

# An Integrated Actuation-Perception Framework for Robotic Leaf Retrieval: Detection, Localization, and Cutting

Merrick Campbell, Amel Dechemi, and Konstantinos Karydis

**Abstract**—Contemporary robots in precision agriculture focus primarily on automated harvesting or remote sensing to monitor crop health. Comparatively less work has been performed with respect to collecting physical leaf samples in the field and retaining them for further analysis. Typically, orchard growers manually collect sample leaves and utilize them for stem water potential measurements to analyze tree health and determine irrigation routines. While this technique benefits orchard management, the process of collecting, assessing, and interpreting measurements requires significant human labor and often leads to infrequent sampling. Automated sampling can provide highly accurate and timely information to growers. The first step in such automated in-situ leaf analysis is identifying and cutting a leaf from a tree. This retrieval process requires new methods for actuation and perception. We present a technique for detecting and localizing candidate leaves using point cloud data from a depth camera. This technique is tested on both indoor and outdoor point clouds from avocado trees. We then use a custom-built leaf-cutting end-effector on a 6-DOF robotic arm to test the proposed detection and localization technique by cutting leaves from an avocado tree. Experimental testing with a real avocado tree demonstrates our proposed approach can enable our mobile manipulator and custom end-effector system to successfully detect, localize, and cut leaves.

## I. INTRODUCTION

Precision agriculture is a farming practice that utilizes sensor networks to help improve the use of agronomic inputs (e.g., water, fertilizers, pesticides) [1]. Robotics research in precision agriculture has largely focused on remote sensing via ground or aerial robots (e.g., [2]–[4]). Besides remote sensing, an increasing number of works has begun addressing interactions with the crop. Such works consider primarily robotic harvesting in both row (e.g., corn and soybean) and tree crops (e.g., citrus and avocado). For example, robots have been deployed to pick peppers, apples, citrus, and tomatoes by wrapping the fruit and twisting it off the stem with either a soft gripper [5]–[7], rigid gripper [8]–[14], or vacuum [15]–[17]. Some robots can pick strawberries, cucumbers, citrus, and peppers by cutting the stem [18]–[23].

This paper focuses on interaction with tree crops and addresses a conceptually-related yet less explored topic compared to robotic harvesting: *robotic leaf sampling*. Leaf sampling is important in agriculture since remote sensing typically provides field-level information without sufficient




Fig. 1. We develop a custom-built end-effector attached to an off-the-shelf 6-DOF robotic arm and a visual perception algorithm to detect, localize and cut leaves at their stem. (The supplementary video demonstrates the end-effector’s operation and overall system testing.)

resolution to accurately diagnose problems. Agronomists utilize specialized instruments that can be difficult to transport to the field and thus rely upon sample retrieval for later lab analysis. While this has been mostly a manual process to date, some work has been performed using aerial and ground robots. Mueller-Sim et al. demonstrated a robotic platform for rapid phenotyping and capable of manipulating leaves for in-situ measurements [24], [25]. Orol et al. developed a tele-operated aerial robot for cutting and collecting leaves from trees [26]. Ahlin et al. presented an algorithm for selecting and grasping tree leaves using a robotic arm [27]. The latter work demonstrates a high level of control using monoscopic depth analysis (MDA) and image-based visual servoing, but focuses on grasping and pulling the leaf instead of cleanly cutting the stem of the leaf, which is the focus of our work.

Our work is motivated by the need to perform leaf water potential measurements, an important process performed by agronomists to estimate tree stress levels and hence optimize irrigation patterns [28]. A leaf cut at its stem is placed inside a pressurized chamber instrument with its cut end exposed, and then the pressure at which water begins to escape from the cut stem is used to determine the leaf water potential [29], [30]. Agronomists use this measurement as a proxy for tree stress levels to optimize irrigation patterns. Though effective, these instruments can be tedious and potentially dangerous to operate.<sup>1</sup> As a result, a single tree is often used to quantify the health of the entire orchard leading to infrequent measurements and undersampled regions. Enabling robotic leaf sampling (this paper’s focus) for future use in robotic leaf water potential analysis can help improve measurement coverage and frequency while reducing human fatigue, and

The authors are with the Dept. of Electrical and Computer Eng., Univ. of California, Riverside, 900 University Avenue, Riverside, CA 92521, USA. Email: {mcamp077, adech003, karydis}@ucr.edu. We gratefully acknowledge the support of USDA-NIFA under grant # 2021-67022-33453 and a UC MRPI Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

<sup>1</sup>[www.pmsinstrument.com/maintenance/safety/](http://www.pmsinstrument.com/maintenance/safety/)

risk of bodily injury. Our work joins a growing body of works on robotic means for monitoring crop health and improving irrigation management practices [31], [32].

Compared to existing robotic leaf sampling methods [24]–[27] and harvesting systems that cut the stem of a fruit/vegetable [18]–[23], we are interested in performing clean cuts at leaves’ stems and retaining leaves for stem water potential analysis. As with related works, we also incorporate a visual perception component (to identify and localize a leaf) and an actuation component (to move the end-effector toward the leaf, and then cut it). Collecting a leaf sample from a tree presents unique challenges in perception and actuation, distinct from robotic fruit harvesting. Similarly, finding a motion plan to retrieve a physical sample needs to account for the presence of other leaves and branches which can also interfere with the extraction process. Yet, identifying a leaf sample involves not only segmenting the canopy, but also selecting an unblemished leaf suitable for stem water potential analysis [28].

To this end, we propose a leaf-cutting end-effector combined with a visual perception system that detects the center of a leaf and estimates its 6D pose (Fig. 1). The end-effector can cut and capture leaves of several common tree crops, such as avocado, clementine, grapefruit, and lemon. Unlike the MDA approach [27], we use a depth camera and a 3D point cloud to identify the centroid of the leaf and then estimate its 6D pose. This paper outlines our perception and actuation process to detect, localize, and cut leaves at their stem while retaining them, to enable future automated leaf water potential analysis in tree crops.

## II. RELATED WORKS

Development of harvesting end-effectors is an active area of research due to the wide variety of crops. While there are some commonalities across approaches, differences in size, weight, shape, texture, and firmness of specialty crops have led to unique solutions. Apples and citrus require a specific motion to grasp, twist, and pull from the tree without damage [21], [33]. Bell peppers and cucumbers can be directly cut and harvested [19], [20], [22], [34]. More delicate crops like strawberries call for manipulators with force feedback and flexible pneumatic actuators [18], [35]–[37]. Despite their unique applications, harvesting end-effectors generally have three primary components: the gripping mechanism (mechanical, pneumatic or hybrid), the removal mechanism (mechanical or electrical), and the sensing modality (monocular camera, stereo camera, time-of-flight) [23], [38].

At the same time, there has been development of perception techniques to monitor crop growth [39], [40], help prevent disease through early detection [41], [42], assist with quality control [43], [44], and help automate harvesting [33]. Success of these tasks depends on the visual perception subsystem’s ability to provide precise and accurate information about the target crop and relevant environmental context [45], including segmentation and localization of targets of interest. Most approaches have focused on fruit/vegetable targets by harnessing distinct colors and/or shapes [27], [46]–[49].

In this paper we are targeting identification and pose estimation of individual tree-crop leaves. This presents similar yet unique challenges compared to fruit (and broader canopy) identification. Instead of filtering out the leaves to focus on the fruit, our objective is to retain the leaves and segment the tree canopy further to obtain individual leaf poses. Leaf segmentation has been considered in current research using both classical computer vision tools [50]–[52] as well as machine learning [53]–[55]. However, classical methods are prone to changes in the environment, such as light, occlusions or overlapping surfaces, whereas learning-based methods require large training datasets and may still generalize poorly as environmental factors vary [56].

Furthermore, these techniques have rarely been employed online on onboard computers as part of a robotic manipulation system to identify, localize and physically cut the leaf. Although a leaf’s 3D position can be readily obtained, it is not sufficient to successfully accomplish the task as orientation plays an important role as to how a robotic arm approaches the leaf to cut it. Thus, obtaining at least an estimate of the 6D pose (position and orientation) is critical. Traditional 6D pose estimation approaches usually perform local keypoint detection and feature matching, and then a RANSAC-based PnP algorithm on the established 3D-to-2D correspondences to estimate the pose of an object [57], [58]. Still, they typically fail to perform with heavily occluded and poorly textured objects. On the other hand, learning-based methods use a deep neural network (DNN) to obtain the correspondences between 3D objects points and their 2D image projections [59]–[61]. Use of synthetic data generators [62], [63] can relieve in part the challenge of acquiring large labeled datasets; however, it requires realistic models that take into account the variations of the detected object e.g., shape, size, orientation or curvature which can be hard to develop.

Our developed end-effector focuses on the actuation and perception techniques to cut and retain a leaf. This task has received much less attention by existing robotic harvesting/leaf sampling technology yet an important aspect toward enabling future robotic leaf water potential measurements.

## III. TECHNICAL APPROACH

Picking a leaf requires two key components: actuation and perception. For actuation, we design a custom-built leaf-cutting end-effector (Section III-A) and retrofit it on a mobile manipulation base platform (Kinova Gen-2 six degree of freedom [6-DOF] robot arm mounted on a Clearpath Robotics Husky wheeled robot). For perception, we utilize point cloud data from a depth camera (Intel RealSense D435i) for the leaf detection and localization algorithm developed herein (Section III-B). The point cloud data is processed using Open3D [64] running on an Intel i7-10710U CPU, without any additional GPU acceleration. Figure 2 highlights how our contributions interact in a leaf-cutting system, which is further evaluated in Section IV.

Identified and segmented leaves serve as target for the arm to move and align the end-effector along a viable leaf (to be

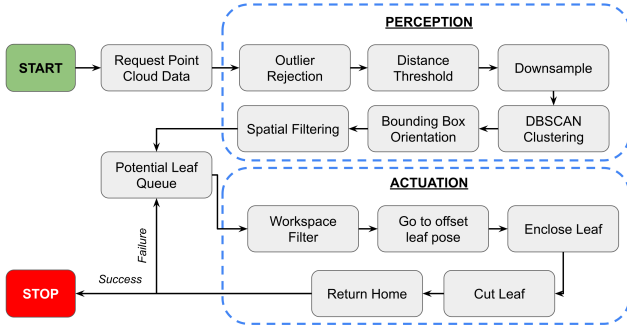


Fig. 2. Our approach jointly considers perception and actuation. The perception module processes point cloud data to segment leaves and deposit leaf candidates into a queue. Candidate leaves are then passed to the robot arm controller to actuate the end-effector. If a cut is successful, the routine ends. If unsuccessful, the arm controller requests the next leaf in the queue.

defined in Section IV), at an offset position from the center of the leaf. The offset distance is equivalent to the length of the leaf. Once at the offset position, the arm moves linearly toward the leaf to capture it. When the leaf is enclosed, the end-effector cuts the leaf. Then, the arm returns home.

#### A. Actuation

The stem-cutting end-effector developed herein utilizes two 4-bar linkages to actuate a set of sliding gates, one of which contains a razor blade to remove the leaf from the tree (Fig. 3). The gates also help retain the leaf within the end-effector’s chamber after removal from the tree. These 4-bar mechanisms are connected via a gear train to achieve synchronized motion. A low-cost, high-torque R/C servo (FEETECH FT5335M) drives the gear train while being amenable to position control. An Arduino Due microcontroller controls the servo motor and receives serial commands from a ROS control node. A breakout board connected to the Arduino contains a “safe/armed” switch along with LED indicators to reduce the risk of accidental injury.

Stem water potential analysis requires the test leaf’s stem to be cleanly cut; a damaged specimen would negatively impact the analysis [29]. Organic matter such as leaf stems exhibit visco-elastic properties. Based on visco-elastic material principles, faster cuts will require less force and result in less deformation of the leaf stem. Our prototype end-effector is able to cut leaf stems with a design target force of 20 N at 1.1 m/s. The end-effector’s chamber has an opening of 110 mm by 45 mm and a depth of 185 mm to accommodate typical avocado leaves. The end-effector is constructed with miniature aluminum extrusions, lightweight 3D printed parts, and laser-cut acrylic panels. The assembly weighs 1.091 kg, which is 42% of the robotic arm’s 2.6 kg payload. The end-effector is powered separately from the arm to enable stand-alone testing with a 7.4 V 2S LiPo battery.

To determine the types of leaves that can be cut by the mechanism, we performed testing with a variety of trees in a local orchard. The end-effector was manually placed around leaves and activated. Four different crops were selected (avocado, clementine, grapefruit, and lemon) for evaluation. For each crop, ten cutting attempts were performed. Results are shown in Table I. The end-effector was able to cut 95%




Fig. 3. The end-effector contains the components necessary to cut a leaf from a tree. The servo motor (red) actuates a double four-bar mechanism (yellow) that closes a set of gates (blue) with a razor blade to cut and capture a leaf. An Intel RealSense camera D435i is mounted on the top of the end-effector for perception. A microcontroller is mounted on the arm for controlling the motor. This end-effector can be mounted to a robotic arm using an adaptor plate (green). (Figure best viewed in color.)

of the leaves (38 out of 40). Lower success rates were observed for the lemon and grapefruit leaves. This is due to these particular leaves having shorter stems which made it harder to position the end-effector around the stem without interference from branches or other leaves. The end-effector worked consistently on clementine and avocado leaves.

TABLE I  
LEAF CUTTING TESTS

Crop	Successful Cuts	Attempts	Rate
Avocado	10	10	100%
Clementine	10	10	100%
Grapefruit	9	10	90%
Lemon	9	10	90%
Total	38	40	95%

#### B. Perception

We propose a leaf detection and localization algorithm using 3D point cloud and processed through the Open3D library. Our approach is outlined in Fig. 2. The detection phase seeks to obtain the 3D bounding box of leaves candidates from point cloud captured from the depth camera. First, we remove outliers considered as noise resulting from sensor measurement inaccuracies and segment out the background at a specific distance threshold from the camera frame. Then, downsampling is applied to optimize the performance of the upcoming step. Next, we group the remaining point cloud segments into clusters using the Density Based Spatial Clustering of Applications with Noise (DBSCAN) approach [65]. It relies on two parameters, the minimum distance between two points to be considered as neighbors (*eps*) and the number of minimum points to form a cluster (*MinPoints*).

Each resulting cluster is considered a potential leaf and described by a 3D bounding box defined by center  $C = [c_x, c_y, c_z]^T$ , dimensions  $D = [h, w, d]$ , and orientation  $R(\theta, \Phi, \alpha)$ . Then, filtering is applied on the clusters using geometric features of the bounding box: number of points, volume, leaf ratio. Finally, the pose of the center of each bounding box is returned as the 6D pose of a potential leaf.





Fig. 4. Key steps in our proposed leaf detection and localization process. The sample here corresponds to an outdoor point cloud: (a) corresponding RGB image of the tree, (b) raw point cloud, (c) distance filtered ROI, (d) downsampled point cloud, (e) segmented clusters, and (f) detected candidate leaves without 6D pose bounding boxes.

To validate our approach, we conducted offline tests for detection and localization separately. For the detection step, ROSbags were collected both in indoor and outdoor settings. Indoors (lab with constant light conditions), we used the Kinova arm with the camera placed at different distances (0.2 – 0.3 m) from a potted tree. Outdoors (local orchard with varying light conditions), we collected data manually. We considered a wide range (0.5 – 1.6 m) of distances from trees; an example is shown in Fig. 4.a. A total of 25 point clouds were collected (10 indoor and 15 outdoor). and tested offline with different combinations for *eps* and *MinPoints* parameters, to determine optimal values for later use.

Table II shows the outcome of our experiments on the 10 indoor point clouds and 15 outdoor point clouds. We attain an average of 80.0% of detection with a maximum of 90% for indoor dataset, and an average of 79.8% with a maximum 85% for outdoor. Further, we observed that the distance between the camera and the tree impacts the optimal values for the point cloud processing. The greater the distance from the camera, the higher *eps* while *MinPoints* decreases.

TABLE II  
LEAF POINT CLOUD DETECTION

	Point Clouds	Total # Leaves	Average Detection	Percentage
Indoor	10	20	16	80.0%
Outdoor	15	99	79	79.8%

To validate the localization phase, we compare several 6D poses obtained via our proposed approach against ground truth data obtained from a VICON motion capture camera system. Retroreflective markers were placed around the center of leaves, as shown in Fig. 5, to estimate their pose.

Table III summarizes the results obtained for 12 random leaves positions. Our approach provides an estimation with



Fig. 5. We used motion capture to establish a ground truth for determining the leaf 6D pose. Markers were placed on a target leaf (left) with origin at the base of our 6-DOF robot (right). (A real avocado tree was used.)

mean error of 8.28 mm, 14.38 mm, and 15.54 mm along x-axis, y-axis, and z-axis, respectively, for avocado leaves of width ranging between 24 – 86 mm and length ranging between 54 – 150 mm. Based on the average leaf size ( $48 \times 91$  mm), estimation errors represent nearly 15% of the width and 17% of the length. We evaluated the orientation by calculating the Euclidean distance between the two provided values using the definition in [66]. We obtained a mean error of 5.3deg. We observe that the obtained 6D pose may drift from the physical center of the leaf mainly on the y-axis and z-axis due to human-induced error and the non-rigid nature of the leaf which impacts marker placement.

TABLE III  
LEAF 6D POSE ERROR

Error	$\Delta x$ (mm)	$\Delta y$ (mm)	$\Delta z$ (mm)	Orientation (deg)
Mean	8.28	14.38	15.54	5.3
Std dev	7.46	5.46	6.69	15.5

The proposed approach provides an initial 6D pose along with useful information of potential leaves using a processed 3D point cloud and obtained up to 80% of detection and a mean error less than 16mm and 5.3 deg. Both detection and localization steps were performed without the need of collection or storage of large data including 3D models, and training process. Furthermore, all tests were run using a CPU configuration, without any additional GPU acceleration.

#### IV. EXPERIMENTAL VALIDATION OF LEAF CUTTING

To evaluate our overall leaf detection, localization and cutting pipeline, we tested with a real potted avocado tree indoors (lab). The mobile manipulator and end-effector system was positioned at random poses near the base of the tree so that the end-effector was at distances ranging between 0.2 – 0.3 m from the edge of the tree canopy. An experimental trial consisted of collecting a point cloud, storing the identified and localized potential leaves in a queue, and then sending the queued leaves to the arm for a retrieval attempt. Each trial concluded once the queue was depleted and the tree was repositioned for the next trial. Figure 6 outlines this process.




Fig. 6. Overall leaf retrieval process. During the perception phase, (a) the point cloud is processed to determine a potential leaf. If a viable leaf is detected, (b) the arm will move to an offset position. (c) The arm will then perform a linear motion to capture the leaf. Once in position, (d) the arm will cut the leaf and (e) the leaf will fall into the enclosed chamber. (f) After completing the cut, the arm will return to the home position.

For each retrieval attempt, leaf candidates and viable leaves are determined. *Leaf candidates* are leaves that have a pose within the arm’s workspace. *Viable leaves* are leaf candidates that have a retrieval path within the arm’s workspace. For testing our point cloud detection, we are interested in monitoring both successful captures and successful cuts of the leaf. A *successful capture* occurs when the end-effector is placed around a viable leaf while a *successful cut* occurs when the enclosed leaf is removed from the tree. A *clean cut* occurs when the leaf is severed cleanly at the stem such that it could be used for stem water potential analysis.

Out of 46 trials, 63 potential leaves were detected by the point cloud. (Note that each point cloud in the trial could produce a variable amount of leaves, hence a higher number of potential leaves than trials.) After filtering the potential leaves to remove the leaves outside of the work space, 39 viable leaves remained. Out of these leaves, 27 were captured successfully (69.2%) while 21 of the 27 captured leaves were cut (77.8%). Table IV summarizes retrieval results while Table V highlights the process times. The mean point cloud processing (perception) time was 5.6 sec and the mean cutting (actuation) time was 10.6 sec. The mean total retrieval time was 16.2 sec.

TABLE IV  
LEAF RETRIEVAL NUMBERS & RATES

Stage	Number	Rate
Potential Leaves	63	N/A
Candidate Leaves	51	81.0%
Viable Leaves	39	76.5%
Successful Captures	27	69.2%
Successful Cuts	21	77.8%
Clean Cuts	4	19.0%
Near Misses	7	30.0%

Our system was able to remove a total of 21 leaves from the tree. However, not all leaves were clean cuts on the stem; four were classified as clean cuts for use in stem water potential analysis. The majority of the leaves were severed at the top of the leaf and not at the stem (Fig. 7). Our system produced seven near-misses where the leaf was cut within an average of 9.58 mm from the stem (std dev: 6.1 mm). The remaining 10 leaves were severed closer to the middle of the leaf, largely due to collisions with the branches. Similar branch interference also lead to four out of the six missed cuts from the captured leaf. These two problems could be solved in future work through a refined end-effector design, more robust path planning to account for branches, and implementing visual servoing for continuous stem alignment as the end-effector approaches a viable leaf.

TABLE V  
LEAF RETRIEVAL PERFORMANCE TIME (SECONDS)

Metric	Perception Part	Actuation Part	Overall Retrieval
Min	0.5	4.6	6.1
Max	11.0	61.7	62.5
Mean	5.6	10.6	16.2
Median	7.7	8.1	15.3
Std dev	3.9	10.4	10.2

## V. CONCLUSIONS

Our work develops a co-designed actuation and perception method for leaf identification, 6D pose estimation and cutting. Our developed leaf-cutting end-effector can cut leaves of various types of trees (avocado, clementine, grapefruit and lemon) cleanly at their stem with a 95% success rate on average. Our proposed 3D point cloud technique can be successful for detecting an average of 80.0% of leaves






Fig. 7. Sample leaves cut from our avocado tree during automated indoor tests. (a) The four leaves represent clean cuts suitable for stem water potential analysis. (b) The system also cut seven leaves that were classified as near-misses, which removed the leaf without the stem. (c) The remaining leaves were cut closer to the center, due to interference between the end-effector and the branches. (d) In two cases, collateral damage occurred when a second leaf was removed along with the target leaf. These instances were classified as a single successful cut, but not a clean cut since the two leaves would need to be separated for stem water potential analysis.

indoors and 79.8% outdoors, and localizing them with less than 17% error along the leaf's length or width. Experimental testing of the overall proposed framework for leaf cutting reveals that our system can capture 69.2% of viable leaves and cut 77.8% of those captured leaves.

These results offers a promising initial step toward automated stem water potential analysis, nonetheless several steps remain and are exciting avenues for future work. The end-effector can effectively cut the leaves, but its size presents a challenge when cutting certain leaves like those from lemon and grapefruit trees which in turn calls for further design optimization. The current path planning approach works well for leaves that are on the periphery of the tree's canopy. Alternate path planning strategies can be explored to reach leaves within the canopy closer to the trunk, and integrated with visual servoing to better align the cutter with the stem of the leaf as it is about to cut it. Furthermore, the system will need to be robust to disturbances such as wind before deployment in an outdoor orchard environment. Finally, to enable automated stem water potential analysis, the captured leaf will need to be transferred from the end-effector into a pressure chamber.

## REFERENCES

- [1] N. Zhang, M. Wang, and N. Wang, "Precision agriculture—a world-wide overview," *Computers and Electronics in Agriculture*, vol. 36, no. 2, pp. 113–132, 2002.
- [2] W. H. Maes and K. Steppe, "Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture," *Trends in Plant Science*, vol. 24, no. 2, pp. 152–164, 2019.
- [3] P. Radoglou-Grammatikis, P. Sariannidis, T. Lagkas, and I. Moscholios, "A compilation of uav applications for precision agriculture," *Computer Networks*, vol. 172, p. 107148, 2020.
- [4] J. Kim, S. Kim, C. Ju, and H. I. Son, "Unmanned aerial vehicles in agriculture: A review of perspective of platform, control, and applications," *IEEE Access*, vol. 7, pp. 105 100–105 115, 2019.
- [5] C. Lehnert, A. English, C. McCool, A. W. Tow, and T. Perez, "Autonomous sweet pepper harvesting for protected cropping systems," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 872–879, 2017.
- [6] C. J. Hohimer, H. Wang, S. Bhusal, J. Miller, C. Mo, and M. Karkee, "Design and field evaluation of a robotic apple harvesting system with a 3d-printed soft-robotic end-effector," *Transactions of the ASABE*, vol. 62, no. 2, pp. 405–414, 2019.
- [7] G. Chowdhary, M. Gazzola, G. Krishnan, C. Soman, and S. Lovell, "Soft robotics as an enabling technology for agroforestry practice and research," *Sustainability*, vol. 11, no. 23, p. 6751, 2019.
- [8] S. Mehta and T. Burks, "Vision-based control of robotic manipulator for citrus harvesting," *Computers and Electronics in Agriculture*, vol. 102, pp. 146–158, 2014.
- [9] S. S. Mehta, W. MacKunis, and T. F. Burks, "Robust visual servo control in the presence of fruit motion for robotic citrus harvesting," *Computers and Electronics in Agriculture*, vol. 123, pp. 362–375, 2016.
- [10] Z. De-An, L. Jidong, J. Wei, Z. Ying, and C. Yu, "Design and control of an apple harvesting robot," *Biosystems Engineering*, vol. 110, no. 2, pp. 112–122, 2011.
- [11] J. R. Davidson, C. J. Hohimer, C. Mo, and M. Karkee, "Dual robot coordination for apple harvesting," in *ASABE Annual International Meeting*. American Society of Agricultural and Biological Engineers, 2017.
- [12] J. R. Davidson, A. Silwal, C. J. Hohimer, M. Karkee, C. Mo, and Q. Zhang, "Proof-of-concept of a robotic apple harvester," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 634–639.
- [13] T. T. Nguyen, E. Kayacan, J. De Baedemaeker, and W. Saeys, "Task and motion planning for apple harvesting robot," *IFAC Proceedings Volumes*, vol. 46, no. 18, pp. 247–252, 2013.
- [14] N. K. Uppalapati, B. Walt, A. Havens, A. Mahdian, G. Chowdhary, and G. Krishnan, "A berry picking robot with a hybrid soft-rigid arm: Design and task space control," *Robotics: Science and Systems*, p. 95, 2020.
- [15] J. Schupp, T. Baugher, E. Winzeler, M. Schupp, and W. Messner, "Preliminary results with a vacuum assisted harvest system for apples," *Fruit Notes*, vol. 76, no. 4, pp. 1–5, 2011.
- [16] J. Baeten, K. Donné, S. Boedrij, W. Beckers, and E. Claesen, "Autonomous fruit picking machine: A robotic apple harvester," in *Field and Service Robotics*. Springer, 2008, pp. 531–539.
- [17] K. Zhang, K. Lammers, P. Chu, Z. Li, and R. Lu, "System design and control of an apple harvesting robot," *Mechatronics*, vol. 79, p. 102644, 2021.
- [18] S. Hayashi, K. Shigematsu, S. Yamamoto, K. Kobayashi, Y. Kohno, J. Kamata, and M. Kurita, "Evaluation of a strawberry-harvesting robot in a field test," *Biosystems Engineering*, vol. 105, no. 2, pp. 160–171, 2010.
- [19] E. Van Henten, D. Van't Slot, C. Hol, and L. Van Willigenburg, "Optimal manipulator design for a cucumber harvesting robot," *Computers and Electronics in Agriculture*, vol. 65, no. 2, pp. 247–257, 2009.
- [20] E. Van Henten, B. v. Van Tuijl, J. Hemming, J. Kornet, J. Bontsema, and E. Van Os, "Field test of an autonomous cucumber picking robot," *Biosystems Engineering*, vol. 86, no. 3, pp. 305–313, 2003.
- [21] C. Aloisio, R. K. Mishra, C.-Y. Chang, and J. English, "Next generation image guided citrus fruit picker," in *IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*, 2012, pp. 37–41.
- [22] B. Arad, J. Balendonck, R. Barth, O. Ben-Shahar, Y. Edan, T. Hellström, J. Hemming, P. Kurtser, O. Ringdahl, T. Tielen, and B. van Tuijl, "Development of a sweet pepper harvesting robot," *Journal of Field Robotics*, vol. 37, no. 6, pp. 1027–1039, 2020.
- [23] R. R. Shamshiri, C. Weltzien, I. A. Hameed, I. J. Yule, T. E. Grift, S. K. Balasundram, L. Pitonakova, D. Ahmad, and G. Chowdhary, "Research and development in agricultural robotics: A perspective of digital farming," *International Journal of Agricultural & Biological Engineering*, vol. 11, no. 4, pp. 1–14, 2018.
- [24] T. Mueller-Sim, M. Jenkins, J. Abel, and G. Kantor, "The robotanist: A ground-based agricultural robot for high-throughput crop phenotyping," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3634–3639.
- [25] J. Abel, "In-field robotic leaf grasping and automated crop spectroscopy," Master's thesis, Carnegie Mellon University: Pittsburgh, PA, USA, 2018.

- [26] D. Orol, J. Das, L. Vacek, I. Orr, M. Paret, C. J. Taylor, and V. Kumar, "An aerial phytobiopsy system: Design, evaluation, and lessons learned," in *International Conference on Unmanned Aircraft Systems (ICUAS)*, 2017, pp. 188–195.
- [27] K. Ahlin, B. Joffe, A.-P. Hu, G. McMurray, and N. Sadegh, "Autonomous leaf picking using deep learning and visual-servoing," *IFAC-PapersOnLine*, vol. 49, pp. 177–183, 2016.
- [28] "Using the pressure chamber for irrigation management in walnut almond and prune," <https://ucanr.edu/datastoreFiles/391-761.pdf>, accessed: 2022-02-28.
- [29] P. F. Scholander, E. D. Bradstreet, E. Hemmingsen, and H. Hammel, "Sap pressure in vascular plants: negative hydrostatic pressure can be measured in plants," *Science*, vol. 148, no. 3668, pp. 339–346, 1965.
- [30] M. Tyree and H. Hammel, "The measurement of the turgor pressure and the water relations of plants by the pressure-bomb technique," *Journal of Experimental Botany*, vol. 23, no. 1, pp. 267–282, 1972.
- [31] S. Carpin, K. Goldberg, S. Vougioukas, R. Berenstein, and J. Viers, "Use of intelligent/autonomous systems in crop irrigation," in *Robotics and automation for improving agriculture*, 2019.
- [32] D. Tseng, D. Wang, C. Chen, L. Miller, W. Song, J. Viers, S. Vougioukas, S. Carpin, J. A. Ojea, and K. Goldberg, "Towards automating precision irrigation: Deep learning to infer local soil moisture conditions from synthetic aerial agricultural images," in *IEEE Conference on Automation Science and Engineering (CASE)*, 2018, pp. 284–291.
- [33] L. Bu, G. Hu, C. Chen, A. Sugirbay, and J. Chen, "Experimental and simulation analysis of optimum picking patterns for robotic apple harvesting," *Scientia Horticulturae*, vol. 261, p. 108937, 2020.
- [34] B. Lee, D. Kam, B. Min, J. Hwa, and S. Oh, "A vision servo system for automated harvest of sweet pepper in korean greenhouse environment," *Applied Sciences*, vol. 9, no. 12, p. 2395, 2019.
- [35] W. Simonton, "Robotic end effectors for handling greenhouse plant material," *American Society of Agricultural and Biological Engineers*, pp. 2615–2621, 1991.
- [36] Y. Xiong, Y. Ge, L. Grimstad, and P. J. From, "An autonomous strawberry-harvesting robot: Design, development, integration, and field evaluation," *Journal of Field Robotics*, vol. 37, no. 2, pp. 202–224, 2020.
- [37] G. Bao, P. Yao, S. Cai, S. Ying, and Q. Yang, "Flexible pneumatic end-effector for agricultural robot: Design & experiment," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2015, pp. 2175–2180.
- [38] C. Morar, I. Doroftei, I. Doroftei, and M. Hagan, "Robotic applications on agricultural industry. a review," in *IOP Conference Series: Materials Science and Engineering*, vol. 997, no. 1, 2020, p. 012081.
- [39] Y. Zhu, Z. Cao, H. Lu, Y. Li, and Y. Xiao, "In-field automatic observation of wheat heading stage using computer vision," *Biosystems Engineering*, vol. 143, pp. 28–41, 2016.
- [40] P. Sadeghi-Tehran, K. Sabermanesh, N. Virlet, and M. J. Hawkesford, "Automated method to determine two critical growth stages of wheat: heading and flowering," *Frontiers in Plant Science*, vol. 8, p. 252, 2017.
- [41] T. Akram, S. R. Naqvi, S. A. Haider, and M. Kamran, "Towards real-time crops surveillance for disease classification: exploiting parallelism in computer vision," *Computers & Electrical Engineering*, vol. 59, pp. 15–26, 2017.
- [42] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, "Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks," *IEEE Access*, vol. 7, pp. 59 069–59 080, 2019.
- [43] Q. Su, N. Kondo, M. Li, H. Sun, D. F. Al Riza, and H. Habaragamuwa, "Potato quality grading based on machine vision and 3d shape analysis," *Computers and Electronics in Agriculture*, vol. 152, pp. 261–268, 2018.
- [44] G. Jahns, H. M. Nielsen, and W. Paul, "Measuring image analysis attributes and modelling fuzzy consumer aspects for tomato quality grading," *Computers and Electronics in Agriculture*, vol. 31, no. 1, pp. 17–29, 2001.
- [45] K. Kapach, E. Barnea, R. Mairon, Y. Edan, and O. Shahar, "Computer vision for fruit harvesting robots—state of the art and challenges ahead," *International Journal of Computational Vision and Robotics*, vol. 3, pp. 4–34, 2012.
- [46] L. Fu, F. Gao, J. Wu, R. Li, M. Karkee, and Q. Zhang, "Application of consumer rgb-d cameras for fruit detection and localization in field: A critical review," *Computers and Electronics in Agriculture*, vol. 177, p. 105687, 2020.
- [47] S. W. Chen, S. S. Shivakumar, S. Dcunha, J. Das, E. Okon, C. Qu, C. J. Taylor, and V. Kumar, "Counting apples and oranges with deep learning: A data-driven approach," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 781–788, 2017.
- [48] T. T. Nguyen, K. Vandevoorde, N. Wouters, E. Kayacan, J. G. De Baerdemaeker, and W. Saeys, "Detection of red and bicoloured apples on tree with an rgb-d camera," *Biosystems Engineering*, vol. 146, pp. 33–44, 2016.
- [49] Qiu Quan, Tian Lanlan, Qiao Xiaojun, Jiang Kai, and Feng Qingchun, "Selecting candidate regions of clustered tomato fruits under complex greenhouse scenes using rgb-d data," in *International Conference on Control, Automation and Robotics (ICCAR)*, 2017, pp. 389–393.
- [50] Y. Chen, S. Baireddy, E. Cai, C. Yang, and E. J. Delp, "Leaf segmentation by functional modeling," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019, pp. 2685–2694.
- [51] T. Miao, C. Zhu, T. Xu, T. Yang, N. Li, Y. Zhou, and H. Deng, "Automatic stem-leaf segmentation of maize shoots using three-dimensional point cloud," *Computers and Electronics in Agriculture*, vol. 187, p. 106310, 2021.
- [52] B. Elnashef, S. Filin, and R. N. Lati, "Tensor-based classification and segmentation of three-dimensional point clouds for organ-level plant phenotyping and growth analysis," *Computers and Electronics in Agriculture*, vol. 156, pp. 51–61, 2019.
- [53] R. Guo, L. Qu, D. Niu, Z. Li, and J. Yue, "Leafmask: Towards greater accuracy on leaf segmentation," in *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2021, pp. 1249–1258.
- [54] D. Kuznichov, A. Zvirin, Y. Honen, and R. Kimmel, "Data augmentation for leaf segmentation and counting tasks in rosette plants," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019, pp. 2580–2589.
- [55] H. Scharr, M. Minervini, A. P. French, C. Klukas, D. M. Kramer, X. Liu, I. Luengo, J.-M. Pape, G. Polder, D. Vukadinovic, X. Yin, and S. A. Tsafaris, "Leaf segmentation in plant phenotyping: a collation study," *Machine Vision and Applications*, vol. 27, pp. 585–606, 2015.
- [56] Z. He, W. Feng, X. Zhao, and Y. Lv, "6d pose estimation of objects: Recent technologies and challenges," *Applied Sciences*, vol. 11, no. 1, p. 228, 2021.
- [57] F. Michel, A. Kirillov, E. Brachmann, A. Krull, S. Gumhold, B. Savchynskyy, and C. Rother, "Global hypothesis generation for 6d object pose estimation," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 115–124, 2017.
- [58] E. Brachmann, F. Michel, A. Krull, M. Y. Yang, S. Gumhold, and C. Rother, "Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3364–3372, 2016.
- [59] Y. Hu, J. Hugonot, P. V. Fua, and M. Salzmann, "Segmentation-driven 6d object pose estimation," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3380–3389, 2019.
- [60] Y. Hu, P. Fua, W. Wang, and M. Salzmann, "Single-stage 6d object pose estimation," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2927–2936, 2020.
- [61] K. Park, T. Patten, and M. Vincze, "Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation," *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7667–7676, 2019.
- [62] M. V. Giuffrida, H. Scharr, and S. A. Tsafaris, "Arigan: Synthetic arabidopsis plants using generative adversarial network," *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pp. 2064–2071, 2017.
- [63] Y. Zhu, M. Aoun, M. Krijn, and J. Vanschoren, "Data augmentation using conditional generative adversarial networks for leaf counting in arabidopsis plants," in *British Machine Vision Conference (BMVC)*, 2018, p. 324.
- [64] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3D: A modern library for 3D data processing," *arXiv:1801.09847*, 2018.
- [65] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *International Conference on Knowledge Discovery and Data Mining*, 1996, pp. 226–231.
- [66] D. Q. Huynh, "Metrics for 3d rotations: Comparison and analysis," *Journal of Mathematical Imaging and Vision*, vol. 35, pp. 155–164, 2009.