

Monitoring and mapping of crop fields with UAV swarms based on information gain

Carlos Carbone¹, Dario Albani^{1,2,7}, Federico Magistri³, Dimitri Ognibene^{5,6},
Cyrill Stachniss³, Gert Kootstra⁴, Daniele Nardi¹, and Vito Trianni²

¹ DIAG, Sapienza University of Rome, Italy
`carbone@uniroma1.it`

² ISTC, National Research Council, Rome, Italy

³ University of Bonn, Germany

⁴ Wageningen University and Research, The Netherlands

⁵ University of Milano-Bicocca, Milan, Italy

⁶ University of Essex, Colchester, UK

⁷ Technology Innovation Institute, Abu Dhabi, UAE

Abstract. Monitoring of crop fields to map features like weeds can be efficiently performed with unmanned aerial vehicles (UAVs) that, owing to their privileged perspective and motion speed, can cover large areas in a short time. However, the need for high-resolution images for precise classification of features (e.g., detecting even the smallest weeds in the field) contrasts with the limited payload and flight time that characterises current UAVs, and requires several flights to uniformly cover a large field. However, the assumption that the whole field must be observed with the same precision is unnecessary when features are heterogeneously distributed, like weeds appearing in patches over the field. In this case, an adaptive approach that focuses only on relevant areas can perform better, especially when multiple UAVs are employed at the same time. Leveraging on a swarm-robotics approach, we propose a monitoring and mapping strategy that adaptively chooses the target areas based on the expected information gain, which measures the potential for uncertainty reduction due to further observations. The proposed strategy scales well with group size and overall leads to smaller mapping errors than optimal pre-planned monitoring approaches.

Keywords: swarm robotics, precision farming, information gain, UAV

1 Introduction

Precision farming requires high-quality data from the field in order to support operational and strategic decisions [11]. In this respect, unmanned aerial vehicles (UAVs) present a very flexible tool for remote sensing, as they can be deployed on demand and can quickly monitor large areas [26]. However, the platform limitations in terms of payload and flight time are often constraining, limiting the resolution of the acquired images (e.g., in terms of ground-sample distance) and requiring multiple flights to cover extensive fields. At the same time, the

heterogeneity of agricultural fields often requires high-resolution data only in certain areas where relevant features are present, while less interesting areas could be monitored on coarser resolution. This is often the case, for instance, for weed-management practices, in which the density of weeds is not uniform across the field, as weeds appear in patches. In order to support efficient management of the field, a larger effort should be dedicated only to those areas where relevant features are actually present. In this way, the limited energy budget of each UAV is properly allocated avoiding to dissipate it by flying over uninteresting areas. Given that the feature distribution over the field is a priori unknown, the monitoring and mapping strategies should be adaptive, responding to the observed field features in order to determine the best flight pattern, therefore ruling out predefined mission plans that do not take into account the actual distribution of relevant features over the field [8, 5].

UAV swarms have been proposed to address the adaptive monitoring and mapping of extensive fields. A swarm-robotics approach can indeed improve efficiency thanks to parallel monitoring of the field by different UAVs, reducing the operation to a fraction of the time required by a single-UAV approach [10, 3, 1]. Additionally, accuracy can be improved thanks to collaboration among UAVs in their feature-detection task [17]. Finally, self-organised deployment strategies can be envisaged to leverage the ability of UAVs to estimate which region within the field is of greater interest (e.g., estimating weed density from high-altitude flights [15]) and perform accurate monitoring and mapping only where relevant (e.g., by collecting high-resolution images while flying at low altitude [2]). With such an approach, UAVs can autonomously decide to enter or leave a given region on the basis of estimates of the completion level of the monitoring activity and energy constraints. As a consequence, the monitoring and mapping strategy of a single region needs to be flexible and scalable to adapt to changes in the actual number of UAVs that are concurrently operating.

Improving over previous work [3, 1], we propose a fully decentralised strategy to monitor a region or field based on reinforced random walks (RRW [25]), which maximises the monitoring effort only on areas that likely provide relevant information. To this end, we exploit the Information Gain (IG), an information-theoretic measure of the expected reduction in uncertainty from additional observations of a specific area. The usage of information theory for exploration and mapping has been demonstrated across application domains [6, 16, 22, 21, 20], and precision agriculture in particular [14, 24]. Here, we exploit IG to support exploration and coordination among robots. To this end, we build a model of the weed-density uncertainty in a given area that accounts for detection errors of a convolutional neural network (CNN) trained on a real dataset. To enable real-time on-board execution, we reduce the complexity of the neural classifier, and compensate for increased error by allowing multiple observations of the same area of the field [17]. On such basis, we compute the expected reduction in uncertainty as the IG from repeated independent observations made by UAVs within the swarm on a same area. Then, we exploit IG both to prioritize areas of interest to be observed next, and to determine if these areas are likely to be targeted

by other UAVs within the swarm. These two aspects are combined to guide the random selection of the next location to visit. Our results with multi-UAV simulations show that the swarm manages to quickly monitor those areas of the field that require more attention, minimising the observation error faster than approaches based on potential fields [3, 1], and better than a baseline approach based on a predefined flight plan that uniformly covers the entire field.

2 Problem description

We consider a field-monitoring and weed-mapping problem in which areas of high weed infestation need to be identified, creating a prescription map that can be exploited for weed control (e.g., variable-rate application of herbicides). Specifically, we focus on the identification of volunteer potatoes that infest sugar-beet fields—a common benchmark for precision agriculture [19, 15]—and use automatic object classification to inform the monitoring and mapping strategy. Our goal is to deploy a swarm of UAVs that can minimize the error in detection of weeds within the field. To this end, we exploit a simulated scenario to evaluate the effectiveness and scalability of the proposed strategy.

2.1 World model and UAV swarm simulation

We consider the field as divided in small areas forming a grid of cells, and each cell is fully contained within the camera field-of-view of a UAV hovering over its centre. Without loss of generality, we consider here a square field divided in a grid of $C \times C$ square cells, each with side l_c . Each UAV travels at a cruise speed of $v = 0.1l_c$ m/s at an altitude h sufficient to observe the whole cell given the camera footprint (e.g., $h \geq l_c/2$ if the camera aperture is $\frac{\pi}{2}$). Whenever moving over a cell, a UAV takes a RGB image used for crop/weed classification, leading to an estimation of the number of weeds present in the cell (see Section 2.2). After each observation, the UAV updates its local world representation, that is, a $C \times C$ map of the field (see Section 2.3). Additionally, UAVs can communicate with each other by broadcasting short messages exploiting a radio link, with range Rl_c . When the communication range is sufficiently large, any UAV can receive the information shared by any other UAV. Otherwise, UAVs apply a simple re-broadcasting protocol to maximise the reach of information shared within the swarm: upon reception of a message, a UAV re-broadcasts the message once and then puts it in a blacklist, hence avoiding to overload the communication channel. UAVs exploit communication to share information about the observations made on visited cells, and to also share their absolute position—available from some GNSS positioning system—hence allowing collision avoidance (here implemented using ORCA [4]) as well as collaboration for monitoring and mapping.

While crops are uniformly distributed over the field, weeds mostly appear in patches. We consider here C_p patches, each extending on a square of $n_p \times n_p$ cells, more densely distributed in the center than in the periphery following a Gaussian distribution. Some isolated weeds are also present within C_i additional

cells. Overall, the number of cells with some weed is $C_p n_p^2 + C_i$. Each cell can contain at most N_W weeds, i.e., the maximum value observed over the field.

2.2 Model of weed-classification uncertainty

To build a model of the weed-classification uncertainty, we consider a CNN to detect individual plants and to label them either as crops or weeds. In other words, given an input image, the CNN returns a list of bounding boxes that enclose the detected plants, together with the identified class of the plant. We use the state-of-the-art framework *Faster R-CNN* [23], but we limit its computational requirements by employing a shallow backbone (instead of the deep backbones normally used, such as ResNet-101) to match the constraints imposed by the limited power available on UAVs (see Figure 1a). The input image is passed through a convolutional block followed by 4 residual blocks [9]. The convolutional block is composed of a 7×7 convolutional layer followed by a 3×3 convolutional layer (represented in light blue in Figure 1a), each layer having 64 filters with stride 2 and ReLU activation functions. Each residual block is composed of a residual connection and by two 3×3 convolutional layers with ReLU activation functions (represented in orange in Figure 1a). The first residual block has 64 filters. Afterwards, each block doubles the previous number of filters. The head of the network is the same as in the original work, predicting the location, size and class of object detections. We exploit a dataset collected at Wageningen University, composed of 500 images—400 used as the training set and 100 as the test set—labelled to represent sugar beets in the crop class and volunteer potatoes in the weed class (see Figure 1b). We train our network for 10,000 epochs, using the ADAM optimizer [12] with a learning rate of 0.01. We evaluate our CNN using the average precision (AP) metric as defined in the MS-COCO challenge [13] obtaining an AP score of 89.6 for the crop and of 56.1 for the weed. While the AP for the crop class is in line with the literature [7, 18], the lower accuracy for the weed class is mainly due to miss-classifications of plants as background (Figure 1b).

This model is exploited during a simulation to both obtain a realistic observation—in terms of detected weeds—every time a UAV processes the information of a cell, as well as to update the internal knowledge about the world, as discussed in the following. Starting from the classification output of the trained CNN, our model associates a probability distribution to all possible observations given the actual number of weeds present in a cell. Considering that each cell can contain a discrete and small number of weeds (from the available data, we estimated $N_W = 12$ weeds for $l_c = 4$ m), we build a table that associates the actual number of weeds w with each possible observed number o , and gives for each combination a probability of occurrence $P(o|w)$. This table is estimated from the available data exploiting the trained CNN. Specifically, we compare the number of detected weeds against the true number of weeds. The element $T(o, w) = P(o|w)$ of the table is the relative frequency of detecting o weeds when the true number is w . Having several false positive cases, we also consider into the table the case of detecting more weeds than the given maximum N_W , that

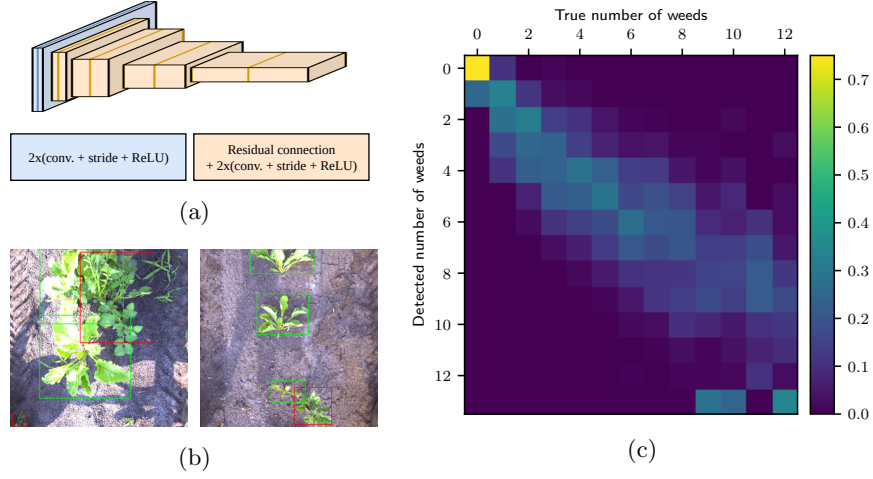


Fig. 1: (a) The Faster R-CNN framework with a shallow backbone was used while the head of the CNN remained the same as in the original work. (b) a few samples of CNN classification of volunteer potatoes (red) and sugar beets (green). (c) Sensor model derived from the CNN detection performance.

is, $o \geq N_W + 1$, N_W^+ represents the case in which the number of weeds detected is larger than N_W . The model obtained from the data is presented in Figure 1c: the CNN results in both false positives and false negatives, with errors becoming more frequent when the number of weeds is larger, as one would expect.

2.3 Uncertainty reduction from multiple observations

For each cell c , each UAV maintains a knowledge vector $p_c = [p_c(0), \dots, p_c(N_W)]$ that represents the expected probability distribution of all possible number of weeds considered. Taking a conservative approach, when no observation for a cell has been performed and no prior knowledge is available, the probability distribution is assumed to be uniform, with every element $p_c(w) = \frac{1}{N_W+1}$. Whenever a UAV observes cell c at time t , it detects a number of weeds o_c^t that depends on the actual value of weeds w_c present in c . Following the observation o_c^t , the probability $P_c(w|o_c^t)$ represents the updated knowledge vector for each possible value w . This can be easily computed exploiting the the Bayes theorem:

$$\forall w, p_c(w) \leftarrow P_c(w|o_c^t) = \frac{p_c(w)P(o_c^t|w)}{P(o_c^t)} = \frac{p_c(w)T(o_c^t, w)}{\sum_{j=0}^{N_W} p_c(j)T(o_c^t, j)} \quad (1)$$

The residual uncertainty about the weeds present in the cell corresponds to the information entropy of the knowledge vector:

$$H_c(W) = \sum_{w=0}^{N_W} -p_c(w) \log(p_c(w)) \quad (2)$$

The residual uncertainty is updated at every additional observation. We consider that sufficient observations have been performed for a given cell c when the entropy decreases past a low threshold that we heuristically set at $\hat{H} = -\frac{N_W-1}{N_W} \log \frac{N_W-1}{N_W} - \frac{1}{N_W} \log \frac{1}{N_W}$, which corresponds to the case in which there are only two remaining alternatives, one of which is much more likely than the other. When the threshold is reached, cells are marked as mapped and are not considered anymore for further observations. Note that observations can also be shared by other UAVs, allowing to update the residual uncertainty about a cell also when others have visited it. In this way, UAVs try to maintain aligned their local representation of the field.

The residual uncertainty of a cell is also the first component for the calculation of the IG of a cell, which represents the expected reduction in entropy from any possible additional observation o . A UAV can compute the IG of a cell c as follows:

$$\text{IG}_c(W) = H_c(W) - H_c(W|O) \quad (3)$$

where $H_c(W|O)$ is the conditional entropy of the same cell given that additional observations will be performed next:

$$H_c(W|O) = - \sum_{o=0}^{N_W^+} P_c(o) \sum_{w=0}^{N_W} [P_c(w|o) \log (P_c(w|o))] = \quad (4)$$

$$= - \sum_{o=0}^{N_W^+} \left[\sum_{j=0}^{N_W} T(o, j) p_c(j) \right] \sum_{w=0}^{N_W} [P_c(w|o) \log (P_c(w|o))] \quad (5)$$

Hence, on the basis of the available knowledge, each UAV can compute the IG for each cell of the field, and use it to quantify the information that could be gathered from a new observation, that is, the utility of visiting it.

3 Reinforced random walks for monitoring and mapping

A RRW is an exploration strategy that exploits available information to bias the random selection of the targets. The available information can derive from world knowledge (e.g., avoid to visit cells that are marked as mapped) or from other UAVs (e.g., others' location to avoid interference). A wise combination of both aspects can produce efficient monitoring and mapping strategies that focus on relevant areas, minimising the detection error.

3.1 Neighbourhood selection strategy

Whenever a UAV i has to determine the next cell to visit, it considers only those in its neighbourhood \mathcal{N}_i as target, which proved to be the best strategy to avoid long relocations that may be costly [3, 1]. Additionally, a UAV tries to select cells that are *valid*, i.e., not yet marked as mapped and not currently targeted by other UAVs to avoid interference. We define the neighbourhood \mathcal{N}_i^d of UAV

i as the cells at Chebyshev distance d . A UAV first considers the neighbourhood \mathcal{N}_i^1 , and only if there are no valid cells available, it considers the neighbourhood \mathcal{N}_i^2 . If also this does not contain valid cells, the UAV selects a cell in \mathcal{N}_i^2 to move away from the current position, limiting the choice among those cells that are already marked as mapped. With this strategy, a set of neighbouring cells \mathcal{V}_i is selected for a decision to be taken.

IG-based RRW The selection of the cell to target is based on a utility measure that is computed starting from the IG. Given a cell c_k , each UAV i determines the $\text{IG}_{c_k}^i$ on the basis of the currently available knowledge—different among UAVs in case of constrained communication—and, independently from any other UAV, assigns a probability of selecting the cell proportional to the relative IG:

$$P_{i,c_k}^{\text{IG}} = \frac{\text{IG}_{c_k}^i(W)}{\sum_{z \in \mathcal{V}_i} \text{IG}_{c_z}^i(W)} \quad (6)$$

To take into account the presence of other UAVs that are concurrently monitoring the field, the utility of selecting a cell is computed from P_{i,c_k}^{IG} considering the likelihood that the cell is *not* chosen by other UAVs:

$$u_{i,c_k} = P_{i,c_k}^{\text{IG}} \prod_{j \neq i} [1 - P_{j,c_k}^{\text{IG}}] \quad (7)$$

Here, P_{j,c_k}^{IG} is computed considering the neighbourhood \mathcal{V}_j of UAV j , although such computation is performed exploiting i 's private knowledge. To limit the computational complexity of this calculation, the product is extended only to those UAVs j that could potentially target the cell c_k at the same time, that is, the UAVs within the neighbourhood $\mathcal{N}_{c_k}^d$, $d \leq 2$ from cell c_k . By reducing the likelihood of choosing cells that are in reach from other UAVs, the proposed RRW strategy implements an implicit coordination mechanism that allows UAVs to share the monitoring burden and efficiently map the relevant areas. Should UAVs decide to target the same cell, the ORCA method would prevent collisions and a timeout mechanism allows to resolve possible—albeit unlikely—deadlocks. Given eq. (7), we propose different heuristics to choose the next cell to visit:

G: a cell c_k is chosen greedily selecting the one with highest utility u_{i,c_k} . In case of cells with identical utility, one is chosen at random.

R: a cell c_k is randomly chosen proportionally to the utility u_{i,c_k} .

S_γ: a cell c_k is chosen according to a softmax function of u_{i,c_k} with base e^γ .

3.2 Baseline strategy based on optimal pre-planned trajectories

As a baseline to confront the proposed strategy, we assume here an optimal reference point \mathbf{B} based on the uniform coverage of the field, i.e., N UAVs visit each cell of the field a fixed number of times. This optimal benchmark could be approximated by pre-planning the trajectories of all UAVs, although in reality a number of factors reduce the ideal performance proposed here. Considering the

speed of the UAV and the distance between cells, the time for a single UAV to fully cover the field is given by $\mathcal{T}_1 = C^2 l_c / v = 10C^2$. We assume that N UAVs can optimally partition the field, hence $\mathcal{T}_N = \mathcal{T}_1 / N$. To allow for independent observations of the field by different UAVs, we consider repeated passages that therefore entail longer times, that is, $M\mathcal{T}_N$ for M independent passages. After each observation, the residual uncertainty of a cell c is computed following eq. (2).

3.3 Baseline strategy based on potential fields

In addition, we consider a second baseline **B_{PF}** that exploits a RRW based on potential fields (PF). Here, the target cell is selected by UAV i according to a directional bias given by an attraction vector \vec{a}_i toward areas where weed was detected, and a repulsion vector \vec{r}_i from other agents to avoid overcrowding. This strategy is adapted from previous versions [3, 1] to reduce its computational demands, limiting the information exploited for the computation of potential fields. Attraction and repulsion vectors are computed as follows:

$$\vec{r}_i = \sum_{j \neq i} S(\vec{x}_i - \vec{x}_j, \sigma_r), \quad \vec{a}_i = \sum_c \frac{\hat{w}_c}{N_W} S(\vec{x}_c - \vec{x}_i, \sigma_a), \quad S(\vec{v}, \sigma) = 2e^{i \angle \vec{v}} e^{-\frac{|\vec{v}|}{2\sigma^2}}, \quad (8)$$

where \vec{x} represents the position of an agent/cell, and $S(\vec{v}, \sigma)$ returns a vector in the direction of \vec{v} with a Gaussian length with spread σ . With respect to [1], we reduce the number of agents considered to those belonging to the neighbourhood \mathcal{N}_i^4 , which are those ultimately considered also for the IG-based RRW. Additionally, we consider attraction to the cells $c \in \mathcal{N}_i^d$, $d \leq 2$ but discounting the force by the number \hat{w}_c of weeds detected in the last observation. Cells are not considered for attraction when they are marked as mapped. The selection of the next cell to visit by UAV i is performed randomly using the vector $\vec{v}_i = \vec{r}_i + \vec{a}_i$ as a bias. Specifically, the cell $c_k \in \mathcal{V}_i$ is selected randomly proportionally to the utility u_{i,c_k} , which is computed according to the angular difference $\theta_{i,c_k} = \angle(\vec{x}_{c_k} - \vec{v}_i)$:

$$u_{i,c_k} = C\left(\theta_{i,c_k}, 0.9 \left(1 - e^{\beta|\vec{v}_i|}\right)\right), \quad C(\theta, p) = \frac{1}{2\pi} \frac{1 - p^2}{1 + p^2 - 2p \cos \theta} \quad (9)$$

where $C(\cdot)$ is the wrapped Cauchy density function with persistence p . In this way, the length of the vector—smoothed through an exponential ceiling—determines the relevance accorded to the bias: the smaller the module, the lower the effect of the directional bias in the cell choice. This strategy has several parameters—namely the Gaussian parameters σ_a and σ_r , and the exponential constant β —which depend on the UAV swarm size N and need to be carefully tuned for appropriate performance [3, 1].

4 Experimental Results

We performed several simulations with a square field of $C \times C = 2500$ cells, having $C_q = 4$ square weed patches of $n_p \times n_p = 49$ cells randomly positioned

within the field, plus additional $C_i = 40$ isolated cells randomly scattered. In total, less than 10% of the field presents cells containing weeds. As already mentioned, we consider $N_W = 12$. At initialisation, UAVs start at random positions within the field, with no prior knowledge of the weed distribution.¹ We focus on the ability to monitor and map the field minimising the mapping error. We perform 50 runs for each experimental condition obtained varying swarm size $N \in \{10, 20, 30, 40, 50\}$, the communication range $R \in \{10, \infty\}$, the heuristics exploited for the information gain (\mathbf{G} , \mathbf{P} and \mathbf{S}_γ with $\gamma \in \{1, 5\}$) and the parameters $\sigma_r, \sigma_a \in \{0, 2, 4, 8, 16, 32\}, l_c$ for the PF-based RRW with $\beta = 1$. With ideal communication ($R = \infty$), each broadcasted message reaches every other UAV in the swarm.

To evaluate the system performance, we consider an aggregated world map resulting from the individual UAV maps by considering for each cell c and UAV i the knowledge vector $p_c = p_{c,i}$ with lowest uncertainty $H_i(c)$. Then, the value $\tilde{w}_c = \arg \max_w p_c$ is considered as the mapped number of weeds in cell c , and the mean squared error (MSE) is computed with respect of the real value w_c for the whole field. Figure 2 shows how this error decreases as further observations are gathered from the field, taking as temporal reference the time $M\mathcal{T}_N$ necessary to gather M independent observations for each cell with N UAVs following the baseline \mathbf{B} . When the communication is perfect, all UAVs share the same map and any new observation contributes in reducing the uncertainty. The mapping error is initially rather high, but decreases as UAVs discover and focus on the areas of interest. Notably, all the IG-based strategies scale well with respect to the swarm size N , and present the best performance when the greedy approach \mathbf{G} is employed. In this case, the MSE gets better than the one of the baseline \mathbf{B} for $M \approx 3$ —when the error is about 0.15—and continues to decrease of another order of magnitude with further observations. The baseline \mathbf{B}_{PF} is instead less efficient in reducing the mapping error, and outperforms \mathbf{B} only when $M > 6$, hence requiring much longer than the best IG-based strategy \mathbf{G} . The other heuristics perform slightly worse, meaning that additional randomness in the selection of the next cell to visit does not increase performance. Specifically, the \mathbf{R} and the \mathbf{S}_5 strategies have very similar profiles, and the worst performance is obtained with the \mathbf{S}_1 strategy.

The difference in performance among the proposed heuristics is visible also for the total completion times, shown in Figure 3. Here, we measure the time \mathcal{T}_C necessary to fully cover the entire field, by passing over every cell at least once, scaled with respect to the time \mathcal{T}_N of the baseline \mathbf{B} . The greedy strategy \mathbf{G} performs better than any other strategy, with a coverage time that scales perfectly with group size and values that roughly correspond to the time required to have an MSE smaller than \mathbf{B} . The other strategies have generally longer coverage times, with slight improvements for larger group size N although never being in par with \mathbf{G} .

¹ We ignore here the initial relocation from a deployment station, and also disregard the need to return to a predefined location.

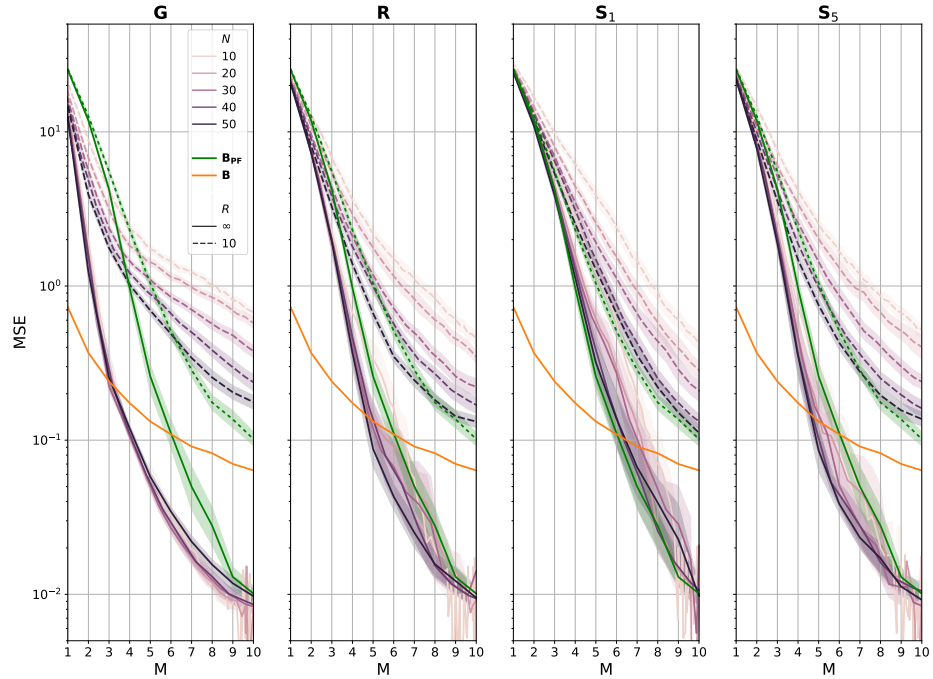


Fig. 2: MSE of the maps generated with different strategies (average and standard deviation), with respect to the optimal pre-planned monitoring baseline \mathbf{B} and the PF-based RRW strategy \mathbf{B}_{PF} . Note that all plots are rescaled in time with respect to the time \mathcal{T}_N required by N agents to fully cover the field once. As a consequence, the average MSEs from the baseline \mathbf{B} for all values of N coincide. The baseline \mathbf{B}_{PF} is shown for both $R = \infty$ and $R = 10$, and corresponds to the case with $\beta = 1$, $\sigma_a = 2$, $\sigma_r = 8$ and $N = 50$, which is the one with lowest MSE among all the tested parameters.

The adaptive approach manages to reduce the uncertainty below the optimal pre-planned strategy by mainly focusing on the areas where weed is present, avoiding to monitor those areas that are devoid of weeds. Indeed, Figure 4 shows the Pearson’s correlation coefficient between the number of weeds present in a cell, and the number of independent observations gathered for it. It is possible to note that, as time goes by, the relationship between these two variables builds stronger, meaning that more time is spent over areas that require more observations to substantially reduce the uncertainty, while areas without weed are quickly abandoned.

The results obtained for a limited communication range tell a different story, however. First and foremost, the swarm does not reduce the error substantially below the baseline \mathbf{B} (see dotted lines in Figure 2). The monitoring and mapping performance is substantially worse, due to an inefficient reduction of uncertainty as the observations from other UAVs are only shared locally, and there is no




Fig. 3: Coverage time \mathcal{T}_C relative to \mathcal{T}_N for the different IG-based strategies, for $R = \infty$ (left) and $R = 10$ (right).

mechanism to ensure that UAVs maintain the local maps aligned. The error gets lower as the group size increases, as an effect of the larger diffusion of messages thanks to the re-broadcasting communication protocol. Also the coverage time is positively affected by group size, as shown in Figure 3. Nevertheless, the density of agents is not sufficiently high to ensure that the communication topology is always connected, leading to a partial loss of communications. As a consequence, different UAVs may map areas already sufficiently observed by others, without any improvement for the collective. This leads also to a reduced correlation between the number of weeds in a cell and number of observations, especially for low values of N as shown in Figure 4.

5 Conclusions

This paper has demonstrated that an improvement in performance is expected when adaptive approaches are employed, leading to a substantial reduction of uncertainty below what can be achieved through a blind acquisition of information by means of pre-planned trajectories. By investing little more exploration time, the mapping error is substantially reduced despite the low accuracy in the classification output of single images. Additionally, the best heuristic based on IG is rigorous and bears no free parameters, being therefore ideal for deployment in a swarm robotics system notwithstanding the group size or the field dimensions, as no tuning is needed.

One observed limitation concerns the loss in performance for constrained communication, which jeopardises the benefits from an adaptive approach with a UAV swarm. This aspect is worth studying in detail, although it does not invalidate the results presented for the ideal communication case. In several scenarios—and in precision agriculture in particular—one can reasonably assume that a good-enough communication channel can be set up in place, thanks to long range radio communication or to the upcoming 5G technology. The limited communication scenario may be addressed by either changing the communication




Fig. 4: Correlation between the number of weeds present and the number of observations performed, plotted over time for the different strategies.

protocol in a way to ensure better alignment among UAVs of the local world representations (e.g., by sharing not just the latest observation, but the entire word representation), or by explicitly maintaining a high degree of connectivity within the swarm.

Another limitation of the proposed approach stands in the assumption of independence between observations, which likely does not hold in real-world scenarios, especially for slowly-changing environmental conditions as in crop fields. Indeed, multiple observations are likely to provide the same errors, if these are related to specific features of the observed area (e.g., a volunteer potato hiding between two sugar beets), especially if observations are made from the same perspective and within a short time-frame. To mitigate this issue and increase the chances that repeated observations lead to substantial uncertainty reduction, a possible approach is to vary the observation position relative to the point of interest, so that different perspectives can lead to more informative observations. Another complementary approach consists in the usage of already available knowledge as a prior for the classification, exploiting specially-trained CNNs [17]. The combination of these two possibilities can lead to substantial reduction in uncertainty that can be modelled and exploited with the approach proposed in

this paper. Future work will investigate this possibility while tackling real-world monitoring and mapping tasks.

Acknowledgements

This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy, EXC-2070 - 390732324 - PhenoRob. Vito Trianni acknowledges partial support from the project TAILOR (H2020-ICT-48 GA: 952215).

References

1. D. Albani, T. Manoni, A. Arik, D. Nardi, and V. Trianni. Field Coverage for Weed Mapping: Toward Experiments with a UAV Swarm. In *BICT 2019: Bio-inspired Information and Communication Technologies*, volume 289 of *LNICST*, pages 132–146. Springer, Cham, 2019.
2. D. Albani, T. Manoni, D. Nardi, and V. Trianni. Dynamic UAV Swarm Deployment for Non-Uniform Coverage. In *17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 523–531, 2018.
3. D. Albani, D. Nardi, and V. Trianni. Field coverage and weed mapping by UAV swarms. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4319–4325, 2017.
4. D. Bareiss and J. van den Berg. Generalized reciprocal collision avoidance. *The International Journal of Robotics Research*, 34(12):1501–1514, 2015.
5. T. M. Cabreira, L. B. Brisolará, and P. R. Ferreira Jr. Survey on Coverage Path Planning with Unmanned Aerial Vehicles. *Drones*, 3(1):1–38, 2019.
6. J. Das, F. Py, J. B. Harvey, J. P. Ryan, A. Gellene, R. Graham, D. A. Caron, K. Rajan, and G. S. Sukhatme. Data-driven robotic sampling for marine ecosystem monitoring. *The International Journal of Robotics Research*, 34(12):1435–1452, 2015.
7. M. Fawakherji, A. Youssef, D. Bloisi, A. Pretto, and D. Nardi. Crop and weeds classification for precision agriculture using context-independent pixel-wise segmentation. In *2019 Third IEEE International Conference on Robotic Computing (IRC)*, pages 146–152. IEEE, 2019.
8. E. Galceran and M. Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013.
9. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition (CVPR)*, pages 770–778, 2016.
10. A. C. Kapoutsis, S. A. Chatzichristofis, and E. B. Kosmatopoulos. DARP: Divide Areas Algorithm for Optimal Multi-Robot Coverage Path Planning. *Journal of Intelligent & Robotic Systems*, 86(3-4):663–680, 2017.
11. A. King. Technology: The Future of Agriculture. *Nature*, 544(7651):21–23, 2017.
12. D. P. Kingma and J. Ba. ADAM: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
13. T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

14. J. Liu and R. K. Williams. Monitoring Over the Long Term: Intermittent Deployment and Sensing Strategies for Multi-Robot Teams. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA 2020)*, pages 7733–7739, 2020.
15. P. Lottes, R. Khanna, J. Pfeifer, R. Siegwart, and C. Stachniss. UAV-based crop and weed classification for smart farming. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3024–3031, 2017.
16. K.-C. Ma, Z. Ma, L. Liu, and G. S. Sukhatme. Multi-robot Informative and Adaptive Planning for Persistent Environmental Monitoring. In R. Groß, A. Kolling, S. Berman, E. Frazzoli, A. Martinoli, F. Matsuno, and M. Gauci, editors, *Distributed Autonomous Robotic Systems*, volume 6 of *Springer Proceedings in Advanced Robotics*, pages 285–298. Springer International Publishing, Cham, 2018.
17. F. Magistri, D. Nardi, and V. Trianni. Using prior information to improve crop/weed classification by mav swarms. In *11th International Micro Air Vehicle Competition and Conference*, pages 67–75, 2019.
18. A. Milioto, P. Lottes, and C. Stachniss. Real-time blob-wise sugar beets vs weeds classification for monitoring fields using convolutional neural networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4:41, 2017.
19. A. T. Nieuwenhuizen, J. W. Hofstee, J. C. v. d. Zande, J. Meuleman, and E. J. v. Henten. Classification of sugar beet and volunteer potato reflection spectra with a neural network and statistical discriminant analysis to select discriminative wavelengths. *Computers and Electronics in Agriculture*, 73(2):146–153, 2010.
20. D. Ognibene and Y. Demiris. Towards active event recognition. In *IJCAI*, volume 13, pages 2495–501, 2013.
21. D. Ognibene, L. Mirante, and L. Marchegiani. Proactive intention recognition for joint human-robot search and rescue missions through monte-carlo planning in pomdp environments. In *International Conference on Social Robotics*, pages 332–343. Springer, 2019.
22. E. Palazzolo and C. Stachniss. Effective exploration for mavs based on the expected information gain. *Drones*, 2(1):9, 2018.
23. S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
24. N. B. Rossello, R. F. Carpio, A. Gasparri, and E. Garone. Information-Driven Path Planning for UAV with Limited Autonomy in Large-scale Field Monitoring. *arXiv*, 2020.
25. P. E. Smouse, S. Focardi, P. R. Moorcroft, J. G. Kie, J. D. Forester, and J. M. Morales. Stochastic modelling of animal movement. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2201–2211, 2010.
26. 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.