Area Inspection by Robot Swarms through Exploitation of Information Gain

Carlos Carbone*, Dario Albani†*, Daniele Nardi*, Dimitri Ognibene‡ and Vito Trianni§

*La Sapienza University of Rome, Rome, Italy
Email: [surname]@diag.uniroma1.it

†Technology Innovation Institute, Abu Dhabi, United Arab Emirates
Email: dario.albani@tii.ae

‡University of Essex, Colchester, UK
Email: dimitri.ognibene@essex.ac.uk

§National Research Council of Italy, Rome, Italy
Email: vito.trianni@istc.cnr.it

***Abstract*—We propose a decentralized, collaborative approach for area coverage and mapping by means of a swarm of robots. The approach is hinged on Information Theory, and builds over a Reinforced Random Walk (RRW) specifically tailored for a precision agriculture scenario, but general enough to accommodate different applications. Here, we improve by considering the estimated uncertainty about the features present in a target area, and by the expected reduction in uncertainty that visiting the target area could provide, that is, the information entropy and information gain, respectively. The latter is exploited to weight the random selection of the next area to explore, taking also into account the presence of nearby agents that could visit the same target area. The proposed approach features no configuration parameters related to the number of agents employed and the size of the field, opening to direct implementation without preliminary tuning and configuration steps.**

***Keywords*–Swarm Robotics; Entropy; Information Gain; Random Walk.**

I. INTRODUCTION

Many monitoring and mapping applications require to fully cover a wide region of space detecting the presence of points of interest and mapping their exact position. This is a common task, especially for precision agriculture, which has been approached in many different ways: pheromone-based approaches [1] [2], evolutionary path planning [3], and random walk based approach [4] [5] to cite some. In this research, we rely on Information Theory and propose a new algorithm for coverage and mapping of large areas. We consider precision agriculture as the target application, whereas a swarm of Unmanned Aerial Vehicles (UAVs) is required to detect the presence of weeds in the field, but the proposed algorithm is general enough to accommodate a variety of scenarios. The swarm inspection strategy is informed by a heuristic computed according to Information Theory concepts such as Entropy and Information Gain (IG). Moreover, we exploit the presence of multiple robots scattered throughout the field and propose a decentralized collaborative approach that improves accuracy and reduces the time needed for exploration. Indeed, by quantifying and including the knowledge of an agent about the area, it is possible to develop a cooperative behavior that focuses on points of interest and reduce the mapping time, i.e., the time needed to recognize all the relevant features within the field. In this work, we assume the following simplified world model. The work area is partitioned in a 4-connected grid which can be configured to represent spaces of different complexity. Agents are not limited to orthogonal motion and can move in continuous space. A grid cell $c_k \in \mathcal{C}$ represents a region of the field, where a robot can move, and might contain a certain number of points of interest (e.g., weeds). All robots are identical and each robot is identified by its unique id i and its position in the environment. The robots move at constant speed and are able to avoid collision thanks to an on-board collision avoidance algorithm [6]. Robots can communicate with each other by using broadcast communication that might be subject to range limitations.

II. INFORMATION GAIN FOR EXPLORATION

The swarm strategy aims at maximizing the expected information that could be gathered from an area after inspection: the IG, that is, the expected reduction in entropy. When used for exploration and mapping tasks, the IG can be used to quantify how much knowledge would be obtained if an observation in a certain location occurs [7]. To this end, we exploit the IG to quantify the information that could be gathered from a new observation performed by the agents in a specific cell c_k and to represent the utility of visiting it. In its simplest form, at a specific time instant, the agent i can compute the IG of a cell c_k according to the following equation:

$$\mathrm{IG}_i(c_k) = H_i(c_k) - H_i(c_k|o_k(i)) \tag{1}$$

where $H_i(c_k)$ express the residual uncertainty that robot i has about cell c_k, and $H_i(c_k|o_k(i))$ is the conditional entropy of the same cell given the observation o_k performed by robot i at a specific time instant. The residual uncertainty of a cell is computed as follows.

$$H_i(c_k) = -\sum p_i(c_k)log(p_i(c_k)), \tag{2}$$

with $p_i(c_k)$ representing the knowledge of robot i—i.e., the current knowledge about the number of points of interest existing in cell c_k, which are in a discrete number and are represented by a vector that associates to each value $c \in [0, C]$

 ISBN: 978-1-61208-787-0

the probability of having c points of interest. Lastly, replacing $o_k(i)$ with $\tilde{o}_k$, the conditional entropy is:

$$H_i(c_k|\tilde{o}_k) = -\sum_o p_i(\tilde{o}_k) \sum_c \left[p_i(c_k|\tilde{o}_k) \log\left(p_i(c_k|\tilde{o}_k)\right)\right]. \tag{3}$$

The knowledge vector is calculated by means of the probability of having a certain observation as $p_i(c_k|\tilde{o}_k) = p_i(c_k)p_i(\tilde{o}_k|c_k)/p_i(\tilde{o}_k)$, where $p_i(c_k|\tilde{o}_k)$ is the probability of robot i performing an observation o for the cell c_k. Thus, each new observation increases the confidence about the points of interest present in a cell. When the uncertainty decreases below a fixed threshold, the cell is considered as mapped and no further observation is required. Note that observations can also be shared by neighbouring robots, allowing to update the residual uncertainty about a cell also when others have visited it. In this way, the robots keep a local model of the entire area exploiting both own and others' observations.

III. INFORMATION THEORY ENRICHED RANDOM WALK

The IG of a cell can now be used as a proxy of the expected quantity of information gathered from a new observation of a cell. This is computed separately from each agent in the swarm and in a completely decentralized way since the units in the swarm only rely on their local knowledge for the computation of these values—e.g., the knowledge vector $p(c_k)$. Computed values are then used for next target selection. In particular, we propose a distance-aware collaborative strategy that assigns higher probabilities of inspection to closer cells—avoiding big jumps that proved to be detrimental for exploration [4]—and that weights the decisions according to other agents expected behaviors. We start by assigning probabilities to each cell thanks to (1):

$$P_i(c_k) = \frac{\mathrm{IG}_i(c_k)}{\sum_z \mathrm{IG}_i(c_z)} \tag{4}$$

where P_i is the probability of selecting cell c_k computed from the perspective of agent i with respect to all cells considered for inspection at this stage. Nonetheless, (4) alone is not enough since it does not consider the presence of other agents, hence, collaboration. To this end, we rewrite (1) has:

$$P_i(c_k) = \frac{\mathrm{IG}_i(c_k)}{\sum_z \mathrm{IG}_i(c_z)} \prod_{j \neq i} \left[1 - \frac{\mathrm{IG}_j(c_k)}{\sum_z \mathrm{IG}_j(c_z)}\right] \tag{5}$$

where, the probability $P_i(c_k)$ is now weighted by the probability that the cell c_k will be selected by inspection by another agent $j \neq i$. This allows for direct inclusion of other agents operations and in a completely decentralized way since all the probabilities in the right-most term are computed relying only on the local knowledge of agent i. Lastly, we introduce distance and upgrade (5) as:

$$P_i(c_k) = \frac{d_i(c_k)^{-1}\mathrm{IG}_i(c_k)}{\sum_z d_i(c_z)^{-1}\mathrm{IG}_i(c_z)} \prod_{j \neq i} \left[1 - \frac{d_j(c_k)^{-1}\mathrm{IG}_j(c_k)}{\sum_z d_j(c_z)^{-1}\mathrm{IG}_j(c_z)}\right] \tag{6}$$

with the terms $d_i(c_k)$ representing the euclidean distance between agent i and the cell c_k. From a computational point of view, (5) is expensive and does not scale well with hundreds of agents. To mitigate this issue, the choices are constrained to the local neighbourhood of the agent (see Figure 1). At first, the algorithm computes the probabilities only for immediate neighbors cells—i.e., those composing the 3×3 neighbourhood. If no valid cell is found—i.e., all cells are already mapped or targeted by other agents—the algorithms proceeds with the 5×5 neighbourhood. In case there is still no valid cell, a random choice is made among the cells of the outer border that are not targeted by other agents. Note that the set of agents taken into account in (6) is constrained to those agents that can potentially move to the target cell c_k, hence those that are within a 5×5 neighbourhood of c_k.

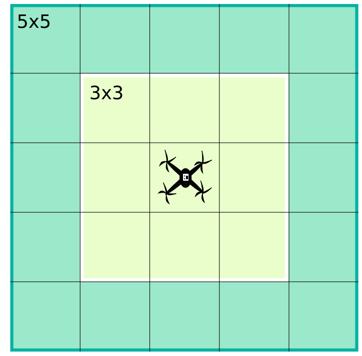

Figure 1. Graphical illustration of cells consideration for random selection based on IG.

IV. CONCLUSION AND FUTURE WORK

We presented an algorithm for area inspection that relies on IG to chose the next area to inspect. It does not present free parameters, exempting the user from pre-operational tuning and making it suitable for application such as search and rescue and precision agriculture, where the deployment speed is an important requirement. Next, it is completely decentralized, robust to failure and noisy communication since the cell selection procedure relies only on local knowledge. Nonetheless, the latter is built over information received during operation and local communications are required for good performance. To this end, we are currently working on introducing direct sharing of the IG through belief propagation. We believe that this would help reducing the overall communication overhead and will greatly boost the performances. Last, the algorithm scales up to swarm size of hundreds and, if limiting the agents considered in (6) to local neighbors, even more.

REFERENCES

[1] B. Hrolenok, S. Luke, K. Sullivan, and C. Vo, "Collaborative foraging using beacons." in AAMAS, vol. 10, 2010, pp. 1197–1204.

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

[3] 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.

[4] D. Albani, D. Nardi, and V. Trianni, "Field coverage and weed mapping by uav swarms," in Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on. IEEE, 2017, pp. 4319–4325.

[5] D. Albani, T. Manoni, A. Arik, D. Nardi, and V. Trianni, "Field coverage for weed mapping: Toward experiments with a uav swarm," in Bio-inspired Information and Communication Technologies, A. Compagnoni, W. Casey, Y. Cai, and B. Mishra, Eds. Cham: Springer International Publishing, 2019, pp. 132–146.

[6] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Optimal reciprocal collision avoidance for multi-agent navigation," in Proc. of the IEEE International Conference on Robotics and Automation, Anchorage (AK), USA, 2010.

[7] E. Palazzolo and C. Stachniss, "Effective exploration for mavs based on the expected information gain," Drones, vol. 2, no. 1, 2018, p. 9.

