

A Deep Learning Driven Algorithmic Pipeline for Autonomous Navigation in Row-Based Crops

Original

A Deep Learning Driven Algorithmic Pipeline for Autonomous Navigation in Row-Based Crops / Cerrato, Simone; Mazza, Vittorio; Salvetti, Francesco; Martini, Mauro; Angarano, Simone; Navone, Alessandro; Chiaberge, Marcello. - In: IEEE ACCESS. - ISSN 2169-3536. - ELETTRONICO. - 12:(2024), pp. 138306-138318. [10.1109/access.2024.3465873]

Availability:

This version is available at: 11583/2993047 since: 2024-10-03T10:54:23Z

Publisher:

IEEE

Published

DOI:10.1109/access.2024.3465873

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)

RESEARCH ARTICLE

A Deep Learning Driven Algorithmic Pipeline for Autonomous Navigation in Row-Based Crops

SIMONE CERRATO, VITTORIO MAZZIA¹, FRANCESCO SALVETTI¹,
MAURO MARTINI¹, (Member, IEEE), SIMONE ANGARANO¹,
ALESSANDRO NAVONE¹, (Graduate Student Member, IEEE),
AND MARCELLO CHIABERGE¹

Department of Electronics and Telecommunications (DET), Politecnico di Torino, 10129 Turin, Italy

Corresponding author: Mauro Martini (mauro.martini@polito.it)

This work was supported in part by the Contribution of the Politecnico di Torino, Interdepartmental Centre for Service Robotics (PIC4SeR) (<https://pic4ser.polito.it/>) and SmartData@Polito (<https://smartdata.polito.it/>); and in part by the Project Nord-Ovest Digitale E Sostenibile (NODES) through the MUR—M4C2 1.5 of Piano Nazionale Ripresa Resilienza (PNRR) funded by European Union—NextGenerationEU under Grant ECS00000036.

ABSTRACT Expensive sensors and inefficient algorithmic pipelines significantly affect the overall cost of autonomous machines. However, affordable robotic solutions are essential to practical usage, and their financial impact constitutes a fundamental requirement for employing service robotics in most fields of application. Among all, researchers in the precision agriculture domain strive to devise robust and cost-effective autonomous platforms in order to provide genuinely large-scale competitive solutions. In this article, we present a complete algorithmic pipeline for row-based crops autonomous navigation, specifically designed to cope with low-range sensors and seasonal variations. Firstly, we build on a robust data-driven methodology to generate a viable path for the autonomous machine, covering the full extension of the crop with only the occupancy grid map information of the field. Moreover, our solution leverages on latest advancement of deep learning optimization techniques and synthetic generation of data to provide an affordable solution that efficiently tackles the well-known Global Navigation Satellite System unreliability and degradation due to vegetation growing inside rows. Extensive experimentation and simulations against computer-generated environments and real-world crops demonstrated the robustness and intrinsic generalizability to different factors of variations of our methodology that open the possibility of highly affordable and fully autonomous machines.

INDEX TERMS Autonomous navigation, robotics, artificial intelligence, precision agriculture.

I. INTRODUCTION

Agriculture 3.0 and 4.0 have gradually introduced autonomous machines and interconnected sensors into several agricultural processes [1], [2], trying to introduce robust and cost-effective novel solutions into the overall production chain. For instance, precision agriculture has progressively innovated tools for automatic harvesting [3], vegetative assessment [4], crops yield estimation [5], smart and sustainable pesticide spraying robotic system [6] and many others [7], [8]. Indeed, the pervasiveness of precision agriculture techniques has such a huge impact

The associate editor coordinating the review of this manuscript and approving it for publication was Yangmin Li¹.

on certain fields of application that their adoption has become increasingly essential to achieve high product quality standards [9]. Moreover, the introduction of robots in agriculture will increasingly have a stronger impact on economic, political, social, cultural, and security, [10], and will be the only tool to satisfy the future food demand of our society [11].

Nevertheless, research on autonomous machines still requires further developments and improvements to meet the necessary industrial conditions of robustness and effectiveness. Global path planning [1], [2], mapping [12], localization [13], and decision-making [14] are only some of the required tools that each year undergo heavy research from the scientific community to achieve the necessary



FIGURE 1. Field tests with the Jackal platform in different seasonal periods of the same crop. Lush vegetation and thick canopies reduce the GPS accuracy, affecting its reliability and the overall navigation pipeline. Our proposed segmentation-based algorithm exploits semantic segmentation properties to provide a proportional controller that drives the robotic platform along the whole row.

requirements for full automation. Among all requirements, a low financial impact constitutes a fundamental goal in order to provide genuinely large-scale competitive solutions [15]. Indeed, expensive sensors and demanding computational algorithms significantly impact the actual usefulness of robotics platforms, essentially preventing their large-scale adoption and introduction into the agricultural world.

Recently, deep learning methodologies [16] revolutionized the entire computer perception field, endowing machines with an unprecedented representation of the surrounding environment [17]. Moreover, the intrinsic robustness of representation learning techniques to different factors of variations opens the possibility to achieve noteworthy perception capabilities with low-cost and low-range sensors, greatly relieving the overall cost of the target machine [18], [19]. Finally, the latest advancement in deep learning optimization techniques, [20], has progressively reduced the latency, inference cost, memory, and storage footprint of its algorithms. That effectively scales down computational requirements enabling low-power computational devices boosted by hardware accelerators such as visual processing units (VPUs), tensor processing units (TPUs), and embedded GP-GPUs, [21].

Building on the latest deep-learning research for computer perception and exclusively making use of low-range sensors, we present a complete algorithmic pipeline for autonomous navigation in row-based crops. The proposed robust solution is explicitly designed to adapt to seasonal variation, as shown in Fig. 1, ensuring complete coverage of the field in different situations. Moreover, the low financial impact of our methodology greatly reduces maintenance and production costs and enables a large-scale adoption of fully autonomous machines for precision agriculture.

Firstly, we build on a robust data-driven methodology to generate a viable path for the autonomous machine, covering the full extension of the crop with only the occupancy grid

map information of the field. Successively, depending on the vegetation growth status of the crop, we adopt a purely Global Navigation Satellite System (GNSS) local path planning or a vision-based algorithm that exclusively makes use of a low-cost RGB-D camera to navigate inside the inter-row space. Indeed, meteorological conditions, especially lush vegetation and thick canopies, significantly affect GNSS reliability, degrading its precision and consequently the overall navigation pipeline, [22], [23]. Conversely, our vision-based system exploits semantic information of the environment to navigate between rows and depth information to refine the underlying control smoothness, disentangling from the necessity of precise localization. Moreover, we make exclusive usage of synthetic data and domain randomization, [24], to enable affordable supervised learning and simultaneously bridge the domain gap between simulation and reality. Such a technique allows us to easily construct and train a deep learning model able to efficiently generalize on different row-based crops.

The overall proposed methodology guarantees to autonomously navigate through row-based crops without expensive sensors and with every seasonal variation. All the algorithmic pipelines have been developed ROS-compatible in order to make easier communication among different software modules and to be easily deployed on our developing platform, the Jackal Unmanned Ground Vehicle (UGV) by Clearpath Robotics.¹ Extensive experimentation and simulations against computer-generated environments and diverse real-world crops demonstrated the robustness and intrinsic generalizability of our solution. All of our data² are open source and publicly available.

A. RELATED WORK

Over the past years, autonomous systems designed to navigate in row-crops fields make large usage of high-precision GNSS receivers in combination with laser-based sensors [25], [26]. Nevertheless, the canopies on the sides of the row reduce the GNSS accuracy, affecting its reliability and forcing the adoption of more expensive sensors [22]. Indeed, more robust solutions fuse multiple sensor information (GNSS, inertial navigation systems, wheel encoders) to obtain a better estimation of the mobile platform location in the presence of thick canopies and adverse meteorological conditions [27], [28]. However, the high adoption of high-range sensors leads to higher system and maintenance costs, preventing a large-scale adoption of self-driving agricultural machinery.

On the other hand, visual odometry (VO), [29], [30], and computer vision-based solutions, [31] and [32], have been proposed as more affordable approaches. Nonetheless, VO systems show poor performance on long distances due to the accumulating errors and struggle with highly similar patterns [33] and unpredictable lighting conditions [34]. Moreover, purely vision-based solutions cannot deal with

¹<https://clearpathrobotics.com/>

²<https://naspic4ser.polito.it/files/sharing/P4ZwMHt7n>




FIGURE 2. A representation of the overall pipeline. From left to right, an occupancy grid of the crop is provided as input to the DeepWay neural network, which estimates the waypoints at the start/end of the vineyard's row. Then, a custom algorithm is responsible for ordering the generated waypoints and computing a global path, maintaining a safe distance from crops. Finally, a local path planning policy chooses the right navigation algorithm according to UGV's position with respect to waypoint in order to drive the mobile platform along the whole path.

seasonal variations, and generalizability to different crops is difficult to achieve.

In this state-of-the-art landscape, our team started in 2019 a research project with the precise aim to develop a complete working pipeline to autonomously navigate in vineyard rows without relying on multiple expensive sensors. As already introduced, computer vision and deep learning-based algorithms, [16], have demonstrated particular robustness in solving problems with noisy signals. Moreover, optimization and edge AI techniques have progressively made inference computational affordable, opening the usage of deep learning methodologies to diverse practical applications [35]. Consequently, in [36] and [37] were addressed navigation in vineyard rows with deep learning-based algorithm and the exclusive usage of an RGB-D camera. Furthermore, starting with [38], it has addressed the global path generation automation problem, which is commonly neglected by the research community. However, a suitable path generator is crucial for obtaining a complete autonomous navigation performance and its absence prevents the control of the platform on the field. For that reason, in [1], moving from the clustering solution proposed in [38] DeepWay, a robust data-driven approach for global path generation, was presented to detect the rows of the vineyards. Besides being more robust and easier to adopt, DeepWay is a highly general approach that can be extended to every kind of row-based crop.

B. NOVELTIES

Based upon the aforementioned previous work, the main contributions of the presented approach herein are as follows:

- Introduction of a highly affordable complete algorithmic pipeline for autonomous navigation in row-based crops. The methodology makes use of only low-cost, low-range sensors to drive a platform for the full extension of a crop in the different seasonal periods.
- Building on [37], we extend the local navigation between vineyard rows to a general row-based crop. We propose a domain randomization-based training

procedure to easily obtain a segmentation network with only synthetic data.

- We introduce a novel global path planning procedure to connect two successive rows, avoiding the usage of too general algorithms that could easily introduce jerky and inefficient paths.
- Improved usage of semantic information of the crop to navigate between rows in case of lush vegetation and thick canopies.

Compared to other existing autonomous navigation algorithms, our solution makes use of low-cost sensors such as a cheap GNSS receiver with Real-Time Kinematic (RTK) corrections, an RGB-D camera, an Inertial Measurement Unit (IMU), and encoders. Our approach aims to fill the gaps between usage of low-cost sensors and robustness of autonomous navigation in the precision agriculture context exploiting the collaboration between Artificial Intelligence and standard navigation algorithms. On the other hand, existing solution exploits high cost and very accurate sensors, as: 3D LiDAR, expensive RTK-GNSS receiver to achieve a reliable and robust complete autonomous solution.

The remainder of this paper is organized as follows. Section II describes the overall system framework. The detailed explanation of the proposed full pipeline is introduced in Section III. Section IV presents the pipeline evaluation either against computer-generated environments or real-world crops. Finally, Section V draws conclusions and suggests future work.

II. SYSTEM OVERVIEW

The proposed work is intended to present a complete autonomous navigation system for general row-based crops. The designed autonomous system is organized into several software modules that should collaborate with each other in order to obtain effective and reliable driverless navigation throughout the whole field. A visualization of the complete pipeline is shown in Fig. 2.

Firstly, the system takes as input a georeferenced occupancy grid map of the considered crop to compute a global path made of geographic coordinates; in particular, it exploits the DeepWay network, [1], to estimate the start/end waypoints of each row, then a custom global path planner, [39], computes the desired path maintaining a safe distance from crops. Secondly, according to the season period and the amount of crop vegetation, it is possible to choose the kind of navigation to perform: only GNSS-based or GNSS and AI-assisted. More in detail, the AI-assisted solution is exploited during summer and periods with thick vegetation that obstructs the GNSS signal.

In case a good view of the sky is available both outside and inside the row space, the system exploits only Real Time Kinematic (RTK) corrections, GNSS signals, and inertial data to autonomously guide the mobile platform throughout the crops. On the other hand, when lush vegetation is present, the proposed navigation system makes use of RTK corrections, GNSS signals, and inertial data to perform the row switch since outside the row space, a good sky view is available, while along the rows it exploits the camera, the deep neural network, and the segmentation-based control to overcome GNSS signals unreliability. In both scenarios, the system strongly relies on the estimated global positions of the UGV in order to follow the provided global path. In our approach, the localization problem is tackled by fusing the positioning information coming from an RTK-enabled GNSS receiver and the inertial data provided by an IMU. All the data is loosely fused, exploiting the well-known Extended Kalman Filter (EKF), which uses an omnidirectional model for prediction and outputs position estimations in the form of x, y, yaw since the navigation happens in 2-dimension. x and y are global spatial information represented in the East-North-Up (ENU) reference frame, while yaw is the absolute orientation with respect to magnetic north, corrected with the actual magnetic declination. Once the localization filter is set up, in case a clear view of the sky is available, the navigation algorithm uses the estimated UGV positions, the global path, and a local planner based on the Dynamic Window Approach (DWA) to autonomously guide the mobile platform; otherwise, it makes use of both GNSS-based and AI-assisted navigation, that exploit GNSS signals and semantic information of crops to safely navigate in the whole field without colliding with the crops.

In the latter case, the navigation type selection occurs by comparing the estimated UGV global positions and an ordered waypoints list computed by the DeepWay neural network and successively refined. Finally, it is important to underline that the DWA navigation scheme is only one possible solution. It is adopted in the presented pipeline for its simplicity, flexibility, and online collision avoidance capability. However, further experimentation with even simpler algorithms has been performed but not presented for the sake of conciseness. For instance, the Pure Pursuit controller [40] demonstrated very promising results between and outside rows, bringing possible

advantages in presence of more packed rows or larger vehicles.

III. METHODOLOGY

In this section, each block of the navigation system is presented, detailing all the aspects of the path planning and navigation processes. Our methodology requires as input a georeferenced occupancy grid of the target field $X_{occ} \in \mathbb{R}^{H \times W}$, obtained by segmenting an aerial view of the environment. The system is developed and tested on satellite imagery, but the very same methodology can be applied to images obtained by drones flying over the target field. The global path is computed by the first two blocks of the pipeline and is represented as an ordered set of points $P = \{(x, y) | x, y \in \mathbb{R}\}$ to be followed by the UGV in order to reach full coverage of the target field.

Fig. 3 shows an example of an occupancy grid with the predicted global path in red superimposed on the aerial image of the field. Once the global path P has been generated, the system exploits a local path planning policy to autonomously navigate in the considered crop choosing the proper control algorithm according to seasonal period and the presence of thick canopies on crops. Indeed, full autonomous navigation happens by exploiting the GNSS information and the semantic information obtained by employing the segmentation network.




FIGURE 3. Aerial view of a row crop field, together with the occupancy grid (white), the estimated waypoints (yellow), and the global path (red). The meters/pixel resolution is 0.1 m/px, and the end-row distance margin $d_{er} = 20$ px, which results in a 2 m real-world margin.

A. WAYPOINTS ESTIMATION

The first block aims at predicting the list of l waypoints $W \in \mathbb{N}^{l \times 2}$ in the occupancy grid reference frame that represents the begin and end of each row of the target field. Since classical clustering methods fail with real-world conditions such as rows of different lengths, holes, and outliers, we adopt the DeepWay framework [1], which frames the waypoints prediction as a regression problem. DeepWay is a fully convolutional neural network that takes as input the occupancy grid X_{occ} of dimension $H \times W$ and outputs a map

\hat{Y} of dimension $U_H \times U_W \times 3$:

$$\hat{Y} = f_{DeepWay}(X_{occ}) \quad (1)$$

The first channel of \hat{Y} is a confidence map that outputs for each cell u the probability $P(u)$ that a waypoint falls inside the cell itself. The output map dimensions are obtained by subsampling the input space of a factor k :

$$U_H = H/k \quad U_W = W/k \quad (2)$$

Thus, each cell u represents a square region of $k \times k$ pixels of the original occupancy grid. The other two channels of the output map \hat{Y} predict for each cell u two compensation factors Δx and Δy used to localize the waypoint inside the u cell, as shown in Fig. 4. Those factors are normalized to $[-1, 1]$ range so that they represent a distance from the center of the cell. Thus, a factor of 1 means a positive deviation on the corresponding axis of half the length of the cell. Eventually, the final location of the predicted waypoints in the input space can be recovered as:

$$\hat{y}_O = k \left(\hat{y}_U + \frac{\Delta + 1}{2} \right) \quad (3)$$

where \hat{y}_O and \hat{y}_U are the vectors of (x, y) coordinates of a generic waypoint in the input and output reference systems, respectively; Δ is the vector of the $(\Delta x, \Delta y)$ normalized compensation factors.

The prediction confidences stored in the first channel of the output map \hat{Y} are compared to a confidence threshold c_{thr} , and all the positions with $P(u) > c_{thr}$ are selected and projected in the input space as in Eq. 3. Furthermore, a suppression mechanism is adopted as in standard object detection algorithms in order to avoid multiple predictions of the same waypoint: all the points falling within a distance threshold d_{thr} from each other are replaced with the one with maximum confidence $P(u)$. The final waypoints are stored in the list W .

The network is characterized by a stack of N Residual Reduction modules that are based on 2D convolutions with Mish activation [41] and implement both channel and spatial attention [42]. Each module halves the spatial dimension with a Reduction block based on 2D convolutions with strides of two. After N modules, a 2D Transpose Convolution increases the spatial dimension by a factor of two. A last 2D convolution projects a concatenation of the features of the last two modules to the 3-dimensional space of the output. Since the first and last convolutional layers also have strides of two, the output tensor spatial dimension is reduced by a factor $k = (N + 1)^2$ with respect to the input. The final layer uses sigmoid activation for the first channel, which encodes the waypoint probability, and tanh activation for the other two channels, which encodes the normalized compensation factors. All code related to DeepWay is open-source and can be found online.³

³<https://github.com/fsalv/DeepWay>




FIGURE 4. Example of an output map \hat{Y} of DeepWay [1]. Since $k = 8$, each cell u represents a square area of 8×8 pixels of the original occupancy grid. The local reference system R_u is at the center of the cell, and the compensation factors Δx and Δy are normalized to the $[-1, 1]$ range as a fraction of the semi-cell length.

B. GLOBAL PATH PLANNING

The output list of waypoints W should be ordered in order to plan a global path that reaches a full coverage of the field. We adopt the same post-processing as in [1] with the following steps:

- 1) the row crops orientation is estimated from the occupancy grid X_{occ} with the progressive probabilistic Hough transform [43].
- 2) the waypoints are clustered using the density-based algorithm DBSCAN [44] that creates a variable number of clusters depending on the space density of the waypoints.
- 3) the points in each cluster are ordered, projecting them along the normal to the direction estimated in step 1).
- 4) the clusters are merged with a heuristic approach based on their position and size until two main groups representing the two sides of the field are reached.
- 5) the final ordered list of waypoints $W_{ord} \in \mathbb{N}^{l \times 2}$ is obtained by selecting the points from the two main clusters following an A-B-B-A scheme.

The global path is generated from the ordered list W_{ord} in two steps. The intra-row paths are obtained with [39], which exploits a gradient-based planner between the starting and the ending waypoints of each row. On the other hand, the inter-row paths are generated with a circular pattern in order to keep a safe margin from the end of the rows to avoid collision during the turns. Considering an end-row waypoint p_i and the successive point that starts the following row p_{i+1} , the waypoints are firstly moved along the row direction to get an end-row margin d_{er} :

$$\begin{aligned} p_i^{shifted} &= p_i + d_{er} \begin{bmatrix} \cos\alpha \\ \sin\alpha \end{bmatrix} \\ p_{i+1}^{shifted} &= p_{i+1} + (d_{er} + \Delta d) \begin{bmatrix} \cos\alpha \\ \sin\alpha \end{bmatrix} \end{aligned} \quad (4)$$

where α is the angle estimated during the post-processing steps and Δd is the distance between the two points along the row direction:

$$\Delta d = (p_i - p_{i+1}) \cdot \begin{bmatrix} \cos\alpha \\ \sin\alpha \end{bmatrix} \quad (5)$$




FIGURE 5. Graphical representation of the architecture of the segmentation network. Features at different resolutions feed the segmentation head that combines them, producing the output binary map \hat{Y} .

The end-row margin d_{er} can be selected depending on the meters/pixel resolution of the occupancy grid in order to have a target margin in meters in the real environment. A circular interpolation is adopted to connect the shifted points by linearly interpolating the angles considering the mean point as the center of the circumference. The whole sequence of points obtained by the intra-row and inter-rows planning creates the global path $P = \{(x, y) | x, y \in \mathbb{R}\}$, defined in the reference system of the occupancy grid X_{occ} . If the field map is georeferenced, it is possible to convert the global path into a list of geographic coordinates that can be directly used in the local planning phase to control the UGV motion. In Fig. 3, an aerial view of a field is shown, together with the predicted waypoints in yellow and the global path in red.

The overall segmentation network acts as a function H_{seg} , parameterized by Θ , that at each temporal instant t takes as input the RGB frame from the onboard camera $X_{rgb} \in \mathbb{R}^{h \times w \times c}$ and produces a binary map, $\hat{X}_{seg} \in \mathbb{R}^{h \times w}$ with h , w and c as height, width and channels, respectively. The output positive class segments the crops and the foliage in the camera view. Ideally, it should be equally split on the sides of the frame for a perfectly centered path. Successively, the semantic information of the row, \hat{X}_{seg} , is used in conjunction with its corresponding depth map to control all movements of the platform inside the crop rows.

Among all recent real-time semantic segmentation models, we carefully select an architecture that guarantees high accuracy levels by also containing hardware costs, optimization simplicity, and computational load. Indeed, the segmentation-based control does not considerably benefit from fine-grained predictions and elaborated encoder-decoder networks, [45], or two-pathway backbones, [46], does not bring any considerable improvement. Therefore, we adopt a very lightweight backbone, MobileNetV3 [47], followed by a reduced version of the Atrous Spatial Pyramid Pooling module, [48], to capture richer contextual information with minimal computational impact. Indeed, the output of the last layer of the backbone can not be used

directly to predict the segmentation mask due to the lack of spatial details.

The overall architecture is depicted in Fig. 5. The backbone, with repeated spatial reductions, extracts contextual information, and two of its branches at different resolutions feed the segmentation head. One layer applies atrous convolution to the 1/16 resolution to extract denser features, and the other one is used to add a skip connection from the 1/4 resolution to work with more detailed information. Finally, in order to maintain real-time performance even without hardware accelerators, we employ a 224×224 low-resolution input. So, we rescale the global average pooling layer setting the kernel size to 12×12 with strides (4,5). Additionally, to have equal input and output dimensions, we add a final bilinear upsampling with a factor of 8 at the end of the segmentation head.

Algorithm 1 Segmentation-Based Algorithm

Input: X_{ctrl}^t : Pre-processed segmented image

Output: v_x, ω_z : Continuous control commands

- 1: noise_reduction_function()
- 2: **for** $i=0, \dots, w$ **do**
- 3: $c \leftarrow \text{sum_columns}(X_{ctrl}^t)$
- 4: **end for**
- 5: $\text{zeros} \leftarrow \text{list_zero_clusters}(c)$
- 6: $\text{max_cluster} \leftarrow \text{find_max_cluster}(\text{zeros})$
- 7: **if** $\text{cluster_length}(\text{max_cluster}) \geq \text{anomaly}_{th}$ **then**
- 8: $v_x, \omega_z \leftarrow 0, 0$
- 9: **else**
- 10: $\text{compute_cluster_center}()$
- 11: $v_x, \omega_z \leftarrow \text{control_function}()$
- 12: **end if**

C. SEGMENTATION-BASED CONTROL

The segmentation masks $\hat{X}_{seg} \in \mathbb{R}^{h \times w}$ provided by the deep neural network are post-processed and fed into a custom control algorithm in order to generate consistent velocity commands to drive the UGV inside the inter-row space and maintain as much as possible the inter-row centrality. As in [37], we compute a sum of S segmentation maps along with an intersection with depth information provided by an RGB-D camera in order to obtain a more stable control. First, we pick S consecutive segmentation maps at times $\{t - S, \dots, t\}$ and we fuse them

$$\hat{X}_{cumSeg}^t = \sum_{n=0}^S \hat{X}_{seg}^{t-n} \quad (6)$$

then, we join the depth information $X_{depth}^t \in \mathbb{R}^{h \times w}$ to reduce the line of sight of the actual scene and remove some background noise. The line of sight is limited of a fixed value generating a binary map $X_{depthT}^t \in \mathbb{N}^{h \times w}$, as follows:

$$X_{depthT}^t_{i=0, \dots, h, j=0, \dots, w}(i, j) = \begin{cases} 0, & \text{if } (X_{depth}^t)_{i,j} \geq d_{depth} \\ 1, & \text{if } (X_{depth}^t)_{i,j} < d_{depth} \end{cases} \quad (7)$$

where d_{depth} is a fixed experimental scalar. Finally, exploiting an interception operation between the cumulative output \hat{X}_{cumSeg}^t , computed in equation (6) and the binary map X_{depthT}^t previously generated, we obtain the pre-processed input $X_{ctrl}^t \in \mathbb{R}^{h \times w}$ for the control algorithm:

$$X_{ctrl}^t = \hat{X}_{cumSeg}^t \cap X_{depthT}^t \quad (8)$$

In the X_{ctrl}^t binary map 1 stands for obstacles and 0 free-space. The segmentation-based control algorithm is developed building over the SPC algorithm presented in [37]. Indeed, we propose a simplified version of that control function to avoid useless conditional blocks and obtain a more real-time control algorithm. Algorithm 1 contains the pseudo-code of the proposed custom control, that starting from a pre-processed segmented image X_{ctrl}^t , is responsible for computing the driving velocity commands. First, a noise reduction function gets rid of undesired noise in the bottom part of the cumulative segmentation mask X_{ctrl}^t due to grass on the terrain, which may be wrongly segmented by the neural network. To perform such an operation, we compute the sum over rows of X_{ctrl}^t obtaining an array $g_{noise} \in \mathbb{R}^h$, then we set $X_{ctrl}^t(indices,:) = 0$, where *indices* contains the matrix-row indices such that $g_{noise} < th_{noise}$, with $th_{noise} = 0.03 \cdot \max(g_{noise})$ as threshold. We perform such an operation because, in an ideal segmentation mask, there are no 1s at the top of the image and on the bottom, whilst the majority of them are supposed to be in the central belt. After the noise reduction phase, we store the sum over columns of the obtained matrix X_{ctrl}^t in the array $c \in \mathbb{R}^w$, that contains the number of segmented trees for each column. Therefore, every zero in c is a potential empty space where to route the mobile platform. Then, we select the clusters of zeros in c , which are the groups of consecutive zeros, in order to store them in the list *zeros*. Next, we look for the largest cluster of zeros *max_cluster* and in case of the length of such cluster is over an empirically chosen threshold, *anomaly_th* = $0.8 \cdot w$, the driving commands are set to zero value because it means the provided cumulative segmentation mask X_{ctrl}^t has more zeros than ones, that is an anomaly, so for safety reason the mobile platform is stopped. While, in case of no anomalies, we compute the cluster center that is given as input to the control function. The identified cluster contains the obstacle-free space information that can be exploited to safely drive the mobile platform; as a consequence, the linear and angular velocities are computed using the center of the selected cluster, which ideally corresponds to the center position of the row in front of the UGV. The desired velocities are obtained by means of two custom functions:

$$\omega_z = -\omega_{z,gain} \cdot d \quad (9)$$

$$v_x = v_{x,max} \cdot \left[1 - \left[\frac{d^2}{\left(\frac{w}{2}\right)^2} \right] \right] \quad (10)$$

where $\omega_{z,gain} = 0.01$ and $v_{x,max} = 1.0$ are two constants that define the angular gain and maximum linear velocity of the mobile platform, respectively, w is the width of X_{ctrl}^t and d is

defined as:

$$d = x_c - \frac{w}{2} \quad (11)$$

with x_c center coordinate of the selected cluster. Equation (9) represents the angular velocity control function, while equation (10) has been used to compute the linear velocity, as in [37] and [49]. Eventually, the control velocity commands sent to the actuators are smoothed using the Exponential Moving Average (EMA), formalized in equation (12), in order to prevent the mobile platform from sharp motion.

$$EMA_t = EMA_{t-1} \cdot (1 - \alpha_{EMA}) + \begin{bmatrix} v_x \\ \omega_z \end{bmatrix} \cdot \alpha_{EMA} \quad (12)$$

where t is the time step and $\alpha_{EMA} = 0.18$ is the multiplier for weighting the EMA, which value is found experimentally.

D. LOCAL PATH PLANNING POLICY

Once the global path P and the start/end rows waypoints W_{ord} have been correctly generated, we exploit the provided information to locally navigate throughout the whole field.

In the case of lush vegetation and thick canopies that may distort GNSS signals inside the inter-row space, the local navigation problem is solved using the synergy of two different algorithms according to the kind of navigation requested in a determined place of the considered crop:

- 1) Inside the inter-row space: we exploit the custom control algorithm, described in Section III-C, based on the segmentation information provided by the deep neural network, in order to overcome localization inaccuracies due to blocked GNSS signals by overgrown plant vegetation.
- 2) Switch between different rows: we use the standard DWA, well described in [50], with fine-tuned parameters to safely switch between two rows following the circular path generated in Section III-B, since outside the inter-row space a clear view of satellites and sky should be available.

The choice of the suitable algorithm happens by comparing the estimated position by the EKF localization filter and the provided start/end rows waypoints; in the case of start row recognition, the local path planning policy selects the segmentation-based control, otherwise, it uses the DWA, as shown in Fig. 2. The comparison happens by computing a simple Euclidean distance between the waypoint of the start/end rows and the estimated positions. In case the calculated distance is lower than a threshold, *waypoint_th* = 0.5, the algorithm considers the waypoint as reached and selects the right local controller.

On the other hand, during specific year periods, when plant vegetation is not so dense and overgrown, the local path planning policy exploits only DWA to navigate throughout the whole field, thanks to a good view of both the sky and satellites in every place of the crop, following the complete global path generated in Section III-B.

IV. EXPERIMENTAL RESULTS AND DISCUSSIONS

In this section, we describe the main experiments conducted in simulation and real environments in order to better validate the algorithms. First, we discuss the creation of synthetic datasets for the training of the two deep neural networks. Then, we illustrate the training process and the optimization techniques adopted to minimize inference costs. Finally, we conclude with the simulation and real environment evaluations.

All the code has been developed ROS-compatible in order to exploit some of the most used ROS packages in robotics research, as *move_base*⁴ and *robot_localization*⁵ packages, that offer ready-to-use local planners and basic localization methods, respectively. Moreover, all tests have been performed using Ubuntu 18.04 and ROS Melodic.

A. DATASETS CREATION

As far as the two presented deep neural networks are concerned, they require to be trained on specific datasets according to the desired final application. Supervised learning training algorithms are the easiest to adopt and usually have the best final results. However, they all require supervision by means of a labeled training dataset. That greatly affects costs and makes data collection complex and time-consuming. Therefore, we make exclusive use of synthetic data and domain randomization [24] to enable affordable supervised learning and simultaneously bridge the domain gap between simulation and reality.

For DeepWay, a dataset of random occupancy grids is generated. All parameters, such as number, orientation, depth, and length of the rows, are randomly selected. The coordinates of the starting and ending points of each row are generated by geometrical reasoning, and the occupancy grid is obtained with circles of different radii for each location in between two corresponding points. Random holes in the rows are also created to increase the variability of the images. Ground truth waypoints are obtained to always lay inside the rows, since we experimentally found it helps the network prediction. Given two starting/ending points a and b of two successive rows, we compute the target waypoint as follows:

$$p_{a,b} = \begin{bmatrix} 0 & \mp 1 \\ \pm 1 & 0 \end{bmatrix} \frac{a-b}{2} + \frac{a+b}{2} \quad (13)$$

that corresponds to a ± 90 degrees rotation around the mean point, where the sign is selected to make the waypoint inside the row. We train DeepWay with a total of 3000 synthetic images, and we validate it with 100 satellite images taken from the Google Maps database and manually annotated.

On the other hand, the segmentation neural network requires a set of RGB images along with segmentation masks as inputs in order to be correctly trained. As a consequence, inspired by the work of Tejaswi Digumarti et al. [32], we generate synthetic RGB images coupled with the corresponding




FIGURE 6. An example of synthetic RGB image (a) and the corresponding segmentation mask (b).

segmentation masks. For such purpose, we exploit Blender⁶ 2.8, which is an open-source 3D computer graphics software compatible with Python language, and the Modular Tree⁷ addon to speed up the tree generation. We design four main scenarios: two single different trees, one group of heterogeneous trees, and one row-based scenario with various backgrounds and soils. Then, exploiting Python language compatibility of Blender, we write a script able to automatically capture RGB images and the corresponding segmentation masks of the scene from different positions with respect to the central reference frame and with different illuminations conditions in order to obtain as many as possible a random and complete synthetic dataset. An example of a synthetic RGB image, along with the corresponding segmentation mask, is shown in Fig. 6. Every single rendering takes about 30 seconds, multiplied by a total of 2776 rendering, which is about 23 hours of continuous work on an RTX 2080 GPU. That total number of images in conjunction with transfer learning allows to train a segmentation network with high generalization capabilities while minimizing generation data costs. The overall segmentation training dataset is composed of 2776 RGB synthetic images for training and 100 manually annotated images for testing, acquired in a real environment (Italy, Valle San Giorgio di Baone).

B. NETWORKS TRAINING AND OPTIMIZATION

For training both networks, we employ the TensorFlow⁸ 2 framework on a PC with 32-GB RAM, an Intel i7-9700K CPU, and an Nvidia 2080 Super GP-GPU.

DeepWay is trained following the methodology presented in [1], with an input dimension of $H = W = 800$ and $N = 2$ Residual Reduction modules that result in a subsampling factor of $k = 8$ and output dimensions of $U_H = U_W = 100$. We select a kernel size of 5 and 16 filters for all the convolutional layers, except the first and last ones, which have kernel sizes of 7 and 3, respectively. These hyperparameters have been selected by performing a grid search over reasonable sets of values and adopting those that experimentally provided the best convergence. As a loss function, a weighted mean squared function (L_2) is used to compensate for the higher number of negative cells (i.e., with

⁶<https://www.blender.org/>

⁷https://github.com/MaximeHerpin/modular_tree/tree/blender_28

⁸<https://www.tensorflow.org>

⁴http://wiki.ros.org/move_base

⁵http://docs.ros.org/en/melodic/api/robot_localization/html/index.html

no waypoint in the target image) with respect to positive ones. We set these weights to 0.7 for the positive cells and 0.3 for the negative. The default distance threshold for the waypoints suppression algorithm is set to $d_{thr} = 8$ pixels, that is the minimum inter-row distance of our dataset, and the confidence threshold to $c_{thr} = 0.9$, in order to select the most confident predictions only. As in standard object detection algorithms, we adopt the Average Precision (AP) as a metric for the prediction quality. We compute the AP at different distance ranges d_r . A prediction is considered a True Positive (TP) only if it falls within a distance d_r from the target waypoint. On the 100 real-world images, we reach an AP of 0.9794 with $d_r = 8$ pixels, 0.9558 with 4 pixels and 0.7500 with 2 pixels.

TABLE 1. Comparison between different devices' energy consumption and inference performances with graph optimization (G.O.) and weight precision (W.P).

Device	GO	WP	Latency [ms]	E_{net} [mJ]	Size [MB]
RTX 2080	N	FP32	28 ± 109	819	9.3
	Y	FP32	0.1 ± 0.3	52	7.4
	Y	FP16	0.1 ± 0.2	39	4.9
Cortex-A57	Y	FP32	111 ± 0.9	166	4.2
	Y	FP16	111 ± 2.3	165	2.2
Cortex-A76	Y	FP32	55.4 ± 10.6	210	4.2
	Y	FP16	65.3 ± 9.5	248	2.2

Regarding the segmentation network, we train our model applying transfer learning to the selected backbone. Indeed, rather than using randomly initialized weights, we exploit MobileNetV3 variables derived from an initial training phase on the 1k classes and 1.3M images of the ImageNet dataset [51]. Moreover, we pre-trained the overall segmentation network with Cityscapes [52], a publicly available dataset with 30 different classes and 5000 fine annotated images. Finally, we adopt a strong data augmentation with random crops, brightness, saturation, contrast, rotation, and flips. All together, those techniques largely improve the final robustness of the model and its final generalization capability with a reduced number of training samples. We train the network with stochastic gradient descent with a learning rate of 0.03 and Intersection over Unit (IoU) as a loss function. The accuracy over the test set is 0.8, with an IoU of 0.46. In comparison, the accuracy of the validation set with 0.1 of the synthetic dataset is 0.86 with a domain gap of 0.08. Moreover, our experimentation shows that larger input sizes improve segmentation over the synthetic dataset but greatly reduce accuracy over real images.

The trained network is optimized to reduce latency, inference cost, memory, and storage footprint. That is obtained with two distinct techniques: model pruning and quantization. The first simplifies the topological structure, removing unnecessary parts of the architecture, and favors a more sparse model introducing zeros to the parameter tensors. Subsequently, with quantization, we reduce the precision of the numbers used to represent model parameters from float32 to float16. That can be accomplished with a post-training




FIGURE 7. A visual representation of the simulation environment.

quantization procedure. In Table 1, experimentation results with some reference architectures are summarized.

C. PLATFORM HARDWARE AND SENSORS SETUP

As a mobile platform, we select the Jackal UGV by Clearpath Robotics, which can be briefly described as a small and weatherproof rover (IP62 code) with a 4×4 high-torque drivetrain. It is highly customizable and ROS-compatible, allowing fast deployment and algorithm testing. All the algorithms run on Jackal's onboard Mini-ITX PC with a CPU Intel Core i3-4330TE @2.4GHz and 4GB DDR3 RAM. For what concern the localization sensors, the RTK enabled GNSS receiver is the Piksi Multi by Swift Navigation⁹ mounted on an evaluation board, which provides easy input/output communication with the receiver (used acquisition rate: 10 Hz), while the inertial measurements are provided by the MPU-9250 IMU, with an acquisition rate of about 100 Hz. In addition, to get a front view of the environment, we selected the Intel Realsense D455 RGBD camera, which provides frames at 30 FPS and is mounted on the front part of Jackal's top plate. Finally, the odometry is provided by the onboard quadrature encoders that are able to run at 78000 pulses/m. As mentioned in Section II, the IMU and GNSS receiver data are fused by means of an EKF in order to obtain a global position estimate of the mobile platform time by time, described in terms of x, y, yaw . However, GNSS positioning is highly inaccurate, about $3m - 5m$, without implementing any corrections technique. As a consequence, we provide RTK corrections to Piksi Multi receiver, coming from the SPIN3 GNSS¹⁰ of Piemonte, Lombardia, and Valle d'Aosta, through the Internet. Then, the GNSS receiver directly uses such corrections to obtain more reliable and accurate global position estimates, with an error range of $[0.05, 0.10]m$, in a clear view of the sky and a good antenna position. We exploit such corrections service because it is entirely free prior to an online subscription, and it can send out RTK corrections through the Internet.

D. SIMULATION ENVIRONMENT EVALUATION

All the presented pipelines are tested in a simulation environment prior to real-world tests in order to check the basic performances and perform the first experimental setting

⁹<https://www.swiftnav.com/>

¹⁰<https://www.spingnss.it/spiderweb/frmIndex.aspx>

of various gains and thresholds. First, we build a custom simulation environment made of vine plants organized in rows, and a bumpy and uneven terrain, as shown in Fig. 7, using the Gazebo¹¹ simulator, that is ROS-compatible and open-source. Moreover, it provides advanced 3D graphics, dynamics simulation, and several plugins to simulate sensors, such as GNSS, IMU, and cameras.

Then, we compare the UGV trajectory obtained with the proposed methodology with a ground truth line in order to evaluate different error metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Standard Deviation (σ), as shown in Table 2. The ground truth is computed in two steps:

- 1) manual annotation of GNSS simulated positions that correspond to an ideal global path: centered in the inter-row space and with a safe distance from crops switching between two different rows.
- 2) linear interpolation of such points.

We have performed six different tests in two different simulation environments (varying the vine plants positions and maintaining a row-based organization). The obtained performances are promising in terms of MAE, RMSE, and σ , as shown in Table 2. The worst attained results are stated by an RMSE=0.265 m and an MAE=0.217 m in the sixth test, while the best one is achieved in the first test with an RMSE=0.089 m and an MAE=0.068 m. Finally, the mean time to complete a test is about 4 minutes with a maximum speed of 0.5 m/s. All considered, the overall achieved performances are satisfactory, taking into account the worse results obtained in the last two tests where the vineyard rows are organized with a slight curvature shape.

TABLE 2. Comparison of error metrics in six tests performed in two different simulation environments. The first row describes the number of visited rows in the corresponding test.

	T1	T2	T3	T4	T5	T6
N. rows	4	4	4	4	4	4
Min. Error [m]	0.001	0.002	0.002	0.001	0.002	0.002
Max. Error [m]	0.726	0.678	0.600	0.633	1.21	1.21
MAE [m]	0.068	0.085	0.077	0.082	0.215	0.217
RMSE [m]	0.089	0.100	0.092	0.096	0.263	0.265
σ [m]	0.057	0.053	0.050	0.050	0.152	0.152

E. REAL ENVIRONMENT EVALUATION

The overall system is extensively tested in two real environment scenarios with multiple experiments in different seasonal periods (Fig. 1): a vineyard and a pear orchard, shown in Fig. 8, for a total of more than 80 hours of experimentation from 9 a.m. to 6 p.m., but without particular adverse weather conditions (e.g., rain, snow, fog). Moreover, all the tests were performed using the same hardware and software setup to obtain consistent data. The vineyard is located in Grugliasco and managed by the Department of Agricultural, Forestry, and Food Sciences of Università degli Studi di Torino (UNITO). In contrast, the pear orchard is located in Montegrosso d'Asti and managed by Mura Mura

farm. The first scenario has an inter-row space of about 2.80m and a height of about 2.0 m, while the second is organized in rows with an inter-row space of 4.50m and a height of about 3.0m. The terrain of the fields at the time of the tests did not present major unevenness like significant obstacles. However, terrain slopes, irregularities, and holes, together with small obstacles such as stones or fallen fruits, can suddenly drift the robot's trajectory. Thus, the visual-based control algorithm aims to mitigate these unexpected disturbances that are not easy for an avoidance planner to detect and handle.




FIGURE 8. A visual representation of the real-world testing environments.

TABLE 3. Comparison of different error metrics in three different tests performed in the pear orchard. The second column describes the number of visited rows in the corresponding test.

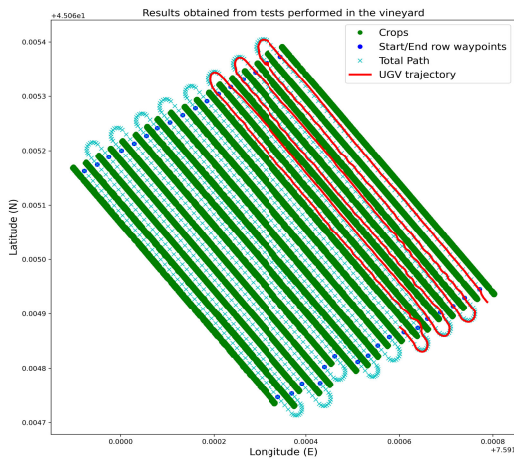
Test	N. rows	Min. Error [m]	Max. Error [m]	MAE [m]	RMSE [m]	σ [m]
Test n. 1	4	0.008	1.395	0.523	0.627	0.351
Test n. 2	4	0.002	1.105	0.457	0.551	0.315
Test n. 3	2	0.007	1.320	0.659	0.755	0.375

TABLE 4. Comparison of error metrics in three tests performed in the vineyard. The second column describes the number of visited rows in the corresponding test.


Test	N. rows	Min. Error [m]	Max. Error [m]	MAE [m]	RMSE [m]	σ [m]
Test n. 1	4	0.013	0.621	0.296	0.332	0.160
Test n. 2	6	0.006	0.598	0.218	0.240	0.119
Test n. 3	6	0.003	0.720	0.204	0.246	0.145

The errors, described in Table 3 and Table 4, are computed comparing the RTK-GNSS positions provided by the Piksi Multi receiver and the global path provided to the navigation system. All of this is possible due to the high accuracy of GNSS estimated positions, thanks to a clear view of the sky and a good position of the high-end antenna on the UGV. All the collected data are represented in latitude and longitude coordinates. However, for analysis purposes, they have been transformed in meters with respect to a known GNSS coordinate of the georeferenced occupancy grid map: the top left corner pixel.

¹¹<http://gazebosim.org/>



(a) Test n. 3 in the vineyard scenario



(b) Test n. 2 in the orchard scenario

FIGURE 9. A visual representation of the obtained results. Both images contain the path followed by the UGV (red line), the provided global path (cyan x), the start/end row waypoints (blue dots) and the crop (green dots).

Table 3 and Table 4, together with the visual representation of Fig. 9, shows some numerical results obtained by the proposed novel approach and demonstrates that a methodology that exploits semantic segmentation along with a standard navigation approach based on the GNSS is able to provide complete and reliable navigation throughout the whole row-based crop. The results obtained in the pear orchard (Table 3) are slightly worse than the vineyard ones; however, this effect may be due to the greater inter-row space of the pear orchard with respect to the vineyard one. Eventually, the mean time of a single test is about 25 minutes, while the maximum velocity of the UGV is 0.5 m/s. All considered, the overall performance achieved in terms of MAE and RMSE is adequate for the setup of low-cost sensors and demonstrates the effectiveness of our approach.

V. CONCLUSION AND FUTURE WORK

We presented a novel, affordable algorithmic pipeline for autonomous navigation in row-based crops. Our methodology is a complete solution that covers all navigation stages,

from global to local path planning, only relying on low-cost, low-range sensors. Moreover, the system efficiently tackles GNSS unreliability in presence of lush vegetation and thick canopies, allowing the platform to autonomously navigate in all seasonal periods. The experimental results in real fields show that the proposed solution is reliable and precise, safely guiding the autonomous platform for several hours of testing through orchards and vineyards. The control algorithm maintains a better central trajectory in rows with an inter-row space smaller than 3 m compared to larger rows. Finally, the adopted domain generalization and optimization techniques greatly make training and inference of deep neural network less time and computational costly. This work paves the way for further research on low-cost navigation solutions for agricultural vehicles of small and medium size, demonstrating that AI-based solutions can achieve, in combination with standard techniques, robust performance also in real-world conditions. Further works will aim at assessing our proposed algorithmic pipeline onto a more cumbersome vehicle and introducing a collision avoidance algorithm in the segmentation-based control.

REFERENCES

- [1] V. Mazzia, F. Salvetti, D. Aghi, and M. Chiaberge, "DeepWay: A deep learning waypoint estimator for global path generation," *Comput. Electron. Agricult.*, vol. 184, May 2021, Art. no. 106091.
- [2] F. Salvetti, S. Angarano, M. Martini, S. Cerrato, and M. Chiaberge, "Waypoint generation in row-based crops with deep learning and contrastive clustering," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*. Cham, Switzerland: Springer, 2022, pp. 203–218.
- [3] J. J. Roldán et al., "Robots in agriculture: State of art and practical experiences," in *Service Robots*, vol. 4. InTech, Jan. 2018, doi: 10.5772/intechopen.69874.
- [4] G. Zhang, T. Xu, Y. Tian, H. Xu, J. Song, and Y. Lan, "Assessment of rice leaf blast severity using hyperspectral imaging during late vegetative growth," *Australas. Plant Pathol.*, vol. 49, no. 5, pp. 571–578, Sep. 2020.
- [5] A. Feng, J. Zhou, E. D. Vories, K. A. Sudduth, and M. Zhang, "Yield estimation in cotton using UAV-based multi-sensor imagery," *Biosyst. Eng.*, vol. 193, pp. 101–114, May 2020.
- [6] D. Deshmukh, D. K. Pratihari, A. K. Deb, H. Ray, and N. Bhattacharyya, "Design and development of intelligent pesticide spraying system for agricultural robot," in *Proc. Int. Conf. Hybrid Intell. Syst.* Springer, Dec. 2020, pp. 157–170.
- [7] A. Khaliq, V. Mazzia, and M. Chiaberge, "Refining satellite imagery by using UAV imagery for vineyard environment: A CNN based approach," in *Proc. IEEE Int. Workshop Metrol. Agricult. Forestry (MetroAgriFor)*, Oct. 2019, pp. 25–29.
- [8] P. Radoglou-Grammatikis, P. Sarigiannidis, T. Lagkas, and I. Moscholios, "A compilation of UAV applications for precision agriculture," *Comput. Netw.*, vol. 172, May 2020, Art. no. 107148.
- [9] L. Comba, P. Gay, J. Primicerio, and D. Riccauda Aimonino, "Vineyard detection from unmanned aerial systems images," *Comput. Electron. Agricult.*, vol. 114, pp. 78–87, Jun. 2015.
- [10] R. Sparrow and M. Howard, "Robots in agriculture: Prospects, impacts, ethics, and policy," *Precis. Agricult.*, vol. 22, no. 3, pp. 818–833, Jun. 2021.
- [11] T. Duckett et al., "Agricultural robotics: The future of robotic agriculture," 2018, *arXiv:1806.06762*.
- [12] S. Garg, N. Sünderhauf, F. Dayoub, D. Morrison, A. Cosgun, G. Carneiro, Q. Wu, T.-J. Chin, I. Reid, S. Gould, P. Corke, and M. Milford, "Semantics for robotic mapping, perception and interaction: A survey," 2021, *arXiv:2101.00443*.
- [13] S. Saeedi et al., "Navigating the landscape for real-time localization and mapping for robotics and virtual and augmented reality," *Proc. IEEE*, vol. 106, no. 11, pp. 2020–2039, Nov. 2018.

- [14] T. Mota, M. Sridharan, and A. Leonardis, "Integrated commonsense reasoning and deep learning for transparent decision making in robotics," *SN Comput. Sci.*, vol. 2, no. 4, p. 242, 2021.
- [15] S. J. LeVoi, P. A. Farley, T. Sun, and C. Xu, "High-accuracy adaptive low-cost location sensing subsystems for autonomous rover in precision agriculture," *IEEE Open J. Ind. Appl.*, vol. 1, pp. 74–94, 2020.
- [16] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015, doi: [10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [17] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *J. Field Robot.*, vol. 37, no. 3, pp. 362–386, Apr. 2020.
- [18] A. Boschi, F. Salvetti, V. Mazzia, and M. Chiaberge, "A cost-effective person-following system for assistive unmanned vehicles with deep learning at the edge," *Machines*, vol. 8, no. 3, p. 49, Aug. 2020.
- [19] M. Martini, S. Cerrato, F. Salvetti, S. Angarano, and M. Chiaberge, "Position-agnostic autonomous navigation in vineyards with deep reinforcement learning," in *Proc. IEEE 18th Int. Conf. Autom. Sci. Eng. (CASE)*, Aug. 2022, pp. 477–484.
- [20] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2704–2713.
- [21] V. Mazzia, A. Khaliq, F. Salvetti, and M. Chiaberge, "Real-time apple detection system using embedded systems with hardware accelerators: An edge AI application," *IEEE Access*, vol. 8, pp. 9102–9114, 2020.
- [22] M. S. N. Kabir, M.-Z. Song, N.-S. Sung, S.-O. Chung, Y.-J. Kim, N. Noguchi, and S.-J. Hong, "Performance comparison of single and multi-GNSS receivers under agricultural fields in Korea," *Eng. Agricult., Environ. Food*, vol. 9, no. 1, pp. 27–35, Jan. 2016.
- [23] S. Marden and M. Whitty, "GPS-free localisation and navigation of an unmanned ground vehicle for yield forecasting in a vineyard," in *Proc. 13th Int. Conf. Intell. Auton. Syst. (IAS)* (Recent Advances in Agricultural Robotics), Jul. 2014.
- [24] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *Proc. IEEE/RSS Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2017, pp. 23–30.
- [25] O. Ly, H. Gimbert, G. Passault, and G. Baron, "A fully autonomous robot for putting posts for trellising vineyard with centimetric accuracy," in *Proc. IEEE Int. Conf. Auto. Robot Syst. Competitions*, Apr. 2015, pp. 44–49.
- [26] S. J. Moorehead, C. K. Wellington, B. J. Gilmore, and C. Vallespi, "Automating orchards: A system of autonomous tractors for orchard maintenance," in *Proc. IEEE Int. Conf. Intell. Robots Syst., Workshop Agricult. Robot.*, Oct. 2012.
- [27] S. Bonadies and S. A. Gadsden, "An overview of autonomous crop row navigation strategies for unmanned ground vehicles," *Eng. Agricult., Environ. Food*, vol. 12, no. 1, pp. 24–31, Jan. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S188183661730188X>
- [28] W. Winterhalter, F. Fleckenstein, C. Dornhege, and W. Burgard, "Localization for precision navigation in agricultural fields—Beyond crop row following," *J. Field Robot.*, vol. 38, no. 3, pp. 429–451, May 2021.
- [29] S. Zaman, L. Comba, A. Biglia, D. Riccauda Aimonino, P. Barge, and P. Gay, "Cost-effective visual odometry system for vehicle motion control in agricultural environments," *Comput. Electron. Agricult.*, vol. 162, pp. 82–94, Jul. 2019.
- [30] I. Nevliudov, S. Novoselov, O. Sychova, and S. Tesliuk, "Development of the architecture of the base platform agricultural robot for determining the trajectory using the method of visual odometry," in *Proc. IEEE XVII Int. Conf. Perspective Technol. Methods MEMS Design (MEMSTECH)*, May 2021, pp. 64–68.
- [31] Y. Ma, W. Zhang, W. S. Qureshi, C. Gao, C. Zhang, and W. Li, "Autonomous navigation for a wolfberry picking robot using visual cues and fuzzy control," *Inf. Process. Agricult.*, vol. 8, no. 1, pp. 15–26, Mar. 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214317319303269>
- [32] S. Tejaswi Digumarti, L. M. Schmid, G. M. Rizzi, J. Nieto, R. Siegwart, P. Beardsley, and C. Cadena, "An approach for semantic segmentation of tree-like vegetation," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 1801–1807.
- [33] P. Kim, B. Coltin, and H. J. Kim, "Low-drift visual odometry in structured environments by decoupling rotational and translational motion," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 7247–7253.
- [34] D. Anthony and C. Detweiler, "UAV localization in row crops," *J. Field Robot.*, vol. 34, no. 7, pp. 1275–1296, Oct. 2017.
- [35] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Comput. Electron. Agricult.*, vol. 147, pp. 70–90, Apr. 2018.
- [36] D. Aghi, V. Mazzia, and M. Chiaberge, "Autonomous navigation in vineyards with deep learning at the edge," in *Proc. Int. Conf. Robot. Alpe-Adria Danube Region*. Cham, Switzerland: Springer, Jun. 2020, pp. 479–486.
- [37] D. Aghi, S. Cerrato, V. Mazzia, and M. Chiaberge, "Deep semantic segmentation at the edge for autonomous navigation in vineyard rows," in *Proc. IEEE/RSS Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 3421–3428.
- [38] J. Zoto, M. A. Musci, A. Khaliq, M. Chiaberge, and I. Aicardi, "Automatic path planning for unmanned ground vehicle using UAV imagery," in *Proc. Int. Conf. Robot. Alpe-Adria Danube Region*. Springer, 2019, pp. 223–230.
- [39] S. Cerrato, D. Aghi, V. Mazzia, F. Salvetti, and M. Chiaberge, "An adaptive row crops path generator with deep learning synergy," in *Proc. 6th Asia-Pacific Conf. Intell. Robot Syst. (ACIRS)*, Jul. 2021, pp. 6–12.
- [40] M. Samuel, M. Hussein, and M. B. Mohamad, "A review of some pure-pursuit based path tracking techniques for control of autonomous vehicle," *Int. J. Comput. Appl.*, vol. 135, no. 1, pp. 35–38, 2016.
- [41] D. Misra, "Mish: A self regularized non-monotonic activation function," 2019, [arXiv:1908.08681](https://arxiv.org/abs/1908.08681).
- [42] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis.*, Sep. 2018, pp. 3–19.
- [43] J. Matas, C. Galambos, and J. Kittler, "Robust detection of lines using the progressive probabilistic Hough transform," *Comput. Vis. Image Understand.*, vol. 78, no. 1, pp. 119–137, Apr. 2000.
- [44] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proc. KDD*, 1996, vol. 96, no. 34, pp. 226–231.
- [45] P. Hu, F. Perazzi, F. C. Heilbron, O. Wang, Z. Lin, K. Saenko, and S. Sclaroff, "Real-time semantic segmentation with fast attention," *IEEE Robot. Autom. Lett.*, vol. 6, no. 1, pp. 263–270, Jan. 2021.
- [46] C. Yu, C. Gao, J. Wang, G. Yu, C. Shen, and N. Sang, "BiSeNet v2: Bilateral network with guided aggregation for real-time semantic segmentation," 2020, [arXiv:2004.02147](https://arxiv.org/abs/2004.02147).
- [47] A. Howard, M. Sandler, B. Chen, W. Wang, L.-C. Chen, M. Tan, G. Chu, V. Vasudevan, Y. Zhu, R. Pang, H. Adam, and Q. Le, "Searching for MobileNetV3," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1314–1324.
- [48] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017, [arXiv:1706.05587](https://arxiv.org/abs/1706.05587).
- [49] D. Aghi, V. Mazzia, and M. Chiaberge, "Local motion planner for autonomous navigation in vineyards with a RGB-D camera-based algorithm and deep learning synergy," *Machines*, vol. 8, no. 2, p. 27, May 2020.
- [50] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Autom. Mag.*, vol. 4, no. 1, pp. 23–33, Mar. 1997.
- [51] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [52] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.




SIMONE CERRATO received the B.Sc. degree in computer engineering, in 2018, and the M.Sc. degree in mechatronic engineering, in 2020, with the thesis "GPS-Based Autonomous Navigation of Unmanned Ground Vehicles in Precision Agriculture Applications." He has been a Researcher with the Interdepartmental Centre on Service Robotics (PIC4SeR), Politecnico di Torino. His research interests include autonomous guidance algorithms for unmanned ground vehicles in precision agriculture applications, conducting an on-field testing activity in orchards and vineyards.



language processing and their embedded implementation for AI at the edge.

VITTORIO MAZZIA received the master's degree in mechatronics engineering from the Politecnico di Torino and the joint Ph.D. degree in electrical, electronics and communications engineering from PIC4SeR and SmartData, Politecnico di Torino, with the dissertation "Machine Learning Algorithms and their Embedded Implementation for Service Robotics Applications." His current research interests include deep learning applied to different tasks of computer vision and natural



interests include efficient deep learning models for robot perception and control, particularly highlighting key aspects of real-world applications like generalization and robustness.

SIMONE ANGARANO received the bachelor's degree in electronic engineering and the master's degree in mechatronic engineering from the Politecnico di Torino, in 2018 and 2020, respectively, where he is currently pursuing the Ph.D. degree in electrical, electronics, and communications engineering with the Interdepartmental Center for Service Robotics. He is also a member of the Interdepartmental Center for Service Robotics, Politecnico di Torino. His research



niques." He is currently working on machine learning applied to computer vision and image processing in robotics applications.

FRANCESCO SALVETTI received the bachelor's degree in electronic engineering and the master's degree in mechatronics engineering from the Politecnico di Torino, Italy, in 2017 and 2019, respectively, and the joint Ph.D. degree in electrical, electronics and communications engineering from PIC4SeR and SmartData, Politecnico di Torino, in 2023, with the thesis "Machine Learning Algorithms for Robotic Navigation and Perception and Embedded Implementation Techniques."



perception methods for service robotics, with a particular focus on precision agriculture applications. He is also a member of the Interdepartmental Center for Service Robotics, Politecnico di Torino.

ALESSANDRO NAVONE (Graduate Student Member, IEEE) received the master's degree in mechatronic engineering from the Politecnico di Torino, in 2022, where he is currently pursuing the Ph.D. degree in electrical, electronics, and communications engineering with the Interdepartmental Center for Service Robotics. His master's thesis "Learning Odometric Error in Mobile Robots With Neural Networks." His research



Service Robotics (PIC4SeR), investigating machine learning solutions for autonomous navigation in service mobile robots, with a particular focus on precision agriculture and social navigation.

MAURO MARTINI (Member, IEEE) received the master's degree (cum laude) in mechatronic engineering from the Politecnico di Torino, in 2020, and the Ph.D. degree in electrical, electronics and communications engineering from the Politecnico di Torino, in 2024, with the thesis "Machine Learning for Perception and Autonomous Navigation of Service Mobile Robots." He is currently carrying out his research activity in collaboration with the Interdepartmental Centre for



research interests include hardware implementation of neural networks and fuzzy systems and the design and implementation of reconfigurable real-time computing architectures.

MARCELLO CHIABERGE is currently an Associate Professor with the Department of Electronics and Telecommunications, Politecnico di Torino, Turin, Italy, where he is also the Co-Director of the Mechatronics Laboratory and the Director and the Principal Investigator of the Interdepartmental Center for Service Robotics (PIC4SeR), Turin. He has authored more than 100 articles accepted in international conferences and journals. He is the co-author of nine international patents. His

...

Open Access funding provided by 'Politecnico di Torino' within the CRUI CARE Agreement