Cost-Aware Sequential Bayesian Tasking and Decision-Making for Search and 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 sequential Bayesian decision-making strategy for the search and classification of multiple unknown objects within a task domain. Search and classification of multiple objects of unknown numbers are competing tasks under limited vehicle and sensory resources. This is because sensor equipped vehicles in the system can perform either the search or classification task but not both at the same time. The decision of one task over the other may result in missing other, more important objects not yet found or missing the opportunity to classify a found critical object. Building on previous Bayesian-based probabilistic work by the authors, in this paper we develop a cost-aware sequential Bayesian decision-making strategy for search and classification. It will be demonstrated that the proposed approach results in detection and satisfactory classification of all objects in the domain.**

I. Introduction

In a search and classification mission, an autonomous sensor-equipped vehicle searches for and classifies multiple objects distributed over a domain. Search and classification are two competing demands. The objective in a search task is to find each object and fix its position in space, while the objective in a classification task is to observe each found object and collect enough information to classify it. Therefore, a sensor vehicle can perform either the search task or the classification task but not both at the same time (search requires mobility and classification constrains the motion of the vehicle to that of the object). Hence, a sensor vehicle has to decide on whether to continue searching or stop and characterize once it finds an object. This decision may be very critical in some applications as in search and rescue, where, for example, finding and analyzing a nonhuman object may come at the cost of delaying or altogether missing a live human victim. Conversely, a vehicle may come across a human victim and, at the cost of missing it, decides to continue the search task.

When we consider the observation cost of taking each new measurement, we may soon make a decision for either the search/classification process using very limited observations at hand with high uncertainty. This is because, at the outset of the mission, it is costly to take observations at one specific spot while ignoring a large amount of unsurveyed regions in the domain. Building on previous work in the deterministic framework [1], and the Bayesian-based probabilistic framework in [2], in this work a cost-aware Bayesian sequential search and classification decision-making algorithm is developed to guarantee the detection and satisfactory classification of all objects in the domain with minimum Bayes risk.

We first review some of the related literature. Inspired by work on particle filtering, in [3] the authors develop a strategy to dynamically control the relative configuration of sensor teams under a probabilistic framework. The goal is to get optimal estimates for target tracking through sensor fusion. In [4], the authors use the Beta distribution to model the level of confidence of target existence for an unmanned aerial vehicle (UAV) search task in an uncertain environment. The Beta distribution defined for each cell is a function of the prior probabilities which is updated through Bayes' theorem. In [5], the above uncertainty measurements are extended by using the Modified Bayes Factor, and prediction of future measurement is also taken into account to calculate the possible uncertainty reduction in UAV search operations. An alternate approach for searching in an uncertain environment is simultaneous localization and mapping (SLAM) [6].

Coordinated search and tracking in a probabilistic framework has been studied mainly for optimal path planning in the literature. In [7], the authors investigate search-and-tracking using recursive Bayesian filtering with foreknown targets' positions with noise. A vehicle will keep searching until the target detection probability is above some preset threshold. However, the target might be lost and need to be found again due to measurement noise. The results are extended in [8] for dynamic search spaces based on forward reachable set analysis. In [9], the author proposes a Bayesian-based multisensor-multitarget sensor management scheme. The approximation strategy, based on probability hypothesis densities, maximizes the square of the expected number of targets. With the same objective, in [10] the authors seek to maximize the probability of finding a target with some foreknown location information in the presence of uncertainty. In the above literature, there is no explicit decision-making strategy for search and tracking proposed in the above literature.

Sequential detection [11], also known as quickest detection [12], allows the number of observation samples to vary in order to achieve an optimal decision. Due to the randomness of observations we may get at each time

*Mechanical Engineering Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609. E-mail: {yuewang, ihussein}@wpi.edu.

†Electrical and Computer Engineering Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609. E-mail: drb@wpi.edu.

‡Air Force Research Laboratory, Space Vehicles Directorate AFRL/RV, 3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117.

step, a decision may be made with a few observation samples, whereas for other cases we would rather take more samples for a possibly better decision. In our problem, with a relatively high observation cost at the beginning of the mission, it is wise to make a crude decision first and return later when the cost is low. The Baysian sequential detection method used in this paper is such that the Bayes risk is minimized in each time step. Other sequential detection method includes Sequential Probability Ratio Test (SPRT) [11] based on Neyman-Pearson formulation where no prior probability information is needed.

In this paper, we will employ the sequential Bayesian decision-making strategy for both the search and classification processes, which considers the cost of taking observations. We first introduce a Bernoulli type sensor model in Section II, combined with a real time observation, we update the posterior probability according to the Bayes rule in Section III-A. In Section III-B, an uncertainty map for the search process is built based on the probability of object presence over the domain. A classification uncertainty function is also defined for each found object. Based on the uncertainty function, we propose the metrics for the search and classification processes in Section IV. In Section V, we study the cost-aware sequential Bayesian decision-making method. We then make sequential decisions based on both the posterior probabilities and the minimum Bayes risk curve. To illustrate the performance of our approach, in Section VI, we provide a simulation result for a single cell in a search task. The algorithms for motion control of the autonomous vehicle is developed in Section VII and a full-scale simulation with the autonomous mobile vehicle over the entire mission domain for search versus classification is presented in Section VIII. We conclude the paper by current and future work in Section IX.

II. Setup and Sensor Model

Let $\mathcal{D} \subset \mathbb{R}^2$ be a domain in which objects to be found and classified are located. Let $\tilde{\mathbf{q}}$ be an arbitrary point in $\mathcal{D}$. Assume there are N_{tot} cells in $\mathcal{D}$ and let $1 \leq N_{\text{o}} \leq N_{\text{tot}}$ be the number of objects in $\mathcal{D}$. Both N_{o} and the positions of the objects in $\mathcal{D}$ are unknown beforehand. We assume that the probabilities of object presence are spatially i.i.d. over $\mathcal{D}$. Let the position of the static object $\mathcal{O}_j$, $j \in \{1, 2, \ldots, N_{\text{o}}\}$, be $\mathbf{p}_j$. $\mathcal{P}$ is the set of all object positions (unknown and randomly generated). N_{o} is a binomial random variable with parameters N_{tot} and $\text{Prob}(\tilde{\mathbf{q}} \in \mathcal{P})$, where $\text{Prob}(\tilde{\mathbf{q}} \in \mathcal{P})$ is the probability of object presence at point $\tilde{\mathbf{q}}$ (identical for all $\tilde{\mathbf{q}}$ and independent). Hence, the expectation of N_{o} equals to the number of total points multiplied by $\text{Prob}(\tilde{\mathbf{q}} \in \mathcal{P})$, that is,

$$E[N_{\text{o}}] = N_{\text{tot}} \text{Prob}(\tilde{\mathbf{q}} \in \mathcal{P}).$$

We will assume that there exists a single autonomous sensor-equipped vehicle (denoted by $\mathcal{V}$) that performs the search or classification tasks. The vehicle $\mathcal{V}$ satisfies the following simple first order discrete-time equation of motion

$$\mathbf{q}(t+1) \quad = \quad \mathbf{q}(t) + \mathbf{u}(t),$$

where $\mathbf{q} \in \mathcal{D} \subset \mathbb{R}^2$ represents the position of $\mathcal{V}$, and $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^2$ is the control input, $\mathcal{U}$ is the set of allowable controls. At any time t, the vehicle can either perform the search task or the classification task, but it is not capable of both at the same time.

In this work, for both the search and classification processes, we use a sensor model with Bernoulli distribution, which gives binary outputs for a single observation, however, with different observation contents: object "present" or "absent" for search, and property "F" or "G" for classification.

A. Sensor Model for Search

In the search process, let $X_s(\tilde{\mathbf{q}})$ be a binary state variable, where 0 corresponds to object absent, and 1 corresponds to object present. Note that the realization of $X_s(\tilde{\mathbf{q}})$ depends on the position of the observed point $\tilde{\mathbf{q}}$, that is,

$$X_s(\tilde{\mathbf{q}}) = \begin{cases} 1 & \tilde{\mathbf{q}} \in \mathcal{P}, \\ 0 & \text{otherwise.} \end{cases}$$

Since $\mathcal{P}$ is unknown and random, $X_s(\tilde{\mathbf{q}})$ is a random variable with respect to every $\tilde{\mathbf{q}} \in \mathcal{D}$. Define the observation indicating object present as a positive observation and the observation indicating object absent as a negative observation. Let $Y_s(\tilde{\mathbf{q}})$ be a binary observation variable, where 0 corresponds to a negative observation, and 1 corresponds to a positive one taken at point $\tilde{\mathbf{q}}$.

Because the states $X_s(\tilde{\mathbf{q}})$ are spatially i.i.d., the observations $Y_s(\tilde{\mathbf{q}})$ taken at every point $\tilde{\mathbf{q}}$ within the mission domain $\mathcal{D}$ are spatially i.i.d. Conditioned on the state $X_s(\tilde{\mathbf{q}})$ at a particular point $\tilde{\mathbf{q}}$, let t be time index, the observations $Y_{s,t}(\tilde{\mathbf{q}})$ taken along time are temporally i.i.d. Therefore, if we take an observation at each time step at $\tilde{\mathbf{q}}$, for a window of L time steps, we have $L+1$ different combinations of unordered scalar observations, that is, ranging from zero positive observation to L positive ones. Let the variable $Z_s(\tilde{\mathbf{q}})$ be the number of positive observations at point $\tilde{\mathbf{q}}$, which is a number in the set $\{0, \cdots, L\}$. The following $(L+1) \times 2$ matrix gives the general conditional probability matrix for the search task:

$$B_s = \begin{bmatrix} \text{Prob}[Z_s = 0 | X_s(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_s = 0 | X_s(\tilde{\mathbf{q}}) = 1] \\ \text{Prob}[Z_s = 1 | X_s(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_s = 1 | X_s(\tilde{\mathbf{q}}) = 1] \\ \vdots & \vdots \\ \text{Prob}[Z_s = L | X_s(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_s = L | X_s(\tilde{\mathbf{q}}) = 1] \end{bmatrix},$$

with $\sum_{l=0}^{L} \text{Prob}[Z_s = l | X_s(\tilde{\mathbf{q}}) = j] = 1, j = 0, 1$. Because the sensor follows the Bernoulli distribution for a single observation, $\text{Prob}[Z_s = l | X_s(\tilde{\mathbf{q}}) = j]$ will follow a binomial distribution with parameter $\beta_{i,s}, i = 0, 1$ and L, which describes the probability of having l positive observations given state $X_s(\tilde{\mathbf{q}}) = j$.

For the sake of simplicity, we can assume that the

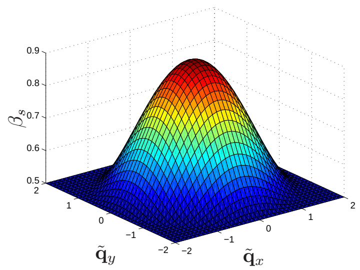

Fig. 1. β_s as a function of distance from sensor location to an arbitrary point in the domain $\mathcal{D}$ with $\mathbf{q} = 0$, $M = 0.4$ and $r_s = 2$.

sensor probabilities of making a correct observation are the same. That is, we have $\beta_{1,s} = \beta_{0,s} = \beta_s$. Hence, the conditional probability matrix can be written as follows:

$$B_s = \begin{bmatrix} \beta_s^L & (1-\beta_s)^L \\ L(1-\beta_s)\beta_s^{L-1} & L\beta_s(1-\beta_s)^{L-1} \\ \vdots & \vdots \\ (1-\beta_s)^L & \beta_s^L \end{bmatrix}. \quad (1)$$

The value of β_s depends on the range between the sensor and the observed point. Here we assume a simple model for β_s that is a fourth order polynomial function of $s = \|\mathbf{q}(t) - \tilde{\mathbf{q}}\|$ within the sensor range r_s and $b_n = 0.5$ otherwise,

$$\beta_s(s) = \begin{cases} \frac{M_s}{r_s^4}\left(s^2 - r_s^2\right)^2 + b_n & \text{if } s \le r_s \\ b_n & \text{if } s > r_s \end{cases}, \quad (2)$$

where $M_s + b_n$ gives the peak value of β_s if $\tilde{\mathbf{q}}$ being observed is located at the sensor vehicle's location.The sensing capability decreases with range and becomes 0.5 outside of the limited sensory range $\mathcal{W}$. Figure 1 shows the β_s function over a square domain of size 4×4. Other models for β_s can also be used and the analysis still holds.

B. Sensor Model for Classification

In the classification process, let $X_c(\tilde{\mathbf{q}}_k)$ be a binary state variable for object $\tilde{\mathbf{q}}_k, k = 1, 2, \cdots, N_o$, where 0 corresponds to object having Property "G", and 1 corresponds to Property "F". Let $Y_c(\tilde{\mathbf{q}}_k)$ be a binary observation variable, where 0 corresponds to observation showing Property "G", and 1 corresponds to observation showing Property "F".

Assume the observations $Y_{c,t}(\tilde{\mathbf{q}}_k)$ taken along time are temporally i.i.d. For a window of L time steps, there are $L+1$ different combinations of unordered scalar observations. Let $Z_c(\tilde{\mathbf{q}}_k)$ be the number of observations showing Property "F" for object k. Assume that the sensor probabilities of making a correct observation are the same, the following $(L+1)\times 2$ matrix gives the general conditional probability matrix for the classification task:

$$B_c = \begin{bmatrix} \text{Prob}[Z_c = 0|X_c(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_c = 0|X_c(\tilde{\mathbf{q}}) = 1] \\ \text{Prob}[Z_c = 1|X_c(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_c = 1|X_c(\tilde{\mathbf{q}}) = 1] \\ \vdots & \vdots \\ \text{Prob}[Z_c = L|X_c(\tilde{\mathbf{q}}) = 0] & \text{Prob}[Z_c = L|X_c(\tilde{\mathbf{q}}) = 1] \end{bmatrix}$$

$$= \begin{bmatrix} \beta_c^L & (1-\beta_c)^L \\ L(1-\beta_c)\beta_c^{L-1} & L\beta_c(1-\beta_c)^{L-1} \\ \vdots & \vdots \\ (1-\beta_c)^L & \beta_c^L \end{bmatrix},$$

with $\sum_{l=0}^{L} \text{Prob}[Z_c = l|X_c(\tilde{\mathbf{q}}) = j] = 1, j = 0, 1.$

The value of β_c follows a some form as β_s and is given by

$$\beta_s(c) = \begin{cases} \frac{M_c}{r_c^4}\left(s^2 - r_c^2\right)^2 + b_n & \text{if } s \le r_c \\ b_n & \text{if } s > r_c \end{cases},$$

where $M_c + b_n$ gives the maximum sensing capacity and r_c is the sensing range for classification.

III. Bayesian Updates and Uncertainty Map

In this section, we summarize the main algorithms presented by the authors in [2].

A. Bayesian Updates for Search and Classification

For the search process, we define the general update equation for the probability of object presence $P_s(\tilde{\mathbf{q}}, t+1)$ as follows:

$$P_s(\tilde{\mathbf{q}}, t+1) = y_{s,t}(\tilde{\mathbf{q}}) \frac{\beta_s P_s(\tilde{\mathbf{q}}, t)}{2\beta_s P_s(\tilde{\mathbf{q}}, t) - \beta_s - P_s(\tilde{\mathbf{q}}, t) + 1} + (1 - y_{s,t}(\tilde{\mathbf{q}})) \frac{(1-\beta_s) P_s(\tilde{\mathbf{q}}, t)}{-2\beta_s P_s(\tilde{\mathbf{q}}, t) + \beta_s + P_s(\tilde{\mathbf{q}}, t)}, \quad (3)$$

where $y_{s,t}(\tilde{\mathbf{q}})$ is the actual realization of the random variable $Y_{s,t}(\tilde{\mathbf{q}})$. Note that the probability of object absence is given by $1 - P_s(\tilde{\mathbf{q}}, t+1)$.

For the classification process, we use a similar update equation as (3) to express the posterior probability of a found object k at location $\tilde{\mathbf{q}}_k$ having property "F":

$$P_c(\tilde{\mathbf{q}}_k, t+1) = y_{c,t}(\tilde{\mathbf{q}}) \frac{\beta_c P_c(\tilde{\mathbf{q}}_k, t)}{2\beta_c P_c(\tilde{\mathbf{q}}_k, t) - \beta_c - P_c(\tilde{\mathbf{q}}_k, t) + 1} + (1 - y_{c,t}(\tilde{\mathbf{q}})) \frac{(1-\beta_c) P_c(\tilde{\mathbf{q}}_k, t)}{-2\beta_c P_c(\tilde{\mathbf{q}}_k, t) + \beta_c + P_c(\tilde{\mathbf{q}}_k, t)}, \quad (4)$$

where $\tilde{\mathbf{q}}_k$ is the position of object k. The probability of having property "G" is $1 - P_c(\tilde{\mathbf{q}}_k, t+1)$.

B. Uncertainty Map

For the search process, we use the information entropy function of the probability distribution of object presence to construct the uncertainty map for every $\tilde{\mathbf{q}}$ within the search domain. The uncertainty map will be used to guide the vehicle in the search domain. We define the information entropy distribution for discrete probability distribution $P_{H_s} = \{P_s(\tilde{\mathbf{q}}, t), 1 - P_s(\tilde{\mathbf{q}}, t)\}$ (i.e, P_{H_s} is the probability distribution for object present and absent) at $\tilde{\mathbf{q}}$ at each time step t as

$$H_s(P_{H_s}, \tilde{\mathbf{q}}, t) = -P_s(\tilde{\mathbf{q}}, t) \ln P_s(\tilde{\mathbf{q}}, t) - (1 - P_s(\tilde{\mathbf{q}}, t)) \ln(1 - P_s(\tilde{\mathbf{q}}, t)). \quad (5)$$

For the classification process, we define a similar entropy function $H_c(P_{H_c}, \tilde{\mathbf{q}}_k, t)$, with $P_{H_c} = \{P_c(\tilde{\mathbf{q}}_k, t), 1 - P_c(\tilde{\mathbf{q}}_k, t)\}$, for every found object k (located at $\tilde{\mathbf{q}}_k$) to evaluate classification uncertainty.

$$H_c(P_{H_c}, \tilde{\mathbf{q}}_k, t) = -P_c(\tilde{\mathbf{q}}_k, t)\ln P_c(\tilde{\mathbf{q}}_k, t) - (1 - P_c(\tilde{\mathbf{q}}_k, t))\ln(1 - P_c(\tilde{\mathbf{q}}_k, t)). \quad (6)$$

There are as many scalar H_c's as there are found objects k up to time t.

IV. Search and Classification Metrics

In this section we develop metrics to be used for the search versus classification decision-making process. We define the cost of not carrying on further search as

$$\mathcal{J}(t) = \frac{\int_{\mathcal{D}} H_s(P_{H_s}, \tilde{\mathbf{q}}, t)\mathrm{d}\tilde{\mathbf{q}}}{H_{s,\max} A_{\mathcal{D}}}. \quad (7)$$

The cost $\mathcal{J}$ is proportional to the total integral of the search uncertainty over $\mathcal{D}$. We divide the integral by the area of the domain $A_{\mathcal{D}}$ multiplied by $H_{s,\max}$ in order to normalize $\mathcal{J}(t)$. According to this definition, we have $0 \le \mathcal{J}(t) \le 1$. If for some t_s we have $H_s(P_{H_s}, \tilde{\mathbf{q}}, t_s) = 0$ for all $\tilde{\mathbf{q}} \in \mathcal{D}$, then $\mathcal{J}(t_s) = 0$ and the entire domain has been satisfactorily covered and we know with 100% certainty that there are no objects yet to be found.

For the classification process, let $\bar{N}_o(t)$ be the number of objects found by the autonomous sensor vehicle up to time t. For each found object $j \in \{1, 2, \cdots, \bar{N}_o(t)\}$, define the classification metric $H_d(\tilde{\mathbf{q}}_j, t)$ to be

$$H_d(\tilde{\mathbf{q}}_j, t) = \epsilon_c \mathcal{J}(t), \quad (8)$$

where ϵ_c is a preset upper bound on the desired uncertainty level for classification.

Let us define the classification conditions as follows:

$$\begin{cases} \|\mathbf{q}(t) - \tilde{\mathbf{q}}_j\| \le r_c & \text{(a)} \\ H_c(P_{H_c}, \tilde{\mathbf{q}}_j, t) > H_d(\tilde{\mathbf{q}}_j, t) & \text{(b)} \\ H_s(P_{H_s}, \tilde{\mathbf{q}}_j, t) \le \epsilon_s & \text{(c)} \\ \text{No Decision at } \tilde{\mathbf{q}}_j \text{ at } t & \text{(d)} \end{cases}, \quad (9)$$

where ϵ_s is some upper bound on the search uncertainty to be met before a classification task can be carried on. Only when all the classification conditions are satisfied, i.e., (a) the object $\mathcal{O}_j$ is within the vehicle's sensory range, (b) the classification uncertainty of j is larger than the desired uncertainty, (c) the search uncertainty of j is relatively low (we are to some extent sure that j is an object), and (d) no decision has been made about the property of j yet at previous time step, then the vehicle will start to classify j. If any one of the above condition fails, the vehicle $\mathcal{V}$ stop classifying the found object and switch to searching again. It can resume classifying an object that has been detected and completely or partially classified in the past if it finds it again during the search process. When this occurs, the value of H_d will be smaller than the last time the objected has been detected.

V. Cost-Aware Sequential Decision-Making

Assuming a Uniform Cost Assignment (UCA), we define the decision cost matrix as

$$C_{ij} = \begin{cases} 0 & \text{if } i = j \\ 1 & \text{if } i \neq j \end{cases},$$

where $i = 0, 1$ represent 0: object absent and 1: object present, $j = 0, 1$ correspond to state $X_s(\tilde{\mathbf{q}}) = 0$ and $X_s(\tilde{\mathbf{q}}) = 1$. Hence C_{ij} is the cost of deciding i when the state is $X_s(\tilde{\mathbf{q}}) = j$. C can be written in the matrix form as

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

Let $\tilde{R}_0(\tilde{\mathbf{q}}, L, \Delta), L \ge 1$, be the conditional risk of deciding there is an object at $\tilde{\mathbf{q}}$ given that there is actually none over at least one observation,

$$\tilde{R}_0(\tilde{\mathbf{q}}, L, \Delta) = c_0^T \Delta b_0, \quad (10)$$

where c_0 is the first column of the cost matrix C and contains the costs of deciding object absence or presence when there is nothing at $\tilde{\mathbf{q}}$. The quantity $b_0 = [\beta_s^L, \cdots, (1-\beta_s)^L]^T$ is the first column of conditional probability matrix B_s and contains the probabilities of having zero ($Z_s = 0$) to L ($Z_s = L$) positive observations when there is nothing at $\tilde{\mathbf{q}}$. The quantity Δ is a deterministic decision rule and used to evaluate the Bayes risk r to help decide the minimum Bayesian risk $r^*_{\min}$ (to be defined in Section V). For $L \ge 1$, Δ is a $2 \times (L+1)$ matrix. The number 2 is the number of possible final decisions in this case, corresponding to "object absent" and "object present" respectively. The quantity $L \ge 1$ is the number of observations that one can make over a window of L time steps, and the first to $(L+1)_{\text{th}}$ columns in the Δ matrix correspond to zero to L positive observations. $\Delta_i^l, i = 0, 1,\ l = 0, 1, \cdots, L$ can be either 0 or 1, and $\sum_{i=0}^{1} \Delta_i^l = 1$. If $\Delta_i^l = 1$, it means that we will make decision i given there are l positive observations. Therefore, Δ can have 2^{L+1} different matrix values. When $L = 0$, i.e., there is no observation taken, Δ could be either "always decide there is an object" or "always decide there is no object", regardless of the observations, but there will be no explicit matrix form for Δ.

In this paper, we assume that the sensor is a "good" one, that is to say, the detection probability is higher than the error probability of the sensor, i.e., $\beta_s > 0.5$. Therefore, there are only a small number of "reasonable" deterministic decision rules. Given L observations, the set of "reasonable" deterministic decision rules is the set of all rules of the type

$$\Delta_1^l = \begin{cases} 1 & l \ge v \\ 0 & \text{otherwise} \end{cases}$$

where $l \in \{0, \ldots, L\}$ is the total number of positive observations and $v \in \{0, \ldots, L+1\}$ is the threshold where we make a positive decision. This means we only need to consider decision rule matrices that look like

$$\Delta = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix}$$

and not like

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

When the threshold $v = 0$, we will always decide object present and ignore the observations. Similarly, when $v =$

$L+1$, we will always decide object absent. Note that "reasonable" decision rules grows linearly with L and dominates any other type of decision rules with the same value of L.

Similarly, the conditional risk

$$\tilde{R}_1(\tilde{\mathbf{q}}, L, \Delta) = c_1^T \Delta b_1 \quad (11)$$

gives the cost of deciding that there is no object at $\tilde{\mathbf{q}}$ given that there is actually something over $L \geq 1$ observations.

Therefore, under UCA, there is no cost if the decision is the actual state, and the conditional risk $\tilde{R}_0(\tilde{R}_1)$ can be interpreted as the error probability of deciding there is an (no) object given that there is actually none (one) under a certain decision rule Δ over L observations for point $\tilde{\mathbf{q}}$.

Assume that we could collect one observation from the sensor at each time step. Now each observation is assigned a cost $c_{\text{obs}}(t)$, which may be based on energy, time etc. For both the search and classification processes, we couple the search task with the classification task using a dynamic observation cost c_{obs} that depends on how much more uncertainty there remains to be eliminated

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

where $\gamma > 0$ is some constant. That is to say, at the outset of the mission with a relatively high value of $\mathcal{J}$ (high uncertainty), the observation cost is high since there are still many uncovered points in the domain and it is expensive to only make observations at a particular point. The cost-aware Bayesian sequential decision-making strategy tends to make a decision with a few observation sample in that case, which may yield large number of false detections and miss detections. When the vehicle has surveyed most points in the domain, both the uncertainty and the observation cost decrease. The vehicle is able to come back to some already searched/classified points to take observations again and make better decisions with lower uncertainty. The process is repeated until $\mathcal{J}(t) \to 0$ for search or $H_c < H_d$ for classification.

At every time step, the vehicle has to choose to decide object present, decide object absent or take one more observation and postpone making any decisions regarding object presence to the following time step. 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 sequential Bayesian decision-making strategy is such that the Bayes risk at each time step is minimized. Denote $\phi = \{\phi_k\}_{k=0}^{\infty}$ to be the stopping rule and $\delta = \{\delta_k\}_{k=0}^{\infty}$ the terminal decision rule. If $\phi_k = 0$, we continue to take another measurement, if $\phi_k = 1$, we stop taking further observations. The quantity δ_k can be either one of three possibilities: decide object present, decide object absent or take one more observation. Note that while the final decision is either object present or absent, at every time step k before we could make a final decision, δ_k takes up one of three possibilities: make a decision of object present, make a decision of object absence, and make a decision of taking one more observation.

Define the stopping time as

$$N(\phi) = \min\{k : \phi_k = 1\},$$

which is a random variable due to the randomness of the observation we may get at each time step. The expected stopping time under state $X_s(\tilde{\mathbf{q}}) = j$ is then given by

$$E_j[N(\phi)] = E[N(\phi) | X_s(\tilde{\mathbf{q}}) = j].$$

Since now we assign a cost c_{obs} for each observation, the conditional Bayesian risk (10,11) under UCA over $L \geq 0$ observations can be modified to be:

$$\begin{aligned} R_0(\tilde{\mathbf{q}}, L, \Delta) &= \text{Prob(decide object presence}|X_s(\tilde{\mathbf{q}}) = 0) \\ &\quad + c_{\text{obs}} E_0[N(\phi)], \quad (12) \\ R_1(\tilde{\mathbf{q}}, L, \Delta) &= \text{Prob(decide no object}|X_s(\tilde{\mathbf{q}}) = 1) \\ &\quad + c_{\text{obs}} E_1[N(\phi)]. \quad (13) \end{aligned}$$

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

$$R_0(\tilde{\mathbf{q}}, L, \Delta) = c_0^T \Delta b_0 + c_{\text{obs}} E_0[N(\phi)], \quad (14)$$

$$R_1(\tilde{\mathbf{q}}, L, \Delta) = c_1^T \Delta b_1 + c_{\text{obs}} E_1[N(\phi)]. \quad (15)$$

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

$$\begin{aligned} &r(\tilde{\mathbf{q}}, L, \pi_1, \Delta) \\ &= (1-\pi_1) R_0(\tilde{\mathbf{q}}, L, \Delta) + \pi_1 R_1(\tilde{\mathbf{q}}, L, \Delta), \ L \geq 0 \quad (16) \end{aligned}$$

where $\pi_1 = \text{Prob}[\tilde{\mathbf{q}} \in \mathcal{P}]$ is the prior probability of object presence at point $\tilde{\mathbf{q}}$, which ranges from 0 to 1. Similarly, $\pi_0 = 1 - \pi_1$ gives the prior probability of object absence. Note that the prior is used to construct the minimum Bayes risk curve over all possible priors and lengths of observations. Fix a $\pi_1 \in [0, 1]$, the minimum Bayes risk curve at this particular prior has the minimum r value over all possible choices of Δ with $L \geq 0$. This is not the same prior probability $P_s(\tilde{\mathbf{q}}, t)$ defined in Section III-A, which evolves along with time for a specific $\tilde{\mathbf{q}}$.

We want to find a terminal decision rule δ based on all decision rules Δ with $L \geq 0$ such that the Bayes risk r is minimized at each time step. This sequential Bayesian-based decision making strategy is optimal in every time step, but not historical. All the cost paid in the past steps is forgotten in the current step, the best decision is chosen based on both the current observation $Y_{s,t}$ and the time-evolving prior probabilities $P_s(\tilde{\mathbf{q}}, t)$.

If we do not take any observations ($L = 0$) and directly make a decision, the Bayes risks of $2^1 = 2$ different decision rules Δ are as follows

$$\begin{aligned} &r(\tilde{\mathbf{q}}, L = 0, \pi_1, \Delta = \text{always decide there is an object}) \\ &= 1 - \pi_1, \\ &r(\tilde{\mathbf{q}}, L = 0, \pi_1, \Delta = \text{always decide there is no object}) \\ &= \pi_1. \end{aligned}$$

If we do not stop at $t = 0$ but take one observation ($L = 1$), the minimum Bayes risk over all possible choices of Δ with $L = 1$ is

$$r_{\min}(\tilde{\mathbf{q}}, L=1, \pi_1) = \min_{\Delta \in \mathcal{G}_L} (1-\pi_1) R_0(\tilde{\mathbf{q}}, L=1, \Delta) + \pi_1 R_1(\tilde{\mathbf{q}}, L=1, \Delta) \geq c_{\text{obs}}$$

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

Following the same procedure, we compute the minimum Bayes risk functions under different observation numbers and find the overall minimum Bayes risk over all decision rules ($L \geq 0$).

$$r^*_{\min}(\tilde{\mathbf{q}}, \pi_1) = \min_{L=0,1,2,\ldots} r_{\min}(\tilde{\mathbf{q}}, L, \pi_1).$$

The basic idea of the cost-aware sequential Bayesian decision-making method is as follows: With an initial prior probability of object presence $P_s(\tilde{\mathbf{q}}, t)$, we check its corresponding $r^*_{\min}$ value in the overall minimum Bayes risk curve. If $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$ is given by the line with $L \geq 1$, the Bayes risk is lowered by taking an observation $Y_{s,t+1}$, computing the posterior probability $P_s(\tilde{\mathbf{q}}, t+1)$ according to Equation (3) and again checking its corresponding minimum Bayesian risk $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$ to make decision. That is, an observation is taken if and only if the prior $\pi_1 = P_s(\tilde{\mathbf{q}}, t)$ is such that

$$r_{\min}(\tilde{\mathbf{q}}, L \geq 1, \pi_1) < \min(\pi_1, 1-\pi_1).$$

The left hand side of this inequality corresponds to the minimum risk of all decision rules that take at least one observation. The right hand side of this inequality corresponds to the minimum risk of the decision rules that take no observation. Hence, an observation is only taken if, given the prior, the risk is reduced with respect to not taking an observation. If an observation is taken, the prior is updated to a posterior according to (3) and the process is repeated using this posterior as the new prior. For example, at time step $t+1$, $P_s(\tilde{\mathbf{q}}, t)$ is regarded as the prior probability in the update equation (3). If the decision is to take one more observation, the same procedure is repeated until the Bayes risk of the taking one more observation is higher than the cost of making a wrong decision. That is to say, $r^*_{\min}$ is given by $r(\tilde{\mathbf{q}}, L=0, \pi_1, \Delta)$.

Let us illustrate the detailed scheme by the following first-step simulation.

VI. Simulation for a Single Cell Observation

In this simulation, we fix a point $\tilde{\mathbf{q}}$, choose $\beta_s = 0.8$ (i.e., $M = 0.3$ and the sensor is right at this point), and set the observation cost as a fixed number $c_{\text{obs}} = 0.05$ to demonstrate the sequential Bayesian-based decision rule. Figure 2(a) shows all the Bayes risk functions r under 0 (black lines), 1 (blue lines) or 2 (green lines) observations with $\pi_1 \in [0, 1]$. In Figure 2(b), the red lines indicate the overall minimum Bayes risk $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$. The overall minimum Bayes risk curve $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$ is constructed by taking the smallest value of all $r_{\min}(\tilde{\mathbf{q}}, L, \pi_1), L = 0, 1, 2, \cdots$ under each fixed prior probability π_1. Figure 2(c) shows the construction of the minimum Bayes risk (the red dot) under a fixed prior π^*_0. Here, we only show the lines of decision rules that constitute these red lines and list the equation of these lines. The Bayes risk functions under more than 3 observations ($L \geq 3$) have larger r values and do not contribute to $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$ for the particular choice of β_s and c_{obs} here.

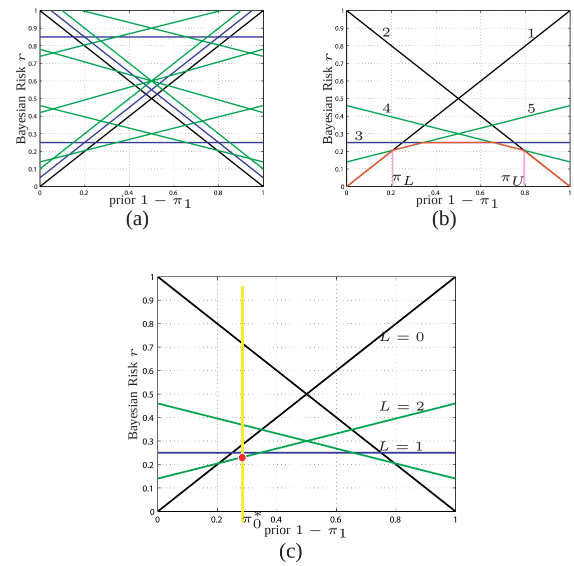

Fig. 2. (a) Bayes Risk Functions under $0, 1, 2$ observations; (b) Minimum Bayes Risk Function; (c) Demonstration of the construction of the minimum Bayes risk curve under a fixed prior π^*_0.

Let us interpret each of the lines.

Line 1. This line represents the decision rules without any observation. Always decide there is an object at the cell regardless of the observations. According to Equation (16),

$$r(\tilde{\mathbf{q}}, L=0, \pi_1, \Delta = \text{always decide there is an object}) = (1-\pi_1) \times 1 + \pi_1 \times 0 = 1 - \pi_1;$$

Line 2. This line also represents the decision rules without any observation. Always decide there is no object regardless of the observations:

$$r(\tilde{\mathbf{q}}, L=0, \pi_1, \Delta = \text{always decide there is no object}) = (1-\pi_1)(0 + c_{\text{obs}} \times 0) + \pi_1(1 + c_{\text{obs}} \times 0) = \pi_1.$$

Line 3. The blue line corresponds to the decision rule 3 after taking one observation: decide the actual state according to the only one observation, that is, if $Z_s = 1$, we decide there is actually an object. We have

$$r(\tilde{\mathbf{q}}, L=1, \pi_1, \Delta = \Delta_{11}) = (1-\pi_1)(1-\beta_s+c_{\text{obs}}) + \pi_1(1-\beta_s+c_{\text{obs}}) = 1-\beta_s+c_{\text{obs}}.$$

Line 4. This line gives the decision rules after two observations. Line 4 corresponds to the decision rule that decides there is actually an object if and only if all the two observations are positive ($Z_s = 2$). Following the same procedure as above, we have

$$r(\tilde{\mathbf{q}}, L=2, \pi_1, \Delta = \Delta_{21}) = (1-\beta_s)^2(1-\pi_1) + (2\beta_s(1-\beta_s) + (1-\beta_s)^2)\pi_1 + 2c_{\text{obs}};$$

Line 5. This line also gives the decision rules after two observations. Line 5 corresponds to the decision rule that decides there is no object if and only if none of the two observations is object present,

$$r(\tilde{\mathbf{q}}, L = 2, \pi_1, \Delta = \Delta_{22}) =$$
$$(2\beta_s(1-\beta_s) + (1-\beta_s)^2)(1-\pi_1) + (1-\beta_s)^2\pi_1 + 2c_{\text{obs}}.$$

Thus, the red line gives the minimum Bayesian risk $r^*_{\min}(\tilde{\mathbf{q}}, \pi_1)$ over 0,1,2 observations.

The intersection of lines 1, 5 is the lower prior probability $\pi_L = 0.2059$. When the posterior probability $1 - P_s(\tilde{\mathbf{q}}, t)$ updated through Equation (3) is below π_L, we stop taking observation and decide that the actual state is object present. This is because the minimum Bayesian risk is determined by line 1 instead of line 5 when $1 - P_s(\tilde{\mathbf{q}}, t) \in [0, \pi_L]$. The intersection of lines 2, 4 is the upper prior probability $\pi_U = 0.7941$. When $1 - P_s(\tilde{\mathbf{q}}, t)$ is above π_U (i.e., $P_s(\tilde{\mathbf{q}}, t) \leq \pi_L$), we decide that there is actually no object. We start with the prior $P_s(\tilde{\mathbf{q}}, 0) = 1 - P_s(\tilde{\mathbf{q}}, 0) = 0.5$, and decide to take one observation according to the red line. Next, we take an observation, and if the observation is $Y_{s,t=1}(\tilde{\mathbf{q}}) = 1$ indicating there is an object, then we update the posterior probabilities according to the new observation and the Bayes update rules (3). The posterior probability is $P_s(\tilde{\mathbf{q}}, 1) = 0.8, 1 - P_s(\tilde{\mathbf{q}}, 1) = 0.2 \leq \pi_L$. We decide not to take any more observation and determine there is actually an object at this point with Bayes risk $r = 0.2$ according to line 1.

VII. Vehicle Motion Control

In this section, we summarize the main results of the search control strategy presented in [2]. We will consider a motion control strategy for the vehicle that guarantees finding all objects in $\mathcal{D}$ (i.e., achieve $\mathcal{J} \to 0$) with the minimum Bayes risk at every time step.

Let the control $\mathbf{u}(t)$ be restricted to a set $\mathcal{U}$. Based on this constraint on the control, we define the set $\mathcal{Q}_{\mathcal{W}}(t)$ of points in $\mathcal{W}$ reachable from the current location of the vehicle at time t:

$$\mathcal{Q}_{\mathcal{W}}(t) = \{\tilde{\mathbf{q}} \in \mathcal{W} : \tilde{\mathbf{q}} - \mathbf{q}(t) \in \mathcal{U}\}. \qquad (17)$$

We use a control law that drives the vehicle to some point $\tilde{\mathbf{q}} \in \mathcal{Q}_{\mathcal{W}}(t)$ that has the highest uncertainty, and switch to a perturbation control law when the vehicle is trapped in a region where no such point exists. Let us first consider the following condition, whose utility will become obvious shortly.

Condition C1. $H_s(P_{H_s}, \tilde{\mathbf{q}}, t) \leq \epsilon,\ \forall \tilde{\mathbf{q}} \in \mathcal{Q}_{\mathcal{W}}(t)$, where ϵ is a preset threshold of some small value.

Consider the following control law

$$\mathbf{u}^*(t) = \begin{cases} \bar{\mathbf{u}}(t) & \text{if } \mathbf{C1} \text{ does not hold} \\ \bar{\bar{\mathbf{u}}}(t) & \text{if } \mathbf{C1} \text{ holds} \end{cases} \qquad (18)$$

where $\bar{\mathbf{u}}(t)$ is the *nominal control law*, and $\bar{\bar{\mathbf{u}}}(t)$ is the perturbation control law.

Let $\tilde{\mathbf{q}}_\star$ be the point that has the highest uncertainty within $\mathcal{Q}_{\mathcal{W}}(t)$, that is,

$$\tilde{\mathbf{q}}_\star(t+1) = \text{argmax}_{\tilde{\mathbf{q}} \in \mathcal{Q}_{\mathcal{W}}(t)} H_s(P_{H_s}, \tilde{\mathbf{q}}, t). \qquad (19)$$

The nominal control law is then set to be

$$\bar{\mathbf{u}}(t) = \tilde{\mathbf{q}}_\star(t+1) - \mathbf{q}(t) \in \mathcal{U}. \qquad (20)$$

This choice for the nominal control law is inspired by the nominal control law in [13].

If Condition **C1** holds, then the perturbation controller $\bar{\bar{\mathbf{u}}}(t)$ is used:

$$\bar{\bar{\mathbf{u}}}(t) = -\bar{\bar{k}}(\mathbf{q}(t) - \tilde{\mathbf{q}}^*)$$

where $0 < \bar{\bar{k}} \leq 1$ is the controller gain, and $\tilde{\mathbf{q}}^* \in \mathcal{Q}_{\mathcal{D}}(t) := \{\tilde{\mathbf{q}} \in \mathcal{D} : \tilde{\mathbf{q}} - \mathbf{q}(t) \in \mathcal{U}\}$ such that $H_s(P_{H_s}, \tilde{\mathbf{q}}^*, t) > \epsilon$. We assume that $\mathcal{U}$ is such that $\mathcal{Q}_{\mathcal{D}}(t) = \mathcal{D}$ for all time t. The controller is used to drive the vehicle out of the region with low uncertainty ϵ to some $\tilde{\mathbf{q}}^* \in \mathcal{Q}_{\mathcal{D}}(t)$ such that $H_s(P_{H_s}, \tilde{\mathbf{q}}^*, t) > \epsilon$, if such a point exists.

Here we present an efficient energy-wise way to choose a point $\tilde{\mathbf{q}}^*$. Let

$$\mathcal{D}_\epsilon(t) := \{\tilde{\mathbf{q}} \in \mathcal{Q}_{\mathcal{D}}(t) : H_s(P, \tilde{\mathbf{q}}, t) > \epsilon\},$$

which is an open set of all $\tilde{\mathbf{q}}$ for which $H_s(P, \tilde{\mathbf{q}}, t)$ is larger than a preset value ϵ. Let $\overline{\mathcal{D}}_\epsilon(t)$ be the closure of $\mathcal{D}_\epsilon(t)$. Let $\overline{\mathcal{D}}_{\epsilon,\mathcal{V}}(t)$ be the set of points in $\overline{\mathcal{D}}_\epsilon(t)$ that minimize the distance between the position vector of vehicle $\mathcal{V}$, $\mathbf{q}$, and the set $\overline{\mathcal{D}}_\epsilon(t)$:

$$\overline{\mathcal{D}}_{\epsilon,\mathcal{V}}(t)$$
$$= \left\{\tilde{\mathbf{q}}^* \in \overline{\mathcal{D}}_\epsilon(t) : \tilde{\mathbf{q}}^* = \text{argmin}_{\tilde{\mathbf{q}} \in \overline{\mathcal{D}}_\epsilon(t)} \|\tilde{\mathbf{q}} - \mathbf{q}(t)\|\right\}.$$

VIII. Simulation

In this simulation, we consider all the points $\tilde{\mathbf{q}}$ within a 20×20 square domain $\mathcal{D}$. For each $\tilde{\mathbf{q}} \in \mathcal{D}$, we assume an i.i.d. prior probability of object presence equals to $P_s(\tilde{\mathbf{q}}, 0) = \frac{E[N_o]}{N_{\text{tot}}} = 0.2$, where $E[N_o]$ is the expected number of object. The number and locations of the objects are randomly generated. The number of objects generated for this simulation turns out to be 82 with locations as indicated by the magenta dots in Figure 3. The radius r_s of the search sensor is chosen to be 8 and the classification radius r_c is chosen to be 6, as shown by the magenta and green circle in Figure 3. The black dot represents the position of the vehicle. Figure 3 shows the evolution of H_s. From Figure 3(d), we can conclude that at most $H_s = 1.1 \times 10^{-6}$ has been achieved everywhere within $\mathcal{D}$. In this case, we set the maximum sensing capacity as $M = 0.5$, and the cost-aware decision rule will stop and make decision if and only if the point is under at least sensory capability 0.95. The parameter $\gamma = 0.05$.

For the classification process, let the desired upper bound for classification uncertainty be $\epsilon_c = 0.01$ and $\epsilon_s = 0.3$ for search. The priors $P_c(\tilde{\mathbf{q}}, 0) = 0.5$ and all the objects with even number have property “F”.

Here we use the control law in equation (18) with control gain $\bar{\bar{k}} = 0.2$. The set $\mathcal{U}$ is chosen to be $\mathcal{D}$.

Figure 4(a) records the number of false detections and miss detections versus time. We can see from the figure that the number of miss detections (18) is much larger than that of the false detection (2) at the beginning of the task. This is because the initial prior probability $P_s(\tilde{\mathbf{q}}, 0)$

we start with is closer to zero, which makes it easier to have a wrong decision with one negative observation given that the actual state is object present. Figure 4(b) compares the number of deciding Property "F" given Property "G", and deciding Property "G" given Property "F" over all detected objects. These two numbers are similar since we have $P_c(\tilde{\mathbf{q}}, 0) = 0.5$. In both figures, it can be shown that as time increases, the number of miss detections and false detections decrease. Both of the error numbers go to zero with zero uncertainty at the end of the mission. This implies that we can balance between the error numbers within the tolerance range and the limited time we have to decide when to stop.

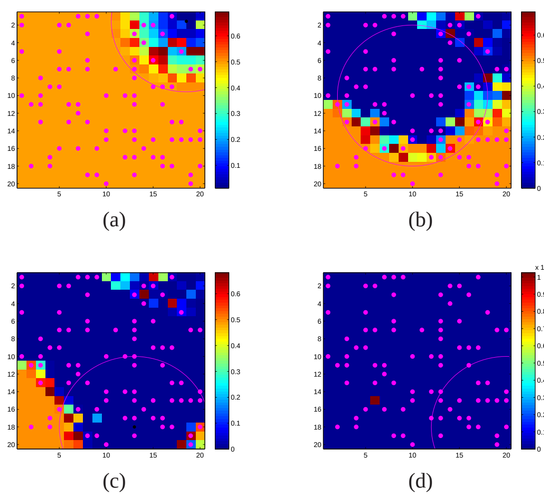

Fig. 3. Uncertainty map (dark red for highest uncertainty and dark blue for lowest uncertainty) at (a) $t = 1$, (b) $t = 200$, (c) $t = 400$, and (d) $t = 800$ (with initial uncertainty $H_s(P_s(\tilde{\mathbf{q}}, 0))$.

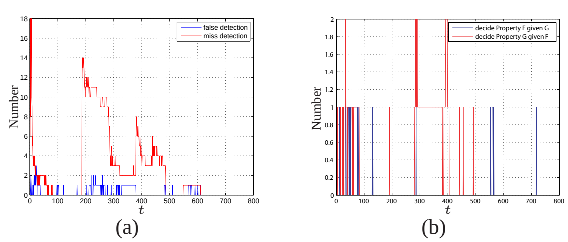

Fig. 4. (a) Posterior probabilities of object presence for every $\tilde{\mathbf{q}}$ within $\mathcal{D}$ at $t = 800$, (b) Number of false detections and miss detections, and (c) Number of deciding Property "F" given Property "G", and deciding Property "G" given Property "F".

Figure 5(a) shows the property of object 1, which has zero probability of having Property "F" with zero uncertainty, i.e., we are 100% sure that object 1 has Property "G". Figure 5(b) shows that object 2 has Property "F" with zero uncertainty.

IX. Conclusion

Based on a Bayesian-based probabilistic framework, a decision-making strategy was developed to guarantee the detection and classification of all objects in a domain using Bayesian risk analysis. Numerical simulations demonstrated the performance of the strategies. Future research will focus on locating and classifying dynamic objects with multiple autonomous sensor vehicles. The question of unknown environment geometries (i.e., unknown $\mathcal{D}$) will also be addressed. Objects with uniform distributions over the domain will be investigated, where the decision-making at one point is affected by all the decisions made at other points. SPRT method will also be investigated for the cases where no prior information is available.

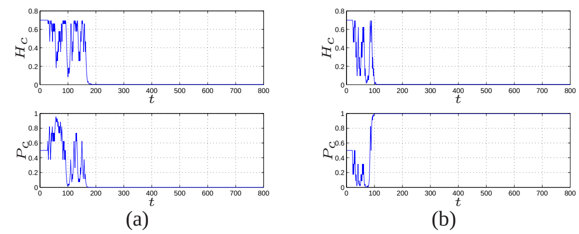

Fig. 5. Probability of object having Property "F" and the corresponding uncertainty function H_c for (a) object 1, and (b) object 2.

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," *American Control Conference*, 2009.

[3] J. R. Spletzer and C. J. Taylor, "Dynamic Sensor Planning and Control for Optimally Tracking Targets," *The International Journal of Robotics Research*, no. 1, pp. 7–20, January 2003.

[4] L. F. Bertuccelli and J. P. How, "Robust UAV Search for Environments with Imprecise Probability Maps," *Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference*, December 2005.

[5] ——, "Bayesian Forecasting in Multi-vehicle Search Operations," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, August 2006.

[6] J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous Map Building and Localization for an Autonomous Mobile Robot," in *IEEE/RSJ International Workshop on Intelligent Robots and Systems IROS '91*, Osaka, Japan, November 1991, pp. 1442 – 1447.

[7] 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.

[8] 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.

[9] R. Mahler, "Objective Functions for Bayeisan Control-Theoretic Sensor Management, I: Multitarget First-Moment Approximation," *Proceedings of IEEE Aerospace Conference*, 2003.

[10] 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.

[11] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer-Verlag, 1994.

[12] H. V. Poor and O. Hadjiliadis, *Quickest Detection*, 1st ed. Cambridge University Press, December 2008.

[13] I. I. Hussein, "A Kalman-filter based control strategy for dynamic coverage control," *Proceedings of the American Control Conference*, pp. 3271–3276, 2007.