

Cost-Aware Bayesian Sequential Decision-Making for Domain Search and Object Classification

Y. Wang*, I. I. Hussein*, D. R. Brown III[†], and R. S. Erwin[‡]

Abstract—This paper focuses on the development of a cost-aware Bayesian sequential decision-making strategy for the search and classification of multiple unknown objects over a given domain using a sensor with limited sensory capability. Under such scenario, it is risky to allocate all the available sensing resources at a single location of interest, while ignoring other regions in the domain that may contain more critical objects. On the other hand, for the sake of finding and classifying more objects elsewhere, making a decision regarding object existence or its property based on insufficient observations may result in miss-detecting or miss-classifying a critical object of interest. Therefore, a decision-making strategy that balances the desired decision accuracy and tolerable risks/costs is highly motivated. The strategy developed in this paper seeks to find and classify all unknown objects within the domain with minimum risk under limited resources.

I. INTRODUCTION

In many domain search and object classification problems, the effective management of sensing resources is key to mission success. In a search task, the objective is to find every unknown object in a domain and fix its position. In a classification task, the objective is to take enough measurements to determine the nature of the object. On one hand, a sensor may give a false alarm while there is actually none, or may miss detecting a critical object. Similarly, the sensor might report incorrect classification results. On the other hand, taking exhaustive observations at one particular location may miss the opportunity to find and classify possibly more critical objects. This is especially true when the mission domain is large-scale, or the number of objects is far more than that of sensors [1]–[3]. Under these scenarios, a sensor has to decide whether to move and look for other objects, or stop and keep taking observations at a specific location. To accomplish these competing tasks with minimum risks under limited sensory resources, there is a strong motivation to develop a real-time decision-making strategy that chooses the task to perform based on an overall risk assessment associated with the decision. This is the problem addressed in this paper.

We first review some related literature. Coordinated search and tracking in probabilistic frameworks has been studied mainly for optimal path planning. In [4], the

authors investigate the search-and-tracking problem using recursive Bayesian filtering with foreknown targets' positions with noise. The results are extended in [5] for dynamic search spaces based on forward reachable set analysis. In [6], the author proposes a Bayesian-based multisensor-multitarget sensor management scheme. The approximation strategy maximizes the square of the expected number of targets. In [7], the problem of finding a target with some foreknown location information in the presence of uncertainty and limited communication channels is discussed. The probability of target existence is defined as a cost function to determine the vehicle's optimal path.

It is worth noting that in the above literature there is no explicit decision-making strategy for search and tracking/classification. To remedy this, the authors developed a deterministic decision-making strategy for search versus tracking in [1], and extended the results to a probabilistic framework in [2] and [3]. The deterministic strategy proposed in [1] guarantees the detection of all objects and the tracking of each object's state for a minimum amount of time τ_c . In [2] and [3], a probabilistic Bayesian version of this algorithm was developed for unknown object search versus classification as two competing tasks. In this work, we focus on a cost-aware decision-making strategy for search and classification under the probabilistic framework. A Bayesian sequential decision-making strategy is proposed to minimize the Bayes risk (to be formally defined in Section IV), which is based on both the error probabilities of the decisions made and the observation cost. The proposed strategy guarantees a desired detection and classification uncertainty level everywhere within the domain with minimum risks.

Sequential detection [8] allows the number of observation samples to vary in order to achieve an optimal decision. Due to the randomness of observations at each time step, a decision may be made with a few observation samples, whereas for other cases one would rather take more samples for a possibly better decision. The Bayesian sequential detection method used in this paper is such that the Bayes risk is minimized at each time step.

The paper is organized as follows. We introduce the sensor model in Section II. In Section III, the Bayes updates are developed. We propose the cost-aware Bayesian sequential decision-making strategy in Section III and provide a simulation for a single cell in Section V. In Section VI, uncertainty maps are built over the entire domain, which are used to define search and classification metrics.

*Mechanical Engineering Department, [†]Electrical and Computer Engineering Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609 {yuewang, ihussein, drb}@wpi.edu, [‡]Air Force Research Laboratory, Space Vehicles Directorate AFRL/RV, 3550 Aberdeen Ave, SE, Kirtland Air Force Base, NM 87117.

The search metric is related to a dynamic observation cost used for the decision-making strategy. In Section VII, a sensor motion control strategy is developed for multi-cell domains and a full-scale simulation is presented. We conclude the paper with a summary of our current and future work in Section VIII.

II. SETUP AND SENSOR MODEL

A. Problem Setup

Let $\mathcal{D} \subset \mathbb{R}^2$ be the domain in which objects to be found and classified are located. We discretize the domain into N_{tot} cells. Let \tilde{c} be an arbitrary cell in \mathcal{D} and \tilde{q} is the centroid of cell \tilde{c} . Define $1 \leq N_o \leq N_{\text{tot}}$ as the total number of objects and \mathbf{p}_j as the position of object j , both of which are unknown beforehand. Assume that the objects are i.i.d distributed over \mathcal{D} , and the partition of the domain is fine enough so that at most one object can exist in a cell. Without loss of generality, we assume that an object can have one of two properties, either Property ‘F’ or Property ‘G’. Let \mathcal{S}_F be the set of all cells containing an object having Property ‘F’ and \mathcal{S}_G be the set of all cells containing an object having Property ‘G’. Let P_F and P_G be the initial probability of a cell having an object with Property ‘F’ and Property ‘G’, respectively. Hence, N_o is a binomial random variable with parameters N_{tot} and $P_p = P_F + P_G$, where P_p gives the initial total probability of object present at a cell \tilde{c} and is i.i.d for all $\tilde{c} \in \mathcal{D}$. The expected number of objects in \mathcal{D} is given by

$$E[N_o] = N_{\text{tot}} P_p. \quad (1)$$

Let $X(\tilde{c})$ be a ternary state random variable, where 0 corresponds to object absent, 1 corresponds to object having Property ‘F’, and ‘2’ corresponds to object having Property ‘G’. Note that the realization of $X(\tilde{c})$ depends on the cell being observed:

$$X(\tilde{c}) = \begin{cases} 1 & \tilde{c} \in \mathcal{S}_F \\ 2 & \tilde{c} \in \mathcal{S}_G \\ 0 & \text{otherwise} \end{cases}.$$

Since both \mathcal{S}_F and \mathcal{S}_G are unknown and random, $X(\tilde{c})$ is a random variable with respect to every $\tilde{c} \in \mathcal{D}$. This paper focuses on the case where objects are immobile, therefore, $X(\tilde{c})$ is invariant with respect to time.

B. Sensor Model

In this work, we assume a sensor is able to observe only one cell at a time. Other sensor models that are capable of observing multiple cells at the same time (e.g., the sensor models with limited sensory range proposed in [1]–[3], [9]–[16]) can be used. We consider the extreme case in which the resources available are at a minimum (a single sensor as opposed to multiple cooperating ones).

To be consistent with the state model, we define a ternary observation random variable $Y(\tilde{c})$, where 0 corresponds to an observation indicating object absent at cell \tilde{c} , 1 corresponds to an observation indicating that there exists an object having Property ‘F’, and 2 corresponds to an observation indicating an object having Property ‘G’.

Given a state $X(\tilde{c}) = i$, $i = 0, 1, 2$, the probability mass function f of the observation distribution is given

by

$$f_Y(y|X(\tilde{c}) = i) = \begin{cases} \beta_{i0} & \text{if } y = 0 \\ \beta_{i1} & \text{if } y = 1 \\ \beta_{i2} & \text{if } y = 2 \end{cases}, \quad (2)$$

where $\sum_{j=0}^2 \beta_{ij} = 1$, Y corresponds to the ternary random variable and y is the dummy variable. Because the states $X(\tilde{c})$ are spatially i.i.d., the observations $Y(\tilde{c})$ taken at every cell \tilde{c} within the mission domain \mathcal{D} are spatially i.i.d. and hence the probability distribution for every $\tilde{c} \in \mathcal{D}$ follows the same structure.

Conditioned on the actual state $X(\tilde{c})$ at a particular cell \tilde{c} , let t be the time index, the observations $Y_t(\tilde{c})$ taken along time are temporally i.i.d. Let $Z_0(\tilde{c})$, $Z_1(\tilde{c})$, and $Z_2(\tilde{c})$ be the number of times that observation $Y(\tilde{c}) = 0, 1$, and 2, respectively, appears during a window of L time steps. The quantities $Z_0(\tilde{c})$, $Z_1(\tilde{c})$, and $Z_2(\tilde{c})$ are integer random variables that satisfy $\sum_{k=0}^2 Z_k(\tilde{c}) = L$, $Z_k(\tilde{c}) \in [0, L]$. Therefore, given an actual state $X(\tilde{c}) = i$, the probability of having observation z_0, z_1, z_2 in a window of L time steps follows a multinomial distribution

$$\begin{aligned} & \text{Prob}(Z_0(\tilde{c}) = z_0, Z_1(\tilde{c}) = z_1, Z_2(\tilde{c}) = z_2 | X(\tilde{c}) = i) \\ &= \frac{L!}{z_0! z_1! z_2!} \beta_{i0}^{z_0} \beta_{i1}^{z_1} \beta_{i2}^{z_2}, \quad \sum_{k=0}^2 z_k = L. \end{aligned} \quad (3)$$

The sensor probabilities of making a correct observation are β_{00}, β_{11} and β_{22} . For the sake of simplicity, here we assume that the values are some constants greater than $\frac{1}{2}$. For the sensor’s probability of making an erroneous observation β_{ij} , $i \neq j$, we use a simple linear combination model $\beta_{ij} = \lambda_j(1 - \beta_{ii})$, $i \neq j$, $\sum_{j=0}^2 \beta_{ij} = 1$, where λ_j is some weighting parameter that satisfies $\sum_{j \neq i} \lambda_j = 1$, $0 \leq \lambda_j \leq 1$.

III. BAYES UPDATES FOR SEARCH & CLASSIFICATION

Based on the sensor model, in this section, we employ Bayes’ rule to update the probability of object absence, and its classification property at a single cell \tilde{c} . Under the i.i.d. assumption, the Bayesian updates equations developed in this section will be deployed to the multi-cell domain in Section VI and VII.

Given a single observation $Y_t(\tilde{c}) = j$, $j = 0, 1, 2$ at time step t , according to Bayes’ rule, for each \tilde{c} , we have $P(X(\tilde{c}) = i | Y_t(\tilde{c}) = j; t + 1) = \alpha_j \beta_{ij} P(X(\tilde{c}) = i; t)$, (4) where $P(X(\tilde{c}) = i | Y_t(\tilde{c}) = j; t + 1)$, $i = 0, 1, 2$ is the posterior probability of the actual state being $X(\tilde{c}) = i$ at time step $t + 1$, $P(X(\tilde{c}) = i; t)$ is the prior probability of being state type $X(\tilde{c}) = i$ at t , and α_j serves as a normalizing function that ensures that the posterior probabilities $\sum_{i=0}^2 P(X(\tilde{c}) = i | Y_t(\tilde{c}) = j; t + 1) = 1$.

IV. COST-AWARE BAYESIAN SEQUENTIAL DECISION-MAKING

In this section, we will extend the standard binary Bayesian sequential detection method [8], [17], [18] from signal detection theory [8], [19]–[21] into a ternary cost-aware Bayesian sequential decision-making strategy.

Assuming a Uniform Cost Assignment (UCA) [8], we define the decision cost components as $C_{ij} = 1$ if $i \neq j$, and $C_{ii} = 0$. Here, C_{ij} is the cost of deciding i when the state is j . Let $i = 0, 1, 2$ represent ‘deciding object absent’, ‘deciding object having Property ‘F’’, and ‘deciding object having Property ‘G’’, respectively, and j corresponds to state $X(\tilde{c}) = j$.

Let $\tilde{R}_0(\tilde{c}, L, \Delta)$, $L \geq 1$, be the conditional Bayes risk of deciding there is an object having Property ‘F’ or ‘G’ at \tilde{c} given that there is actually none over at least one observation,

$$\tilde{R}_0(\tilde{c}, L, \Delta) = c_0 \Delta b_0, \quad c_0 = [C_{00} \ C_{10} \ C_{20}], \quad (5)$$

where c_0 contains the costs of deciding object absent, having Property ‘F’ and ‘G’ when there is actually nothing at \tilde{c} . The quantity b_0 is the first column of the general conditional probability matrix B for $L \geq 1$. The dimension of B is $N \times 3$. The number 3 is the number of possible decisions. The quantity N is the total number of different observation combinations (z_0, z_1, z_2) that the sensor can take according to the multinomial distribution (3) over a window of L time steps. The element B_{ij} gives the probability of having the i th kind of observation combination out of N given state j . Note that $\sum_{i=0}^{N-1} B_{ij} = 1$. The quantity Δ is a deterministic decision rule. For $L \geq 1$, Δ is a $3 \times N$ matrix. The matrix element Δ_i^n can be either 0 or 1, and $\sum_{i=0}^2 \Delta_i^n = 1$. When $\Delta_i^n = 1$, it means decision i is made given that the observation type is in the n th column. When $L = 0$, i.e., there are no observations taken, decisions will be made regardless of observations and there is no explicit matrix form for Δ .

Similarly, the conditional Bayes risk $\tilde{R}_i(\tilde{c}, L, \Delta) = c_i \Delta b_i$, $c_i = [C_{0i} \ C_{1i} \ C_{2i}]$, $i = 1, 2$, (6) gives the cost of making an erroneous decision at \tilde{c} given that the actual state is either $X(\tilde{c}) = 1$ or $X(\tilde{c}) = 2$ over $L \geq 1$ observations.

Therefore, under UCA, there is no cost if the decision is the actual state, and the conditional Bayes risk $\tilde{R}_0, \tilde{R}_1, \tilde{R}_2$ can be interpreted as the error probability of making a wrong decision under a certain decision rule Δ over L observations at cell \tilde{c} .

Now let us assign an observation cost c_{obs} each time the sensor makes a new observation. In this section, we assume it is a constant. In Section VI, a dynamic $c_{\text{obs}}(t)$ is developed to relate the observation cost with the task metrics for multi-cell domains.

For each cell at every time step, the sensor has to choose among: (i) deciding object absent, (ii) deciding object having Property ‘F’, (iii) deciding object having Property ‘G’, or (iv) taking one more observation. This same decision procedure is repeated until the cost of making a wrong decision based on the current observation is less than that of taking one more observation for a possibly better decision. The cost-aware Bayesian sequential decision-making strategy is such that the Bayes risk at each time step is minimized. Let $\phi = \{\phi_k\}_{k=0}^{\infty}$ be the stopping rule. If $\phi_k = 0$, the sensor takes another

measurement, if $\phi_k = 1$, the sensor stops taking further observations. Define the stopping time as $N(\phi) = \min\{k : \phi_k = 1\}$, which is a random variable due to the randomness of the observations. The expected stopping time under state $X(\tilde{c}) = i$ is then given by $E_i[N(\phi)] = E[N(\phi)|X(\tilde{c}) = i]$.

Since now we assign a cost c_{obs} for each observation, the conditional Bayes risks (5,6) under UCA over $L \geq 0$ observations can be modified to be:

$$R_i(\tilde{c}, L, \Delta) = \text{Prob}(\text{decide } X(\tilde{c}) \neq i | X(\tilde{c}) = i) + c_{\text{obs}} E_i[N(\phi)], \quad i = 0, 1, 2. \quad (7)$$

If $L \geq 1$, Δ has explicit matrix form and we can further rewrite the above equations as:

$$R_i(\tilde{c}, L, \Delta) = c_i \Delta b_i + c_{\text{obs}} E_i[N(\phi)], \quad i = 0, 1, 2. \quad (8)$$

Define the Bayes risk as the expected conditional Bayes risk of making a wrong decision under decision rule Δ :

$$r(\tilde{c}, L, \pi_1, \pi_2, \Delta) = (1 - \pi_1 - \pi_2) R_0(\tilde{c}, L, \Delta) + \pi_1 R_1(\tilde{c}, L, \Delta) + \pi_2 R_2(\tilde{c}, L, \Delta), \quad L \geq 0, \quad (9)$$

where $\pi_1 = P(X(\tilde{c}) = 1; t = t_v)$ and $\pi_2 = P(X(\tilde{c}) = 2; t = t_v)$ are the prior probabilities of object having Property ‘F’ and Property ‘G’, respectively, at cell \tilde{c} and $\pi_0 = 1 - \pi_1 - \pi_2$ gives the prior probability of object absent. Here, t_v is the time instant whenever an observation is taken at cell \tilde{c} . Fix a pair of (π_1, π_2) under the constraints $\pi_i \in [0, 1]$ and $\sum_{i=1}^2 \pi_i \leq 1$, the minimum Bayes risk surface at this particular cell has the minimum r value over all possible choices of Δ with $L \geq 0$. We want to sequentially make optimal decisions based on all Δ under $L \geq 0$ such that the Bayes risk r is minimized at each time step.

If the sensor does not take any observations ($L = 0$) and directly make a decision, the Bayes risks of 3 different decision rules Δ are as follows: $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide object absent}) = \pi_1 + \pi_2$, $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide Property ‘F’}) = 1 - \pi_1$, $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide Property ‘G’}) = 1 - \pi_2$.

If the sensor takes an observation at $t = 0$ ($L = 1$), the minimum Bayes risk over all possible choices of Δ with $L = 1$ is

$$r_{\min}(\tilde{c}, L = 1, \pi_1, \pi_2) = \min_{\Delta \in \mathcal{G}_L} \pi_0 R_0(\tilde{c}, L = 1, \Delta)$$

$$+ \pi_1 R_1(\tilde{c}, L = 1, \Delta) + \pi_2 R_2(\tilde{c}, L = 1, \Delta) \geq c_{\text{obs}},$$

where \mathcal{G}_L is defined as the set of all deterministic decision rules that are based on exactly L observations (Here, $L = 1$).

Following the same procedure, we compute the minimum Bayes risk functions $r_{\min}(\tilde{c}, L, \pi_1, \pi_2)$ under different observation numbers $L \geq 0$ and then find the overall minimum Bayes risk,

$$r_{\min}^*(\tilde{c}, \pi_1, \pi_2) = \min_{L=0,1,2,\dots} r_{\min}(\tilde{c}, L, \pi_1, \pi_2).$$

The basic idea of the cost-aware Bayesian sequential decision-making strategy is as follows: With initial priors $\pi_1 = P(X(\tilde{c}) = 1; t = 0)$ and $\pi_2 = P(X(\tilde{c}) = 2; t = 0)$, check the corresponding r_{\min}^* value in the overall minimum Bayes risk surface. If r_{\min}^* is given by the risk

plane with $L \geq 1$, the Bayes risk is lowered by taking an observation $Y_{t=0}(\tilde{c})$. Compute the posterior probabilities $P(X(\tilde{c}) = i | Y_{t=0}(\tilde{c}); t = 1)$ according to Equation (4) and again check the corresponding minimum Bayes risk r_{\min}^* to make decisions. The process is repeated using these posteriors as the new priors until the Bayes risk of taking one more observation is higher than the cost of making a wrong decision.

Let us illustrate the details of the above scheme via the following preliminary simulation for a single cell.

V. SIMULATION FOR A SINGLE CELL

In this simulation, we fix a cell \tilde{c} , and assume that the sensor is located at the centroid of this cell. The sensing parameters are chosen as follows:

$$\begin{aligned} \beta_{00} &= 0.8, \beta_{01} = 0.1, \beta_{02} = 0.1, \\ \beta_{10} &= 0.2, \beta_{11} = 0.7, \beta_{12} = 0.1, \\ \beta_{20} &= 0.1, \beta_{21} = 0.15, \beta_{22} = 0.75. \end{aligned} \quad (10)$$

The observation cost is set as $c_{\text{obs}} = 0.05$. Figure 1 shows the overall minimum Bayes risk $r_{\min}^*(\tilde{c}, \pi_1, \pi_2)$. It is constructed by taking the smallest value of all $r_{\min}(\tilde{c}, L, \pi_1, \pi_2)$, $L \geq 0$ under each fixed prior probability pair (π_1, π_2) . Here, we only list the expressions for the risk planes of decision rules that constitute r_{\min}^* as annotated by the numerals 1 – 10 in Figure 1. The Bayes risk functions under more than 3 observations ($L \geq 3$) have larger r values and do not contribute to r_{\min}^* for the particular choice of β and c_{obs} here.




Fig. 1. The overall minimum Bayes risk surface r_{\min}^* is composed of the enumerated risk planes described in this section.

Risk Plane 1. This plane represents the decision rule “decide no object at the cell regardless of observations under $L = 0$ ”. According to Equations (7) and (9), the Bayes risk is $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide no object}) = \pi_1 + \pi_2$.

Risk Plane 2. This plane represents the decision rule “decide there is an object with Property ‘F’ regardless of the observations under $L = 0$ ”: $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide Property ‘F’}) = 1 - \pi_1$.

Risk Plane 3. This plane represents the decision rule “decide there is an object with Property ‘G’ regardless of

the observations under $L = 0$ ”: $r(\tilde{c}, L = 0, \pi_1, \pi_2, \Delta = \text{decide Property ‘G’}) = 1 - \pi_2$.

Risk Plane 4. This plane corresponds to the decision rule under $L = 1$. The general conditional probability matrix for $L = 1$ is given as

$$B(L = 1) = \begin{bmatrix} \beta_{00} & \beta_{10} & \beta_{20} \\ \beta_{01} & \beta_{11} & \beta_{21} \\ \beta_{02} & \beta_{12} & \beta_{22} \end{bmatrix},$$

where the rows correspond to the observations $(z_0 = 1, z_1 = 0, z_2 = 0)$, $(z_0 = 0, z_1 = 1, z_2 = 0)$, and $(z_0 = 0, z_1 = 0, z_2 = 1)$, respectively. Let us consider the following decision rule,

$$\Delta_{11} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

That is, decide the actual state according to the only one observation that was taken. Therefore, $r(\tilde{c}, L = 1, \pi_1, \pi_2, \Delta = \Delta_{11})$ is given directly by Equations (8) and (9).

Risk Plane 5. This plane gives the decision rule under $L = 2$. The general conditional probability matrix is given as

$$B(L = 2) = \begin{bmatrix} \beta_{00}^2 & \beta_{10}^2 & \beta_{20}^2 \\ \beta_{01}^2 & \beta_{11}^2 & \beta_{21}^2 \\ \beta_{02}^2 & \beta_{12}^2 & \beta_{22}^2 \\ 2\beta_{00}\beta_{01} & 2\beta_{10}\beta_{11} & 2\beta_{20}\beta_{21} \\ 2\beta_{00}\beta_{02} & 2\beta_{10}\beta_{12} & 2\beta_{20}\beta_{22} \\ 2\beta_{01}\beta_{02} & 2\beta_{11}\beta_{12} & 2\beta_{21}\beta_{22} \end{bmatrix},$$

where the rows correspond to the observations $(z_0 = 2, z_1 = 0, z_2 = 0)$, $(z_0 = 0, z_1 = 2, z_2 = 0)$, $(z_0 = 0, z_1 = 0, z_2 = 2)$, $(z_0 = 1, z_1 = 1, z_2 = 0)$, $(z_0 = 1, z_1 = 0, z_2 = 1)$, and $(z_0 = 0, z_1 = 1, z_2 = 1)$, respectively. Risk Plane 5 corresponds to the following decision rule,

$$\Delta_{21} = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}.$$

Following the same procedure as above, we can get $r(\tilde{c}, L = 2, \pi_1, \pi_2, \Delta = \Delta_{21})$ according to Equation (9) without difficulty.

Risk Plane 6-10. These planes also give the decision rules under $L = 2$. The corresponding decision rules are,

$$\begin{aligned} \Delta_{22} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} && \text{(Risk Plane 6)} \\ \Delta_{23} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{bmatrix} && \text{(Risk Plane 7)} \\ \Delta_{24} &= \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix} && \text{(Risk Plane 8)} \\ \Delta_{25} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} && \text{(Risk Plane 9)} \\ \Delta_{26} &= \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} && \text{(Risk Plane 10)} \end{aligned}$$

The Bayes risks follow as above.

When r_{\min}^* is given by Risk Plane 1, 2 or 3, the sensor stops taking observation and makes the corresponding decision, otherwise, it always takes one more observation.

VI. THE UNCERTAINTY MAP AND TASK METRICS

In this section, we define the uncertainty maps based on the posterior probabilities derived in Section III as well as the metrics for the search and classification tasks in general multi-cell domains. We also relate the task metrics with a dynamic observation cost for the Bayesian sequential decision-making strategy in multi-cell domains.

A. The Uncertainty Map

For the search task, we use the information entropy to construct an uncertainty map for a multi-cell task domain [22]. Let the probability distribution $P(\tilde{c}, t)$ for object absent and present at cell \tilde{c} at time t be $P(\tilde{c}, t) = \{P(X(\tilde{c}) = 0; t), 1 - P(X(\tilde{c}) = 0; t)\}$. We define the information entropy for the distribution $P(\tilde{c}, t)$ as:

$$H_0(P(\tilde{c}, t)) = -P(X(\tilde{c}) = 0; t) \ln P(X(\tilde{c}) = 0; t) - (1 - P(X(\tilde{c}) = 0; t)) \ln(1 - P(X(\tilde{c}) = 0; t)). \quad (11)$$

$H_0(P(\tilde{c}, t))$ measures the uncertainty level of object existence at cell \tilde{c} at time t . The greater the value of H_0 , the larger the uncertainty is. Note that $H_0(P(\tilde{c}, t)) \geq 0$. The desired uncertainty level is $H_0(P(\tilde{c}, t)) = 0$ and its maximum attainable value is $H_{0,\max} = 0.6931$ when $P(X(\tilde{c}) = 0; t) = 0.5$. See [2], [3] and references therein for more detailed properties of the uncertainty function. The information entropy distribution at time step t over the domain forms an uncertainty map at that time instant.

Similarly, we can build uncertainty maps H_1 and H_2 for state $X(\tilde{c}) = 1, 2, \forall \tilde{c} \in \mathcal{D}$, respectively, to evaluate the sensor's confidence level for classification.

B. Task Metrics

When the observation cost is low, the Bayes risk is minimized by taking more observations, the sensor will decide not to proceed searching for more objects, but to stop and take an observation at the current cell. Under this scenario, we define the cost of not carrying on further search as follows:

$$\mathcal{J}(t) = \frac{\sum_{\tilde{c} \in \mathcal{D}} H_0(P(\tilde{c}, t))}{H_{0,\max} A_{\mathcal{D}}}, \quad (12)$$

where $A_{\mathcal{D}}$ is the area of the domain. The cost \mathcal{J} is proportional to the sum of the search uncertainty over \mathcal{D} . According to this definition, we have $0 \leq \mathcal{J}(t) \leq 1$. If $H_0(P(\tilde{c}, t_s)) = 0$ at some $t = t_s$ for all $\tilde{c} \in \mathcal{D}$, then $\mathcal{J}(t_s) = 0$ and the entire domain has been satisfactorily covered and we know with 100% certainty that there are no more objects yet to be found.

Similarly, We use H_1 and H_2 as the classification metrics. When the classification uncertainty of a cell is within a small neighborhood of zero, the classification task is said to be completed.

Now let us associate a dynamic observation cost $c_{\text{obs}}(t)$ with the search cost function $\mathcal{J}(t)$,

$$c_{\text{obs}}(t) = \gamma \mathcal{J}(t), \quad (13)$$

where $\gamma > 0$ is some constant. At the outset of the mission, the observation cost is high since there are still many uncovered regions in the domain. The cost-aware Bayesian sequential decision-making strategy tends to make a decision with a few observations, which may yield large number of wrong decisions, but increase the potential of rapidly detecting more critical objects. When the sensor has surveyed more regions, the uncertainty level for all the visited cells is reduced, and both the search cost function and the observation cost decrease. The process will be repeated until $\mathcal{J}(t) \rightarrow 0$, $H_1 \rightarrow 0$, and $H_2 \rightarrow 0$, $\forall \tilde{c} \in \mathcal{D}$, i.e., all the unknown objects of interest within the domain have been found and classified with a desired uncertainty level.

VII. FULL-SCALE DOMAIN MOTION CONTROL & SIMULATIONS

In this section, we consider a sensor motion control strategy over the mission domain \mathcal{D} that seeks to find all objects in \mathcal{D} with a desired confidence level (i.e., achieve $\mathcal{J} \rightarrow 0$). For the sake of simplicity, we assume that there is no speed limit on the sensor, i.e., the sensor is able to move to any cell within \mathcal{D} from its current location.

First, we define the set $\mathcal{Q}_H(t) = \{\tilde{c} \in \mathcal{D} : \text{argmax}_{\tilde{c}} H_0(P(\tilde{c}, t))\}$. Next, let $\tilde{q}_c(t)$ be the centroid of the cell that the sensor is currently located at and define the subset $\mathcal{Q}_d(t) \subseteq \mathcal{Q}_H(t)$ as $\mathcal{Q}_d(t) = \{\tilde{c} \in \mathcal{Q}_H(t) : \text{argmin}_{\tilde{c}} \|\tilde{q}_c(t) - \tilde{q}\|\}$, where \tilde{q} is the centroid of \tilde{c} . The set $\mathcal{Q}_d(t)$ contains the cells which have both the shortest distance from the current cell and the highest search uncertainty. When the sensor finishes taking observations in a current cell via the Bayesian sequential decision-making strategy and decides to move to a new cell, it will choose the next cell to go to from $\mathcal{Q}_d(t)$. Note that $\mathcal{Q}_d(t)$ may have more than one cell. Let N_{Hd} be the number of cells in $\mathcal{Q}_d(t)$, the sensor will randomly pick a cell from $\mathcal{Q}_d(t)$ with probability $\frac{1}{N_{\text{Hd}}}$.

Next, let us demonstrate the cost-aware Bayesian sequential decision-making strategy over a full-scale domain via a simulation. We consider a 20×20 square domain \mathcal{D} . For each $\tilde{c} \in \mathcal{D}$, we assume an i.i.d. prior probability distribution with $P(X(\tilde{c}) = 0; t = 0) = 0.8$, $P(X(\tilde{c}) = 1; t = 0) = 0.1$, and $P(X(\tilde{c}) = 2; t = 0) = 0.1$. The number, locations and properties of the objects are randomly generated. The radius of the sensor is shown by the black circle in Figure 2. The black dot represents the position of the sensor. The sensing parameters β_{ij} are the same as in Equation (10). The constant γ for the observation cost in Equation (13) is set as 0.05 and the desired uncertainty for every cell is $\epsilon = 0.02$.

The number of objects turns out to be 72 (the expected number of objects is 80 according to Equation (1)) with locations indicated by the 37 white dots (objects with Property 'F') and 35 magenta dots (objects with Property 'G') in Figure 2. Figure 2 shows the evolution of H_0 . At $t = 1023$, $H_0 = 1.1 \times 10^{-2} < \epsilon$ has been achieved everywhere within \mathcal{D} .




Fig. 2. Uncertainty map (dark red for highest uncertainty and dark blue for lowest uncertainty) at (a) $t = 1$, (b) $t = 200$, (c) $t = 620$, and (d) $t = 1023$ (with initial uncertainty $H_0(P(\bar{c}, 0))$).

Figure 3(a) shows the evolution of $\mathcal{J}(t)$ and can be seen to converge to zero. Figure 3(b) shows that object 6 (located at (2, 11)) has Property ‘F’ with classification uncertainty $H_1 = 0.0074 < \epsilon$ at time step 1030. The result is consistent with the simulation setup. The properties of other objects are also satisfactorily classified and can be shown like Figure 3(b).




Fig. 3. (a) Evolution of $\mathcal{J}(t)$, and (b) Probability of object 6 having Property ‘F’ and the corresponding uncertainty function H_1 .

The total number of missed detections during the entire missions is 16, and that of false detections is 5. The number of misclassifications of Property ‘G’ given ‘F’, and Property ‘F’ given ‘G’ is 1 and 0, respectively. Note that the numbers of erroneous decisions are small relative to the total number of cells within the domain. This suggests that the cost-aware Bayesian sequential decision-making strategy is efficient in making good decisions given limited available observations.

VIII. CONCLUSION

In this paper, a cost-aware decision-making strategy was developed for the detection and satisfactory classification of all objects in a domain via sequential Bayesian risk analysis. Future research will focus on the tracking and classification of mobile objects using multiple autonomous vehicles. The question of unknown environment

geometries will also be addressed. Objects with non i.i.d distributions over the domain will be investigated, where decision-making at one cell is affected by all the decisions made at other cells. Sequential Probability Ratio Test (SPRT) method will also be investigated for the cases where no prior information is available.

REFERENCES

- [1] Y. Wang, I. I. Hussein, and R. S. Erwin, “Awareness-Based Decision Making for Search and Tracking,” *American Control Conference*, 2008, invited Paper.
- [2] Y. Wang and I. I. Hussein, “Bayesian-Based Decision Making for Object Search and Characterization,” *IEEE American Control Conference*, 2009.
- [3] Y. Wang, I. I. Hussein, D. R. Brown III, and R. S. Erwin, “Cost-Aware Sequential Bayesian Tasking and Decision-Making for Search and Classification,” *American Control Conference*, 2010.
- [4] T. Furukawa, F. Bourgault, B. Lavis, and H. F. Durrant-Whyte, “Recursive Bayesian Search-and-Tracking using Coordinated UAVs for Lost Targets,” *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, May 2006.
- [5] B. Lavis, T. Furukawa, and H. F. Durrant-Whyte, “Dynamic Space Reconfiguration for Bayesian Search-and-Tracking with Moving Targets,” *Autonomous Robots*, vol. 24, pp. 387–399, May 2008.
- [6] R. Mahler, “Objective Functions for Bayesian Control-Theoretic Sensor Management, I: Multitarget First-Moment Approximation,” *Proceedings of IEEE Aerospace Conference*, 2003.
- [7] M. Flint, M. Polycarpou, and E. Fernández-Gaucherand, “Cooperative Control for Multiple Autonomous UAV’s Searching for Targets,” *Proceedings of the 41st IEEE Conference on Decision and Control*, December 2002.
- [8] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer-Verlag, 1994.
- [9] I. I. Hussein and D. Stipanović, “Effective Coverage Control using Dynamic Sensor Networks with Flocking and Guaranteed Collision Avoidance,” *2007 American Control Conference*, pp. 3420–3425, July 2007.
- [10] —, “Effective Coverage Control using Dynamic Sensor Networks,” *2006 IEEE Conference on Decision and Control*, pp. 2747–2752, December 2006.
- [11] —, “Effective Coverage Control for Mobile Sensor Networks with Guaranteed Collision Avoidance, special issue on multi-vehicle systems cooperative control with applications,” *IEEE Transactions on Control Systems Technology*, vol. 15, no. 4, pp. 642–657, 2007.
- [12] Y. Wang and I. I. Hussein, “Cooperative Vision-based Multi-vehicle Dynamic Coverage Control for Underwater Applications,” *Proceedings of the IEEE Multiconference on Systems and Control*, 2007, invited paper.
- [13] I. I. Hussein, D. Stipanović, and Y. Wang, “Reliable Coverage Control using Heterogeneous Vehicles,” *Conference on Decision and Control*, 2007.
- [14] Y. Wang and I. I. Hussein, “Underwater Acoustic Imaging using Autonomous Vehicles,” *2008 IFAC Workshop on Navigation, Guidance and Control of Underwater Vehicles*, April 2008.
- [15] —, “Awareness Coverage Control Over Large Scale Domains with Intermittent Communications,” *2008 American Control Conference*, 2008.
- [16] —, “Awareness Coverage Control Over Large Scale Domains with Intermittent Communications,” *IEEE Transactions on Automatic Control*, vol. 55, no. 8, August 2010, to appear.
- [17] H. V. Poor and O. Hadjiladis, *Quickest Detection*, 1st ed. Cambridge University Press, December 2008.
- [18] A. Wald, *Sequential Analysis*. Dover Publications, 2004.
- [19] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Volume I*. John Wiley & Sons, 2003.
- [20] A. D. Whalen, *Detection of Signals in Noise*. New York: Academic Press, 1971.
- [21] T. A. Schonhoff and A. A. Giordano, *Detection and Estimation Theory and Its Applications*. Prentice Hall, 2006.
- [22] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley, 2006.