

# Segmentation-Based 4D Registration of Plants Point Clouds for Phenotyping

Federico Magistri

Nived Chebrolu

Cyrril Stachniss

**Abstract**—Plant phenotyping, i.e., the task of measuring plant traits to describe the anatomy and physiology of plants, is a central task in crop science and plant breeding. Standard methods often require intrusive or time-consuming operations involving a lot of manual labor. Cameras or range sensors, paired with 3D reconstructions methods, can support phenotyping but the task yields several challenges in practice such as plant growth over time. In this paper, we address the problem of finding correspondences between plants recorded at different points in time to track phenotypic traits in an automated fashion. Our approach makes use of semantic segmentation and unsupervised clustering to compute keypoints from plant point clouds. We extract a compact representation of the considered scan that encodes both, topology and semantic information. Through our approach, we are able to tackle the data association problem for 4D point cloud data of plants effectively. We tested our approach on different 3D plus time, i.e., 4D, sequences of plant point clouds of different plant species. The experiments presented in this paper suggest that our 4D matching approach allows for non-rigid registration of the plants that change over time. Moreover, we show that our method allows for tracking different phenotyping traits at an organ level, forming a basis for automated temporal phenotyping.

## I. INTRODUCTION

High-resolution monitoring of plants is important in modern agriculture, for crop breeding, and in crop science in general. It forms the basis for analyzing the crop performance and can also provide an indicator for plant stresses. Measuring how individual plants develop and grow over time is a time-consuming process and may even require invasive methods that harm the crop. For example, today's standard approach to measuring leaf areas is to cut off the leaves and scan them with a flatbed scanner. New measurement technologies for measuring and tracking phenotyping features employing robots and robotic sensors can open up new possibilities to support farmers, breeders, and researchers to measure plant performances [5], [6].

In robotics, LiDARs and RGB-D cameras are common sensors to perceive static and dynamic environments in the context of SLAM [23] as well as many other tasks. Extending such approaches to the agricultural setting, however, is not straightforward. One of the challenges in this context is to develop techniques that can robustly deal with growing objects, changing appearance, the development of new organs causing changes in the topology as well as non-rigid deformations caused by external factors such as sunlight, wind, gravity




Fig. 1. 3D scans of the same plant recorded at different points in time. We aim at finding correspondences between temporally separated point clouds by computing meaningful keypoints to obtain a minimal representation of the considered scan to simplify data association and non-rigid registration.

and so forth. From our perspective, especially the task of finding correspondences over time between different scans of plants is challenging. Fig. 1 depicts an example of two 3D point clouds of plants recorded with a laser scanner as well as correspondence information.

However, by exploiting standard learning techniques, it is possible to compute a minimal representation of the plant to capture its topology and semantics. In the computer graphics community, the problem of computing a minimal representation of an object is often referred to as skeletonization and it is used to create animation adapting its textures based on an interpolation scheme. This is a non-trivial problem that has been extensively studied, see Tagliasacchi et al. [25]. Most skeletonization approaches also do not include any semantic information, while using only classification methods is often not sufficient to apply non-rigid registration techniques.

The main contribution of this paper is a fully automatic 4D data association algorithm of plant point clouds. We achieve this leveraging a keypoint extraction pipeline, based on successive learning stages. Starting with a classifier, we classify each point as a leaf or stem point. Then we apply a clustering approach to compute the individual leaf instances. Given the segmented point cloud, we employ self-organizing maps [12] to extract the keypoint of each plant. At this stage, it becomes possible to register point clouds at different points in time casting the task as a minimization problem and to track phenotyping traits such as leaf area index (LAI), leaf length, stem diameter, and stem length over time.

All authors are with the University of Bonn, Germany.

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 - 390732324 - PhenoRob.

In sum, our approach is able to (i) extract a set of keypoints encoding both, semantic and topology of different plant species at different growth stages, (ii) perform a temporal non-rigid registration of the point clouds effectively, and (iii) enables us to track different phenotyping traits at an organ level. These claims are backed up by the paper and our experimental evaluation.

## II. RELATED WORK

Registering point clouds is a common problem in robotics and variants of ICP are part of most laser-based or RGB-D-based SLAM systems [15], [23], [29]. Several approaches exist to register deformable objects such as humans [8], often assuming a given topology of the deformable object [9], [7], [19]. The problem becomes even more difficult if the object changes its appearance and topology over time.

A closely related approach to our work is the recently proposed one by Chebrolu et al. [3]. It uses a hidden Markov model for 4D data associations and registration but requires a given, minimal representation of the plants in the form of a consistent skeleton. This is different to the methods proposed in our paper, which computes this minimal representation explicitly. We address the computation of a minimal representation to register temporally separated point clouds, with the goal of encoding both the topology and the semantic of the plant.

For building such a representation, semantic information about plants is helpful. While the agricultural robotics community produced a large literature about classification algorithm for 2D images, the number of approaches in this domain that operate on 3D point clouds is rather small. Paulus et al. [17] propose an SVM-based classifier that relies on a surface histogram to classify each point in a 3D point cloud as leaf or stem. The recent approach by Zermas et al. [28] uses an iterative algorithm called randomly intercepted node to tackle the same problem. Sodhi et al. [21] use 2D images to extract 3D phytomers, namely fragments of the stem attached to a leaf for leaves detection. Shi et al. [20] propose a multi-view deep learning approach inspired by Su et al. [24] to address the organ segmentation problem, while Zermas et al. [27] uses a skeletonization approach to segment leaf instances. Our approach builds upon the work of Paulus et al. [17] and additionally groups the leaf points with an unsupervised clustering algorithm to extract leaf instances. In this way, we achieve an organ segmentation exploiting labeled data for the leaves. We use this in turn as the basis for our two-stage registration approach across plant point clouds.

The semantic information turns out to be useful to compute a plant skeletonization and the data association. Concerning skeletonization, Tagliasacchi et al. [25] propose the notion of a generalized rotational symmetry axis of an oriented set to extract skeletons from incomplete point clouds, while Huang et al. [11] adapted L1-medians locally to a point set to obtain similar skeletons. Wu et al. [26] uses a Laplacian point cloud contraction to extract skeletons of maize point clouds. We compute the skeleton points individually for each detected

organ based on a self-organizing map (SOM). This leads to an unsupervised approach for computing our keypoints.

The last step before being able to track phenotyping traits over time is to find correspondences between different point clouds recorded at different points in time. Li et al. [14] and Paproki et al. [16] analyze time-series point cloud data to detect topological events such as branching, decay and track the growth of different organs. The above-mentioned approach of Chebrolu et al. [3] performs a non-rigid registration based on an HMM model for finding the data association. In further work on UAV-based field analysis, Chebrolu et al. [2], [4] propose scale-invariant 2D descriptors to register crop rows in fields over time. Carlone et al. [1] estimate crown radius and height for each plant in a field, from a set of dense 3D reconstructions. In contrast, we use a two-step hierarchical approach. First, we match different organs to each other. In the second step, we register the keypoints inside the pair of organs. We solve this assignment task using the Hungarian method [13], a general method to assign a set of  $n$  jobs to a set of  $m$  machines. Although it was initially designed for problems with  $n = m$ , it can be generalized to non-squared problems [22].

## III. OUR APPROACH

Our approach takes as input a time-series of 3D point clouds from a plant and aims at finding correspondences between temporally separated scans. The correspondences then allow for registering the clouds in a non-rigid approach [3] and to track different phenotypic traits such as leaf area index, leaf length, stem diameter, and stem length over time. Our input data consist only of several unordered 3D point clouds, without inside-outside information nor RGB or intensity data. Our approach uses the following main steps, which are detailed in the sections below: (i) point cloud segmentation, (ii) plant keypoint extraction, and (iii) hierarchical matching over time.

### A. Point Cloud Segmentation

The first step of our approach for temporal registration of plant point clouds is to classify each point of a single point cloud as stem or leaf. To achieve this, we use a standard support vector machine (SVM) classifier extending each  $xyz$ -coordinates with the fast point feature histograms (FPFH) [18] yielding an  $m$ -dimensional feature vector. The FPFH computes a simplified histogram of directions for a query point and its neighbors, thus describing surface properties in a compressed form.

With such feature vectors as inputs, the SVM can classify each point of a plant point cloud into stem and leaf points. The SVM operates successfully, even in different growth stages. Considering only the points classified as leaf points, we perform a clustering step in order to find the individual leaves, i.e., the instances. We achieve this in two steps. For each plant scan, we apply the unsupervised clustering algorithm DBSCAN using only  $xyz$ -coordinates to obtain an initial segmentation. The second step is to discard small

clusters and to assign each discarded point to one of the remained clusters based on a  $k$ -nearest neighbor approach.

### B. Plant Registration

Once each point in a plant point cloud is assigned to one organ, namely to the stem or to one leaf instance, we can learn keypoints for each organ using self-organizing maps (SOM) [12]. Standard approach to skeletonization provides less reliable results when the geometry of the object presents inhomogeneities, instead SOMs can easily adapt to different geometries. SOMs are unsupervised neural networks using competitive learning instead of backpropagation. They take as input a grid that organizes itself to capture the topology of the input data. Given the input grid  $\mathcal{G}$  and the input set  $\mathcal{P}$ , the SOM defines a fully-connected layer between  $\mathcal{G}$  and  $\mathcal{P}$ . The learning process is composed of two alternating steps until convergence. First, the winning unit is computed as the  $\text{argmin}_i \|\mathbf{x} - \mathbf{w}_i\|$ , where  $\mathbf{x}$  is a randomly chosen sample from  $\mathcal{P}$  and  $\mathbf{w}_i$  is the weight vector more similar to  $\mathbf{x}$ , also called best matching unit. The second step consists of updating the weights of each unit according to  $\mathbf{w}_n = \mathbf{w}_n + \eta \beta(i) (\mathbf{x} - \mathbf{w}_i)$ , where  $\eta$  is the learning rate and  $\beta(i)$  a function, which weights the distance between unit  $n$  and the best matching unit.

In our case, the grid for each organ is an  $n \times 1$  chain of 3D points that will form the keypoints for that organ. The length of the chain  $n$  is proportional to the size of the organ, such that the keypoints are expected to have a minimum distance between 2 cm and 7 cm among each other. In this way, it is possible to obtain a skeleton-like structure for each plant that we considered. Fig. 2 shows a graphical explanation of the SOM.

Once the keypoints are computed for each organ, we aim at finding correspondences between pairs input point clouds. For determining the correspondences between the individual point clouds, we adopt a two-step registration approach. As a preprocessing step, we align the emerging point of the stem for both point clouds, in this way we have a rough initial guess to start our data association procedure. Here, we first find correspondences between whole organs at a low resolution. This is achieved by computing the best assignment between the mean of each organ in the scan recorded at  $t_i$  and the mean of each organ in the scan recorded at  $t_{i+1}$ , giving higher priority to order organs. We define a cost matrix  $\mathbf{M}$  as  $M_{i,j} = \|\mathbf{m}_i - \mathbf{n}_j\|$ , where  $\mathbf{m}_i$  is the mean of organ  $i$  in the first scan and  $\mathbf{n}_j$  is the mean of organ  $j$  in the second scan. Given this cost matrix, we can compute the optimal assignment using the Hungarian method by Kuhn [13].

Once we have the organ assignment, we perform a keypoint matching *inside each organ* to register the points of each organ with its corresponding ones in the other scan. The cost of an assignment within an organ is given by the cost matrix  $\mathbf{W}$ :

$$\mathbf{W}_{i,j} = \|\mathbf{p}_i - \mathbf{q}_j\|, \quad (1)$$




Fig. 2. Illustration of our approach for keypoint extraction using SOMs. An input chain (left) and a point cloud (right) are the inputs of the algorithm. After convergence, the chain arranges itself on the point cloud. The points of the chain on the cloud could form the keypoints inside each organ of the plants.

where  $\mathbf{W}_{i,j}$  is the Euclidean distance between node  $i$  in the first scan and node  $j$  in the second one. Then we minimize the following assignment problem:

$$\mathbf{X}^* = \underset{\mathbf{X}}{\text{argmin}} \|\mathbf{W} \circ \mathbf{X}\|_F \quad (2)$$

$$\text{s.t.} \quad \sum_i^n \mathbf{X}_{i*} = 1 \quad \text{and} \quad \sum_j^m \mathbf{X}_{*j} \leq 1, \quad (3)$$

where  $\circ$  refers to the Hadamard product, i.e., the element-wise matrix multiplication, and  $\|\cdot\|_F$  is the Frobenius norm of a matrix. The term  $\mathbf{X}$  is a selection matrix in which  $\mathbf{X}_{i,j} = 1$  if the match  $(i,j)$  contributes to the solution, otherwise  $\mathbf{X}_{i,j} = 0$ . Intuitively, the term  $\|\mathbf{W} \circ \mathbf{X}\|_F$  in Eq. (2) computes the sum of all individual assignment costs that contribute to the solution specified through the selection matrix  $\mathbf{X}$ . Additionally, the first constraint in Eq. (3) ensures that each node in the first scan is connected to only one node in the second scan. The second constraint in Eq. (3) guarantees that each keypoint in the second scan is connected to at most one node in the first scan. This minimization problem can also be solved optimally in cubic time in  $n$  using the Hungarian method. Here, we use the variant described by Stachniss [22] for multi-robot target allocation problems during exploration. Once we computed the data association, we apply the non-rigid registration pipeline described by Chebrolu et al. [3] but we replace its HMM-based assignment with the result  $\mathbf{X}^*$  given by Eq. (2).

### C. Computing Basic Phenotypic Traits

Once we aligned the scan, we can compute different phenotypic traits at an organ level. For both, the stem and leaf instances, we leverage the already computed keypoints. For the stem class, we compute stem diameter and stem length. For that, we assign each point in the cloud classified as stem to the closest keypoint. We then compute the main axis of each region using the standard singular value decomposition (SVD) approach. At this point, we can compute both, stem diameter and stem length with respect to the main axis of the considered region.

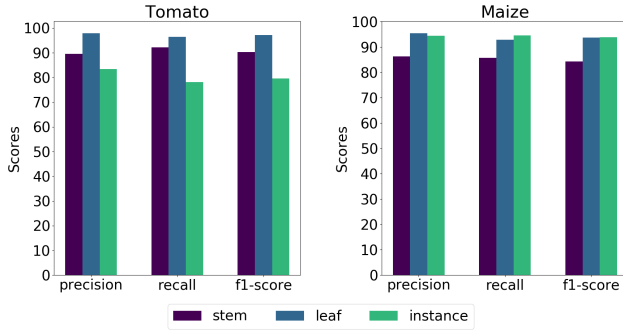


Fig. 3. Overall performances of our organ segmentation approach. The SVM is responsible for the stem and the leaf class, while instance refers to the unsupervised clustering of individual leaves. Best viewed in colors.

For stem diameter  $s_d$ , we use:

$$s_d = \sum_n \frac{1}{k_n} \sum_{k_n} \|\mathbf{p}_k - \mathbf{l}_n\|, \quad (4)$$

where  $n$  is the number of keypoints of the stem,  $k_n$  is the number of points in the point cloud assigned to the  $n$ -th keypoint, and  $\mathbf{l}_n$  is the main axis of the region computed with the SVD as explained above.

To compute the length of the stem, we first project each point in the  $n$ -th region onto the main axis and then compute the stem length  $s_l$  by:

$$s_l = \sum_n \max_{i,j} \|\mathbf{p}_j - \mathbf{p}_i\|. \quad (5)$$

To compute phenotypic traits associated with the leaf, we use a very similar approach. The LAI is defined as the green leaf area per unit ground surface area, namely  $LAI = A_l / \rho(A_l)$ . Where  $A_l$  is the leaf area and  $\rho(A_l)$  is its projection on the ground plane. To compute it, we first compute the leaf area  $A_l$ . Using again an SVD approach, we compute the main plane of the points associated with each node to obtain a 2D representation of the considered region that allows us to compute the leaf area  $A_l$ :

$$A_l = \sum_n \text{hull}(\mathbf{p}_n), \quad (6)$$

where  $\text{hull}(\mathbf{p}_n)$  represents the area of the convex hull computed from the points associated with the  $n$ -th keypoint. In a similar fashion, we compute  $\rho(A_l)$ , projecting all the points of the considered leaf instance on the ground plane and then compute the area of the resulting convex hull.

Finally, we compute the leaf length  $l_l$  as in Eq. (5) through:

$$l_l = \sum_n \max_{i,j} \|\mathbf{p}_j - \mathbf{p}_i\|. \quad (7)$$

Other geometric traits that build upon the area or shape of the leaves or stem can be computed as well based on the segmented, labeled, and aligned clouds.

#### IV. EXPERIMENTAL EVALUATION

##### A. Datasets

The experiments are designed to illustrate the capability of our approach, assess its performance, and support the

TABLE I  
TOMATO MATCHING ACCURACY.

	Precision	Recall	F1-Score
Chebrolu et al. [3] (HMM)	88.30	66.94	76.15
Hungarian w/o semantics	80.01	<b>100.00</b>	88.89
Our approach (Hung. w/ semantics)	<b>93.56</b>	<b>100.00</b>	<b>97.73</b>

TABLE II  
MAIZE MATCHING ACCURACY.

	Precision	Recall	F1-Score
Chebrolu et al. [3] (HMM)	88.70	85.45	87.04
Hungarian w/o semantics	86.89	<b>100.00</b>	92.98
Our approach (Hung. w/ semantics)	<b>97.50</b>	<b>100.00</b>	<b>98.76</b>

claims made at the end of the introduction of this paper. We evaluate our approach on two different 3D time-series datasets of different plant species, on tomato plants (*solanum lycopersicum*) and maize plants (*zea mays*). Both datasets were recorded over more than 10 days using a laser scan mounted over a Romer Absolute Arm, which is manually moved around the plants to obtain high accurate scans. This dataset covers the growth of such plants in an early stage, showing substantial changes in both appearance and topology. The dataset used in this paper can be found at: <https://www.ipb.uni-bonn.de/data/4d-plant-registration/>

##### B. Point Cloud Segmentation

The first experiment is designed to illustrate and assess the performance of our combined classification and clustering approach and to support the claim that it is well-suited for the organ segmentation of different plant species and different growth stages. For each temporal sequence, we randomly sample two scans as the training set and the rest as the test set. Our input is an  $n \times 36$  feature vector, which is composed of the  $xyz$ -coordinates and the FPFH descriptors. We use a radial basis function as the kernel for the SVM classifier. After performing the classification, we apply DBSCAN clustering on the leaf class using the Euclidean distance as metrics to obtain a complete organ segmentation. For the leaf instances segmentation task, we compute the assignment between predictions and ground truth as described in Sec. III-A. We evaluate the performances of our proposed approach computing standard metrics such as precision, recall and f1-score on the temporal sequences. The results are given in Fig. 3 and a visual example of our complete segmentation approach is shown in Fig. 4, best viewed in color.

##### C. Plant Registration Over Time

The second experiment is designed to support the claim that with our approach it is possible to achieve a better data association and, thus, a better non-rigid registration between temporally separated point clouds than with the approach described by Chebrolu et al. [3]. As the first evaluation, we compute the accuracy of our matching pipeline against the



Fig. 4. Sample results of our segmentation approach. In the top row, results on a tomato plant at different growth stages, in the bottom row results on a maize plant. The brown color represents the stem class, while different shades of green represent leaf instances. Best viewed in colors.

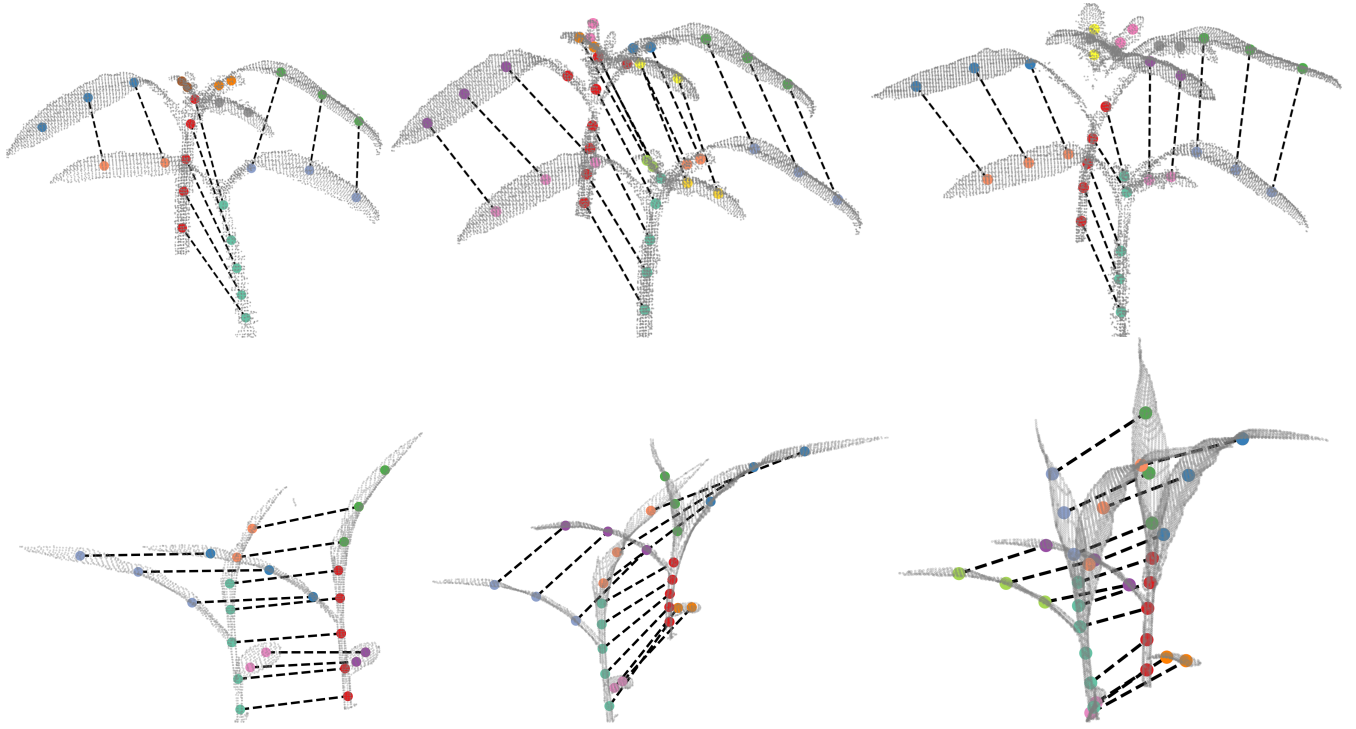


Fig. 5. Results of our complete matching pipeline. Different colors indicate the semantic of the keypoints. Our approach can deal with substantial growth, emerging of new organs and anomalies such as missing data and extreme leaf bending. Best viewed in colors.

correspondences found by a hidden Markov model approach using precision, recall and f1-score. We define true positives as the number of correct matches, false positives as the number of wrong matches, and false negatives as the number of non-assigned keypoints in the first scan. A summary of the performance given these metrics is shown in Tab. I for the tomato sequence and in Tab. II for the maize sequence. The tables also provide a comparison of the accuracy of our proposed approach against the matching algorithm based on the HMM and, as ablation study of our approach, we also show the results of the Hungarian method-based matching without exploiting semantic information (labeled as “Hun-

garian w/o semantics”). As can be seen, our approach clearly outperforms alternative solutions. The perfect recall of 100% of our approach is a result of our Hungarian matching described in Sec. III-B. Stating the matching problem as an optimization task including the semantic information performs best.

As second evaluation, we register the point clouds  $\mathcal{P}_i$  against the point cloud  $\mathcal{P}_j$  in a non-rigid manner using our approach as well as the one proposed by Chebrolu et al. By comparing the registration error of both approaches we can make a qualitative statement about the ability to register point clouds of plant taken at different point in time. We indicate



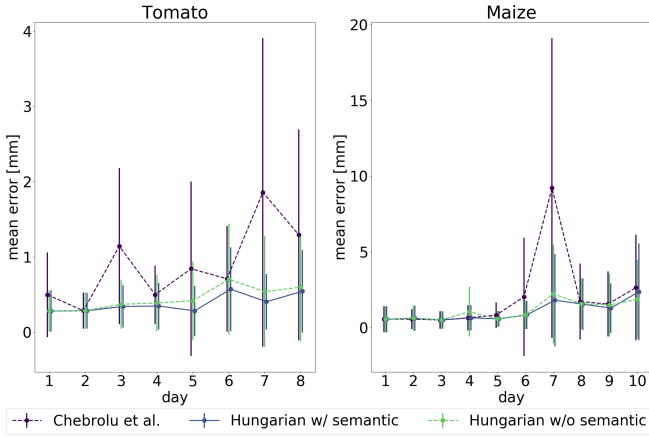


Fig. 6. Mean error and standard deviation of the Hausdorff distance. Bars represent the  $2\sigma$  standard deviation. Our matching approach *Hungarian w/ semantics* results to be the most stable compared to *Hungarian w/o semantics*, i.e. the same method without considering any semantic information, and Chebrolu et al., i.e. the matching found by a HMM approach. Best viewed in colors.

as  $\mathcal{P}_{j|HMM}$  the results of the non-rigid registration described by Chebrolu et al. [3] given the HMM-based matching, whereas,  $\mathcal{P}_{j|our}$  refers to the result of the non-rigid registration when applied to the matches found by our approach, i.e Hungarian method plus semantic information. We also compute  $\mathcal{P}_{j|hung}$ , i.e. the results of the registration given the correspondences computed by the Hungarian method without considering any semantic information.

We can estimate the quality of our matching approach computing the Hausdorff distance between the transformed point clouds and the original point cloud  $\mathcal{P}_j$ :

$$d(\mathbf{p}_j, \mathcal{P}_{j|*}) = \min_{\mathbf{p}' \in \mathcal{P}_{j|*}} \|\mathbf{p}_j - \mathbf{p}'\| \quad (8)$$

Fig. 6 depicts the results for all matching approaches. For the considered temporal sequences, we plot the mean error of the Hausdorff distance as well as its standard deviation. We see that the mean error of the baseline is greater or equal than for our approach. Moreover, the standard deviations of our approach show smaller fluctuations compared to the HMM-based method. It is also clear that including the semantic information in the matching pipeline yields, on average, to more realistic transformations. Fig. 7 shows an illustration of the two matching algorithms, in the first column is shown  $\mathcal{P}_{j|HMM}$ , in the second  $\mathcal{P}_j$ , while in the third  $\mathcal{P}_{j|our}$ . As can be seen also qualitatively, our approach produces a better alignment.

#### D. Phenotypic Traits

The third experiment is designed to illustrate the potential of our approach to derive phenotypic traits. Exploiting our complete pipeline, namely the organ segmentation followed by the data association between temporally separated scans, we are able to automatically compute a diverse number of plant traits, such as the leaf area index, leaf length, stem diameter, and stem length. In Fig. 8, we show the tracking of such traits over the acquisition period. For visualization

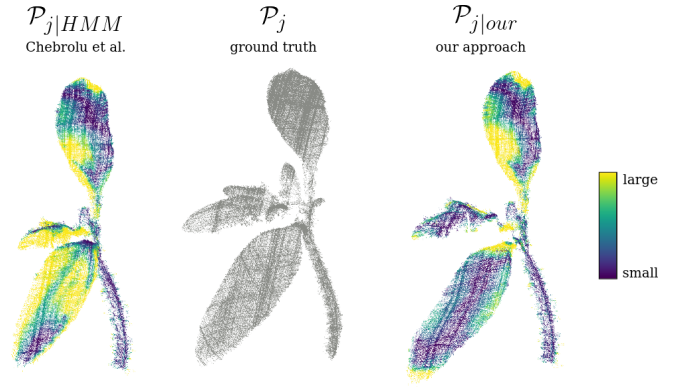


Fig. 7. A qualitative overview of the two matching algorithms, brightest regions correspond to highest errors. Our matching approach results to be most accurate compared to Chebrolu et al. Best viewed in colors.

reasons, we plot the leaf area instead of the LAI. However, the latter can be easily derived as described in Sec. III-C by computing the area of the convex hull associated with the projection of the considered leaf instance on the ground plane. We can observe similar behaviors in leaves that belongs to the same canopy region, as an example the first two emerging leaves in both tomato and maize plant show a lower growth rate with respect to the other leaves, likely due to the fact that such leaves capture less light than the ones at the top of the canopy. Based on the leaf information, our approach can also be used to obtain the BBCH index [10], which is used to describe the phenological development stages of plants.

#### V. CONCLUSION

In this paper, we presented a novel approach to compute data association between temporally separated plant point clouds subject to plant growth. We proposed a two-stage registration approach to tackle the organ segmentation task. It consists of a stem vs. leaf classification, followed an unsupervised clustering for the leaf instance estimation task. Given the segmented organs, we can extract plant keypoints, encoding both, semantic and topology information based on an unsupervised approach using self-organizing maps. This allows us to successfully develop a two-stage hierarchical matching algorithm that first matches organs to organs and then performs the registrations inside the organs. Our approach computes the best assignment between scans recorded at different points in time for our given cost matrix using the Hungarian method.

We implemented and evaluated our approach on two different plant species, tomato, and maize plants, and provided comparisons to other existing techniques. Furthermore, we supported all claims made in this paper through our evaluation. The experiments suggest that our approach yields a better data association between plants point cloud at different growth stages and allows for computing multiple phenotypic traits of the plant at each point in time.

Despite these encouraging results, there is further space for improvements. For example, some of the recent deep




Fig. 8. Growth estimation: the days are plotted on the x-axis, on the y-axis we show our estimations expressed in *mm*. The first two rows show the growth estimation, respectively, of leaf area and leaf length, where each color represents a leaf instance. The last row shows the stem parameters estimation, for visualization reasons the diameter values are multiplied by a factor of 10. Best viewed in colors.

learning methods might be better suited in more challenging scenes with serious overlaps or when the point clouds are not complete. Moreover, we plan to extend our work to deploy it in crop fields directly on a mobile platform, thus going beyond greenhouse conditions.

#### ACKNOWLEDGMENTS

We thank Heiner Kuhlmann and his group at the University of Bonn as well as Stefan Paulus for providing the 3D scans of the plants used in this paper. We also thank Lasse Klingbeil and Christopher McCool for fruitful discussions.

#### REFERENCES

- [1] L. Carlone, J. Dong, S. Fenu, G.G. Rains, and F. Dellaert. Towards 4d crop analysis in precision agriculture: Estimating plant height and crown radius over time via expectation-maximization. In *ICRA Workshop on Robotics in Agriculture*, 2015.
- [2] N. Chebrolu, T. Labe, and C. Stachniss. Robust long-term registration of uav images of crop fields for precision agriculture. *IEEE Robotics and Automation Letters*, 3(4):3097–3104, 2018.
- [3] N. Chebrolu, T. Labe, and C. Stachniss. Spatio-temporal non-rigid registration of 3d point clouds of plants. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.
- [4] N. Chebrolu, P. Lottes, T. Labe, and C. Stachniss. Robot localization based on aerial images for precision agriculture tasks in crop fields. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 1787–1793, 2019.
- [5] F. Fiorani, U. Rascher, S. Jahnke, and U. Schurr. Imaging plants dynamics in heterogenic environments. *Current opinion in biotechnology*, 23(2):227–235, 2012.
- [6] F. Fiorani and U. Schurr. Future scenarios for plant phenotyping. *Annual review of plant biology*, 64:267–291, 2013.
- [7] J. Gall, C. Stoll, E. De Aguiar, C. Theobalt, B. Rosenhahn, and H.P. Seidel. Motion capture using joint skeleton tracking and surface estimation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1746–1753, 2009.
- [8] D. Haehnel, S. Thrun, and W. Burgard. An extension of the icp algorithm for modeling nonrigid objects with mobile robots. In *Proc. of the Intl. Conf. on Artificial Intelligence (IJCAI)*, volume 3, pages 915–920, 2003.
- [9] L. Herda, P. Fua, R. Plankers, R. Boulic, and D. Thalmann. Skeleton-based motion capture for robust reconstruction of human motion. In *Proceedings Computer Animation 2000*, pages 77–83, May 2000.
- [10] M. Hess, G. Barralis, H. Bleiholder, L. Buhr, T.H. Eggers, H. Hack, and R. Stauss. Use of the extended bbch scale general for the descriptions of the growth stages of mono- and dicotyledonous weed species. *Weed Research*, 37(6):433–441, 1997.
- [11] H. Huang, S. Wu, D. Cohen-Or, M. Gong, H. Zhang, G. Li, and B. Chen. L1-medial skeleton of point cloud. *ACM Transactions on Graphics*, 32(4):65–1, 2013.
- [12] T. Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990.
- [13] H.W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [14] Y. Li, X. Fan, N.J. Mitra, D. Chamovitz, D. Cohen-Or, and B. Chen. Analyzing growing plants from 4d point cloud data. *ACM Transactions on Graphics*, 32(6):157, 2013.
- [15] E. Palazzolo, J. Behley, P. Lottes, P. Giguere, and C. Stachniss. ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [16] A. Paproki, X. Sirault, S. Berry, R. Furbank, and J. Fripp. A novel mesh processing based technique for 3d plant analysis. *BMC plant biology*, 12(1):63, 2012.
- [17] S. Paulus, J. Dupuisand A. Mahlein, and H. Kuhlmann. Surface feature based classification of plant organs from 3d laserscanned point clouds for plant phenotyping. *BMC Bioinformatics*, 14(1):238, 2013.
- [18] R.B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fph) for 3d registration. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 3212–3217, 2009.
- [19] L.A. Schwarz, A. Mkhitarian, D. Mateus, and N. Navab. Human skeleton tracking from depth data using geodesic distances and optical flow. *Image and Vision Computing*, 30(3):217 – 226, 2012.
- [20] W. Shi, R. van de Zedde, H. Jiang, and G. Kootstra. Plant-part segmentation using deep learning and multi-view vision. *Biosystems Engineering*, 187:81–95, 2019.
- [21] P. Sodhi, S. Vijayarangan, and D. Wettergreen. In-field segmentation and identification of plant structures using 3d imaging. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5180–5187, 2017.
- [22] C. Stachniss. *Exploration and Mapping with Mobile Robots*. PhD thesis, University of Freiburg, Department of Computer Science, 2006.
- [23] C. Stachniss, J. Leonard, and S. Thrun. *Springer Handbook of Robotics, 2nd edition*, chapter Chapt. 46: Simultaneous Localization and Mapping. Springer Verlag, 2016.
- [24] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proc. of the IEEE Intl. Conf. on Computer Vision (ICCV)*, pages 945–953, 2015.
- [25] A. Tagliasacchi, T. Delame, M. Spagnuolo, N. Amenta, and A. Telea. 3d skeletons: A state-of-the-art report. In *Computer Graphics Forum*, volume 35, pages 573–597, 2016.
- [26] S. Wu, W. Wen, B. Xiao, X. Guo, J. Du, C. Wang, and Y. Wang. An accurate skeleton extraction approach from 3d point clouds of maize plants. *Frontiers in plant science*, 10, 2019.
- [27] D. Zermas, V. Morellas, D. Mulla, and N. Papanikolopoulos. Estimating the leaf area index of crops through the evaluation of 3d models. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 6155–6162, 2017.
- [28] D. Zermas, V. Morellas, D. Mulla, and N. Papanikolopoulos. Extracting phenotypic characteristics of corn crops through the use of reconstructed 3d models. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 8247–8254, 2018.
- [29] M. Zollhofer, P. Stotko, A. Gorlitz, C. Theobalt, M. Niener, R. Klein, and A. Kolb. State of the Art on 3D Reconstruction with RGB-D Cameras. In *Eurographics - State-of-the-Art Reports (STARs)*, volume 37, 2018.