Monitoring and Mapping with Robot Swarms for Agricultural Applications

Dario Albani
ISTC-CNR
Rome, Italy.
dario.albani@istc.cnr.it

Joris IJsselmuiden
Wageningen University
Wageningen, Netherlands
joris.ijsselmuiden@wur.nl

Ramon Haken
Avular B.V.
Eindhoven, Netherlands.
r.haken@avular.com

Vito Trianni
ISTC-CNR
Rome, Italy.
vito.trianni@istc.cnr.it

Abstract

Robotics is expected to play a major role in the agricultural domain, and often multi-robot systems and collaborative approaches are mentioned as potential solutions to improve efficiency and system robustness. Among the multi-robot approaches, swarm robotics stresses aspects like flexibility, scalability and robustness in solving complex tasks, and is considered very relevant for precision farming and large-scale agricultural applications. However, swarm robotics research is still confined into the lab, and no application in the field is currently available. In this paper, we describe a roadmap to bring swarm robotics to the field within the domain of weed control problems. This roadmap is implemented within the experiment SAGA, founded within the context of the ECORD++ EU Project. Together with the experiment concept, we introduce baseline results for the target scenario of monitoring and mapping weed in a field by means of a swarm of UAVs.

1. Introduction

Despite being studied for about 20 years, swarm robotics is still confined into laboratory settings [6, 7] and no commercial application can be acknowledged to date. This has several reasons, including the need for cost-effective hardware, the lack of established user-swarm interaction methodologies, and the need of convincing use-cases and business models. Nevertheless, research in swarm robotics has produced a large knowledge base that can be exploited to deliver concrete applications [4, 27]. Additionally, several proposals have been advanced to provide engineering methods for swarm robotics, independently of the specific domain [9, 20]. Hence, time is mature for the last missing step: approaching a concrete real-world problem with a genuine swarm robotics approach. Precision agriculture is one such domain where day-to-day activities, such as monitoring and mapping of large fields, could be tackled through intelligent robotic technologies. Generally speaking, agriculture represents a very challenging and increasingly important domain, and constitutes the second highest-impact market (after defence) for mobile service robot applications, according to the World Robotic Report 2016 of the International Federation of Robotics[1]. Beyond mere automation, robots offer additional means to truly implement a precision agriculture approach [1, 8]. Agricultural problems are characterised by unstructured environments, large spatial distributions and heterogeneities that naturally call for flexible and robust multi-robot approaches. The application of swarm intelligence to agricultural robotics can lead to disruptive innovation, thanks to the miniaturisation of hardware and the cooperation within a highly redundant system. Redundancy and cooperation within a distributed robotic system can provide resilience and robustness to faults, and can result in super-linear performance, so as to maximise the effectiveness of the group as a whole beyond the sensing and information-processing abilities of individual units (e.g., exploiting biological models of information retrieval and integration [2, 20]). Automated weed control is certainly a priority to reduce labor and operation costs, while maximising yield and minimising/avoiding the usage of chemicals. Autonomous weed control systems require (i) efficient navigation within the field, (ii) automatic detection and identification of weeds, (iii) mechanisms for individual weed removal and control, and (iv) field mapping to support decision-making at a global scale [25]. Out of the above issues, weed recognition and field mapping still represent important challenges for an automatic weed con-

[1] *www.ifr.org*

978-1-5386-2939-0/17/$31.00 ©2017 IEEE

trol system [24, 25]. In this respect, various robotic solutions are being developed, with unmanned aerial vehicles (UAVs) recently having a large share, thanks to (i) the reduced costs and the increased reactivity and resolution with respect to satellite or other aerial photogrammetry technologies, and (ii) the ability to move over and rapidly map the field at a higher speed with respect to ground vehicles [3, 29]. The use of UAVs is acknowledged in several environmental monitoring scenarios, such as waste detection or fire prevention, and agriculture represents one of the most important targets for commercial drones deployment. However, the perceptual abilities of UAVs may not be sufficient when inspection must be performed from close by. Also, efficiency issues come into play for coverage of large areas, where the battery lifetime represents a tough constraint. A possible solution is provided by robust and scalable multi-robot approaches studied within swarm robotics [4]. By exploiting the power of collective intelligence, it is possible to overcome the individual perceptual limitations, parallelise activities and deal with uncertain environmental conditions leading toward various applications related to precision agriculture and beyond. Indeed, solutions tailored for the agricultural domain can be easily exported to other environmental monitoring problems without much effort, owing to the flexibility, scalability and robustness provided by a genuine swarm robotics approach. Starting from these insights, we have developed a roadmap for the application of swarm robotics to the weed monitoring/mapping problem, which is discussed in Section 2. This roadmap is implemented within the Swarm Robotics for Agricultural Applications - SAGA project[2], founded by the ECHORD++ (GA: 601116) EU project. SAGA aims to demonstrate the usage of a group of small UAVs to collectively monitor a field and distributedly map the presence of weeds.

The paper is organised as follows. In Sect. 2, we discuss the concept and the starting point of the roadmap in relation to the state of the art, and we indicate the activities planned to tackle the weed monitoring and mapping problem described above. In Sect. 3, we introduce an abstract model for multi-robot field monitoring, and we describe baseline results obtained with a simple decentralised approach. Sect. 4 closes the paper with some discussions and perspective on the commercial application of swarm robotics solutions for the agricultural domain.

2. Concept and background

Weed monitoring and mapping is a tough problem that determines the daily activities in a farm. Fig. 1 (left) shows the example of sugar beets that are infested by volunteer potatoes—a common benchmark in agricultural robotics [3, 13]. Volunteer potatoes originate from tubers that remained in the soil after harvesting. In the next season, when sugar beets are grown in the same field, volunteer potatoes are a major threat because they spread diseases (e.g., late blight) and facilitate harmful soil nematodes. Furthermore, they compete for resources with the sugar beets. Regulations obligate farmers to control volunteer potatoes, a very costly operation that involves a lot of human labour. In the Netherlands, the costs have been estimated between €50 and €300 per hectare per growing season [16]. An autonomous weed control system would drastically reduce the costs of removing volunteer potatoes, a task that is largely performed manually. The system proposed within the SAGA experiment aims to take over the field monitoring task, and to generate task maps for future autonomous weeding robots, telling them which areas to work on and how to plan their paths. More specifically, SAGA will provide an automatic weed monitoring and mapping system by means of a swarm of UAVs able to patrol the field, recognise the presence of weeds, dedicate resources to the most interesting areas and collectively build a field map indicating areas with different urgency of intervention. All this has to be obtained through a genuine swarm robotics approach, featuring decentralised control and flexible and scalable behaviour. The SAGA concept represents a novelty within the agricultural robotics domain, despite significant effort and resources being dedicated to agricultural robotics research, including multi-robot approaches. We propose a solution to the monitoring and mapping problem that is completely decentralised, so that desired properties like robustness and scalability are taken into account at design time.

2.1. Collective-level monitoring and mapping

The collective-level control is responsible for the overall mission accomplishment. Instead of a-priori planning the mission for the whole group, we will exploit swarm robotics techniques in which the group behaviour emerges from self-organisation, hence providing flexibility, robustness to faults and scalability with group size. Our goal is to devise collective strategies with an optimal trade-off between distributed exploration and timely weed recognition.

The study of the collective monitoring and mapping behaviour will be initially performed in simulation (see also Section 3), and different bio-inspired algorithms will be evaluated. In particular, we will consider honeybee foraging and collective decision-making as sources of inspiration, and will exploit design patterns to implement such behaviours in UAV swarms [19, 20]. This will provide the mechanisms to explore the field and allocate resources during the monitoring activities: UAVs will be recruited to monitor those areas in the field that have been identified as potentially containing weed patches, while weedless areas are quickly abandoned by the swarm. In this way, resource allocation is adapted to the field heterogeneities, and

[2] *laral.istc.cnr.it/saga*

error-prone individual inspection will be compensated for through collaborative re-sampling. Additionally, we will consider the emergence of a categorisation system from peer-to-peer interactions [2], and implement a collective mapping behaviour as a categorisation problem of different areas of the field, so that labelling of different areas will result from a consensus process among the UAVs that individually estimated the presence of weeds.

2.2. On-board vision for weed detection and navigation

The on-board vision system of each individual UAV has to perform object detection or semantic segmentation on images like in Fig. 1 (left), to count the number of weeds above a specific size or otherwise measure their development. The results of each image have to be mapped to real-world coordinates, using absolute and relative pose estimates from other sensors (e.g., GPS, IMU). This provides the basis for timing and path planning as described above.

Current robotics approaches employ unmanned vehicles (e.g., Bonirob [3]) for weed detection and removal [17, 13]. As demarcation strategy, Bonirob and similar platforms use a protective cover with artificial lighting and cameras underneath. This is a suboptimal solution that simplifies the vision problem with constant, shadowless light conditions. Previous work demonstrated the use of SURF features, bag-of-visual-words clustering, and support vector machines (SVMs) to classify image patches as crop or weed [26]. This method can handle strong light variations and shadows from direct sunlight. The approach can be combined with a sliding window approach or selective search to detect objects in the whole image. State-of-the-art solutions for object recognition make use of convolutional neural networks [12, 10] as in Fig. 1 (left). Such methods are currently being imported with some success also in the agricultural domain [18, 22], but it is also necessary to recognise that the field has specific requirements that needs a deeper investigation of the suitability of such methodologies, Considering the computing and power limitations, convolutional neural networks may be too demanding, hence a trade-off shall be considered between accuracy of the vision algorithm and computational resources available. In this sense, within SAGA we aim at considering accuracy at the swarm level, compensating for individual deficiencies through collaboration and re-sampling by multiple drones. On-board vision can also support the individual motion control. In [28], modified excessive green, Otsu's method, and the Hough transform were combined to find crop rows in a field, and dynamic extrinsic camera calibration and PID-controllers were used to make the UAV follow the detected crop rows. These methods can be used to support GPS-based navigation, or navigation based on ultra-wideband beacons (UWB, see below).

2.3. Hardware enhancement for swarm operations

UAV research is a hot topic to date, pushed by the huge development of commercial flying drones, mainly thanks to the establishment of multi-rotor helicopters bringing to the consumer and industrial market stable and easily controllable platforms. However, UAVs are usually not conceived for group operation, and hardware adaptations are required to have UAVs communicate with each other and coordinate their operations. Additionally, if we consider autonomous flight with on-board sensing, research is still ongoing due to the size and payload constraints associated with aerial vehicles, as well as to the short battery lifetime [11, 21]. Specific control issues emerge for multi-robot settings, which render approaches with individual robot labelling impractical due to the exponential explosion of the state space, and call for low-dimensional abstractions of the group [15]. Additionally, specific sensory systems are required for collision-free flight and networked operations [5, 21]. In summary, several state-of-the-art technologies need to be integrated in a single platform to support swarm operations [30]. Within SAGA, hardware development to enable swarm operation will start from the UAV platform shown in Fig. 1 right. The drone is a quadcopter able to fly up to 30 minutes on a single charge. Key features include a triple redundant autopilot, five inertial measurement units (IMUs) and RTK-GPS. The location and orientation data will be synchronized with the imagery from the RGB camera and the corresponding object detections/semantic segmentation. The standard drone needs to be equipped with several additional hardware modules as well as software communication protocols. The hardware modules include radio-communication between multiple UAVs, based on UWB technology, which will provide at the same time self-localisation with respect to stationary beacons and communication abilities between UAVs. Additionally, the UAV must be enhanced with on-board vision and processing power so as to run the monitoring and motion algorithms. A design based on the Nvidia Jetson[3] will be developed allowing to use the same processor for both machine vision and motion control.

3. Baseline simulation of collective monitoring

Field monitoring is a fundamental activity. It consists of patrolling the field and detecting the weed presence and location. This activity is generally supported by absolute positioning systems (e.g., RTK-GPS) which allows geo-refencing and planning of the optimal path. The most common approach is a "sweeping" strategy, in which a ground or aerial vehicle follows a zigzag course. With multiple vehicles, the field can be decomposed in non-overlapping areas to be assigned to different UAVs [14]. It is often the case that coverage strategies allow to capture large amounts

[3] *www.nvidia.com/object/embedded-systems.html.*

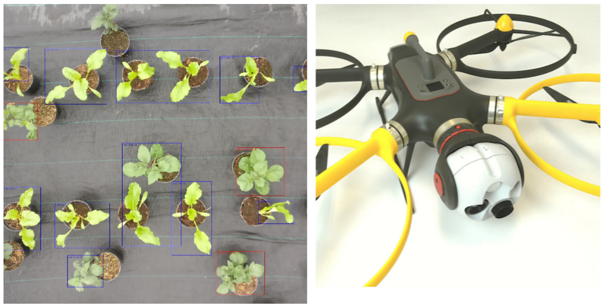

Figure 1: Left: Detection and classification of sugar beets and volunteer potatoes, the latter representing a serious weed problem for large sugar beet cultivations. Right: a close-up view of the PrecisionScout, the UAV produced by Avular B.V. and exploited within the SAGA experiment.

of images to be stitched together and analysed offline [29]. These strategies do not provide robustness against failures of UAVs, neither do they deal optimally with high error-detection rates. Indeed, *a priori* path planning and *a posteriori* analysis do not allow to adapt the monitoring strategy to the actual weed distribution and to exploit online visual processing to influence the field coverage. We aim to produce a completely decentralised solution that exploits online visual feedback to direct the individual search strategy. In this section, we introduce a weed monitoring model and multi-agent simulations developed to test different strategies.

3.1. Weed monitoring model

We consider an abstract scenario in which a square field of side L needs to be monitored The field is virtually divided in square cells of side ℓ, for a total of L/ℓ cells per side, and each cell i can contain one or more weed units, resulting in the weed density ρ_i. We consider here N_w units that are distributed either uniformly in the field, or heterogeneously in N_p patches. Each patch p is obtained as a gaussian spread of items around the patch center $\mathbf{x}_p$ (standard deviation, σ_p, see Fig. 2, top row). Each cell can be visited by an UAV—hereafter, agent—several times. At each visit $k > 0$, an agent a inspects the cell i for τ_v seconds and iteratively updates the locally estimated weed density $\hat{\rho}_{i,a}$ as follows:

$$\begin{aligned} \hat{\rho}_{i,a}(k+1) = (1-\phi_w)\hat{\rho}_{i,a}(k) + \phi_w \rho_i, \\ 0 \le \phi_w \le 1, \\ \hat{\rho}_{i,a}(0) = 0 \end{aligned} \tag{1}$$

where ϕ_w represents the percentage of weed that can be correctly detected in one visit: when $\phi_w = 1$ there is no detection error and one visit is sufficient; when $\phi_w < 1$ more than one visit is necessary. The exponential average models detection from multiple visits as being independent from each other. We compute the globally estimated weed density $\hat{\rho}_i$ of cell i by aggregating information from multiple agents, in a similar way as in equation (1). Additionally, we record the number of visits κ_i that each cell i receives. We consider the field completely covered and correctly inspected when $\forall i,\ \kappa_i > 0 \wedge \hat{\rho}_i = \rho_i$. We evaluate the efficiency of the monitoring activities by looking at the time t_c in which the field is first completely covered, and at the time t_w in which all weed items are correctly detected.

3.2. Baseline collective monitoring strategy

The division of the field in cells allows to simplify the motion strategy of each agent. As a baseline approach, we implement a random-walk-like strategy in which each agent a decides the next cell to visit according to a 2D gaussian distribution. More specifically, the likelihood to choose cell i by agent a is computed as follows:

$$\begin{aligned} F_a(i,j;\sigma_j,\gamma_i) = \gamma_i e^{\frac{-d_{ij}^2}{2\sigma_j^2}}, \\ \sigma_j = \frac{\hat{\sigma}}{1+\hat{\rho}_{j,a}}, \\ \gamma_i = \begin{cases} 1 & \kappa_i = 0 \\ \delta_{i,a} & \kappa_i > 0 \end{cases} \end{aligned} \tag{2}$$

where d_{ij} is the Euclidean distance of cell i from the current cell j, and $\hat{\sigma}$ represents the base spread of the gaussian function. Given the likelihood value for each cell i, a roulette-wheel selection is performed to choose the next cell to visit. In this process, the current cell is excluded, as well as the cells targeted by other agents, which are available thanks to agent-agent communication.[4] By choosing cells according to equation (2), we ensure local exploration thanks to the gaussian spread, we promote longer displacements when the locally estimated weed density $\hat{\rho}_{j,a}$ is low (i.e., high values of σ_j), and we scale the likelihood according to the latest detection improvement $\delta_{i,a}$, which goes to zero the more the estimated weed density approaches the real value, so as to avoid to often revisit the same cells. To summarise, the above strategy implements an isotropic random walk, giving lower importance to areas that have already been sufficiently covered. For comparison, we developed a sweeping strategy in which a single agent $\hat{a}$ covers the whole field by moving through adjacent cells every time the detection improvement $\delta_{j,\hat{a}}$ on the current cell j falls to zero. As mentioned above, the sweeping strategy—although optimal from the efficiency point of view—is not resilient and robust against failures, hence the motivation to develop a collective monitoring system.

[4] We assume here for simplicity a fully connected communication network among agents.

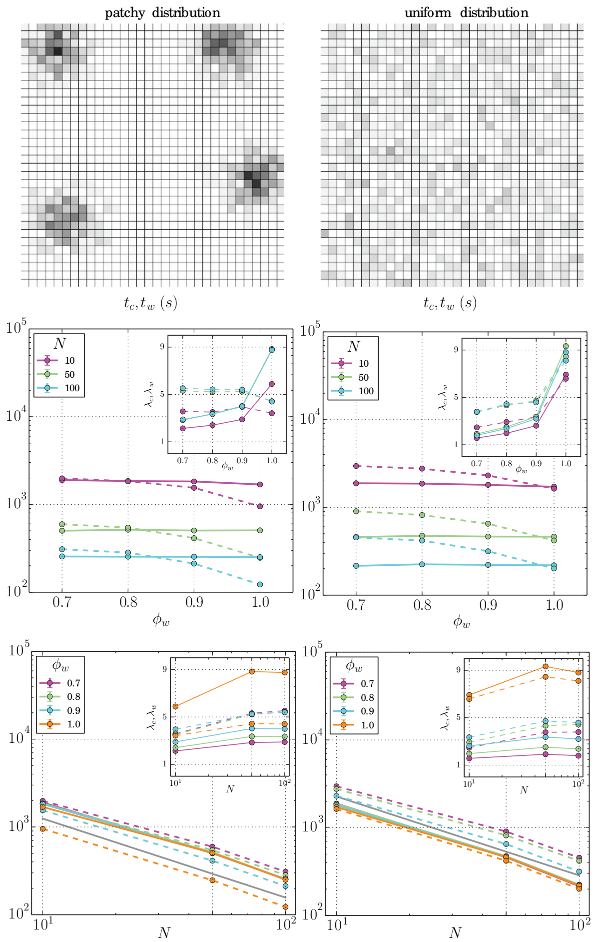

Figure 2: Results of the baseline collective monitoring strategy for patchy (left) or uniform (right) weed distribution. Top: example density map of weed distribution, darker areas correspond to higher density. Middle: Average value for the coverage time t_c (solid lines) and for the monitoring time t_w (dashed) plotted against ϕ_w, for various values of N. Bottom: Scaling of t_c (solid lines) and t_w (dashed) with the group size N. The black solid line is a guide for the eye corresponding to a power-law scaling $t \propto N^{-0.9}$. Insets: relative performance computed against the optimal sweeping behaviour.

3.3. Experimental results

Given the described system, we have performed preliminary investigations to understand the influence of the different parameters on the global outcome. We vary the group size $N \in \{10, 50, 100\}$ and the weed detection rate $\phi_w \in \{0.7, 0.8, 0.9, 1.0\}$. For each configuration we execute 200 evaluation runs in randomly generated fields with either uniform or the patchy weed distribution. We observe that the coverage time t_c is independent of the weed detection rate for both uniform and patchy weed distribution (solid lines in the center plots of Fig. 2). This is expected given that the choice to visit new cells is not affected by ϕ_w. The average values of t_c are also similar for uniform and patchy distributions, as coverage requires to visit every cell in the field at least once. What changes significantly is the course of t_w (dashed lines in Fig. 2), which is decreasing for increasing values of ϕ_w, and above all presents lower values for the patchy weed distribution. Indeed, for high values of ϕ_w, only few visits are required per cell to completely inspect it (at most 2 visits per agent when $\phi_w = 1$), and this has a positive impact on reducing the detection time: in the uniform distribution case, t_w becomes comparable to t_c, as weeds can be found anywhere in the field; in the patchy distribution case, t_w gets significantly lower as weed patches are completely detected before the entire field is fully covered. For what concerns the scaling with the group size N, it is possible to appreciate a power-law decay $t \propto N^{\alpha}$ for both uniform and patchy weed distribution, with exponent $\alpha \approx -0.9$, not too distant from the ideal case of $\alpha = -1$ (bottom plots in Fig. 2). This confirms that the provided solution—although improvable exploiting agent-agent interactions—scales very well with the group size. The performance relative to the optimal weeding strategy is shown in the insets, where $\lambda_c = t_c/t_c^{\star}$ and $\lambda_w = t_w/t_w^{\star}$ are plotted. It is possible to notice that the collective monitoring strategy is in general slower by a factor of 2 to 9, with best performance for low values of ϕ_w, especially for the coverage time, and for the monitoring time in the uniform distribution. For the patchy distribution, monitoring is 3 to 5 times slower than the optimal strategy, but performance slightly improves for $\phi_w = 1$.

4. Conclusions

In this paper, we have presented the SAGA concept and the roadmap it implements to propose swarm robotics as a viable technology for monitoring and mapping applications. We have also proposed an abstract model for weed monitoring and preliminary results exploiting a simple random walk strategy, which constitute a baseline against which to test improved collective monitoring approaches. The baseline monitoring strategy is efficient against a patchy weed distribution, deals well with low rates of weed detection and present good scalability with the group size. It exploits multiple visits from different agents to obtain a complete monitoring. Given the absence of interaction among agents, future work within the SAGA experiment will be dedicated to the engineering of a suitable strategy that minimises the gap with the optimal one, and that at the same time guarantees properties like resilience, robustness and scalability. Possible extensions take into account the ability of the UAVs to move in a 3D space, thus observing the field from different altitudes [23]. This allows to inspect the environment at different resolutions (high for low altitude and vice-versa), whereby monitoring from distance provides a coarse estimate of the weed density to be exploited to allocate resources only towards the most interesting regions. This seems to be a powerful strategy, especially if applied to scenarios that present patches of interest or only a limited number of points of interest over a large area (e.g. census of animals, forest inspection, waste detection and more).

References

[1] F. A. Auat Cheein and R. Carelli. Agricultural Robotics: Unmanned Robotic Service Units in Agricultural Tasks. *IEEE Industrial Electronics Magazine*, 7(3):48–58, Sept. 2013.

[2] A. Baronchelli, T. Gong, A. Puglisi, and V. Loreto. Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences of the United States of America*, 107(6):2403–2407, Feb. 2010.

[3] P. Biber, U. Weiss, M. Dorna, and A. Albert. Navigation system of the autonomous agricultural robot Bonirob. In *Workshop on Agricultural Robotics: Enabling Safe, Efficient, and Affordable Robots for Food Production*, 2012.

[4] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, Jan. 2013.

[5] L. Doitsidis and *et al.* Optimal surveillance coverage for teams of micro aerial vehicles in GPS-denied environments using onboard vision. *Autonomous Robots*, 33(1-2):173–188, Mar. 2012.

[6] M. Dorigo and *et al.* Swarmanoid: A Novel Concept for the Study of Heterogeneous Robotic Swarms. *IEEE Robotics & Automation Magazine*, 20(4):60–71, 2013.

[7] M. Gauci, J. Chen, W. Li, T. J. Dodd, and R. Gross. Self-organized aggregation without computation. *The International Journal of Robotics Research*, 33(8):1145–1161, July 2014.

[8] R. Gebbers and V. I. Adamchuk. Precision Agriculture and Food Security. *Science*, 327(5967):828–831, Feb. 2010.

[9] H. Hamann and H. Wörn. A framework of space–time continuous models for algorithm design in swarm robotics. *Swarm Intelligence*, 2(2-4):209–239, 2008.

[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[11] V. Kumar and N. Michael. Opportunities and challenges with autonomous micro aerial vehicles. *The International Journal of Robotics Research*, 31(11):1279–1291, Sept. 2012.

[12] Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

[13] P. Lottes and *et al.* An effective classification system for separating sugar beets and weeds for precision farming applications. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5157–5163. IEEE, 2016.

[14] I. Maza and A. Ollero. Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms. In *Distributed Autonomous Robotic Systems 6*, pages 221–230. Springer Japan, Tokyo, 2007.

[15] N. Michael and V. Kumar. Control of Ensembles of Aerial Robots. *Proceedings of the IEEE*, 99(9):1587–1602, Aug. 2011.

[16] A. T. Nieuwenhuizen, J. W. Hofstee, and E. J. van Henten. Performance evaluation of an automated detection and control system for volunteer potatoes in sugar beet fields. *Biosystems Engineering*, 107(1):46–53, Sept. 2010.

[17] A. T. Nieuwenhuizen, L. Tang, J. W. Hofstee, J. Müller, and E. J. van Henten. Colour based detection of volunteer potatoes as weeds in sugar beet fields using machine vision. *Precision Agriculture*, 8(6):267–278, Nov. 2007.

[18] C. Potena, D. Nardi, and A. Pretto. Fast and accurate crop and weed identification with summarized train sets for precision agriculture. In *Proc. of 14th International Conference on Intelligent Autonomous Systems (IAS-14)*, 2016.

[19] A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102, 2015.

[20] A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A Design Pattern for Decentralised Decision Making. *PLoS ONE*, 10(10):e0140950–18, Oct. 2015.

[21] J. Roberts, T. Stirling, J.-C. Zufferey, and D. Floreano. 3-D relative positioning sensor for indoor flying robots. *Autonomous Robots*, 33(1-2):5–20, Feb. 2012.

[22] I. Sa, Z. Ge, F. Dayoub, B. Upcroft, T. Perez, and C. McCool. DeepFruits: A Fruit Detection System Using Deep Neural Networks. *Sensors*, 16(8):1222–23, 2016.

[23] S. A. Sadat, J. Wawerla, and R. T. Vaughan. Recursive non-uniform coverage of unknown terrains for uavs. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 1742–1747. IEEE, 2014.

[24] D. L. Shaner and H. J. Beckie. The future for weed control and technology. *Pest Management Science*, 70(9):1329–1339, Jan. 2014.

[25] D. C. Slaughter, D. K. Giles, and D. Downey. Autonomous robotic weed control systems: A review. *Computers and Electronics in Agriculture*, 61(1):63–78, Apr. 2008.

[26] H. K. Suh, J. W. Hofstee, J. IJsselmuiden, and E. J. Van Henten. Discrimination between volunteer potato and sugar beet with a bag-of-visual-words model. In *International Conference of Agricultural Engineering*, 2016.

[27] V. Trianni and A. Campo. Fundamental Collective Behaviors in Swarm Robotics. In J. Kacprzyk and W. Pedrycz, editors, *Springer Handbook of Computational Intelligence*, pages 1377–1394. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.

[28] K. van Boheemen. Autonomous UAVs in agriculture, navigation and control using real-time image analysis. *BSc thesis Wageningen University*, June 2015.

[29] C. Zhang and J. M. Kovacs. The application of small unmanned aerial systems for precision agriculture: a review. *Precision Agriculture*, 13(6):693–712, July 2012.

[30] J. Zufferey, S. Hauert, T. Stirling, S. Leven, and J. Roberts. Aerial collective systems. *Handbook of Collective Robotics*, 2013.