases in colonisation in whole root systems over time.

AMFinder adapts to a wide array of experimental conditions. It enables accurate, reproducible analyses of plant root systems and will support better documentation of AM fungal colonisation analyses.

AMFinder can be accessed here: https://github.com/SchornacklabSLCU/amfinder.git

Keywords: classification, ClearSee, ConvNet, image analysis, mycorrhiza, *Rhizophagus*, root.

Introduction

35   Soil fungi establish mutualistic interactions with the roots of more than 85% of vascular land plants (Brundrett & Tedersoo, 2018). These interactions, termed mycorrhizae, lead either to the formation of a dense hyphal sheath surrounding the root surface (ectomycorrhizae) or to fungal hyphae penetrating host tissues (endomycorrhizae) (Brundrett, 2002). The best-characterized type of endomycorrhiza, called arbuscular mycorrhiza (AM), involves species

40   from the subphylum Glomeromycotina (Schüßler *et al.*, 2001; Spatafora *et al.*, 2016). AM fungal hyphae grow toward plant roots following the exchange of diffusible chemical cues (Luginbuehl & Oldroyd, 2017). At root surface penetration points, hyphae differentiate into swollen or branched structures termed hyphopodia. Following entry and crossing of the root epidermis, hyphae spread either between cortical cells (Arum-type colonization) or via

45   intracellular passages of cortical cells (Paris-type colonization) (Dickson, 2004). The differentiation of highly-branched intracellular exchange structures, the arbuscules, accompanies hyphal growth and enables reciprocal transfer of nutrients (Luginbuehl & Oldroyd, 2017). Post-arbuscular development includes the differentiation of vesicles and sporulation. While these successive differentiation events reflect a precise morphogenetic

50   program, the whole hyphal network is not synchronized. As a result, the various types of intraradical hyphal structures occur simultaneously inside plant roots (Montero *et al.*, 2019).

      Rhizophagus irregularis (formerly *Glomus intraradices*) is one of the most extensively characterised mycobionts in endomycorrhiza research. To date, genetic manipulation of *R. irregularis* remains challenging (Helber & Requena, 2008) and the main advances in AM

55   fungal symbiosis research relate to the experimentally more tractable plant hosts. The extent of root colonisation and the relative abundance of intraradical hyphal structures in mutant roots are essential parameters for characterising host genes that underlie mycorrhiza establishment and accommodation (Montero *et al.*, 2019). Mycorrhiza-responsive host genes facilitate the molecular quantification of fungal colonisation. For instance, expression of the

60   *Medicago truncatula Phosphate transporter 4* (*MtPT4*) gene is limited to the root tip in the absence of mycorrhiza (Volpe *et al.*, 2016), while cells with arbuscules express *MtPT4* to enable plant acquisition of inorganic phosphate (Harrison *et al.*, 2002; Maeda *et al.*, 2006; Javot *et al.*, 2007). Likewise, the abundance of transcripts encoding *M. truncatula Blue Copper-Binding Protein 1* and *Lotus japonicus* apoplastic subtilase *SbtM* correlate with stage

65   transitions during arbuscule development (Hohnjec *et al.*, 2005; Takeda *et al.*, 2009; Parádi *et al.*, 2010). Complementary to molecular methods and independent of gene sequence

knowledge is the visual diagnosis of AM fungal colonisation. It consists in differential staining of fungal cell walls (Vierheilig *et al.*, 1998, 2005; Hulse, 2018) followed by random sampling and counting using a grid-intersect method (Giovannetti & Mosse, 1980). This

70   method is considered a standard in mycorrhiza research (Sun & Tang, 2012).

Deep learning encompasses an extensive class of computational models that learn to extract information from raw data at multiple levels of abstraction, thereby mimicking how the human brain perceives and understands information (Voulodimos *et al.*, 2018). In supervised learning problems, where example data labelled with correct outputs are available,

75   these models can be iteratively improved to minimise discrepancies between correct and model-predicted outputs considering all possible interfering factors (O'Mahony *et al.*, 2020). With the increase in computing power over recent years, deep learning has fostered tremendous data analysis advances. Computer vision is one of the most iconic examples, with the development of convolutional neural networks (CNNs), a class of deep learning methods

80   inspired by models of the visual system's structure (LeCun *et al.*, 1998). A typical CNN architecture comprises three types of processing (or neural) layers. First, a convolutional layer uses a set of local receptive fields called filters or kernels to extract elementary visual features (e.g. edges or arbitrary shapes) from a group of neighbouring pixels. The resulting feature map then feeds a pooling layer. Pooling layers are down-sampling steps aiming to reduce

85   feature map width and height for the next convolutional layer. A typical CNN allows for detecting high-order features through several rounds of convolutional and pooling layers (Voulodimos *et al.*, 2018; Dhillon & Verma, 2020). Fully connected layers follow to enable high-level reasoning and decision-making. As their name implies, these layers contain neurons that have full connections to the previous layer. Fully connected layers convert two-

90   dimensional feature maps into a one-dimension feature vector that can be fed forward into categories for classification (Krizhevsky *et al.*, 2017) or used for further processing (Girshick *et al.*, 2014). CNNs underlie breakthrough advances in diverse technological and biomedical domains including face recognition, object detection, diagnostic imaging, and self-driving cars (Matsugu *et al.*, 2003; Szarvas *et al.*, 2005; Bojarski *et al.*, 2016; Yamashita *et al.*, 2018).

95   CNNs are also used in botany (Liu *et al.*, 2017) and plant pathology (Mohanty *et al.*, 2016; Ferentinos, 2018).

We took advantage of CNNs to develop the Automatic Mycorrhiza Finder (AMFinder), an automatic, user-supervised tool suite for *in silico* analysis of AM fungal colonisation and recognition of intraradical hyphal structures. Using AMFinder, we quantified

fungal colonisation dynamics on whole *Nicotiana benthamiana* root systems using low-resolution, flatbed scanner-acquired scans of ink-stained roots. Moreover, AMFinder robustly identified colonised root sections and intraradical hyphal structures on several model species used in mycorrhiza research, including *Medicago truncatula*, *Lotus japonicus*, and *Oryza sativa*. We developed a standalone graphical browser to enable efficient browsing of large images and manual curation of computer predictions. Overall, our work provides a framework for reproducible automated phenotyping of AM fungal colonisation of plant roots.

Materials and Methods

Plant material

Nicotiana benthamiana is a laboratory cultivar obtained from The Sainsbury Laboratory, Norwich, UK, originating from Australia (Bally *et al.*, 2018). *Medicago truncatula* R108 seeds were provided by Giles Oldroyd (The Sainsbury Laboratory, UK). *Lotus japonicus* Gifu seeds were provided by Simona Radutoiu (Aarhus University, Denmark). Rice (*Oryza sativa* subsp. *japonica*) plant material, growth conditions and AM colonisation conditions were described elsewhere (Choi *et al.*, 2020).

Seed germination

N. benthamiana seeds were germinated on Levington F2 compost (ICL, Ipswich, UK) for one week at 24°C with a 16-h photoperiod. *M. truncatula* seeds were scarified in sulphuric acid for 5 min, rinsed in sterile water and surface-sterilized in bleach for 5 min. Seeds were then soaked in water for 30 min and stratified for 3 days at 4°C in the dark. *L. japonicus* seeds were scarified with sandpaper, surface-sterilized in bleach for 15 min and soaked overnight in water at 4°C. Germination was induced at 20°C.

Growth conditions for AM colonisation

One-week-old seedlings were transferred to 6×5 cellular trays containing silver sand supplemented with a 1:10 volume of *R. irregularis* crude inoculum (PlantWorks, Sittingbourne, UK) and grown at 24°C with a 16-h photoperiod. *N. benthamiana* plants were watered with a low-phosphate Long Ashton nutrient solution (Hewitt, 1966), while milliQ water was used for *L. japonicus* and *M. truncatula* plants. Plant roots were harvested at either 4 or 6 weeks post-inoculation and directly used for staining or total mRNA extraction.

Fungal staining

130 A modified ink-vinegar method (Vierheilig *et al.*, 1998) was used to stain fungal structures within plant roots. Briefly, roots from 4- and 6-week-old plants were incubated in 10% (w/v) potassium hydroxide (KOH) for 10 min at 95°C and rinsed in 5% (v/v) acetic acid before staining in staining solution (5% (v/v) Sheaffer Skrip black ink, 5% (v/v) acetic acid) (A.T. Cross Company, Providence, RI, USA) for 10 min at 95°C. Stained roots were rinsed in

135 distilled water, followed by clearing in ClearSee (Kurihara *et al.*, 2015) for 30 s. Cleared roots were mounted in a glycerol-containing mounting medium (20% (v/v) glycerol, 50 mM Tris-HCl pH 7.5, 0.1% (v/v) Tween-20).

Scanning and bright field imaging

Low-magnification images of ink-stained roots were acquired with an Epson Perfection

140 flatbed scanner (Epson UK, Hemel Hempstead, UK) using default settings and a resolution of 3200 dots per inch. High-magnification images were acquired with a VHX-5000 digital microscope (Keyence, Milton Keynes, UK) equipped with a 50-200× zoom lens set to 200× magnification, using transillumination mode.

Generation of modified image datasets

145 Image modifications were achieved using the batch processing tool `convert` from the ImageMagick suite (https://imagemagick.org). Sharpening was achieved using `-unsharp 5`. The parameter `-blur 2x8` was used for image blurring. Gaussian and plasma noise were added by generating composite images (`-compose over -composite`). Gaussian blur was obtained with `+clone -fill gray -colorize 100 +noise gaussian -colorspace gray`

150 `-alpha on -channel a -evaluate multiply 0.3 +channel`. Plasma noise was added using `-size 256x256 plasma: -colorspace gray -alpha set -channel a -evaluate multiply 0.3 +channel`. Hue modifications were achieved with `-define modulate:colorspace=HSB -modulate 100,100,$(seq -180 .01 180 | shuf | head -n1)`. Desaturation was achieved with `-set colorspace Gray -separate -average -`

155 `colorspace sRGB -type truecolor`. Low brightness was achieved with `-modulate 20,10,100`. Colour inversion was achieved with `-channel RGB -negate`. The parameter `-normalize` was used to increase the contrast of modified images, except where brightness had been changed.

Quantitative RT-PCR

Total RNA was extracted from 100 mg root material using RNEasy Mini Kit (Qiagen, Hilden, Germany) according to the manufacturer's instructions. Quality was assessed by electrophoresis on an agarose gel. One microgram was reverse transcribed to generate first-strand complementary DNA (cDNA) using IScript cDNA Synthesis kit (Bio-Rad, Hercules, CA, USA). qRT-PCR experiments were performed with 2.5 µl of a 1:20 dilution of first-strand cDNA and LightCycler 480 SYBR Green I Master mix, according to the manufacturer's instructions (Roche, Basel, Switzerland). Gene-specific oligonucleotides were designed with BatchPrimer3 software (You et al., 2008), and their specificity was validated by analysing dissociation curves after each run. Genes encoding L23 (*Niben101Scf01444g02009*) and FBOX (*Niben101Scf04495g02005*) were selected as constitutive internal controls for *N. benthamiana* (Liu *et al.,* 2012). Transcripts encoding RiEF1α were amplified using the forward primer 5'-TGTTGCTTTCGTCCCAATATC-3' and the reverse primer 5'-GGTTTATCGGTAGGTCGAG-3', while primers 5'-TCGAAGATGGAGTTCCGGCA-3' and 5'-AATCAACATCAATGGTCCAGCC-3' were used to amplify transcripts encoding NbBCP1b (Niben101Scf07438g04015.1). Six biological replicates of the entire experiment were performed. Gene expression was normalised with respect to constitutively expressed internal controls (Vandesompele *et al.*, 2002) and plotted using R (https://www.r-project.org/).

Software implementation, licensing and code availability

The command-line tool `amf` was implemented in Python (https://www.python.org/) using the deep learning library Keras (https://keras.io/). The graphical interface `amfbrowser` was implemented in OCaml (https://ocaml.org/) using the 2D graphics library Cairo (https://www.cairographics.org/) and the cross-platform widget toolkit GTK (https://www.gtk.org/). Source code for both applications is released under the terms of the open-source MIT license (https://opensource.org/licenses/MIT) and is available on Github (https://github.com/SchornacklabSLCU/amfinder.git).

Results

AMFinder uses computer vision for *in silico* analysis of roots colonised by AM fungi

The extent of fungal root colonisation is an important parameter used to characterise mutualistic relationships between AM fungi and plants. We developed the artificial intelligence-based software Automatic Mycorrhiza Finder (AMFinder) to enable a straightforward, automated, and reproducible estimation of this parameter. AMFinder uses a computer vision approach to quantify fungal colonisation and intraradical hyphal structures in root pictures. It comprises a command-line program (`amf`) for automatic root image analysis and a standalone interface (`amfbrowser`) for user supervision of computer predictions **(Fig. 1, S1)**.

The AMFinder prediction pipeline consists of five steps. During the initial pre-processing step, root images are divided into tiles by `amf` using a user-defined tile size depending on image magnification and resolution **(Fig. 1a)**. The first round of predictions follows. A convolutional neural network (CNN1) labels colonised root segments by analysing tiles individually **(Fig. 1a, S1)**. CNN1 comprises four blocks of 3×3 filters (convolutions) interleaved with size-reduction layers (maximum pooling) **(Fig. S1a)**, followed by a classifier made of three fully connected layers which compute the probabilities of each tile belonging to the mutually exclusive classes 'colonised root section' (M+), 'non-colonised root section' (M−), and 'background/not a root/other' **(Fig. 1a, Fig. S1b)**.

The third step consists of the user-supervised conversion of CNN1 predictions to annotations using `amfbrowser` **(Fig. 1a, b)**. CNN1-returned probabilities can be visualised in their image context as pie charts. `amfbrowser` can perform an automatic conversion by using the class with the highest confidence as an annotation. Tools are provided for users to specifically review tiles with low-confidence predictions (i.e. those having all probabilities close to 1/3) **(Fig. 1b, toolbox 6)**. A magnified view of the active tile surrounded by its eight neighbours facilitates manual inspection and identification of fungal structures located at tile edges **(Fig. 1b)**. A fixed-size sliding window (13×12 tiles) gives an overview of a larger image area and ensures optimal tile display irrespective of the overall image dimensions **(Fig. 1b)**. For immediate visual distinction, final annotations are displayed as squares **(Fig. 1a, b)**.

Figure 1. AMFinder prediction pipeline enables semi-automated, user-supervised analysis of AM fungal colonisation *in silico*. **(a)** AMFinder uses a two-stage prediction pipeline for image annotation. First, input images are split into tiles and processed by `amf` neural network 1 (CNN1) to identify colonised root sections (prediction stage 1). If resolution allows, colonised areas can be further analysed to identify intraradical hyphal structures (prediction stage 2, CNN2). **(b)** Representative screenshot of an `amfbrowser` annotation session. **(1)** Buttons to switch the display between prediction stages. **(2)** Clickable buttons to define the annotations present in the active tile. **(3)** Magnified view of the active tile (red square) and eight surrounding tiles. **(4)** Annotation mosaic overview. **(5)** Layer toolbar to filter the display. Numbers indicate annotation counts for the whole image. **(6)** Prediction toolbar to load predictions, fix ambiguous cases and generate annotations. **(7)** Export functions.

After converting CNN1 predictions to annotations and upon user request, `amf` can proceed with a more detailed analysis of AM fungal hyphal structures **(Fig. 1a)**. An independent convolutional neural network (CNN2) then predicts the presence of arbuscules (A), vesicles (V), hyphopodia (H), and intraradical hyphae (IH) on M+ tiles only **(Fig. 1a, b)**. The CNN2 architecture is essentially the same as that of CNN1 **(Fig. S1a)**. However, the probability that each different type of intraradical hyphal structure is present is computed by a separate stack of three fully connected layers atop the convolutional and pooling layers **(Fig.**

S1c). Therefore, each classifier returns a single probability. As for CNN1, CNN2 predictions can be displayed in their image context using `amfbrowser` for manual inspection **(Fig. 1b)**. Since each fungal structure receives an independent probability, CNN2 scores are displayed as radar charts. Radar charts consist of three concentric circles, with the outermost circle

225 corresponding to highest confidence **(Fig. 1a, b)**, overlaid with coloured dots corresponding to individual hyphal structures. Dot positioning reflects prediction confidence. Automatic conversion to annotations (using 0.5 as probability threshold between presence and absence) is also available.

amf and `amfbrowser` communicate through a standard ZIP archive file that stores `amf`

230 probabilities, user annotations, and image settings. Together, the AMFinder pipeline consisting of probability-scoring using `amf` and visual inspection and analysis of results in `amfbrowser` enables semi-automated, user-supervised, high-precision analysis of AM fungal colonisation *in silico*.

Figure 2. AMFinder training pipeline. A set of manually annotated images (green) get pre-processed to generate test (black) and validation (grey) tile sets. *N* training cycles (epochs) follow, where CNN internal parameters are adjusted based on the test set. The validation set allows for independent monitoring of CNN performance by assessing model misfit (loss). An early stopping mechanism (blue) terminates training if loss value does not further decrease in 12 consecutive epochs. Best-performing parameters (blue) are saved for further use.

AMFinder CNNs can be trained to maximise versatility

235 A key prerequisite of AMFinder is having CNNs trained to recognise the desired structures **(Fig. 2, 3)**. The `amf` training pipeline uses a set of images that are split into tiles and manually annotated within `amfbrowser`. First, `amf` randomly assigns annotated tiles to two groups termed training **(Fig. 2, black)** and validation **(Fig. 2, grey)** tile subsets. For optimal training, `amf` compensates for background over-representation by randomly removing excess

240 background tiles and assigns training weights based on tile count in each annotation class to account for any residual imbalance. One hundred training cycles (epochs) follow where the training subset is used to adjust CNN model parameters **(Fig. 2)**. The validation set is used at the end of each epoch to estimate model accuracy on independent data and detect overfitting,

i.e. a performance decrease due to specialisation toward the training dataset. Model performance assessment relies on two evaluation metrics: the accuracy, which is the ratio between the number of correct predictions and the total number of predictions, and the loss, which is a measure of the distance between the correct annotations and the model's predictions. Consistent with their respective output, CNN1 uses categorical cross-entropy as its loss function, while each CNN2 classifier uses binary cross-entropy (Gordon-Rodriguez *et al.*, 2020). To prevent overfitting, an early stopping mechanism prematurely terminates training and restores the best-performing model parameters if the loss does not decrease for twelve training cycles in a row **(Fig. 2)**. These steps result in CNNs best trained to recognise the desired structures and ensure versatility of AMFinder to a range of different types of user-provided images.

Generation of trained networks for ink-stained root images.

We generated pre-trained CNNs to ensure immediate workability of AMFinder. First, we trained CNN1 with 32 images (43855 tiles) of ink-stained *N. benthamiana* roots acquired at low- (tile size: 40×40 pixels), medium- (tile size: 126×126 pixels) and high-resolutions (tile size: 256×256 pixels) **(Fig. 3a, b)**. To ensure optimal brightness and contrast, we cleared roots with ClearSee (Kurihara *et al.*, 2015) following ink staining **(Fig. S2)**. We compromised between scanning speed and resolution when acquiring low-resolution images **(Fig. S3)**. Thirty per cent of the dataset (13055 tiles) belonged to the 'colonised root section' annotation class, while 38% (16536 tiles) corresponded to non-colonised roots. The remaining 33% (14264 tiles) were background tiles. The CNN1 loss function reached a minimum after 20 epochs **(Fig. 3a, left)**, while CNN1 accuracy was 94% **(Fig. 3a, right)**. To better characterise CNN1 performance, we then determined the model precision and specificity **(Fig. S4a-f)**. CNN1 model precision was 92%, 81%, and 95% for classes 'M+', 'M−', and 'Other', respectively **(Fig S4b)**, while specificity was 99%, 95%, and 98% **(Fig S4c)**.

Figure 3. AMFinder consistently labels colonised root sections and intraradical hyphal structures in *Nicotiana benthamiana* roots. (a) CNN1 loss (left) and accuracy (right) plots obtained after training with 32 low- and high-resolution images of ink-stained *N. benthamiana* roots. **(b)** Representative images of CNN1 predictions on low- (top) and high-resolution (bottom) images. **(c)** Loss (left) and accuracy (right) plots obtained after CNN2 training with colonised root sections from 64 high-resolution images of ink-stained *N. benthamiana* roots. **(d)** Representative images of CNN2 predictions made on colonised root sections of a high-resolution image. Asterisks (red) were added to indicate mispredictions. Scale bar is given in micrometers.

Next, we trained CNN2 using colonised (M+) tiles from 64 high-resolution images of

270   ink-stained, ClearSee-treated *N. benthamiana* roots **(Fig. 3c, d)**. Arbuscules occurred in 7607 tiles, while 6356 and 6180 tiles contained vesicles and intraradical hyphae, respectively. Only

100 tiles contained hyphopodia, preventing us from achieving efficient hyphopodia training due to the scarcity of this hyphal structure. CNN2 loss was minimal after 24 epochs, while accuracy was 93%, 88%, and 86% for vesicles, arbuscules, and intraradical hyphae, respectively **(Fig. 3c, d)**. CNN2 model precision was 84%, 86%, and 88% for vesicles, arbuscules, and intraradical hyphae, respectively **(Fig S4g)**, while specificity was 78%, 64%, and 46% **(Fig S4h)**.

We further assessed the AMFinder CNNs' accuracy on an independent dataset comprising 10 low- and 10 high-resolution ink-stained, ClearSee-treated *N. benthamiana* root images **(Fig. 3b, S5-S9)**. AMFinder identified roots and background in all tested images. Colonised and non-colonised root areas were also well delineated. Overall, CNN1's accuracy was 97% on low-resolution images and 96.5% on high-resolution images **(Fig. S5A, S7, S8)**. Consistent with the multiple tile sizes used for training, changes in tile size did not significantly affect CNN1's performance **(Fig. S6)**. Thus, CNN1 consistently labels fungal colonisation of *N. benthamiana* roots irrespective of the image resolution, suggesting it may be compatible with a wide range of acquisition devices.

We reused the high-resolution images to assess CNN2 accuracy **(Fig. S5b, S9)**. We found that CNN2 accuracy was 96%, 99.5%, and 97% for arbuscules, vesicles, and intraradical hyphae, respectively **(Fig. S5b, S9)**. Hence, AMFinder enables a more detailed analysis of fungal hyphal structures, suggesting it may be used to monitor intraradical hyphal structure abundance within host roots. Both AMFinder CNNs converged to high-accuracy scores and thus can robustly identify AM fungal colonisation and intraradical structures. Pre-trained CNNs are available on Github (https://github.com/SchornacklabSLCU/amfinder.git).

Brightness, contrast, and colour are instrumental to pre-trained network performance

Neural network accuracy is affected by discrepancies between images used for predictions, and the training dataset. To provide guidelines about image settings required for optimal performance of our pre-trained models, we assessed CNN1 pre-trained model accuracy on high-resolution images following modification of either edge sharpness, noise levels, brightness, or colours **(Fig. 4, S10)**. The pre-trained model was robust to sharpening and blurring **(Fig. 4a)**, likely due to the smooth edges of low-resolution images used for training. Similarly, an increase in gaussian noise did not significantly alter model performance **(Fig.**

4a). By contrast, low brightness or contrast, changes in hue or saturation, and color inversion, significantly increased misprediction rate up to 73% **(Fig. 4a)**.

Figure 4. AMFinder can adapt to any image type. (a) CNN1 prediction error rate on high-resolution images after pre-processing to either alter edges (sharpen, blur), increase noise (gaussian, plasma), reduce brightness, or change colours (hue, grayscale, invert). Statistical significance was assessed with ANOVA ($p < 0.05$). Letters indicate statistical groupings. Average rates are shown in italics. **(b)** Representative set of tiles used for *de novo* CNN1 training on images bearing the same modifications as in **a**. **(c)** Loss (left) and accuracy (right) plots obtained after CNN1 training (N = 32 images). **(d)** Representative examples of CNN1 predictions obtained with pre-trained and *de novo*-trained models. Asterisks (red) indicate mispredictions.

To assess whether these limitations are due to the training dataset, we trained CNN1 *de novo* using high-resolution images bearing the previously tested modifications **(Fig. 4b, c)**. The loss function reached a minimum at epoch 16, leading to an accuracy of 92% **(Fig. 4c)**. We then used the new model to predict AM fungal colonisation on independent images **(Fig. 4d, S10)**. By contrast with the original pre-trained model, the new version accurately labelled colonised root segments, roots and background in all the tested conditions **(Fig. 4d, S10)**. Hence, adjusting general image parameters is essential to achieve optimal AMFinder performance without the need for computer-intensive calculations. A guide is provided to troubleshoot the most frequent prediction issues **(Table S1)**. Training further allows AMFinder to analyse highly dissimilar datasets, suggesting our software can adapt to any image type independent of fungal staining method and imaging system.

AMFinder performs consistently on multiple host model species.

A wide range of plants is used in endomycorrhiza research including legumes and monocot species with various root size and morphology. We assessed the suitability of pre-trained CNNs models trained on *N. benthamiana* root images to predict AM fungal colonisation and intraradical hyphal structures on colonised root images from *Lotus japonicus* cv. Gifu **(Fig. 5a)**, *Medicago truncatula* ecotype R108 **(Fig. 5b)**, and *Oryza sativa* cv. Nipponbare **(Fig. 5c)**. The image contrast of ClearSee-treated *L. japonicus* and *M. truncatula* roots was similar to *N. benthamiana*. Conversely, large lateral roots of *O. sativa* showed higher background **(Fig. 5c)**. AMFinder correctly identified roots and background in all tested images, with colonised and non-colonised root areas being accurately resolved **(Fig. 5a-c)**, including in cases where colonisation was restricted to inner cortical cell files **(Fig. 5a)**. All types of intraradical hyphal structures were accurately recognised **(Fig. 5a-c)**. In addition, CNN2 identified hyphopodia in four tiles, although only one was a *bona fide* prediction **(Fig. 5b-c)**. In the absence of extensive training for hyphopodia identification, these results should be considered with caution. Together, these data suggest our pre-trained CNNs are compatible with multiple host plants. They can serve as a basis for further model refinement without the need for extensive training datasets.

Figure 5. AMFinder accurately identifies AM fungal colonisation and intraradical structures in roots of various host model species. (a-c) Roots from *Lotus japonicus* cv. Gifu **(a)**, *Medicago truncatula* ecotype R108 **(b)**, and *Oryza sativa* cv. Nipponbare **(c)** were stained with ink and cleared with ClearSee prior to analysis with AMFinder. Dashed square inserts indicate areas used for AMFinder predictions. Scale bar is given in micrometers.

AMFinder enables *in silico* quantification of AM fungal colonisation dynamics

335   We next investigated whether AMFinder could be used to reliably quantify AM colonisation changes of plant roots over time. To that end, we assessed AM fungal colonisation extent on *N. benthamiana* roots harvested after a 4- or 6-week co-cultivation with *R. irregularis* **(Fig. 6)**. First, we monitored the accumulation of transcripts encoding a *N. benthamiana* ortholog of the mycorrhiza-responsive gene *MtBCP1b* (Parádi *et al.*, 2010) **(Fig. 6a)** and quantified

340   fungal biomass by monitoring *R. irregularis EF1α* transcript levels **(Fig. 6b)**. Both methods showed a significant, two- to three-fold increase in fungal content at 6 wpi compared to 4 wpi **(Fig. 6a, b)**. Then, using the grid-line intersect method (Giovannetti & Mosse, 1980) we studied the colonisation extent within randomly sampled root fragments **(Fig. 6c)**. Consistent with the molecular analysis, more colonisation was observed at 6 wpi. The analysis of the

same samples with AMFinder gave similar results **(Fig. 6d)**. We then tested AMFinder's ability to predict fungal colonisation on low-resolution flatbed scanner pictures of whole root systems **(Fig. 6e)**. Consistent with random sampling and molecular data, AM fungal colonisation levels were significantly higher at 6 wpi, although the colonisation extent values were lower than those obtained through random sampling **(Fig. 6e)**. Thus, AMFinder allows for *in silico* quantification of AM fungal colonisation of plant roots over time including in whole root systems.

Figure 6. Computer vision enables quantification of AM fungal colonization on whole root systems. *Nicotiana benthamiana* plants were inoculated with *Rhizophagus irregularis*. Colonization levels were determined 4 and 6 weeks post-inoculation (wpi). **(a-b)** Quantification of transcripts corresponding to the *N. benthamiana* mycorrhizal marker gene *Blue Copper Protein 1* (*BCP1*) **(a)** and the *R. irregularis* elongation factor 1α **(b)**. Data are expressed relative to both *N. benthamiana L23* and *F-box* reference genes. **(c-d)** Quantification of AM fungal colonization on ink-stained root pictures using the gridline intersect method **(c)**, or AMFinder **(d)**. **(e)** Quantification of AM fungal colonisation on whole root systems using AMFinder, and representative images of computer-generated maps featuring colonized (M+, magenta) and non-colonized (M−, grey) root areas at 4 and 6 wpi. Dots correspond to biological replicates. Bars represent standard error. Statistical significance was assessed using Student's *t*-test (*: $p < 0.05$; **: $p < 0.01$).

Discussion

We developed the software AMFinder, which uses two convolutional neural networks to annotate and quantify AM fungi in plant roots. AMFinder performs consistently well on root images of several model plant species used in endosymbiosis research. AMFinder-mediated quantification of AM fungal colonisation gives similar results to those obtained using current standard counting methods. We further show that AMFinder can process whole root systems using low-resolution flatbed scans obtained from an optimised ink-staining protocol which relies on ClearSee as a contrast enhancer. We illustrate the usefulness of this approach to study fungal colonisation dynamics over time in wild-type and mutant plants.

AMFinder deploys on Microsoft Windows, macOS and GNU/Linux and is compatible with an installation in a virtual environment. The training and prediction tool `amf` is implemented in Python (van Rossum & Drake, 2009; Srinath, 2017) to benefit from widely used machine learning libraries (Chollet & others, 2015; Abadi *et al.*, 2016). Its command-line interface is suitable for batch processing and makes it compatible with high-performance computing clusters. The graphical interface `amfbrowser` is implemented in OCaml for enhanced language expressiveness and performance (Leroy *et al.*, 2020). These tools constitute a highly flexible tool suite that fits many computing systems and experimental setups.

AMFinder's design adequately addresses limitations arising from its computer vision approach while still enabling a low- to medium-throughput workflow. Specifically, we implemented a semi-automatic pipeline that requires user supervision of computer predictions. High-throughput AMFinder analyses, such as large-scale field experiments, would require the entire prediction pipeline to be fully automatic, including the conversion to annotations. However, input image parameters can influence pre-trained model accuracy and may require user adjustments. Automatic analyses assuming image suitability without quality control may overestimate CNN model accuracy. Besides, CNN2 predictions on mispredicted M+ tiles without intraradical hyphal structures have not been investigated in this study. This AMFinder implementation does not discriminate between multiple root types. Therefore, the quantification of AM fungal colonisation may be affected by highly colonised contaminants. Image data from experiments relying on crude inoculum (Habte & Byappanhalli, 1998) or nurse root systems such as chives (Demchenko *et al.*, 2004) as an inoculation method may pose problems when contaminating root fragments remain in root images. We have trained

and tested AMFinder on ink-vinegar stained *N. benthamiana* roots. Ink-vinegar is an inexpensive, non-toxic fungal staining method compatible with various plant and mycobiont species (Vierheilig *et al.*, 1998). Thus, pre-trained CNNs generated from ink-stained roots ensure immediate workability of AMFinder for most endosymbiosis host systems without the need to generate manually-annotated training datasets. However, AMFinder can be trained using datasets obtained using other dyes and fluorophores for fungal staining (Vierheilig *et al.*, 2005) or for the annotation of other tissues colonised by fungi such as liverwort thalli (Ligrone *et al.*, 2007; Carella & Schornack, 2018; Kobae *et al.*, 2019). Computer-intensive computations required for *ab initio* training can be avoided by refining the existing pre-trained networks. Thus, this software is highly versatile and can be adapted or a wide range of fungal colonisation; it may also be of interest to researchers of pathogenic fungi.

AMFinder's precision is similar on all intraradical hyphal structures. However, specificity is best on vesicles, likely because AM fungal vesicles are fairly invariant, globular-shaped structures surrounded by a thick, multilayered wall (Jabaji-Hare *et al.*, 1990) that result in high contrast signals within the surrounding plant tissues. By contrast, the arbuscular shape is more diverse, with branching extent and cell volume occupancy increasing during the initial development stages (Toth & Miller, 1984) and a size that is ultimately defined by host cell size. Intraradical hyphae show different diameters, orientations, and staining intensities, and occasionally overlay other intraradical structures. Besides, the limited pixel information of a single tile may not always discriminate between intraradical and extraradical hyphae. An approach using information from a wider area of the original image, rather than treating each tile in isolation, may help address this issue. In particular, it would be interesting to apply deep learning image segmentation techniques (Ghosh *et al.*, 2019) to this problem, as researchers have often found success with this approach in other types of biological imaging. Another possible issue is that convolutional neural networks do not retain relative spatial information (Patrick *et al.*, 2019). Solutions to overcome this limitation include the combination of convolutional neural networks and multi-layer perceptrons (Haldekar *et al.*, 2017), and capsule networks (CapsNets) (Sabour *et al.*, 2017; Patrick *et al.*, 2019). Future work will explore the usefulness of such approaches to achieve even higher prediction accuracy.

Obtaining contrasted fungal structures within root tissues is pivotal for accurate AMFinder predictions. The first report of ink-vinegar staining of AM fungi suggests that black and blue inks allow for high-contrast images in at least four plant species (Vierheilig *et*

al., 1998). Background destaining in tap water with few vinegar droplets requires at least 20 min incubation and is only effective against excess ink (Vierheilig *et al.*, 1998). By contrast, ClearSee treatment works in seconds and allows for both destaining and clearing (Kurihara *et al.*, 2015). Such a feature is of particular interest for thick or pigmented roots, and soil samples. Also, ClearSee preserves fluorescence (Kurihara *et al.*, 2015) and is thus compatible with immunohistochemical fungal labelling techniques such as wheat germ agglutinin-fluorophore conjugates (Bonfante-Fasolo *et al.*, 1990).

AMFinder can improve the robustness and reproducibility of AM fungal quantification. In the gridline-intersect method, gridlines have been primarily used as guides for the systematic selection of observation points (Giovannetti & Mosse, 1980), and the distance between adjacent lines has been studied to estimate the total root length, but not to improve quantification accuracy (Newman, 1966; Marsh, 1971; Giovannetti & Mosse, 1980). As a result, a low number of root fragments is considered prejudicial to quantification accuracy (Giovannetti & Mosse, 1980). Also, the shape of the area surrounding the grid-root intersection used for visual scoring has not been formally described and may account for variations between experimenters. By contrast, AMFinder analyses well-defined tiles, and tile size can adjust to image resolution without impairing prediction accuracy.

Intraradical hyphal structures cannot be identified from flatbed scans due to limited resolution. However, machine learning-based algorithms have been recently developed to achieve data-driven image super-resolution (Park *et al.*, 2003; Wang *et al.*, 2019). By contrast with standard image interpolation techniques, super-resolution algorithms predict missing details by learning common patterns from training datasets. Future AMFinder development will investigate whether these algorithms can enable a detailed analysis of AM fungal hyphal structures from flatbed scans.

Conclusions

We have demonstrated that AMFinder adapts to different plant models, fungal staining methods and acquisition devices. Its design ensures user control over the annotation process and facilitates data visualisation in the context of the root images. As such, it will support better documentation and reproducibility of AM fungal colonisation analyses.

Acknowledgements

We thank Raymond Wightman (SLCU, Cambridge) for assistance with digital microscopy. We are grateful to Johannes Stuttmann (Martin-Luther-Universität Halle-Wittenberg, Germany) for critical comments on the manuscript. This work was funded by the Gatsby Charitable Foundation (GAT3395/GLD), by the European Research Council (ERC-2014-STG, H2020, and 637537), and by the Royal Society (UF110073 and UF160413). C.T. is supported by a Junior Research Fellowship at Gonville & Caius College, Cambridge, and also acknowledges support by the STFC consolidated grant ST/P000681/1.

Author Contribution

E.E. developed the software, conceived the experimental strategy, conducted experiments, acquired data, analysed data and wrote the manuscript. C.T. developed the software and analysed data. A.M. conceived the original idea, conducted experiments and acquired data. L.S., T.Y. and A.G. conducted experiments and acquired data. E.S. and C.Q. contributed material. S.S. acquired funding, conceived the experimental strategy, analysed data and wrote the manuscript. All authors have read and approved the manuscript.

Data Availability

AMFinder source code, pre-trained models and sample images are available on GitHub (https://github.com/SchornacklabSLCU/amfinder.git).

References

Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, *et al.* 2016. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *arXiv*.

Bally J, Jung H, Mortimer C, Naim F, Philips JG, Hellens R, Bombarely A, Goodin MM, Waterhouse PM. 2018. The rise and rise of *Nicotiana benthamiana*: A plant for all reasons. *Annual Review of Phytopathology* **56**: 405–426.

Bojarski M, Del Testa D, Dworakowski D, Firner B, Flepp B, Goyal P, Jackel LD, Monfort M, Muller U, Zhang J, *et al.* 2016. End to End Learning for Self-Driving Cars. *arXiv*.

Bonfante-Fasolo P, Faccio A, Perotto S, Schubert A. **1990**. Correlation between chitin distribution and cell wall morphology in the mycorrhizal fungus *Glomus versiforme*. *Mycological Research* **94**: 157–165.

Brundrett MC. **2002**. Coevolution of roots and mycorrhizas of land plants. *New Phytologist* **154**: 275–304.

Brundrett MC, Tedersoo L. **2018**. Evolutionary history of mycorrhizal symbioses and global host plant diversity. *New Phytologist* **220**: 1108–1115.

Carella P, Schornack S. **2018**. Manipulation of Bryophyte Hosts by Pathogenic and Symbiotic Microbes. *Plant and Cell Physiology* **59**: 651–660.

Choi J, Lee T, Cho J, Servante EK, Pucker B, Summers W, Bowden S, Rahimi M, An K, An G, *et al.* 2020. The negative regulator SMAX1 controls mycorrhizal symbiosis and strigolactone biosynthesis in rice. *Nature Communications* **11**: 1–13.

Chollet F, others. **2015**. Keras.

Demchenko K, Winzer T, Stougaard J, Parniske M, Pawlowski K. **2004**. Distinct roles of Lotus japonicus SYMRK and SYM15 in root colonization and arbuscule formation. *New Phytologist* **163**: 381–392.

Dhillon A, Verma GK. **2020**. Convolutional neural network: a review of models, methodologies and applications to object detection. *Progress in Artificial Intelligence* **9**: 85–112.

Dickson S. **2004**. The Arum-Paris continuum of mycorrhizal symbioses. *New Phytologist* **163**: 187–200.

Ferentinos KP. **2018**. Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture* **145**: 311–318.

Ghosh S, Das N, Das I, Maulik U. **2019**. Understanding deep learning techniques for image segmentation. *ACM Computing Surveys* **52**.

Giovannetti M, Mosse B. **1980**. An Evaluation of Techniques for Measuring Vesicular Arbuscular Mycorrhizal Infection in Roots. *New Phytologist* **84**: 489–500.

Girshick R, Donahue J, Darrell T, Malik J. **2014**. Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 580–587.

Gordon-Rodriguez E, Loaiza-Ganem G, Pleiss G, Cunningham JP. **2020**. Uses and abuses of the cross-entropy loss: Case studies in modern deep learning. *arXiv*.

Habte M, Byappanhalli BN. **1998**. Influence of pre-storage drying conditions and duration of storage on the effectiveness of root inoculum of Glomus aggregatum. *Journal of Plant Nutrition* **21**: 1375–1389.

Haldekar M, Ganesan A, Oates T. **2017**. Identifying spatial relations in images using convolutional neural networks. In: Proceedings of the International Joint Conference on Neural Networks. Institute of Electrical and Electronics Engineers Inc., 3593–3600.

Harrison MJ, Dewbre GR, Liu J. **2002**. A phosphate transporter from *Medicago truncatula* involved in the acquisition of phosphate released by arbuscular mycorrhizal fungi. *Plant Cell* **14**: 2413–2429.

Helber N, Requena N. **2008**. Expression of the fluorescence markers DsRed and GFP fused to a nuclear localization signal in the arbuscular mycorrhizal fungus *Glomus intraradices*. *New Phytologist* **177**: 537–548.

Hewitt EJ. **1966**. *Sand and water culture methods used in the study of plant nutrition*. 2nd Edition England, Commonwealth Agricultural Bureaux, Farnham Royal, Bucks, England.

Hohnjec N, Vieweg MF, Pühler A, Becker A, Küster H. **2005**. Overlaps in the transcriptional profiles of *Medicago truncatula* roots inoculated with two different *Glomus* fungi provide insights into the genetic program activated during arbuscular mycorrhiza. In: Plant Physiology. American Society of Plant Biologists, 1283–1301.

Hulse JD. **2018**. Review of Comprehensive Staining Techniques Used to Differentiate Arbuscular Mycorrhizal Fungi from Plant Root Tissues. *Acta Scientific Agriculture* **2**: 39–44.

Jabaji-Hare SH, Theiren J, Charest PM. **1990**. High resolution cytochemical study of the vesicular-arbuscular mycorrhizal association, Glomus clarum x Allium porrum. *New Phytologist* **114**: 481–496.

Javot H, Penmetsa RV, Terzaghi N, Cook DR, Harrison MJ. **2007**. A *Medicago truncatula* phosphate transporter indispensable for the arbuscular mycorrhizal symbiosis. *Proceedings of the National Academy of Sciences of the United States of America* **104**: 1720–1725.

Kobae Y, Ohtomo R, Morimoto S, Sato D, Nakagawa T, Oka N, Sato S. **2019**. Isolation of native arbuscular mycorrhizal fungi within young thalli of the liverwort *Marchantia paleacea*. *Plants* **8**.

Krizhevsky A, Sutskever I, Hinton GE. **2017**. ImageNet classification with deep convolutional neural networks. *Communications of the ACM* **60**: 84–90.

Kurihara D, Mizuta Y, Sato Y, Higashiyama T. **2015**. ClearSee: A rapid optical clearing reagent for whole-plant fluorescence imaging. *Development* **142**: 4168–4179.

LeCun Y, Bottou L, Bengio Y, Haffner P. **1998**. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**: 2278–2323.

Leroy X, Doligez D, Frisch A, Garrigue J, Rémy D, Vouillon J. **2020**. *The OCaml system, release 4.11 documentation and user's manual*. Technical report, Inria.

Ligrone R, Carafa A, Lumini E, Bianciotto V, Bonfante P, Duckett JG. **2007**.
Glomeromycotean associations in liverworts: A molecular, cellular, and taxonomic analysis.
American Journal of Botany **94**: 1756–1777.

Liu D, Shi L, Han C, Yu J, Li D, Zhang Y. **2012**. Validation of Reference Genes for Gene
Expression Studies in Virus-Infected Nicotiana benthamiana Using Quantitative Real-Time
PCR. *PLoS ONE* **7**: e46451.

Liu Y, Tang F, Zhou D, Meng Y, Dong W. **2017**. Flower classification via convolutional
neural network. In: Proceedings - 2016 IEEE International Conference on Functional-
Structural Plant Growth Modeling, Simulation, Visualization and Applications, FSPMA 2016.
Institute of Electrical and Electronics Engineers Inc., 110–116.

Luginbuehl LH, Oldroyd GED. **2017**. Understanding the Arbuscule at the Heart of
Endomycorrhizal Symbioses in Plants. *Current Biology* **27**: R952–R963.

**Maeda D, Ashida K, Iguchi K, Chechetka SA, Hijikata A, Okusako Y, Deguchi Y, Izui K,
Hata S**. **2006**. Knockdown of an arbuscular mycorrhiza-inducible phosphate transporter gene
of Lotus japonicus suppresses mutualistic symbiosis. *Plant and Cell Physiology* **47**: 807–817.

Marsh B. **1971**. Measurement of Length in Random Arrangements of Lines. *The Journal of
Applied Ecology* **8**: 265.

Matsugu M, Mori K, Mitari Y, Kaneda Y. **2003**. Subject independent facial expression
recognition with robust face detection using a convolutional neural network. In: Neural
Networks. Elsevier Ltd, 555–559.

Mohanty SP, Hughes DP, Salathé M. **2016**. Using deep learning for image-based plant
disease detection. *Frontiers in Plant Science* **7**: 1419.

Montero H, Choi J, Paszkowski U. **2019**. Arbuscular mycorrhizal phenotyping: the dos and
don'ts. *New Phytologist* **221**: 1182–1186.

Newman EI. **1966**. A Method of Estimating the Total Length of Root in a Sample. *The
Journal of Applied Ecology* **3**: 139.

**O'Mahony N, Campbell S, Carvalho A, Harapanahalli S, Hernandez GV, Krpalkova L,
Riordan D, Walsh J**. **2020**. Deep Learning vs. Traditional Computer Vision. In: Arai K,
Kapoor S, eds. Advances in Intelligent Systems and Computing. Advances in Intelligent
Systems and Computing. Cham: Springer International Publishing, 128–144.

Parádi I, Van Tuinen D, Morandi D, Ochatt S, Robert F, Jacas L, Dumas-Gaudot E.
2010. Transcription of two blue copper-binding protein isogenes is highly correlated with
arbuscular mycorrhizal development in *Medicago truncatula*. *Molecular Plant-Microbe
Interactions* **23**: 1175–1183.

Park SC, Park MK, Kang MG. **2003**. Super-resolution image reconstruction: A technical
overview. *IEEE Signal Processing Magazine* **20**: 21–36.

Patrick M, Adekoya F, Mighty A, Edward B. **2019**. Capsule Networks – A survey. *Journal of King Saud University - Computer and Information Sciences*.

van Rossum G, Drake FL. **2009**. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace.

Sabour S, Frosst N, Hinton GE. **2017**. Dynamic routing between capsules. In: Advances in Neural Information Processing Systems. Neural information processing systems foundation, 3857–3867.

Schüßler A, Schwarzott D, Walker C. **2001**. A new fungal phylum, the Glomeromycota: Phylogeny and evolution. *Mycological Research* **105**: 1413–1421.

Spatafora JW, Chang Y, Benny GL, Lazarus K, Smith ME, Berbee ML, Bonito G, Corradi N, Grigoriev I, Gryganskyi A, *et al.* 2016. A phylum-level phylogenetic classification of zygomycete fungi based on genome-scale data. *Mycologia* **108**: 1028–1046.

Srinath KR. **2017**. Python-The Fastest Growing Programming Language. *International Research Journal of Engineering and Technology*.

Sun XG, Tang M. **2012**. Comparison of four routinely used methods for assessing root colonization by arbuscular mycorrhizal fungi. *Botany* **90**: 1073–1083.

Szarvas M, Yoshizawa A, Yamamoto M, Ogata J. **2005**. Pedestrian detection with convolutional neural networks. In: IEEE Intelligent Vehicles Symposium, Proceedings. 224–229.

Takeda N, Sato S, Asamizu E, Tabata S, Parniske M. **2009**. Apoplastic plant subtilases support arbuscular mycorrhiza development in *Lotus japonicus*. *Plant Journal* **58**: 766–777.

Toth R, Miller RM. **1984**. Dynamics of Arbuscule Development and Degeneration in a *Zea mays* Mycorrhiza. *American Journal of Botany* **71**: 449.

Vandesompele J, De Preter K, Pattyn F, Poppe B, Van Roy N, De Paepe A, Speleman F. **2002**. Accurate normalization of real-time quantitative RT-PCR data by geometric averaging of multiple internal control genes. *Genome Biology* **3**: research0034.1.

Vierheilig H, Coughlan AP, Wyss U, Piché Y. **1998**. Ink and vinegar, a simple staining technique for arbuscular-mycorrhizal fungi. *Applied and Environmental Microbiology* **64**: 5004–5007.

Vierheilig H, Schweiger P, Brundrett M. **2005**. An overview of methods for the detection and observation of arbuscular mycorrhizal fungi in roots. *Physiologia Plantarum* **125**: 393–404.

Volpe V, Giovannetti M, Sun XG, Fiorilli V, Bonfante P. **2016**. The phosphate transporters LjPT4 and MtPT4 mediate early root responses to phosphate status in non mycorrhizal roots. *Plant Cell and Environment* **39**: 660–671.

Voulodimos A, Doulamis N, Doulamis A, Protopapadakis E. **2018**. Deep Learning for Computer Vision: A Brief Review. *Computational Intelligence and Neuroscience* **2018**.

Wang Z, Chen J, Hoi SCH. **2019**. Deep learning for image super-resolution: A survey. *ArXiv*.

Yamashita R, Nishio M, Do RKG, Togashi K. **2018**. Convolutional neural networks: an overview and application in radiology. *Insights into Imaging* **9**: 611–629.

Figure legends

Figure 1. AMFinder prediction pipeline enables semi-automated, user-supervised analysis of AM fungal colonisation *in silico*. **(a)** AMFinder uses a two-stage prediction pipeline for image annotation. First, input images are split into tiles and processed by `amf` neural network 1 (CNN1) to identify colonised root sections (prediction stage 1). If resolution allows, colonised areas can be further analysed to identify intraradical hyphal structures (prediction stage 2, CNN2). **(b)** Representative screenshot of an `amfbrowser` annotation session. **(1)** Buttons to switch the display between prediction stages. **(2)** Clickable buttons to define the annotations present in the active tile. **(3)** Magnified view of the active tile (red square) and eight surrounding tiles. **(4)** Annotation mosaic overview. **(5)** Layer toolbar to filter the display. Numbers indicate annotation counts for the whole image. **(6)** Prediction toolbar to load predictions, fix ambiguous cases and generate annotations. **(7)** Export functions.

Figure 2. AMFinder training pipeline. A set of manually annotated images (green) get pre-processed to generate test (black) and validation (grey) tile sets. N training cycles (epochs) follow, where CNN internal parameters are adjusted based on the test set. The validation set allows for independent monitoring of CNN performance by assessing model misfit (loss). An early stopping mechanism (blue) terminates training if loss value does not further decrease in 12 consecutive epochs. Best-performing parameters (blue) are saved for further use.

Figure 3. AMFinder consistently labels colonised root sections and intraradical hyphal structures in *Nicotiana benthamiana* roots. **(a)** CNN1 loss (left) and accuracy (right) plots obtained after training with 32 low- and high-resolution images of ink-stained *N. benthamiana* roots. **(b)** Representative images of CNN1 predictions on low- (top) and high-resolution (bottom) images. **(c)** Loss (left) and accuracy (right) plots obtained after CNN2 training with colonised root sections from 64 high-resolution images of ink-stained *N.*

benthamiana roots. **(d)** Representative images of CNN2 predictions made on colonised root sections of a high-resolution image. Asterisks (red) were added to indicate mispredictions. Scale bar is given in micrometers.

Figure 4. AMFinder can adapt to any image type. (a) CNN1 prediction error rate on high-resolution images after pre-processing to either alter edges (sharpen, blur), increase noise (gaussian, plasma), reduce brightness, or change colours (hue, grayscale, invert). Statistical significance was assessed with ANOVA ($p < 0.05$). Letters indicate statistical groupings. Average rates are shown in italics. **(b)** Representative set of tiles used for *de novo* CNN1 training on images bearing the same modifications as in **a**. **(c)** Loss (left) and accuracy (right) plots obtained after CNN1 training (N = 32 images). **(d)** Representative examples of CNN1 predictions obtained with pre-trained and *de novo*-trained models. Asterisks (red) indicate mispredictions.

Figure 5. AMFinder accurately identifies AM fungal colonisation and intraradical structures in roots of various host model species. (a-c) Roots from *Lotus japonicus* cv. Gifu **(a)**, *Medicago truncatula* ecotype R108 **(b)**, and *Oryza sativa* cv. Nipponbare **(c)** were stained with ink and cleared with ClearSee prior to analysis with AMFinder. Dashed square inserts indicate areas used for AMFinder predictions. Scale bar is given in micrometers.

Figure 6. Computer vision enables quantification of AM fungal colonization on whole root systems. *Nicotiana benthamiana* plants were inoculated with *Rhizophagus irregularis*. Colonization levels were determined 4 and 6 weeks post-inoculation (wpi). **(a-b)** Quantification of transcripts corresponding to the *N. benthamiana* mycorrhizal marker gene *Blue Copper Protein 1* (*BCP1*) **(a)** and the *R. irregularis* elongation factor 1α **(b)**. Data are expressed relative to both *N. benthamiana L23* and *F-box* reference genes. **(c-d)** Quantification of AM fungal colonization on ink-stained root pictures using the gridline intersect method **(c)**, or AMFinder **(d)**. **(e)** Quantification of AM fungal colonisation on whole root systems using AMFinder, and representative images of computer-generated maps featuring colonized (M+, magenta) and non-colonized (M−, grey) root areas at 4 and 6 wpi. Dots correspond to biological replicates. Bars represent standard error. Statistical significance was assessed using Student's *t*-test (*: $p < 0.05$; **: $p < 0.01$).

Supporting Information

Figure S1. Schematic representation of AMFinder ConvNet architecture.

Figure S2. ClearSee enhances the contrast of ink-stained roots.

Figure S3. Optimisation of flatbed scanner resolution for imaging of ink-stained roots.

Figure S4. AMFinder enables a detailed analysis of trained model performance.

Figure S5. AMFinder prediction error rates.

Figure S6. Tile size does not affect AMFinder prediction accuracy.

Figure S7. Low-resolution image dataset used to assess CNN1 prediction accuracy.

Figure S8. High-resolution image dataset used to assess CNN1 prediction accuracy.

Figure S9. High-resolution image dataset used to assess CNN2 prediction accuracy.

Figure S10. AMFinder can label AM fungal colonisation on a wide range of input images.

Table S1. AMFinder troubleshooting guide.