

# Joint Stem Detection and Crop-Weed Classification for Plant-specific Treatment in Precision Farming

Philipp Lottes

Jens Behley

Nived Chebrolu

Andres Milioto

Cyrill Stachniss

**Abstract**—Applying agrochemicals is the default procedure for conventional weed control in crop production, but has negative impacts on the environment. Robots have the potential to treat every plant in the field individually and thus can reduce the required use of such chemicals. To achieve that, robots need the ability to identify crops and weeds in the field and must additionally select effective treatments. While certain types of weed can be treated mechanically, other types need to be treated by (selective) spraying. In this paper, we present an approach that provides the necessary information for effective plant-specific treatment. It outputs the stem location for weeds, which allows for mechanical treatments, and the covered area of the weed for selective spraying. Our approach uses an end-to-end trainable fully convolutional network that simultaneously estimates stem positions as well as the covered area of crops and weeds. It jointly learns the class-wise stem detection and the pixel-wise semantic segmentation. Experimental evaluations on different real-world datasets show that our approach is able to reliably solve this problem. Compared to state-of-the-art approaches, our approach not only substantially improves the stem detection accuracy, i.e., distinguishing crop and weed stems, but also provides an improvement in the semantic segmentation performance.

## I. INTRODUCTION

Agrochemicals such as pesticides, herbicides, and fertilizer are currently needed in conventional agriculture for effective weed control and attaining high yields. Agrochemicals, however, can have a negative impact on the environment and consequently affect human health. Thus, sustainable crop production should reduce the amount of applied chemicals. Today, weed control avoiding agrochemicals, however, is often a manual and labor-intensive task.

Robots for precision farming offer a great potential to address this challenge through a targeted treatment on a per-plant level. Agricultural robots equipped with different actuators for weed control such as selective sprayers, mechanical tools, or even lasers, can execute the treatments only where it is actually needed and also can select the most effective treatment for the targeted plant or weed. For example, mechanical and laser-based treatments are most effective if applied to the stem location of a plant. In contrast, grass-like weeds are most effectively treated by applying agrochemicals on their entire leaf area.

To realize a selective and plant-dependent treatment, farming robots need an effective plant classification system. Such a system needs to reliably identify both, the stem location of dicot weeds (weeds whose seeds having two embryonic

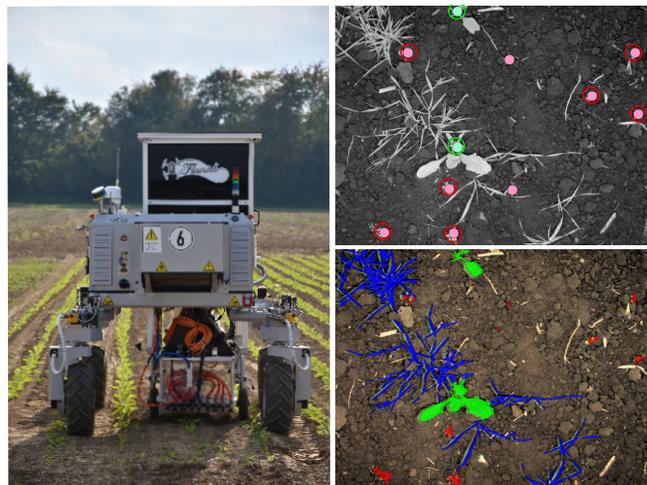


Fig. 1. BoniRob platform operating in a sugar beet field while performing weed control through selective spraying and mechanical weed stamping. Our approach analyzes camera images and provides two classification outputs: the stem positions of the crop plants (green) and dicot weeds (red) as well as a pixel-wise semantic segmentation of the input into the classes crop (green), dicot weed (red), grass (blue) and soil (no color).

leaves) and also the extent of grass weeds given by its leaf area. In this paper, we thus address exactly this problem so that a robot can perform targeted, plant-specific treatments.

The main contribution of this paper is an end-to-end trainable pipeline for joint pixel-wise plant segmentation and plant stem detection enabling plant-specific treatment, for example fertilizing a crop or destroying a weed. We employ a fully convolutional neural network (FCN) architecture sharing the encoded representation of the image content for the specific tasks, i.e., the semantic segmentation of the crops, dicot weeds, and grasses, as well as the stem detection of the individual crops and dicot weeds for mechanical removal. More specifically, we jointly estimate the pixel-wise segmentation into the classes (1) crop, (2) dicot weed, (3) grass weed, and (4) background, i.e., mostly soil, and estimate the stem locations of crops and dicot weeds at the same time.

In sum, we make the following two claims: Our approach is able to (i) determine the stem positions of crop and weed stems, and (ii) accurately separate grass weed from dicot weed for the purpose of a specific treatment in mind. Furthermore, we show that our approach has a superior performance in comparison to other state-of-the-art approaches, such as [5], [15], [19]. All claims are experimentally validated on real-world data. Moreover, we plan to publish our code and the datasets used in this paper.

All authors are with the University of Bonn, Germany.

This work has partly been supported by the EC under the grant number H2020-ICT-644227-Flourish.

## II. RELATED WORK

In recent years, significant progress has been made towards vision based crop-weed classification systems, both using handcrafted features [6], [15], [14] and end-to-end methods based on convolutional neural networks (CNN) [17], [21], [19]. However, none of these methods estimate the stem locations or other information which can be directly used for targeted intervention. With our work, we aim to bridge this gap by developing a system which integrates both the task of plant classification and stem detection in an end-to-end manner with the goal of targeted treatment in mind.

Other approaches have been developed to classify individual plants and identify their stem locations. Most of these approaches are developed based on manually designed heuristics with specific use cases in mind. Kiani and Jafari [10] use hand-crafted shape features selected on the basis of a discriminant analysis to differentiate corn plants from weeds. They identify the stem position of the plant as the centroid of the detected vegetation. This leads to sub-optimal results particularly when the plant shapes are not symmetric or multiple plants are overlapping. Midtiby *et al.* [18] present an approach for sugar beet by detecting individual leaves and use the contours of the leaves for finding the stem locations. This approach fails to locate the stems in the presence of occluded leaves or overlapping plants.

Moving towards a machine learning based approach, Haug *et al.* [5] propose a system to detect plant stems using keypoint-based random forests. They employ a sliding window based classifier to predict stem regions by using several hand-crafted geometric and statistical features. Their evaluation shows that the approach often misses several stems for overlapping plants or generates false positives for leaf areas which locally appear to be stem regions. Kraemer *et al.* [11] aim at addressing this issue by increasing the field of view of the classifier using a fully convolutional networks (FCN) [13].

The goal of their work is to identify plant stems over a temporal period allowing them to use the stem locations as landmarks for localization.

Our work overcomes many of the limitations by taking a holistic approach by jointly detecting stems and estimating a pixel-wise segmentation of the plants based on FCNs. Moreover, we explicitly distinguish crop and dicot weed stems, since it enables plant-specific treatment, for example fertilizing a crop or destroying a weed mechanically.

## III. JOINT STEM DETECTION AND CROP-WEED SEMANTIC SEGMENTATION

The main objective of our plant classification system is to simultaneously provide a semantic segmentation of the visual input into the classes crop, dicot weed, grass weed, and soil as well as the stems positions for dicot weeds and crops. The stem positions are a prerequisite in selective, high precision treatments, e.g., by mechanical stamping or by laser-based weeding. The provided pixel-wise label mask provides the area for more granulated approaches such as selective spraying. We propose an approach for the joint

processing of the plant classification and the stem detection task based on FCN. Our network architecture shares the encoded features for classifying the stem regions as well as for the pixel-wise semantic segmentation using two, task-specific decoder networks.

The input to our network is either given by the raw RGB or by RGB plus a near infra-red (NIR) channel. The output of the proposed network consists of two different label masks representing a probability distribution over the respective class labels. The first output is the plant mask reflecting the pixel-wise semantic segmentation of the vegetation in image space, whereas the second output is the stem mask segmenting regions within the image, which correspond to crop stems and weed stems. Finally, we extract pixel-accurate stem positions from the stem mask.

### A. General Architectural Concept

Fig. 2 depicts the proposed architecture of our joint plant and stem detection approach. The main processing steps of our approach are the (i) preprocessing (red), the (ii) encoder (orange), the (iii) plant decoder (blue), the (iv) stem decoder (green), and (v) the stem extraction (brown).

As common practice in semantic segmentation [1], [9], the encoder uses convolutional layers and downsampling operations to extract a compressed, but highly informative, representation of the image. We use this encoded image representation as input to our task-specific decoders, i.e., a plant decoder and a stem decoder. The plant decoder produces the plant features, for the segmentation of the soil, crops, dicot weeds and grass weeds, whereas the stem decoder produces the stem features for the segmentation of the crop-weed stem regions. Both decoders upsample the shared code volume back to the original input resolution to allow for a pixel-wise segmentation. Finally, we further analyze the stem mask containing the segmentation of potential stem region and extract the stem locations of crops and dicot weeds from it. In the following sections, we describe the different parts of the proposed pipeline in more detail.

### B. Preprocessing

To improve the generalization capabilities and also aid the convergence of training, we first preprocess each channel of the given input images, i.e., red, green, blue, and near infra-red, as follows. First, we apply a Gaussian smoothing using a  $[5 \times 5]$  kernel with weights from a Gaussian with mean  $\mu = 0$  and  $\sigma^2 = 1$ . Then, we standardize each channel by subtracting the mean and divide it by its variance. Finally, we contrast stretch the input values to the range  $[-0.5, 0.5]$ .

### C. FC-DenseNet

The main building block of our FCN architecture is inspired by the so-called Fully Convolutional DenseNet (FC-DenseNet) [9], which combines the recently proposed densely connected CNNs organized as dense blocks [8] with fully convolutional networks (FCN) [13].

A dense block is given by a stack of  $L$  subsequent convolutional layers operating on feature maps with the same

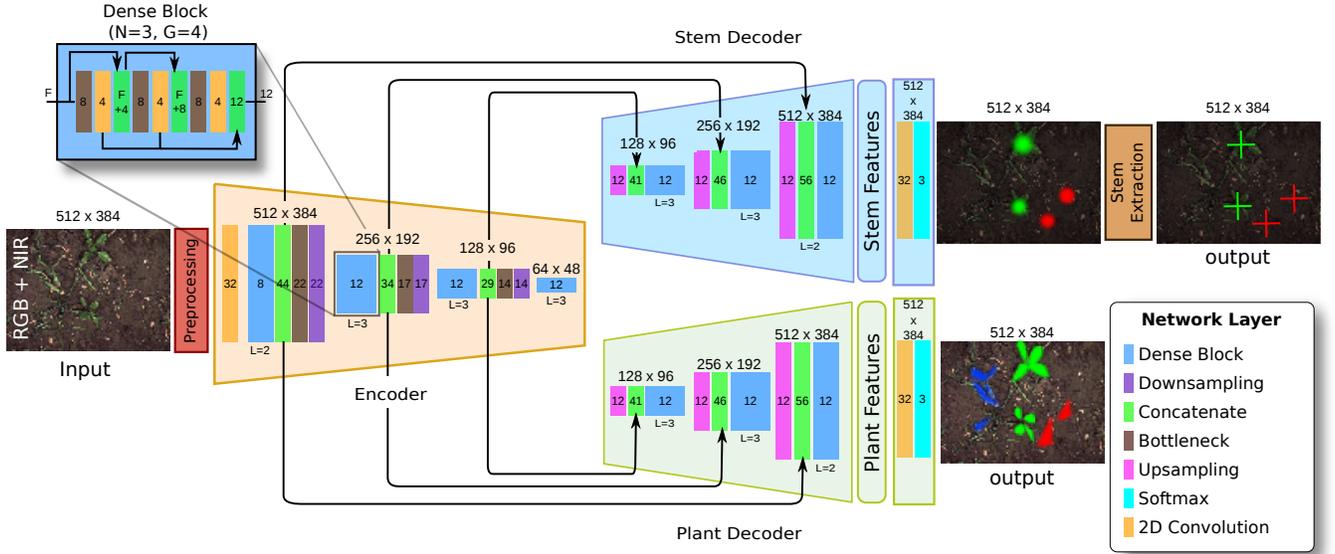


Fig. 2. Architecture of our approach as described in Sec. III-A. We first encode the input image using the encoder. Then pass the encoded feature volume to the task-specific decoder, the stem decoder and the plant decoder. Thus, we obtain two outputs, the plant mask for the semantic segmentation of the plants and the stem mask for the segmentation crop-weed stem regions. Finally, we extract the stem positions from the stem mask in the stem extraction. Note that we denote the size of the feature map above each block of layers. Inside the layers, we show the number of output features maps.

spatial resolution. Here, we define a convolutional layer as composition of a  $[3 \times 3]$  convolution, leaky rectified linear units [16], batch normalization [23], and dropout [22]. Each convolutional layer gets as input a concatenation of the result of the previous layers. For computational efficiency, we use a bottleneck layer before a convolutional layer implemented by a  $[1 \times 1]$  convolution [12]. A dense block produces  $N \cdot G$  feature maps, where  $G$  is the growth rate [8].

We use subsequent dense blocks inside the encoder and concatenate the input to a dense block with its output, which subsequently is compressed again by a bottleneck layer to reduce the growth of feature maps within the encoder. Each dense block is followed by a downsampling operation using strided convolutions employing a convolutional layer with a  $[5 \times 5]$  kernel and a stride of 2.

From the encoded and compressed information, we generate two separate feature volumes specialized for pixel-wise plant classification and stem detection. Thus, we have two decoders, which perform an upsampling using a strided transpose convolution [3] with  $[2 \times 2]$  kernel and a stride of 2. Both decoders also use dense blocks as their main building block and follow the same architectural design to produce the plant features and stem features. Moreover, both task specific decoders use feature maps produced by the encoder through skip connections. We concatenate the corresponding feature maps sharing the same spatial resolution from the encoder before we again use dense blocks for feature computation. Skip connections from the encoder to the decoders facilitates the recovery of spatial information [1]. Finally, we transform the feature maps produced by the stem decoder and the plant decoder into the pixel-wise probability distribution over their respective class labels by a  $[1 \times 1]$  convolution followed by a softmax layer.

For learning, we use a multi-task loss  $\mathcal{L}$  combining the loss for the plant segmentation  $\mathcal{L}_{\text{plant}}$  and for the stem region segmentation  $\mathcal{L}_{\text{stem}}$ , i.e.,

$$\mathcal{L} = (1 - \alpha) \cdot \mathcal{L}_{\text{stem}} + \alpha \cdot \mathcal{L}_{\text{plant}}, \quad (1)$$

where we use  $\alpha = 0.5$ .  $\mathcal{L}_{\text{plant}}$  is a weighted cross entropy loss, where we penalize errors regarding the crops, dicot weeds and grasses by a factor of 10.  $\mathcal{L}_{\text{stem}}$  is a loss based on an approximation of the intersection over union (IoU) metric as it is more stable with imbalanced class labels [20], which is the case in our problem with under-represented stems as compared to the amount of soil. The multi-task loss also enables to share information for learning the encoder, which can use the loss information from both decoders in the backward pass of the backpropagation. For training, we encode the stem locations as blobs with a diameter of 10 mm in object-space.

#### D. Stem Extraction

Given the pixel-wise stem mask prediction from the neural network, i.e.,  $P(y|\mathbf{x})$  with  $y \in \{\text{soil}, \text{crop}, \text{dicot weed}\}$  for each pixel  $\mathbf{x}$ , we want to extract a stem location for the crops and the dicot weeds. To this end, we first determine for each pixel the class with highest label probability, i.e.,  $y^* = \text{argmax}_y P(y|\mathbf{x})$ . Next, we determine the connected components  $\mathcal{X}_j^c$  for each class  $c$  and compute the weighted mean  $\bar{\mathbf{x}}_j^c$  of the pixel locations by

$$\bar{\mathbf{x}}_j^c = \frac{\sum_{\mathbf{x} \in \mathcal{X}_j^c} P(y = c|\mathbf{x}) \cdot \mathbf{x}}{\sum_{\mathbf{x} \in \mathcal{X}_j^c} P(y = c|\mathbf{x})}. \quad (2)$$

The weighted means  $\bar{\mathbf{x}}_j^c$  for class  $c$  are then the stem detections reported by our approach.



Fig. 3. Examples images. The left RGB+NIR image shows data from the BoniRob dataset. The right RGB image belongs to the UAV dataset. Note, the difference in the lighting, but also soil conditions.

#### IV. EXPERIMENTAL EVALUATION

Our experiments are designed to support our two claims: (i) Our approach is able to detect accurately the stem locations of crops and dicot weeds, and simultaneously (ii) is able to accurately segment the images into the classes crops, dicot weed, grasses and soil.

The experiments are conducted on data from two different sugar beet fields located near Bonn. Both datasets contain sugar beet plants, which are our crop, different dicot weeds and grass weeds. The first dataset, called BoniRob dataset, consists of 382 RGB+NIR images recorded under artificial lighting conditions with the BOSCH DeepField Robotics BoniRob and is a subset of a publicly available dataset [2]. It contains 488 crop stems and 1,607 weed stems. The second dataset, called UAV dataset, contains 400 RGB images recorded with an unmanned aerial vehicle (UAV), DJI Inspire II, under natural lighting conditions. In sum, it contains 808 crop stems and 1,451 dicot weed stems. Both datasets represent challenging conditions for our approach as they contain several different dicot weed types of different sizes and multiple grass types and have a substantial amount of inter-plant overlap. We annotated all images of both datasets with pixel-wise class labels and stem locations. Fig. 3 shows examples of each dataset.

We use the mean average precision (mAP) over the per-class average precisions (AP) [4] as metric for our evaluation. The mAP represents the area under the interpolated precision-recall curve. As noted by Everingham *et al.* [4], a method must have a high precision at all levels of recall to achieve a high score with this metric. For the stem detection task, a predicted stem is considered to be a positive detection if its Euclidean distance to the nearest unassigned ground truth stem is below a threshold  $\theta = 10$  mm corresponding to the size of the mechanical stamping tool of the BoniRob. Furthermore, we compute the mean average distance (MAD) for all true positives to show the spatial precision of our approach. For the segmentation task, we evaluate the performance in a pixel-wise manner also using the mAP.

For comparison, we also evaluate the performance with respect to other methods. For the stem detection, we refer to the *baseline-stem* approach, where we apply our proposed architecture as a single encoder-decoder FCN. Analogously, we refer to the *baseline-seg* approach, when using our architecture only for the semantic segmentation task. We furthermore compare the performance with our implementations of state-of-the-art approaches for stem detection and crop-weed classification. For stem detection, we re-implemented

TABLE I  
STEM DETECTION PERFORMANCE ON BONIROB DATASET

Approach	mAP	Crop		Dicot	
		AP	MAD	AP	MAD
ours	79.2	93.5	2.5	65.0	1.8
RF+F [15]	51.1	85.1	1.3	17.1	1.7
RF [6]	48.8	67.5	2.2	28.9	2.1
baseline-stem	73.8	88.9	2.7	58.8	1.9

TABLE II  
STEM DETECTION PERFORMANCE ON UAV DATASET

Approach	mAP	Crop		Dicot	
		AP	MAD	AP	MAD
ours	75.3	78.8	3.2	71.5	2.1
RF+F [15]	51.8	60.6	3.1	70.2	1.8
RF [5]	49.1	59.5	2.7	38.7	1.7
baseline-stem	75.7	78.8	3.2	72.6	2.1

the approach of Haug *et al.* [6] using a random forest and the described features, which we denote by RF. Next, we use the same methodology as Haug *et al.*, but with the visual and shape features proposed by Lottes *et al.* [15] and denote this method by RF+F. For the segmentation task, we use the approach by Lottes *et al.* [15] and a state-of-the-art FCN-based approach [19] for crop-weed detection as a reference.

##### A. Parameters

In all our experiments, we learned our network from scratch using only the training data of the respective dataset, which is 75% of all available images. For validation, we use 5% and a 20%-held out portion for evaluating the test performance, which we report within this section. We downsample the images to a resolution of  $W = 512$  and  $H = 384$  pixels, which yields a ground resolution of approx.  $1 \frac{\text{mm}^2}{\text{pixel}}$ . In all experiments, we use a growth rate  $G = 4$  and dropout probability  $p = \frac{1}{3}$ . Following common best practices for training deep networks, we initialize the weights according to He *et al.* [7], use ADAM for optimization, and a mini-batch size of  $B = 4$ . The initial learning rate is set to 0.01 and divided by 10 after 50, 250 and 1,000 epochs. We stop training after 2,000 epochs. We implemented our approach using Tensorflow. Our approach analyzes an image at inference time with at least 5 Hz.

##### B. Stem Detection Performance

The first experiment is designed to support the claim that our approach detects the stem position for crops and dicot weeds. We compare the performance with the approaches based on random forests (RF and RF+F) as well as with our baseline approach. Tab. I and Tab. II shows the performance for BoniRob dataset and UAV dataset.

In both datasets, we see that our approach outperforms the competing approaches using random forests classifiers in terms of mAP. The difference in mAP is mainly due to the improved performance with respect to the dicot weed stem detection. In both datasets, we observe a gain of around 25 – 30% in the mAP. We also gain around

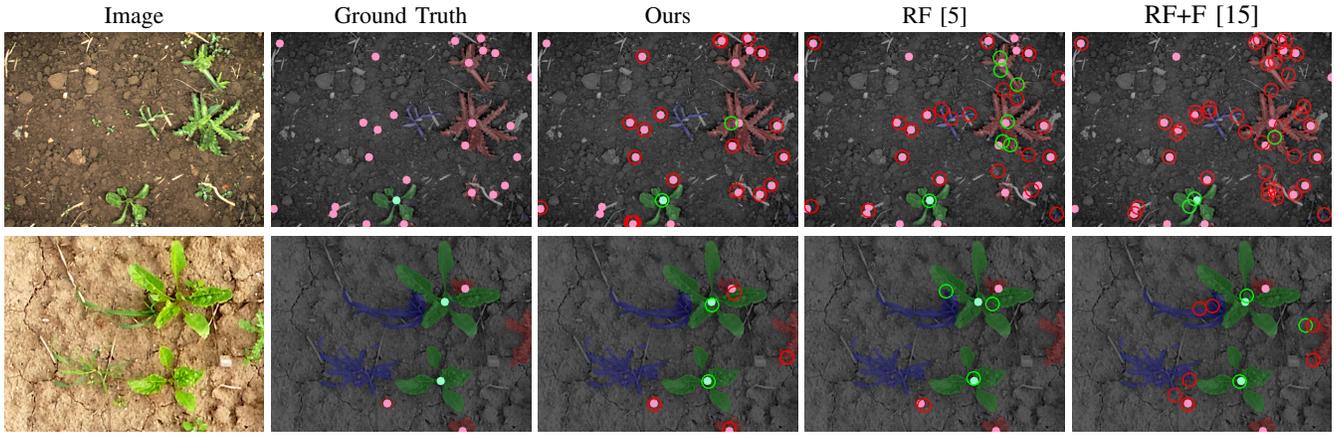


Fig. 4. Qualitative results of the stem detection for BoniRob dataset (top row) and UAV dataset (bottom row). From left to right: RGB input image, ground truth information for the pixel-wise class labels provided for better view, predicted stems and corresponding ground truth by our approach, RF [5] and RF+F [15]. Red refers to the class dicot weed and green refers to crop. The solid circles in the same, but brighter color, refer to the respective ground truth information.

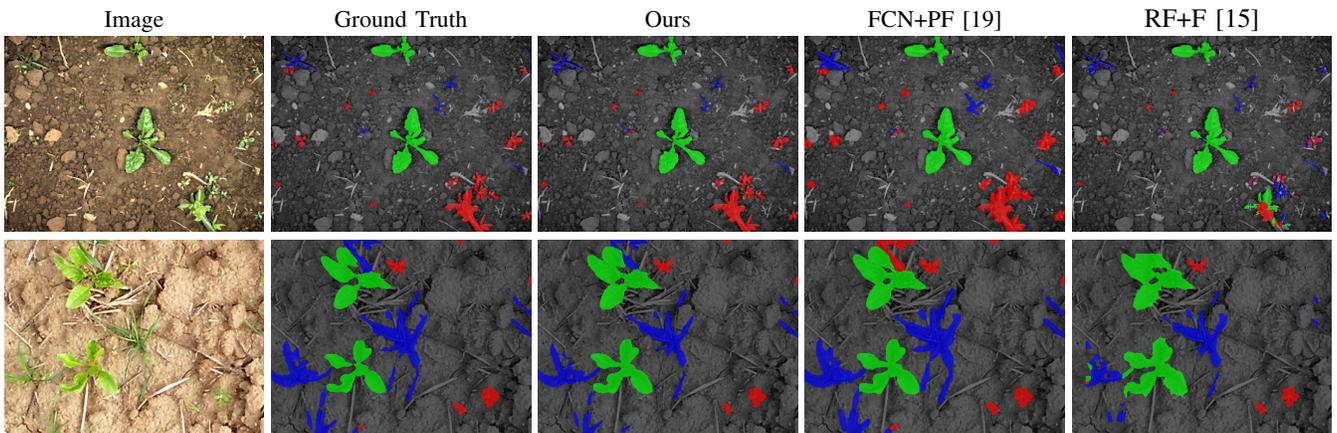


Fig. 5. Qualitative results of the pixel-wise semantic segmentation for BoniRob dataset (top row) and UAV dataset (bottom row). From left to right: RGB input image, the corresponding ground truth information for the pixel-wise class labels, prediction obtained by our proposed approach, FCN+PF [19], and RF+F [15]. Red refers to the class dicot weed, green refers to crop and blue refers to grass.

TABLE III  
SEGMENTATION PERFORMANCE ON BONIROB DATASET

Approach	mAP	Soil AP	Crop AP	Dicot AP	Grass AP
ours	83.8	99.8	91.2	69.4	75.0
RF+F [15]	68.9	97.7	85.1	46.4	46.4
FCN+PF [19]	64.5	99.8	85.3	38.1	34.8
baseline-seg	81.9	99.8	94.2	60.4	73.0

TABLE IV  
SEGMENTATION PERFORMANCE ON UAV DATASET

Approach	mAP	Soil AP	Crop AP	Dicot AP	Grass AP
ours	87.3	99.3	79.8	87.9	82.0
RF+F [15]	80.2	93.0	63.2	89.2	75.3
FCN+PF [19]	85.5	99.9	79.8	86.7	75.6
baseline-seg	87.8	99.4	75.7	88.9	75.7

5% in mAP with respect to the baseline-stem approach on BoniRob dataset. On UAV dataset the performance is comparable. This suggests that the stem detection benefits from using the shared encoder influenced by both the stem

detection loss  $\mathcal{L}_{\text{stem}}$  and the plant segmentation loss  $\mathcal{L}_{\text{plant}}$ . Furthermore, it is computationally more efficient compared to using two separate networks as it saves around  $\frac{1}{4}$  of the parameters by sharing the encoder.

Fig. 4 illustrates qualitative results of the stem detection in comparison to the other approaches. We can see that our approach performs best regarding the stem detection for the dicot weed. The random forest based approaches tend to detect more false positives for crops and dicot weed stems on the image parts containing vegetation. The FCN-based approach most probably benefits from the learned features providing a richer representation for the given task.

Notably, we had to manually fine-tune the vegetation detection for the random forest-based approaches (RF and RF+F), since the automated thresholding for the vegetation detection step did not lead to satisfactory results. This holds especially for the UAV dataset as it does not provide the additional NIR information, which typically aids the vegetation segmentation. In contrast for our approach, we selected only one set of hyper-parameters, such as the training schedule and initialization scheme, for training both datasets.

In terms of the MAD, we see that all approaches provide the stem position within around 2 mm in object-space, which is a sufficient accuracy for precise, plant-specific treatment, like mechanical stamping.

### C. Semantic Segmentation Performance

The second experiment is designed to show the performance of the pixel-wise semantic segmentation and to support our second claim that our approach provides an accurate segmentation of the images into the classes crop, dicot weed, grass, and background. Here, we compare again with RF+F [15], but now let the random forest predict the pixel-wise classification of the image. In addition to that we compare the performance with a state-of-the-art approach [19] employing FCNs, denoted by FCN+PF.

Tab. III summarizes the performance obtained for BoniRob dataset. Here, our approach achieves the best results. With a mAP of 83.8% most of the plants are correctly segmented. Analogous to the stem detection experiment, the better performance is mainly due to the high precision and recall for weed classes, i.e. dicot weed and grass.

Fig. 5 illustrates qualitative results of the semantic segmentation for both datasets. The analysis of the qualitative result shows that our approach properly segments small weeds and grasses, whereas the RF+F [15] has visibly more false detections and the FCN+PF [19], tends to have more “blobby” prediction. In turn, this leads to a high recall for weeds, but decreases the precision for these classes.

Regarding the comparison with the baseline-seg model, we observe a similar behavior as for the stem detection, i.e., a better performance on BoniRob dataset and comparable one for UAV dataset. These results show that our approach provides state-of-the-art performance for the semantic segmentation task outperforming two separate FCNs.

## V. CONCLUSION

In this paper, we presented a novel approach for joint stem detection and crop-weed segmentation using a FCN. We see our approach as a building block enabling farm robots to perform selective and plant-specific weed treatment. Our proposed architecture enables a sharing of feature computations in the encoder, while using two distinct task-specific decoder networks for stem detection and pixel-wise semantic segmentation of the input images. The experiments with two different datasets demonstrates the improved performance of our approach in comparison to state-of-the-art approaches for stem detection and crop-weed classification.

## ACKNOWLEDGMENTS

We thank R. Pude and his team from the Campus Klein Altendorf for their great support as well as A. Kräußling, F. Langer, and J. Weyler for labeling the datasets.

## REFERENCES

- [1] V. Badrinarayanan, A. Kendall, and R. Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(12):2481–2495, 2017.
- [2] N. Chebrolu, P. Lottes, A. Schaefer, W. Winterhalter, W. Burgard, and C. Stachniss. Agricultural Robot Dataset for Plant Classification, Localization and Mapping on Sugar Beet Fields. *Intl. Journal of Robotics Research (IJRR)*, 2017.
- [3] V. Dumoulin and F. Visin. A guide to convolution arithmetic for deep learning. *arXiv preprint*, abs/1603.07285, 2018.
- [4] M. Everingham, L. v. Gool, C.K.I. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *Intl. Journal of Computer Vision (IJCV)*, 88(2):303–338, 2010.
- [5] S. Haug, P. Biber, A. Michaels, and J. Ostermann. Plant stem detection and position estimation using machine vision. In *Workshop Proc. of Conf. on Intelligent Autonomous Systems (IAS)*, pages 483–490, 2014.
- [6] S. Haug, A. Michaels, P. Biber, and J. Ostermann. Plant Classification System for Crop / Weed Discrimination without Segmentation. In *IEEE Winter Conf. on Appl. of Computer Vision (WACV)*, 2014.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proc. of the IEEE Intl. Conf. on Computer Vision (ICCV)*, 2015.
- [8] G. Huang, Z. Liu, L.v.d. Maaten, and K. Q. Weinberger. Densely Connected Convolutional Networks. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] S. Jégou, M. Drozdal, D. Vázquez, A. Romero, and Y. Bengio. The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation. *arXiv preprint*, abs/1611.09326, 2017.
- [10] S. Kiani and A. Jafari. Crop detection and positioning in the field using discriminant analysis and neural networks based on shape features. *Journal of Agricultural Science and Technology*, 14:755–765, 07 2012.
- [11] F. Kraemer, A. Schaefer, A. Eitel, J. Vertens, and W. Burgard. From Plants to Landmarks: Time-invariant Plant Localization that uses Deep Pose Regression in Agricultural Fields. In *IROS Workshop on Agri-Food Robotics*, 2017.
- [12] M. Lin, Q. Chen, and S. Yan. Network In Network. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2014.
- [13] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [14] P. Lottes, M. Hoferlin, S. Sanders, and C. Stachniss. Effective Vision-Based Classification for Separating Sugar Beets and Weeds for Precision Farming. *Journal of Field Robotics (JFR)*, 2016.
- [15] P. Lottes, H. Markus, S. Sander, M. Matthias, S.L. Peter, and C. Stachniss. An Effective Classification System for Separating Sugar Beets and Weeds for Precision Farming Applications. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2016.
- [16] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- [17] C.S. McCool, T. Perez, and B. Upcroft. Mixtures of Lightweight Deep Convolutional Neural Networks: Applied to Agricultural Robotics. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2017.
- [18] H.S. Midtby, T.M. Giselsson, and R.N. Joergensen. Estimating the plant stem emerging points (pseps) of sugar beets at early growth stages. *Biosystems Engineering*, 111(1):83 – 90, 2012.
- [19] A. Millioto, P. Lottes, and C. Stachniss. Real-time Semantic Segmentation of Crop and Weed for Precision Agriculture Robots Leveraging Background Knowledge in CNNs. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2018.
- [20] M. A. Rahman and Y. Wang. Optimizing Intersection-Over-Union in Deep Neural Networks for Image Segmentation. In *Int. Symp. on Visual Computing*, 2016.
- [21] I. Sa, Z. Chen, M. Popvic, R. Khanna, F. Liebisch, J. Nieto, and R. Siegwart. weedNet: Dense Semantic Weed Classification Using Multispectral Images and MAV for Smart Farming. *IEEE Robotics and Automation Letters (RA-L)*, 3(1):588–595, 2018.
- [22] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal on Machine Learning Research (JMLR)*, 15:1929–1958, 2014.
- [23] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.