Suckers Emission Detection and Volume Estimation for the Precision Farming of Hazelnut Orchards

Ciro Potena*, Renzo Fabrizio Carpio, Nico Pietroni, Jacopo Maiolini, Giovanni Ulivi, Emanuele Garone, Andrea Gasparri

***Abstract*— In this work, inspired by the needs of the H2020 European Project PANTHEON[1], we address the hazelnut sucker detection and canopy volume estimation problem on a per-plant basis. Sucker control is an essential but challenging practice in agriculture, given the fact that suckers, i.e., shoots that grow from the tree roots, compete with the tree itself for water and nutrients. This research is motivated by the observation that in current best-practice, sucker control is carried out by applying a non-calibrated amount of chemical inputs to each tree. Indeed, a proper sucker detection and estimation algorithm would represent the enabling technology for an environmentally friendly sucker control approach where the amount of herbicide could be properly calibrated according to the needs of each individual plant. In this work, we propose an end-to-end algorithm for detecting the presence of suckers and for estimating their canopy. First a sparse point cloud-based representation of the suckers is detected, then an approximated canopy estimation is achieved by means of a tailored meshing strategy that performs a leaf-based clustering and an iterative clusters connection. The volume is then estimated by the resulting mesh. Preliminary real-world experiments are provided to corroborate the effectiveness of the proposed canopy estimation strategy.**

I. Introduction

Precision Agriculture is a farming management concept and is based on observing, measuring, and responding to inter and intra-field variability in crops [1]. Modern hazelnut farming is generally carried out using regular layouts to allow the mechanization of many field operations, as it is the use of tractors and specialized agricultural machinery. However, there are still time-consuming and labor-intense agronomic activities that could be clearly improved. In this regard, the European project PANTHEON focuses on the management of large-scale hazelnut orchards, where, to the best of our knowledge, precision farming techniques have not been investigated yet. Among the main field management activities, the sucker control is one of the most labor-intensive.

Nowadays, sucker control in large-scale orchards is commonly carried out by tractors with a pump spraying herbicide with no distinction between a plant that needs the treatment and a plant that does not require it. Additionally, a non-calibrated amount of herbicide is applied to all plants. In this context, the scope of PANTHEON is to design a solution to individually treat each tree according to its needs.

C. Potena, R. F. Carpio, J. Maiolini, G. Ulivi, A. Gasparri are with Roma Tre University, Rome, Italy, 00146. N. Pietroni is with University Of Technology Sydney. Emanuele Garone is with the with the Université Libre de Bruxelles, 1050 Brussels, Belgium. This work has been supported by the European Commission under the grant agreement number 774571 – Project PANTHEON. *Corresponding author e-mail: *cpotena@os.uniroma3.it*.

[1] http://www.project-pantheon.eu

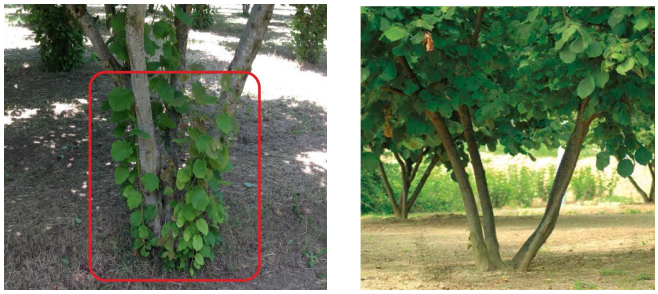

Fig. 1: Example of a hazelnut tree with suckers, in the red square, (left) and a hazelnut tree after suckers removal (right).

The sucker control solution proposed within PANTHEON is composed of two steps: i) in the first step, the characterization and estimation of sucker canopy dimensions of every plant is carried out, and ii) in the second step, tailor-made treatments are computed based on this information and performed to allow the application of different rates of herbicide to each tree. This innovative, environmentally friendly solution is expected to reduce herbicide volumes and enhance the health conditions of the plants.

In this work we focus on the first step of the suckering control solution proposed within the project PANTHEON. To achieve this objective, we propose a method that relies on visual and 3D sensory data. Our approach is based on two major components: (i) a sucker detector and (ii) a sucker canopy estimator. More specifically, the former locates suckers from visual data. In this context, the extremely variable light conditions in open environments play a key role and pose important challenges in the sucker detection process. We achieve a robust and reliable detection performance by fine-tuning a state-of-the-art object detection system with on-field gathered data. The outcome of the sucker detection process is then used to obtain a sparse point-cloud-based representation of the sucker by filtering out redundant 3D information coming from a LiDAR sensor.

The second component, the canopy estimator, starts from the resulting sparse point-cloud and converts it into a triangle mesh to perform an exact volumetric estimation. Although the reconstruction of a 3D mesh from a point cloud is a widely investigated area [2], state-of-the-art methods might perform poorly (see Figure 2) with open-air vegetation. This can be explained by the fact that open-air vegetation shows extremely varying shapes, and it is challenging to acquire using scanning devices usually resulting into a quite sparse point cloud. Hence, we derived a novel strategy to compute

a volumetric approximation, which is designed for the specific characteristics of suckers. In particular, we exploit a leaf-based 3D points clustering and an iterative clusters connection process to derive an approximated mesh. Notably, the design of this novel strategy is made even more challenging by the fact that we intentionally exploited a LIDAR sensor originally designed for navigation purposes, thus with a significantly lower resolution compared to classical LIDAR solutions designed for 3D reality capture. Indeed, this choice was made by observing that reasonably such a LIDAR sensor should already be available on an autonomous vehicle, thus limiting the impact on the economic cost of the proposed sucker control solution. The whole system eventually results in an end-to-end estimation process that can potentially provide benefits to a large variety of canopy estimation applications.

To validate the proposed approach, we report preliminary results from real-world data where the ground-truth volume has been computed through the Archimedes method [3]. We show that the proposed method represents a promising solution for achieving an adequate volumetric estimation that could be used for calibrating the amount of chemical treatment according to the needs of each single plant. We also provide an open-source annotated dataset for sucker detection at the following link

`https://tinyurl.com/y68wdz5j`

A. Related Work

The problem of detecting suckers and estimating their canopy has not been extensively investigated over the years, and relatively few solutions for different cultivars can be found at the state of the art [4], [5]. Kang *et al.* [4] presented a sucker detection system consisting of a laser scanner for vine detection and a color camera for imaging suckers. The main purpose was to trigger a targeted spray for vine-specific sucker control in grape vineyards. In particular, by using a camera installed below the laser scanner, this system would allow to calculate the diameter and the position of the vine trunk. Furthermore, to facilitate the separation between sucker and background the authors used white boards to block background noise. The results of this work show a high level of accuracy. Wang *et al.* [5] presented an identification and location algorithm of vine suckers for real-time, targeted, spraying activities based on a sensor fusion among 2D laser scanners and a camera. The sucker position estimation is performed in the camera image plane by using white calibration board to determine the invisible laser scanning line. The reported results show an average computational time lower than 1 second, while a recognition rate about 90%. Although these methods presented above show promising results, they make use of external components to aid the estimation process. This solution might not be adequate for large-scale orchards.

A branch of the literature close to sucker canopy estimation is the automated plant phenotyping for which a full 3D reconstruction of the crop is essential to measure the plant growth and to estimate the yield. This estimation process is carried out with different sensor modalities, and several approaches have been proposed in the last years [6]–[10]. Paulus *et al.* [6] presented a high throughput phenotyping of barley organs based on 3D laser scanning. This work showed the automatic parameter tracking of the leaves and stem of a barley plant on a temporal scale, with the final goal of growth monitoring in a time course for barley plants. Nguyen *et al.* [7] implemented a full 3D reconstruction system for plant phenotyping, which utilized simultaneous, multi-view, high-resolution color digital imagery. The cameras were mounted on an arc-shaped super structure and organized into 16 stereo pairs in four separate arcs. The results revealed that both can yield satisfactory 3D reconstruction models and are suited for high-throughput phenotyping without destroying any leaves or stems of the plant. Pound *at al.* [10] an approach for the 3D plant reconstruction by means of a single low-cost camera. The reconstruction pipeline exploits a series of small planar sections to model more complex leaf surfaces. The output of this work is a 3D mesh structure that is suitable for modelling applications, in a format that can be imported in most of 3D graphics and software packages. Shim *et al.* [8] presented a multi-view acquisition system composed of Time of Flight (ToF) and RGB color cameras. The proposed system compensates for errors in the acquisition process through an ad-hoc calibration procedure and is well-suited for the 3D reconstruction of dynamic scenes. In Dong et al. [9], the authors implemented a method which exploits the structured environment. Specifically, this method utilizes global features and semantic information to obtain an initial solution aligning the two sides of orchard rows. The 3D model is then refined by using semantic information extracted by using a robust detection and fitting algorithm. The vision system could be used to measure semantic traits from the optimized 3D model in order to obtain an estimate of canopy volume, trunk diameter, tree height and fruit count. Although these works described above represent promising solutions for the automated plant phenotyping, they are focused on the reconstruction of the entire tree, while in our work, inspired by the needs of the H2020 project PANTHEON, we proposed a tailored approach for the detection of suckers and estimation of their canopy volume.

B. Contributions

Briefly, the contributions made in this work can be summarized as follows:

(i) an end-to-end canopy estimation algorithm;
(ii) a novel meshing strategy for open-air vegetation;
(iii) publicly available annotated data for sucker detection.

II. Problem Statement and Assumptions

In this section, first the mathematical formulation of the proposed sucker detection and canopy estimation problem is provided, then some working assumptions are discussed.

Problem 1. Given a 3D sparse point-cloud set $\mathcal{S}$ representing a sucker (Figure 2, left column), our goal is to find a function $F : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ that computes a volumetric estimation of the sucker.

Notably, the set $\mathcal{S}$ can be either collected by one or more 3D sensors (e.g. LiDARs and RGBD cameras) or generated in a post-processing phase by a photogrammetry-based 3D reconstruction software. In this work we aim to perform volumetric estimations in real-time, thus collecting data from 3D sensors.

To derive an effective method, we made the following assumptions on the data collection process and on the structural properties of the sucker canopy:

$i)$ The 3D point-cloud set $\mathcal{S}$ mainly belongs to the target sucker;

$ii)$ The leaves are connected through minimal paths.

Assumption $i)$ is an essential data requirement, and is fulfilled by extracting 3D points inside the "Region Of Interest" (ROI) surrounding the target sucker. In this regard, we also perform a further outlier rejection refinement step by pruning out all those points that might belong to the soil terrain. Assumption $ii)$ is a reasonable simplification assumption since in a real-world scenario the different branches of a plant are individually linked to a single connected component. Indeed, Assumption $ii)$ is explicitly exploited for the sucker canopy volume estimation.

III. Sucker Detection

In this section we focus on the sucker detection problem, i.e., the problem of extracting a set of Regions Of Interest (ROIs) $\mathcal{B} = (B_0, B_1, \ldots, B_i)$ (see Figure 1 for an example of ROI) from the input images, such that they bound the suckers to inspect. As a matter of fact, detecting suckers in an open-environment present two major challenges: (i) the background noise (e.g. grass, or other plants [4]) and (ii) the extremely variable light conditions.

We tackle these issues by fine-tuning the state-of-the-art convolutional neural network YOLOv3 [11]. Notably, this choice was made as the YOLOv3-tiny architecture (depicted in Table I) represents as a trade-off beetween computational time and accuracy. Briefly, the YOLOv3-tiny architecture is composed of successive 3×3 and 1×1 convolutional layers, and it accepts input images with size 608×608.

In particular, the fine tuning is performed on a pre-trained model, trained on the ImageNet [12] dataset, with on-field gathered datasets. To achieve a reliable sucker detection, the datasets have been gathered in different weather and daylight conditions. The accuracy of the proposed neural network architecture will be discussed in section V-B.

IV. Sucker Volume Estimation

In this section, we focus on the estimation of the sucker canopy volume. Note that, this process cannot be performed by solely relying on visual data. Therefore, we exploit sparse 3D information contained in the point cloud $\mathcal{S}$. To map such points to their respective pixel coordinates in the camera image plane, we make use of the calibration procedure described in [13]. Points that do not belong to the sucker are detected and removed by exploiting the ROIs determined through the detection process described in Section III. At the end of this process, it is very likely that a non-negligible portion of the remaining 3D points does not belong to the sucker leaf surface (*e.g.* tree trunk or soil terrain). This could lead the volume estimation algorithm to yield poor results. Therefore, a preliminary step to filter out misleading data is mandatory. We achieve this goal by means of Excess Green (ExG) [14] index threshold operation. The ExG index is an image color transformation process which is commonly adopted in agronomic image analysis since it highlights the presence of vegetation through a green color enhancement. In our case, the exploitation of such index allows to easily distinguish between the sucker canopy and the soil and the trunk of the hazelnut tree. Despite its simplicity, this method can effectively remove most of the outliers. Future work will focus on the investigation of more refined strategies for outliers processing and removal.

Type	Filters	Size	Output
Convolutional	$16 \,/\, 2$	3×3	304×304
Convolutional	$32 \,/\, 2$	3×3	152×152
Convolutional	$64 \,/\, 2$	3×3	76×76
Convolutional	$128 \,/\, 2$	3×3	38×38
Convolutional	$256 \,/\, 2$	3×3	19×19
Convolutional	$512 \,/\, 2$	3×3	10×10
Convolutional	$1024 \,/\, 2$	3×3	10×10
Convolutional	$256 \,/\, 2$	1×1	10×10
Convolutional	$512 \,/\, 2$	3×3	10×10
Convolutional	$256 \,/\, 2$	1×1	10×10
Avgpool		Global	
Connected		1000	
Softmax			

TABLE I: YOLOv3-tiny architecture.

At this point, given the filtered point cloud, we start the estimation procedure by building a 2-manifold triangle mesh, for which the exact volume could be computed by resorting to the simple algorithm described in [15] assuming that the mesh is watertight and free of self-intersections.

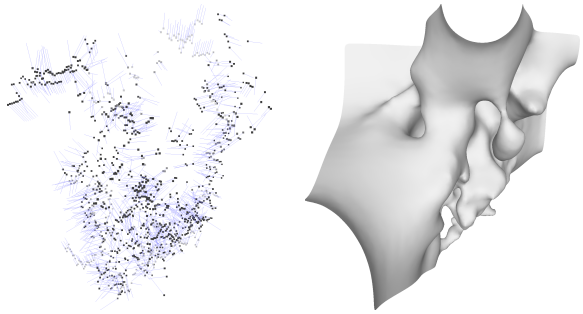

Fig. 2: The poor 3D mesh reconstruction provided by Screened Poisson reconstruction [16] (right) due to a noisy normal computation (shown in blue) and a sparse input point set (left).

Several methods to reconstruct such a 3D mesh from a point cloud can be found at the state of the art. Notably, this research area has seen substantial progress in the past two decades, the reader is referred to [2] for a comprehensive overview of this topic. While most of these methods are actively used in the industry to produce well-structured triangular meshes, they usually require a dense point set as an input, a pre-condition which unfortunately does not hold

in our application scenario. As a matter of fact, open-air vegetation is extremely challenging to acquire using scanning devices, and the result is usually a quite sparse point cloud where state-of-the-art reconstruction methods perform poorly (see Figure 2). Hence, we propose a novel automated strategy to derive a volumetric approximation of the sucker canopy.

input : A set of samples $s_i \in \mathcal{S}$
output: A new set of connected Samples

STEP 1 - AVERAGE : computing average distance α between each sample and the closest;
$\alpha \leftarrow 0$;
forall $s_i \in \mathcal{S}$ **do**
  $s_i^c \leftarrow$ `FindClosest` $(\mathcal{S}, s_i)$;
  $\alpha \leftarrow \alpha + |s_i^c - s_i|$;
end
$\alpha \leftarrow \alpha / |\mathcal{S}|$;

STEP 2 - CLUSTERING : Cluster samples whose distance is below α;
forall $s_i \in \mathcal{S}$ **do**
  $c_i = \{s_i\}$;
end
$merged \leftarrow true$;
while *merged* **do**
  $merged \leftarrow false$;
  forall $c_i \in \mathcal{C}$ **do**
    $c_i^c \leftarrow$ `FindClosestCluster` (c_i) ;
    if (`ClusterDistance` $(c_i, c_i^c) < \alpha$) **then**
      $c_i \leftarrow c_i \cup c_i^c$;
      $\mathcal{C} \leftarrow \mathcal{C} - c_i^c$;
      $merged \leftarrow true$;
    end
  end
end

STEP 3 - LINK : Greedy add links between partitions until a single connected component is created;
while $|\mathcal{C}| > 1$ **do**
  $c_0, c_1 \leftarrow$ FindClosestClusterPair($\mathcal{C}$);
  $c_{link} \leftarrow$ FindLinkSamples(c_0, c_1);
  $c_{new} \leftarrow (c_0 \cup c_1 \cup c_{link})$;
  $\mathcal{C} \leftarrow \mathcal{C} \cup c_{new}$;
  $\mathcal{C} \leftarrow \mathcal{C} - c_0$;
  $\mathcal{C} \leftarrow \mathcal{C} - c_1$;
end

Algorithm 1: A pseudo-code illustrating the different steps of our clustering procedure.

Given a sparse point cloud set $\mathcal{S}$ (see Figure 3.b), we first estimate a distribution factor α, as the average distance between each point $s_i \in \mathcal{S}$ to its closest sample s_i^c (see Figure 3.c). Then we connect the samples whose distance is below α, and we cluster the connected components (see Figure 3.d). Intuitively, if we create a sphere with a radius of α for each sample, then every cluster will result in a connected volume. A pseudo-code of this procedure is described by Algorithm 1. The first two steps of our algorithm are similar to the well-known hierarchical clustering algorithm. However, we perform an additional step (Step 3) to connect the different branches of a plant. Consequently, we perform an extra step to connect the various branches. We use a greedy strategy that favors the connection between closest cluster. Given the closest pair of clusters c_0 and c_1, we connect them to a new cluster c_{new} by creating a new set of samples which are distributed along the shortest segment that connects c_0 with c_1. We iterate this process by repeatedly connecting the closest pair of clusters until all the samples are joined to a unique cluster. Notice that for these new samples, we defined a sphere with a smaller radius $\alpha/5$ as they represent the branches of the sucker (see figure 3.e). We point out that the value of alpha for the branch of the sucker is the result of a trial and error procedure that has been carried out in order to obtain the lowest volumetric estimation error. We finally create a single mesh for every sphere, and we merge them all by using the Boolean operations of [17] implemented in LibIGL [18]. At this stage, as we performed the union of a set of closed, 2-manifold, orientable spheres, we also guarantee a closed, 2-manifold, orientable mesh T as an output, matching the condition for exact volume calculation. However, the mesh resulting from the boolean operation might have badly shaped triangles and geometrical sharp features which can affect the volume estimation. Hence, we improve the meshing by performing a couple of loop subdivision steps [19] using the implementation in the VCG library [20] (see figure 3.f). Future work will be focused on the optimization of this pipeline for the sake of real-time computation. Finally, as we previously stated, the volume of the mesh can be computed as the sum of all the signed volumes of the tetrahedra bounded by the origin and the three vertices of a triangle.

V. Experimental Validation

In this section we provide a preliminary experimental validation to demonstrate that the proposed method represents a promising solution for achieving an accurate volumetric estimation of the sucker canopy. In particular, first the sucker detection part is validated, then the volumetric estimation process is verified.

A. Experimental Setup

Real experiments have been carried out within the real-word hazelnut orchard available for the experimental validation of the PANTHEON project, which is composed of 3 fields selected within the "Azienda Agricola Vignola", a farm located in the municipality of Caprarola, in the province of Viterbo, Italy. In particular, the data collection of the suckers and the validation of the proposed approach were conducted within field 16, which hosts a young orchard (cultivar Nocchione treated as multi stemmed bushes) with a 4.5 m x 3.0 m layout.

Figure 4 shows the ground vehicle prototype (top), namely SHERPA HL robotic platform R-A, together with the sensor equipment used for data acquisition: i) a Velodyne VLP-16 Puck LITE 3D LiDAR, ii) a Genius WideCam F100 webcam (bottom), and iii) an Intel NF697 with an Intel Core i7-7700 Processor. In particular, the Genius webcam is located right below the Velodyne and both are mounted parallel to the ground plane. The Velodyne VLP-16 generates approximately 300.000 points/second and consists into 16 rings with an horizontal field of 360° and a vertical one of 30° (± 15° from the horizon). Given the front-side placement in our setup, we restricted its horizontal F.O.V. to be only at 180°. The Genius webcam has a 120° vision and an image resolution of 1280×720. It supports only manual focus that

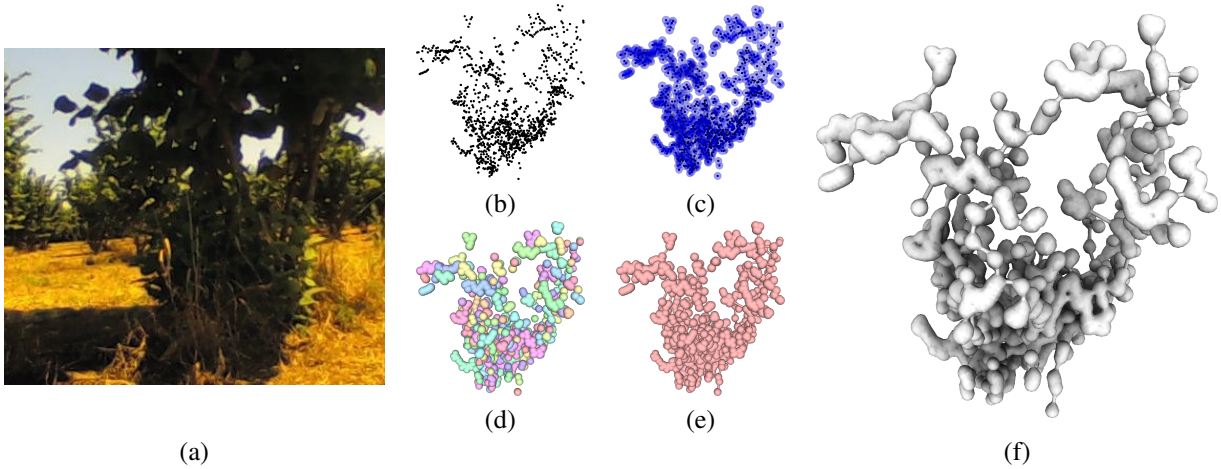

Fig. 3: The 3D reconstruction pipeline: given a sparse set of 3D points (b) sampled through 3D scanning on the real sucker (a) we estimate the average radius α (c) and we cluster samples based on their proximity (d); then we connect the separated branches using a greedy strategy (e) and we gather the final mesh as the union of all the spheres (f).

Fig. 4: The robotic platform (top) and the front-side of the robot where the sensors are mounted (bottom).

was kept fixed for all the duration of the experiments. As pointed out in the introduction, we reiterate that the LIDAR used for the experimental validation is purposely designed for navigation tasks, and therefore it does not provide the typical resolution of classical LIDAR solutions designed for 3D reality capture, such as the FARO S70 laser scanner. Nevertheless, we made this choice motivated by the fact that a sensor such as the Velodyne VLP-16 can be reasonably found in almost any autonomous vehicle, thus significantly reducing the economic impact of the hardware requirements of the proposed approach.

B. Sucker Detection

To test the sucker detection process, we gathered four datasets in different weather and daylight conditions. Specifically, $Data_A$ is composed of 715 images and was acquired in the middle afternoon and late morning; $Data_B$ is composed of 61 images and was acquired in the early morning with cloudy sky; $Data_C$ is composed of 181 images and was acquired at midday; finally $Data_D$ is composed of 213 images and was acquired in the late afternoon. We point out that the heterogeneity in the size of the collected datasets depends on the daylight conditions. Indeed, the most varying light conditions happen in the middle afternoon and in the late morning, leading $Data_A$ to be the largest dataset, while $Data_B$ to be the smallest dataset since the cloudy sky does not involves variable shadowing conditions. The four datasets were recorded across four different weeks in the summer. Each one was collected according to a different robot approaching maneuver towards a plant with suckers.

Note that, while $Data_A$ and $Data_C$ were then divided into train and test datasets, namely $Train/Test_A$ and $Train/Test_C$, $Data_B$ were entirely used for training, while $Data_D$ was split into two individual test sets, namely $Test_{D1}$ and $Test_{D2}$. It is important to remark that the fast growth rate of the suckers potentially leads to sensibly change their appearance in a week. Therefore, our aim is to use $TestSet_{D1,2}$ to test how the neural network performs on a previously unseen suckers condition.

	Sub Set I			Sub Set II		
Test Set	IoU	.5R	.75R	IoU	.5R	.75R
$Test_A$	65.4	87.3	22.1	62.4	100	22.5
$Test_C$	72.7	100	42.3	59.6	76.3	5.2
$Test_{D1}$	51.2	55.8	43.4	64.9	95.9	84.4
$Test_{D2}$	72.7	85.2	39.3	69.3	53.2	21.1

TABLE II: Sucker detection statistics.

Figure 5 reports some outcomes of the convolutional neural network across the different test datasets. We report numerical results in Table II according to the following metrics: (i) the average Intersection over Union (Iou), (ii) the recall rate at an IoU percentage of 50% (.5R), and (iii) the recall rate at an IoU percentage of 75% (.75R). To evaluate the sucker detection process for each approaching maneuver, the datasets have been further divided into $SubSetI$ and $SubSetII$. As reported, the IoU is always kept above the 50%, also in $Test_{D1,2}$. The same trend is also reported in the recall rates. The only exception is the $SubSetII$ in $Test_C$ where the camera was facing the sun, leading the detection network to fail more frequently, thus resulting in a lower recall rate. Future work will focus on investigating a strategy to mitigate this effect. A feasible solution could be to gather more data with similar light conditions, supplementing the current training datasets.

Fig. 5: Examples of network predictions for $TestSet_A$, $TestSet_B$ and $TestSet_{D1,2}$.

C. Volumetric Estimation

In this section, we provide a preliminary experimental validation of the accuracy of the proposed volumetric estimation. To compute an error metric, the first step was to obtain a reliable ground truth. In the case of suckers, standard 3D reconstruction approaches based on 3D dense point clouds, as shown in Figure 2, might suffer the 3D geometric irregularities of the sucker surface. Thus, to get an accurate estimate of the suckers real volume we used a variation of the Archimedes method. Specifically, the target sucker is suspended below the surface of the water in a container placed on an electronic scale. The volume V of the immersed object will simply be the increase in weight Δw divided by the water density ρ, that is $V = \Delta w/\rho$, see [3] for further details. In order to carry out our preliminary experimental validation, we measured three suckers, namely $Sucker_{1,2,3}$.

Sucker ID	RMSE [%]	St. Dev. [%]	Ground Truth [cm^3]
$Sucker_1$	21%	5.1%	631.9
$Sucker_2$	15%	3.8%	834.3
$Sucker_3$	6%	2.5%	372.8

TABLE III: Sucker volumetric estimation statistics.

Notably, even though only a limited number of samples were used for the experimental validation, according to Table III these preliminary results seem very promising. In detail, the RMSE remains below the 21% in all datasets. The entire process takes on average 10 seconds on a 2,9 GHz Intel Core i7 Mac, except for the boolean union step which might take up to 2 minutes.

VI. Conclusions and Future Work

In this paper, we proposed a sucker detection and canopy volume estimation approach composed of 2 steps. The first step detects the suckers and filter the 3D data, while the second one compute an approximate meshing of the target sucker and estimate its volume. A preliminary experimental validation has been conducted to demonstrate that the proposed approach represents a promising solution to develop a more environmentally friendly sucker control methodology. Future work will be mainly focused on: i) the design of a real-time processing pipeline; ii) the evaluation of different and cheaper sensor setups to gather both visual and 3D data; and iii) a more comprehensive experimental validation according to data collected in the next vegetative season.

References

[1] W. Lee, V. Alchanatis, C. Yang, M. Hirafuji, D. Moshou, and C. Li, "Sensing technologies for precision specialty crop production," *Computers and Electronics in Agriculture*, vol. 74, no. 1, pp. 2 – 33, 2010.

[2] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva, "A survey of surface reconstruction from point clouds," *Comput. Graph. Forum*, vol. 36, no. 1, pp. 301–329, Jan. 2017.

[3] S. Hughes, "Archimedes revisited: A faster, better, cheaper method of accurately measuring the volume of small objects," *Physics Education*, vol. 40, 09 2005.

[4] F. Kang, H. Wang, F. Pierce, Q. Zhang, and S. Wang, "Sucker detection of grapevines for targeted spray using optical sensors," *Transactions of the ASABE (American Society of Agricultural and Biological Engineers)*, vol. 55, pp. 2007–2014, 11 2012.

[5] Y. Wang, X. Shasha, L. Wenbin, F. Kang, and Z. Yongjun, "Identification and location of grapevine sucker based on information fusion of 2d laser scanner and machine vision," vol. 10, pp. 84–93, 01 2017.

[6] S. Paulus, J. Dupuis, S. Riedel, and H. Kuhlmann, "Automated analysis of barley organs using 3d laser scanning: An approach for high throughput phenotyping," *Sensors (Basel, Switzerland)*, vol. 14, pp. 12 670–12 686, 07 2014.

[7] T. Nguyen, D. Slaughter, B. Townsley, L. Carriedo, J. Maloof, and N. Sinha, "Comparison of structure-from-motion and stereo vision techniques for full in-field 3d reconstruction and phenotyping of plants: An investigation in sunflower," 07 2016.

[8] H. Shim, R. Adelsberger, J. Kim, S.-M. Rhee, T. Rhee, J.-Y. Sim, M. Gross, and C. Y. Kim, "Time-of-flight sensor and color camera calibration for multi-view acquisition," *The Visual Computer*, vol. 28, pp. 1139–1151, 12 2012.

[9] W. Dong, P. Roy, and V. Isler, "Semantic mapping for orchard environments by merging two-sides reconstructions of tree rows," *Journal of Field Robotics*, 08 2018.

[10] M. Pound, A. French, E. Murchie, and T. Pridmore, "Automated recovery of three-dimensional models of plant shoots from multiple color images," *Plant physiology*, vol. 166, 10 2014.

[11] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv*, 2018.

[12] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.

[13] A. Dhall, K. Chelani, V. Radhakrishnan, and K. M. Krishna, "LiDAR-Camera Calibration using 3D-3D Point correspondences," *ArXiv e-prints*, May 2017.

[14] D. Woebbecke, G. Meyer, K. Von Bargen, and D. Mortensen, "Shape features for identifying young weeds using image analysis," *Transactions of the ASAE*, vol. 38, no. 1, pp. 271–281, 1995.

[15] B. Mirtich, "Fast and accurate computation of polyhedral mass properties," *J. Graph. Tools*, vol. 1, no. 2, pp. 31–50, Feb. 1996.

[16] M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Trans. Graph.*, vol. 32, no. 3, pp. 29:1–29:13, July 2013.

[17] Q. Zhou, E. Grinspun, D. Zorin, and A. Jacobson, "Mesh arrangements for solid geometry," *ACM Trans. Graph.*, vol. 35, no. 4, pp. 39:1–39:15, July 2016.

[18] A. Jacobson, D. Panozzo, *et al.*, "libigl: A simple C++ geometry processing library," 2018, https://libigl.github.io/.

[19] C. Loop, "Smooth subdivision surfaces based on triangles," January 1987.

[20] CNR, "The visualization and computer graphics library," 2013, http://vcg.isti.cnr.it/vcglib/.