

3D U-Net for Segmentation of Plant Root MRI Images in Super-Resolution

Yi Zhao¹, Nils Wandel¹, Magdalena Landl², Andrea Schnepf² and Sven Behnke^{1*}

1 – University of Bonn, Computer Science Institute VI, 53115 Bonn, Germany

2 – FZ Jülich GmbH, Institute of Bio- and Geosciences 3, 52425 Jülich, Germany

Abstract. Magnetic resonance imaging (MRI) enables plant scientists to non-invasively study root system development and root-soil interaction. Challenging recording conditions, such as low resolution and a high level of noise hamper the performance of traditional root extraction algorithms, though. We propose to increase signal-to-noise ratio and resolution by segmenting the scanned volumes into root and soil in super-resolution using a 3D U-Net. Tests on real data show that the trained network is capable to detect most roots successfully and even finds roots that were missed by human annotators. Our experiments show that the segmentation performance can be further improved with modifications of the loss function.

1 Introduction

In order to gain a better understanding of plant growth and its response to environmental influences such as droughts, sinking groundwater levels or climate change, plant scientists investigate root-soil interaction processes. Magnetic resonance imaging (MRI) non-invasively measures both roots and soil [1]. Low resolution of MRI scans—compared to the diameter of thin roots—and high noise (Figure 1) due to ferromagnetic particles or low contrast between root and soil water signal hamper the application of traditional root extraction algorithms [2]. For these reasons, root reconstruction is often done manually. To enable automatic reconstruction, the signal-to-noise ratio (SNR) and resolution of these MRI scans need to be enhanced.

In this work, we propose using a 3D U-Net [3] to segment scans of pots containing plant roots into root and soil in super-resolution. As training data is scarce, we create synthetic MRIs by rendering roots in 3D and combine them with real MRI scans of pure soil. In this way, we improve segmentation performance in soil regions, compared to models that were trained on synthetic data only [4, 5]. Furthermore, we investigated the influence of different loss modifications. By doing this, the segmentation performance gets improved, especially for thin roots which are difficult to detect intactly.

2 Related Work

Recently, CNNs have achieved remarkable performance in 3D image segmentation. Havaei et al. [6] use 2D CNNs to segment horizontal slices of 3D images

*This research was supported by grants BE 2556/15 and SCHN 1361/3-1 of German Research Foundation (DFG).




Fig. 1: 3D visualizations of noisy root images and original MRI image slices. (a)(c) each depicts the 3D visualization of a 3D MRI image, and (b)(d) each shows a horizontal slice of the 3D image on its left.

and combine them into 3D outputs. However, 2D CNNs can only make limited use of spatial information along the vertical dimension. 3D CNNs can overcome this issue by learning 3D features from input images [3].

Since their introduction in 2016, 3D U-Net [3] has become a commonly used architecture for segmentation tasks of 3D data. The structure of 3D U-Net consists of a downsampling encoder and an upsampling decoder. The encoder extracts increasingly abstract features from an increasingly broader context, while the decoder restores the resolution from the downsampled feature maps. Moreover, to share the high-resolution features extracted in the encoder with the decoder, shortcut connections are established between them.

For image super-resolution, deep neural networks show state-of-the-art performance. To achieve a higher resolution in the network output, the input image needs to be upsampled. This can be done either by upsampling the input using interpolation [7], or within the network using methods like transposed convolution [8]. The latter results in better performance because the interpolation is learned directly from the data.

Root extraction algorithms were developed to automatically extract structural models of roots from 3D MRI images [2]. However, their performances degrade significantly in the presence of low SNR or low resolution [2]. Uzman et al. [4] increased the SNR and resolution of MRI images using a 2D RefineNet. Due to the limited number of real images, synthetic MRI data was used during training. While the RefineNet is able to detect most roots, it also showed a non-negligible amount of false positives in some cases. Horn et al. [5] tried to complete the same task with a 3D CNN. However, because of memory constraints of the graphics processing unit (GPU), the network structure was relatively shallow. This resulted in fewer false positives, but at the same time significantly more false negatives for thin roots, leading to large gaps between detected root fragments.

3 Segmentation in Super-Resolution

We segment the MRI scans in double resolution, i.e. the task is to map an input image $I \in \mathbb{R}^{x \times y \times z}$ to a binary segmentation output $S \in \mathbb{B}^{2x \times 2y \times 2z}$. To compute

the mapping, a 3D U-Net architecture was chosen. It is relatively simple but still incorporates a large 3D context in its decision. During training, the network was fed randomly sampled $60 \times 60 \times 60$ crops in order to stay within the limitations of GPU memory.

3.1 3D U-Net

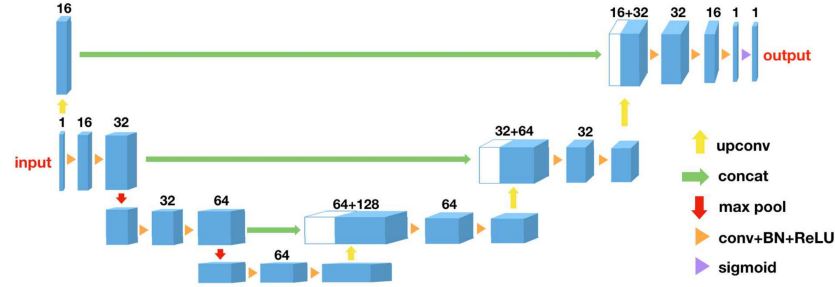


Fig. 2: 3D Superresolution U-Net. The output has twice the input resolution.

Like the original 3D U-Net, our network consists of an encoder and a decoder (Figure 2). The encoder contains three convolutional modules, each with two convolutional layers and an increasing number of channels. To avoid introducing misleading information through padding, only valid convolutions are used. Adjacent convolutional modules are connected by a maxpooling layer for downsampling. The decoder part consists of three convolutional modules which are connected by $2 \times$ transposed convolutional layers for upsampling. Each upsampled tensor is concatenated with an intermediate output from the encoder part. In addition to entering the encoder-decoder path, the input is also directly upsampled with $2 \times$ transposed convolution and concatenated with the super-resolution intermediate output of the decoder. At the segmentation output, the channel number is reduced to one and a sigmoid function returns root probability estimates between 0 and 1.

3.2 Dataset

The dataset used in this work builds on the synthetic dataset generated by Uzman et al. [4]. The images in it are generated by combining augmentations of four reconstructed root structures with synthetic soil images which simulate real soil features.

Combining generated roots with real soil MRI To increase the diversity of root structures, we generated 30 random roots of various shapes and root-diameters. Furthermore, we used MRI scans of pure soil to capture a larger range of 3D features of soil. For training, random crops of the generated roots were combined

with crops of pure soil and augmented by varying parameters such as the contrast between root and soil.

For validation, random crops of a different subset of the dataset were used. The final evaluation of the network’s performance was done on five real MRI images [1] with human annotations.

3.3 Loss

The loss function for training is binary cross-entropy loss averaged among all image voxels. On top of that, we investigated two loss modifications. The first one is applying a higher weight on the root voxels than the soil voxels, forcing the network to focus more on the correct prediction of roots. The reason for trying this is the imbalance of the soil voxel quantity and the root voxel quantity. The second modification is using a don’t-care flag to label the root-soil border area, and ignore the voxels in it when calculating the loss. This is done to check the effect of making the learning task easier, because the this area is hard to segment precisely and also not so important.

3.4 Evaluation

Since the MRI images contain far more soil voxels than root voxels, the F1-Score is used for evaluation because it is robust against class imbalance. The human annotations are slightly misaligned to the real MRI data, which is a common problem that also occurs in 2D root images [9]. Because of that, the Distance Tolerant F1-Score as introduced by Uzman et al. [4] was used. This way, misalignments between the output and the ground truth will be tolerated if the distance between them is not larger than the value of the distance tolerance.

4 Results

Figure 3a shows the Distance Tolerant F1-Scores of the 3D U-Net model trained on the improved dataset, evaluated on the test set with five real MRI scans. For larger tolerances, the score reaches higher values. Figure 3b shows the relative improvement of the score, compared to training with the original data set. The example segmentation results shown in Figure 3c, 3d demonstrate that the number of false positives decreases when using the improved dataset. The model trained with the improved dataset reaches an F1-Score of 0.964 on the validation set, and its segmentation results on the test data detect most root branches correctly.

Figure 4 compares our 3D U-Net qualitatively to the segmentation results of RefineNet [4]. There are fewer false positives in the results of our model, reducing the risk of those false positives being incorrectly extracted as roots by root extraction algorithms. However, there are also more false negatives in the 3D U-Net results, which appear as disconnections in the roots. Closing large disconnections can be difficult for the root extraction algorithms.




Fig. 3: Comparison between the models trained with the original dataset and the improved dataset. (a) shows the Distance Tolerant F1-Scores of the model trained with the improved dataset. (b) shows the percentage of improvement of the Distance Tolerant F1-Scores when trained with the improved dataset compared to the original dataset. Error bar indicates the 95% confidence interval. (c)(d) are example segmentation results of the models trained with the original and improved dataset, respectively. Green, red, blue represent true positives (TPs), false positives (FPs), and false negatives (FNs), respectively.




Fig. 4: Segmentation outputs compared between RefineNet and 3D U-Net on real MRI images. Green, red, blue represent TPs, FPs, and FNs, respectively. (a)(c) are results from RefineNet and (b)(d) are results from 3D U-Net.

Furthermore, we investigated two types of loss modifications to deal with the disconnection issue. Here, all compared models are trained with the improved dataset. The first modification is using a higher root weight of 10, which results in improved Distance-Tolerant F1-Scores (Figure 5a) as well as fewer disconnections in the segmented thin roots (Figure 5c). Some root-like false positive predictions are likely roots missed in the human annotations. The second loss modification is to ignore the root-soil border area using don't-care flag. This also leads to improved F1-Scores (Figure 5a) but the roots in the segmentation result become visibly thicker which is slightly unrealistic (Figure 5d). The improvement of the F1-Scores is probably largely due to the higher recall because of the thicker root predictions.

5 Conclusion

We investigated semantic segmentation in super-resolution of plant root MRI images using a 3D U-Net model. When evaluated on the real MRI images, the




Fig. 5: Effect of loss modifications. (a) shows the percentage of improvement of the Distance Tolerant F1-Scores of the 2 types of loss modification. Error bar indicates the 95% confidence interval. (b)(c)(d) shows the example segmentation results of the models trained with the original loss, root weight of 10, and don't-care flag, respectively. Green, red, blue represent TPs, FPs, and FNs, respectively.

model can detect most root branches correctly, despite some minor disconnections in thin roots. This disconnection problem can be mitigated by applying higher weights on root voxels or ignoring the voxels near the root-soil border when calculating the training loss.

References

- [1] L. Stingaciu, H. Schulz, A. Pohlmeier, S. Behnke, H. Zilken, M. Javaux, and H. Vereecken. In situ root system architecture extraction from magnetic resonance imaging for water uptake modeling. *Vadose Zone Journal*, 12(1), 2013.
- [2] H. Schulz, J. A. Postma, D. van Dusschoten, H. Scharr, and S. Behnke. Plant root system analysis from mri images. In *Computer Vision, Imaging and Computer Graphics. Theory and Application*, pages 411–425. Springer, 2013.
- [3] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432. Springer, 2016.
- [4] A. O. Uzman, J. Horn, and S. Behnke. Learning super-resolution 3D segmentation of plant root mri images from few examples. In *ESANN*, pages 553–558, 2019.
- [5] J. Horn. *Superresolution 3D Image Segmentation for Plant Root MRI*. Bachelor’s thesis, University of Bonn, Germany: Computer Science VI, 2018.
- [6] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle. Brain tumor segmentation with deep neural networks. *Medical image analysis*, 35:18–31, 2017.
- [7] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, pages 184–199. Springer, 2014.
- [8] C. Dong, C. C. Loy, and X. Tang. Accelerating the super-resolution convolutional neural network. In *ECCV*, pages 391–407. Springer, 2016.
- [9] A. G. Smith, J. Petersen, R. Selvan, and C. R. Rasmussen. Segmentation of roots in soil with u-net. *Plant Methods*, 16(1):1–15, 2020.