

Coordinated Search with a Swarm of UAVs

Sonia Waharte and Niki Trigoni

University of Oxford

Computing Laboratory

Oxford, United Kingdom

Email: {sonia.waharte, niki.trigoni}@comlab.ox.ac.uk

Simon J. Julier

University College London

Dept. of Computer Science

London, United Kingdom

Email: S.Julier@cs.ucl.ac.uk

Abstract—Search is a fundamental task for Wilderness Search and Rescue that can greatly benefit from the use of a swarm of autonomous UAVs to survey the environment. The benefits are maximised if the UAVs coordinate their search activities with one another. In this poster, we present our preliminary work on developing coordination strategies for multiple UAVs. It is based on a distributed, grid-based probabilistic environmental model. We discuss the practicalities of the search task, present a simplified mathematical model of the environment and sensors, and present some preliminary simulation-based results. These clearly illustrate, even in a highly simplified case, the great benefits of coordinated search.

I. INTRODUCTION

Arguably the most important task in Wilderness Search and Rescue (WiSAR) is *search* — until a missing person has been found, they cannot be rescued or recovered. Key to the search task is the collection and comprehension of evidence [1]. The benefits of Unmanned Aerial Vehicles (UAVs) in evidence collection are well-established: they can rapidly acquire aerial imagery even in dangerous environments. Multiple UAVs can collect data from multiple vantage points simultaneously, greatly increasing these advantages.

Most systems are equipped with cameras and are controlled by a UAV operator (who flies the UAV) and a sensor operator (who controls the camera and interprets the data) [1]. When multiple UAVs are used, the complexity of coordination means that they are normally flown in a fixed formation relative to one another at a fixed altitude above the ground.

Although UAV systems have been successfully deployed, there are a number of difficulties with them. First, the use of pre-planned trajectories means that the search strategies are not necessarily optimal. Second, the sensor operators take the burden of observing, assessing and integrating all the information received from the sensors. This can be a difficult and error-prone process. Finally, the human resources needed to operate the UAVs can be prohibitively expensive.

One way to mitigate these difficulties is to *automate* the operations of the UAVs. If the UAVs can perform in-flight collaborative self-organisation, they can optimise their strategies to sense the environment in the most efficient manner possible. They can also be responsive to system and sensor failures. If they perform in-flight object detection, the sensor operator need only be informed of “critical” events.

In this poster, we discuss our preliminary work on developing coordinated search strategies. The next section describes the



Fig. 1. An AscTec Hummingbird quadrotor helicopter.

search task together with the platforms used. The mathematical model is described in Section III. The coordination strategies and data fusion algorithms are outlined in Section IV. Results are presented and discussed in Section V, and conclusions and future work in Section VI.

II. SEARCH TASK AND PLATFORMS USED

The initial search task we are conducting is to search for at most a single object of interest (or target) in a planar environment such as a field. Although this is a simplification of our final problem domain (which will be a hilly environment), it exposes many significant issues including problems of coordination, navigation and planning in uncertain environments.

A key requirement for our project is that the UAVs should be small, lightweight and manoeuvrable. Therefore, we have chosen to use the Ascending Technologies Hummingbird quadrotor helicopters (illustrated in Fig. 1) [2]. These platforms can carry up to 200g payload that can consist of additional sensing devices or processing boards. The UAVs have a flight time of 23 minutes without payload and 12 minutes with a payload of 200g. They are also highly stable in gusting wind. The UAVs are equipped with two classes of sensors: navigation sensors and surveillance sensors. The navigation sensors are used by the quadrotor’s autopilot and GPS waypoint following systems. Even without vision-based aiding, position can be measured with an accuracy of 1.5–2.5m. The surveillance sensors are used to detect the potential targets on the ground. We are currently using a single, downward-pointing camera. Two sample images, captured at different altitudes are shown in Fig. 2.

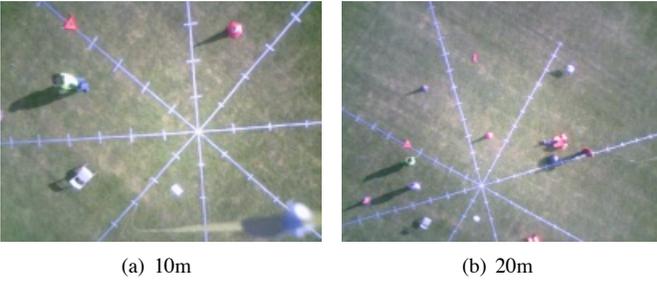


Fig. 2. Sample frames taken from a UAV. The large star-shape is a calibration pattern which is being used to groundtruth the pose of the UAV.

The UAVs will wirelessly communicate with one another via IEEE 802.11 wireless devices.

III. MATHEMATICAL DESCRIPTION OF SEARCH TASK

Following Chung [3], we model the problem as follows:

- 1) We are searching for a single, stationary target x_T that might lie in the search area \mathcal{A} .
- 2) \mathcal{A} is decomposed into a set of $|\mathcal{A}|$ grid cells. The target (if present) occupies at most a single cell.
- 3) The UAV is equipped with navigation sensors (so that it knows its position within the resolution of a cell) and surveillance sensors (which provide a detect / no detect event for the cell the UAV is flying over).
- 4) The surveillance sensor of the k^{th} UAV covers M_{h_k} cells, where h_k is the altitude of the UAV.
- 5) The control input is a waypoint which specifies which cell the UAV will fly to next.

We adopt a Bayesian approach to keep track of the target state probability density function. This approach is sensible in our context, where non-Gaussian sensor measurements are considered.

Each UAV maintains a grid-based probabilistic map (*belief map* or *occupancy grid*) composed of $|\mathcal{A}|$ cells. Let $x_T = a$ denote the event that the target lies in cell a . Each cell contains the probability that the target is present in that cell [3], [4].

Let H be a binary random variable representing the event that there is a target in the search area. The probability $Pr(H = 1) = \delta$ represents the prior belief of the target in the search area. The dimension of the sensing area increases with the height of the UAVs. UAVs at high altitude have a greater sensing coverage than UAVs at low altitude but they also have lower sensing resolution (smaller detection probability). This tradeoff will be investigated in future work in the evaluation and design of search strategies.

The eventual implementation will have to segment the sensing area and recognise individual objects. However, for our preliminary experiments we assume a simple detection model. Let d_a^t be the detection measurement for cell a at time t and $x_T = a$ represents the presence of the target in the a^{th} cell. We can formally express our sensing model at height h by:

$$\bullet Pr_h(d_a^t = 1 | x_T = a) = 1 - \beta_h$$

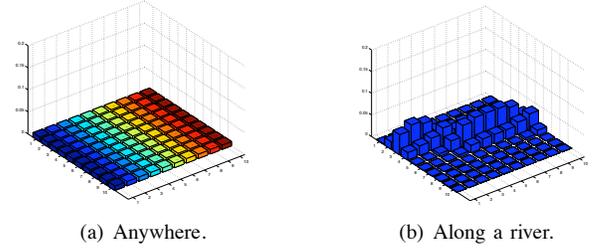


Fig. 3. Different priors on target location.

- $Pr_h(d_a^t = 0 | x_T = a) = \beta_h$
- $Pr_h(d_a^t = 0 | x_T \neq a) = 1 - \alpha_h$
- $Pr_h(d_a^t = 1 | x_T \neq a) = \alpha_h$

α_h ($0 \leq \alpha_h \leq 1$) and β_h ($0 \leq \beta_h \leq 1$) represent the false alarm and missed detection probabilities at height h . As Fig. 2 suggests, these detection probabilities will be strongly dependent upon the altitude of the UAV. Developing an empirical error model is a work in progress. However, here we assume that both the false alarm and missed detection probabilities increase with altitude.

IV. MAPS AND DATA FUSION

The search problem starts with the assumption of a prior probability distribution function that describes the initial belief of the target location. This can be a Gaussian distribution or a coarse estimate of the target location depending on environmental features such as river or mountains. If no prior information is known, we assume a uniform distribution (Fig. 3). After each observation, the probability distribution function of the target state is recomputed. We assume that observations are independent and we consider cells inter-dependence in a single stationary target scenario. The cell update mechanism is described below for different cases.

A. Grid cell independence assumption

A commonly adopted approach is to assume that the grid cells are independent of one another and the belief map is updated only for the cell in which the observation occurs. This assumption can be appropriate in scenarios where the number of targets is unknown, the targets are moving or the uncertainty about the target location evolves over time due to external factors (e.g. addition/removal of landmarks). Let D^t be the set of measurements up to time t and $Pr(x_T = a)$ the probability that a target is located in cell a . Given a cell a and a measurement in a , the probability that the target is located in a is updated as follows:

$$Pr(x_T = a | D^t) = \frac{Pr(d_a^t | x_T = a, D^{t-1}) Pr(x_T = a | D^{t-1})}{Pr(d_a^t | D^{t-1})}$$

B. Grid cells dependence assumption

If we assume the existence of only a single stationary target, an observation in any given cell also affects the probability that the target is located in other cells. To account for this dependence, we update all the cells in the belief map after

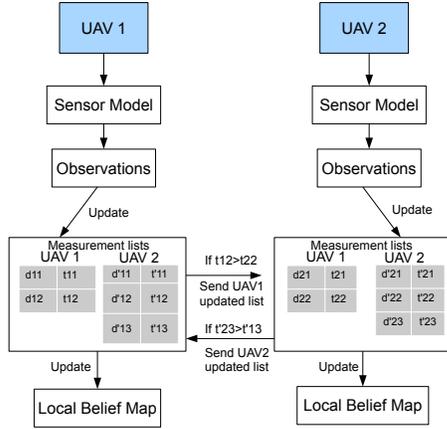


Fig. 4. Illustration of belief maps update mechanism with 2 UAVs. Each UAV maintains measurement lists for all UAVs. When 2 UAVs come within communication range, they update their measurement lists.

each observation in a similar way as described in [3] and [5]. In addition, we also account for the fact that for each measurement, the k^{th} UAV can observe a set of $M_{h,k}$ cells simultaneously, where h_k is the altitude of the UAV. Let $D(t) = \{d_1^t, \dots, d_{M_{h,k}}^t\}$ be the observations at time t . $D(1:t)$ represents the history of all the observations from time 1 to t . With assumption of independence in observations, we have:

$$Pr(x_T = a | D(1:t)) = \frac{\prod_{i=1}^{M_{h,k}} Pr(d_i^t | x_T = a) Pr(x_T = a | D(1:t-1))}{Pr(D(t) | D(1:t-1))}$$

C. Distributed Data Fusion

Since the UAVs have limited communication range, they can only exchange data when they are within communication range. Each UAV therefore maintains a local belief map that can often differ from other UAVs'. Each UAV progressively updates its belief map based on the measurement lists (timestamped) it receives from other UAVs (Fig. 4). A UAV always maintains an up-to-date list of its own measurements. When two UAVs come within communication range, they synchronize their measurement lists and recompute their local belief map. Note that these belief maps might still differ from the rest of the UAVs since only the local measurement lists of the UAVs are guaranteed to be up-to-date. The measurement lists they maintain for the rest of the UAVs can be stale.

V. SIMULATION RESULTS

We performed some simulations using MATLAB/Simulink (with real-time control systems [6]) with two UAVs and a single stationary target. The UAVs evolve in a 10×10 grid. We do not assume any prior knowledge on the target location. The initial probability of the presence of a target in the search area is set to 0.5. Some preliminary results are shown in Fig. 5. The figure shows the evolution of the belief probability of the presence of a target in the search area over time. A UAV decides on the next cell to visit using a steepest gradient

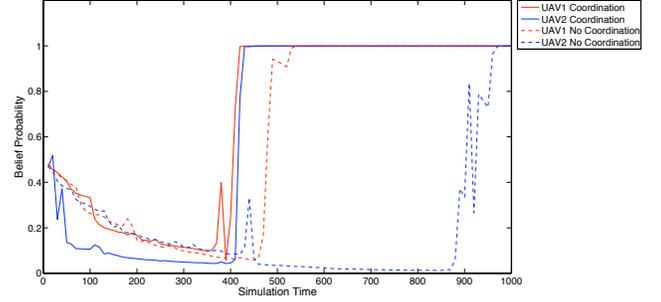


Fig. 5. Performance in simulation trials with 2 UAVs in cooperative and non-cooperative modes. UAVs in cooperative modes converge faster to a decision (on the presence of a target) than UAVs in non-cooperative mode.

ascent method applied to its belief map. Cooperative and non-cooperative strategies have been considered. The probabilities of missed detection and false alarm have been set to $1/4$ and $1/5$ respectively. As we can see and as expected, sharing measurements among UAVs (cooperative strategy) leads to super-linear speed up in the decision time. Considering that UAVs can be deployed at different heights consequently covering areas of different sizes with different sensing accuracies, one research challenge is to determine the optimal tradeoff in order to achieve the defined goal in the most effective manner. Since UAVs have limited communication range, they might not be able to share their measurements if appropriate meeting strategies are not implemented.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented our preliminary results on how information on a target location can be maintained and exchanged in a distributed manner between a swarm of helicopters. We described the update mechanisms of grid-based probabilistic maps based on recursive bayesian processes. We also showed some observations obtained with our UAVs in outdoor environments with embedded downward-facing cameras. The results of these observations will be used as inputs for the design and tests of search strategies. Of interest is the assessment of the impact of the sensing resolution (with UAVs deployed at different heights) on cooperation/coordination strategies.

REFERENCES

- [1] M. A. Goodrich, J. L. Cooper, J. A. Adams, C. Humphrey, R. Zeeman, and B. G. Buss, "Using a Mini-UAV to Support Wilderness Search and Rescue: Practices for Human-Robot Teaming," *Journal of Field Robotics*, vol. 25, 2008.
- [2] A. T. Hummingbird, "http://www.asctec.de."
- [3] T. Chung and J. Burdick, "A Decision-Making framework for control strategies in probabilistic search," in *Robotics and Automation, 2007 IEEE International Conference on*, 2007, pp. 4386–4393.
- [4] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989. [Online]. Available: <http://portal.acm.org/citation.cfm?id=68495>
- [5] T. Chung and J. Burdick, "Multi-agent probabilistic search in a sequential decision-theoretic framework," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 2008, pp. 146–151.
- [6] "Truetime: Simulations of networks and embedded control systems." [Online]. Available: <http://www.control.lth.se/truetime/>