Field Coverage and Weed Mapping by UAV Swarms

Dario Albani,[°▽] Daniele Nardi[°] and Vito Trianni[▽]

***Abstract*— The demands from precision agriculture (PA) for high-quality information at the individual plant level require to re-think the approaches exploited to date for remote sensing as performed by unmanned aerial vehicles (UAVs). A swarm of collaborating UAVs may prove more efficient and economically viable compared to other solutions. To identify the merits and limitations of a swarm intelligence approach to remote sensing, we propose here a decentralised multi-agent system for a field coverage and weed mapping problem, which is efficient, intrinsically robust and scalable to different group sizes. The proposed solution is based on a reinforced random walk with inhibition of return, where the information available from other agents (UAVs) is exploited to bias the individual motion pattern. Experiments are performed to demonstrate the efficiency and scalability of the proposed approach under a variety of experimental conditions, accounting also for limited communication range and different routing protocols.**

I. Introduction

Current advancements in remote sensing technologies—especially exploiting small and efficient unmanned aerial vehicles (UAVs) [1]—are revolutionising the agricultural domain, providing loads of data to be exploited for process optimisation in nearly every possible activity related to agricultural production. In this way, the tenet of precision agriculture (PA) to "produce more with fewer inputs" becomes possible. However, PA requires fine-grained data and intervention abilities to maximise yields and minimise the usage of water, fertilisers and herbicides (e.g., through variable-rate applications) [2], [3]. High-resolution satellite images have been considered the main source of information until recently, but the trend is changing thanks to the extensive deployment of UAVs, which allow to reduce problems related to high costs and limited availability of satellite imagery [4], [5]. As soon as data are needed at the single plant level (e.g., to recognise the weed species and sizes and select the best-performing herbicide type and dose to apply), strong constraints are imposed on the remote sensing technology to be applied. Fixed-wing UAVs become unpractical as they cannot hover and need to fly at a too high altitude, shifting the choice towards rotary-wing UAVs like quad-copters. These are however limited in flight time, given the current state of the art, demanding for time/energy-efficient solutions. PA also requires higher levels of autonomy for UAVs largely beyond the current usage as passive sensors with predefined mission plans. Intelligent, real-time adaptation of the UAVs sampling policies to reflect the information being gathered allows to focus on areas of importance (e.g., where weeds are concentrated), and avoid wasting resources on areas of lower relevance (e.g., devoid of weed). Similar strategies can be implemented on a UAV endowed with higher autonomy, but are not flexible and efficient in terms of time and energy expenditure. Additionally, it is important to note that, despite hardware prices going down with the growth of the drone market, high-end solutions for agricultural applications with onboard autonomy and high precision sensors would still be rather costly. As a consequence, an autonomous UAV is an asset that does not allow to scale costs with different farm sizes, and does not provide robustness against faults during operation accomplishment. The above discussion motivates the present study to exploit a swarm robotics approach [6], [7], using groups of miniaturised low-cost UAVs that can provide—collectively—similar precision as high-end solutions, can be efficiently scaled to different farm sizes without performance loss, and provide intrinsic robustness against individual faults.

In this study, we focus on a field coverage and mapping problem, whereby the presence and density of weeds must be detected in a crop field, therefore creating a infestation map that can be exploited for weed control optimisation. We present a simulation study to compare the efficiency of a swarm of collaborating UAVs capable of on-board weed recognition and online mission adaptation, against the performance of a standard remote-sensing approach by which a predefined image-collection mission is followed by off-board batch processing. Considering that the number and spatial distribution of weeds is unknown, a complete coverage of the field is a necessary condition. Additionally, weed infestation is often heterogeneous in the field, with patches of plants separated by regions devoid of weeds, requiring adaptive information foraging abilities. Finally, the identification of weeds could be prone to errors (e.g., due to varying weather conditions), so much that a large resampling effort may be required to obtain a sufficient accuracy.

Coverage and mapping are interrelated problems largely studied in multi-agent systems and robotics, although not specifically linked to agricultural applications [8], [9], [10]. Generally speaking, a certain number of points of interest (POIs, i.e., weeds) must be reached (coverage) and processed

*This work has been conducted within the project SAGA (Swarm Robotics for Agricultural Applications, see http://laral.istc.cnr.it/saga), an experiment founded within the ECHORD++ project (GA: 601116). Dario Albani and Daniele Nardi acknowledge partial support from the european project FLOURISH (H2020-ICT-23-2014, ID 644227). The authors would like to thank Prof. Sean Luke from the George Mason University for the precious support in developing the simulations with Mason.

°Department of Computer, Control, and Management Engineering "Antonio Ruberti" at Sapienza University of Rome, Via Ariosto 25, 00185 Rome, Italy. name.surname@diag.uniroma1.it

▽Institute of Cognitive Sciences and Technologies (ISTC), National Research Council (CNR), Via San Martino della Battaglia 44, 00185 Rome, Italy. name.surname@istc.cnr.it

(mapping). This makes the problem similar to task allocation in multi-agent systems, which has been approached with different methods like Distributed Constraint Optimization (DCOP [11]), Distributed Pseudotree Optimization (DPOP [12]), the Contract Net Protocol [13] or bounty-based auction methods [14], to cite some. Few studies take a swarm robotics approach to coverage and mapping. In [15], the coverage problem is tackled through a flocking algorithm ensuring permanent connectivity among robots. Similarly in [16], multi-robot coverage is performed to maximise spread while maintaining connectivity among robots and reducing the communication overhead. The focus on connectivity is justified by the need to largely spread information within the group. In this study, we relax this constraint to provide more efficient coverage solutions, and we introduce re-broadcast protocols to account for limited communication ranges.

In order to provide scalability and robustness to the system, we devise a stochastic exploration strategy based on a reinforced random walk [17], [18], [19]. Individual UAVs follow a correlated random walk and interact with neighbours to avoid interferences (see Section II-B). A mechanism to avoid returning on previously visited areas leads to a very efficient coverage strategy that is only slightly inferior to the optimal case of a lawnmower strategy (see Section III). Communication is exploited here to share information between UAVs so as to avoid re-visiting areas that have been already visited before. Finally, mapping weeds in the field is performed by UAVs recruiting each other towards areas of possible interest—exploiting communication to create virtual pheromones or "beacons" [20], [8], [21]—hence increasing the sampling effort only where this is needed (see Section II-C for details). Overall, the results obtained suggest that swarm robotics can provide useful approaches to PA. Finally, Section IV discusses the merits and limitations of the proposed solution, and points to possible future improvements.

II. Experimental setup

A. Field simulation and weed detection model

We consider the case of a open field without obstacles, defined by the field boundary. Without loss of generality, we choose a rectangular field with dimension $L_x \times L_y$. The field is divided in a single 2D $n \times m$ matrix. A single i, j location is referred to as a cell c_{ij}, and may contain zero, one or multiple weeds with overall density ρ_{ij}. We assume a heterogeneous distribution of weeds, which appear in patches across the field, each patch being modelled as a gaussian distribution. UAVs—hereafter also called agents—move at constant speed v and can hover over a cell to take images of the field and distinguish between background, crops and weeds. We assume that the vision algorithm for weed detection returns, for each image acquired for cell c_{ij}, the detected weed density $\hat{\rho}_{ij} \leq \rho_{ij}$, and a confidence value $z_{ij} \in [0, 1]$. We do not consider here the possibility of false positives. Indeed, the vision processing algorithm can be affected by detection errors—e.g., due to variable weather conditions or hovering fluctuations—that we model with a fixed probability P_ϵ of incurring in errors that result in a lower detection ability, so that the estimated weed density and confidence value are obtained as follows:

$$\hat{\rho}_{ij}(t) = \rho_{ij}(1 - \epsilon_w(t)), \qquad \epsilon_w(t) \sim \mathcal{U}(0, \epsilon_M) \tag{1}$$

$$z_{ij}(t) = 1 - \epsilon_z(t), \qquad \epsilon_z(t) \sim \mathcal{U}(0, \epsilon_w(t)) \tag{2}$$

where $\epsilon_M \in [0, 1]$ represents the maximum percentage of missed weeds we consider possible, and $\mathcal{U}(a, b)$ represents a uniform distribution in the interval $[a, b]$.

Each agent maintains a local map of the field, where to store, for each cell c_{ij}, the information acquired from onboard processing or received from other agents through communication. In case multiple observations of the same cell are performed—by the same or different agents—we assume that information is aggregated by discarding the one with the lower confidence:

$$\langle \hat{\rho}, z \rangle_{ij} = \max_z \left(\langle \hat{\rho}_1, z_1 \rangle_{ij}, \langle \hat{\rho}_2, z_2 \rangle_{ij} \right) \tag{3}$$

where $\langle \hat{\rho}_h, z_h \rangle_{ij}$ represents the result of a single detection process on cell c_{ij}. Different approaches can be taken in future work to account for the actual processing performed by the UAVs. We let the UAVs explore the field until complete coverage, that is, when each cell has been visited at least once by any UAV. We therefore record the *coverage time* t_c, which is the principal metric for evaluating the system efficiency. To evaluate the mapping quality, we consider the global detection efficiency, computed as follow:

$$d_{ij}^h = \frac{\hat{\rho}_{ij}^h}{\rho_{ij}} \tag{4}$$

$$D = \frac{1}{nm} \sum_{ij} \max_h d_{ij}^h \tag{5}$$

where d_{ij}^h is the detection efficiency of agent h in cell c_{ij}, and D is the overall, aggregated detection efficiency.

B. Swarm exploration strategy for field coverage

The basic strategy for field coverage adopted by the agents is a simple correlated random walk [18], [19], on top of which we implement several mechanisms to improve the exploration efficiency. Given that the field is divided in cells, the UAV motion planning is performed by selecting a new cell to visit so that weed detection can be performed according to the model presented above. A random walk for each UAV is implemented by selecting, at each decision step, a random cell to be visited in the neighbourhood. To minimise the distance covered, the choice is made within sets of cells at increasing distance, as depicted in Fig. 1 left, where cells belonging to the same set are identified by the same ordinal number R. For instance, the set with $R = 1$ corresponds to the cells directly adjacent to the agent's current one. To choose the next cell, an agent makes a two-step decision. In the first step, it looks for a sufficient number of *valid* cells, that is, cells that have not been previously visited or that are not occupied/targeted by other agents, to the best of the local knowledge available (see also Sect. II-D). Validity of cells is checked sequentially for sets of increasing distance R, until a given number V of cells is discovered (in

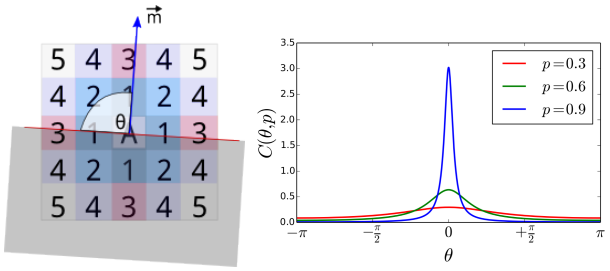

Fig. 1. Left. Distance-based exploration pattern. In the image the central cell marked with an A represents the agent's position. Numbers indicate the priority of each location assigned according to the distance. Vector $\vec{m}$ is the momentum of the agent, representing a possible directional bias. Cells that fall in the shadowed region are accepted only if no valid cell is found in the other semi-plane. Right: The wrapped Cauchy density function for different values of persistence p.

this study, $V = 1$). At this point, a set $\mathcal{V}$ is defined including all valid cells within the maximum distance reached. Note that, in general, $|\mathcal{V}| > V$, as all the cells within a given distance are included, not limited to those that allow reaching the threshold V. To avoid that the selection of valid cells interferes with possible directional biases for cell selection (i.e., when the first valid cells lay in the opposite direction), an agent divides the exploration plane in two parts (semi-planes), hence choosing valid cells first in the semi-plane complying with the directional bias, and only if unsuccessful, in the opposite semi-plane. More precisely, given the agent h position $\vec{x}_h$, a directional bias $\vec{b}$ and the relative position of the cell $\vec{r}_{ij} = \vec{x}_{ij} - \vec{x}_h$, the semi-plane with priority is given by all cells c_{ij} for which $\vec{b} \cdot \vec{r}_{ij} \geq 0$. If no valid cell is available satisfying this conditions, then also the remaining cells are evaluated for inclusion in $\mathcal{V}$ (see Fig. 1 left).

In the second step, a random choice is performed within $\mathcal{V}$ to select the target cell. To implement a correlated random walk, we consider a unit vector $\vec{m}_h$ representing the momentum of the agent h. Each cell $c_{ij} \in \mathcal{V}$ is assigned a utility according to the angular difference θ_{ij} as follows:

$$u_{ij} = C(\theta_{ij}, p), \quad C(\theta, p) = \frac{1}{2\pi} \frac{1 - p^2}{1 + p^2 - 2p\cos\theta} \tag{6}$$

where θ_{ij} is the angle between the cell and the momentum $\vec{m}_h$, while $C(\cdot)$ is the wrapped Cauchy density function with persistence $p \in [0, 1[$ which determines the function skewness, as shown in Fig. 1 right. For high values of persistence, the utility is very high only for cells aligned with the momentum vector, while for a low value of persistence the utility is more uniformly assigned despite the angular deviation. Additionally, the influence from neighbouring agents shall be considered to avoid interferences and to bias random movements in areas with low agent density. To this end, we compute for each agent h a repulsion vector $\vec{r}_h$ as follows:

$$\vec{r}_h = \sum_{k \neq h} S(\vec{x}_h - \vec{x}_k, \sigma_a), \; S(\vec{v}, \sigma) = 2e^{\mathrm{i}\angle\vec{v}} e^{-\frac{|\vec{v}|}{2\sigma^2}}, \tag{7}$$

where the sum extends over all agents k in the neighbourhood that are known to agent h (see also Sect. II-D), and the function S returns a vector pointing away from agent k, with the module decaying according to a gaussian function with width σ_a. To account for both momentum and repulsion, a bias vector is computed as $\vec{b}_h = \vec{m}_h + \vec{r}_h$ and a utility is assigned to each cell $c_{ij} \in \mathcal{V}$ according to eq. (6) using the angle between the cell and $\vec{b}_h$. A utility-proportional random choice is then made among the cells in the valid set $\mathcal{V}$.

C. Swarm Mapping Strategy

Similarly to field coverage, weed mapping is performed by UAVs exploiting a reinforced random walk strategy. Here, the decision about the next cell to visit is also influenced by the mapping activity of other UAVs, that can recruit each other towards areas of interest. We exploit the concept of virtual pheromones or "beacons" [20], [8], [21], that is, attractive points that bias the individual motion. A beacon is activated within a cell by an agent, following the recognition of some weeds. More precisely, an agent activates a beacon when the confidence value $z_{ij} < 1$, meaning that additional monitoring effort is required, and communicates its position to the neighbouring agents. Agents maintain a list of active beacons $\mathcal{B}$—characterised by position $\vec{x}_b$ and activation time t_b—and remove items from the list after a period T_b from activation. To limit the proliferation of beacons, agents are restrained from activating new beacons if their local list contains more than M_b items. Whenever an agent broadcasts some message, it also attaches the local list of available beacons $\mathcal{B}$, so that receiving agents can update their local list accordingly (see also Sect. II-D). This ensures an efficient diffusion of information with a small communication overhead, given the limited number of active beacons M_b.

The motion strategy for mapping is similar to the coverage case, with the following differences. The set of valid cells $\mathcal{V}$ is built in a similar way, but now includes all cells c_{ij} that have a confidence value $z_{ij} < 1$. We assume that uncovered cells have null confidence. As for coverage, cells that are occupied or targeted by other agents are excluded from $\mathcal{V}$.

Given the set of valid cells, a random choice is performed by agent h taking into account the agent momentum $\vec{m}_h$, the repulsion vector from other agents $\vec{r}_h$ and an attraction vector resulting from the known active beacons:

$$\vec{a}_h = \sum_b S(\vec{x}_b - \vec{x}_h, \sigma_b), \tag{8}$$

where the gaussian function S from equation (7) is now characterised by the width parameter σ_b. Hence, the resulting bias vector is $\vec{b}_h = \vec{m}_h + \vec{r}_h + \vec{a}_h$, which is used to assign the utility to each cell $c_{ij} \in \mathcal{V}$ according to eq. (6).

Every time an agent reaches a new cell, it maps weeds and updates its local information following the model described in Sect. II-A. Then, it shares the updated information with its neighbours following the protocol described in the following.

D. Agents' Communication

For activity coordination and efficient operations, communication among UAVs is essential. By broadcasting their absolute position, UAVs can implement collaborative avoidance strategies (e.g., the hybrid reciprocal velocity obstacle

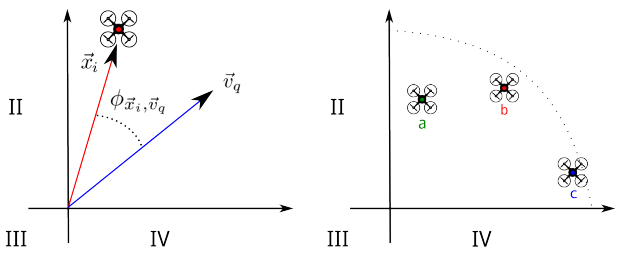

Fig. 2. Left: A graphical representation of the variables considered in the broadcast utility function. Right: A visual representation of a possible displacement of UAVs in a single quadrant. The dotted circle represents the communication range (how far the transmitting agent can send the data). In this example, b (the red UAV) is selected for rebroadcast while the other two UAVs act as simple listeners.

method [22]). Additionally, UAVs broadcast information about the visited cells and active beacons to support field coverage and weed mapping. In this work, we neglect collision avoidance issues, and we focus on the information required for monitoring and mapping. We assume that an agent h broadcasts a message after visiting a new cell. Such a message contains: (i) the current cell c_{ij}, (ii) the estimated weed density ρ_{ij} and confidence z_{ij}, (iii) the next cell c'_{ij} chosen for visiting, (iv) the list of known active beacons $\mathcal{B}$, and (v) a unique identifier of the message.

The ability to spread such information widely within the swarm is key for efficient coordination. It is therefore important to study how communication affects both the behaviour of the agents and the overall swarm efficiency. Information spreading depends on the communication range R_c, and on the chosen re-broadcasting protocol. The communication range and the density of agents in the field determines the features of the static interaction network, and the speed at which information propagates [23]. The re-broadcasting protocol determines how efficient the communication is in reaching agents via multi-hop forwarding. In this work, we adopt concepts from information-centric mobile ad-hoc networks (ICMANET [24]), in which the host-centric paradigm is abandoned in favour of data content. To understand the effect of the chosen communication protocol, we implemented three different strategies:

- **Simple**—A single-broadcast communication protocol without re-broadcast, in which a message reaches only the neighbours within range R_c.
- **Flooding**—A multi-broadcast protocol where messages are forwarded as long as the same content—based on the unique identifier of the message—was not already delivered before [25].
- **Geo-aware**—A variant of the "geo-aware energy efficient approach" [26], adapted to account for a custom utility function. This function is computed by the origin of the message, for each quadrant of the plane and for each known agent, and selects a target agent on the basis of the distance reached and the centrality with respect to the quadrant bisector:

$$\arg\max_{i\in\mathcal{A}_q} u_q(i) = |\vec{x}_i| \cdot (1 - |\frac{\phi_{\vec{x}_i,\vec{v}_q}}{\frac{\pi}{4}}|) \tag{9}$$

where $\mathcal{A}_q$ is the subset of known agents in the quadrant q, the versor $\vec{v}_q$ splits the quadrant q in two identical parts, and ϕ represents the angle between the agent position vector and $\vec{v}_q$. The agent with maximum utility is chosen for rebroadcast (see Fig. 2). Also in this case, agents never forward the same message more than once.

Note that, despite identical messages are not re-broadcast, an agent processes any received message to update the information about the neighbourhood. Here, we do not consider any communication error, to focus on the effects of the communication range and protocol over the efficiency. To this end, we define two simple metrics: (i) the total number of messages exchanged and (ii) the average utility of received messages. The former is used as a proxy of the load on the communication channel, and is computed as follows:

$$E_{tx} = \sum_i N_{tx}(i) \tag{10}$$

where $N_{tx}(i)$ is the number of transmitted messages by agent i. The utility encodes the amount of redundant information among all received messages, and is computed as follows:

$$E_{rx} = \frac{1}{NE_{tx}} \sum_i \sum_m \frac{1}{N_{rx}(m,i)} \tag{11}$$

where N is the number of agents and $N_{rx}(m,i) \geq 1$ the number of times a message m was received by agent i.

E. Reference lawnmower strategy

As a demarcation strategy, we consider here a lawnmower agent that moves from one cell to a neighbouring one, sweeping the whole field. Such a strategy provides an optimal coverage time as it minimises the distance covered from cell to cell, and sequentially visits all cells in the field. To compare with the proposed approach, we simply divide the coverage time $t_c^\star$ of a single lawnmower agent by the number N of available agents, as if each agent was *a priori* assigned an equal, non-overlapping portion of the field to be covered. By comparing with such an optimal coverage time, we can appreciate how good is the devised stochastic strategy.

For what concerns weed mapping, we assume here that the reference lawnmower agent does not have any ability of onboard processing, so that a number N_s of pictures are shot for every cell of the field for off-board processing, each shot taking a single simulation step. This constitutes a reasonable assumption considering the range of commercial applications for remote sensing with UAVs. We therefore compare the proposed approach for mapping with the coverage time of N lawnmower agents (possibly increased by the N_s steps necessary to take additional images) and the global detection efficiency $D^\star$, which results from exploiting the same vision processing model described in Sect. II-A on N_s images, ignoring possible correlations of errors between subsequent samples.

III. Results

To evaluate the merits and limitations of the proposed approach, we use Mason [27], a professional framework for multi-agent simulations useful to test our approach largely varying the parameter set. For each experimental conditions, we have performed 100 runs always using a field divided in 50×50 cells. Agents are positioned uniformly random within the field at the beginning of each simulation run.

A. Efficiency in field coverage

We consider a swarm of $N \in \{10, 50, 100\}$ agents, using an infinite communication range and a single-broadcast protocol, which is sufficient to ensure that each agent has perfect knowledge of the system state. We investigate how the coverage efficiency is influenced from both presence of neighbours and the persistence of the biased random walk. The former is determined by σ_a in eq. (7): the higher the value, the wider the influence range of an agent, so that even very distant agents can influence each other. The persistence parameter p in eq. (6) instead determines the importance of the directional bias in determining the choice of the next cell.

Overall, the results presented in Fig. 3 indicate a good efficiency for the proposed coverage strategy, with about 20% overhead with respect to the optimal lawnmower coverage for the best parameterisations. The best conditions generally correspond to high persistence p and low to medium repulsion σ_a among agents. For persistence $p = 0$, the directional bias has no effect and the next cell to visit is chosen in the whole neighbourhood, disregarding the agent momentum and the repulsion from other agents. In this case, we observe that coverage efficiency is higher for smaller groups, because larger groups tend to leave small portions of the field uncovered, hence requiring large displacements and longer times as the number of valid cells diminishes. When $\sigma_a = 0$, there is no repulsion between agents, and the directional bias is limited to the momentum $\vec{m}$, resulting in a correlated random walk. Larger persistence values produce a kind of lawnmower strategy in which an agent goes straight as long as valid cells are present in the motion direction, otherwise choses a different direction of motion. Very efficient coverage is observed in such condition for $N = 10$, as the agents are initially spread in the field and do not interfere strongly. A lower advantage is observed for larger groups, in which more coordination is required to limit interferences. Indeed, for larger groups more coordination is required to limit interferences. Higher values of persistence and repulsion allow to account for the presence of other agents. The best conditions are observed for high persistence ($p \geq 0.6$) and medium repulsion ($\sigma_a \in \{8, 16\}$), for which agents divide the field in non-overlapping territories within which a lawnmower-like motion is observable. Higher values of repulsion lead to an excessive segregation on the borders and/or to the cancellation of the repulsion effects when the density of agents is high (e.g., for $N \geq 50$).

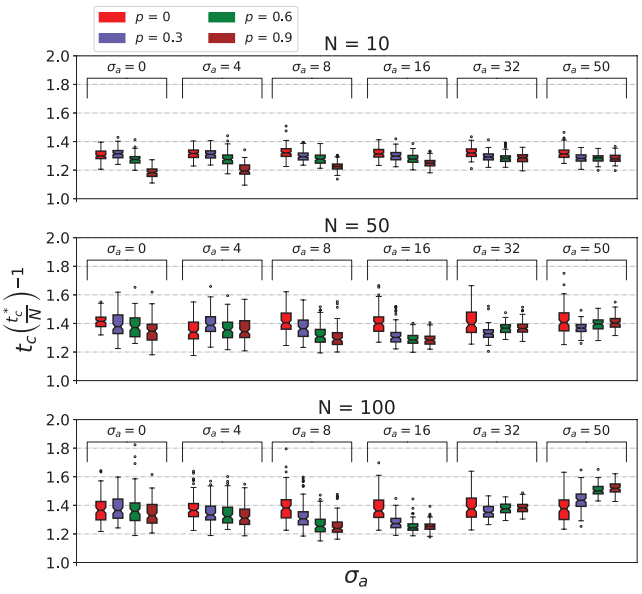

Fig. 3. Results obtained varying the number of agents $N \in \{10, 50, 100\}$, the spread of the repulsion among agents $\sigma_A \in \{0, 4, 8, 16, 32, 50\}$ and finally persistence $p \in \{0, 0.3, 0.6, 0.9\}$.

B. Effects of communication range and protocol

The coverage efficiency so far discussed assumes an infinite range for communication ($R_c = \infty$). When reducing the communication range, the coverage efficiency significantly decreases, due to the fact that agents have a partial information of the world and tend to revisit areas that have been already covered by other agents. We performed additional simulations with $R_c \in \{5, 10, 20, 40\}$ cells, and we also

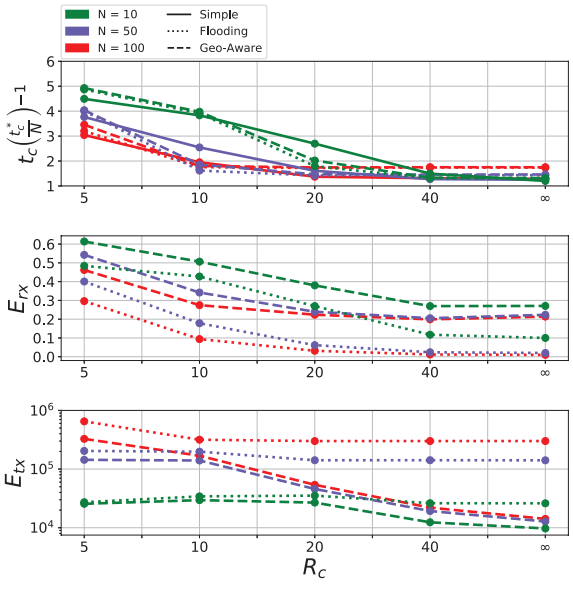

Fig. 4. Results obtained varying the communication range $R_c \in \{5, 10, 20, 40, \infty\}$, the type of communication protocol and number of agents $N \in \{10, 50, 100\}$. These simulations are performed setting $\sigma_a = 8$ and $p = 0.6$. Top: coverage time relative to the reference lawnmower strategy. Center: Utility of delivered message. Bottom: Total number of messages transmitted.

tested the effect of different multi-broadcast protocols, which can help diffusing information within the swarm (Fig. 4). With a single-broadcast protocol, performance quickly degrades with shorter communication ranges, especially for small groups ($N = 10$). The usage of a multi-broadcast protocol alleviates significantly the problem, having efficient coverage for R_c as small as 10 cells for $N \geq 50$. The advantage of the multi-broadcast protocols comes at the cost of more messages transmitted and a lower overall utility (see the center panel in Fig. 4). Utility decreases with larger R_c because of the larger amounts of re-broadcasts from the agents that are reached by a message. The "Geo-aware" protocol proves much better in minimising the number of messages transmitted (see the bottom panel in Fig. 4) and maximising the utility E_{rx}, without affecting the coverage efficiency.

C. Efficiency in weed mapping

The weed mapping results build over those obtained for coverage. Here, we introduce the possibility of agents to use beacons to attract other agents towards relevant areas. We run experiments for varying attractiveness of beacons as determined by $\sigma_b \in \{4, 8, 16, 32\}$ and for varying swarm size $N \in \{10, 50, 100\}$. To also evaluate the effects of communication, we perform tests with the Geo-aware protocol with $R_c \in \{10, \infty\}$. We compute the coverage time t_c and the detection efficiency D, and we compare the obtained values with the time taken by the reference lawnmower strategy with $N_s = 1$ (i.e., a single image captured per cell) and with $N_s = 10$ (i.e., multiple images to account for detection errors). We use here a maximum number of beacons linked to the number of agents: $M_b = N$. Additionally, we consider the following detection errors: $P_\epsilon = 0.75$ and $\epsilon_M = 0.25$.

Generally speaking, the detection accuracy is very high, as the agents re-visit multiple times the cells containing weed items, therefore substantially obtaining a mapping efficiency close to 1 in all tested conditions (the average over all the performed runs amounts to $D = 0.9998 \pm 0.0001$ for $R_c = 10$ and $D = 0.9996 \pm 0.00015$ for $R_c = \infty$). The reference strategy instead has an average error of 0.12 with $N_s = 1$, while the average error decreases to 0.04 for $N_s = 10$. The high performance obtained for the detection efficiency derives mainly from the ability to attract agents towards cells that have not been correctly mapped. In this way, a very good performance can be reached within the time the whole field gets covered. The high performance comes however at the cost of a longer coverage time t_c (see Fig. 5), given that agents tend to focus on unmapped areas before visiting cells in zones devoid of weed. The best conditions is for an infinite communication range and small groups. Here, the proposed approach has comparable performance with the lawnmower strategy with $N_s = 1$, and provides higher detection accuracy. Coherently with the coverage results, performance decreases for the small communication range $R_c = 10$, where the best results are found for $N = 50$ agents. It is worth noting that the system is faster for lower attractiveness of the beacons, which means that local attraction is sufficient for good mapping. The comparison with $N_s = 10$ produces similar results, although more favourable to the swarm robotics approach of about 30%, as the difference is limited to a longer time taken by the lawnmower agents to terminate the coverage.

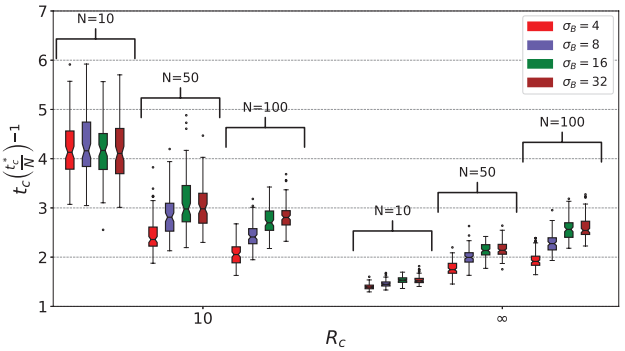

Fig. 5. Results obtained varying the number of agents $N \in \{10, 50, 100\}$, attractiveness of beacons $\sigma_B \in \{4, 8, 16, 32\}$ and the communication range $R_c \in \{10, \infty\}$. These simulations are performed setting $\sigma_a = 8$ and $p = 0.6$. The communication protocol is the geo-aware.

IV. Conclusions and Future Works

We presented a distributed, stochastic field coverage and weed mapping strategy for a UAV swarm, and analysed its efficiency in relation to a reference lawnmower strategy. We deeply analysed the deployed system by systematically varying the parameters that have a bearing on the efficiency of the collective coverage and mapping behaviour, explaining the merits and limitations of different choices. We accounted for limited communication ranges and for different broadcasting protocols, hence providing valuable information for future implementations with real drones. The obtained results are satisfactory and confirm our expectations, that a swarm of UAVs can represent an efficient solution for monitoring and mapping problems compared to current approaches based on image collection and off-board processing. We have shown that the proposed system scales reasonably well with the group size without requiring specific tuning. The decentralised, self-organising nature of the system also leads to intrinsic robustness against faults, which is a valuable property for realistic applications. Note that the presented results have been obtained under fair assumptions with respect to the detection accuracy, as the reference lawnmower strategy employs exactly the same methodology. Additionally, assumptions about coverage in the lawnmower strategy are rather optimistic, postulating a perfect linear scaling. On such basis, the achieved performance can be positively judged.

One important limitation of the proposed methodology, which requires further investigation, concerns the decrease in performance for small communication ranges. Considering that crop fields can be very extensive, the communication range may be an issue especially when connectivity among UAVs is sporadic and newly gathered information may get lost or not efficiently spread within the group. Further studies are required to determine how and with what frequency

the information available to each agent should be shared with neighbours, so as to align as much as possible the local knowledge about the world to the actual conditions, and average out possible errors. This will also need to take into account the requirements for efficient communication, so as to avoid overloading the communication channel. A possible solution can be given by trading data rate for longer communication ranges. Indeed, by allowing for better information sharing, the swarm robotics approach has very good properties, as shown within this study when $R_c = \infty$.

Future work will aim to improve the system in view of deployment on real UAVs. To this end, the inclusion of collision avoidance strategies is a requirement for testing with real UAVs. Further improvements are sought to improve the coverage efficiency and the mapping accuracy. An interesting possibility in this direction corresponds to introducing non-uniform coverage strategies [28], [29]. This can be obtained by introducing different observation layers (e.g., different flying altitudes for the UAVs), which provide different detection accuracy. Coverage can be performed mainly at the coarser level, hence providing basic information for mapping, which can be instead performed with higher accuracy by descending into fine-grained observation layers (e.g., lowering the hovering altitude to obtain more precise information). In this way, it could be possible to further increase the coverage efficiency and the detection ability. Finally, our approach has been introduced as specific for the precision agriculture domain, but it can be relaxed and modified to fit a great variety of contexts. A generalised model for monitoring and mapping could be introduced and adapted to different application domains.

REFERENCES

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

[2] J. Gago, C. Douthe, R.E. Coopman, P.P. Gallego, M. Ribas-Carbo, J. Flexas, J. Escalona, and H. Medrano. UAVs challenge to assess water stress for sustainable agriculture. *Agricultural Water Management*, 153:9–19, 2015.

[3] H. Hoffmann, R. Jensen, A. Thomsen, H. Nieto, J. Rasmussen, and T. Friborg. Crop water stress maps for an entire growing season from visible and thermal UAV imagery. *Biogeosciences*, 13(24):6545–6563, 2016.

[4] J. Wu, D. Wang, and M.E. Bauer. Assessing broadband vegetation indices and quickbird data in estimating leaf area index of corn and potato canopies. *Field crops research*, 102(1):33–42, 2007.

[5] T. Van der Wal, B. Abma, A. Viguria, E. Prévinaire, P.J. Zarco-Tejada, P. Serruys, E. van Valkengoed, and P. van der Voet. Fieldcopter: unmanned aerial systems for crop monitoring services. In *Precision agriculture '13*, pages 169–175. Wageningen Academic Publishers, Wageningen, 2013.

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

[7] 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 Verlag, Berlin, Germany, 2015.

[8] B. Hrolenok, S. Luke, K. Sullivan, and C. Vo. Collaborative foraging using beacons. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 3-Volume 3*, pages 1197–1204. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[9] M. Paradzik and G. İnce. Multi-agent search strategy based on digital pheromones for uavs. In *Signal Processing and Communication Application Conference (SIU), 2016 24th*, pages 233–236. IEEE, 2016.

[10] M. Popovic, G. Hitz, J. Nieto, I. Sa, R. Siegwart, and E. Galceran. Online informative path planning for active classification using uavs. *arXiv preprint arXiv:1609.08446*, 2016.

[11] M. Yokoo. *Distributed constraint satisfaction: foundations of cooperation in multi-agent systems*. Springer Science & Business Media, 2012.

[12] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, IJCAI'05, pages 266–271. Morgan Kaufmann Publishers Inc., San Francisco, CA, 2005.

[13] R.G. Smith. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.

[14] D. Wicke, D. Freelan, and S. Luke. Bounty hunters and multiagent task allocation. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 387–394. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[15] E. Mathews, T. Graf, and K.S.S.B. Kulathunga. Biologically inspired swarm robotic network ensuring coverage and connectivity. In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, pages 84–90. IEEE, 2012.

[16] Z. Laouici, M.A. Mami, and M.F. Khelfi. Cooperative approach for an optimal area coverage and connectivity in multi-robot systems. In *Advanced Robotics (ICAR), 2015 International Conference on*, pages 176–181. IEEE, 2015.

[17] A. Stevens and H. Othmer. Aggregation, Blowup, and Collapse: The ABC's of Taxis in Reinforced Random Walks. *Siam Journal on Applied Mathematics*, 57(4):1044–1081, 1997.

[18] E.A. Codling, M.J. Plank, and S. Benhamou. Random walk models in biology. *Journal of The Royal Society Interface*, 5(25):813–834, 2008.

[19] C. Dimidov, G. Oriolo, and V. Trianni. Random Walks in Swarm Robotics: An Experiment with Kilobots. In M. Dorigo, M. Birattari, X. Li, M. López-Ibáñez, K. Ohkura, C. Pinciroli, and T. Stützle, editors, *Proceedings of the 10th International Conference on Swarm Intelligence (ANTS 2016)*, pages 185–196. Springer International Publishing, 2016.

[20] D. Payton, M. Daily, R. Estowski, M. Howard, and C. Lee. Pheromone Robotics. *Autonomous Robots*, 11(3), 2001.

[21] J.P. Hecker and M.E. Moses. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence*, 9(1):1–28, 2015.

[22] J. Snape, J. van den Berg, S.J. Guy, and D. Manocha. The Hybrid Reciprocal Velocity Obstacle. *IEEE Transactions on Robotics*, 27(4):696–706, 2011.

[23] V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of Consensus in a Multi-Robot Network: From Abstract Models to Empirical Validation. *IEEE Robotics and Automation Letters*, 1(1):348–353, 2016.

[24] X. Liu, Z. Li, P. Yang, and Y. Dong. Information-centric mobile ad hoc networks and content routing: A survey. *Ad Hoc Networks*, 58:255–268, 2017.

[25] D. Cokuslu and K. Erciyes. A flooding based routing algorithm for mobile ad hoc networks. In *2008 IEEE 16th Signal Processing, Communication and Applications Conference*, pages 1–5. IEEE, 2008.

[26] Y. Lu, B. Zhou, L.-C. Tung, M. Gerla, A. Ramesh, and L. Nagaraja. Energy-efficient content retrieval in mobile cloud. In *Proceedings of the second ACM SIGCOMM workshop on Mobile cloud computing*, pages 21–26. ACM, 2013.

[27] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan. Mason: A multiagent simulation environment. *Simulation*, 81(7):517–527, 2005.

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

[29] S.A. Sadat, J. Wawerla, and R.T. Vaughan. Fractal trajectories for online non-uniform aerial coverage. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2971–2976. IEEE, 2015.