

Lavender autonomous navigation with semantic segmentation at the edge

Original

Lavender autonomous navigation with semantic segmentation at the edge / Navone, A., Romanelli, F., Ambrosio, M., Martini, M., Angarano, S., Chiaberge, M.. - ELETTRONICO. - 2135:(2025), pp. 280-291. (Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD) 2023 Torino (Ita) September 18–22, 2023) [10.1007/978-3-031-74633-8_18].

Availability:

This version is available at: 11583/2996293 since: 2025-01-17T09:40:28Z

Publisher:

Springer Nature

Published

DOI:10.1007/978-3-031-74633-8_18

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Lavender Autonomous Navigation with Semantic Segmentation at the Edge

Alessandro Navone¹[0009-0001-3768-3832] (✉),
Fabrizio Romanelli²[0000-0002-1888-7004], Marco Ambrosio¹[0000-0002-2120-9837],
Mauro Martini¹[0000-0002-6204-3845], Simone Angarano¹[0000-0002-4445-9783], and
Marcello Chiaberge¹[0000-0002-1921-0126]

¹Department of Electronics and Telecommunications, Politecnico di Torino, 10129, Torino, Italy
{name.surname}@polito.it

²Department of Civil Engineering and Computer Engineering, University of Rome Tor Vergata,
00133, Rome, Italy
fabrizio.romanelli@uniroma2.it

Abstract. Achieving success in agricultural activities heavily relies on precise navigation in row crop fields. Recently, segmentation-based navigation has emerged as a reliable technique when GPS-based localization is unavailable or higher accuracy is needed due to vegetation or unfavorable weather conditions. It also comes in handy when plants are growing rapidly and require an online adaptation of the navigation algorithm. This work applies a segmentation-based visual agnostic navigation algorithm to lavender fields, considering both simulation and real-world scenarios. The effectiveness of this approach is validated through a wide set of experimental tests, which show the capability of the proposed solution to generalize over different scenarios and provide highly-reliable results.

Keywords: Autonomous Navigation · Semantic Segmentation · Precision Agriculture

1 Introduction

In recent times, the food and farming industries have been facing a significant rise in global demand for food production, resulting in an increased need for resources [5]. As a result, there is a growing requirement for new methods and technologies to enhance efficiency and productivity while also ensuring sustainability throughout the process. As a consequence, farming industries focused their resources on developing new techniques to boost productivity and lower production costs, thereby reducing the need for labor-intensive human work [23].

The use of Deep Learning (DL) plays a significant role in the shift towards autonomous agents carrying out tasks in Agriculture 3.0 and 4.0 scenarios. Its ability to analyze data from various sources and reduce the likelihood of human errors makes it a valuable tool that can generalize and reduce tedious work [15]. DL has been proven to be beneficial for autonomous guidance systems, which include automating tasks such as navigation along the fields [6], including localization [21], global path planning [19]

and local motion planning [13], [18]. Moreover, among the different kinds of farming techniques, it emerged how row-organized crops are the most widely adopted farming arrangement, representing around 75% of the USA farmland [4].

The aim of this research is to address issues that arise when traditional localization methods, like GPS, are inaccurate due to factors such as weather conditions, obstructions, or signal interference caused by tall vegetation. Robust navigation in fields can be considered the starting point to carry out more complex and structural tasks such as crop monitoring [7], diseases detection [10] [20], spraying pesticides in a more localized and efficient way [8], fruit counting [16], harvesting [9] and monitoring the crop status and eventual diseases [10] [20]. Within this context, autonomous guidance systems and artificial intelligence have a primary role in achieving the objectives above.

The aim of this project is to create a motion planner that can navigate through plant rows of medium to high height. Many existing examples face challenges such as high costs and limited scalability. Typically, autonomous systems that navigate through row crops use high-precision GPS receivers and accuracy enhancement techniques [12], or a combination of sensors including lasers and GPS [17]. However, vegetation can interfere with the GPS signal reducing its accuracy and reliability [11]. Recent solutions involve the employment of multiple sensors such as GPS, inertial navigation systems (INS) wheels' encoder and LIDARs to improve localization accuracy [2]; however, equipment composed of multiple sensors can lead to an increase in the costs of the platform. A visual odometry system that uses a downward-looking camera to reduce costs was proposed in the study cited as [22]. Nonetheless, the authors noted that accuracy tends to decrease when the path is longer due to the accumulation of odometric error, and therefore, to keep the error bounded, integration with an absolute reference is necessary.

The main contribution of this work can be identified as the extensive experimentation carried out to validate the proposed method, which proves its generalization capabilities from a simulation to a real-world environment.

The rest of this work can be summarized as follows: Section 2 introduces the adopted solution for visual-based navigation in row crops, starting from the segmentation model to the adopted control algorithm. Section 3 presents the simulation environment and the experimental setup and, later, reports the obtained results for both simulation and real-world tests. Finally, Section 4 redraws and comments on the main points of this work.

2 Methodology

This study presents a new method for navigating without relying on position, using RGB-D data. The control algorithm created for this approach utilizes only real-time visual information to move through the rows of medium vegetation crops, like lavender. Additionally, this solution incorporates advancements in AI that enable "edge inferencing," where the robot's computer, often resource-limited, is able to process the input data in real time. The basic idea behind this work consists of the generalization of the visual control algorithms proposed for high-vegetation crops in [18] to lower vegetation fields, obtaining a reliable and continuous control using a cost-effective set of sensors.

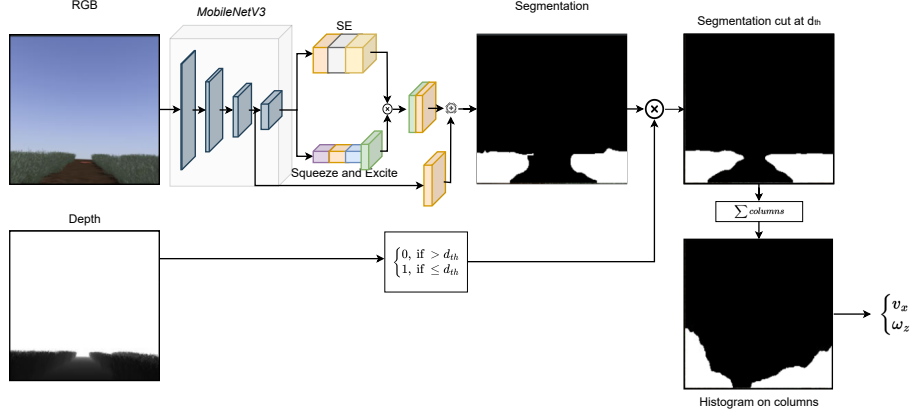


Fig. 1. Overall scheme of the pipeline of the proposed approach.

In fact, the proposed system only exploits RGB and depth images, overcoming the problems experimented with using GPS-based solutions. Furthermore, the system can seamlessly integrate into a comprehensive navigation configuration, which includes a waypoints generator, a global path planner, and a classic navigation system based on a Dynamic Window Approach (DWA) that controls the robots outside of the rows [6].

At instant t camera located in the front of the UGV platform captures the RGB frame, denoted as $\mathbf{X}_{rgb}^t \in \mathbb{R}^{h \times w \times c}$, and the depth frame denoted as $\mathbf{X}_{depth}^t \in \mathbb{R}^{h \times w}$. Their dimensions are determined by the number of rows, represented by h , the number of columns, represented by w , and the number of channels, represented by c . The RGB frame is later segmented semantically through a segmentation neural network, indicated as $H(\cdot)$. Its output, namely $\hat{\mathbf{X}}_{seg}^t$, is a binary segmentation mask that indicates the segmentation of the crop vegetation with ones, while the remaining areas of the frame are represented by zeroes.

$$\hat{\mathbf{X}}_{seg}^t = H(\mathbf{X}_{rgb}^t) \quad (1)$$

To achieve a more accurate determination of the end of the row's position, we eliminate noise from the segmented frame. We assess each column of the frame, denoted as $\hat{\mathbf{X}}_{seg}^t(:, j)$, and if over 97 % of the column is recognized as background, we set the remaining portion to zero. Moreover, to further increase the algorithm's robustness, the segmentation masks of the last N , i.e. $\{t - N, \dots, t\}$ time instants, are super-imposed, obtaining a cumulative segmentation mask.

$$\hat{\mathbf{X}}_{CumSeg}^t = \bigcup_{j=t-N}^t \hat{\mathbf{X}}_{seg}^t \quad (2)$$

where $\hat{\mathbf{X}}_{CumSeg}^t$ is the cumulative segmentation mask, and \bigcup represents the bitwise OR operation between the several segmentation masks.

Later on, the depth frame \mathbf{X}_d^t is used to cut the cumulative segmented frame at a threshold distance d_{th} in order to ignore the further segmentation data and better identify the continuation of the row.

$$\hat{\mathbf{X}}_{SegDepth}^t(i, j) = \begin{cases} 0, & \text{if } \hat{\mathbf{X}}_{CumDepth}^t(i, j) \cdot \mathbf{X}_{d(i, j)}^t > d_{th} \\ 1, & \text{if } \hat{\mathbf{X}}_{CumDepth}^t(i, j) \cdot \mathbf{X}_{d(i, j)}^t \leq d_{th} \end{cases} \quad (3)$$

where $i = 0, \dots, h$, $j = 0, \dots, w$ and $\hat{\mathbf{X}}_{SegDepth}^t$ is the segmentation frame cut with the depth information.

In order to determine the center of the row and generate velocity commands, the columns of the depth-cut segmented image are summed, obtaining a histogram, \mathbf{h}^t , obtained as in Equation 4. The practical idea behind this is to estimate the amount of vegetation per column.

$$\mathbf{h}_j^t = \sum_{i=1}^w \hat{\mathbf{X}}_{SegDepth}^t(i, j) \quad (4)$$

Therefore, after obtaining the histogram \mathbf{h}_j^t , it is evident how empty regions, namely clusters of zeros, represent regions in the field of view of the camera where no vegetation is present. Thus, identifying the widest cluster of zeros equals finding the continuation of the row. Therefore, the desired cluster is identified with the following steps:

1. The zeroes in the histogram are grouped in several clusters if they are in contiguous positions.
2. If clusters are smaller than a threshold (i. e. smaller than 3 elements), they are discarded.
3. the largest cluster is considered.
4. If the largest cluster occupies more than 80% of the space, it is considered an end-of-row condition.

Once the desired cluster is identified, the distance d between the center of the frame and the cluster center is considered and employed to calculate velocity commands.

2.1 Segmentation Network

Our real-time crop segmentation approach is based on previous works [18], [1], that utilized a network consisting of a MobilenetV3 backbone for feature extraction and an efficient LR-ASPP segmentation head, as represented in Figure 2. The LR-ASPP utilizes depth-wise convolutions, channel-wise attention, and residual skip connections to ensure a balance between accuracy and inference speed.

2.2 Velocity Command Generation

Once the center of the largest cluster is identified and the distance d from the center of the frame is evaluated, the velocity control commands can be generated. Furthermore, parabolic functions are used to generate the velocity functions.




Fig. 2. Sample frames and relative segmentation masks from AgriSeg dataset including synthetic data (a) and real data (b).

$$v_x = v_{x,max} \cdot \left(1 - \frac{d^2}{\left(\frac{w}{2}\right)^2} \right) \quad (5)$$

$$\omega_z = -\omega_{z,max} \cdot \text{sign}(d) \cdot \frac{d^2}{\left(\frac{w}{2}\right)^2} \quad (6)$$

In 5 and 6, $v_{x,max}$ and $\omega_{z,max}$ indicate the maximum linear and angular velocities respectively. The function $\text{sign}(\cdot)$ is used to determine the sign of a value, resulting in a value of 1 if the argument is greater than or equal to zero and -1 if it is less than zero.

3 Tests and Results

In this section, we present the results of experimental tests conducted both in simulation and in a real lavender field using a robot equipped with semantic segmentation capabilities at the edge. The objective of these tests was to evaluate the performance of the lavender autonomous navigation system and assess its effectiveness in identifying and avoiding obstacles. Specifically, the system has been tested on the field in order to assess its capability to be centered with respect to the rows of lavender as it travels along the row autonomously.

3.1 Evaluation Metrics

To quantitatively assess the performance of the lavender autonomous navigation system, we employed the following evaluation metrics:

- Navigation Success Rate: This metric measures the success rate of the autonomous robot in navigating through the lavender field without colliding with any obstacles. It is calculated as the percentage of successful navigation trials out of the total number of trials conducted.
- Root Mean Square Error (RMSE) between the actual trajectory and the center of the rows: This metric measures the autonomous system’s capability to maintain an equal distance from the rows while navigating toward the end of the field.

Regarding the navigation success rate metrics, we will report the percentage of successful navigation trials out of the total number of trials for the experiments in the next section.

While for the simulation test, we have an absolute ground truth provided by the simulator, in the real-world experiments, we considered the effective ground truth of the robot path to be the data acquired from the GPS RTK sensor during the experimental runs. Furthermore, in order to get a more accurate ground truth, we resorted to a Visual Odometry system fusing the trajectory estimated by an optical flow algorithm and the one estimated by the ORB-SLAM2 algorithm (a modified version of the one presented in [3]). The GPS RTK and Visual Odometry paths are finally fused and used as the trajectory ground truth for all the experiments. In order to assess how close the robot trajectory is from the center of the lavender rows, a set of measurements have been performed, measuring the GPS RTK coordinates, with 1 m step along the path, at the middle of the rows.

3.2 Segmentation Network Training

Frames with dimensions $w = 224$ and $h = 224$ were considered as input of the Segmentation Neural network. Moreover, the number of channels c is equal to 3 since they are RGB images. The model was trained on a combination of synthetic and real images from the lavender section of the AgriSeg dataset [14]. The dataset consisted of 4800 synthetic images and 1100 real ones. The model underwent 50 epochs of training with an ADAM optimizer, utilizing a learning rate of $3 \cdot 10^{-4}$. The data was augmented through cropping, flipping, grayscaling, and random jitters. The training of the model was conducted using TensorFlow 2 environment, starting from an ImageNet pretrained network on a single Nvidia RTX 3090 GPU.

3.3 Simulation Tests

Simulation Setup The software used to perform the simulation is Gazebo¹ since it is one of the most supported and diffused simulators for robotics applications. Blender² is used to create realistic models of plants and terrain that are exported and assembled in a Gazebo world using a procedural tool [14]. The robot model used in the simulation is the Husky UGV, the same that is used in real tests to reduce the gap between simulation and reality. Simulations are performed in rows of 8 – 10 m length, and at least three runs are performed to assess the repeatability of the control algorithm.

¹ <https://classic.gazebosim.org/>

² <https://www.blender.org/download/lts/3-3/>

Simulation Results The results obtained from the simulation are evaluated according to the metrics reported in Section 3.1. The Navigation Success Rate is 1.0 since, in all the runs, the robot managed to reach the end of the row without hitting plants. The overall behavior of the robot was robust and reliable navigation until the end of the row. The RMSE of the robot trajectory with respect to the center for all the performed tests is reported in Table 1. In addition, Figure 3 represent the trajectory followed by the robot during a simulation and the central line between the plants.




Fig. 3. Robot trajectories for one simulation run. The green dashed lines represent the plants, the blue line is the ideal central line, and the golden line represents the actual robot trajectory.

3.4 Experimental Tests



Fig. 4. A first-person view of the lavender rows with the path followed by the robot during the experimental tests.

Experimental Setup The experimental tests were conducted in a real lavender field with varying degrees of complexity in terms of terrain, vegetation density, and lighting conditions. The autonomous robot used in the experiments was equipped with state-of-the-art sensors, including cameras for capturing RGB images and depth information,

and a powerful onboard processing unit capable of performing semantic segmentation at the edge. The lavender field, where the experimental tests have been run, is located in Tuscany, Italy. In the same estate, two lavender fields are located. One field is composed of 22 lavender rows with an average length of 35 m with a distance between rows of about 80 cm. The other field is composed of 16 lavender rows with an average length of 25 m with a distance between rows of about 65 cm. The terrain was even, and there were weeds along the rows, even taller than the robot itself. The lighting conditions varied a lot between the experiments: from sunny to cloudy, also changing during every single run. An image of the environment is reported in Figure 4, where a first-person view of the lavender rows is visible, together with the cloudy sky and weeds along the path.

The robot used in the tests is a ClearPath Robotics Husky UGV, shown in Figure 5. Husky is a medium-sized robotic development platform. It has a large payload capacity and power systems. Stereo cameras, LIDAR, GPS, and IMUs are mounted on the UGV in order to achieve a high level of autonomy. The Husky UGV has a rugged construction and high-torque drivetrain that enables the robot to work in harsh environments and uneven terrain. Husky is also fully supported in ROS, where the algorithms have been developed, tested, and used in real experiments. Its external dimensions are $990 \times 670 \times 390$ mm, while its internal dimensions are $296 \times 411 \times 155$ mm. It weighs 50 Kg, and it can support a payload of 75 Kg at 1 m/s maximum speed with a 3 hours battery autonomy.




Fig. 5. An aerial view of the ClearPath Husky UGV, where the GPS antennas, IMU, LIDAR, and Intel RealSense D435 Camera are visible.

The real experiments consist of the robot running autonomously between the lavender rows with the aim of maintaining a proper distance from the lavender plants exploiting the semantic segmentation algorithms presented in Section 2, while moving towards the end of the row.

Table 1. Results of simulation and real-world tests. In simulation tests, the RMSE is computed between the geometrical center of the row and the actual trajectory; in real-world tests, it is computed between the actual robot trajectory (from the fused VO and GPS trajectory) and the center of the row.

Test number	Test type	Path length [m]	RMSE [m]
test 1	simulation	8.26	0.077
test 2	simulation	8.36	0.073
test 3	simulation	8.14	0.082
overall	simulation	24.76	0.077 ± 0.004
test 1	real world	23.27	0.289
test 2	real world	21.62	0.259
test 3	real world	11.88	0.049
test 4	real world	10.52	0.051
test 5	real world	18.55	0.152
test 6	real world	17.98	0.113
test 7	real world	21.12	0.235
overall	real world	124.94	0.164 ± 0.098

Experimental Results During the experimental tests, the lavender autonomous navigation system exhibited robust performance and demonstrated promising results. Here, we present the key findings based on the evaluation metrics mentioned in Section 3.1.

The autonomous robot achieved a remarkable navigation success rate of 70%. This indicates that the lavender autonomous navigation system effectively planned and executed collision-free paths, successfully maneuvering through the lavender field while avoiding obstacles. In order to assess this metric, we performed 10 runs within different lavender rows with different lengths, and the robot successfully moved without any collision on 7 tests. The 3 runs where the robot failed to perform its task were characterized by strong changes in lighting conditions (e.g., the sky conditions changed during the experiment passing from cloudy to fully sunny). However, even in these cases, the robot was able to complete about 80% of the path.

Furthermore, in order to assess the system’s ability to maintain a centered position with respect to the lavender rows, we measured the RMSE between the actual trajectory and the center of the rows. Here we report three significant results: one test for a long path (about 20 m) that presented the highest RMSE, one for a long path (about 20 m) that presented the lowest RMSE, and one for a short path of about 10 m. These results are summarized in Table 1. From these results, we noticed that the accuracy decreases when the path is longer due to the accumulation of odometric errors; also, tests 1 and 2 were performed along rows where the lavender plants were blooming, and therefore the rows were narrower, compared to those of test 3 where the lavender plants were lower. Finally, we present the details about the trajectories for the three tests just introduced. Figure 6 shows the experimental results for the three tests; there, we reported, for each test, the Visual Odometry estimated trajectory, the GPS RTK trajectory, the VO and GPS fused trajectory, and the center of the row. In the same figure, we also reported the




Fig. 6. Robot trajectories for three tests. The red solid line represents the center of the row, the green dashed line is the Visual Odometry trajectory, the blue dashed line is the GPS trajectory whilst the black dashed line represents the VO and GPS fused trajectory. For clarity's sake, we reported the boundaries of the lavender rows as red dotted lines.

row boundaries as red dotted lines. These boundaries are taken with a rough estimation of the lavender plants' center and help to determine the surface that can be traveled between the rows. As previously pointed out, test number 3 shows the best results in terms of the ability to stay close to the row's center because of two factors: the path length and the lavender plant state (lower and smaller than in cases 1 and 2).

4 Conclusions

Overall, the experimental results demonstrate the efficacy of the lavender autonomous navigation system with semantic segmentation at the edge. The high accuracy in semantic segmentation showcases the system's potential for autonomous operation in real-world lavender farming environments. These results validate the effectiveness of the proposed system in enabling autonomous robots to navigate through lavender fields efficiently, reducing manual intervention, and enhancing productivity in lavender cultivation. However, further testing and refinement are necessary to ensure robustness and adaptability across a wider range of lavender field conditions and environmental variations.

Acknowledgements This work has been developed with the contribution of Politecnico di Torino Interdepartmental Center for Service Robotics PIC4SeR³.

³ www.pic4ser.polito.it

References

1. Aghi, D., Cerrato, S., Mazzia, V., Chiaberge, M.: Deep semantic segmentation at the edge for autonomous navigation in vineyard rows. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3421–3428 (2021). <https://doi.org/10.1109/IROS51168.2021.9635969>
2. Astolfi, P., Gabrielli, A., Bascetta, L., Matteucci, M.: Vineyard autonomous navigation in the echord++ grape experiment**this work has been conducted under the “ground robot for vineyard monitoring and protection (grape)” experiment funded by the european commission under the echord++ project (fp7-601116). <http://echord.eu/grape/>. IFAC-PapersOnLine **51**(11), 704–709 (2018). <https://doi.org/10.1016/j.ifacol.2018.08.401>, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018
3. Bianchi, L., Carnevale, D., Masocco, R., Mattogno, S., Oliva, F., Romanelli, F., Tenaglia, A.: Efficient visual sensor fusion for autonomous agents. In: 2023 International Conference on Control, Automation and Diagnosis (ICCAD). pp. 01–06 (2023). <https://doi.org/10.1109/ICCAD57653.2023.10152399>
4. Bigelow, D., Borchers, A.: Major uses of land in the united states, 2012. Economic Information Bulletin Number 178 **178**(1476-2017-4340), 69 (2017)
5. Calicioglu, O., Flammini, A., Bracco, S., Bellù, L., Sims, R.: The Future Challenges of Food and Agriculture: An Integrated Analysis of Trends and Solutions. Sustainability **11**(1), 222 (Jan 2019). <https://doi.org/10.3390/su11010222>
6. Cerrato, S., Mazzia, V., Salvetti, F., Chiaberge, M.: A deep learning driven algorithmic pipeline for autonomous navigation in row-based crops (2021)
7. Comba, L., Biglia, A., Ricauda Aimonino, D., Barge, P., Tortia, C., Gay, P.: 2d and 3d data fusion for crop monitoring in precision agriculture. Proceedings of the IEEE pp. 62–67 (11 2019). <https://doi.org/10.1109/MetroAgriFor.2019.8909219>
8. Deshmukh, D., Pratihari, D., Deb, A., Ray, H., Bhattacharyya, N.: Design and development of intelligent pesticide spraying system for agricultural robot. In: Advances in Intelligent Systems and Computing. pp. 157–170 (04 2021). https://doi.org/10.1007/978-3-030-73050-5_16
9. Droukas, L., Dougeri, Z., Tsakiridis, N.L., Triantafyllou, D., Kleitsiotis, I., Mariolis, I., Giakoumis, D., Tzovaras, D., Kateris, D., Bochtis, D.: A survey of robotic harvesting systems and enabling technologies. Journal of Intelligent Robotic Systems **107**(2) (jan 2023). <https://doi.org/10.1007/s10846-022-01793-z>
10. Ferentinos, K.P.: Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture **145**, 311–318 (2018). <https://doi.org/10.1016/j.compag.2018.01.009>
11. Kabir, M.S.N., Song, M.Z., Sung, N.S., Chung, S.O., Kim, Y.J., Noguchi, N., Hong, S.J.: Performance comparison of single and multi-gnss receivers under agricultural fields in korea. Engineering in Agriculture, Environment and Food **9**(1), 27–35 (2016). <https://doi.org/10.1016/j.eaef.2015.09.002>
12. Ly, O., Gimbert, H., Passault, G., Baron, G.: A fully autonomous robot for putting posts for trellising vineyard with centimetric accuracy. In: 2015 IEEE International Conference on Autonomous Robot Systems and Competitions. pp. 44–49 (2015). <https://doi.org/10.1109/ICARSC.2015.42>
13. Martini, M., Cerrato, S., Salvetti, F., Angarano, S., Chiaberge, M.: Position-agnostic autonomous navigation in vineyards with deep reinforcement learning. In: 2022 IEEE 18th International Conference on Automation Science and Engineering (CASE). pp. 477–484 (2022). <https://doi.org/10.1109/CASE49997.2022.9926582>

14. Martini, M., Eirale, A., Tuberga, B., Ambrosio, M., Ostuni, A., Messina, F., Mazzara, L., Chiaberge, M.: Enhancing navigation benchmarking and perception data generation for row-based crops in simulation. In: Precision agriculture'23, pp. 451–457. Wageningen Academic Publishers (2023)
15. Mavridou, E., Vrochidou, E., Papakostas, G.A., Pachidis, T., Kaburlasos, V.G.: Machine vision systems in precision agriculture for crop farming. *Journal of Imaging* **5**(12) (2019). <https://doi.org/10.3390/jimaging5120089>
16. Mazzia, V., Khaliq, A., Salvetti, F., Chiaberge, M.: Real-time apple detection system using embedded systems with hardware accelerators: An edge ai application. *IEEE Access* **8**, 9102–9114 (2020). <https://doi.org/10.1109/ACCESS.2020.2964608>
17. Moorehead, S., Wellington, C., Gilmore, B., Vallespi, C.: Automating orchards: A system of autonomous tractors for orchard maintenance. In: Proceedings of the IEEE international conference of intelligent robots and systems, workshop on agricultural robotics (01 2012)
18. Navone, A., Martini, M., Ostuni, A., Angarano, S., Chiaberge, M.: Autonomous navigation in rows of trees and high crops with deep semantic segmentation. *arXiv preprint arXiv:2304.08988* (2023)
19. Salvetti, F., Angarano, S., Martini, M., Cerrato, S., Chiaberge, M.: Waypoint generation in row-based crops with deep learning and contrastive clustering. In: Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2022, Grenoble, France, September 19–23, 2022, Proceedings, Part VI. pp. 203–218. Springer (2023)
20. Shruthi, U., Nagaveni, V., Raghavendra, B.: A review on machine learning classification techniques for plant disease detection. In: 2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS). pp. 281–284 (2019). <https://doi.org/10.1109/ICACCS.2019.8728415>
21. Winterhalter, W., Fleckenstein, F., Dornhege, C., Burgard, W.: Localization for precision navigation in agricultural fields—beyond crop row following. *Journal of Field Robotics* **38**(3), 429–451 (2021)
22. Zaman, S., Comba, L., Biglia, A., Ricauda Aimonino, D., Barge, P., Gay, P.: Cost-effective visual odometry system for vehicle motion control in agricultural environments. *Computers and Electronics in Agriculture* **162**, 82–94 (2019). <https://doi.org/10.1016/j.compag.2019.03.037>
23. Zhai, Z., Martínez, J.F., Beltran, V., Martínez, N.L.: Decision support systems for agriculture 4.0: Survey and challenges. *Computers and Electronics in Agriculture* **170**, 105256 (Mar 2020). <https://doi.org/10.1016/j.compag.2020.105256>