

Probabilistic Search with Agile UAVs

Sonia Waharte, Andrew Symington, Niki Trigoni

Abstract— Through their ability to rapidly acquire aerial imagery, Unmanned Aerial Vehicles (UAVs) have the potential to aid target search tasks. Many of the core algorithms which are used to plan search tasks use occupancy grid-based representations and are often based on two main assumptions. Firstly, the altitude of the UAV is constant. Secondly, the onboard sensors can measure the entire state of an entire grid cell. Although these assumptions are sufficient for fixed-wing, high speed UAVs, we do not believe that they are appropriate for small, lightweight, low speed and agile UAVs such as quadrotors. These platforms have the ability to change altitude and their low speed means that multiple measurements may easily overlap multiple cells for substantial periods of time.

In this paper we extend a framework for probabilistic search based on decision making to incorporate multiple observations of grid cells and changes in UAV altitude. We account for observation areas that completely and partially cover multiple grid cells. We show the resultant impact on a number of simulation examples.

Index Terms— Unmanned aerial vehicle, search, exploration, target.

I. INTRODUCTION

For many applications ranging from surveillance to search and rescue, the ability to monitor an environment and find a target of interest is of paramount importance [4], [11]. UAVs have the potential to aid this task through rapidly collecting aerial imagery. In the Wilderness Search and Rescue (WiSAR), for example, the search task often consists of finding evidence and using it to constrain the location of a missing person [12].

UAVs are active systems and therefore the trajectory of the UAV must be controlled to optimise information collection. Many path planning algorithms use occupancy-grid based representations of the environment [4]–[6], [8]. Such representations are advantageous because they can incorporate both *positive information* (detection of the target) and *negative information* (no detection of the target) and may also maintain complicated spatial distributions of where the target might be.

However, most of these algorithms utilise two common assumptions. The first is that the altitude of the UAV remains fixed. As a consequence, the sensor coverage region and sensor properties are the same everywhere. Secondly, the UAV sensors monitor the state of a *single* grid cell in its entirety to determine occupancy. Both of these assumptions are relevant for fixed-wing UAVs, where straight and level

flight is often desired and the movement of the UAV is sufficiently fast that successive measurements lie in separate grid cells. The advent of small, lightweight, agile and low speed UAVs has meant that changing altitude becomes a valid control strategy. Furthermore, low speed means that many observations can lie within the same grid cell. Changing altitude also means that the sensor coverage at one altitude could not cover a complete number of cells at a different altitude. Although these difficulties may be overcome by refining the decomposition of the environment into more and smaller cells, this has consequences both in terms of computational and storage costs.

In this paper we propose a generalized probabilistic search framework that takes into account the following cases:

- 1) The observation region of a sensor can *completely* cover multiple grid cells, and not just a single grid cell.
- 2) The observation region of a sensor can *partially* cover multiple grid cells, and not just a single grid cell.
- 3) Observations can be performed at different heights, with different sensing qualities.

The structure of this paper is as follows. In Section II we lay out the probabilistic search framework and discuss its limitations. Section III extends the framework to consider the case in which the UAV’s sensor completely observes multiple grid cells. In Section IV we extend this analysis to include cells which are *partially* observed. Although the update solution at a single timestep can be readily formulated, it cannot be readily formulated over multiple timesteps because of unmodelled dependency issues between observed and unobserved regions in a single cell. We do not address this issue in this paper. An exploration algorithm, which utilises height as a control variable, is discussed in Section V and results for a simulation scenario are presented in Section VI. We present related work in Section VII. The summary and conclusions are discussed in Section VIII.

II. PROBLEM STATEMENT

A. Occupancy Grid Representation

The objective is to search for a single, stationary target x_T which is suspected to lie in a two-dimensional search region \mathcal{A} [5]. The environment itself is exhaustively decomposed into a set of $|\mathcal{A}|$ disjoint (non-overlapping) regions or *cells*, where the a th cell is C_a . Because the decomposition is both exhaustive and disjoint, we have:

$$\mathcal{A} = \bigcup_{a=1}^{|\mathcal{A}|} C_a. \quad (1)$$

This work was supported by the SUAAVE project. More information on this project can be found at <http://www.suaave.org>.

Andrew Symington, Niki Trigoni and Sonia Waharte are with the Oxford Computing Laboratory, Wolfson Building, Oxford, OX1 3QD, UK. niki.trigoni@comlab.ox.ac.uk

The search task attempts to achieve two goals: to determine if $x_T \in \mathcal{A}$ and, if so, to determine in which cell the target lies. Let $Pr(x_T \in \mathcal{A})$ be the probability that the target lies within the region \mathcal{A} . Using the law of disjoint probabilities and the decomposition given in Eq. 1,

$$Pr(x_T \in \mathcal{A}) = \sum_{a=1}^{|\mathcal{A}|} Pr(x_T \in \mathcal{C}_a)$$

$$Pr(x_T \notin \mathcal{A}) = 1 - Pr(x_T \in \mathcal{A})$$

Therefore, the solution to the search task entails computing $Pr(x_T \in \mathcal{C}_a)$, $\forall a = [1, \dots, |\mathcal{A}|]$. These probabilities are updated by a UAV which traverses the area and is equipped with a sensor.

B. UAV Pose and Observations

We introduce in this section the notations for the UAV pose and the observation model we use in the remainder of the paper. The pose k of the UAV (which includes its position and altitude) at time t is represented by the vector k^t . The UAV is equipped with an on-board sensor, which has an *observation region* $\mathcal{O}(k^t)$. This corresponds to the region of \mathcal{A} which is visible to the sensor at the current time.

A sensor return d_t taken at time t for a UAV position k^t is binary and corresponds to either a target detection or target no detection event. To account for clutter and missed detections, we use an observation model similar to the ones described in [5] and [6]:

$$\begin{aligned} Pr(d_t = 1 | x_T \in \mathcal{O}(k^t)) &= 1 - \beta(k^t), \\ Pr(d_t = 0 | x_T \in \mathcal{O}(k^t)) &= \beta(k^t), \\ Pr(d_t = 0 | x_T \notin \mathcal{O}(k^t)) &= 1 - \alpha(k^t), \\ Pr(d_t = 1 | x_T \notin \mathcal{O}(k^t)) &= \alpha(k^t). \end{aligned} \quad (2)$$

where $\alpha(k^t)$ is the probability of false alarm, and $\beta(k^t)$ represents the probability of missed detection for a UAV position k^t .

C. Single Cell Observation

Chung et al. [5] considered the special case where:

- 1) The UAV position k^t is the middle of a single grid cell.
- 2) The observation region directly maps to a single cell.
- 3) The sensor characteristics do not change and so $\alpha(k^t) = \alpha$ and $\beta(k^t) = \beta$.
- 4) Observations are assumed to be independent.

Under these conditions, the recursive formulation to update the occupancy probability for each grid cell can be written as follows. Let $D_t = \{d_1, \dots, d_t\}$ be the set of observations from time 1 to time t . After a sensor measurement by a UAV located at k at time t , the probability of target presence in each grid cell \mathcal{C}_a of the search area \mathcal{A} is updated as follows:

$$Pr(x_T \in \mathcal{C}_a | D_t) = \frac{Pr(d_t | x_T \in \mathcal{C}_a) Pr(x_T \in \mathcal{C}_a | D_{t-1})}{Pr(d_t | D_{t-1})}. \quad (3)$$




Fig. 1. Multiple complete and partial cell observations. The grid cells are the white rectangles, the observation region is the grey square.

where:

- $Pr(d_t | x_T \in \mathcal{C}_a)$ is obtained from the observation model (Eq. 2),
- $Pr(x_T \in \mathcal{C}_a | D_{t-1})$ represents the prior probability of target presence in cell \mathcal{C}_a ,
- $Pr(d_t | D_{t-1})$ is a normalization factor such that: $Pr(d_t | D_{t-1}) = Pr(d_t | x_T \in \mathcal{A}) Pr(x_T \in \mathcal{A} | D_{t-1}) + Pr(d_t | x_T \notin \mathcal{A}) Pr(x_T \notin \mathcal{A} | D_{t-1})$.

Although these equations provide a Bayes optimal formulation for the search problem, they introduce a number of constraints. The most important of these is that the observation is of a *single cell*. This creates many limitations as it restricts the altitude at which the UAV should fly to cover a fixed size region. With agile UAVs, one prefers a more flexible solution. Therefore, a key challenge is to remove the condition that the observation region aligns with a single cell. Consider two cases: first the case where the observation region completely covers multiple cells, and second the case where it partially covers multiple cells.

III. MULTIPLE COMPLETE CELL OBSERVATIONS

Consider the situation shown in Figure 1(a): the observation region consists of the *union* of a set of grid cells,

$$\mathcal{O}(k^t) = \bigcup_{\mathcal{C}_a \in \mathcal{O}(k^t)} \mathcal{C}_a. \quad (4)$$

To update the probability of target presence in a grid cell, we need to distinguish the case where the grid cell is directly observed, and where it is not directly observed.

A. Updating Completely Observed Grid Cells

The update rule is a straightforward extension to Eq. 3, but with the modification that the *same* likelihood is applied over all cells in the observation area. Then, we only need to compute the probability of target presence over the observation area $\mathcal{O}(k^t)$ and redistribute this probability for all cells $\mathcal{C}_a \in \mathcal{O}(k^t)$.

Using Bayes' Rule, the update occupancy in the *observation region* is given by:

$$Pr(x_T \in \mathcal{O}(k^t) | D_t) = \frac{Pr(d_t | x_T \in \mathcal{O}(k^t)) Pr(x_T \in \mathcal{O}(k^t) | D_{t-1})}{Pr(d_t | D_{t-1})}. \quad (5)$$

Since the grid cells do not overlap, the prior probability of the target lying in the observation region is:

$$Pr(x_T \in \mathcal{O}(k^t) | D_{t-1}) = \sum_{\mathcal{C}_a \in \mathcal{O}(k^t)} Pr(x_T \in \mathcal{C}_a | D_{t-1}).$$

$Pr(d_t | x_T \in \mathcal{O}(k^t))$ is given by the sensor observation model (Eq. 2) and $Pr(d_t | D_{t-1})$ is a normalization factor.

B. Updating Unobserved Grid Cells

For the grid cells not in the observation area, we apply Eq. 3 directly. Nonetheless, we need to prove that $Pr(d_t | x_T \in \mathcal{C}_a)$, the probability of target detection given that the target is in the grid cell \mathcal{C}_a , can be obtained from the observation model described in Eq. 2 (Theorem 1).

Theorem 1: Given that the target lies in \mathcal{C}_a , the probability of detection is given by:

$$Pr(d_t | x_T \in \mathcal{C}_a) = \begin{cases} Pr(d_t | x_T \in \mathcal{O}(k^t)) & \text{for } \mathcal{C}_a \in \mathcal{O}(k^t) \\ Pr(d_t | x_T \notin \mathcal{O}(k^t)) & \text{for } \mathcal{C}_a \notin \mathcal{O}(k^t) \end{cases}$$

Proof: We only provide the proof of the first part of the theorem, as the second part can be derived in a similar manner. Marginalising, the sensor likelihood can be written as:

$$Pr(d_t | x_T \in \mathcal{O}(k^t)) = \sum_{\mathcal{C}_a \in \mathcal{O}(k^t)} \{Pr(d_t | x_T \in \mathcal{C}_a, x_T \in \mathcal{O}(k^t)) \times Pr(x_T \in \mathcal{C}_a | x_T \in \mathcal{O}(k^t))\}. \quad (6)$$

Now, given that \mathcal{C}_a is a subset of $\mathcal{O}(k^t)$,

$$Pr(d_t | x_T \in \mathcal{C}_a, x_T \in \mathcal{O}(k^t)) = Pr(d_t | x_T \in \mathcal{C}_a).$$

Substituting into Eq. 6, the update is:

$$Pr(d_t | x_T \in \mathcal{O}(k^t)) = \sum_{\mathcal{C}_a \in \mathcal{O}(k^t)} \{Pr(d_t | x_T \in \mathcal{C}_a) \times Pr(x_T \in \mathcal{C}_a | x_T \in \mathcal{O}(k^t))\}.$$

From the assumption that the detection properties are constant throughout the detection region, $Pr(d_t | x_T \in \mathcal{C}_a)$ is the same for all $\mathcal{C}_a \in \mathcal{O}(k^t)$. Therefore,

$$\begin{aligned} Pr(d_t | x_T \in \mathcal{O}(k^t)) &= Pr(d_t | x_T \in \mathcal{C}_a) \\ &\times \sum_{\mathcal{C}_a \in \mathcal{O}(k^t)} \{Pr(x_T \in \mathcal{C}_a | x_T \in \mathcal{O}(k^t))\} \\ &= Pr(d_t | x_T \in \mathcal{C}_a). \end{aligned}$$

Finally, the normalisation term $Pr(d_t | D_{t-1})$ must be computed. Using the Chain Rule,

$$\begin{aligned} Pr(d_t | D_{t-1}) &= Pr(d_t | x_T \notin \mathcal{A}) Pr(x_T \notin \mathcal{A} | D_{t-1}) \\ &+ Pr(d_t | x_T \in \mathcal{A}) Pr(x_T \in \mathcal{A} | D_{t-1}). \end{aligned}$$

Although this can be applied with multi-resolution occupancy grids (and thus approximate changes in altitude), it does not fundamentally address the case in which the boundaries of the observation region do not align with the occupancy grid. We now consider this case.

IV. UNALIGNED CELL OBSERVATIONS

Consider the case illustrated in Figure 1(b): the boundaries of the observation region do not align with those of the grid cells. We incorporate this information by first splitting the partially observed grid cells into regions that overlap and regions that do not overlap with the observation region. Then we update the overlapping regions using the multiregion update expression, and finally combine the split grid cells.

A. Splitting the Observation Region

Let $\mathcal{O}(k^t)$ be the set of indices of all grid cells that intersect with $\mathcal{O}(k^t)$ when the UAV is located at k at time step t . This can be decomposed into $\mathcal{O}(k^t) = \mathcal{O}_1(k^t) \cup \mathcal{O}_2(k^t)$, where $\mathcal{O}_1(k^t)$ are the indices of all cells which lie *completely* within $\mathcal{O}(k^t)$ and $\mathcal{O}_2(k^t)$ are the indices of cells which only *partially* lie within $\mathcal{O}(k^t)$.

Each partially overlapped cell can be divided into the following two regions,

$$\mathcal{C}_a = \mathcal{O}_a(k^t) \cup \mathcal{O}'_a(k^t), \quad \forall a \in \mathcal{O}_2(k^t).$$

where $\mathcal{O}_a(k^t)$ is the part of cell \mathcal{C}_a that is observed when the UAV is located at k at time step t , and $\mathcal{O}'_a(k^t)$ is the part of cell a that is not observed.

Therefore, the observation region can be written as:

$$\begin{aligned} \mathcal{O}(k^t) &= \bigcup_{a \in \mathcal{O}(k^t)} (\mathcal{C}_a \cap \mathcal{O}(k^t)) \\ &= \bigcup_{a \in \mathcal{O}_1(k^t)} \mathcal{C}_a + \bigcup_{a \in \mathcal{O}_2(k^t)} \mathcal{O}_a(k^t). \end{aligned} \quad (7)$$

B. Updating Cells in Unaligned Observation Region

We need consider two cases: a *single update* from a non-aligned sensing region, and the effects of fusing multiple unaligned measurements over time.

The single step update case applies the multiple complete cell update equation (5) to the regrided cell. Once the update has been performed, the split cells are combined, and the probabilities computed. Specifically, consider a cell $a \in \mathcal{O}_2(k^t)$. This has been decomposed into the cells $\mathcal{O}_a(k^t)$ and $\mathcal{O}'_a(k^t)$. Given that the regions are disjoint, the probability of target presence in \mathcal{C}_a can be expressed as follows:

$$\begin{aligned} Pr(x_T \in \mathcal{C}_a | D_t) &= Pr(x_T \in \mathcal{O}_a(k^t) | D_t) \\ &+ Pr(x_T \in \mathcal{O}'_a(k^t) | D_t). \end{aligned} \quad (8)$$

Therefore, using Bayes' Rule,

$$Pr(x_T \in \mathcal{O}_a(k^t) | D_t) = \frac{A(\mathcal{O}_a(k^t))}{A(\mathcal{O}(k^t))} Pr(x_T \in \mathcal{O}(k^t) | D_t). \quad (9)$$

$$Pr(x_T \in \mathcal{O}'_a(k^t) | D_t) = \left(1 - \frac{A(\mathcal{O}_a(k^t))}{A(\mathcal{C}_a)}\right) Pr(x_T \in \mathcal{C}_a | D_t). \quad (10)$$

where the function A returns the size of the region that it takes as a parameter.

However, although this straightforward generalisation for the partially observed case is correct for a single update, it is not correct for the case in which multiple observations of

partially observed cells are carried out. Suppose a stationary UAV continued to view the same part of the environment with a sensor and that the sensor always returned that no object is detected. Although the probability of occupancy for all cells in $O_1(k^t)$ will (correctly) tend to be zero, the probability of occupancy for all the cells in $O_2(k^t)$ will also tend to zero. However, this is incorrect: only parts of the cells observed in $O_2(k^t)$, and not the entire cells themselves, have been observed. The reason is that in Eq. 8, the occupancy probabilities for $\mathcal{O}_a(k^t)$ and $\mathcal{O}'_a(k^t)$ are mixed together across the entire cell.

The optimal solution is to propagate the decomposed grid structure. However, as explained above, this leads to significant computational and storage costs. We are currently seeking a more principled approximation that is based on the assumption that underestimating target existence probability is a more costly mistake than overestimating it. From (8), we can obtain that

$$Pr(x_T \in C_a | D_t) \leq \frac{A(\mathcal{O}_a(k^t))}{A(\mathcal{O}(k^t))} + \frac{A(\mathcal{O}'_a(k^t))}{A(C_a)}$$

$$Pr(x_T \in C_a | D_t) \leq \frac{A(\mathcal{O}_a(k^t))}{A(\mathcal{O}(k^t))} + (1 - \frac{A(\mathcal{O}_a(k^t))}{A(C_a)}). \quad (11)$$

Therefore, if a cell is completely observed and if there is a target in this cell, $\frac{A(C_a)}{A(\mathcal{O}(k^t))}$ represents an upper bound on the probability of target presence in this cell.

In our experiments presented below we found that, given the speed of the UAV, the detrimental effects were insignificant and did not impact the overall algorithm's performance. However, this is largely attributable to the fact that the movement of the UAV was relatively fast.

V. SEARCH AND RESCUE EXPLORATION ALGORITHM

To illustrate how partial cell observations may be exploited for optimizing search and rescue operations, let us consider the following case scenario. We assume that the search space is discretized into a 10 by 10 grid, with each square cell having 5-meter sides. To obtain observation areas of size 5x5m, 5.8x5.8m and 6.6x6.6m, we can either use camera sensors of 2.4x2.4mm, 2.8x2.8mm and 3.2x3.2mm, with a focal length of 4.8mm with the UAV flying at 10m, or alternatively use a camera sensor of 2.4x2.4mm with a focal length of 4.8mm and the UAV flying at 10m, 11.6m and 13.2m. The variations in altitude are however not significant enough to obtain a noticeable difference in the probabilities of false alarm and missed detection at each altitude. We therefore use the same values for all three altitudes. Each cell in the grid is identified by discrete coordinates (i, j) , with (0,0) being the starting point of the UAV. We implement a control algorithm based on a 1-step look ahead gradient ascent strategy. Basically, the UAV chooses to move to the neighboring cell for which the probability of presence of the target is the highest. Although many other control strategies can be implemented, we use this simple approach for illustration purposes.

Figure 2 shows the evolution of the maximum probability of target presence across the cells of the search area for three different UAV altitudes. Initially (up to the first 23 steps), the UAV that flies at the highest altitude ($h_3=13.2\text{m}$) and reaches a higher maximum detection probability than the other two UAVs operating at lower altitudes. However, the situation is gradually reversed as UAVs continue to fly. The UAV flying at $h_2=11.6\text{m}$ outperforms the other two, between steps 23 and 98, whereas the lower flying UAV ($h_1=10\text{m}$) becomes the best after step 98. We can see that, although the belief




Fig. 2. Evolution of the maximum probability of target presence for different UAV altitudes.

of the presence of a target increases faster when a larger observation area is used (at a higher altitude), the maximum confidence on the presence of the target is lower.

We therefore derive a search and rescue exploration algorithm (Alg. 1) using partial cell observation and varying altitude. We assume a conservative approach and uniformly distribute the target location probability over the search area if no prior knowledge of the target location is given. Let h_0, \dots, h_m be the set of heights at which a UAV may fly. For each given height h_i ($0 \leq i \leq m$) we define a threshold for the detection probability $Threshold_{h_i}$ that should be less than $\frac{A(C_a)}{A(\mathcal{O}(k^t))}$. This threshold represents the maximum confidence that can be achieved on the presence of target when the UAV is flying at this particular height. After each observation, the UAV evaluates from its occupancy grid the maximum probability $MaxBelief$ on where the target lies. If this probability is greater than a probability $MaxThreshold$, which represents the upper bound on the probability at which a target is declared as detected, or if it is less than a probability $MinThreshold$, which represents the lower bound on the probability at which a target is declared as absent from the search area, the mission is aborted. If $MaxBelief$ does not provide conclusive evidence of target presence or absence, we use the values of the thresholds $Threshold_{h_k}$ to determine at which altitude the UAV should fly.

Algorithm 1: Advanced search algorithm

```
Init: start exploration at height  $h_k(0 \leq k \leq m)$   
while true do  
  if ( $MaxBelief > MaxThreshold$ ) ||  
    ( $MaxBelief < MinThreshold$ ) then  
    | terminate mission  
  else if ( $MaxBelief < Threshold_{h_k}$ ) &&  
    ( $MaxBelief > Threshold_{h_{k+1}}$ ) then  
    | keep exploring at the same height  $h_k$   
  else if ( $MaxBelief > Threshold_{h_k}$ )&&(  $h_k > h_0$ )  
  then  
    | reduce altitude from  $h_k$  to  $h_{k-1}$  to reduce  
    | observation area  
  else if ( $MaxBelief < Threshold_{h_{k+1}}$ ) &&  
    ( $h_{k+1} < h_m$ ) then  
    | increase altitude  $h_k$  to  $h_{k+1}$  to increase  
    | observation area  
  end  
end
```

VI. EXPERIMENTS AND RESULTS

We compare two approaches for target search: a basic search strategy where UAVs operate at a fixed height h_1 and where only completely covered cells are considered; and an advanced search strategy where UAVs can operate at two heights h_1 and h_2 , and exploit partial cell observations (Alg. 1). In our simulations, we set $h_1 = 10m$ and $h_2 = 11.6m$. We fixed $Threshold_{h_1}$ to 0.3 and $Threshold_{h_2}$ to 0.95. These values have been selected as they are between the prior (0.01 for a uniform distribution in a 10x10 grid) and the maximum threshold value ($\frac{A(C_a)}{A(O(k^h))} = 1$ for h_1 , and 0.79 for h_2). $MaxThreshold$ was set to 0.95. A careful study of how threshold values must be selected will be part of future work. We also consider that the UAV position is perfectly known and that the UAV is able to move from the center of a cell to the center of another cell. We assume that exactly one target exists in the search region but without any prior information on its actual possible location. The prior probabilities of target position are equal over the search area. We evaluate the time to target detection in terms of the number of moves from one cell to the next adjacent cell. In each cell, a picture is taken and a sensing algorithm is used to determine the presence of an object of interest. The feature-based sensing algorithm we implemented makes use of the Speeded-up Robust Features (SURF) transform [1]. This algorithm calculates a collection of points in the image that are most likely to be robust against changes in lighting, scale, rotation and perspective. SURF features were obtained for the template set images from a sample video frame. The features in each template image were individually paired with the best matching feature in the video frame. Weak correspondences were culled and those template images with a low total number of correspondences were discarded. The correspondence set was then passed to an implementation of Random Sample Consensus (RANSAC), which calculated the best-fitting projection of the template image into the scene [9]. The template object with the greatest number




Fig. 3. Distribution of the number of moves required to find the target for both the basic and advanced search strategy. The data was acquired by exhaustively testing all possible target positions.

of correspondences was chosen as the best fitting object, with its centre point being given by the planar projection. We derived our observation model from videos taken at 10m using Hummingbird quadrotor UAV from Ascending Technologies¹ and Point Grey Camera. In our observation model, we derived 0.2 and 0.3 as probabilities of false alarm and missed detection respectively.

Figures 4 and 5 show the evolution of the probability of the target presence at different time steps. They show how the belief of target presence decreases in regions that have been observed. When the target has been detected, the probability of target presence progressively increases until a predefined threshold is reached.

Figure 3 shows the distribution of the number of moves required when the target is positioned in each of the cells of the search area. It shows that the advanced search strategy typically finds a target faster than the basic strategy.

Hence, we conclude that a hierarchical approach to search and rescue operations with UAVs operating at different heights drastically improves the time to target discovery. In the advanced search strategy, a UAV starts at a higher altitude where it obtains a rough approximation of the target location and gradually decreases its altitude to refine its estimate. This can be accomplished either by changing a single UAV's position or by collaboration between multiple UAVs.

VII. RELATED WORK

Similarly to our work, Goodrich et al. have also considered the tradeoff between coverage and detection quality as the altitude of the UAV varies [12]. The higher the UAV flies the larger the observation area, but the less accurately one can detect the presence or absence of the target. Our work differs in three respects: 1) We introduce a different notion of degradation of sensing quality as a function of altitude. As altitude increases, the observation area includes more cells, and it becomes more difficult to distinguish the actual cell where the target lies. 2) Using a feature-based sensing algorithm we derive realistic probabilities of missed detection and false alarm, which we take into account in the search task. 3) We propose a search algorithm in which UAVs can

¹Ascending Technologies Hummingbird, <http://www.ascotec.de>




Fig. 4. Evolution of the updates of the probabilistic occupancy grid with basic search strategy (fixed UAV altitude and fully observed cells).




Fig. 5. Evolution of the updates of the probabilistic occupancy grid with advanced search strategy (varying UAV altitude and partially observed cells).

dynamically and autonomously alter their altitude in response to their belief of the target presence. Goodrich et al. selected a fixed UAV altitude so that human operators can clearly see the target in the pictures that UAVs generate.

The use of occupancy grids to maintain and share information has been very popular to address localization or navigation problems due to its scalability to large-scale environments [3], [8]. They have also been used for search and detection of individuals. For instance, Berclaz et al. proposed an approach that uses a synthetic model of a background to estimate the presence of pedestrians [2], [10].

Closer to our work, Chung et al. proposed a probabilistic framework to the search problem for multiple UAVs searching for potentially multiple targets [5]–[7]. We extended this framework to support observations of multiple grid cells by a single UAV, observations of partial grid cells (for a single update step), and to support changes in UAV altitudes.

VIII. CONCLUSION AND FUTURE WORK

In this paper we addressed the problem of extending a probabilistic search framework to account for the properties of agile UAVs. Specifically, we considered observations of multiple grid cells by a single UAV, observations of partial grid cells (for a single update step), and to support changes in UAV altitudes. We showed that developing a search strategy exploiting the tradeoff between size of the observation area and confidence in a target presence can improve the average time-to-target detection.

In the future, we shall examine alternative, non-grid based approaches such as the Probabilistic Hypothesis Density Filter. We should also consider changes in detection probabilities and coverage size as a function of altitude and we will evaluate the performance of our algorithm extensively with non-uniform prior distributions of the target location.

ACKNOWLEDGMENTS

The work in this paper was supported under the EPSRC-funded project “SUAAVE: Sensing Unmanned Autonomous Aerial Vehicles” (EP/F064179/1). The authors would also like to thank Dr. Simon J. Julier for his valuable inputs and Dr. Luke Teacy for comments on an earlier draft of this paper.

REFERENCES

- [1] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool. Speeded-up robust features (SURF). *Comput. Vis. Image Underst.*, 110(3):346–359, 2008.
- [2] J. Berclaz, A. Shahrokni, F. Fleuret, J. Ferryman, and P. Fua. Evaluation of probabilistic occupancy map people detection for surveillance systems. In *PETS*, pages 55–62, Miami, Florida, 2009.
- [3] A. Birk and S. Carpin. Merging occupancy grid maps from multiple robots. *Proceedings of the IEEE*, 94(7):1384–1397, July 2006.
- [4] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Coordinated decentralized search for a lost target in a bayesian world. In *IROS*, volume 1, pages 48–53 vol.1, 2003.
- [5] T.H. Chung and J.W. Burdick. A Decision-Making framework for control strategies in probabilistic search. In *ICRA*, 2007.
- [6] T.H. Chung and J.W. Burdick. Multi-agent probabilistic search in a sequential decision-theoretic framework. In *ICRA*, 2008.
- [7] T.H. Chung, M. Kress, and J. O. Royset. Probabilistic Search Optimization and Mission Assignment for Heterogeneous Autonomous Agents. In *ICRA*, 2009.
- [8] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.
- [9] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
- [10] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua. Multicamera people tracking with a probabilistic occupancy map. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):267–282, 2008.
- [11] T. Furukawa, F. Bourgault, B. Lavis, and H.F. Durrant-Whyte. Recursive bayesian search-and-tracking using coordinated uavs for lost targets. In *ICRA*, pages 2521–2526, 2006.
- [12] M. Goodrich, J. Cooper, J. Adams, C. Humphrey, R. Zeeman, and B.Buss. Supporting Wilderness Search and Rescue using a Camera-Equipped Mini UAV. *Journal of Field Robotics*, 25(1–2):89–110, 2008.