

Observation Modelling for Vision-Based Target Search by Unmanned Aerial Vehicles

W. T. Luke Teacy
University of Southampton
Southampton, SO17 1BJ
United Kingdom
wltl@ecs.soton.ac.uk

Simon J. Julier
Renzo De Nardi
University College London
London, WC1E 6BT
United Kingdom
S.Julier@ucl.ac.uk
r.denardi@cs.ucl.ac.uk

Alex Rogers
Nicholas R. Jennings
University of Southampton
Southampton, SO17 1BJ
United Kingdom
{acr,nrj}@ecs.soton.ac.uk

ABSTRACT

Unmanned Aerial Vehicles (UAVs) are playing an increasing role in gathering information about objects on the ground. In particular, a key problem is to detect and classify objects from a sequence of camera images. However, existing systems typically adopt an idealised model of sensor observations, by assuming they are independent, and take the form of maximum likelihood predictions of an object's class. In contrast, real vision systems produce output that can be highly correlated and corrupted by noise. Therefore, traditional approaches can lead to inaccurate or overconfident results, which in turn lead to poor decisions about what to observe next to improve these predictions.

To address these issues, we develop a Gaussian Process based observation model that characterises the correlation between classifier outputs as a function of UAV position. We then use this to fuse classifier observations from a sequence of images and to plan the UAV's movements. In both real and simulated target search scenarios, we show that this can achieve a decrease in mean squared detection error of up to 66% relative to existing state-of-the-art methods.

Categories and Subject Descriptors

I.2.9 [Computing methodologies]: Artificial Intelligence—Robotics

General Terms

Algorithms, Theory, Experimentation

Keywords

Active Sensing; Target Search; Unmanned Aerial Vehicles; Gaussian Processes

1. INTRODUCTION

In many applications, such as environmental monitoring and disaster response, an aerial view is invaluable for gathering information about the situation on the ground. For example, in disaster response, there is a need to know the loca-

tion of civilians in need of rescue, and to assess the damage to vital infrastructure to aid logistical planning. Until recently, manned flight has been the only option for gathering such information quickly from the air. However, with recent advances in technology, small light-weight Unmanned Aerial Vehicles (UAVs) are increasingly offering a viable alternative, as they are relatively cheap to maintain; can fly through areas too small or too hazardous for manned flight; and potentially, can cover more ground by allowing a single operator to control multiple platforms simultaneously. However, current systems are labour intensive. Fatigue is a common problem for UAV operators [19], and increased autonomy has the potential to mitigate it by allowing the UAV to decide for itself what to observe [10] to maximise information gain.

Solving such *active sensing* problems typically requires three components: (1) an *observation* (or *measurement*) *model*; (2) a *predictive model*, which predicts the target state given sensor measurements; and (3) a *planner* to decide what actions to take to maximise the informative value of future observations. Unfortunately, the first of these components is over simplified in many existing systems. This, in turn, limits the ability of the other components in the system to perform well in real-world environments. This is particularly true in systems aimed at target search problems, such as search and rescue, in which three assumptions are prevalent [23, 27, 4, 24, 12]: (1) target detections take the form of binary responses, indicating that a target is either present or not present at some location; (2) the probability of misdetection and false positives is a constant irrespective of the relative pose between the target and the platform, or the type of terrain; and (3) successive detections are conditionally independent of one another given the target state.

However, these assumptions are rarely true in practice. Real computer vision systems are not limited to binary responses, instead they provide rich information in the form of class probabilities or scores, indicating how well an object's features match the target class; reliability *depends* on the relative camera position; and observations are, by nature, highly correlated, because an object's appearance will be similar between images taken close together in space and time. As a result, the above assumptions can lead to inaccurate or over confident results, especially when searching for stationary targets that blend into their environment, and whose appearance does not change significantly over time.

In contrast, the problem of fusing information from mul-

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.
Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

multiple dependent measurements has been widely studied for target localisation. In Simultaneous Localisation and Mapping (SLAM), arguably the closest problem to our one here, a platform attempts to construct a map of features in the environment and use it, at the same time, to localise the platform. The joint structure between all the features in the map and the vehicle is vital for the success of the approach [11]. Therefore, the full covariance matrix (or suitable encoding of it) must be maintained, and any algorithm which fuses information from multiple observations must take the dependency structure into account.

Here, however, our primary concern is not target localisation, but how such targets are *detected* in the first place. To date, dependencies between target detections have only been considered by Velez et al. [25]. However, their work is aimed at problems in which objects are detected by chance while on route to a primary goal destination. Moreover, although they do account for correlations between previous target detections, future detections are still treated as independent for the purpose of planning (see Section 3.3). As such, their work is not applicable to target search problems, in which a UAV must actively seek out previously undetected targets, taking into account the dependent nature of future observations. With this in mind, we propose a new approach for target search, that accounts for correlations and dependencies between both future and previously observed target detections. By so doing, we make the following three contributions to the state-of-the-art.

1. We present the first observation model for target search to account for correlations between classifier scores, both for previously observed target detections, and for predicting the information value of future observations.
2. We show how, in a tractable way, this model can be used to predict likely target locations, and provide a measure of information value to guide active sensing planners.
3. We demonstrate empirically that these features have a significant beneficial impact on performance, when used to select observations from a real dataset collected from a camera-equipped UAV, and in a simulated target search scenario.

In the next section, we introduce a motivating scenario to provide context for our work, along with our key modelling assumptions. Following this, we describe the model in detail and how it can be used for inference (Section 3); show how it can be trained using a real-world dataset (Section 4); evaluate its performance empirically against other models representative of the state-of-the-art (Section 5); and finally, we conclude and discuss future work in Section 6.

2. PROBLEM DEFINITION

We focus on finding an unknown number of static targets within a bounded search region on the ground, using a camera-equipped UAV. Although a human operator must specify the search region, both target detection and decisions about where to search are carried out autonomously by the UAV. Observations made by the UAV are then used to update a probability distribution in real-time, which both informs the UAV’s search pattern, and highlights likely target locations to the human operator for further investigation.

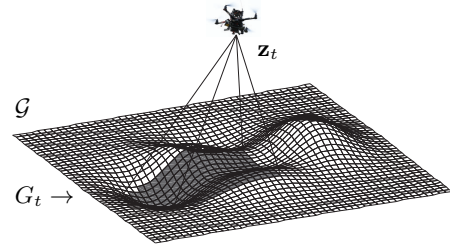


Figure 1: The UAV’s field of view, G_t , (shaded grey) determined by projecting from the UAV’s current position, $\mathbf{z}_t \in \mathcal{Z}$, onto the search region, \mathcal{G} .

We model the search region, \mathcal{G} , as a two-dimensional *occupancy grid* [27], in which each cell, $g \in \mathcal{G}$ may be either occupied by a target or not. Denoting the unknown *occupation* of cell g as δ_g , our goal is then to determine the probability that each cell is occupied ($\delta_g = 1$) or not ($\delta_g = 0$), and to continually refine these probabilities given visual observations made by the UAV. Although targets can occupy any set of continuous points on the ground, partitioning these points into a finite grid greatly simplifies the problem of maintaining probability distributions over possible target locations, and so is commonly adopted in the literature [27, 6, 4]. This does not require that the ground is flat, nor does it restrict the size of each grid cell. On the contrary, \mathcal{G} can map onto any two-dimensional manifold representing the topology of the ground; and grid cells can be sufficiently small for a single target to occupy multiple adjacent cells.¹ The grid scale may be selected to provide the required precision for any given application. For example, in search and rescue, it is clear that 1–2m² grid cells should be sufficiently small to accurately locate a missing person’s position.

To determine likely target locations as quickly and as accurately as possible, the UAV must choose its own flight path to maximise the information gained about target locations. Specifically, we assume that the UAV uses its onboard camera to capture a series of images of the ground at regular time intervals along its flight path. For each image, we assume that the UAV can determine its position, $\mathbf{z}_t \in \mathcal{Z}$, where $\mathcal{Z} \subset \mathbb{R}^3$ is a predefined flight zone bounded below by the ground, and $t \in \mathbb{Z}^+$ is the time at which the image was captured. Moreover, we assume that the UAV can determine the set of grid cells, $G_t \subset \mathcal{G}$, that are visible in each image, by projecting the camera field of view onto the ground, given its current position \mathbf{z}_t (see Figure 1).

In reality, neither G_t nor \mathbf{z}_t can be determined exactly due to imperfect sensors and incomplete information about the ground’s topography. However, we do not deal with this type of uncertainty explicitly in this paper, because this can be handled by existing techniques that can be readily incorporated into our model (see Section 6). Instead, we focus on the probability that observed cells are occupied, given the output of an onboard visual classifier. Specifically, for each $g \in G_t$ and for every t at which an image is captured, we receive a *score*, $s_g^t \in \mathbb{R}$, from a classifier, indicating the likelihood that g is occupied. For example, if the classifier is Bayesian [2] then s_g^t may be the probability that g is occu-

¹As a result, we cannot directly determine the number of targets from the number of occupied cells. However, our focus is on locating rather than counting targets.

ped, given its visual appearance from \mathbf{z}_t at time t ; whereas, if the classifier is a Support Vector Machine (SVM) [8], then s_g^t may be the distance between the cell’s visual features (as input to the classifier) and the class decision boundary (see Section 4).

Putting these together, we define an *observation* as a pair, $d_g^t = (s_g^t, \mathbf{z}_t)$, and denote the set of all such observations of a given cell, g , up to time t as $D_g^t = \{d_g^r | r \leq t\}$. Given this, our goal is to determine the posterior probability function (p.f.) for the occupation of each cell, $P(\delta_g | D_g^t)$; and to choose the UAV’s flight path so that the accuracy of this p.f. is maximised. The next section describes this in detail.

3. MODELLING AND INFERENCE

In this section, we describe how to model the distribution of the target locations given previously observed data, and use this to guide observation selection. This requires three components (see Section 1): the *observation model*, a *predictive model* and a *planner*. We discuss each in turn.

3.1 Predictive Modelling

The predictive model computes the posterior probability that each cell, g , is occupied. Using Bayes rule, this is calculated according to (1) as follows.

$$P(\delta_g | D_g^t) = \frac{P(\delta_g)P(D_g^t | \delta_g)}{P(D_g^t)} \quad (1)$$

Here, $P(D_g^t)$ is the normalising constant (i.e. the marginal likelihood) required to ensure that $\sum_{\delta_g \in \{0,1\}} P(\delta_g | D_g^t) = 1$; $P(\delta_g)$ is the prior probability that grid cell g is occupied; and $P(D_g^t | \delta_g)$ is the data likelihood, which is provided by the observation model described in Section 3.2.

To simplify this further, [6] not only assumes independence between cells, but also assumes independence between observations, and that instead of scores, we receive only the maximum likelihood class for each observation. However, there are two significant disadvantages to this approach. First, by observing only the predicted class with a fixed probability for false or missed detections, any differences in the confidence of each observation are ignored. This assumption is relaxed by some authors, including [23] and [12], who allow the probability of correct classification to vary with camera height. However, this still ignores other effects caused by visual appearance, such as target pose or the amount of background clutter.

Second, assuming independence between observations can lead to misleading and overconfident results (i.e. $P(\delta_g | D_g^t)$ close to 0 or 1, when this should be less certain). This is because each new observation is considered to be conditionally independent. However, as discussed in Section 1, images of the same scene taken at similar times or positions are highly likely to have the same classifications and are not conditionally independent. This is *not* the same as assuming independence between cells (as we do here): although, in reality, cells may be dependent (e.g. due to spatial correlations in terrain) this information can be safely ignored, provided there is not significant restriction on gathering information about cells by observing them directly.² In contrast, assuming independence between observations implies that more information is available, rather than less, which is in general a

²Dependencies between cells can be modelled [13] and are outside the scope of this paper.

much riskier strategy. One approach is to use suboptimal fusion rules, such as Generalized Covariance Intersection [1] or those based on Transferable Belief Models, Imprecise Probabilities or Random Sets [3]. These approaches often assume that no dependence information is available and, in consequence, the posteriors can be very uniform. In this paper, we seek a solution which can be used to model the *joint* data likelihood. By doing so, Bayes rule can be applied directly and rapid classification can be made.

3.2 Observation Modelling

We develop our observation model using a *Gaussian Process* (GP) [20] over classifier scores. Unlike previous approaches, we can take into account differences in classifier performance under different circumstances, and to model correlations between observations, so avoiding overconfident predictions due to inappropriate independence assumptions. Specifically, for any given cell, g , we model the classifier score, s_g^t , as a GP dependent on the UAV’s position at the time of observation:

$$s_g^t = f_{\delta_g}(\mathbf{z}_t) + \epsilon_\delta \quad (2)$$

$$f_{\delta_g}(\mathbf{z}_t) \sim \mathcal{GP}(\mu_{\delta_g}(\mathbf{z}_t), \mathbf{k}_{\delta_g}(\mathbf{z}_t, \mathbf{z}_t)) \quad (3)$$

$$\epsilon_\delta \sim \mathcal{N}(0, h_\delta) \quad (4)$$

Here, f_{δ_g} is a GP that captures the relationship between classifier scores and camera position, while ϵ_δ is zero-mean Gaussian noise, with variance, h_δ , which captures any additional variation between scores over time. These are defined separately for each possible class $\delta_g = \{0, 1\}$, since (in general) we expect the classifier to respond differently depending on whether a cell is occupied or not.

As with any GP, f_{δ_g} is specified by a *mean* function, μ_{δ_g} , and a *covariance* function, \mathbf{k}_{δ_g} . μ_{δ_g} is defined as the expected score, $E[s_g^t | \mathbf{z}_t, \delta_g]$, given the cell’s class and UAV’s position. It accounts for any expected difference in classifier performance dependent on camera position. As discussed above, we would expect it to vary with altitude, although it may also depend on horizontal position if we have any prior knowledge about the classifier’s performance in different parts of the environment, for example, due to different types of terrain. Similarly, \mathbf{k}_{δ_g} , is the *covariance* function between scores observed at any given pair of positions:

$$\mathbf{k}_{\delta_g}(\mathbf{z}_r, \mathbf{z}_t) = E[(s_g^r - \mu_{\delta_g}(\mathbf{z}_r))(s_g^t - \mu_{\delta_g}(\mathbf{z}_t))].$$

Again, this may depend on the cell’s class, and captures the expected similarity between observations of a given cell from different positions. Although this may take a variety of forms, we adopt the commonly used *squared exponential* function [20], which captures the intuition that the covariance between observed scores s_g^t and s_g^r will tend to increase as the distance between their corresponding camera positions decreases:

$$\mathbf{k}_\delta(\mathbf{z}_r, \mathbf{z}_t) = \sigma_\delta \exp \left[-\frac{1}{2} (\mathbf{z}_t - \mathbf{z}_r)^T L_\delta^{-1} (\mathbf{z}_t - \mathbf{z}_r) \right], \quad (5)$$

where σ_δ is the magnitude, which determines the maximum covariance between scores observed at different times; and L_δ is a diagonal *length scale* matrix, which determines how quickly the covariance decreases with distance along each dimension of \mathbf{z}_t . Together with the mean function and noise variance, h_δ , these parameters can be selected to best fit the performance of an actual classifier (see Section 4).

According to this specification, the joint data likelihood for a given cell is then defined as a joint normal distribution as follows.

$$P(D_g^t | \delta_g = \delta) = \mathcal{N}(S_g^t \boldsymbol{\mu}_\delta(Z_g^t), h_\delta I + K_\delta(Z_g^t)) \quad (6)$$

where the mean and covariance are specified as

$$\boldsymbol{\mu}_\delta(Z_g^t) = \begin{bmatrix} \mu_\delta(\mathbf{z}_g^r) \\ \vdots \\ \mu_\delta(\mathbf{z}_g^t) \end{bmatrix} \quad (7)$$

$$K_\delta(Z_g^t) = \begin{bmatrix} \mathbf{k}_\delta(\mathbf{z}_g^r, \mathbf{z}_g^r) & \cdots & \mathbf{k}_\delta(\mathbf{z}_g^r, \mathbf{z}_g^t) \\ \vdots & \ddots & \vdots \\ \mathbf{k}_\delta(\mathbf{z}_g^t, \mathbf{z}_g^r) & \cdots & \mathbf{k}_\delta(\mathbf{z}_g^t, \mathbf{z}_g^t) \end{bmatrix} \quad (8)$$

By substituting (6) into (1), we are thus able to calculate the posterior probability that any given grid cell is occupied, without assuming independence between observations. In the following sections, we demonstrate that this not only leads to more accurate results compared to existing techniques, but also enables better decisions about what to observe next to further improve predictions.

3.3 Measuring Information Value

At this point, we now have a complete model for predicting likely target locations given previous observations. However, to actively search for targets, we need to measure how much different candidate observations will improve predictions about target locations. We do not focus on how candidate observations are proposed. For this, we assume that some suitable planner is available [21, 22, 5], which can propose observations to be collected along potential flight paths. Instead, we provide a measure, based on our model, so that a planner can choose between candidate paths. Existing approaches generally focus on maximising the probability of detection [12, 24], which incentivises the UAV to observe cells with high probability of occupation, even if we are sure these are occupied. However, we seek to minimise the overall uncertainty in the presence and location of targets. This is directly measured by *mutual information* [15].

We define the path planning problem as follows. First, we denote the set of all previous observations for all grid cells up to time t as $D^{-t} = \{d_g^r | g \in \mathcal{G}, r \leq t\}$, and the set of all future observations up to some finite planning horizon t' as $D^{+t} = \{d_g^r | g \in \mathcal{G}, t < r \leq t'\}$. Our objective is to choose a flight path that maximises the mutual information between the future observations, D^{+t} , and the unknown cell occupations, $\Delta = \{\delta_g | g \in \mathcal{G}\}$, given the previous observations, D^{-t} . Following [18], this is defined as the expected reduction in entropy for Δ after observing D^{+t} ,

$$I(\Delta; D^{+t} | D^{-t}) = H(\Delta | D^{-t}) - H(\Delta | D^{-t}, D^{+t}) \quad (9)$$

Under our assumption that cell occupations are independent, this can be decomposed in terms of individual cells according to Equation (10), in which the conditional entropies, $H(\delta_g | D_g^t)$ and $H(\delta_g | D_g^{t'})$, are calculated w.r.t. the posterior distribution in Equation (1).

$$I(\Delta; D^{+t} | D^{-t}) = \sum_{g \in \mathcal{G}} H(\delta_g | D_g^t) - H(\delta_g | D_g^{t'}) \quad (10)$$

Importantly, this equation calculates the informative value of a set of candidate observations *as a whole*, taking into

account the correlations between them, as yet, unobserved values as a consequence of our model. In this respect, our approach differs from that of Velez et al. [25]. Although their model accounts for spatial correlations between previously observed values, they estimate the information value of future observations using an approximate technique based on the Posterior Belief Distribution algorithm [14]. This technique uses a Kalman filter, based on the assumption that future observations *are* independent, even though this is later revealed to be false once they have been observed. As such, planners using this technique will overestimate the value of candidate trajectories, since each new observation is assumed to provide completely new information, uncorrelated with any other future observations collected on route.

In contrast, we calculate mutual information directly based on our model, and in this way, account for correlations between both previous and future observations. Given this, we can now measure the information value for any set of candidate observations, which can then be used by a planner to select the the most informative trajectory for a UAV. In the following sections, we apply and evaluate this in both simulation, and with real data.

4. MODEL AND CLASSIFIER TRAINING

To put this theoretical framework into practice, we applied it to a search and rescue scenario, using images collected from a quadrotor UAV. This required the following three steps.

Data Collection

A UAV fitted with a downward facing camera captured images of people in different poses in different types of wilderness environments (peatland, rough pastoral land, and rocky coastal areas) at 5 frames per second at altitudes between 1m and 40m (see Figure 2). We registered these images to their corresponding camera and ground positions using Visual Structure From Motion [28], which could estimate the position of known landmarks within 5m. Using a grid of 2×2 metre cells, this meant that observations of fixed points could be wrongly spread across several adjacent cells.

Although this breaks our simplifying assumption that the localisation of observed points on the ground can be established perfectly, such inaccuracy is to be expected in real applications. Moreover, as discussed previously, this can readily be dealt with using existing techniques that can be easily integrated into our model. Nevertheless, as demonstrated in Section 5, our model still outperforms existing approaches despite being applied without any mechanism to deal explicitly with this uncertainty.

With this in mind, we classified any grid cell as occupied, provided it contained the estimated centre position of any person on the ground. The entire dataset of around 10,000 images was then randomly divided into three equally sized subsets, with no overlapping images: one to train a classifier to detect people, a second to train our observation model, and a third for empirical evaluation (see Section 5.2).

Classifier Training

To detect targets on the ground, we used an SVM classifier trained to detect people using Histograms of Oriented Gradients (HOG) features. Although, our model is agnostic to the choice of classifier, the combination of these two techniques has demonstrated good performance for people



Figure 2: Example images collected from quadrotor for model training and evaluation: (top row) original images, (middle row) negative cases used to train classifier, (bottom row) positive cases used to train classifier.

detection in other domains [9], and so is representative of the type of system that could be used in practice. To apply this to our scenario, we trained the classifier using images of people and general background (see Figure 2) extracted from the classifier training set mentioned above, and used this to generate classifier scores for all grid cells visible in the remaining images.

Observation Model Training

To train our observation model, we first used the training data to fit the mean function, μ_δ , to the mean classifier scores observed for each class (cell occupied or not). In general, this can be viewed as a regression problem, in which we must estimate the score given the UAV’s altitude. For example, this may be solved by fitting a separate Gaussian Process to the training data, and using this to predict the expected value of any potential observation as a function of altitude. However, given the maximum safe altitude of our UAV was 40m, we found no significant dependence between altitude and classifier scores in our dataset. Although we do not expect this to hold for higher altitudes, for our experiments it was therefore sufficient to set μ_δ to the mean score in the training set for each class (person present, or not present).

Given these means, we then set the noise and covariance parameters to their maximum likelihood values across the training set. In particular, these values result in a significant spatial correlation of up to 0.42 for observations taken 1 metre apart, and 0.14 at 5 metres apart. Incorporating this correlation in the observation model is important, since assuming zero correlation would result in overconfident results (see Section 3). In the next section, we demonstrate this by comparing the empirical performance of our model to others, representative of the state-of-the-art.

5. EMPIRICAL EVALUATION

For our initial experiments, we evaluated the performance in

two different scenarios: (1) a simulated on-line setting, and (2) an off-line setting, using real data from the evaluation set described in Section 4. In both cases, we measured performance using the Mean Squared Detection Error (MSDE), as it is standard in many estimation problems [7]. This is calculated by estimating the probability of occupation of all cells in \mathcal{G} , and then calculating the MSDE between these estimates, and each cell’s true class:

$$\text{MSDE} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} [\delta_g - P(\delta_g = 1 | D_g^t)]^2 \quad (11)$$

For comparison, this was repeated for all experiments, using three different models. The first is *correlated*, which denotes our proposed model, taking into account correlations between classifier scores. The second is *independent*, which uses Naïve Bayes on the raw classifier scores. It is identical to correlated, but assumes zero covariance (independence) between scores. The final is *class only*, which also adopts a Naïve Bayes approach, but only receives binary maximum likelihood classifications for each observation, and is the prevailing approach adopted in the existing literature (including [12], [23], [6]).

This uses a fixed confusion matrix, learned from the training data, to specify the likelihood of the observed class (target present or not present), given the true class. Based on these assumptions, all three models use Bayes rule (1) to fuse all available observations into a single probability of target presence for each cell. In the following subsections, we describe the results of these experiments in both on-line and off-line settings.

5.1 On-line Selection in Simulation

In these experiments, the aim is to simulate an on-line active sensing scenario in which a UAV must choose its flight path to maximise the mutual information between the resulting observations and the unknown target locations. For this purpose, we simulated a 100×100 grid of which 10–200 cells were randomly selected to be occupied by a target in each episode. When observed, each cell generates a classifier score that is correlated with its previous observations according to the same type of GP used in our observation model. However, for the purpose of these experiments, the mean score was fixed at -1 for unoccupied cells, and 1 for occupied cells. The level of difficulty was then controlled by varying the covariance parameters between each experiment. As a result, the optimal decision boundary used for the *class only* model was fixed at zero, while the probability of a false detection was determined by the magnitude of the score variance, relative to their mean.

To observe this grid, we simulated a UAV with a simple myopic planner,³ which explores in a sweep search pattern, observing 8×8 blocks of cells at a time. In each timestep, the UAV can choose to either move onto the next block, or observe the current block from possible altitudes of 1, 2, 4 or 8 metres. In each case, a change of altitude would bring about a corresponding change in the field of view. Specifically, at 8 metres, the UAV could observe all 64 cells in the current block simultaneously, while at low-altitudes, sub-blocks of only 1×1 , 2×2 or 4×4 cells could be observed at once. This introduces a trade-off between flying high to observe more cells at once, or flying low to observe fewer

³Other planners which could be used include [17, 5].

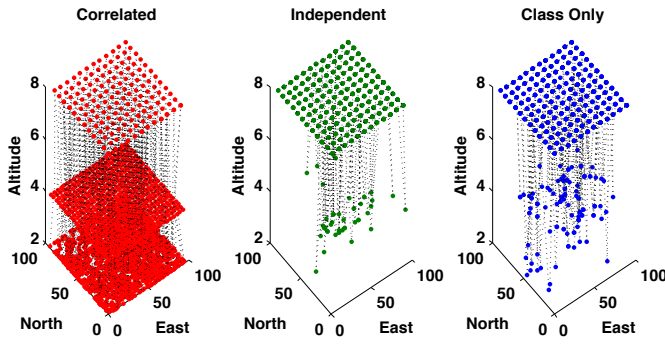


Figure 3: UAV trajectories in the simulated scenarios, shown as coloured dots at visited locations connected by dashed lines. Notice that, relative to the other two approaches, the correlated model covers more unique positions (at different altitudes) to gain more information.

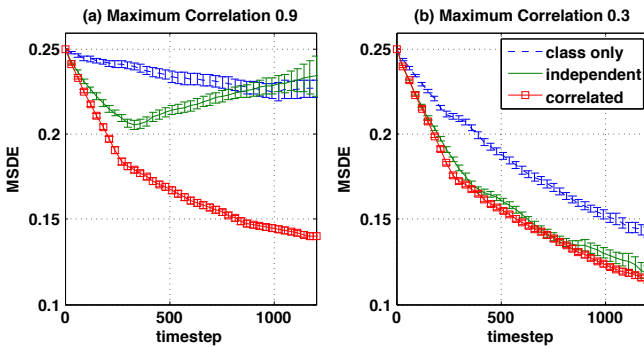


Figure 4: MSDE in probability of target presence for simulated active sensing scenario.

cells from a significantly different perspective, while keeping them within the field of view. This is because choosing observations from significantly different perspectives generally results in less correlated observations, that thus result in high information gain — even though the field of view is reduced at lower altitudes.

Figure 4 shows the results for these experiments in two different cases, with error bars representing 95% confidence intervals for the MSDE. In each case, the covariance length scale was fixed⁴ at $L_\delta = I$, while h_δ and σ_δ were chosen to fix the prior variance at 3.61. The difference is that for plot (a) $h_\delta = 0.361$, resulting in a maximum correlation between scores of 0.9; while for plot (b) $h_\delta = 2.527$, resulting in a maximum correlation of 0.3. For plot (a) this means that little new information can be gained from observing a cell from the same position, as this would result in very similar observations; whereas for plot (b) a significant amount of information can be gained by repeated observation from the same location.

The key effect is that both Naïve Bayes approaches perform poorly on highly correlated scores for two reasons: (1) by assuming independence, they over estimate the information content of each new observation; and (2) they fail to realise the benefit of changing altitude to gain a different

⁴In other experiments, not shown here, we tried different parameter values, but observed similar results.

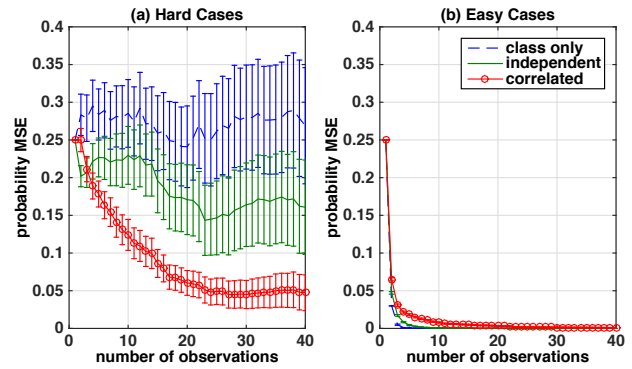


Figure 5: MSDE in probability of target presence predicted for (a) hard cases, and (b) easy cases, selected from real dataset.

perspective and thus less correlated observations. In order to maximise their field of view, they instead adopt a sweep search pattern at the maximum altitude (see Figure 3), only dropping to a lower altitude when they can detect no significant difference in mutual information. By accounting for observed scores, the *independent* model is still able to outperform the *class only* model. However, this advantage disappears as the number of observations grows, due to over-confidence in its results.

In contrast, the *correlated* model dominates in all cases, since it correctly estimates the mutual information between observations. Moreover, although it initially chooses to fly high to maximise field of view, it subsequently reduces altitude, to gain additional information about each subblock of cells, prioritising those with high class uncertainty. This advantage is less profound when correlation between scores is low, but even in these cases, it does no worse than the other evaluated methods.

5.2 Off-line Selection of Real Data

Although the results above demonstrate the potential benefits of our model for on-line active sensing problems, they do not guarantee that the same benefits apply to problems involving real data. For this reason, we also compared the performance of all three models when applied to the real data described in Section 4. However, since this data was collected in advance, the UAV’s flight path was predetermined and could not be changed. Instead, the goal in these experiments was to make predictions by selecting observations (of grid cells) from within images captured along the UAV’s flight path. To maximise information gain, this is done by selecting individual observations in order of their mutual information with each grid cell’s hidden occupation. Although this is a slightly different problem, this scenario may still occur in practice in cases in which the flight path must be predetermined (e.g. due to legal constraints). In such cases, selecting only the most informative observations can still be an advantage, due to the significant computational cost of processing and classifying every part of every image.

With this in mind, Figure 5 shows the results of these experiments. In particular, part (a) shows the results for *hard* cases, which we define as any cell for which any of the three models predicted the wrong class as most likely

at least once; while (b) shows the results for *easy* cases, for which all three models always predicted the correct class as most likely. This shows that, for easy cases, the choice of model has little effect: on average, all three models have near zero error with as few as 10 observations. However, our model does take slightly longer to approach zero, because we take a more conservative view on the amount of new information provided by each observation. That is, although all three models always report the correct class as most likely in these cases, the correlated model places a slightly lower probability on the correct class, because it avoids unrealistic independence assumptions and thus requires more evidence to report a given level of confidence in its predictions.

Although such easy cases make up the majority of our dataset (e.g. grass and other natural features that do not resemble a person) what matters more are the hard cases (e.g. people in dense clutter, or linear features that could be mistaken for a person). In these cases, an accurate assessment of uncertainty is essential for both deciding what to observe next to improve predictions, and for highlighting areas of uncertainty to the UAV's operators. In this respect, the *independent* model generally performs better than *class only*. Although not statistically significant, this is inline with the simulated experiments, which do suggest some benefit in observing scores, rather than classes only. However, as before, our model significantly outperforms both benchmarks, by also taking into account correlations between these scores. Unlike the *independent* benchmark, this avoids placing too much weight in observations made from similar camera positions, which can otherwise lead to misleading and overconfident results. As a result, our model is able to achieve a decrease in MSDE for hard cases of up to 66% relative to the *independent* and *class only* models.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we have highlighted the importance of adopting realistic observation models in active sensing problems, and in particular, target search applications using an autonomous camera-equipped UAV. In such applications, previous authors have typically assumed that observations are always independent, and take the form of binary classifications, specifying whether a target is present at a given location or not. Unfortunately, neither of these assumptions are true in practice: real vision systems often produce valuable information in the form of scores, indicating the likelihood that a target is present given a position's visual appearance. Moreover, such observations typically exhibit strong spatial correlations, since viewing an object from the same camera position will typically produce the same classification.

To address these limitations, we have presented a novel observation model, that achieves up to 66% greater accuracy than existing approaches at detecting targets. This is achieved by modelling spatial correlations between classifier scores, which enables more accurate assessment of the information value of each observation, and better decisions about what to observe next to improve future predictions.

Although these results represent a significant first step, there are three main aspects that we are yet to address. First, although we have considered the affect of camera position, we have not explicitly considered the impact of other factors, such as occlusions, moving targets and lighting conditions. Second, by focusing on classifier uncertainty, we have not addressed the additional impact of position uncer-

tainty w.r.t. both the camera and its field of view. Finally, although we have demonstrated the benefits of our model in both a simulated on-line setting, and an off-line setting using data collected from a real UAV, we are yet to apply the model to autonomous on-line path planning using a real UAV. In future work, we plan to address these limitations, by integrating our techniques with existing state-of-the-art approaches for modelling positional uncertainty [16, 26]. We then plan to use these techniques for active sensing with multiple UAVs, working together in coordination to locate targets on the ground.

Acknowledgements

This work was funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under the following project grants: Distributed sensing, control and decision making in multiagent autonomous systems (EP/J011894/2), the SAAVE project (EP/F06358X/1 and EP/F064179/1), and the ORCHID project (EP/I011587/1). The authors also acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

REFERENCES

- [1] T. Bailey, S. Julier, and G. Agamennoni. On conservative fusion of information with unknown non-gaussian dependence. In *Proceedings of the 15th International Conference on Information Fusion*, pages 1876–1883, 2012.
- [2] D. Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- [3] A. Benavoli and B. Ristic. Classification with imprecise likelihoods: A comparison of tbm, random set and imprecise probability approach. In *Proceedings of the 14th International Conference on Information Fusion*, pages 1–8, 2011.
- [4] F. Bourgault, T. Furukawa, and H. Durrant-Whyte. Coordinated decentralized search for a lost target in a bayesian world. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 48–53, 2003.
- [5] C. P. C. Chanel, F. Teichteil-Königsbuch, and C. Lesire. Multi-target detection and recognition by uavs using online pomdps. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 1381–1387, 2013.
- [6] T. Chung and J. Burdick. A decision-making framework for control strategies in probabilistic search. In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, pages 4386–4393, 2007.
- [7] P. R. Cohen. *Empirical Methods for Artificial Intelligence*. M.I.T. Press, 1995.
- [8] N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press, 2000.
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 886–893, 2005.
- [10] M. E. Dempsey. Eyes of the army. u. s. army roadmap

- for unmanned aircraft systems 2010–2035. Technical report, U. S. Army UAS Center of Excellence, 2010.
- [11] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: Part i. *IEEE Robotics and Automation Magazine*, 13(2):99–110, June 2006.
- [12] S. K. Gan and S. Sukkarieh. Multi-uav target search using explicit decentralized gradient-based negotiation. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation*, pages 751–756, 2011.
- [13] A. E. Gelfand, P. J. Diggle, M. Fuentes, and P. Guttorp, editors. *Handbook of Spatial Statistics*. Chapman & Hall, 2010.
- [14] R. He, E. Brunskill, and N. Roy. Efficient planning under uncertainty with macro-actions. *Journal of Artificial Intelligence Research*, 40:523–570, 2011.
- [15] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.
- [16] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g^2o : A general framework for graph optimization. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation*, pages 3607–3613, 2011.
- [17] D. Levine, B. Luders, and J. How. Information-Rich Path Planning with General Constraints Using Rapidly-Exploring Random Trees. In *AIAA Infotech@Aerospace*, 2010.
- [18] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [19] J. A. Ouma, W. L. Chappelle, and A. Salinas. Facets of occupational burnout among u.s. air force active duty and national guard/reserve mq-1 predator and mq-9 reaper operators. Technical report, Air Force Research Laboratory, 2011.
- [20] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [21] P. Scerri, R. Grinton, S. Owens, D. Scerri, and K. Sycara. Geolocation of RF Emitters by Many UAVs. In *AIAA Infotech@Aerospace*, 2007.
- [22] M. T. J. Spaan. Cooperative active perception using pomdps. In *Proceedings of the AAAI 2008 Workshop on Advancements in POMDP Solvers*, 2008.
- [23] A. Symington, S. Waharte, S. Julier, and N. Trigoni. Probabilistic target detection by camera-equipped uavs. In *Proceedings of the 2010 IEEE International Conference on Robotics and Automation*, pages 4076–4081, 2010.
- [24] J. Tisdale, A. Ryan, Z. Kim, D. Tornqvist, and J. Hedrick. A multiple uav system for vision-based search and localization. In *Proceedings of the 2008 American Control Conference*, pages 1985–1990, 2008.
- [25] J. Velez, G. Hemann, A. S. Huang, I. Posner, and N. Roy. Modelling observation correlations for active exploration and robust object detection. *Journal of Artificial Intelligence Research*, 44:423–453, 2012.
- [26] B.-N. Vo and W.-K. Ma. The Gaussian Mixture Probability Hypothesis Density Filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006.
- [27] S. Waharte, A. Symington, and N. Trigoni. Probabilistic search with agile uavs. In *Proceedings of the 2010 IEEE International Conference on Robotics and Automation*, pages 2840–2845, 2010.
- [28] C. Wu. Towards linear-time incremental structure from motion. In *Proceedings of the 2013 International Conference on 3D Vision*, pages 127–134, 2013.