A Dataset and Benchmark for Shape Completion of Fruits for Agricultural Robotics

Federico Magistri, Thomas Läbe, Elias Marks, Sumanth Nagulavancha, Yue Pan, Claus Smitt, Lasse Klingbeil, Michael Halstead, Heiner Kuhlmann, Chris McCool, Jens Behley, Cyrill Stachniss

Abstract—As the world population is expected to reach 10 billion by 2050, our agricultural production system needs to double its productivity despite a decline of human workforce in the agricultural sector. Autonomous robotic systems are one promising pathway to increase productivity by taking over labor-intensive manual tasks like fruit picking. To be effective, such systems need to monitor and interact with plants and fruits precisely, which is challenging due to the cluttered nature of agricultural environments causing, for example, strong occlusions. Thus, being able to estimate the complete 3D shapes of objects in presence of occlusions is crucial for automating operations such as fruit harvesting. In this paper, we propose the first publicly available 3D shape completion dataset for agricultural vision systems. We provide an RGB-D dataset for estimating the 3D shape of fruits. Specifically, our dataset contains RGB-D frames of single sweet peppers in lab conditions but also in a commercial greenhouse. For each fruit, we additionally collected high-precision point clouds that we use as ground truth. For acquiring the ground truth shape, we developed a measuring process that allows us to record data of real sweet pepper plants, both in the lab and in the greenhouse with high precision, and determine the shape of the sensed fruits. We release our dataset, consisting of almost 7,000 RGB-D frames belonging to more than 100 different fruits. We provide segmented RGB-D frames, with camera intrinsics to easily obtain colored point clouds, together with the corresponding high-precision, occlusion-free point clouds obtained with a high-precision laser scanner. We additionally enable evaluation of shape completion approaches on a hidden test set through a public challenge on a benchmark server.

I. INTRODUCTION

Our agricultural production system needs to double its production of food, feed, fiber, and fuel to cope with an ever-growing population. A promising solution to support us in achieving this goal is the usage of autonomous systems that can continuously monitor the field, selectively harvest fruits and crops, or spot diseases early enough to avoid loss of yield. Such systems need to estimate important features of fruits and crops such as size, shape, and health status from perceived sensor data. In agricultural environments, this is extremely challenging due to their cluttered nature leading to plant and fruits being partially occluded.

In this paper, we aim to facilitate the development of agricultural vision systems in presence of such occlusions.



Fig. 1: With our dataset, we tackle the problem of estimating the shape of fruits as a mesh (shown in grey) given a partial observation of an RGB-D sensor providing a colored point cloud. The estimated shape of the fruit is essential to allow safe grasping in situation with severe occlusions, like a greenhouse environment.

Occlusions are a key issue in developing vision systems in agricultural environments. A correct handling of occlusions could benefit a diverse number of agricultural tasks such as high-throughput phenotyping, yield estimation and forecasting, and robotic harvesting.

The main solutions for handling occlusions in agricultural environments have been constrained to 2D images, either by amodal instance segmentation methods that produce masks combining the visible and occluded part of crops [4] or by generative adversarial networks that estimate the appearance of crops as if there are no occlusions [17]. In contrast to the 2D case, fewer efforts have been made to tackle the occlusion problem in a 3D scene. We are taking a different avenue to address occlusions in the context of fruit picking using 3D data provided by commonly employed RGB-D sensors. In particular, manipulation tasks in robotics commonly employ 3D models of the objects to perform grasp prediction, but usually employ CAD models of known objects for this purpose, which cannot be pre-determined in an agricultural context, as each fruit has a different shape. Therefore, we propose a dataset for promoting research on predicting a complete shape model of real-world fruits that can be used in manipulation scenarios [24]. Thereby, we are addressing the challenging task of uncovering the full 3D shape of partially occluded fruits. Achieving this allows us to perform more

All authors are with the Center for Robotics at the University of Bonn, Germany. Chris McCool and Cyrill Stachniss are additionally with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany. Cyrill Stachniss is furthermore with the Department of Engineering Science at the University of Oxford, UK.

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 – 390732324 – PhenoRob and under STA 1051/5-1 within the FOR 5351 (AID4Crops).

reliable fruit grasping, and, can be used to derive phenotypic traits about the fruit, such as size, shape, and quality.

The main contribution of this paper is the release of the first publicly available 3D shape completion dataset consisting of RGB-D frames and corresponding high-precision, occlusion-free point clouds obtained with a high-precision laser scanner. The prediction task targets the estimation of the complete fruit shape from partial observations, as shown in Fig. 1. In this paper, we also describe the process of generating the ground truth shape models for the sweet peppers in a real-world commercial greenhouse. We specifically design a measuring process that allows us to acquire highly accurate and dense point cloud models of the sweet peppers.

In sum, our contributions can be summarized as follows: (i) we release the first public dataset¹ for 3D shape completion in agricultural environments, (ii) we provide a CodaLab competition with a hidden test set to foster research in this area, (iii) we provide a description of the process to obtain the dataset to facilitate other researcher's data collection efforts, (iv) we release a Python-based development toolkit² for handling the dataset and computing metrics.

II. RELATED WORK

Dataset and associated benchmarks have a long history in robotics and computer vision. They enable researchers to quantitatively and qualitatively measure the progress of research providing the foundation of a reproducible evaluation. The datasets provide a starting point for investigation of novel research areas, where it is hard to acquire data or the data generation process requires specialized equipment.

In recent years, there has been an increasing interest in semantic interpretation of images in the agricultural context, both on arable crops [20], [38], [32] and horticulture [13], [33], [2]. These works rely on large image datasets with pixel-accurate labels designed for various tasks: crop/weed segmentation [6], [34], leaf segmentation [16], [37], fruit counting [36], [31], and fruit size estimation [10], [12]. We refer to Lu et al. [21] for an overview of image datasets in agricultural environments. In contrast to these datasets, we provide 3D ground truth shapes of fruits in the agricultural context rather than segmentation labels.

While there are several image datasets available, few provide point cloud data and none of them tackle the occlusion problem in agriculture to the best of our knowledge. For instance, Chaudhury et al. [5] proposed a synthetically generated point cloud dataset based on plant models with semantic labels. Furthermore, the Pheno4D dataset [35] provides point clouds of tomato and maize plants captured with a high-precision laser scanner with labels of individual leaves that are consistent over time. James et al. [14] provide a temporally consistent dataset of strawberry point clouds. Dutagaci et al. [9] release a rosebush plants 3D dataset acquired through X-ray tomography with semantic labels. Khanna et al. [15] present a dataset containing color images, infra-red stereo image pairs, and multi-spectral camera images along with applied treatments and weather conditions of the surroundings. Finally, Marks et al. [25] published a sugar beet point cloud dataset with leaves instance labels from real breeding plots. In contrast to these works, our dataset is specifically designed to tackle the occlusion problem in agricultural environments by providing complete shapes of sweet peppers with sub-millimeter accuracy.

Recently, a few datasets have been released for mapping applications in agricultural environments. The MAgro dataset [26] consists of robotic sensor data, such as 3D LiDAR, data from an inertial measurement unit, and wheel encoders, gathered in apple and pear orchards with calibrated RTK GNSS to evaluate localization methods. The Bacchus dataset [30] captures the whole canopy growth of a vineyard tailored for mapping and localization algorithms for longterm autonomous robotic operation. Instead, we designed our dataset to estimate precisely the 3D shape of individual fruits, which could be exploited for robotic grasping applications.

In summary, our dataset complements available point cloud datasets by providing measurements from a greenhouse using an RGB-D sensor with occlusions caused by leaves and other fruits. By providing the accurately measured shapes of the sensed fruits, we release a dataset for shape completion from partial observations and furthermore enable unbiased and reproducible evaluation on a hidden test set.

III. OUR DATASET

We propose a dataset of fruits collected in two scenarios. First, data collected in a lab environment for controlled experiments and, second, data collected in a commercial greenhouse showing the full complexity of the application scenario. In our benchmark, we use data collected in the lab for training set, while we use data collected in the greenhouse for testing set. The validation set contains data coming from both environments. The reason for this separation lies in the difficulty in collecting greenhouse data, having a procedure to collect data in the lab allows us to have a large training set.

As inputs in both scenarios, we choose to use RGB-D frames collected with an Intel RealSense d435i, which is a commonly used camera in agricultural robotics research. For collecting ground truth data capturing the shape of the fruits, we use a high-precision LiDAR, a Perceptron V5 laser scanner, mounted on a non-actuated measuring arm, a Hexagon ROMER Infinite 2.0 arm that provides sub-millimeter measurements. In this way, we can manually position the laser scanner to cover the fruits as good as possible. For more details on the LiDAR and the measuring arm, we refer to Schunck et al. [35]. For both scenarios, we designed a specific measurement procedure to align the measurements from the RGB-D sensor with the accurate ground truth point cloud of the LiDAR such that we have a pose in the coordinate frame of the LiDAR.

The recording setup in the lab scenario allows us to register the different point clouds from the RGB sensor and

¹Further details on downloading the dataset can be found at:

https://www.ipb.uni-bonn.de/data/shape_completion ²Our development toolkit including a data loader is available at:

https://github.com/PRBonn/shape_completion_toolkit



Fig. 2: Registration procedure for the lab scenario. Given the TSDFaligned RGB-D frames with the corresponding point cloud (on the left side) and the dense point cloud of the LiDAR (on the right side), we can estimate planes in each point cloud (shown in the middle). With an initial pose estimated via the extracted planes, we can register both point clouds using ICP automatically resulting in the final registration.

the LiDAR in a common reference frame. Here, we exploit a known environment structure to automatically determine the pose of the RGB-D sensor in the coordinate system of the LiDAR as described in Sec. III-A.

In the greenhouse scenario, we have to adapt the measurement procedure as we cannot scan the fruits with the high-precision scanner on the trees, but still want to capture the plant in a natural environment including the occlusions caused by leaves and other parts of the plants. Therefore, we designed a process that allows to associate the real scanned fruits with the measured fruit shapes of the ground truth point cloud, which we describe in Sec. III-B. Note that additional sensors such as an IMU can be used to improve the mapping accuracy. Despite having access to one within the camera, we did not use it as we found the mapping accuracy satisfactory.

A. Lab Scenario

In the lab scenario, we are able to record data of different fruits from more viewpoints compared to the greenhouse scenario. For each fruit, we collect the RGB-D frames by manually moving the camera around the fruit and trying to obtain as many viewpoints as possible. Afterwards, we repeat the process with the LiDAR covering the complete fruit. We want to highlight that it was not possible to scan the fruit with the different sensors at the same time, since the camera has a minimum range of 30 cm while the LiDAR operates at a closer range, i.e., 10 cm on average.

Our goal is twofold. First, we want to use the laser scanner data from the test set as ground truth, while using the RGB-D data as input. Second, we want to use the laser scanner data to allow learning shape priors from high-quality point clouds that can be helpful when working on noisy point clouds obtained with a RGB-D camera. To achieve this, we need to register the data coming from the two sensors into a single coordinate system. A high-level overview of the registration process is shown in Fig. 2.

To register the sensor data, we start by registering each RGB-D image belonging to the same fruit to each other with a standard TSDF fusion pipeline [8], [27], [39]. This provides us with locally consistent poses of each recorded RGB-D image. The results of this step are the camera poses for each RGB-D frame in a local coordinate system of the first RGB-D image and a mesh of the object extracted via marching cubes [19]. To register the point cloud of the LiDAR to the points of the mesh from the RGB-D sensor, we apply iterative closest point (ICP) [3], where we approximate the initial transformation between the two point clouds using the environment structure. To initialize the transformation between the point clouds, we scan the fruits inside three perpendicular planes and exploit this structure in the registration.

To determine the initial transformation, we extract the planes using a RANSAC approach which determines the three most dominant planes. To match the planes, we simply assume that for the first RGB-D frame we are approximately looking at the intersection point of the three planes. This results in three pairs of planes with normals $(n_i, n'_i), i = 1, 2, 3$, where we have normals $n_i \in \mathbb{R}^3$ from the planes in the first point cloud and $n'_i \in \mathbb{R}^3$ of the planes in the second point cloud. The rotation matrix $\mathbf{R} \in \mathbb{R}^{3\times 3}$ between the two sets with known associations between the normals (n_i, n'_i) can be calculated as follows [11]:

$$\boldsymbol{R} = \frac{\boldsymbol{n}_1'(\boldsymbol{n}_2 \times \boldsymbol{n}_3)^\top + \boldsymbol{n}_2'(\boldsymbol{n}_3 \times \boldsymbol{n}_1)^\top + \boldsymbol{n}_3'(\boldsymbol{n}_1 \times \boldsymbol{n}_2)^\top}{\det(\boldsymbol{n}_1^\top \quad \boldsymbol{n}_2^\top \quad \boldsymbol{n}_3^\top)},$$
(1)

where det(·) corresponds to the determinant of the matrix $(\mathbf{n}_1^{\top} \ \mathbf{n}_2^{\top} \ \mathbf{n}_3^{\top})$ with the normals \mathbf{n}_i as column vectors. The matrix \mathbf{R} is usually not an orthonormal matrix because of the measurement noise, but a true rotation matrix, can be computed using a singular value decomposition of \mathbf{R} .

The translation between the two sets of planes is the difference between the intersection point of the planes in the RGB-D point cloud and the intersection point of the planes in the LiDAR point cloud. After a fine-registration using ICP, we can compute the pose of every RGB-D frame in the coordinate system of the LiDAR point cloud.

B. Greenhouse Scenario

For the greenhouse scenario, we collected data in a greenhouse at Campus Klein-Altendorf near Bonn, Germany, which we provide as validation and test sets. As before, the collected data is comprised of measurements of two different sensors. First, RGB-D frames collected with an autonomous robot equipped with three Intel RealSense d435i cameras mounted on a robotic platform as described by Smitt et al. [36] and, second, high-precision scans of individual sweet



Fig. 3: Measuring procedure to align greenhouse RGB-D frames with the corresponding ground truth point cloud generated by the LiDAR. Given two recordings with and without markers, we first align the photogrammetric point clouds via a transformation T, which allows us to associate the sweet peppers without and with markers. Based on the markers and manually identified pins, we are able to determine the transformation of the scanned fruit in the RGB-D frame yielding the final registration.

peppers obtained with the aforementioned LiDAR as ground truth measurement.

In the lab scenario, we exploit the environment structure for aligning the data from different sensors. However, the same procedure is not possible as the sensors are located in different places. To register the individual sweet peppers in the RGB-D frames with their respective ground truth, we rely on a different measurement procedure.

We first perform a data collection run without any markers to get data as seen by a robot operating in the greenhouse. For getting the ground truth shape via high-precision point cloud using the LiDAR, we augment the environment with markers to associate the harvested fruits with the already collected data. See Fig. 3 for a visual impression of our greenhouse scenario and the measurement procedure.

More specifically, we manually tag individual sweet peppers in two ways: first, we place a label with an ID close to each fruit and, second, we insert multiple 3D markers, similar to paper pins, on each sweet pepper. The ID ensures that we can identify the fruits after harvesting and scanning them with the LiDAR. The attached 3D markers allow us to register the scanned fruits in the RGB-D frames. After augmentation with markers, we deploy our robot again to collect data with the markers attached to the sweet peppers.

After these procedures we obtain: (i) a clean set of RGB-D frames that will be used as inputs for the shape completion challenge, (ii) a tagged set of RGB-D frames that we use for registration purposes, and (iii) high-precision point clouds for each sweet pepper that represents our ground truth for each individual fruit. To register each fruit in the RGB-D frame they appear, we first conduct a 3D reconstruction for the complete plant row. The TSDF fusion pipeline of the RGB-D frames that we used for the lab scenario did not yield convincing results, due to larger noise in the depth measurements under natural lighting conditions. To solve such an issue, we opted for an estimation of camera poses via an RGB bundle adjustment using a commercial software

for photogrammetric 3D reconstruction [1]. Together with the estimated camera poses, we can further derive a dense point cloud and a 3D mesh.

We scaled the two reconstructions, i.e., reconstruction of the first clean recording run and the recording with the markers, to the real scene scale using the data from the wheel odometry of the robot. As we need the relative transformation between these two 3D reconstructions for transforming fruit poses from the tagged dataset to the original dataset, we apply ICP to the dense point clouds for alignment. The initial transformation for ICP is estimated by manually identifying a single pair of frames of the two recording runs which show the same part of the row and computing their relative pose.

For aligning a point cloud of a fruit from the LiDAR to a frame of the tagged dataset, we manually measure three corresponding 3D markers in the laser scan and on the image frame. Then we intersect the viewing rays constructed with the measured image coordinates with the reconstructed 3D mesh to get the 3D coordinates of the 3D markers in the tagged mesh. Given these 3D point correspondences, we can estimate the relative pose of the scanned fruit in the tagged dataset. As we know the transformation to the original dataset, we concatenate all transformations such that we end up in the pose of the scanned fruit in the RGB-D camera frame. Finally, we apply ICP between the 3D mesh and a single scanned fruit to improve the pose estimate further.

With this, we obtain the ground truth shape of each sweet pepper and also ground truth poses with respect to the fruit canonical pose, i.e., with the peduncle pointing upwards.

C. Curating the Benchmark Dataset

The objective of our dataset is to foster research in 3D shape completion and reconstruction in agricultural environments. To simplify the task, we concentrate on the shape reconstruction and completion task for a segmented fruit, thus, factor out the process of generating an instance

TABLE I: Statistics of the provided data. L indicates data collected in the lab, G in the greenhouse.

Split	#Images	#Fruits	Ground truth	Scenario
Training	4580	66	√	L
Validation	1387	25	√	L+G
Test	980	38	×	G

segmentation of a complete RGB-D frame and the pose estimation of the fruits.

To this end, we employed an instance segmentation approach trained on sweet peppers [36] to extract the parts of the RGB-D frame $\mathbf{I} \in \mathbb{R}^{H \times W \times 4}$, where each pixel $\mathbf{p}_i = (u, v)$ of the image frame of width W and height H corresponds to an RGB value (r_i, g_i, b_i) and it's corresponding metric depth $d_i \in \mathbb{R}$, that corresponds to the fruit \mathcal{F}_j , i.e., $\mathcal{F}_j = {\mathbf{p}_i | \mathbf{p}_i \in \mathbf{S}_j}$, where $\mathbf{S}_j \in {0, 1}^{H \times W}$ is the binary mask corresponding to the *j*-th fruit inside the image \mathbf{I} .

For each fruit point cloud \mathcal{F}_j in the RGB-D point cloud, we then apply the know transformation from the RGB-D image to the LiDAR point cloud to transform the points of \mathcal{F}_j into the LiDAR's canonical pose, i.e., with the peduncle pointing upwards. Thus, for each fruit we provide: (i) the original RGB-D data, \mathcal{P}_j , and (ii) the corresponding complete ground truth point cloud $\hat{\mathcal{P}}_j$ obtained with the LiDAR.

For reproducible experimental evaluation in our benchmark, we split the data of the two scenarios into train, validation, and test set, where we ensured that fruits in the validation and test set are spatially separated. We want to highlight that greenhouse data is not present in the training set. Such a setting introduces a domain gap enabling us to evalute generalization capabilities of tested algorithms. Tab. I provides a summary of the provided data for the task of shape completion of sweet peppers. Note that it is possible that the same fruit is visible from multiple viewpoints but share the same ground truth point cloud. We refer to Fig. 5 for an qualitative impression of the registration results and refer to our dataset website³ for further details on the dataset organization.

IV. SHAPE COMPLETION BENCHMARK

Together with the dataset, we provide a benchmark for the reproducible and unbiased evaluation of shape completion approaches in the agricultural domain. Therefore, we concentrate here on the shape completion in a setting of reconstructing the fruit in a canonical orientation, where we factor out the process of determining the location of the fruits and provide RGB-D information in canonical fruitspecific local reference frame as described in Sec. III-C. In our benchmark, we expect the estimation of a complete 3D mesh \mathcal{M} given partial observation of fruits \mathcal{P} provided by an RGB-D sensor. The desired mesh $\mathcal{M} = (\mathcal{V}, \mathcal{T})$, represented by N vertices $\mathcal{V} = {\mathbf{v}_1, \ldots, \mathbf{v}_N}$ and M triangles $\mathcal{T} = {(i, j, k) | i, j, k \in {1, \ldots, N}}$ defined as tuples with the corresponding vertex indexes, should represent a closed



Fig. 4: We show few example of input point clouds (left) and ground truth point clouds (right). The shape completion tasks involve estimating a complete 3D mesh from a partial and noisy point cloud.

surface of the complete fruit. Thus, the task can be described as the problem of finding the complete mesh \mathcal{M} of a given colored point cloud \mathcal{P} in a canonical frame, see Fig. 4.

A. Metrics

We report four different metrics: f-score, precision, recall, and Chamfer distance for evaluating 3D shapes, but rely on the f-score integrating precision and recall as our main metric. Precision, recall, and Chamfer distance are evaluated on point cloud \mathcal{R} generated by densely sampling of the reconstructed mesh \mathcal{M} and the corresponding point cloud \mathcal{G} of the LiDAR.

In line with Knapitsch et al. [18], we compute the precision $p(\rho)$, i.e., which is the proportion of the points of the reconstruction \mathcal{R} that are close to the ground truth \mathcal{G} and recall $r(\rho)$, i.e., the proportion of the point of the ground truth \mathcal{G} are close to the reconstructed mesh point cloud \mathcal{R} , where closeness is given by a threshold distance ρ :

$$p(\rho) = \frac{100}{|\mathcal{R}|} \sum_{\boldsymbol{r} \in \mathcal{R}} \left[\min_{\boldsymbol{g} \in \mathcal{G}} ||\boldsymbol{r} - \boldsymbol{g}|| < \rho \right], \quad (2)$$

$$r(\rho) = \frac{100}{|\mathcal{G}|} \sum_{\boldsymbol{g} \in \mathcal{G}} \left[\min_{\boldsymbol{r} \in \mathcal{R}} ||\boldsymbol{g} - \boldsymbol{r}|| < \rho \right], \quad (3)$$

where $\boldsymbol{g} \in \mathbb{R}^3$ and $\boldsymbol{r} \in \mathbb{R}^3$ are points from \mathcal{G} and \mathcal{R} , respectively. Here, the operator $[\![c]\!]$ is the Iverson bracket, i.e., if the condition c within the brackets is satisfied it evaluates to 1, otherwise to 0. Rather than choosing a fixed threshold ρ like Knapitsch et al. [18], we evaluate precision and recall at different ρ values, i.e., $\rho \in \{0.01 \text{ m}, 0.02 \text{ m}, \dots, 0.1 \text{ m}\}$ with interval $\Delta = 0.01 \text{ m}$. We compute the area under the curve for the precision \bar{p} and recall \bar{r} as follows:

$$\bar{p} = \frac{1}{\eta} \sum_{\rho \in \{0.01 \text{ m}, 0.02 \text{ m}, \dots, 0.1 \text{ m}\}} \Delta p(\rho)$$
(4)

$$\overline{r} = \frac{1}{\eta} \sum_{\rho \in \{0.01 \text{ m}, 0.02 \text{ m}, \dots, 0.1 \text{ m}\}} \Delta r(\rho),$$
(5)

where η is a normalizer that corresponds to the precision or recall of a perfect estimator evaluated with aforementioned thresholds to arrive at $\bar{p} \in [0, 1]$ and $\bar{r} \in [0, 1]$.

ī

³Our dataset website provides further details on the provided data: https://www.ipb.uni-bonn.de/data/shape_completion



Fig. 5: Qualitative example of our registration results in the lab (top) and greenhouse (bottom). Where we show the ground truth point cloud, in red, aligned with the corresponding RGB-D frames.

Finally, the f-score f is given by:

$$f = \frac{2\bar{p}\,\bar{r}}{\bar{p} + \bar{r}}.\tag{6}$$

We additionally evaluate the Chamfer distance $D_{\rm C}$. It is the average symmetric squared distance \bar{d}^2 of each point to its nearest neighbor in the other point cloud:

$$D_{\rm C}(\mathcal{G},\mathcal{R}) = \frac{\bar{d}^2(\mathcal{G},\mathcal{R})}{2} + \frac{\bar{d}^2(\mathcal{R},\mathcal{G})}{2},\tag{7}$$

with

$$\bar{d}^2(\mathcal{P}_i, \mathcal{P}_j) = \frac{1}{|\mathcal{P}_i|} \sum_{\boldsymbol{x}_i \in \mathcal{P}_i} \min_{\boldsymbol{x}_j \in \mathcal{P}_j} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2^2.$$
(8)

For the ranking of the entries submitted to our competition, we use the F1-score as a primary metric and the Chamfer distance as secondary metric in case of a tie.

B. CodaLab Competition

To evaluate predictions on the hidden test set, we expect a specific directory structure in our CodaLab competition [29]. We ask the participants to submit a zip file that contains one .ply file for each fruit id. Note that we expect the fruit to be in the canonical pose, i.e., with the peduncle pointing upwards, this is trivial to obtain given the provided pose for each RGB-D frame. More details can be found on the CodaLab submission page ⁴.

To speed up the usage of our dataset, we provide a development kit providing a PyTorch-based data loader and a small library for computing the metrics that we use in the CodaLab competition.

C. Baselines

To facilitate the comparisons with previous work, namely CoRe [23], HoMa [28], and T-CoRe [22], we report the performances on the test set in Tab. II. Additionally, we report the performances on the validation set in Tab. III to enable comparison of novel approaches in ablation studies using the validation set. From our experiments, T-CoRe provides better predictions on the test set reaching 5.06 mm of Chamfer distance and 58.0 % of f-score. In contrast, HoMa

⁴The CodaLab competition is available at: https://codalab.lisn.upsaclay.fr/competitions/18987

TABLE II: Baseline results on the test set. The \downarrow and \uparrow indicate that lower or higher values mean better performance.

Approach	$egin{array}{c} m{D_{C}} \ [mm] \ \downarrow \ avg \end{array}$	f [%] \uparrow avg	$ar{m{p}}$ [%] \uparrow avg	$ar{m{r}}$ [%] \uparrow avg
CoRe [23]	6.84	45.61	43.87	48.56
HoMa [28]	5.39	57.11	55.63	58.88
T-CoRe [22]	5.06	58.04	58.60	57.76

TABLE III: Baseline results on the validation set. The \downarrow and \uparrow indicate that lower or higher values mean better performance.

Approach	$m{D_{C}}$ [mm] \downarrow avg	f [%] \uparrow avg	$ar{m{p}}$ [%] \uparrow avg	$ar{m{r}}$ [%] \uparrow avg
CoRe [23]	7.18	42.48	38.92	47.52
HoMa [28]	4.19	65.65	63.58	68.01
T-CoRe [22]	4.69	60.61	60.59	60.82

is the best-performing method on the validation set, obtaining 4.2 mm of Chamfer distance and 65.7% of f-score. We believe this decrease in performance on the test set can be a symptom of overfitting on the recording conditions. Note that the validation set contains data from the lab and from the greenhouse, while the test set only contains samples collected in the greenhouse.

Finally, HoMa and T-CoRe yield better results than CoRe across all metrics, which is expected given that it is the only semi-supervised method. Notably, CoRe is the only approach operating on images, while T-CoRe and HoMa use point clouds. We believe this also contributes to the final performances.

V. CONCLUSION

In this paper, we presented a novel dataset for 3D shape completion of sweet peppers in agricultural environments. Our dataset consists of RGB-D frames collected in the lab and a commercial greenhouse, together with dense highprecision point clouds as ground truth of the sweet peppers, where we aligned the RGB-D data and the ground truth point cloud with a tailored measurement procedure. We additionally provide a CodaLab competition with a hidden test set to ensure fair comparisons and automatic evaluations to foster research into agricultural vision systems.

Our competion has already been used by more than 50 participants with a substantial increase in performances [7]. We believe this further demonstrates that the task used in the benchmark of this paper is not solved yet with plenty of room for improvements and at the same time the benefits of curating public benchmark with hidden test sets.

In future, we also plan to provide ground truth for other applications, such as 3D object pose estimation of fruits, where we leverage our measurement procedure to determine an accurate pose of visible sweet peppers based on the alignment between the marker point clouds and our ground truth measurements.

REFERENCES

 Agisoft LLC. Agisoft Metashape Professional. https://www. agisoft.com/. Accessed: April 2024.

- [2] R. Barth, J. IJsselmuiden, J. Hemming, and E. J. Van Henten. Data Synthesis Methods for Semantic Segmentation in Agriculture: A Capsicum Annuum Dataset. *Computers and Electronics in Agriculture*, 144:284–296, 2018.
- [3] P. Besl and N. McKay. A Method for Registration of 3D Shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 14(2):239–256, 1992.
- [4] P. M. Blok, E. J. van Henten, F. K. van Evert, and G. Kootstra. Image-based size estimation of broccoli heads under varying degrees of occlusion. *Biosystems Engineering*, 208:213–233, 2021.
- [5] 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.
- [6] 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.
- [7] Z. Chen, T. Wei, Z. Zhao, J. S. Lim, Y. Luo, H. Zhang, X. Yu, S. Chapman, and Z. Huang. Cf-prnet: Coarse-to-fine prototype refining network for point cloud completion and reconstruction. *arXiv preprint*, arXiv:2409.08443, 2024.
- [8] B. Curless and M. Levoy. A Volumetric Method for Building Complex Models from Range Images. In Proc. of the Intl. Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH), 1996.
- [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] M. Ferrer-Ferrer, J. Ruiz-Hidalgo, E. Gregorio, V. Vilaplana, J.-R. Morros, and J. Gené-Mola. Simultaneous Fruit Detection and Size Estimation Using Multitask Deep Neural Networks. *Biosystems Engineering*, 233:63–75, 2023.
- [11] W. Förstner and B. Wrobel. Photogrammetric Computer Vision Statistics, Geometry, Orientation and Reconstruction. Springer Verlag, 2016.
- [12] J. Gené-Mola, R. Sanz-Cortiella, J. R. Rosell-Polo, A. Escolà, and E. Gregorio. PFuji-Size Dataset: A Collection of Images and Photogrammetry-Derived 3D Point Clouds with gruth Annotations for Fuji Apple Detection and Size Estimation in Field Conditions. *Data in Brief*, 39:107629, 2021.
- [13] N. Häni, P. Roy, and V. Isler. MinneApple: a Benchmark Dataset for Apple Detection and Segmentation. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):852–858, 2020.
- [14] K. M. F. James, K. Heiwolt, D. J. Sargent, and G. Cielniak. Lincoln's Annotated Spatio-Temporal Strawberry Dataset (LAST-Straw). arXiv preprint, arXiv:2403.00566, 2024.
- [15] 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.
- [16] J. Kierdorf, L. V. Junker-Frohn, M. Delaney, M. D. Olave, A. Burkart, H. Jaenicke, O. Muller, 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.
- [17] J. Kierdorf, I. Weber, A. Kicherer, L. Zabawa, L. Drees, and R. Roscher. Behind the leaves: Estimation of occluded grapevine berries with conditional generative adversarial networks. *Frontiers* in Artificial Intelligence, 5:830026, 2022.
- [18] A. Knapitsch, J. Park, Q. Zhou, and V. Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Trans. on Graphics, 36(4):1–13, 2017.
- [19] W. Lorensen and H. Cline. Marching Cubes: a High Resolution 3D Surface Construction Algorithm. In Proc. of the Intl. Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH), pages 163–169, 1987.
- [20] 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)*, 37:20–34, 2020.
- [21] 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.
- [22] F. Magistri, R. Marcuzzi, E. Marks, M. Sodano, J. Behley, and C. Stachniss. Efficient and Accurate Transformer-Based 3D Shape Completion and Reconstruction of Fruits for Agricultural Robots. In Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA), 2024.

- [23] F. Magistri, E. Marks, S. Nagulavancha, I. Vizzo, T. Läbe, J. Behley, M. Halstead, C. McCool, and C. Stachniss. Contrastive 3d shape completion and reconstruction for agricultural robots using rgb-d frames. *IEEE Robotics and Automation Letters (RA-L)*, 7(4):10120– 10127, 2022.
- [24] F. Magistri, Y. Pan, J. Bartels, J. Behley, C. Stachniss, and C. Lehnert. Improving Robotic Fruit Harvesting Within Cluttered Environments Through 3D Shape Completion. *IEEE Robotics and Automation Letters (RA-L)*, 9(8):7357–7364, 2024.
- [25] E. Marks, J. Bömer, F. Magistri, A. Sah, J. Behley, and C. Stachniss. BonnBeetClouds3D: A Dataset Towards Point Cloud-based Organlevel Phenotyping of Sugar Beet Plants under Field Conditions. In Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2023.
- [26] M. Marzoa Tanco, G. Trinidad Barnech, F. Andrade, J. Baliosian, M. LLofriu, J. Di Martino, and G. Tejera. MAgro dataset: A Dataset for Simultaneous Localization and Mapping in Agricultural Environments. *Intl. Journal of Robotics Research (IJRR)*, 43(5):591– 601, 2024.
- [27] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinect-Fusion: Real-Time Dense Surface Mapping and Tracking. In Proc. of the Intl. Symposium on Mixed and Augmented Reality (ISMAR), 2011.
- [28] Y. Pan, F. Magistri, T. Läbe, E. Marks, C. Smitt, C. McCool, J. Behley, and C. Stachniss. Panoptic Mapping with Fruit Completion and Pose Estimation for Horticultural Robots. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2023.
- [29] A. Pavao, I. Guyon, A.-C. Letournel, D.-T. Tran, X. Baro, H. J. Escalante, S. Escalera, T. Thomas, and Z. Xu. CodaLab Competitions: An Open Source Platform to Organize Scientific Challenges. *Journal on Machine Learning Research (JMLR)*, 24(198):1–6, 2023.
- [30] R. Polvara, S. Molina, I. Hroob, A. Papadimitriou, K. Tsiolis, D. Giakoumis, S. Likothanassis, D. Tzovaras, G. Cielniak, and M. Hanheide. Bacchus Long-Term (BLT) Data Set: Acquisition of the Agricultural Multimodal BLT Data Set with Automated Robot Deployment. *Journal of Field Robotics (JFR)*, 2023.
- [31] I. Pérez-Borrero, D. Marín-Santos, M. E. Gegúndez-Arias, and E. Cortés-Ancos. A Fast and Accurate Deep Learning Method for Strawberry Instance Segmentation. *Computers and Electronics in Agriculture*, 178:105736, 2020.
- [32] 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. *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [33] I. Sa, Z. Ge, F. Dayoub, B. Upcroft, T. Perez, and C. McCool. Deepfruits: A fruit detection system using deep neural networks. *Sensors*, 16(8):1222, 2016.
- [34] 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, 2018.
- [35] 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.
- [36] C. Smitt, M. Halstead, T. Zaenker, M. Bennewitz, and C. McCool. PATHoBot: A robot for glasshouse crop phenotyping and intervention. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2021.
- [37] J. Weyler, F. Magistri, E. Marks, Y. 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 Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 46(12):9583–9594, 2024.
- [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] Q. Zhou, J. Park, and V. Koltun. Open3D: A modern library for 3D data processing. arXiv:1801.09847, 2018.

CERTIFICATE OF REPRODUCIBILITY

The authors of this publication declare that:

- 1) The software related to this publication is distributed in the hope that it will be useful, support open research, and simplify the reproducability of the results but it comes without any warrenty and without even the implied warranty of merchantability or fitness for a particular purpose.
- 2) *Federico Magistri* primarily developed the implementation related to this paper. This was done on Ubuntu 22.04.
- 3) *Jens Behley* verified that the code can be executed on a machine that follows the software specification given in the Git repository available at:

https://github.com/PRBonn/shape_completion_toolkit

4) *Jens Behley* verified that the experimental results presented in this publication can be reproduced using the implementation used at submission, which is labeled with a tag in the Git repository and can be retrieved using the command:

git checkout ICRA