BURLEIGH DODDS SERIES IN AGRICULTURAL SCIENCE

Advances in agri-food robotics

Edited by Professor Eldert van Henten, Wageningen University, The Netherlands and Professor Yael Edan, Ben-Gurion University of the Negev, Israel

E-CHAPTER FROM THIS BOOK

burleigh dodds
SCIENCE PUBLISHING

Implementing a digital twin for flexible operation of agricultural robotics

Frits K. van Evert, Trim Bresilla, Ard Nieuwenhuizen, Bram Veldhuisen and Jochen Hemming, Wageningen University & Research, The Netherlands; Giannis Avgoustakis, Spyros Fountas, Michael Koutsiaras and Nikos Mylonas, Agricultural University of Athens, Greece; Suzanne Baron, Agricultural Robotics and Automatization, France; Thanos Dritsopoulos and Markos Legas, Pegasus, Greece; Han Hilbrands and Jeroen Wolters, Smart Agri Technology BV, The Netherlands; Klaas Jan Hommes, Abemec BV, The Netherlands; Panagiotis Karagiannis, Sotiris Makris and George Michalos, University of Patras, Greece; Dimitrios Paraforos and Galibjon Sharipov, University of Hohenheim, Germany; Søren Marcus Pedersen, University of Copenhagen, Denmark; Bertrand Pinel, Terrena Societe Cooperative Agricole, France; Mladen Radisic, Foodscale Hub, Serbia; Juan F. Rascón and Carlos Rizzo, Eurecat, Centre Tecnològic de Catalunya, Spain; Raul Sanchez, Serrater, Spain; Alea Scovill, AgroIntelli, Denmark; Oriol Serra, Giropoma, Spain; and Josep Vidal, Teyme, Spain

1 Introduction

One of the most important challenges in European agriculture is the cost and scarcity of labour. The cost of labour affects the profitability of farming while the scarcity of labour threatens its very existence. The latter was illustrated vividly during the COVID-19 pandemic when many EU countries were not able to welcome field workers from abroad. As a result, some crops could not be harvested and had to be left on the field.[1]

[1] https://www.nytimes.com/2020/03/27/business/coronavirus-farm-labor-europe.html.

http://dx.doi.org/11.19103/AS.2023.0124.28

Chapter taken from: van Henten, E. and Edan, Y. (ed.), *Advances in agri-food robotics*, pp. 185–222, Burleigh Dodds Science Publishing, Cambridge, UK, 2024, (ISBN: 978 1 80146 277 8; www.bdspublishing.com)

The cost and shortage of labour may be addressed by adopting robotics. Robots have replaced labour in several sectors of the economy, but they have not yet reached widespread acceptance in agriculture. Robots have been proposed for all activities during the cropping cycle, from tillage to sowing, fertilization, crop protection and harvesting. In this chapter, we focus on mobile robots for spraying and mechanical weeding because these are dull, dirty, dangerous tasks that need to be done several times per season.

Agricultural robots are available commercially from vendors including Naïo,[2] Agreenculture,[3] AgroIntelli,[4] AgXeed[5] and FarmDroid.[6] Robotic research platforms include Flourish[7] and Deepfield[8]. Autonomous tractors are another class of robotic vehicles in agriculture, available commercially (Kubota,[9] only for rice, only in Japan) or as a concept (John Deere[10]). Alternatively, after-market kits are available to convert a conventional farm tractor into a robot (e.g. JCA Technologies[11] and Blue White Robotics[12]).

All these vehicles can execute basic autonomous behaviour: follow a crop row to the end of the field, make a turn and return in the next crop row. Path planning is a topic that has received ample attention, especially in the context of covering irregularly shaped fields with a minimum of overlap (Oksanen and Visala, 2009; Mier et al., 2022). It is relevant for human-operated machinery but will also be needed by autonomous vehicles. However, full autonomous behaviour is by far not realized. These vehicles must be brought to the field and pointed in the right direction. Planning and scheduling are capabilities that are not offered. Working autonomously and robustly, during a whole growing season, in a realistic environment, has not been demonstrated. Advanced route planning and dynamic re-planning were proposed 15 years ago and demonstrated in simulation (Bochtis et al., 2007); as far as we are aware, this system was never implemented. More recently, a route planning algorithm for an electrical agricultural vehicle was developed that not only takes into account hill slope and remaining battery charge but also here robust operation and dynamic response to unforeseen events are not considered (Hizatate and Noguchi, 2023).

Agricultural implements have also been equipped with autonomy. The main objective of these smart implements is to close the loop from perception to decision and action. Examples include the WEED-IT spot sprayer (Cloete, 2020) which prevents runoff of crop protection chemicals by applying them

[2] https://www.naio-technologies.com/en/.
[3] https://www.agreenculture.net/.
[4] https://agrointelli.com/.
[5] https://www.agxeed.com/.
[6] https://farmdroid.dk/.
[7] http://flourish-project.eu/ NOTE better to cite scientific papers.
[8] https://www.deepfield-robotics.com/en/ website no longer works but we can cite the scientific paper.
[9] https://www.youtube.com/watch?v=PG3vd3AsdfY.
[10] https://www.deere.co.uk/en/agriculture/future-of-farming/.
[11] https://jcatechnologies.com/.
[12] https://www.bluewhite.co/.

Published by Burleigh Dodds Science Publishing Limited, 2024.

only where weeds are detected by sensors in real-time, the Ecorobotix ARA spot sprayer (Anken and Latsch, 2022; Anken et al., 2022) and See & Spray (ASABE Staff, 2022). The Robocrop system (Tillett and Hague, 2006) and the Steketee IC-Weeder use cameras to guide a hoe around crop plants (Hemming et al., 2018). The 'Hill Control' control system from Rauch improves the distribution accuracy by changing the drop point, disc speed and metering quantity when spreading fertilizers with a centrifugal disk spreader, especially in hilly terrain.[13] Advances in sprayers involve not only boom control systems[14] but also more advanced systems that control the full functionality of a sprayer[15] are important towards achieving an even distribution of the spray deposit. However, the information gathered by the sensors in these cases is not used to monitor the quality of the action and it is not used to alter or optimize the operation of the robotic agricultural operation within the farming system as a whole.

Agricultural robots exist but are not widely used. An important reason is that agricultural robots are not flexible. As a first example, agricultural robots do not have a mechanism that adjusts or stops the robot when the process is not going well (insufficient weed control, crop is being damaged). As a second example, agricultural robots are not flexible in dealing with unforeseen events (weather, obstacles, tank empty, working lane not being wide enough, soil trafficability, etc.). As a result, agricultural robots need a lot of supervision to check their operation and to get them going again when they get stuck. This means that the real amount of labour saved will be much less than expected.

In contrast to agriculture, robots are widely used in industrial manufacturing. Digital Twin (DT) is a concept that is important to understand how robots function in industrial manufacturing. Stark and Damerau (2019) provide the following definition of a DT:

> *A digital twin is a digital representation of an active unique product (real device, object, machine, service, or intangible asset) or unique product-service system (a system consisting of a product and a related service) that comprises its selected characteristics, properties, conditions, and behaviours by means of models, information, and data within a single or even across multiple life cycle phases.*

One of the most common applications of the DT is for the design and configuration of an assembly line (Kousi et al., 2021). The DT enables the modelling of the different levels of a production system, namely the assembly process, production station, and line level, updating its parameters in real-time synthesisng data from multiple 2D–3D sensors about the actual production process. Then an AI-based logic is deriving alternative configurations of the

[13] Rauch, 2019. HillControl: Mehr Präzision beim Düngerstreuen in Hanglagen, auf Kuppen und in Senken!
[14] Raven, 2019. AutoBoom XRT - Raven Applied Technology. https://ravenprecision.com/products/boom-controls/autoboom.
[15] Müller Electronics, 2019. ISOBUS SPRAYER-Controller MAXI. https://www.mueller-elektronik.de/en/products/spray-controller-maxi-2/.

production system, suggesting the most efficient based on a set of predefined criteria.

Another common use of DTs is to control the resources of a production cell and adapt their behaviour based on the production requirements (Kousi et al., 2019). In other words, a model of the real system is used to re-plan production in response to unforeseen events, such as a temporary unavailability of parts due to supply line issues or a breakdown; or people stepping in front of a mobile robot and impeding its motion.

Predictive maintenance addresses production alterations that occur due to unexpected events such as breakdowns and unavailability of resources. Methodologies have been developed to calculate the Remaining Useful Life (RUL) of a machine by using physics-based simulation models and the DT concept (Aivaliotis et al., 2019). In the design phase, a model of a gearbox was created and parameterized using data from the real system. This model was able to simulate the friction inside the gearbox. During the execution phase, the model was fed with data from the machine's controllers as well as from external sensors that monitor their activity. As a result, the increase in gearbox friction could be simulated in real-time and the gearbox RUL could be calculated.

A DT may be used for process control. In additive manufacturing, the quality of processes as well as the robustness and the performance are posing several challenges due to a number of uncertainties, such as the material properties. A DT was designed and implemented for uncertainty management and robust process control (Stavropoulos et al., 2021). Using Linear Matrix Inequalities, within defined uncertainty limits, the process control was made and compared to a typical control approach that uses a fine-tuned conventional proportional–integral–derivative (PID). As a result, the robust control design achieved a 68% faster response in the settling time metric, while a well-calibrated PID only achieved 38% compared to the initial model.

A physics-based model is often complex and a significant amount of time may be needed to realize simulations with it, which prevents using the DT in real-time. Creating a meta-model by using the output of a large number of simulations with the DT ('synthetic data') may address this problem. Of especial interest is the use of AI-based meta-models (Alexopoulos et al., 2020).

A DT implementation has been proposed as part of a general cyber-physical system (CPS) to enable the optimization of the planning and commissioning of human-based production processes using simulation-based approaches (Nikolakis et al., 2019). The benefit of this development is that the above activities are optimized through (1) sensor data fusion and human activity recognition in the shopfloor and (2) knowledge management mechanism for capturing the implicit knowledge of the task execution.

Finally, DTs have been used in human–robot collaboration (HRC) applications as well, for teleoperation and coordination of multiple industrial robot operations, allowing operators from different manufacturing domains to control the robot. Li et al. (2022) describe a multi-robot teleoperation system in which the DT of the real robots is visualized in augmented reality (AR) glasses. Experimental studies with a single robot and with multiple robots working on collaborative tasks have shown that the proposed system can manage the multi-robot teleoperation tasks with high efficiency.

In summary, a DT always involves a model of a real 'thing'. The model can be static or dynamic, but the defining characteristic of a DT model is that the model is parameterized for a specific real-world system.

The DT can be used in two major ways. First, it can be used to configure and optimize the system, even before the real system is created. The DT can also be used to train personnel that will operate the real system. When a DT is used in this fashion, it is sometimes called a 'digital shadow'. Second, the DT can be used as a state estimator in real-time. In this case, sensors on the real system (joint positions, proximity, temperatures, pressures, cameras, etc.) are used to obtain information about the dynamic state of the system. This information may be fused with simulation outcomes. Also future states may be simulated and control actions may be taken in the present time to ensure that undesirable future states are avoided for the real system.

We hypothesize that agriculture has requirements similar to industrial manufacturing in terms of DTs. DT modelling of agricultural fields and agricultural processes could help to optimize agricultural activities. Using a digital version of the field, optimized vehicle routes can be generated both within and outside the field. Additionally, visualization tools for the above-mentioned information could enable the farmer to better understand and adopt the new technologies, especially when these are provided through web and/or mobile UIs. On top of that, in combination with simulation libraries, the farmers could be trained with simulated fields on how they could use the DTs and generate optimized operations, before using them in real conditions. DT can be combined with machine learning and AI-based algorithms to predict farming operations before robots enter the field, as well as optimize them during their execution in the field.

ROBS4CROPS is a 4-year EU-funded project (2021–24). The project aims to develop an agricultural robotic system that uses DT technology to react to unforeseen events autonomously. The objective of our work is to create an agricultural robot system that can operate with much less supervision than current practice. At a minimum, our system will detect problems and stop operating in a safe manner, but ideally it will be able to respond to changing conditions and unforeseen events.

2 Agricultural system under study

ROBS4CROPS focuses on the needs of four specific real-world farming scenarios (Table 1). For each scenario, the project has established a large-scale pilot which serves to elicit requirements, to develop the system and to demonstrate the developed robotic system in real operating conditions. Two pilots focus on mechanical weeding: in vineyards in France, and in arable crops in the Netherlands; and two pilots focus on spraying: in vineyards in Greece, and in apple orchards in Spain.

In the pilot in France, the aim is to reduce the environmental impact of wine growing by replacing chemical weed control in vineyards with mechanical weeding. Mechanical weeding with a tractor driver is economically not attractive (five to seven passes per year, total cost €800 per ha). Also, it is difficult to attract and retain experienced tractor drivers who can operate the several different weeding implements that are used. The goal is to replace the tractor and human driver with a robot. In this pilot, the Agreenculture robot, developed specifically for vineyards, is used (Fig. 1).

The pilot in Greece focuses on spraying table grape vineyards (Fig. 2). Supermarkets and consumers impose very high-quality standards and the fruit is highly sensitive to pests and diseases. Table grapes need up to 30 applications of biocides (preventive and curative) and foliar fertilizer each growing season. Labour costs are high and access to the workforce is difficult, especially given the limited time window available for each spraying application. In this pilot, the Agreenculture robot and a retrofitted tractor (autonomous) are used to carry a lift-mounted sprayer. The amount of chemicals used is reduced by adjusting the spray rate to the density of the canopy. An extra challenge in this pilot is the accessibility of the hilly terrain and the dynamic behaviour of the robots in this pilot.

The pilot in Spain focuses on apple orchards (Fig. 3). The cultivation of apples requires many fungicide applications. Farmers typically do not consider the growing stage of the trees and the disease severity when applying chemicals (uniform application). Using a robot addresses the labour challenge, while the amount of chemicals used is reduced by adjusting the spray rate to the density of the canopy.

Table 1 Overview of the large-scale pilots

Country	Crop	Action	Vehicle
France	Vineyard (wine grapes)	Mechanical weeding	Agreenculture robot
Greece	Vineyard (table grapes)	Spraying	Retrofitted tractor and Agreenculture robot
Spain	Apple orchard	Spraying	Retrofitted tractor
The Netherlands	Sugar beet and pumpkins	Mechanical weeding	Robotti

Figure 1 The Agreenculture robot with a weeding implement attached.

Figure 2 Conventional tractor and sprayer in a table grape vineyard in Greece.

In the pilot in the Netherlands (Fig. 4), the aim is to reduce the environmental impact of arable farming by replacing chemical weed control with mechanical weeding. Sugar beets and pumpkins are both crops that grow slowly at first and therefore need several weeding passes to control the weeds. The Robotti robot, which was developed for arable farming, is used in this pilot.

The four pilots have in common that they use an autonomous vehicle with an implement attached, that field operations consist of visiting each predefined lane once, and that travel to, from and between lanes takes place in predefined

Figure 3 The EOLO sprayer towed by a New Holland tractor in an apple orchard in Spain.

Figure 4 The weeding implement used in the pilot in the Netherlands, mounted on the Robotti robot.

areas. Spraying and weeding are operations that must be performed several times during the growing season on a given field. In addition, a typical farm has several fields that need to be sprayed and/or weeded and the robot must be moved from one field to the next.

3 Agricultural robotic system

The ROBS4CROPS agricultural robotic system consists of three major elements: autonomous vehicles, smart implements and DT-based software to visualize, optimize and control the robots, which we call the Farming Controller. In addition, there is a communication protocol through which information is exchanged between vehicles, implements and the Farming Controller.

At the time of writing, the project is half-way. Consequently, some parts of the project vision have been demonstrated as proof of concept only; they will be further developed and demonstrated in the field in the two remaining years of the project.

3.1 Vehicles

3.1.1 Agreenculture robot

The robot from Agreenculture is a tracked vehicle that specifically targets vineyards. It weighs 800 kg, is powered by a 12.7 kW diesel engine and can operate continuously for 10 h on its 30 L fuel tank. The robot is driven by 48 V DC electric motors. A single electronic unit, the AGCBox, contains an RTK-GNSS receiver as well as the guiding software. The robot is equipped with bumpers for obstacle detection. A three-point linkage[16] allows the use of standard agricultural implements.

A distinguishing element of the robot is its safe fencing, a certified safety feature which ensures that no part of the robot or its tool protrudes beyond a predefined safety contour. Certification is possible because the robot makes use of a dedicated RTK base station, GNSS signals are processed by proprietary software and boundary files are stored in an encrypted format.

When the safety boundary and all lanes to be travelled have been defined, the working path for the robot (the 'mission') can be generated. The mission is validated in a simulator before the robot is deployed in the field.

In order to accommodate the smart implements, ISOBUS and TIM were implemented. Messages for GNSS-based position and velocity were added on the CAN using the ISOBUS protocol allowing the different components (of several partners) to communicate using a common framework.

The robot software was enhanced to allow it to respond to messages requesting a stop, a change of speed and a change of hitch position. This allows

[16] https://www.iso.org/standard/41233.html.

Published by Burleigh Dodds Science Publishing Limited, 2024.

the implement to control some aspects of robot functioning. A stop request received from the weeding implement is routed to the high-level software of the robot. Responding to it does not compromise the safety system of the robot.

3.1.2 Robotti

The Robotti 150D is a field robot for power-intensive operations. It is AgroIntelli's second Robotti model and is being used in more than 15 countries. The robot weighs 3150 kg, measures 2.44 m (length) × 3.05–4.90 m (width) × 2.15 m (height). Robotti has two diesel engines with a total power of 104 kW. The wheels are driven via hydraulic power and the robot features four-wheel driving and two-wheel steering. The robot has a traditional power take-off (PTO), a three-point linkage with a lift capacity of 750 kg, and it provides one double-acting hydraulic outlet (maximum 40 L/min) with free return.

Field operations can be planned and monitored via the Robotti website. Users enter the implement's specifications, define the borders and headlands of the field and determine how the operation is to be carried out. The user can either let the system optimize the route or specify the route manually. In addition the user can configure the type of turns to be made (U-turn or zero-turn), corner paths, transit areas, and non-work areas. When the route is fully specified, it is downloaded to the robot. For safety reasons and to meet ISO standards, the robot can only be started when a person pushes a button on the computer screen on the physical robot. Once the robot has been started, its performance and the progress of the operation can be followed in real-time on the website. Footage from the front and rear surveillance cameras can also be monitored. A third camera, ImplementCam, can be put where the user wants to monitor within the 5 m cable length.

The route planning facility on the website is the only way that ROBOTTI can be programmed. This architecture is not designed to allow interaction with the Farming Controller. However, the route planning facility is accessible via an API. The desired functionality could be achieved when the Farming Controller accesses the routing API to create a new route, which is then uploaded to Robotti. This is not a step which we plan in the project.

Robotti has a CAN bus. For the project, the speed and position of the robot were made available on this bus. Additionally, the software driving the robot was modified such that it can respond to a TIM request_stop message.

3.1.3 Retrofitted tractor

One of the aims of the project is to develop a retrofit kit that can be used to convert a conventional farm tractor into an autonomous vehicle. Such an autonomous tractor can be used both in conventional farming and as part of a robotic system, thus allowing for a gradual transition to agricultural robotics. Two New Holland tractors were purchased: a T4.80F for the pilot in Greece and a T4.110F for the pilot in Spain.

In early 2022, the tractor in the pilot in Spain has been fitted with CANBUS functionality and it was used for operating the smart sprayer during the remainder of the year.

In winter 2022/23, both tractors were instrumented with actuators on the steering wheel, the accelerator and the brake pedal. An angle sensor mounted on the front axle measures the position of the front wheels. A revolution counter on the rear wheel axle measures wheel speed. Closed loop control of wheel speed and steering angle is realized through the same hardware, and similar software, that drives the Agreenculture robot. The autonomous tractor can realize a commanded linear and angular speed. It reports realized values on the CANBUS. The tractor is connected to ISOBUS-compliant smart implements through the IBBC connector.

3.1.4 Localization without relying on GNSS

In 2022, it was found that GNSS signals cannot be reliably received in two of the pilot sites. In Spain, the apple trees are sufficiently high that an RTK fix can often not be obtained. In Greece, a metal frame is used to support plastic covers to protect the grapes from rain; the GNSS signal is seriously affected by this construction. To address this issue, we have implemented a solution based on Simultaneous Localization and Mapping (SLAM).

We first build a map using a SLAM module (Grisetti et al., 2005) while teleoperating the robot, and subsequently the map is used for localization using a Monte Carlo approach (Hornung et al., 2014). This approach enables multi-waypoint navigation in a local frame (map), without relying on GNSS.

SLAM has been intensively addressed by the robotics community in the last two decades, mainly in indoor environments. In these cases, the environment is well-structured, predictable and limited in size. In contrast, in agricultural fields the environment is unstructured or semi-structured, with large extensions of land, the features evolve (e.g. the structure of the trees changes according to the seasons), and the terrain is irregular.

To address these issues, in ROBS4CROPS we have implemented an agricultural-persistent-features-based SLAM. In the case of the orchards, for example, as the structure of the trees evolve, we have focused on lower part of the trunks, which are present and visible across all seasons. To that end, we have implemented a trunk detector module, and with a graph-based approach, a persistent-features map is built, which is subsequently used for localization purposes (Fig. 5).

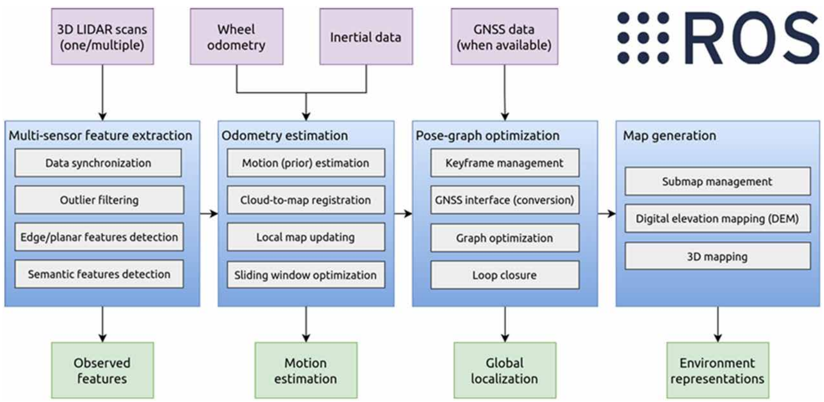

Figure 5 System diagram of the ROBS4CROPS SLAM system.

The ROBS4CROPS SLAM implementation does not rely on GNSS measurements, although these will be used if available. We use a bi-directional conversion between global coordinates (latitude, longitude and altitude) and the SLAM map frame. In this way any robot position (x, y and z) on the local map can be expressed in geodetic (latitude, longitude and altitude) coordinates and sent to the farming controller, even in the absence of GNSS data.

We rely on map-based navigation, in which the waypoints are referenced in the local frame of the map. As the localization precision of the SLAM algorithm depends on the features, we have implemented two navigation modules:

- Map-based navigation: which uses particle filtering over the map built with the SLAM module, suitable for navigation between rows or from arbitrary points of the map, where centimetric precision is not needed.
- Along row reactive navigation: as the clearance between the robot and the side trees of a row is about 0.2 m, we have developed a more robust reactive navigation module, which estimates the central path with centimetre-based precision. To this end, we extract the lines of the left and right sides given the aligned trunks using the Hough transform and derive the centre of the path. Given this information, precise linear and angular velocities corrections are executed to keep the robot centred.

3.2 Smart implements

3.2.1 Smart weeding implement for arable farming in the Netherlands

The weeding implement is a standard weeder with knives that are pulled through the soil a few centimetres below the soil surface (Fig. 4). This is a simple

Published by Burleigh Dodds Science Publishing Limited, 2024.

operation but there are several things that can reduce the efficacy of weeding. Here we focus on two common problems. First, the blades may become misaligned. This can happen when the hoe is thrown out of alignment when it hits a rock or when the fastening becomes loose. It may also be that the crop rows are not evenly spaced due to a problem during seeding. In either case, the result is that the crop plants are damaged. Second, it often happens that clods, weeds and/or crop residue collect on the vertical hoe supports and that crop plants are pulled out of the ground by this jam. A tractor driver who is weeding will frequently check the weeder to ensure that the above-mentioned problems are not happening.

To enable automatic evaluation of the performance of the weeding implement, two cameras were mounted on the Robotti, one at the front looking forward and one at the back looking rearward. The idea is to determine the quality of weeding by considering the differences between the images before and after the implement has passed. We used OAK-D Power-over-Ethernet (PoE) cameras which are designed for use in rugged environments and are rated IP67 (protect against weather and dust). Timing and speed are used to synchronize the images captured by the two cameras.

The initial algorithm to determine whether crop plants are being removed by the weeding implement relies on the assumption that only a small fraction of the ground is covered by weeds (normal situation). In this case, it is expected that the fraction of the soil surface covered by green material is substantially the same before and after weeding. Green soil cover is easily measured. First, each image is converted from RGB to HSV. Pixels with a hue between predetermined limits are considered to represent green material, resulting in a binary image in which green plants are indicated by white pixels. Only a small region of interest (ROI) close to the robot is being considered; the size of this ROI is chosen such that there is almost no overlap between the ROIs in successive images. If the fraction of green measured by the front camera is substantially larger than that measured by the rear camera, this is a strong indication that crop plants have been removed by the weeding operation. Output from the algorithm is illustrated in Fig. 6.

3.2.2 Smart weeding implement for vineyards in France

Four types of weeding implements are used in the large-scale pilot in France. All four concern weeding between the vines: serrated discs create a mound under the row (the 'ridging'), inter-vine knives cut the weeds' roots, Kress fingers scrape the soil surface and lump-breaker discs hoe the soil (Fig. 7).

The most common error encountered during the operation of these implements is that they are jammed with soil, weeds, crop residue and other debris encountered in the field. When this happens to the Kress fingers or

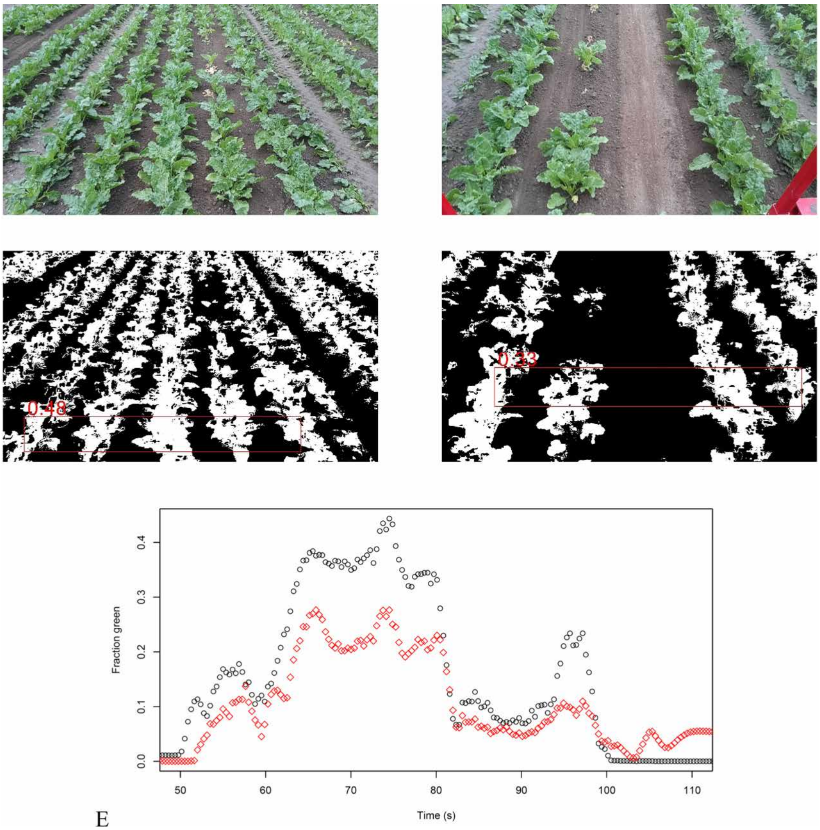

Figure 6 Smart weeder. Top left: forward-looking camera. Top right: rearward-looking camera. Middle row: binary images after transforming to HSV colour space and thresholding on hue; corresponding regions of interest in forward- and rearward-looking images are indicated, as well as the fraction of plant pixels in the ROI. Bottom: rolling median of the fraction of plant pixels in front (black circles) and rear (red diamonds) camera images. The fraction of plant pixels in the rear camera image is lower than in the front image, which is a strong indicator that crop plants have been removed.

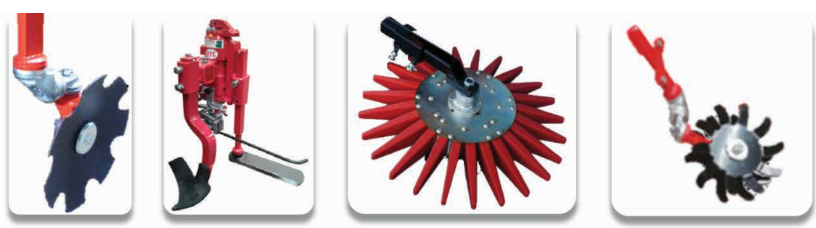

Serrated disc | Ploughshare and inter-vine knife | Kress-fingers | Lump-breaker disc

Figure 7 Weeding implements used in the pilot in France.

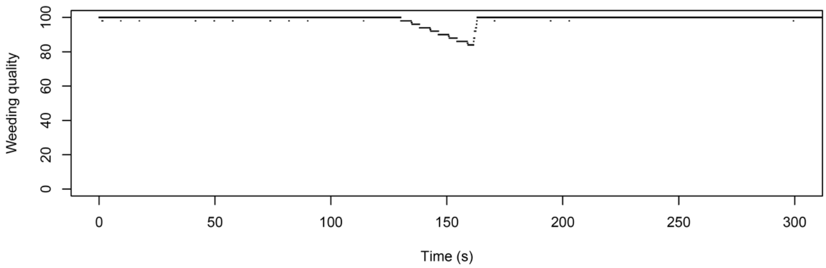

Figure 8 Weeding quality as measured by the weeding implement in the pilot in France. In this case, poor weeding quality was simulated by throwing a handful weeds into one of the Kress fingers at t = 130 s.

the discs their rotation is reduced or stopped altogether. We use revolution counters to measure the rotation of the Kress fingers and the discs. Pulses sent by the sensors are counted by a microcontroller, which sends the information to a PC mounted on the implement. Weeding quality is assumed to be 100% at the beginning. Whenever one of the Kress fingers does not send a pulse in a 1-s interval, weeding quality is decreased by 1. If weeding quality is less than 100 and all Kress fingers rotate properly, then weeding quality is increased. Example results are shown in Fig. 8.

3.2.3 Smart sprayers for vineyards in Greece and orchards in Spain

The smart sprayers consist of two elements: a variable-rate sprayer that can be controlled through ISOBUS and a perception unit (PU) that measures canopy density on-the-go and commands a spraying rate through ISOBUS.

3.2.3.1 Sprayer

Two different spraying units perform crop protection tasks in the spraying pilots. A 2000-L orchard sprayer, EOLO model, to be towed by a tractor, is used in the pilot in Spain and a 200 L lift-mounted vineyard sprayer, ASM model, is used in the pilot in Greece. Both spraying units share the same operating concepts with minor changes on sensors and actuators.

Smart sprayers gather intelligent and autonomous robotic capabilities to operate together with robotics vehicles with minimum human intervention required. Not only the spraying task but also derivative procedures are automated, as mixing phytosanitary when tank filling, tank cleaning, other maintenance and surveillance tasks as nozzle blockage control. An alarm, warnings and status dataset messages generated by the sprayer sensors are implemented on the communication bus of the FMIS-Vehicle-Sprayer architecture.

Variable rate technology (VRT) embedded in the sprayers allows an increase in efficiency and performance compared to traditional sprayers with mechanical-driven components.

The orchard (EOLO) sprayer has two horizontal sections, divided into three vertical layers (Fig. 3). The applied liquid flow rate is adjusted precisely through PWM nozzles independently for every one of the six sections according to commands received from the PU. Otherwise, if the sprayer follows commands from a prescription map, as-applied flow rate setpoint determined by the task controller for every one of the horizontal sections is applied homogeneously through the nozzles in the three vertical layers. The flow air rate assistance for canopy stirring is variable and independently controlled in the three vertical layers by hydraulic proportional valves which adjust hydraulic-driven fan speed. In the same way as liquid flow rate, air flow rate can be determined by the PU or configured through the human interface.

The vineyard (ASM) sprayer has two horizontal sections, divided in four vertical layers (Fig. 9). Liquid variable rate application operates in a similar manner as EOLO model. Flow air rate assistance for canopy stirring is adjusted according to PU commands depending on vineyard stage of vegetation, and commanded by the Sprayer ECU, which actuates the fan hydraulic proportional valve for rotating speed adjustment. Unlike EOLO, the ASM model pump is proportional hydraulic driven, providing more efficiency in adjusting rotation speed depending on liquid flow rate demand.

Figure 9 The ASM sprayer lifted by a New Holland tractor.

An ECU (electronic control unit) from EPEC (Seinäjoki, Finland) is used to control the fan and all the nozzles. ISOBUS functionality is implemented on the ECU with the Library V3 software stack (no peer-to-peer control).

3.2.3.2 Perception unit for canopy density

The amount of spraying liquid needed for effective disease control depends (among other factors) on the leaf area per unit ground surface (Duga et al., 2015). In practice, this means that early in the season, when grape vines and orchard trees carry few leaves, a low spray rate is sufficient to achieve full protection.

The PU developed in ROBS4CROPS measures the leaf area (Fig. 10). It adopts a modular camera design and mounting solution, in order to accommodate different canopy geometries and heights. This vision-based system uses depth and colour information to estimate the canopy density and generate real-time spraying recommendations. An HSV filter is used to keep only the green pixels in the depth image.

There are two main features considered when processing depth information, namely (1) the total pixel count on the depth image, which resembles to the canopy information that is present in the camera field of view (FOV) and (2i) the relative proximity index of neighbouring pixels, which corresponds to different leaf layers at different depths.

A machine learning model (multi-layer perceptron) has been trained to predict the density profile in the image, using three categorical classes based on the canopy presence: (1) no canopy, (2) sparse canopy and (3) full

Figure 10 The canopy density perception unit mounted on the hood of the tractor. The unit has two depth cameras (RealSense D435i, Intel, Santa Clara CA, USA) (one camera is facing the viewer). Also visible are two antennas for the wifi router that is part of the unit.

Figure 11 The effectiveness of the autonomous spraying system in table grapes was assessed in the field using strips of water-sensitive paper. See main text for explanation.

canopy. Annotation of depth data was performed by field experts categorizing each image in the corresponding class, using the Point Quadrant measuring technique and domain knowledge where necessary.

As a result, based on colour and depth information, the PU can modulate the default/maximum tree row spraying volume. If no canopy is detected, the spraying rate is set to 0. If a sparse canopy is detected, the spraying rate is set to 50% of the maximum rate. If a dense canopy is detected, the spraying rate is set to 100% of the maximum rate.

The PU was designed to send real-time spraying recommendations in the sprayer using the current velocity of the tractor/robot. All sending and receiving messages were based on ISOBUS and J1939 protocol communication. The software implementation uses python-can[17] for sending and receiving messages on a CAN bus line.

During the first season of deploying the autonomous spraying system in table grapes in Greece, an assessment protocol was established to measure the spraying effectiveness. Water-sensitive paper was placed on the leaves of three positions with different canopy densities. Canopy positions were selected in order to resemble grape canopies with (1) no canopy, (2) sparse canopy and (3) full canopy coverage. Each canopy was split into nine areas positioned in the upper (U) and lower (L) parts, front (F), middle (M) and back (B) (Fig. 11).

Water-sensitive papers have a specially coated yellow side that changes to blue when exposed to moisture. So, the wet area gets blue and the dry area remains yellow. We collected sprayer performance data in three contrasting settings, namely (1) a conventional sprayer, as currently used in table grape farms, (2) the ASM sprayer without PU, and (3) the ASM sprayer with PU activated (Table 2).

After the field session, a photo of each sprayed water-sensitive paper was taken in order to calculate the percentage of area covered by spray liquid on each paper, using image analysis software.

17 https://python-can.readthedocs.io/en/stable/.

Table 2 Experimental set-up for testing the sprayer. The conventional sprayer used hollow cone nozzles (Magnojet MGA 03). The variable rate sprayer used flat fan nozzles (Arag ASJ SF11004)

Treatment	Sprayer	Nozzles	Pressure (bar)	Speed (km/h)	Repetitions	Variable rate	Spraying rate (L/ha)
1	Conventional	Hollow cone	10	4	2	No	800
2	ASM	Flat fan	4	4	2	No	800
3	ASM	Flat fan	4	4	2	Yes	0, 400, 800

The PU software accurately quantified variations in foliage density in real-time. The results showed that the sprayer achieved satisfactory coverage and droplet size while saving from 10% to 50% of spray liquid compared to the conventional sprayer. It is also noteworthy that the innovative spraying operated with lower pressure and, thus, drift was reduced.

3.3 Communication

The robotic system consists of three main classes of components: autonomous vehicles, smart implements and the Farming Controller. The robotic system requires that these components are in constant communication with each other. Here we describe how this communication is effectuated.

3.3.1 Communication between implement and vehicle via ISOBUS

Communication between the implement and the vehicle is implemented with ISOBUS (ISO 11783). ISOBUS is based on the SAE J1939 communication protocol, which in turn is based on the physical CAN bus (ISO 11898).

ISOBUS is a standard for communication between agricultural vehicles and implements. However, it is in most cases not used by agricultural robots and there are still many implements without ISOBUS. Therefore, we have adopted an approach that allows the gradual introduction of ISOBUS. In this approach, ISOBUS connectivity is emulated for those vehicles and implements that do not yet offer ISOBUS connectivity. Emulation is provided by CANoe 15 SP2 software (Vector Informatik GmbH, Germany) running on an industrial PC. See Fig. 12.

The smart sprayers (Section 3.2.3) are fully ISOBUS compliant. The sprayers are designed with an ECU (EPEC, Seinäjoki, Finland) which integrates the full ISO 11783 protocol. As an input, the ECU of the sprayer receives GNSS information to geolocate the applied points and the ground-based speed to calculate the point-to-point distance from the Agreenculture robot and the autonomous tractor (AGCBox) through IBBC connector. For this, the TECU functionality of the ISOBUS is integrated into the CANoe software to enable address-claiming with the ECU of the sprayer. Another input to the ECU is the prescribed/recommended rate from the PU that is packed as a 'Process Data' (PD) message (DDI36 – Ap*pRateVolumePerTime_Set*). The message PD is defined with the ID of *0xCCBFEFE*. The ID is usually dynamic due to the dependency on the source address (SA) and destination address (DA). The PU (Section 3.4) is partially ISOBUS compliant because of the implemented Task Controller Library Version 3.0 of the ECU Sprayer. In this library, there is no possibility to enable the 'peer-to-peer' control function of the ISOBUS. Therefore, the PU cannot directly claim an address with the Sprayer ECU as

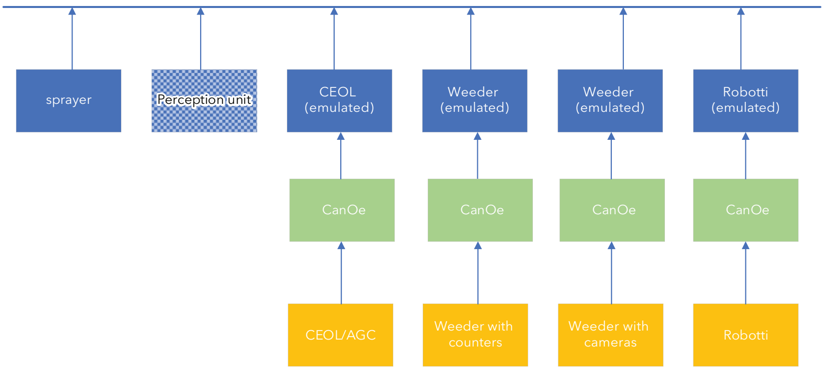

Figure 12 Schematic representation of connectivity between vehicles and implements via ISOBUS. See main text for explanation.

a fully ISOBUS compliant component. The Sprayer ECU sends the as-applied information (PD message, DDI2 – *AppRateVolumePerArea_Act*) as an output.

ISOBUS functionality for the weeding implements and the autonomous vehicles is achieved through emulation. The entire ECU of the weeding implements and ISOBUS functionality are emulated in the CANoe software. For the weeding implement in combination with the Agreenculture robot, the AGCBox emulates the GNSS information to geolocate the weeding points and the ground-based speed. The weeder in France calculates weeding quality by considering the number of revolutions of the discs or Kress fingers per metre travelled by the robot. The weeder in the Netherlands needs geolocation to correlate images taken by its forward- and rearward-looking cameras.

Tractor Implement Management (TIM) is an ISOBUS-based solution for a system where the implement can control certain tractor functions. While TIM is developed for tractor-based systems, it can be expected to be very useful for any situation in which smart implements gather information about their operation and information about the environment. For example, when a smart weeder detects that weed quality is very poor, it would be reasonable to stop the operation.

The vehicles used in our project (and indeed most tractors sold today) do not yet implement TIM. As a first step, we implemented on the vehicles used in our project the ability to respond to a single TIM command, namely the emergency stop command. The vehicles use ISOBUS message (TIM21) with an ID of *0x1C24FF21* to recognize the stop command incoming from the TIM Client node of the middleware. The ISOBUS address claiming between the TIM Client node and the implement ECU to send and receive signals is left to CANoe and defined under the 'acknowledgement' message with and ID of *0x18EEFF21*.

The weeding implement used in the pilot in France counts the number of revolutions for each of the Kress fingers or discs. A microcontroller sends this information via CAN to CANoe which claims an ISOBUS address and puts a corresponding message on the bus. The weeding implement used in the pilot in the Netherlands uses image processing software to determine the weeding quality. Here a Python script sends this information via CAN to CANoe which claims an ISOBUS address and puts a corresponding message on the bus.

Both the Robotti and the Agreenculture robot are connected via CAN to CANoe. Messages sent include ground-based speed and GNSS position data. Full details about these robots can be found in Paraforos et al. (2022) and Sharipov et al. (2023).

3.3.2 Communication between Farming Controller and implement/vehicle via Robot Operating System 2

The Robot Operating System (ROS) is a set of open-source software libraries and tools for robotics applications (Macenski et al., 2022). ROS comprises a message-based communication framework for passing information between the components of a robotic system. ROS has always enabled distributed systems but the latest version (ROS 2) greatly simplifies setting up communication between components running on different computers. All communication between the Farming Controller on the one hand and vehicles and implements on the other hand is implemented through ROS 2.

We used mobile internet and VPN to set up a Software-defined Wide Area Network (SD-WAN). This WAN is used to connect all vehicles and implements.

The Farming Controller and vehicle are on the same ROS 2 domain and therefore discover each other automatically. All topics published by all robots are always visible. We only want to log data when the robot is working, not when it is on a road or when it is in the barn. This is achieved by the Farming Controller only logging when a flag has been raised. Specifically, the robot in the pilot in Spain initiates logging by sending a message with content 'true' on the ROS 2 topic named 'lsps/trigger_flag'. When the flag is up, the Farming Controller scans the list of topics and subscribes to all topics with names starting with '/lsps/'.

3.4 Farming Controller

The Farming Controller is a DT-based software that allows the orchestration of the resources described in the previous sections (vehicles, implements and communication facilities). The Farming Controller must allow to select combinations of vehicles and implements in order to carry out different tasks. It must also allow the planning and scheduling of field operations (spraying and weeding) in a series of fields during an entire growing season; in doing so it must take into account transporting the robot from one field to the next. Once

Published by Burleigh Dodds Science Publishing Limited, 2024.

the robot is in a field, there are decisions to be made in terms of the order in which working lanes are visited (vehicle routing problem (VRP)). Once the robot is working in a field, the system must be able to detect problems with weeding and spraying, with obstacles, with fuel running out or batteries running empty and so on and replan the work if necessary.

The requirements for the Farming Controller call for the implementation of several functions of DTs as used in industrial manufacturing (Table 3). The first function, 'Design and configure', will allow farmers to generate and evaluate work schedules before any work is done in the real world. Schedules may be evaluated for the available set of fields, vehicles and implements, as well as for fields, vehicles and implements under consideration for purchasing. The second function, 'respond in real-time to changing conditions/unforeseen events', will, for example, allow the robotic system to use an alternative route if an obstacle is detected. Also, the system may respond to changes in weather: for example, if the wind increases while a spraying operation is in progress, it may be prudent to switch operations to a field where spray drift is less likely to damage an adjacent crop. The third function, 'predictive maintenance', is technically close to the second in the sense that data collected in real-time about the system is used to inform decision-making. In

Table 3 Digital twin-supported functionalities from industrial manufacturing and their application in agricultural robotics

Function	Industrial manufacturing example	ROBS4CROPS example
Design and configure (before deployment)	Optimize location and number of mobile robots and human workers in vehicle front axle assembly line (Kousi et al., 2021)	Select vehicles and implements; assign them to fields; create a work schedule to perform all necessary operations on each field within the set time frame; identify the work schedule that is optimal in some sense
Respond in real-time to changing conditions/ unforeseen events	Reconfigure assembly line when parts supply is interrupted. Re-plan robot motion path when human worker steps in front of a mobile robot (Kousi et al., 2019)	Reroute vehicle to avoid an obstacle or to accommodate unexpected field conditions
Predictive maintenance	Monitor gearbox wear and signal the need for repairs when they are needed (but not before then) (Aivaliotis et al., 2019)	Monitor battery charge (electric vehicle) or amount of fuel remaining (vehicle with combustion engine), as well as amount of biocide remaining. Reroute vehicle to filling/ charging station when needed

predictive maintenance, for example, the robotic system may monitor over time how well the battery of an electric vehicle is holding up, or how much biocide is used on average in a variable-application-rate modus, and then take appropriate action.

At a high level, an example of a working schedule could be the following. The farmer would like to find the optimal path to perform the spraying and weeding activities in different fields with the following structure:

- Field 1 for spraying;
- Field 2 for spraying; and
- Field 3 for weeding.

The farmer possesses the following resources which can be used by the decision-making algorithm to generate alternative path and operations/field assignments to each resource:

- Robot 1;
- Robot 2;
- Implement 1 for weeding; and
- Implement 2 for spraying.

Firstly, a combination of the possible vehicles is generated (robot with implement), based on the compatible combinations between each robot with each implement that can be made. The suitable combinations that have been generated by the decision-making algorithm are as follows:

- Vehicle 1 (Robot 1 + Implement 1);
- Vehicle 2 (Robot 2 + Implement 1); and
- Vehicle 3 (Robot 2 + Implement 2).

The decision-making algorithm generates many alternative working schedules by allocating the above vehicles to the different fields, taking into account that some fields need spraying and others need weeding (Fig. 13). For each allocation of a vehicle to a field, the time needed for the operation is simulated (see later). The amount of time needed to complete a working schedule depends on the number of vehicles used and on how the vehicles are allocated to fields. The most appropriate alternative is selected. Since this algorithm is based in heuristics and doesn't investigate all the alternatives exhaustively, the solution is considered to be acceptable but it is not necessarily the best.

An algorithm based on the Traveling Salesman Problem (TSP) is used to optimize the manoeuvre time of a vehicle in the field, taking into consideration the turning constraints. The algorithm results in a path that visits all working

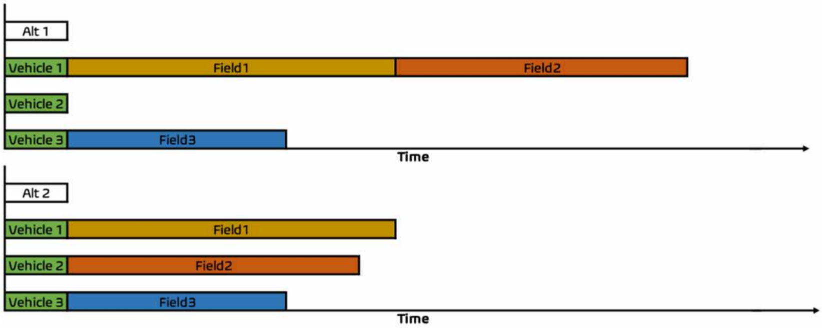

Figure 13 Gantt Chart that illustrates an example of two alternatives that are generated and evaluated by the decision-making algorithm.

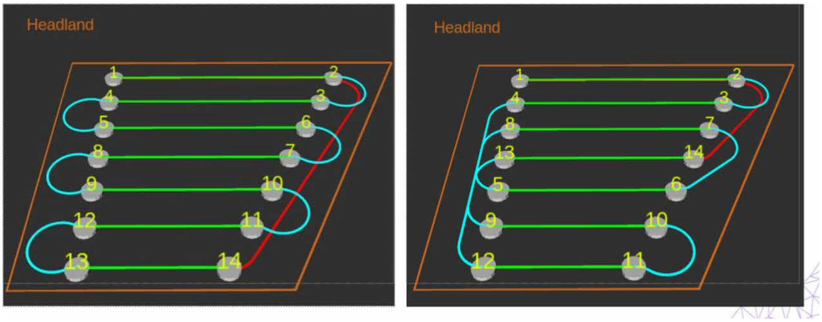

Figure 14 Two possible routes through a given field. Left: an initial route in which all working lanes are visited in sequence. Right: a route proposed by the TSP algorithm that takes into account the time needed to turn from one lane into another and minimizes the time to complete the entire route.

lanes at least once. In this algorithm, the parameters can be adapted to choose a path that is either shortest in length or shortest in travel time (Fig. 14).

The TSP algorithm gives a path that is optimal in terms of distance or travel time. Travel time is calculated by multiplying the length of each path segment by the speed with which that segment is traversed; working lanes will have a different speed than transport lanes and headland turns.

TSP cannot easily be extended to take into account spatially variable conditions (of the field), time-varying conditions (state of the vehicle and/or the field) and interactions between them. An example of a spatially variable condition of the field is that there may be zones where something is limited, such as vehicle speed, weight, width and/or height. An example of a time-varying condition is the weight of the vehicle, which may be high when a spraying operation starts (full tank) and low once the tank has been partially emptied. An example of an interaction is that a vehicle with a full tank may not be able to pass over a bridge with a low load carrying capacity, whereas the same bridge may be passable with a tank that is almost empty. Another example is that travelling uphill with a full tank requires more power and leads to more soil damage than either travelling downhill with a full tank or travelling uphill with a partially empty tank.

We use a dynamic (discrete event) simulation model to evaluate performance indicators associated with a given route. State variables of the simulation model include (for the vehicle) pose (location and orientation), battery status or fuel remaining and biocide level. The environment is monitored through indicators such as fuel consumption rate or engine power uptake (proxies for state of the soil), crop canopy size and crop disease level.

The simulated vehicle moves along the given route and activates the implement in working lanes. The speed is determined as the minimum of target speed for the operation and the maximum speed permitted by the terrain (input to the model). The model can easily be extended to include, for example, soil water content and trafficability and energy consumption as a function of vehicle weight, slope and soil condition.

An example of the model's output is given in Fig. 15 for a field which has four working lanes of 100 m length. Time and distance are simulated for two routes that start at the bottom left corner (East = North = 0). The first route starts with by going east and enters the first working lane. The second route starts by going north and then enters the second working lane. Both routes have the same length and both routes traverse twice a part of the headland. The difference is that the first route doubles up on the western headland and the second route does so on the eastern headland. In this example, speed on the western headland is restricted and therefore the first route takes longer to complete.

The full process of optimizing optimal use of resources is shown schematically in Fig. 16. Modelling takes place at three levels: Gantt chart to allocate vehicles and implements to fields during a growing season, TSP to select the sequence and direction in which working lanes are visited and physics-based simulation model to track changes in state variables as the vehicle moves along the field. Re-planning and -scheduling in response to changing conditions/unforeseen events could take place at any of the three levels, here we focus on the last two levels, i.e. responding while the vehicle is moving inside one field.

When the robot is moving in a field, a change in conditions or an unforeseen event may make it expedient to calculate a new vehicle route. A change in condition may be that the wind is picking up and that it is prudent to prioritize spraying in working lanes where there is no adjacent crop that could be damaged by spray drift. Another change in condition may happen during variable rate spraying where the application rate depends on some spatially varying characteristic of the crop (e.g. canopy density or disease infestation

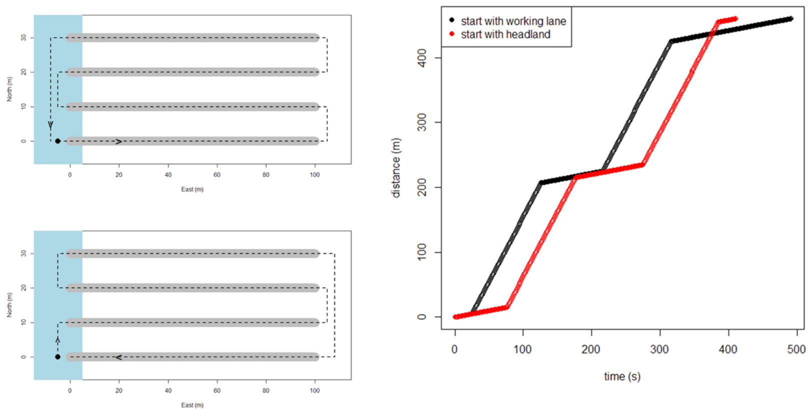

Figure 15 Example output of the agricultural simulation model, where based on working lane conditions the distance travelled and time spent can be determined during operation of the route in the field and during simulation. In this example, the terrain imposes a speed limit on the part of the field where East ≤ 5 m (indicated by the blue background). The 'start with working lane' route has a greater length in this area than the 'start with headland' route and therefore takes longer to finish even though the total length of both routes is the same.

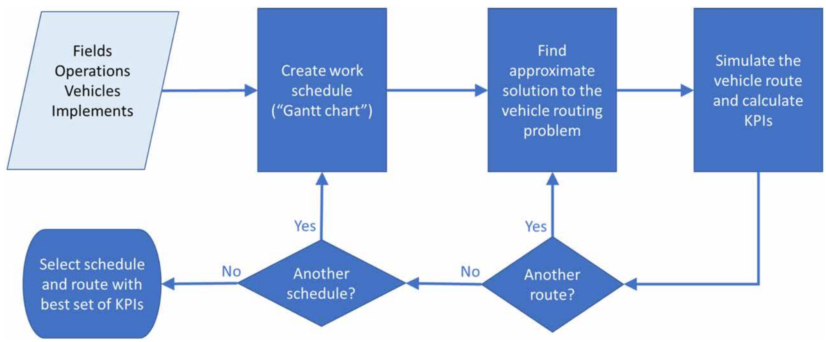

Figure 16 Flow chart of the three-level modelling process inside the Farming Controller. Given the set of fields, operations, vehicles and implements, a feasible working schedule is created. For a given working schedule, an approximate solution for the robot to move through field is created. Finally, for a given vehicle route, a dynamic simulation is performed to evaluate several key performance indicators (KPIs). Each process is repeated until a large set of alternatives is available, from which the most suitable one is selected in the last step of the flow chart.

level). In this situation, it may be that the biocide is used at a higher rate than foreseen. When the DT model is re-run with the new (observed) average application rate, a prediction can be made as to when the robot will run out of biocide and thus, at what point it must return to base for a refill.

The action communication type of ROS 2 is used to implement the robot route control mechanism (Fig. 17). The logic of this control mechanism is to

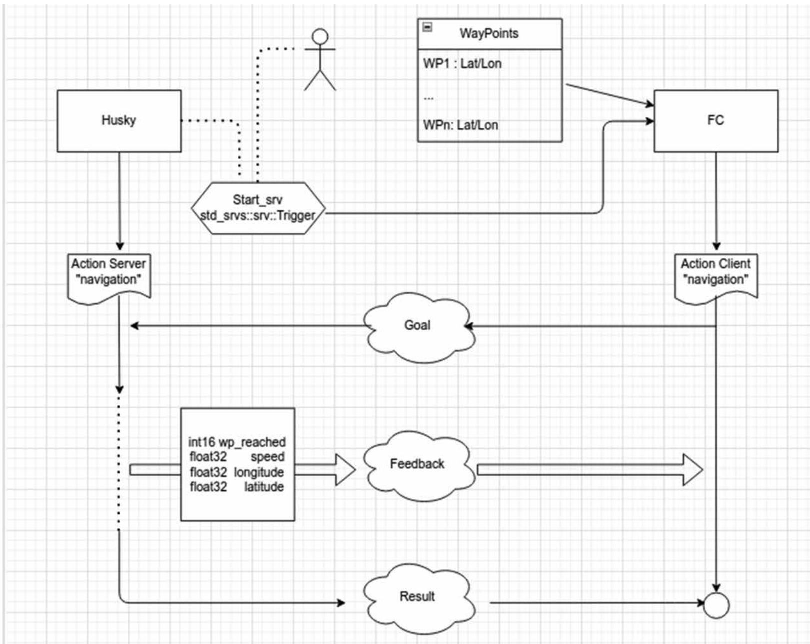

Figure 17 Schema of the interaction between a robot and the Farming Controller. The robot (here denoted as 'Husky') implements the ROS 2 Action Server interface and the Farming Controller implements the Action Client interface. In a loop, the Farming Controller sends an ordered set of waypoints to be visited (the goal) to the robot, and the robot sends feedback to the Farming Controller whenever a waypoint is reached.

```
2
    # Goal
    nav_msgs/Path          path
    ---
    # Result
    std_msgs/Empty         plan_result
    ---
    #Feedback
    int16                  waypoint_reached
    float32                speed
    float32                longitude
    float32                latitude
```

Figure 18 The ROS 2 Action message that is used to send a route to the robot and receive feedback. Here 'path' is the ordered series of waypoints that define the route to be followed; 'waypoint_reached' is the sequence number of the last waypoint reached; 'speed' is the current speed of the robot (m/s); longitude and latitude indicate the current position of the robot in the WGS84 coordinate system.

get an action signal from the robot to the Farming Controller that it is ready to get navigation points and the Farming Controller is responding with the set of waypoints that the robot should visit. Once the robot confirms reaching the first point, the Farming Controller sends the next set of waypoints and this loop is iterated until the whole path is followed. In ROS 2 terms, the Farming Controller is the Action Client and the robot is an Action Server; communication is effectuated through an Action message (Fig. 18). The action message contains the list of waypoints expressed in the WGS84 coordinate system. While the robot is navigating, it sends a continuous stream of position messages (with current speed, latitude, longitude, etc.) as well as the number of each waypoint that it reaches.

When it is necessary to send the robot on a different route, a new goal (Action message) is sent to the robot. The new goal pre-empts the earlier goal and the robot starts following the new route.

The first hardware implementation of the route control mechanism was realized using a Husky (Clearpath Robotics Inc., Kitchener ON, Canada). The

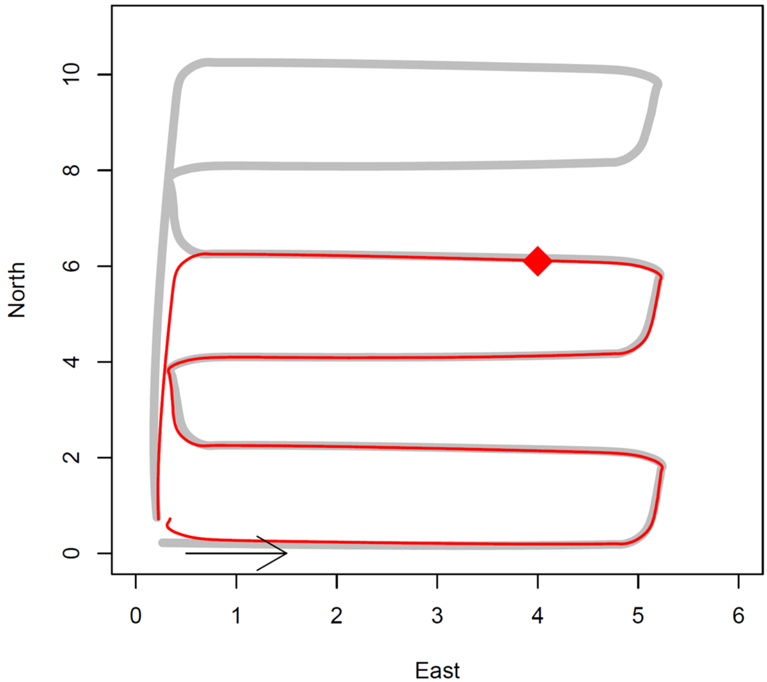

Figure 19 Results of a route control experiment. Starting near (0,0), and initially travelling in the direction indicated by the arrow, the thick grey line shows the path followed by the robot during a normal field operation. The thin red line shows the path followed by the robot during a field operation when an unexpected event occurs (indicated by the diamond symbol). The kind of the event is not specified, it could be a change in the weather, the observation that the battery is almost empty or any other non-emergency event that results in a new route being sent to the robot. When the new route is sent, the robot continues to follow the lane until it reaches the headland; from there it can continue on the alternative route.

small size of this robot makes it easier for development work than a full-sized tractor. A standard set-up using the Navigation 2 package from ROS 2 was put in place. The Action Server was implemented in Python. The experiment was conducted with the Husky located in the Netherlands and the Farming Controller running on a server in Greece. A VPN was used to establish a network connection (see Section 3.3). Results of the experiment are shown in Fig. 19. The experiment itself is trivial but it was instrumental in developing the protocol.

4 Case study results

In the context of the large-scale pilot in France, the robot operated in multiple series of weeding sessions in the field, proving its robustness and ability to perform. However, the main challenge encountered has been the frequency of stops, which posed the need for the user to intervene and restart the robot or perform troubleshooting. Thus, the user was required to always be on-site in order to supervise and act accordingly. The communication of the robot with the Farming Controller has been excellent during field trials. In 2022, nine plots with a total area of 14 ha were weeded. Each plot was weeded 6 or 7 times, with each pass taking approximately 1 h. The most common error was that the robot's sensitive front bumper was activated by a vine; in all cases, the robot could continue working after a manual override.

In the pilot, in the Netherlands, the Robotti proved to be able to perform seeding and weeding throughout the season, even though some adjustments were needed to perfectly fit all implements on the robotic platform. A human supervisor was present at all times to intervene in cases of the robot stopping. The communication of Robotti with the Farming Controller has been excellent during field trials. In 2022, the robot worked in a field of 0.8 ha; seeding, two tine weeding passes and four hoe weeding passes were done with the robot.

Regarding spraying application trails in the large-scale pilot in Spain, spraying using a retrofitted tractor and the PU was satisfactory. During the trials, the tractor was operated by a human driver, while the goal for 2023 is to apply autonomous navigation. The ISOBUS worked efficiently and the communication with the Farming Controller proved to be performing well. Sprayings took place for many hours during the entire growing season from June to September 2022 in a field of 12 ha.

As for the large-scale pilot in Greece, trials during 2022 were performed using a temporary solution of a pre-existing tractor that was modified and retrofitted. As in Spain, the tractor was driven by a human and the communication with the Farming Controller was very stable. Experimental operations lasted for 2 days in a field of 0.2 ha.

The VPN connection with the Farming Controller was very stable in all pilots. Logging data worked without fail.

5 Discussion

We have presented a vision and a preliminary implementation of an agricultural robotics system centred around a DT. This robotic system consists of smart implements and autonomous vehicles which are orchestrated by a Farming Controller. The system is developed to serve the needs of four specific farming situations and is being tested in these farms.

The smart weeder for sugar beets consists of a conventional weeding implement which is made smart by a forward-looking and a rearward-looking camera. Comparing the images captured by these cameras allows to determine the quality of the weeding operation. The algorithm developed is designed to check one specific condition, namely where weeds and/or crop residue are jammed in the hoe, leading to 'collecting' the crop plants, a situation in which the crop plants are pulled out of the ground. The algorithm is based on basic image processing techniques and is sensitive to the threshold values that are used. The algorithm may fail, for example, when dark shadows are present in the image, when a lot of weeds are present and when it is deployed on different soils. However, this initial algorithm can easily be replaced with another algorithm, possibly one based on deep learning. It is also conceivable to have a stack of algorithms running, each checking for a different error condition.

Likewise, the smart weeder for vineyards also checks for a specific (and common) error condition, namely the Kress fingers or rotating discs being jammed by weeds or soil.

The sprayers (orchard in Spain, vineyard in Greece) are smart by relying on the PU which adjusts the spraying rate based on canopy density. Limitations of the current system include step-wise instead of continuous adjustment of the spray rate and not measuring infestation level which would allow to reduce spraying rate on healthy but dense canopies. A further limitation is that the sprayer does not measure flow rate and is thus not able to detect if one or more nozzles are blocked. These shortcomings can be addressed at a later time without invalidating the robotic system itself. In the meantime, the system is already an improvement over current practices.

The robots used in this project have proven to be robust and able to work for many hours. However, working in large, open field with sugar beets, in which no obstacles are present, is easier than working in a vineyard. In the latter case, we found that the robot frequently stopped because a vine touched its safety bumper. In all cases, the robot was commanded by the supervisor to continue driving. This points to a need to be able to differentiate between vine and a real obstacle touching the bumper.

In order to transform a conventional tractor into an autonomous tractor, it is necessary to get accurate and precise feedback about the wheel speed and the angle of the front wheels. These measurements are available on the tractor's internal messaging bus, but the protocol is not documented. We have overcome this problem by installing additional sensors on the tractors.

A major, unexpected obstacle encountered was that GNSS signals were not reliable in two of our pilot sites. In the orchards in Spain, trees are sufficiently high that an RTK fix can often not be obtained. In the vineyard in Greece, we were surprised to discover that a metal frame is used to support plastic covers intended to keep rain from the grapes; naturally, the GNSS signal is seriously affected by this Faraday cage. This problem was overcome by resorting to localization based on LIDAR and a map of the orchard or vineyard.

Implementing the ability to respond to a 'stop' command has turned out to be a complex issue for the robot manufacturers in our project. The robots were designed to operate as stand-alone units and the software for navigation and for low-level control is tightly integrated. This has allowed us to provide the desired level of safety: the robots are unable to start moving without a person being physically present.

The Farming Controller has only partially been developed. However, we have delivered a proof of concept of calculating a route off-site, sending it to a robot in the field and pre-empting that route in real-time by sending an alternative route to the robot.

The smart implements, the autonomous vehicles and the Farming Controller must be in constant communication with each other for the robotic system to work. We use ISOBUS for communication between vehicles and implements. The bandwidth offered by ISOBUS is insufficient to support the delivery of a stream of images; thus we used ethernet for the cameras. We use ROS 2 over a mobile ethernet connection for communication between the Farming Controller and the vehicles. We found this to work very well, but of course this is predicated on having high-quality mobile ethernet in the field.

The work presented here will be elaborated and expanded during the remaining 2 years of the project. For implements, it is expected that camera-based monitoring of weeding will benefit from leveraging developments in deep learning to become more robust and more versatile. The bottleneck here will be collecting sufficient training data.

For the Farming Controller, it is expected that task optimization will become more important. Optimizing a task will typically be a multi-objective goal, taking into account time, fuel use, timeliness of operations and so on. The information gathered by a variety of sensors and sent to the Farming Controller can be collected in a digital twin of the farm, which gives up-to-date, precise, accurate information about the health and size of the crops; the presence of weeds, pests and diseases and trafficability of the field (Paraforos et al., 2019). Using this information, future versions of the Farming Controller may be able to initiate tasks. For example, when measurements in the field or a disease prediction model (Been et al., 2023; Kessel et al., 2018) indicate that a potato crop is at risk of being infected with late blight, the Farming Controller could start a spraying task.

But widespread adoption of agricultural robotics will require more than just technology. In this chapter, we have not addressed that a support system of training, dealers, maintenance provides, banks and insurers is needed before large-scale adoption will be reached.

In summary, the preliminary results presented here support the hypothesis that the proposed robotic system will work. Widespread adoption of an agricultural robotic system may still be some distance away but an immediately relevant contribution of our work is the data sent by the implements and the vehicles to the Farming Controller in a DT framework. This opens the way to create a rich user interface with which an operator (farmer, farm worker or contractor) can follow in real-time the location of the agricultural robots in his or her fleet of robots. Such a user interface can also provide information about the status of each robot, the quality of the work being done by the robot and a live stream of the robot at work. The user interface can also present information about how long it will take to complete each task and it will offer alternative schedules when conditions change. We expect that the willingness of operators to 'trust' robots will greatly increase when they can monitor the performance of the machines on their smartphones, in real-time, wherever they are located. This will be an important step towards adopting a complete robotic system.

6 Conclusion

Our work is based on the hypothesis that digital twins will be just as useful to optimize activities in agriculture as they have proven to be helpful in industrial manufacturing. We have presented a vision of an agricultural robotics system centred around a digital twin. In addition, we have been able to provide an implementation of this system using off-the-shelf hard- and software in combination with some custom-made components. The system is being tested in four different farming systems throughout Europe. Based on initial test results, we conclude that a DT-based agricultural robot system requires less supervision than current systems and that it will be able to respond to changing conditions and unforeseen events.

7 Where to look for further information

- Field robots catalogue: https://www.futurefarming.com/dossier/field-robots/.
- Human supervision and national and EU rules and regulation is a challenge (Lowenberg-DeBoer et al., 2022; Maritan et al., 2023).
- Farmers adoption of autonomous systems are modest due to a range of factors (Gabriel and Gandorfer, 2023; Gil et al., 2023).
- Robots not only address labour shortages but can also lead to reduced CO_2 emissions (Gonzalez-de-Soto et al., 2015) and soil compaction (Duckett et al., 2018).
- Several studies indicate positive economic perspectives from using robots (Lowenberg-DeBoer et al., 2020).

7.1 Key journals/conferences

- Journal of Field Robotics. A Special Issue on agricultural robotics will appear in 2023. The 2020 Special Issue on agricultural robotics contains many relevant papers (but digital twins are not mentioned).
- FIRA (https://www.fira-agtech.com/) and SIMA (https://en.simaonline.com/) showcase many agricultural robots.
- IEEE International Conference on Robotics and Automation (ICRA) and International Conference on Intelligent Robots and Systems (IROS) are leading conferences.

8 Acknowledgements

The ROBS4CROPS project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101016807. Wageningen University & Research has received additional funding from the Netherlands Ministry of Agriculture (project BO-67-101-023).

9 References

Aivaliotis, P., Georgoulias, K. and Chryssolouris, G. 2019. The use of digital twin for predictive maintenance in manufacturing. *International Journal of Computer Integrated Manufacturing* 32(11): 1067–1080.

Alexopoulos, K., Nikolakis, N. and Chryssolouris, G. 2020. Digital twin-driven supervised machine learning for the development of artificial intelligence applications in manufacturing. *International Journal of Computer Integrated Manufacturing* 33(5): 429–439.

Anken, T. and Latsch, A. 2022. Detection rate and spraying accuracy of Ecorobotix ARA. Lecture notes in informatics (LNI). In *Lecture Notes in Informatics (LNI) - Proceedings. Series of the Gesellschaft fur Informatik (GI).* Gesellschaft für Informatik. P-317: 45–50. ISBN 978-3-88579-711-1. Gesellschaft für Informatik.

Anken, T., Latsch, A. and Hatt, M. 2022. Spot spraying broad-leaved dock plants in meadows Performance, accuracy and sprayer test of Ecorobotix ARA. *VDI-Berichte* 2022: 279–284. https://doi.org/10.51202/9783181023952-279.

ASABE Staff. 2022. See & SprayTM select by John Deere. *Resource Magazine* 29(3): 7–8.

Been, T. H., Kempenaar, C., van Evert, F. K., Hoving, I. E., Kessel, G. J. T., Dantuma, W., Booij, J. A., Molendijk, L. P. G., Sijbrandij, F. D. and van Boheemen, K. 2023. Akkerweb and farmmaps: development of open service platforms for precision agriculture. In Cammarano, D., van Evert, F. K. and Kempenaar, C. (Eds). *Precision Agriculture: Modelling*. Springer International Publishing, Cham, pp. 269–293.

Bochtis, D., Vougioukas, S., Ampatzidis, Y. and Tsatsarelis, C. 2007. On-line co-ordination of combines and transport carts during harvesting operations. *Precision Agriculture 2007 - Papers Presented at the 6th European Conference on Precision Agriculture*, ECPA, pp. 715–721.

Cloete, S. 2020. Weed management using sensor technology. *Farmer's Weekly* 2020: 60–61.

Duckett, T., Pearson, S., Blackmore, S., Grieve, B., Chen, W.-H., Cielniak, G., et al. 2018. Agricultural robotics: the future of robotic agriculture. arXiv preprint arXiv:1806.06762.

Duga, A. T., Ruysen, K., Dekeyser, D., Nuyttens, D., Bylemans, D., Nicolai, B. M. and Verboven, P. 2015. Spray deposition profiles in pome fruit trees: effects of sprayer design, training system and tree canopy characteristics. *Crop Protection* 67: 200–213. https://doi.org/10.1016/j.cropro.2014.10.016.

Gabriel, A. and Gandorfer, M. 2023. Adoption of digital technologies in agriculture—an inventory in a European small-scale farming region. *Precision Agriculture* 24: 68–91. https://doi.org/10.1007/s11119-022-09931-1.

Gil, G., Casagrande, D. E., Cortés, L. P. and Verschae, R. 2023. Why the low adoption of robotics in the farms? Challenges for the establishment of commercial agricultural robots. *Smart Agricultural Technology* 3: 100069.

Gonzalez-de-Soto, M., Emmi, L., Garcia, I. and Gonzalez-de-Santos, P. 2015. Reducing fuel consumption in weed and pest control using robotic tractors. *Computers and Electronics in Agriculture* 114: 96–113. https://doi.org/10.1016/j.compag.2015.04.003.

Grisetti, G., Stachniss, C. and Burgard, W. 2005. Improving grid-based slam with rao-blackwellized particle filters by adaptive proposals and selective resampling. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE Publications, USA, pp. 2432–2437.

Hemming, J., de Jong, H., Struik, L., van Meer, J. and Van Henten, E. 2018. Field performance evaluation method for a machine vision based intra row hoeing machine. In *Proceedings of the European Conference on Agricultural Engineering AgEng2018*. Wageningen University & Research, Wageningen, The Netherlands, pp. 188–196.

Hizatate, T. and Noguchi, N. 2023. Work schedule optimization for electric agricultural robots in orchards. *Computers and Electronics in Agriculture* 210: 107889. https://doi.org/10.1016/j.compag.2023.107889.

Hornung, A., Oßwald, S., Maier, D. and Bennewitz, M. 2014. Monte Carlo localization for humanoid robot navigation in complex indoor environments. *International Journal of Humanoid Robotics* 11(2): 1441002.

Kessel, G. J. T., Mullins, E., Evenhuis, A., Stellingwerf, J., Cortes, V. O., Phelan, S., van den Bosch, T., Förch, M. G., Goedhart, P., van der Voet, H. and Lotz, L. A. P. 2018. Development and validation of IPM strategies for the cultivation of cisgenically modified late blight resistant potato. *European Journal of Agronomy* 96: 146–155. https://doi.org/10.1016/j.eja.2018.01.012.

Kousi, N., Gkournelos, C., Aivaliotis, S., Giannoulis, C., Michalos, G. and Makris, S. 2019. Digital twin for adaptation of robots' behavior in flexible robotic assembly lines. *Procedia Manufacturing* 28: 121–126.

Kousi, N., Gkournelos, C., Aivaliotis, S., Lotsaris, K., Bavelos, A. C., Baris, P., Michalos, G. and Makris, S. 2021. Digital twin for designing and reconfiguring human–robot collaborative assembly lines. *Applied Sciences* 11(10): 4620.

Li, C., Zheng, P., Li, S., Pang, Y. and Lee, C. K. M. 2022. AR-assisted digital twin-enabled robot collaborative manufacturing system with human-in-the-loop. *Robotics and Computer-Integrated Manufacturing* 76: 102321. https://doi.org/10.1016/j.rcim.2022.102321.

Lowenberg-DeBoer, J., Huang, I. Y., Grigoriadis, V. and Blackmore, S. 2020. Economics of robots and automation in field crop production. *Precision Agriculture* 21(2): 278–299.

Lowenberg-DeBoer, J., Behrendt, K., Ehlers, M. H., Dillon, C., Gabriel, A., Huang, I. Y., Kumwenda, I., Mark, T., Meyer-Aurich, A., Milics, G., Olagunju, K. O., Pedersen, S. M., Shockley, J. and Rose, D. 2022. Lessons to be learned in adoption of autonomous equipment for field crops. *Applied Economic Perspectives and Policy* 44(2): 848–864.

Macenski, S., Foote, T., Gerkey, B., Lalancette, C. and Woodall, W. 2022. Robot Operating System 2: design, architecture, and uses in the wild. *Science Robotics* 7(66): eabm6074.

Maritan, E., Lowenberg-DeBoer, J., Behrendt, K. and Franklin, K. 2023. Economically optimal farmer supervision of crop robots. *Smart Agricultural Technology* 3: 100110.

Mier, G., Valente, J. and de Bruin, S. 2022. Fields2Cover: an open-source coverage path planning library for unmanned agricultural vehicles. arXiv preprint arXiv:2210.07838.

Nikolakis, N., Alexopoulos, K., Xanthakis, E. and Chryssolouris, G. 2019. The digital twin implementation for linking the virtual representation of human-based production tasks to their physical counterpart in the factory-floor. *International Journal of Computer Integrated Manufacturing* 32(1): 1–12. https://doi.org/10.1080/0951192X.2018.1529430.

Oksanen, T. and Visala, A. 2009. Coverage path planning algorithms for agricultural field machines. *Journal of Field Robotics* 26(8): 651–668.

Paraforos, D. S., Aubé, C., Athanasakos, L., Avgoustakis, I., Baron, S., Bresilla, T., Fountas, S., Hemming, J., Karagiannis, P., Mylonas, N., Nieuwenhuizen, A. T., Garcia, F. R., Pavlenko, T., Scovill, A., Sharipov, G. M., Vidal, J. and van Evert, F. K. 2022. Connecting agricultural robots and smart implements by using ISO 11783 communication. *IFAC-PapersOnLine* 55(32): 200–205. https://doi.org/10.1016/j.ifacol.2022.11.139.

Paraforos, D. S., Sharipov, G. M. and Griepentrog, H. W. 2019. ISO 11783-compatible industrial sensor and control systems and related research: a review. *Computers and Electronics in Agriculture* 163: 104863.

Sharipov, G., Bresilla, T., Nieuwenhuizen, A., Hemming, J., Van Evert, F. K., Aubé, C., et al. 2023. *Communication between Agricultural Robot and Mechanical Weeding Machine Based on ISO 11783 Network*. IFAC, Japan.

Stark, R. and Damerau, T. 2019. Digital twin. In Chatti, S. and Tolio, T. (Eds). *CIRP Encyclopedia of Production Engineering*. Springer, Berlin Heidelberg, Berlin, Heidelberg, pp. 1–8.

Stavropoulos, P., Papacharalampopoulos, A., Michail, C. K. and Chryssolouris, G. 2021. Robust additive manufacturing performance through a control oriented digital twin. *Metals* 11(5): 708.

Tillett, N. D. and Hague, T. 2006. Increasing work rate in vision guided precision banded operations. *Biosystems Engineering* 94(4): 487–494.