

Supporting Search and Rescue Operations with UAVs

Sonia Waharte and Niki Trigoni

University of Oxford

Computing Laboratory

Oxford, United Kingdom

e-mail: {sonia.waharte, niki.trigoni}@comlab.ox.ac.uk

Abstract—Search and rescue operations can greatly benefit from the use of autonomous UAVs to survey the environment and collect evidence about the position of a missing person. To minimize the time to find the victim, some fundamental parameters need to be accounted for in the design of the search algorithms: 1) quality of sensory data collected by the UAVs; 2) UAVs energy limitations; 3) environmental hazards (e.g. winds, trees); 4) level of information exchange/coordination between UAVs.

In this paper, we discuss how these parameters can affect the search task and present some of the research avenues we have been exploring. We then study the performance of different search algorithms when the time to find the victim is the optimization criterion.

I. INTRODUCTION

Search and rescue operations are often characterized by a similar set of constraints: time is critical and any delay can result in dramatic consequences – potentially human losses; operational environments are unfriendly, e.g. disaster scenes, forests, etc. Using Unmanned Aerial Vehicles (UAVs) can provide a critical support to search and rescue operations. UAVs are agile, fast, can exhibit autonomous behaviours and hence perform operations hard to execute by human operators, at low operating costs. In a typical scenario, UAVs will be deployed in an area of interest, perform sensory operations to collect evidence of the presence of a victim, and report their collected information to a remote ground station or rescue team (Fig. 1). UAVs have already demonstrated their benefit

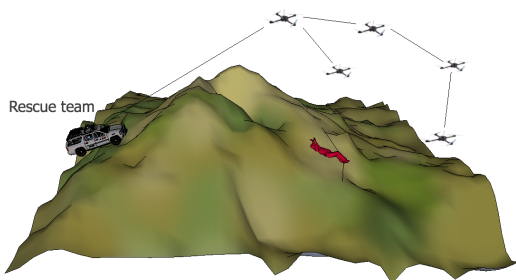


Fig. 1. Rescue scenario with UAVs: the UAVs fly over an area the victim is believed to be located, gather information on its potential location, and transmit it back to a remotely located rescue team.

in search and rescue operations by helping responders to focus their search efforts while avoiding hazards. In 2006, in the

aftermath of Hurricane Katrina, two UAVs were used to survey the damaged area in the search for trapped survivors [1]. The growing recognition of the potential of using UAVs for search and rescue applications is supported by an increasing number of works in the areas of image recognition for victim detection, path planning and task allocation [2]–[4].

In this paper, our focus is the analysis of the performance of different search techniques when the time to find the victim (also referred to as *target*) is the optimization criterion. In particular, we compare the benefit of sharing data between UAVs with different search techniques based on greedy heuristics, potential fields, and partially observable Markov decision process.

The remainder of the paper is organized as follows. We first describe our hardware platforms in Sec. II. We emphasize some research challenges of search and rescue applications with UAVs in Sec. III. We then present the performance evaluation of different search strategies in Sec. IV and Sec. V. We conclude this paper in Sec. VI.

II. UAV QUADROTOR

The hardware platforms we are using are off-the-shelf, electrically-powered quadrotors. We are currently using several models from Ascending Technologies (Fig. 2). These vehicles have the advantage of being easy to prepare, control and fly. The payload they can carry is in the order of a few hundred grammes. With the current battery capacity available, the flight time of the UAVs is restricted to a few tens of minutes. Each of these platforms has been redesigned to accommodate a 1.6 GHz Intel Atom board with 1GB of RAM and 16GB SSD that can provide the computation abilities.

III. RESEARCH CHALLENGES

To optimize the tasks of the UAVs during the search operation, several factors need to be accounted for in the design of the search strategy:

- *Quality of sensory data*: In search and rescue operations, to properly assess the quality/trustworthiness of the information reported to the rescue team is of paramount importance. If a UAV is surveying an area, a victim should not be missed, i.e. the probability of false negative should remain low; at the same time, the probability of false positive, i.e. the probability that a victim is




Fig. 2. UAVs in use in the SUAAVE (Sensing Unmanned Autonomous Aerial VEHicles) project.

considered as detected when it is actually not there, should remain low to avoid sending a ground rescue team to a place of no interest. The challenge is then to establish accurate models of the quality of the sensory data obtained from the UAVs.

- *Energy limitations:* One of the critical aspects of the UAV platforms we are using is their limited flight time. In preliminary experiments, we have established that the current consumption when the UAV is flying and when it is hovering is similar. One way to save energy is then to have the UAV remain idle (rotors off). Depending on the tasks to accomplish and on the capabilities of the UAVs, several search strategies can be envisioned such as alternating periods of flights and periods of rest, or sending only a subset of the UAVs to accomplish certain tasks whilst the rest remains idle (recharging batteries for instance).
- *Environmental hazards.* During its flight, a UAV should be capable of avoiding environmental hazards (trees, building, etc.) as well as avoiding collisions with other UAVs. It is also necessary to account for the fact that depending on the position of the UAV, some areas on the ground might be occluded.
- *Information sharing:* Two aspects of the information sharing process need to be considered: *data fusion* and *network connectivity*.
 - *Data fusion:* When multiple UAVs are deployed, the sensory data they collect can be shared and fused to generate a complete picture of the environment – which can in turn guide the search process. This task is all the more challenging as any solution that will be proposed needs to account for limitations in terms of processing, memory storage, energy consumption, network availability and so on.
 - *Network Connectivity:* During deployment, communications can occur between UAVs and between UAV-Ground Station. The connections can happen opportunistically or can be scheduled. A trade-off needs to be made between reward of establishing a connection (resulting in exchanges of information)

and cost of connection (necessity to travel to a given rendezvous point). Since it is also critical to regularly inform the ground station of the evolution of the mission, it is important to evaluate the benefit to maintain the network fully or partially connected. Other questions that require further investigation are: How frequently should the connections be established between the UAVs? How frequently should the UAVs report to the ground station?

IV. SEARCH ALGORITHMS

In this work, we consider that UAVs are equipped with downward-pointing cameras to detect victims on the ground. By changing altitudes, UAVs can change the size of their observation areas. The higher they fly, the bigger their observation area. But as the UAV flies higher up, the level of detail decreases. Using these observations, we have shown that with a single UAV, changing altitudes is a valid control strategy that can speed up the search process [5], [6]. A question we are investigating in this work is the impact of altitude on the search strategy. When several UAVs are deployed, the complexity of the problem increases. In an initial work [7], we investigated a simple fusion algorithm in which all observations made by all UAVs are exchanged when the UAVs come into communication range and are locally maintained. We use this mechanism in this work when communication between UAVs occur. We also consider differences in quality of sensory data, as well as we account for the presence of obstacles during the search operations. We implicitly account for energy constraints by looking at search strategies that minimize the time to find the victim.

Search algorithms for search and rescue operations should be able to cope with the uncertainties of real-world deployment. Hence real-time approaches are more appropriate and can be divided into three main categories:

- 1) Greedy heuristics
- 2) Potential-based heuristics
- 3) Partially Observable Markov Decision Process (POMDP) based heuristics

To maintain the information on the probability of the victim location, each UAV maintains a grid-based probabilistic map (*belief map*) composed of cells that represent the discretization of the search space. Each cell contains the probability that the

target is present in it. These maps are commonly used for applications such as surface mapping, exploration or navigation (e.g. [8]).

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. 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 rivers or roads (Fig. 3). If no prior information is known, we assume a uniform distribution. 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.

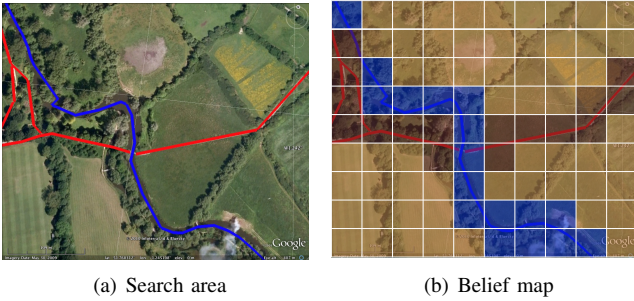


Fig. 3. Example of scenario in which the victim is believed to be along a road (in red) and not in the river (in blue). The search area is mapped into a grid in which the cells in which the victim is located with high probability are shaded in red and the ones with low probability are shaded in blue.

A. Greedy heuristics

The first search strategies we implemented are based on a greedy choice of the cell towards which a UAV should fly based on a predefined criterion. The general structure of the heuristics we implemented is described in Alg. 1. The UAV moves to the neighboring cell s' that has the highest belief (to increase our confidence of the presence of a victim after the sensing operation).

```

Start=s'
Initialize BeliefMap
while true do
  s=s'
  GetObservation(s)
  if target detected then
    Report Base Station
  end
  UpdateBeliefMap(s)
  s'=NextMove(s)
end

```

Algorithm 1: General structure of greedy heuristics

Different tie-breaking strategies can be considered:

- *Rand*: Pick s' at random (the version of Rand where UAVs can communicate is referred to as *MS_Rand*).
- *Cov*: Pick s' that provides the highest coverage.
- *MaxSum*: Pick s' for which the sum of the belief of the cells that the UAV would cover by moving at this location is maximum. The goal is to maximize the information gain.
- *MaxMin*: Pick s' that covers a cell for which the minimum belief is the highest.
- *MaxMax*: Pick s' that covers a cell that has the highest belief.
- *MinNeigh*: Pick s' for which the difference between the maximum and minimum belief is minimum in order to increase our information on the target presence for a greater number of cells at once.

We also study variants of the greedy heuristics (referred to as *1-look ahead*) in which the UAV bases its control decision not on the value of the belief of the victim presence on the neighboring cell but on the overall gain of moving into that cell. The optimization criteria are defined as before and the heuristics we consider are:

- *LA_MaxSum*: Pick s' for which the sum of the belief of the cells that the UAV would cover by moving at this location is maximum. The goal is to maximize the information gain.
- *LA_MaxMin*: Pick s' that covers a cell for which the minimum belief is the highest.
- *LA_MaxMax*: Pick s' that covers a cell that has the highest belief.
- *LA_MinNeigh*: Pick s' for which the difference between the maximum and minimum belief is minimum.

B. Potential-based Algorithm

Potential field methods are based on the following simple concept: goals are associated with attractive potentials and obstacles are associated with repulsive potentials. At each point in the free configuration space (set of configurations that avoid collisions with obstacles), the potential represents the sum of the attractive potentials and repulsive potentials. More formally, let us consider a UAV at position q . Let $U_{obs}(q)$ be the potential created by an obstacle at the UAV location q , and $U_{goal}(q)$ be the potential created by the goal. The potential function U at q represents the sum of the potentials as shown in (1).

$$U(q) = U_{obs}(q) + U_{goal}(q) \quad (1)$$

The force exerted at q is obtained by computing the negative of the gradient of the potential U at q . For example, if we consider a UAV navigating in a 2-D world with fixed orientation, ($q = (x, y)$ in \mathbb{R}^2), we obtain (2):

$$\vec{F}(q) = -\nabla U(q) = \begin{bmatrix} \frac{\partial U}{\partial x} \\ \frac{\partial U}{\partial y} \end{bmatrix} \quad (2)$$

In the design of the attractive and repulsive potential fields, one should consider the following factors:

- *Attractive potentials* should increase as the UAV moves away from the goal. When the UAV reaches the goal, no forces from the goal should act upon it.
- *Repulsive potentials* should increase as the UAV comes closer to an obstacle. But repulsive potentials should not affect the UAV motion when it is far from the obstacle.

Practically, in our search and rescue scenario, we assigned to each cell a weight that depends on its nature: obstacles and already-visited locations. The weight of each cell that is visited is increased by K_v after each visit. A potential is computed for each grid cell according to Alg. 2.

Initialization

For each cell that is an obstacle, assign maxWeight
 $s = \text{start location}$

while true do

 assign weight $w(s) = \min(w(s) + K_v, \text{maxWeight})$

for each cell do

 compute distance d_v to visited cells

$U(\text{cell}) = w(\text{cell})/d_v^2$

end

$s' = \text{neighboring cell with minimum potential}$

 move UAV to location s'

if target in s' then

 Report Base Station

end

$s = s'$

end

Algorithm 2: Potential-based search algorithm

C. Partially Observable Markov Decision Process

Partially Observable Markov Decision Processes (POMDPs) are a generalisation of a Markov Decision Process (MDP) to situations where the system can be modelled by a MDP, but the underlying states are unobservable. Instead, the state distribution must be inferred based on a model of the world and local observations where, typically, sensors provide partial and noisy information about the world. The POMDP framework is general enough to model a variety of real-world sequential decision processes. An exact solution to a POMDP yields the optimal action for each possible belief over the world states. The optimal action maximizes (or minimizes) the expected reward (or cost) of the agent over a time horizon. The sequence of optimal actions is known as the optimal policy and is chosen so as to optimise the expected reward (or cost).

Formally, a POMDP is represented by the following n-tuple: $\{S, A, O, b_0, T, \Omega, R, \gamma\}$, with:

- S : finite set of discrete states,
- A : set of actions,
- O : set of observations,
- b_0 : initial belief states distribution,
- $T(s, a, s')$: probability of transition from s to s' when taking action a ,
- $\Omega(o, s', a)$: probability of observing o from state s' after taking action a ,
- $R(s, a)$: reward when executing action a in state s ,

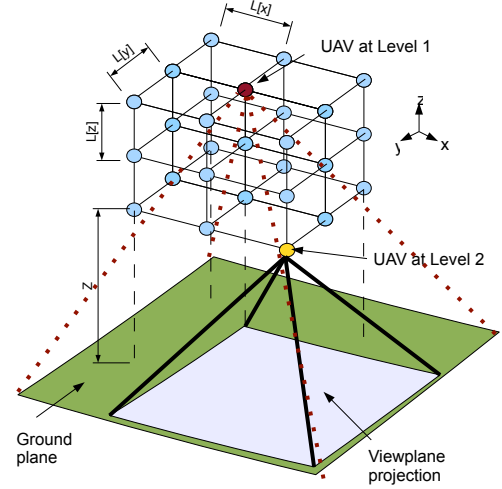


Fig. 4. Scenario for the implementation of Hierarchical POMDP with 2 UAVs flying at 2 different altitudes.

- γ : discount factor.

The goal is to find a set of actions that maximizes the expected sum of rewards over an infinite horizon $E[\sum_t \gamma^t R(s_t, a_t)]$.

In our implementation, we considered that each UAV is able to compute its policy based on information locally maintained. Time constraints have been enforced due to the size of the problem and the number of states that could potentially lead to hours of computation. Since UAVs can fly at different altitudes, we can differentiate the tasks of the UAVs as a function of their altitude (Fig. 4). This can be modelled using a POMDP with different observations models and set of actions for each set of UAVs (referred to as Hierarchical POMDP).

For simplicity, we consider that the UAVs can fly at only two distinct altitudes.

We evaluate three strategies:

- 1) POMDP: UAVs fly at the same altitude but do not exchange any information.
- 2) MS_POMDP: UAVs fly at the same altitude and exchange information when in communication range.
- 3) Hierarchical POMDP (HPOMDP): UAVs fly at different altitudes and exchange information when in communication range.

V. SIMULATIONS

We consider 3D terrains discretized into cells of unit length. In our first evaluations, we only consider trees as obstacles that we model in 3D space by cylinders (Fig. 5). Each cylinder is further discretized into an integer number of grid cells. In our model, we account for the fact that, depending on its position relative to the obstacles, a UAV may not be able to observe a certain number of cells on the ground (shadowing effect).

We consider a sensing model accounting for false positive and false negative:

- Probability(sensing target at height h —target)= $1 - \beta_h$

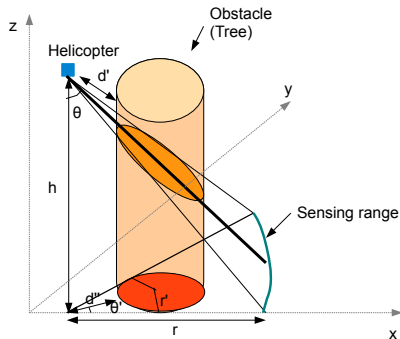


Fig. 5. Example of the effect of the presence of obstacles on the sensing range: a tree (cylinder) blocks a UAV field of view, and makes some areas on the ground not observable.

TABLE I
OBSERVATION MODEL

Altitude (h)	α_h	β_h
5m	0.243599	0.000000
10m	0.028369	0.000000
15m	0.026099	0.046211
20m	0.001110	0.046745

- Probability(not sensing target at height h —target)= β_h
- Probability(not sensing target at height h —no target)= $1 - \alpha_h$
- Probability(sensing target at height h —no target)= α_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 and vary as a function of the height h of the helicopter. From our field experiments [5], we obtained the values of the observation model summarized in Table I.

The sensing areas on the ground cover a set of $M(h)$ cells, where h is the altitude of the UAV considered. 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).

To first assess the performance of search algorithms before trials with real hardware, we used a Matlab/Simulink-based simulator for real-time control systems. This allows to accurately simulate distributed systems with realistic radio transmission models. The characteristics of the wireless channel are:

- Transmit power = -30 dBm
- Receiver threshold = -48 dBm
- Path loss coefficient = 3.5

Please note that the wireless settings have been adjusted to scale with the test areas. The grid cells are 1x1x1 cube meters and the maximum wireless transmission range is 2.27m. We considered squared coverage areas that correspond to a 3x3 grid for the UAVs evolving at level 2, and 4x4 grid for the UAVs evolving at level 1. The maximum simulation time is

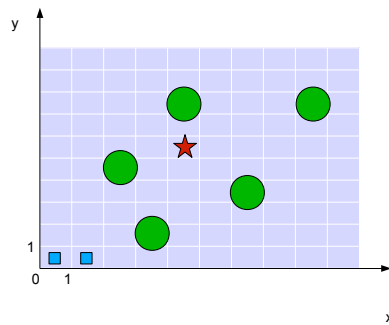


Fig. 6. Test configuration: top view of the test area with 2 UAVs at the bottom left corner, 5 trees and 1 goal (victim) centrally located.

200sec. A UAV periodically broadcasts its position and other relevant information (such as belief maps). Decision on the next step to go towards to is taken when the UAV arrives close to the center of a cell. All UAVs fly at the same speed regardless of their altitude.

To solve POMDPs, we used the ZMDP software [9].

The topology we simulated is a 10x10 grid, with 5 trees randomly positioned. 2 UAVs start their search at nearby position. One target is located in the middle of the test area (Fig. 6).

We assume that the mission is successfully completed when the target has been positively identified. The result of the simulations is depicted in Fig. 7.

We can see that the approach based on POMDP and look-ahead strategies performs the best compared to the greedy heuristics and the potential-based approaches. Greedy heuristics perform in general poorly as the control decision of the UAVs is based only on local information which can lead them to stay longer in areas that have already been explored instead of going to remote areas that have not been explored yet. Potential-based approaches also generally perform poorly due to the known problem of local minima that these search techniques suffer from. Additional heuristics are required to help the UAVs escape from these local minima but this unavoidably impacts the time to locate the target.

This simple simulation scenario demonstrates that trying to estimate the best search paths by sharing information from sensory operations and by exploiting the acquired or prior knowledge of the possible victim location can help speed up the search task. But this result has a computational cost that can not be neglected. With greedy heuristics, the computational cost is minimal as the control decision is based on the value of the belief of the victim position in direct neighboring locations. With potential-based searches, the computational cost is linear with the number of cells. Solving a POMDP in an exact fashion is an intractable problem [10]. To find a solution for a POMDP requires to consider each possible action and each possible observation, which implies that the cost of a solution grows exponentially with the desired planning horizon. Rather than advocating for the use of POMDP or look-ahead strategies for search tasks, this simple case scenario

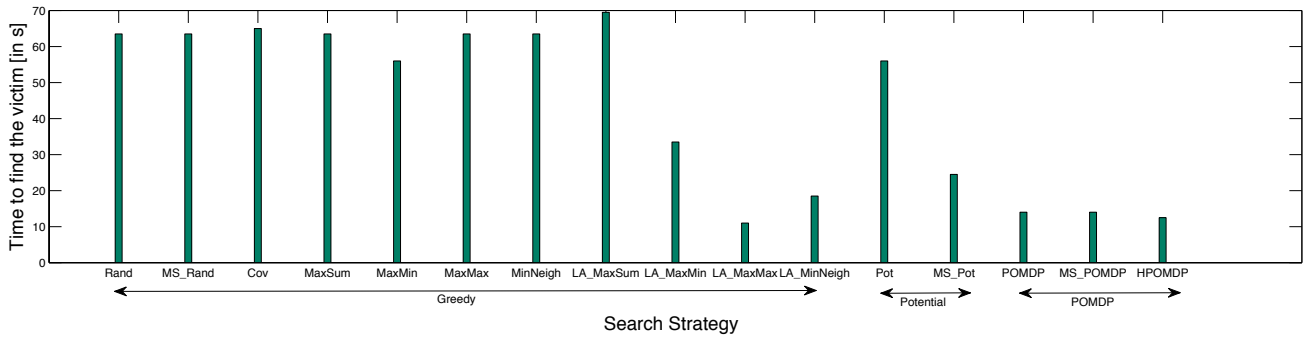


Fig. 7. Simulation results showing the time to find the victim (completion time) with different search techniques.

illustrates the potential of POMDP when the number of states remains limited, and the importance of exploiting as much as possible the information acquired by one or multiple UAVs to predict more accurately a victim location.

VI. CONCLUSION

Unmanned aerial vehicles can provide a critical support for search and rescue operations. However in order to achieve their full potential, it is necessary to properly account for all the parameters that can affect the flight of the UAVs such as quality of sensory operations (that can depend on the position of the UAVs for instance), energy constraints, environmental hazards or data sharing constraints between UAVs and rescue teams.

In this paper, we studied different search strategies based on greedy heuristics, potential-based algorithms and partially observable Markov decision process to design the control strategy of several UAVs. The evaluation criterion was the time taken by the UAVs to find the victim. Our preliminary results show the importance of exploiting as much as possible the information obtained during the search operations and illustrate the potential of POMDP in a simple scenario. However this comes at a high computational cost that needs to be accounted for.

In future work, we intend to investigate more complex scenarios and to account for energy and connectivity constraints. We also intend to study how search algorithms based on POMDP scale when the number of states increases.

REFERENCES

- [1] NSF - Hurricane Katrina, "<http://www.nsf.gov/news/>."
- [2] P. Doherty and P. Rudol, "A uav search and rescue scenario with human body detection and geolocalization," in *In 20th Joint Conference on Artificial Intelligence (AI07)*, 2007.
- [3] M. Goodrich, B. Morse, D. Gerhardt, J. Cooper, M. Quigley, J. Adams, and C. Humphrey, "Supporting wilderness search and rescue using a camera-equipped mini uav: Research articles," *J. Field Robot.*, vol. 25, no. 1-2, pp. 89–110, 2008.
- [4] B. Morse, C. Engh, and M. Goodrich, "Uav video coverage quality maps and prioritized indexing for wilderness search and rescue," in *HRI*. New York, NY, USA: ACM, 2010, pp. 227–234.
- [5] A. Symington, S. Waharte, S. J. Julier, and N. Trigoni, "Probabilistic target detection by camera-equipped uavs," in *ICRA*, 2010.
- [6] N. T. S. Waharte, A. Symington, "Probabilistic search with agile uavs," in *ICRA*, 2010.

- [7] S. Waharte, N. Trigoni, and S. J. Julier, "Coordinated search with a swarm of uavs," in *SECON*, 2009.
- [8] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [9] ZMDP, "<http://www.cs.cmu.edu/~trey/zmdp/>."
- [10] O. Madani, S. Hanks, and A. Condon, "On the undecidability of probabilistic planning and infinite-horizon partially observable markov decision problems," in *AAAI '99*, 1999, pp. 541–548.