

BonnBeetClouds3D: A Dataset Towards Point Cloud-Based Organ-Level Phenotyping of Sugar Beet Plants Under Real Field Conditions

Elias Marks Jonas Bömer Federico Magistri Anurag Sag Jens Behley Cyrill Stachniss

Abstract—Agricultural production is facing challenges in the next decades induced by climate change and the need for more sustainability by reducing its impact on the environment. Advances in field management through robotic intervention, monitoring of crops by autonomous unmanned aerial vehicles (UAVs) supporting breeding of novel and more resilient crop varieties can help to address these challenges. The analysis of plant traits is called phenotyping and is an essential activity in plant breeding; it however involves a great amount of manual labor. With this paper, we provide means to better tackle the problems of instance segmentation to support robotic intervention and automatic fine-grained, organ-level geometric analysis needed for precision phenotyping. As the availability of real-world data in this domain is relatively scarce, we provide a novel dataset that was acquired using UAVs capturing high-resolution images of real breeding trials containing 48 plant varieties and therefore covering a relevant morphological and appearance spectrum. This enables the development of approaches for instance segmentation and autonomous phenotyping that generalize well to different plant varieties. Based on overlapping high-resolution images taken from multiple viewing angles, we provide photogrammetric dense point clouds and provide detailed and accurate point-wise labels for plants, leaves, and salient points as the tip and the base in 3D. Additionally, we include measurements of phenotypic traits performed by experts from the German Federal Plant Variety Office on the real plants, allowing the evaluation of new approaches not only on segmentation and keypoint detection but also directly on actual traits. The provided labeled point clouds enable fine-grained plant analysis and support further progress in the development of automatic phenotyping approaches, but also enable further research in surface reconstruction, point cloud completion, and semantic interpretation of point clouds.

I. INTRODUCTION

Agricultural production of food, feed, and fiber is facing big challenges in the next decades. While the world population is steadily growing, inducing a higher demand for agricultural products, the availability of arable land is decreasing due to climate change [28]. At the same time, we have to reduce the effects of conventional agricultural practices on our environment calling for a more sustainable agricultural production in order to preserve biodiversity [34]

E. Marks, F. Magistri, A. Sag, J. Behley, and C. Stachniss are with the Center for Robotics, University of Bonn, Germany. J. Bömer is with Institute of Sugar Beet Research, Göttingen, Germany. Stachniss is additionally with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany.

This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 – 390732324 – PhenoRob, under STA 1051/5-1 within the FOR 5351 (AID4Crops), and by the Federal Ministry of Food and Agriculture (BMEL) based on a decision of the Parliament of the Federal Republic of Germany via the Federal Office for Agriculture and Food (BLE) under the innovation support programme under funding no 28DK108B20 (RegisTer).




Fig. 1: Our dataset, called BonnBeetClouds3D consists of high-resolution point clouds extracted via bundle adjustment of high-resolution images captured from multiple viewing angles by a UAV. Together with the point clouds, we also provide accurate per-point labels of the plants, leaves, and salient points of the leaves, such as leaf tips, corners and bases. Furthermore, we provide reference values of measurements that are relevant for phenotyping, such as leaf lengths and widths.

and the productivity of arable lands [8], [27], [40]. As one measure, agricultural production needs to significantly reduce the amount of applied agrochemicals, such as pesticides and herbicides.

A potential pathway to enable a more sustainable agricultural production is the breeding of more resilient crop varieties, novel field management techniques including more targeted interventions at a plant level, and an increase of non-chemical weeding practices to reduce the amount of needed herbicides. These areas benefit from artificial intelligence-based solutions, especially advances in robotics and computer vision [33] to obtain a better estimate of the crop status in the field. In particular, the development of more capable perception systems for robotic platforms or automatized tractors could lead to a higher automation enabling plant-specific treatments, but also novel ways of analyzing data from the field could lead to advancement towards more sustainable agricultural production.

In this paper, we provide a dataset to target plant trait assessment in the field by observing the phenotype of a plant, which is an expression of its genotype, influenced by environmental conditions, and affected by the field management.



Fig. 2: Experts estimating phenotypic traits in the fields. This task is performed by visual inspection and, thus, depends also on the subjective observation of the experts.

An example of the provided data can be seen in Fig. 1.

The assessment of the morphological characteristics of a plant is to this date often performed by manual measurements or visual scoring of a trained expert. These experts evaluate different traits such as growth stage, leaf attitude, plant height, and leaf length as shown in Fig. 2. However, this assessment is a tedious and labour-intensive task, and also subjective to a certain extent. All this affects the spatial extent, reliability, and repeatability of phenotyping activities.

A potential solution to the aforementioned limitations of conventional phenotyping is the use of algorithmic solutions to automate the process, which also leads to repeatable and objective measurements of phenotypic traits. In particular, a combination of a robotic solution, which automatically acquires the relevant data in combination with a robot perception or computer vision-based approach to extract such traits from the recorded sensor data is a way to tackle the so-called "phenotyping bottleneck" [10].

The significant progress in image-based perception, but also the development of increasingly capable perception systems in autonomous driving, has been propelled by the availability of domain-specific datasets [2], [3], [6], [11], [15], [35] to a large extent. While there has been recently an increasing interest in agricultural robotics, this progress is stifled by the availability of data to study and advance perception in this domain. There is an increasing number of datasets available in the agricultural domain for semantic interpretation of 2D field images [5], [30], [39], not many, however, contain 3D plant data, which is highly relevant for morphological trait estimation [32].

The main contribution of this paper is a large dataset covering 48 different varieties of sugar beets that we recorded on real breeding trials with over 3,000 plants, with point-

wise annotations for 186 individual plants and 2,661 individual leaves enabling the development and evaluation of segmentation, detection and actual in-field phenotyping algorithms. To allow for direct evaluation of tasks such as extraction of phenotypic traits, we also provide reference values for commonly evaluated traits such as leaf length, leaf width and stem length alongside the position of over 10,000 salient points such as the leaf tips, leaf corners and plant centers. To obtain the data, data loading utilities, evaluators, or submit your results to the public challenge please visit <https://bonnbeetclouds3d.ipb.uni-bonn.de>

In sum we make the following contributions. Our dataset enables: (i) studying phenotyping using photogrammetric point clouds by exploiting geometric information; (ii) instance segmentation of plant organs using 3D point clouds; (iii) estimation of key locations on the leaves that are useful for phenotyping.

II. RELATED WORK

There has been increasing interest in semantic interpretation of field images targeting semantic [16], [17], [18], [19], [25], instance [38], and panoptic segmentation [29] of crops and weeds. These approaches mainly use large image datasets with pixel-wise annotations of crops and weeds [5], [30], but also leaf annotations [13], [39]. For a broader overview of agricultural datasets for semantic interpretation, we refer to the survey by Lu et al. [20].

While there are several image datasets available, few datasets provide point cloud data revealing the geometric structure of the plants that seems mandatory for a fine-grained structural analysis of phenotypic traits. Chaudhury et al. [4] proposed a synthetic dataset generated based on plant models with the addition of noise. While being able to generate large amounts of data, the gap between simulation and reality poses a challenge for the deployment of approaches developed on this data to real fields. The Pheno4D dataset [32] provides point clouds of real tomato and maize plants at different growth stages captured with a high-precision laser scanner with accurate labels of individual plant organs. In contrast to our dataset, these point clouds were acquired in a lab setting under controlled conditions. Dutagaci et al. [9] instead provide a dataset with real world data, but it contains rose plants. Khanna et al. [12] provide a dataset containing sugar beet plants, but they do not provide point-wise instance labels nor reference measurements for leaf measures. Tab. I summarizes the key characteristics of the point cloud datasets in the agricultural domain, which shows that our dataset has unique characteristics not covered by the available datasets. Our dataset covers real breeding trials showing real-world lighting conditions, but also occlusions induced by the planting of the crops. Furthermore, it is augmented with real measurements by the phenotyping experts from the German Federal Plant Variety Office providing reference measurements of per-plot phenotypic traits.

Dataset	Real data	Field	Semantic labels	Instance labels	Phenotypic measurements	Varieties
Dutagaci et al. [9]	✓		✓	✓		Rose
Khanna et al. [12]	✓	✓				Sugar beet
Pheno4D [32]	✓		✓	✓		Tomato, Maize
BonnBeetClouds3D (Ours)	✓	✓	✓	✓	✓	Sugar beet

TABLE I: Overview of the key characteristics of the point cloud datasets in the agricultural domain, which shows that our dataset has unique characteristics not covered by the available datasets in the agricultural domain.

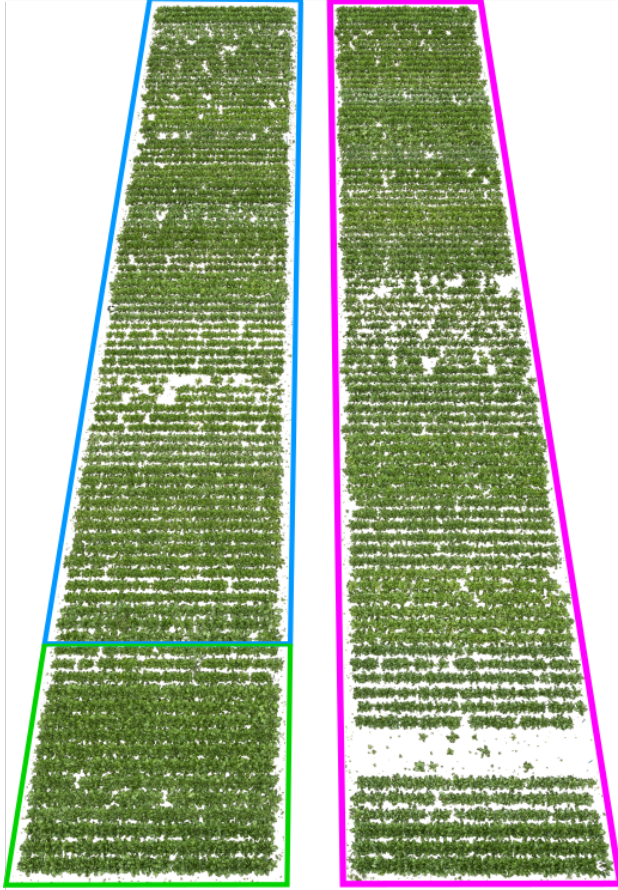


Fig. 3: Overview of the captured field. The data used for training of deep learning approaches is highlighted in purple, the validation area is highlighted in green and the test data is in the blue rectangle.

III. THE BONNBEETCLOUDS3D DATASET

A. Field Setup

We recorded the data of the sugar beet trials of the Federal Plant Variety Office in Magdeburg, Germany (N52.112028°, E011.558448°). The trial was set up in a single factorial block design with the factor variety which contained two repetitions. The plot size was 1.5 m by 7.0 m and contained three rows of plants with a row distance of 0.5 m. We cover a total of 48 varieties in our dataset.

The field technicians of the German Federal Plant Variety Office sowed the plants at a distance of 6 cm inside the row and later manually thinned them out to a final distance of 18 cm by manual hoeing. Plant protection products like herbicides were used in the normal course of business. Any remaining weeds were regularly removed by manual hoeing.




Fig. 4: In the top left part (a) we show an unlabeled point cloud with the original leaf colors, while in the bottom right (b) we show the color-coded labeled leaf instances, where different colors correspond to different leaves.

In addition, the experiment was artificially irrigated when the weather conditions were too dry. These labor-intensive tasks were made to create ideal conditions for the development of the varieties.

B. Sensor Setup

We recorded our field data using a UAV equipped with a PhaseOne iXM-100 camera with a 80 mm RSM prime lens mounted on a gimbal to obtain motion-stabilized RGB 100 MP images at a resolution of 11,664 px \times 8,750 px per image. The UAV was flying at a height of approx. 21 m, resulting in a ground sampling distance (GSD) of 1 $\frac{\text{mm}}{\text{px}}$. We flew three missions on the same field with a camera angle of 45°, 90°, and 135° degrees from the ground plane, which allows us to obtain good coverage of the crops including the lower parts even in the tight row spacing.

C. Data Processing and Labeling

We process the recorded, overlapping 100 megapixel raw camera images of different view angles to determine the relative poses between the images via photogrammetric bundle adjustment, where we used the GNSS positions as initial camera positions. Using the poses, we compute a 3D point cloud of the complete field as shown in Fig. 3.

To obtain a reliably accurate segmentation of plants and leaves, we manually labeled the individual plants and leaf instances in the point clouds using the Semantic Segmentation Editor by Hitachi Automotive and Industrial Laboratory. An example of the annotated point cloud is shown in Fig. 4.

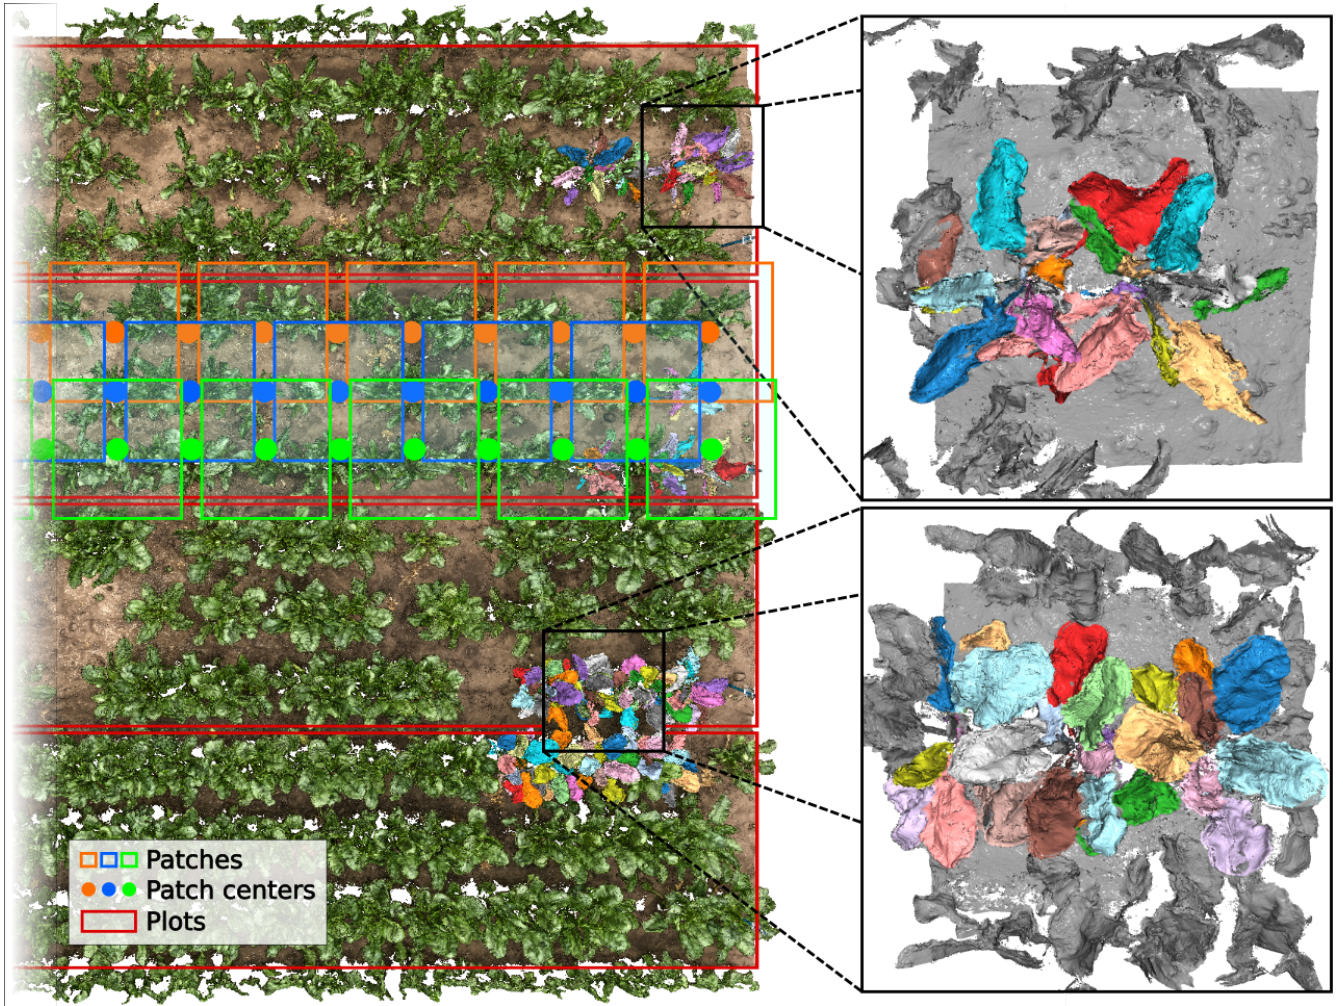


Fig. 5: The patch extraction process from the point cloud of the field. As can be seen in the field representation, a subset of plant has been annotated, here the leaf color encode the respective leaf instances. From the whole field point cloud we then extract patches, as the entire point cloud is too big for being processed at once. The colored dots represent the centers of all the patches extracted from that plot. The colored boxes, instead, show only a subset of the extracted patches for better visualization. On the right there are two example patches with color-coded leaf ids and grey for the unlabeled part.

The annotated plants were carefully checked by a person different from the one that annotated them in the first place, so we ensured that multiple people checked each label in order to obtain a higher accuracy in the resulting annotations. There are many plants present for every variety, making it intractable and also unnecessary to label all of them. Therefore, we selected one group of adjacent plants per variety and annotated them. As processing the whole field point cloud at once is impossible due to the amount of points, we extracted regularly spaced and overlapping patches covering the crop rows. The patching process is shown in Fig. 5. The whole labeling process required substantial manual labor, but resulted in 186 annotated plants and 2,661 annotated leaves, covering all 48 varieties in the breeding trial. We furthermore manually annotated different relevant keypoints of each individual leaf as shown in Fig. 6. The leaf base represents the point where the leaf blade joins the petiole, the leaf tip represent the point that is furthest from the petiole and the lateral corners represent the points of the leaf that

are furthest away from the central axis of the leaf. The stem base, instead, is the point where the stem joins the beet crown, which we also define as the plant center. In sum, we annotated over 10,000 such keypoints.

The German Federal Plant Variety Office performed measurements of different variety traits of 30 leaves for every breeding plot. These are the leaf blade length and width, and the petiole length and width. Additionally, they provided us the average leaf angle with respect to the ground, average plant height, the curliness and the curviness of the leaf border for each plot, which was scored by trained phenotyping experts.

D. Data Organization

To ensure a fair evaluation of the tested approaches, we divided the dataset into a training set consisting of 1,782 leaves from 128 plants to develop and train approaches, a validation set of 260 leaves from 17 plants to tune the hyperparameter of the approaches and a test set of 619 leaves

TABLE II: Attributes of the points in the dataset. Each point has an associated coordinate, normal, color, leaf id, plant id, plot id and keypoint id.

Attributes	Data Type	Range	Notes
point coordinates	$\mathbb{R}^{N \times 3}$	$[-\infty, \infty]$	Point coordinates in m
point normal	$\mathbb{R}^{N \times 3}$	$[-1, 1]$	Unit vector representing point normal
point color	$\mathbb{Z}^{N \times 3}$	$[0, 255]$	8 bit point color
leaf id	$\mathbb{Z}^{N \times 1}$	$[-2, \text{number of leaves}]$	-2: unlabeled, -1: plant center
plant id	$\mathbb{Z}^{N \times 1}$	$[-1, \text{number of plants}]$	-1: unlabeled
plot id	$\mathbb{Z}^{N \times 1}$	$[0, \text{number of plots}]$	Number of the plot the point belongs to
keypoint id	$\mathbb{Z}^{N \times 1}$	$[-1, 3]$	-1: no keypoint, 0: base, 1: left corner, 2: tip, 3: right corner




Fig. 6: Labeled keypoints. The leaf base is the point where the stem joins the leaf blade, the left and right corners are the two points furthest apart from the main leaf axis, which therefore can be used for the computation of the maximum leaf width, the leaf tip is the topmost point of the leaf and the stem base is the point where the stem joins the beet crown, which we also define as the plant center.

from 41 plants to validate the performance of the developed approaches. We provide labels for the train and validation sets. The point clouds are divided into patches of 1 m by 1 m by 1 m to make them easy to use for training deep neural networks. We use the ply format to provide these patches, where each point has the attributes shown in Tab. II.

The keypoint ids are labeled as subsets of the leaves, with multiple points of the same leaf having the same keypoint id. To get the most precise keypoint position, we find the centroid all points of one leaf having the same keypoint id. The ground truth measurements of the phenotypic traits are delivered for every plot in csv-format. These include:

- leaf blade length
- leaf blade width
- petiole length
- petiole width

We show some examples of the data in Fig. 7.

To obtain the data, data loading utilities, evaluators, or submit your results to the public challenge please visit <https://bonnbeetclouds3d.ipb.uni-bonn.de>

TABLE III: Leaf trait detection in the field. In the table we report the errors in the estimation of the 3 leaf traits in field conditions. We present the offsets in millimeters of the estimations from the manually measured ground truth values for each approach.

Approach	leaf length [mm] ↓ avg	blade length [mm] ↓ avg	blade width [mm] ↓ avg
CPD [26]	205	139	109
PF-SGD [22]	35	52	15
CurvLeaf [21]	38	18	13

TABLE IV: Instance segmentation of plant organs. In the table we report the performance of the three baselines. We report panoptic quality (PQ), segmentation quality (SQ), and recognition quality (RQ). The metrics are introduced by Kirillov et al. [14].

Approach	PQ [%] ↑	SQ [%] ↑	RQ [%] ↑
SoftGroup [37]	72.07	79.45	90.69
Mask3D [31]	18.47	67.71	25.80
HP-LIS [23]	75.58	80.97	93.17

IV. EXPERIMENTS

In sum our dataset supports: (i) studying phenotyping using photogrammetric 3D point clouds, by exploiting geometric information; (ii) instance segmentation of plant organs using 3D point clouds; (iii) estimation of key locations on the leaves that are useful for phenotyping.

We show how different existing approaches perform on the different tasks above, so that they can be used as a baseline in the evaluation of new approaches.

A. Estimating Geometric Phenotypic Traits in the Fields

As we provide precise ground truth measurements for the geometric phenotypic traits, i.e., leaf length, leaf blade length, and leaf width, the dataset can be used to evaluate how well new approaches are able to extract those features from the photogrammetric point clouds. This kind of approaches are important for improving the throughput and accuracy of phenotyping practices. We evaluated the approaches based on the prediction error for the different traits in mm. The results of the three approaches are shown in Tab. III.

The first approach consists in using the coherent point drift algorithm (CPD) [26] to deform a triangular mesh to the point cloud and estimating the leaf parameters based on the obtained deformed mesh. The algorithm reconstructs




Fig. 7: The point cloud patches, shown both with RGB colors and color-coded leaf ids based on the annotations. The grey points represent the unlabeled part. The colored circles in the annotated view represent the position of the annotated keypoints. To avoid the figure to be too cluttered we show only the leaf tips and bases and omit the two corners.

completely observed leaves quite well, but it fails for the ones presenting occlusions. The mesh resulting from CPD always collapses on the visible part of the leaves. This leads to a bad performance in the estimation of the leaf dimensions.

Another approach that suffers less from this issue is the partiality filtered stochastic gradient descent, denoted as PF-SGD [22]. Here we used gradient descent to optimize the position of the vertices of a triangular mesh that represents our leaf model. The vertex positions are optimized in order to deform the mesh onto the point cloud of the leaf, and then the leaf dimensions are measured on the deformed model. This approach improves substantially in comparison to CPD.

To improve the robustness of the reconstruction we de-

veloped a curve based leaf model fitting algorithm, denoted as CurvLeaf [21]. Instead of optimizing the position of individual vertices, we here optimize the parameters of parametric curves to best fit the leaf model to the point clouds. This has the advantage that the deformations can be better constrained, leading to more robust estimation of leaf lengths, widths and blade lengths.

B. Instance Segmentation of Plant Organs using 3D Point Clouds

Another important task needed for autonomous high throughput phenotyping is the detection and segmentation of individual plant organs such as leaves and stems in the

TABLE V: Detection of leaf keypoints. In the table we report the errors in the estimation of the 4 leaf keypoints in field conditions. We present the offsets in millimeters of the estimations from the manually annotated ground truth keypoints.

Approach	leaf tip [mm]	leaf base [mm]	left corner [mm]	right corner [mm]
	↓ avg	↓ avg	↓ avg	↓ avg
SVD	15	33	43	46
CPD [26]	30	68	52	50
PF-SGD [22]	17	24	43	41
CurvLeaf [21]	21	19	49	48

field. Only after detecting which part of the the point cloud corresponds to a specific leaf one can estimate the leaf dimensions. We compared different approaches for instance segmentation in point clouds that can be used as baselines for the evaluation of new approaches. As metric for comparison we used panoptic quality [14], which is defined as:

$$PQ = \underbrace{\frac{\sum_{(p,g) \in TP} \text{IoU}(p,g)}{|TP|}}_{SQ} \underbrace{\frac{|TP|}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}}_{RQ}. \quad (1)$$

The performance in terms of panoptic quality [14] is shown in Tab. IV.

The first approach that we evaluated is SoftGroup [37], which is a geometric learning approach that predicts offset vectors pointing to the corresponding leaf center for each point of the point cloud. These offset vectors are then clustered to obtain sets of point corresponding to individual leaves. The resulting leaves are separated quite well, but there are some points assigned to the wrong leaf, especially in the border regions.

To cover different kinds of neural networks, we evaluated a query-based transformer approach, which is performing very well on existing scene segmentation datasets like ScanNet [7], S3DIS [1], and SemanticKITTI [2]. The approach by Schult et al. [31], called Mask3D, learns query vectors during the training phase that are then decoded by the transformer decoder heads into masks covering single instances. Mask3D’s performance suffers a lot from the fact that labels are not present for all points in the point cloud patches while training, leading to worse performance then the other baselines.

Finally, we evaluated a domain-specific approach [23], denoted as HP-LIS. As SoftGroup, it estimates offset vectors pointing to the centroid of the corresponding leaf for each point and cluster them into leaf instances. We however use KPConv [36] as backbone, which implements a convolution defined directly in continuous space, instead of voxelizing the point cloud into a 3D grid. For clustering we use HDBSCAN [24] which is slower than the clustering algorithm used in SoftGroup but leads to better segmentation performance.

C. Estimation of Key Locations on the Leaves

The identification of key locations on the leaves such as the leaf tip, the leaf base and the leaf lateral extremities are useful to detect parts of plant organs. These key locations are also useful for computing phenotypic traits relying on

them, e.g., the distance between different parts or the relative orientation in space. This problem can be approached in different ways. One such way is direct detection of the salient points, which we evaluated by Eigenvector analysis of the point cloud, which we report as SVD. Additionally, we detected the keypoints based on the models fitted by CPD, PF-SGD and CurvLeaf, as those points are known positions on the predefined model.

To compare the performance of the different leaf keypoint detection approaches we compute the L2 distance in mm between the predicted and the manually labeled keypoints, which we show in Tab. V. The leaf tips and bases are detected well by some approaches, while the corners are harder to detect. The main reasons for this are the heavy occlusions and irregular leaf shapes. CPD performs the worst for all keypoints, as the meshes are not well fitted to the point clouds. SVD works quite well, especially for the leaf tip detection. CurvLeaf is best at detecting the leaf base, while PF-SGD best detects the corners.

V. CONCLUSION

This paper presents a novel dataset for agricultural robotics, which enables research on automated morphological parameter extraction for phenotyping. We provide reference values both for intermediate tasks like instance segmentation and for downstream tasks like performing automated leaf measurements. Additionally, we evaluated the performance of existing methods, which can be used to solve the different tasks, to provide an insight on the current state of the art and to provide baselines for further research on the topic. A vision system that is able to segment individual plant organs can enable highly automated robotic interaction with the plants in real fields, but also improve grasping performance, view planning, and mapping.

ACKNOWLEDGMENTS

We thank the German Plant Variety Office (Bundessortenamt), especially Dr. Manthey, Dr. Lichthardt, Dr. Beukert, Dr. Borg, Ms. Zickenrott, and Mr. Bauske, for maintaining the experimental fields on which the data was collected, for supporting the data collection, for collecting the reference data for the phenotypic traits, for the technical advice on the current phenotyping practices, and for a pleasant collaboration. We also thank Facundo Ispizua Yamati and Dirk Koops for providing a picture of them while estimating phenotypic traits in the field.

REFERENCES

- [1] I. Armeni, O. Sener, A.R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [2] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2019.
- [3] M. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays. Argoverse: 3D Tracking and Forecasting With Rich Maps. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [4] A. Chaudhury, F. Boudon, and C. Godin. 3D Plant Phenotyping: All You Need is Labelled Point Cloud Data. In *ECCV Workshop on Computer Vision Problems in Plant Phenotyping*, 2020.
- [5] 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)*, 36(10):1045–1052, 2017.
- [6] M. Cordts, S.M. Omran, Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] A. Dai, A. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] M. Diacono and F. Montemurro. Long-Term Effects of Organic Amendments on Soil Fertility. A Review. *Agronomy for Sustainable Development*, 30:1–23, 2010.
- [9] H. Dutağacı, P. Rasti, G. Galopin, and D. Rousseau. ROSE-X: an annotated data set for evaluation of 3D plant organ segmentation methods. *Plant Methods*, 16, 2020.
- [10] R.T. Furbank and M. Tester. Phenomics – technologies to relieve the phenotyping bottleneck. *Trends in Plant Science*, 16(12):635–644, 2011.
- [11] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [12] R. Khanna, L. Schmid, A. Walter, J. Nieto, R. Siegwart, and F. Liebisch. A spatio temporal spectral framework for plant stress phenotyping. *Plant Methods*, 15, 2019.
- [13] J. Kierdorf, L.V. Junker-Frohn, M. Delaney, M.D. Olave, A. Burkart, H. Jaenicke, O. Müller, U. Rascher, and R. Roscher. Growliflower: An image time series dataset for growth analysis of cauliflower. *Journal of Field Robotics (JFR)*, 40(2):173–192, 2022.
- [14] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár. Panoptic Segmentation. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C.L. Zitnick. Microsoft COCO: Common Objects in Context. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2014.
- [16] P. Lottes, J. Behley, N. Chebrolu, A. Milioto, and C. Stachniss. Robust Joint Stem Detection and Crop-Weed Classification using Image Sequences for Plant-Specific Treatment in Precision Farming. *Journal of Field Robotics (JFR)*, 37(1):20–34, 2020.
- [17] P. Lottes, M. Höferlin, S. Sander, M. Mütter, P. Schulze-Lammers, 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.
- [18] P. Lottes, M. Höferlin, S. Sander, and C. Stachniss. Effective Vision-based Classification for Separating Sugar Beets and Weeds for Precision Farming. *Journal of Field Robotics (JFR)*, 34(6):1160–1178, 2017.
- [19] P. Lottes, R. Khanna, J. Pfeifer, R. Siegwart, and C. Stachniss. UAV-Based Crop and Weed Classification for Smart Farming. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2017.
- [20] Y. Lu and S. Young. A survey of public datasets for computer vision tasks in precision agriculture. *Computers and Electronics in Agriculture*, 178:105760, 2020.
- [21] E. Marks, J. Bömer, F. Magistri, M. Sodano, A. Mahlein, J. Behley, and C. Stachniss. Autonomously estimating leaf parameters of individual plants in real agricultural environments. *Currently under review*, 2024.
- [22] E. Marks, F. Magistri, and C. Stachniss. Precise 3D Reconstruction of Plants from UAV Imagery Combining Bundle Adjustment and Template Matching. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.
- [23] E. Marks, M. Sodano, F. Magistri, L. Wiesmann, D. Desai, R. Marcuzzi, J. Behley, and C. Stachniss. High Precision Leaf Instance Segmentation in Point Clouds Obtained Under Real Field Conditions. *IEEE Robotics and Automation Letters (RA-L)*, 8(8):4791–4798, 2023.
- [24] L. McInnes and J. Healy. Accelerated hierarchical density based clustering. In *Proc. of the IEEE Intl. Conf. on Data Mining Workshops*, 2017.
- [25] A. Milioto, 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.
- [26] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 32(12):2262–2275, 2010.
- [27] P. Mäder, A. Fliessbach, D. Dubois, L. Gunst, P. Fried, and U. Niggli. Soil fertility and biodiversity in organic farming science. *Science*, 296:1694–7, 2002.
- [28] R. Právělie, C. Patriche, P. Borrelli, P. Panagos, B. Roşca, M. Dumitraşcu, I.A. Nita, I. Săvulescu, M.V. Birsan, and G. Bandoc. Arable lands under the pressure of multiple land degradation processes. a global perspective. *Environmental Research*, 194:110697, 2021.
- [29] G. Roggiolani, M. Sodano, T. Guadagnino, F. Magistri, J. Behley, and C. Stachniss. Hierarchical Approach for Joint Semantic, Plant Instance, and Leaf Instance Segmentation in the Agricultural Domain. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [30] I. Sa, M. Popovic, R. Khanna, Z. Chen, P. Lottes, F. Liebisch, J. Nieto, C. Stachniss, and R. Siegwart. WeedMap: A Large-Scale Semantic Weed Mapping Framework Using Aerial Multispectral Imaging and Deep Neural Network for Precision Farming. *Remote Sensing*, 10(9):1423, 2018.
- [31] J. Schult, F. Engelmann, A. Hermans, O. Litany, S. Tang, and B. Leibe. Mask3D for 3D Semantic Instance Segmentation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [32] D. Schunck, F. Magistri, R. Rosu, A. Cornelißen, N. Chebrolu, S. Paulus, J. Léon, S. Behnke, C. Stachniss, H. Kuhlmann, and L. Klingbeil. Pheno4D: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis. *PLOS ONE*, 16(8):1–18, 2021.
- [33] R. Sparrow and M. Howard. Robots in agriculture: prospects, impacts, ethics, and policy. *Precision Agriculture*, 22:818–833, 2021.
- [34] K. Stein-Bachinger, F. Gottwald, A. Haub, and E. Schmidt. To what extent does organic farming promote species richness and abundance in temperate climates? A review. *Organic Agriculture*, 11:1–12, 2021.
- [35] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [36] H. Thomas, C. Qi, J. Deschaud, B. Marcote, G. Goulette, and L. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2019.
- [37] T. Vu, K. Kim, T. Luu, T. Nguyen, and C. Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [38] J. Weyler, F. Magistri, P. Seitz, J. Behley, and C. Stachniss. In-Field Phenotyping Based on Crop Leaf and Plant Instance Segmentation. In *Proc. of the IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2022.
- [39] J. Weyler, F. Magistri, E. Marks, Y.L. Chong, M. Sodano, G. Roggiolani, N. Chebrolu, C. Stachniss, and J. Behley. PhenoBench: A Large Dataset and Benchmarks for Semantic Image Interpretation in the Agricultural Domain. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [40] S.E. Wortman, T.D. Galusha, S.C. Mason, and C.A. Francis. Soil fertility and crop yields in long-term organic and conventional cropping systems in eastern nebraska. *Renewable Agriculture and Food Systems*, 27(3):200–216, 2012.