Environmental Information Improves Robotic Search Performance

Harun Yetkin, Collin Lutz, and Daniel Stilwell
The Bradly Department of Electrical and Computer Engineering
Virginia Polytechnic Institute and State University, Blacksburg, VA 24060, USA
Email: {yetkinh, collin, stilwell}@vt.edu

***Abstract*—We address the problem where a mobile search agent seeks to find an unknown number of stationary objects distributed in a bounded search domain, and the search mission is subject to time/distance constraint. Our work accounts for false positives, false negatives and environmental uncertainty. We consider the case that the performance of a search sensor is dependent on the environment (e.g., clutter density), and therefore sensor performance is better in some locations than in others. For applications where environmental information can be acquired, we derive a decision-theoretic cost function to compute the locations where the environmental information should be acquired. We address the cases where environmental characterization is performed either by a separate vehicle or by the same vehicle that performs the search task.**

I. Introduction

We address search applications where a robotic system is to find an unknown number of objects in a bounded environment and within bounded time. We assume that the environment affects sensor performance, and that variation in the environment throughout the search area causes search performance in some locations to be better than in other locations. The principal contributions of our work show how guidance algorithms for search missions can incorporate stochastic knowledge of the environment to improve search performance, and for the case the environmental information can be acquired, we show where environmental information should be acquired in order to improve overall search performance. We address several cases: (1) environmental characterization is performed prior to a search mission, (2) environmental characterization is performed at the same time as the search mission by a separate asset than the search vehicle, and (3) environmental characterization is performed at the same time as search by the same vehicle that performs the search task.

We use a decision-theoretic value function that is associated with the accuracy of our estimate of the number of objects in the environment. Because search performance is dependent on the environment, knowledge of the environment can improve search performance due to better search plans. For example, one may choose to avoid searching areas that are known to contain excessive clutter and many false positives in favor of environments with few false positives. In situations where the environment is poorly known, efforts to acquire environmental information may lead to improved search effectiveness. We address the case that stochastic knowledge of the environment can be acquired, and we describe where the environment should be surveyed in order to improve overall search performance. One approach for selecting where to acquire environmental information is simply to characterize locations that yield the greatest reduction of uncertainty about the environment. In other words, one might seek to maximize change in entropy, which is often employed in similar applications [1–3]. In contrast, a primary contribution of this work is to show that environmental information should be acquired at the locations where the greatest reduction of uncertainty in anticipated search performance will occur, where we define search performance as the probability that the estimate for the number of objects in the environment is correct.

The remainder of this paper is organized as follows. A brief history of search theory and the benefit of acquiring environmental information in some search missions is provided in Section II. In Section III, we formulate the search problem and define the observation model. In Section IV, we define the objective function that maximizes the estimation accuracy. In Section V and Section VI, we describe our proposed cost function compute the locations where environmental characterization is performed. Section VII provides the numerical results that illustrate our approach.

II. Related Work

Search theory has its roots in numerous civilian and military applications. One of the tasks that arises often in the search literature is to find the optimal coverage paths where the search agent visits every location in the search environment exactly once. The goal is typically to minimize time-to-completion to achieve complete coverage of an environment (eg, [4]). In this study, however, we consider that the search mission is subject to a time or distance constraint. The practical interpretation of this constraint could be the limited battery capacity or the presence of time window to perform the mission. Due to this constraint, optimal search paths may not visit every location. Indeed, in some scenarios it is possible that some locations are visited more than once while other locations are never visited at all. Another task that arises often in the search literature is to find the most likely location of a single object or declare that the object is absent. The goal is either to maximize the probability of detecting the object or to minimize the expected time until a decision about presence or absence of the object is made ([5–7]). We note that in our problem we seek to find an

unknown number of objects as opposed to the case of finding a single target.

In a realistic search problem, there are certain limitations on successfully locating the objects such as imperfect sensor measurements or uncertain knowledge of the search environment. Noisy sensor measurements often include missed detections, i.e. failing to detect an object that is present, and false alarms, i.e. detection of an object that is not present. Local environmental conditions may also affect the number of false alarms and missed detections the sensor observes. All papers surveyed in [8] and [9] consider the effect of missed detections. However, the issue of false alarms is less often addressed, see for example [10–14]. In contrast, false alarms are addressed in [15] and [16], but uncertainty in the environment is not accounted for in these studies. We build upon prior work by accounting for false alarms, missed detections and uncertainty in the environment.

The effect of the environment on search performance is well-known. In subsea applications where sonar is used for search, variations in the seabed induce significant variation in probability of detection and probability of false alarm (see, for example, [17], [18]). For terrestrial applications using ground penetrating radar, search effectiveness is dependent on background clutter and soil properties (see, for example, [19–21]). A few studies in the literature aim to evaluate the benefit of reducing the uncertainty in the environment. In our prior work [22], we show that inaccurate estimate of sensor performance can lead to inaccurate estimate of search performance. For example, when the presumed probability of detection is higher than the actual probability of detection, the probability that all objects have been found during a search mission is exaggerated, and the search mission might be terminated too early. For the particular trial in [23], the experimental results show that the mine-hunting mission takes 40% less time when the environment is known compared to when there is no prior environmental information. These studies show that acquiring environmental information at some locations may significantly improve search performance. For the practical case that environmental characterization is not exhaustive, the challenge is to determine the locations where the environmental information should be acquired. In this study, we consider different cases where search and environmental characterization tasks are performed by separate assets or the same asset and derive a decision-theoretic cost function to compute the optimal paths for each case.

III. Problem Formulation

Search and environmental characterization are accomplished using different sensors that can be mounted on different vehicles or on the same vehicle. When the sensors are placed on different vehicles, the vehicle that possesses the search sensor is called *the search vehicle* and the vehicle that possesses the environmental characterization sensor is called *the environmental characterization vehicle*. When the search and environmental characterization sensors operate simultaneously on a single vehicle, we informally refer to the vehicle as *the search/environmental characterization vehicle*.

A. Preliminaries

Given a bounded search area $S \subset \mathbb{R}^2$ partitioned into K disjoint cells, $\{s_1, s_2, \dots, s_K\} = S$. We associate with each cell random variables X and E that represent the number of objects and the environmental conditions in the cell, respectively. We presume X_i is independent of X_j and E_i is independent of E_j when $i \neq j$. The objective of the search mission is to estimate $X_i, \dots, X_K$ by using a sensor to detect objects in each cell. We assume that sensor performance is dependent on the environment, and we use a stochastic description of sensor performance in the environment.

We assume that the environment in each cell is from a finite set of possible environments $\mathbf{e} = \{b_1, b_2, \dots, b_m\}$. That is, for all $s_i \in S$, the environment is $e_i \in \mathbf{e}$. We presume that the actual environmental condition in each cell is not known, but that a probability distribution is known for each cell. The environment probability distribution for each cell $s_i \in S$ is expressed $\Pi_i = [p_1(i), p_2(i), \dots, p_m(i)]$ where $p_j(i) = P(E_i = b_j)$ is the probability that the environment in cell s_i is b_j. We note that sum of probabilities for each cell is unity,

$$\sum_{j=1}^{m} p_j(i) = 1 \tag{1}$$

B. Sequential Bayesian update for the search vehicle

When the search vehicle visits a cell, it acquires a noisy observation $Z = z$ of the number of objects in the cell. The observation z may be less than the true number of objects because of missed detections, or it might be larger due to false alarms. In this study, we assume that the number of false alarms f and correct detections d are probabilistically independent. Hence, the value of the measurement z can be expressed

$$z = f + d$$

We model the likelihood of observing $Z = z$ objects when x is the true number of objects given that the environment is $b_j \in \mathbf{e}$. The sensor model is

$$P\big(z \mid x, b_j\big) = \sum_{l=0}^{\min(x,z)} P_D\big(d = l \mid x, b_j\big) P_F\big(f = z - l \mid b_j\big) \tag{2}$$

where $P_D(d = l \mid x, b_j)$ is the probability that the sensor detects l objects, and $P_F(f = k \mid b_j)$ is the probability that the sensor returns k false alarms. The sensor model is also described in [22], and is briefly presented here for clarity. Convolution of P_D and P_F in (2) follows from the assumption that the number of missed detections and false alarms are statistically independent. For numerical examples in Section VII, we model the probability of false alarms with a

geometric distribution, and the probability of correct detections with a Binomial distribution,

$$P_F\big(f = k \mid b_j\big) = (1-\alpha_j)\alpha_j^k \qquad k \geq 0 \tag{3}$$

$$P_D\big(d = l \mid x, b_j\big) = \binom{x}{l} D_j^l (1-D_j)^{x-l} \qquad 0 \leq l \leq x \tag{4}$$

where $0 < \alpha_j \leq 1$ denotes the probability of one or more false alarms, and $0 < D_j \leq 1$ denotes the probability of detection. Note that both α_j and D_j are assumed to vary as functions of the environment type b_j. Then, the likelihood is expressed

$$P\big(z \mid x, b_j\big) = \sum_{k=0}^{\min(x,z)} \binom{x}{k} D_j^k (1-D_j)^{x-k} (1-\alpha_j)\alpha_j^{z-k} \tag{5}$$

In subsea applications, probability of false alarm and probability of detection are sometimes modeled through a receiver operating characteristics (ROC) curve which describes the probability of detection as a function of probability of false alarm (see, for example, [24] and [25]). We note that our intention in this study is not to model the characteristics of a specific sensor type. We believe the geometric distribution in (3) efficiently models the intuition that fewer false alarms are more likely to occur than a greater number of false alarms. However, other expressions are also possible for the false alarm model, and our results do not depend on this specific false alarm model except for numerical illustrations.

We assume the number of objects in each cell is statistically independent of the environmental conditions in that cell. Thus, $P(x \mid b_j) = P(x)$. We use Bayesian update law to update the distribution $P(x \mid z, b_j)$ when z is observed.

$$P\big(x \mid z, b_j\big) = \frac{P\big(z \mid x, b_j\big) P\big(x\big)}{P\big(z \mid b_j\big)} \tag{6}$$

where $P(x)$ is our prior belief on X, $P(z \mid x, b_j)$ is the sensor characteristics as in (5), and $P(z \mid b_j)$ can be computed

$$P\big(z \mid b_j\big) = \sum_x P\big(z \mid x, b_j\big) P\big(x\big) \tag{7}$$

C. Sequential Bayesian update for the environmental characterization vehicle

When the characterization vehicle characterizes the environment at a location, it acquires the noisy observation $Y = y$ of the true environment in the cell. We assume the likelihood $P(Y = y \mid E = b_j)$ of observing a particular environment given the true environment $b_j \in \mathbf{e}$ is known before the characterization starts and it does not change. Insight on the form of the likelihood function arises from research on subsea bottom-type characterization, such as in [26]. We use a Bayesian update law to update the distribution $P(E = b_j \mid Y = y)$ when $Y = y$ is observed,

$$P\big(E = b_j \mid Y = y\big) = \frac{P\big(Y = y \mid E = b_j\big) P\big(E = b_j\big)}{P\big(Y = y\big)} \tag{8}$$

where $P(E = b_j)$ is the prior probability that the environment at the location is b_j, and

$$P\big(Y = y\big) = \sum_{j=1}^{m} P\big(Y = y \mid E = b_j\big) P\big(E = b_j\big) \tag{9}$$

D. Sequential Bayesian update for the search/environmental characterization vehicle

When the search sensor and the environmental characterization sensor operate simultaneously on a single vehicle, the noisy observations z and y are acquired simultaneously. Given z and y measurements acquired at a location, we represent the updated belief on the number of objects unconditioned on the environment

$$P\big(x \mid z, y\big) = \sum_{j=1}^{m} P\big(x \mid z, b_j\big) P\big(b_j \mid y\big) \tag{10}$$

where the posterior distribution $P(x \mid z, b_j)$ follows from (6), and $P(b_j \mid y)$ follows from (8).

IV. Path Planning For The Search Vehicle

We perform environmental characterization to improve the results of a search mission. To better understand the value of acquiring an environment measurement at a location, we seek to quantify effect of the acquired environment measurements on search results. Thus, in this section, we briefly present the value of searching a location and the objective function to compute the optimal search paths. For more details on path planning for the search vehicle, we refer the reader to our prior work in [22].

The goal of a search mission is to maximize the probability that the estimated number of objects in a cell is correct. Thus we seek to maximize the estimation accuracy. After the search vehicle visits a location, we compute the estimate $\delta(z)$ of the number of objects x at the location, based on the measurement z. When $\delta(z)$ is greater than x, we overestimate the number of objects, i.e. we declare more than the actual number of objects are present. When $\delta(z)$ is less than x, we underestimate the number of objects, i.e. we fail to declare some of the objects that are present. Both overestimation and underestimation may degrade the utility of the search results. Given the measured data z, we define the utility of the estimate $\delta(z)$ when x is the true number of objects

$$U\big(x, \delta(z)\big) = \begin{cases} 1 & \text{if } x = \delta(z) \\ 0 & \text{if } x \neq \delta(z) \end{cases} \tag{11}$$

which penalizes the deviations from true number of objects. The zero-one function in (11) emphasizes the fact that in some search missions, such as mine hunting, an incorrect estimate

has no utility regardless of how close the estimate is. The posterior expected utility of computing the estimate $\delta(z)$ when the environment is b_j is

$$\mathbb{E}\Big[U\big(x,\delta(z)\big) \mid z, b_j\Big] = \sum_x P\big(x \mid z, b_j\big) U\big(x,\delta(z)\big) \quad (12)$$

where the expectation is taken over the parameter space X with respect to the posterior distribution $P(x \mid z, b_j)$. Let $\delta^\star(z)$ be the estimator that maximizes the expected utility in (12). Such an estimator is called the Bayes estimator, and it is a function of the acquired measurement z

$$\delta^\star(z) = \arg\max_{\delta(z)} \mathbb{E}\Big[U\big(x,\delta(z)\big) \mid z, b_j\Big] \quad (13)$$

Then, the expected utility in (12) is the estimation accuracy conditioned on environment b_j that we seek to maximize, when the estimator $\delta(z)$ is the Bayes estimator in (13). Thus, the estimation accuracy conditioned on the environment b_j after acquiring the measurement z is

$$\mathbb{E}\Big[U\big(x,\delta^\star(z)\big) \mid z, b_j\Big] = \max_x P\big(X = x \mid z, b_j\big) \quad (14)$$

In order to assess the benefit of searching a cell, we compute estimation accuracy in (14) for each possible measurement $z \in Z$. This yields *expected* estimation accuracy of searching a cell conditioned on environment b_j,

$$\mathbb{E}\Big[U\big(x,\delta^\star(z)\big) \mid b_j\Big] = \sum_z P\big(z \mid b_j\big) \max_x P\big(X = x \mid z, b_j\big) \quad (15)$$

where $P\big(z \mid b_j\big)$, the probability of observing a particular measurement given the true environment, is defined in (7). Since deterministic knowledge on the environment is assumed to be unavailable, we compute the expected estimation accuracy unconditional on the environment. Averaging over the environments yields

$$\mathbb{E}\Big[U\big(x,\delta^\star(z)\big)\Big] = \sum_{j=1}^{m} P\big(b_j\big) \mathbb{E}\Big[U\big(x,\delta^\star(z)\big) \mid b_j\Big] \quad (16)$$

and we call this *anticipated* estimation accuracy.

Measurements from different cells are independent, and thus estimation accuracy for a path that passes through multiple cells is simply a product of the estimation accuracy for each cell in the path. The only challenge is book-keeping associated with the case that measurements are acquired from a single cell more than once. Let γ be a path of length N. A path may visit a cell more than once, and we say that the total number of distinct cells visited by path γ is $M \leq N$. Let $\gamma_d = \{q_1, \dots, q_M\}$ denote the set of distinct cells in γ where $q_1, \dots, q_M \in S$ are the search locations. For each $q_i \in \gamma_d$, the multiplicity of q_i, denoted m_i, is the number of occurrences of q_i^{th} cell in γ. We note that $\sum_{i=1}^{M} m_i = N$. The expected utility of traversing γ is

$$\mathbb{E}\Big[U\big(x,\delta^\star(\mathbf{z}_\gamma)\big)\Big] = \prod_{q_i \in \gamma_d} \mathbb{E}\Big[U\big(x_{q_i},\delta^\star(\mathbf{z}_{q_i})\big)\Big] \times \prod_{i \in S \setminus \gamma_d} \max_{x_i} P\big(x_i\big) \quad (17)$$

where $\mathbf{z}_{q_i}$ is the set of m_i independent search measurements acquired at the q_ith cell, and $\max_{x_i} P\big(x_i\big)$ is the certainty in the number of objects in cell i prior to acquiring new measurements. Let $\Gamma(k)$ denote the finite collection of N-length paths available to the search vehicle at time step k. Then, the optimal path is

$$\gamma^\star(k) = \arg\max_{\gamma \in \Gamma(k)} \mathbb{E}\Big[U\big(x,\delta^\star(\mathbf{z}_\gamma)\big)\Big] \quad (18)$$

V. Path Planning For The Environmental Characterization Vehicle

The primary objective of environmental characterization is to improve search performance. With additional information about the environment at a few locations, it might be possible to avoid searching locations where the sensor performs poorly in favor of places where the sensor performs well. We consider two specific cases: 1) environmental characterization is performed prior to search, 2) environmental characterization and search are performed simultaneously. In all cases, we assume that environmental characterization cannot be performed exhaustively due to limited resources. We derive the cost function for case 1 and then show that case 2 is a slight modification of case 1.

Suppose environmental characterization precedes search. It is intuitively appealing that locations for environmental characterization are selected to directly increase the probability that the estimate of the number of objects at a location is correct. That is, we assess the benefit of obtaining a particular environmental measurement $Y = y$ at a location by computing the conditional expected utility $\mathbb{E}\Big[U\big(x,\delta(z)\big)|Y = y\Big]$. To assess the expected benefit of a future environmental measurement, we average over all possible environmental measurements. However, we see directly that the result is the iterated expectation, and that the effect of environmental samples has been averaged out

$$\mathbb{E}_y\Big[\mathbb{E}_z\Big[U\big(x,\delta(z)\big) \mid Y = y\Big]\Big] = \mathbb{E}_z\Big[U\big(x,\delta(z)\big)\Big] \quad (19)$$

Thus, plans for environmental characterization do not directly improve the expected performance of a search plan. In order to assess the value of environmental characterization when planning search missions, a fundamentally different approach is needed.

A. Environmental loss function

Due to the uncertainty in the environment and the noise in environmental observations, estimation accuracy after visiting a cell in (16) may be different than actual estimation accuracy if the true environment were unambiguously known. In [22], we discuss the effect of environment uncertainty on search

results, and show that deviations from true environment result in deviations from actual estimation accuracy and degrade the search performance. In this paper, we extend our findings in [22] to select the best locations to conduct environmental surveys. Our approach is to define a linear loss function that penalizes deviations from the actual expected estimation accuracy for each cell.

To formally define the loss function, we first introduce a preference ordering $\preceq$ on environments. Suppose there is a finite set of environments $\mathbf{e} = \{b_1, b_2, \dots, b_m\}$. For the notational convenience, let $V\big(b_j\big)$ denote the expected estimation accuracy conditioned on environment b_j

$$V\big(b_j\big) = \mathbb{E}\Big[U\big(x, \delta^\star(z)\big) \mid b_j\Big] \tag{20}$$

in (15). We say the environment b_i is more preferred for the search than the environment b_j if the expected estimation accuracy conditioned on b_i is greater than the expected estimation accuracy conditioned on b_j. That is, we say that $b_i \preceq b_j$ if and only if $V(b_i) \leq V(b_j)$. If for some $b_i, b_j \in \mathbf{e}$ when $i \neq j$ we have $V(b_i) = V(b_j)$, then $b_i = b_j$. If $V(b_i) \neq V(b_j)$, we say b_i and b_j are distinct environments. Suppose the environments $b_1, \dots, b_m$ are distinct and *ordered* so that $b_1 \prec b_2 \prec \dots \prec b_m$, let e be the true environment in a cell, and let $d(y)$ be an estimate of the environment based on measurement y and the prior distribution on the number of objects $P(x)$. When the true environment is $e \in \mathbf{e}$, the loss due to the estimate $d(y)$ is defined

$$W\big(e, d(y)\big) = \begin{cases} c_1\Big(V\big(e\big) - V\big(d(y)\big)\Big) & \text{if } d(y) \preceq e \\ c_2\Big(V\big(d(y)\big) - V\big(e\big)\Big) & \text{if } d(y) \succ e \end{cases} \tag{21}$$

where $c_1, c_2 > 0$ are the relative costs of over and underestimation. Underestimating the environment, $d(y) \prec e$, may result in unnecessary extra visits to improve the belief on the number of objects at a location. However, overestimating the environment, $d(y) \succ e$, may yield to inaccurate estimates on the number of objects. In some search applications, such as mine-hunting, overestimation is less preferred to underestimation. Thus, we may assign the relative costs such that $c_1 < c_2$.

Given the environment measurement y, the posterior expected loss of computing the environment estimate $d(y)$ is

$$\mathbb{E}\Big[W\big(e, d(y)\big) \mid y\Big] = \sum_{j=1}^{m} P\big(b_j \mid y\big) W\big(b_j, d(y)\big) \tag{22}$$

where $P\big(b_j|y\big)$ is the updated probability that b_j is the environment at the location after observing the environmental measurement y. We choose the estimator that minimizes the expected loss in (22). Let $d^\star(y)$ be the Bayes estimator such that

$$d^\star(y) = \arg\min_{d(y)} \mathbb{E}\Big[W\big(e, d(y)\big) \mid y\Big] \tag{23}$$

For the specific loss function in (21), the Bayes estimator that minimizes the expected loss can be computed from the cumulative distribution (see [27]),

$$\begin{aligned} F_{E|y}(b_n) &:= P\big(E \preceq b_n \mid Y = y\big) \\ &= \sum_{i=1}^{n} P\big(E = b_i \mid Y = y\big) \end{aligned} \tag{24}$$

for $n \in \{1, 2, \dots, m\}$ where (24) follows since $b_1 \prec b_2 \prec \dots \prec b_m$. For

$$\ell = \max\left\{ n \in \{1, 2, \dots, m-1\} : F_{E|y}(b_n) \leq \frac{c_1}{c_1 + c_2} \right\}, \tag{25}$$

the optimal estimate is

$$d^\star(y) = b_{\ell+1} \tag{26}$$

B. Path planning

A benefit of environmental surveys is to reduce the gap between actual search performance if the true environment is known which we refer to as the expected estimation accuracy (15), and anticipated search performance when only probabilistic information on the environment is known, that we refer to as the anticipated estimation accuracy in (16). When a location is not visited by the search vehicle during a search mission, acquiring an environment measurement at that location will not tighten this gap. Thus, there is no benefit in characterizing the environment at locations where search will not occur. From the loss function in (21), computing the estimate (23) of the environment $e \in \{b_1, \dots, b_m\}$ after acquiring environment measurement y yields the conditional expected loss

$$\mathbb{E}\Big[W\big(e, d^\star(y)\big) \mid y\Big] = \sum_{j=1}^{m} P\big(E = b_j \mid Y = y\big)\ W\big(b_j, d^\star(y)\big) \tag{27}$$

that quantifies the amount of uncertainty in anticipated estimation accuracy after acquiring $Y = y$. Informally speaking, the prior loss before acquiring an environment measurement represents the prior uncertainty and the conditional expected loss in (27) represents the posterior uncertainty in search performance. We denote the gain of acquiring environment measurement y_i in cell i by $G(y_i)$

$$G\big(y_i\big) = \mathbf{1}_{\gamma^\star(y_i)}(s_i)\ \Big(\mathbb{E}\Big[W\big(e_i, d\big)\Big] - \mathbb{E}\Big[W\big(e_i, d(y_i)\big) \mid y_i\Big]\Big) \tag{28}$$

which is the reduction of uncertainty in anticipated estimation accuracy given that cell i is visited by the search vehicle. In (28), the notation $\gamma^\star(y_i)$ denotes the best path for the search vehicle after acquiring the environment measurement y_i, and the indicator function $\mathbf{1}_{\gamma(y_i)}(s_i) : s_i \to \{0, 1\}$ is defined

$$\mathbf{1}_{\gamma(y_i)}(s_i) = \begin{cases} 1 & s_i \in \gamma(y_i) \\ 0 & s_i \notin \gamma(y_i) \end{cases} \tag{29}$$

Let η be a length N path of the environmental characterization vehicle and $\eta_d = \{q_1, q_2, \dots, q_M\}$ be the set of distinct cells in η. Let $\mathbf{y}(q_i)$ be the set of independent environment measurements acquired at q_ith cell, and let $\mathrm{H}(k)$ denote the finite collection of paths available to the characterization vehicle at time step k. Then, the expected gain of traversing η is

$$\mathbb{E}\Big[G\big(\mathbf{y}_\eta\big)\Big] = \sum_{q_i \in \eta_d} \sum_{\mathbf{y}_\eta} \mathbf{1}_{\gamma^\star(\mathbf{y}_\eta)}(s_{q_i})\ P\big(Y_\eta = \mathbf{y}_\eta\big) \times \Big(\mathbb{E}\Big[W\big(e_{q_i}, d\big)\Big] - \mathbb{E}\Big[W\big(e_{q_i}, d(\mathbf{y}_{q_i})\big) \mid \mathbf{y}_{q_i}\Big]\Big) \tag{30}$$

and the optimal path is

$$\eta^\star(k) = \arg\max_{\eta \in \mathrm{H}(k)} \mathbb{E}\Big[G\big(\mathbf{y}_\eta\big)\Big] \tag{31}$$

C. Entropy change maximization

Our approach to select locations for environmental characterization minimizes uncertainty in anticipated search performance. A common approach to address similar applications is to maximize change in entropy, see for example, [1–3]. Let $H(E)$ denote the prior entropy of the environmental distribution.

$$H\big(E\big) = -\sum_{j=1}^{m} P\big(E = b_j\big) \log P\big(E = b_j\big) \tag{32}$$

The amount of change in the entropy for future environment measurement y can be computed by

$$J\big(E\big) = H\big(E\big) - \sum_{y} H\big(E = e \mid Y = y\big) \tag{33}$$

where $H(E = e|Y = y)$ is the posterior entropy for E after acquiring y. Then, the best path computed via entropy change maximization method is

$$\eta^\star = \arg\max_{\eta \in \mathrm{H}} J_\eta(E) \tag{34}$$

D. Both vehicles operate simultaneously

Path planing for environmental characterization for the case that environmental characterization and search are accomplished by different vehicles that operate at the same time is addressed similarly to the case in Section V-B where environmental characterization precedes search. We again compute the gain of acquiring environment measurement y_i in cell s_i as in (28). However, the indicator function in (29) is modified to account for the possible situations where the search vehicle visits a location prior to the environmental characterization vehicle, in which case environmental characterization cannot influence search plans. The indicator function in (29) is unity when the location to be characterized is in the search vehicle's trajectory even if the environmental characterization vehicle arrives *after* the search vehicle. To overcome this problem, we introduce an indexing for each location. Let cell s_i appear in both vehicles's path and let $r_\eta(i)$ and $r_\gamma(i)$ denote when the cell appears in the characterization vehicle's path and in the search vehicle's path, respectively. For example, if cell s_5 is the 4th cell the characterization vehicle visits and the 2nd cell the search vehicle visits, then $r_\eta(5) = 4$ and $r_\gamma(5) = 2$. We note that when $r_\eta(i) > r_\gamma(i)$, there is no gain of characterizing the cell since the increase in the uncertainty of the environment will not affect the search results. Thus, the modification yields

$$\mathbf{1}_{\gamma(\mathbf{y}_\eta)}(s_i) = \begin{cases} 1 & s_i \in \gamma(\mathbf{y}_\eta) \text{ and } r_\eta(i) \le r_\gamma(i) \\ 0 & \text{otherwise} \end{cases} \tag{35}$$

VI. Path Planning For The Search/environmental characterization Vehicle

We lastly consider the case that a single vehicle is equipped with an environmental characterization sensor and a search sensor, and that both sensors can operate simultaneously. We again seek to maximize estimation accuracy. Unlike Section III where the search vehicle aims to maximize the estimation accuracy with only the search measurements, we now acquire a search measurement z and an environmental measurement y when the vehicle visits a location. Thus, the path strategies that do not address the acquisition of environmental measurements do not apply to this case.

Let $\mathcal{V}\big(b_j\big)$ denote estimation accuracy conditioned on the environment b_j in (14),

$$\mathcal{V}\big(b_j\big) = \max_x P\big(x \mid z, b_j\big) \tag{36}$$

and let $\mathbf{e} = \{b_1, b_2, \dots, b_m\}$ be a set of environments. We say $b_i \preceq b_j$ if and only if $\mathcal{V}\big(b_i\big) \le \mathcal{V}\big(b_j\big)$. Note $\mathcal{V}\big(b_j\big)$ in (36) is the accuracy of the estimate of the number of objects at a location while $V\big(b_j\big)$ in (20) is the expected accuracy when a measurement z has not yet been acquired. Suppose the environments $b_1, \dots, b_m$ are distinct and *ordered* (as defined in Section V-A) so that $b_1 \prec b_2 \prec \dots \prec b_m$, let e be the true environment in a cell, and let $d(y)$ be an estimate of the environment based on the environment measurement y and the prior distribution on the number of objects $P(x)$. When the true environment is $e \in \mathbf{e}$, the loss due to the estimate $d(y)$ is defined

$$\mathcal{W}\big(e, d(y)\big) = \begin{cases} c_1\big(\mathcal{V}\big(e\big) - \mathcal{V}\big(d(y)\big)\big) & \text{if } d(y) \preceq e \\ c_2\big(\mathcal{V}\big(d(y)\big) - \mathcal{V}\big(e\big)\big) & \text{if } d(y) \succ e \end{cases} \tag{37}$$

where $c_1, c_2 > 0$ are again the relative costs of over and underestimation. Then, the posterior expected loss of computing the environment estimate $d(y)$, and the corresponding Bayes estimator $d^\star(y)$ are

$$\mathbb{E}\Big[\mathcal{W}\big(e, d(y)\big) \mid z, y\Big] = \sum_{j=1}^{m} P\big(b_j \mid y\big) \mathcal{W}\big(b_j, d(y)\big) \tag{38}$$

$$d^\star(y) = \arg\min_{d(y)} \mathbb{E}\Big[\mathcal{W}\big(e, d(y)\big) \mid z, y\Big] \tag{39}$$

Given measurement z and the estimate $d^\star(y) \in \{b_1, \dots, b_m\}$ from measurement y, the probability that the estimate of the number of objects at a location is correct is computed from

$$\mathbb{E}\Big[U\big(x, \delta(z)\big) \mid z, d^\star(y)\Big] = \max_x P\big(x \mid z, d^\star(y)\big) \tag{40}$$

In order to assess the benefit of visiting a location, we compute the *estimated* estimation accuracy in (40) for each possible set of observations $z \in Z,\ y \in Y$. Then, the expected estimation accuracy before visiting a location can be computed

$$\mathbb{E}\Big[\mathcal{V}\big(d(y)\big)\Big] = \sum_z \sum_y P\big(z, y\big)\ \max_x P\big(x \mid z, d^\star(y)\big) \tag{41}$$

where

$$P\big(z, y\big) = \sum_x \sum_{b_j} P\big(x \mid z, b_j\big) P\big(y \mid b_j\big) P\big(x\big) P\big(b_j\big) \tag{42}$$

We again consider the $N-$length path γ and the set γ_d of distinct cells in γ. Let $\mathbf{y}_{q_i}$ be the set of independent environment measurements acquired at q_ith cell. The expected estimation accuracy for traversing γ is

$$\mathbb{E}\Big[\mathcal{V}\big(d(\mathbf{y}_\gamma)\big)\Big] = \prod_{q_i \in \gamma_d} \mathbb{E}\Big[\mathcal{V}\big(d(\mathbf{y}_{q_i})\big)\Big] \times \prod_{i \in S \setminus \gamma_d} \max_{x_i} P\big(x_i\big) \tag{43}$$

Let $\Gamma(k)$ denote the finite collection of N-length paths available to the vehicle at time step k. Then, the optimal path is

$$\gamma^\star(k) = \arg\max_{\gamma \in \Gamma(k)} \mathbb{E}\Big[\mathcal{V}\big(d(\mathbf{y}_\gamma)\big)\Big] \tag{44}$$

VII. Numerical Results

In this section, we present simulation results that show the efficacy of the proposed search and environmental characterization strategies. Our numerical illustrations aim to evaluate search performance when environmental measurements are available for simplistic scenarios that are inspired by subsea mine-hunting missions. We present numerical illustrations for two scenarios. In one case, search and environmental characterization sensors are on different vehicles and environmental characterization is performed prior to search. In the other case, search and environmental characterization sensors are on the same vehicle and both activities occur simultaneously. When each sensor operate on separate vehicles, our proposed approach maximizes the reduction of uncertainty in search performance (30). Thus, our approach should, on average, display less anticipated estimation accuracy error than other approaches.

We divide the bounded search area S into a grid with 10×10 non-intersecting cells. For each cell $s_i \in S$, we assume there is $0 \leq x_i \leq L$ number of objects bounded above by L. In mine hunting missions, x_i represents the number of mines residing in cell s_i. It is assumed that no prior information exists about the number of objects in any cell. We note that L is typically not known beforehand; however, letting L be a sufficiently large number will capture all realistic scenarios. In our simulations, $L = 2$. The performance of the search sensor is dependent on the environmental conditions. The particular sensor model that we use for the numerical illustrations is (5). We assume there are three candidate environments in the search area, $\mathbf{e} = \{b_1, b_2, b_3\}$. The probability of detection, D, and the probability of at least one false alarm, α, for each environment are $D = 0.65$ and $\alpha = 0.4$ for environment b_1, $D = 0.8$ and $\alpha = 0.3$ for environment b_2, and $D = 0.95$ and $\alpha = 0.05$ for environment b_3. Note that the information about the number of objects revealed after searching a cell increases with increasing probability of detection and decreases with increasing probability of false alarm. Thus, environment b_1 is the least and environment b_3 is the most informative. Let the sensor model for environmental characterization be such that

$$P\big(E = b_i \mid Y = b_j\big) = a_{ij} \quad \text{for all} \quad i, j \in \{1, 2, 3\}$$

where a_{ii} is the probability of observing the true environment b_i. For convenience, we consider that $a_{ij} = a_{ik}$ for $j, k \neq i$.

The mission for a vehicle is subject to a time or distance constraint. We call this constraint the *mission length*, i.e. total number of cells a vehicle can travel through during a mission. For a large mission length, the computational expense of finding the optimal trajectory may not be practically feasible. We instead define the *path length*, i.e. total number of lookahead cells considered for planning the path. As in a typical receding horizon approach, we plan a path with path length N, move part way along that path, and then replan a new N-length path.

In the subsea applications that inspire our numerical illustrations, autonomous underwater vehicles (AUVs) are typically equipped with a side scan-sonar. Because side-scan sonar works poorly while the vehicle is turning, we associate a cost with vehicle turns and constrain the motion of the vehicle in a way that the vehicle can only move forward towards the next grid cell in the row. In order to account for the effect of turns, the vehicle passes through cells that are outside of the search area when transitioning between rows. Passing through cells that are outside the search area requires time but does not improve estimation accuracy since no measurements are acquired.

A. Numerical illustrations

Fig. 1 shows a search area that is partitioned into parts A1 through A6. For each part, the corresponding probability distribution $\Pi = [p_1, p_2, p_3]$ is given, where p_j is the probability that the environment is b_j. For example, for the cells labeled A2, there is a 0.5 probability that the environment is b_1, a 0.3 probability that the environment is b_2, and a 0.2 probability that the environment is b_3. For the observation model of the environmental characterization sensor, we assign $a_{11} = 0.82, a_{22} = 0.84$, and $a_{33} = 0.88$. When

$j, k \neq i$, we assign $a_{ij} = a_{ik}$. The relative costs of over and underestimating the environmental conditions are $c_1 = 1$ and $c_2 = 3$ so that overestimation is penalized more than underestimation. The mission length is 60 for the search and search/environmental characterization vehicles and 40 for the characterization vehicle.

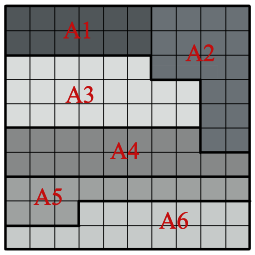

A1	$\rightarrow$	$[1.00, 0.00, 0.00]$
A2	$\rightarrow$	$[0.50, 0.30, 0.20]$
A3	$\rightarrow$	$[0.10, 0.20, 0.70]$
A4	$\rightarrow$	$[0.30, 0.40, 0.30]$
A5	$\rightarrow$	$[0.25, 0.20, 0.55]$
A6	$\rightarrow$	$[0.35, 0.00, 0.65]$

Figure 1: Search area and cell-wise environment distributions

We consider two scenarios. In one scenario the search and the environmental characterization sensors operate on the same vehicle, and in the other scenario they operate on separate vehicles. When the sensors operate on separate vehicles, the objective of the search vehicle is to maximize anticipated estimation accuracy in (17), and the objective of the characterization vehicle is to maximize the expected gain of characterization in (30). On the other hand, when both sensors operate on the same vehicle, the objective of the vehicle is to maximize expected estimation accuracy in (43). Recall that estimation accuracy is the probability that our estimate of the number of objects is correct. For both scenarios, after acquiring environmental information, we obtain an estimate of the environmental conditions that minimizes the expected loss due to the prior uncertainty in the environment (23).

We define the *error in search performance* after a mission as the difference between the actual estimation accuracy when the true environment is known and the anticipated estimation accuracy when the environment is uncertain. We use the error in search performance as a measure to evaluate the efficacy of the proposed approaches in each scenario, and show that the proposed approach yields smaller search performance error, which is predicted by our selection of cost function. We also show that search performance (probability of correct estimate) increases modestly, although our approach does not directly seek to increase estimation accuracy.

When the sensors are on separate vehicles and characterization precedes search, we compare the proposed approach in (30) with the entropy change maximization method described in Section V-C. We consider that the path length for the characterization vehicle is 8. Fig. 2e shows the trajectory for the environmental characterization vehicle when using our proposed approach in (30), which seeks to characterize the environment in locations that are expected to yield the greatest reduction of uncertainty in anticipated estimation accuracy. In contrast, Figure 2f shows the path of an environmental characterization vehicle when the path is selected by maximizing the change in entropy of the environmental distributions in (34). Neither environmental characterization path visits A1 because the environments in those locations are completely known. We note that the environmental characterization path in Figure 2e that was selected using our approach does not visit the most uncertain environments. We find in practice that it tends to visit environments that are both uncertain *and* likely to be where follow-on search missions will occur.

When both sensors operate on the same vehicle, we compare the proposed approach in (44) with entropy change maximization method and with a mowing-the-lawn approach. The latter arises often in subsea applications such as mine-hunting. The path length is equal to the mission length, path length is 60. We note that the entropy change maximization method described in Section V-C accounts only for the entropy change of the environmental distributions. However, when both sensors are placed on the same vehicle, the vehicle acquires environmental measurement and search measurement simultaneously. Thus, we modify (34) as

$$\eta^\star = \arg\max_{\eta \in \mathrm{H}} J_\eta(X) + \beta J_\eta(E)$$

where $J(X)$ denotes the entropy change in X, the number of objects, and β is the relative weight of the entropy change in E compared to the entropy change in X. Since the objective is to reduce the uncertainty in the number of objects, we choose $0 < \beta < 1$. Fig. 2c shows the mowing-the-lawn trajectory where the vehicle travels through the search area back and forth without planning the path until the mission length is met. Fig. 2a shows the trajectory for the proposed approach and Fig. 2b shows the trajectory for the entropy change maximization method with $\beta = 0.5$. We also compute the optimal search trajectory when there is no environmental characterization to show the value of acquiring environmental information. The corresponding trajectory for this case is shown in Fig. 2d.

We expect that for both scenarios our proposed approach yields better search performance compared to the other path planning strategies in Fig. 2. That is, when search and environmental characterization missions are performed on the same vehicle, if the search locations are selected using our approach as in Fig. 2a, the search performance is expected to be better compared to selecting the locations using entropy change maximization as in Fig. 2b or mowing the lawn as in Fig. 2c. When the sensors operate on separate vehicles, we expect that selecting the characterization locations using our approach as in Fig. 2e will yield greater improvement in the performance of a follow-on search mission compared to selecting the locations using entropy change maximization as in Fig. 2f.

Search performance after a mission depends on the acquired observations during the mission. Thus, we conduct Monte Carlo simulations to assess the effects due to random nature of observations. For each cell in the search area, we randomly

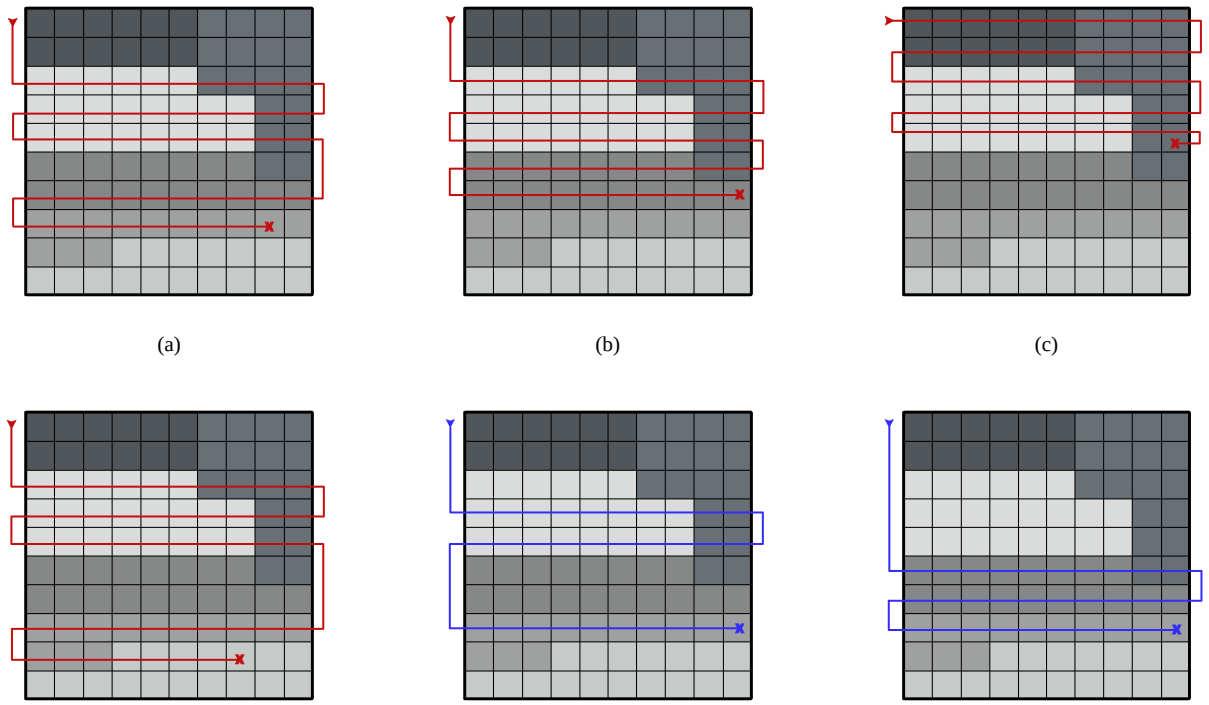

(a) (b) (c)

(d) (e) (f)

Figure 2: Optimal trajectories for search and characterization. Figures (a-c): trajectories for the case both sensors operate on the same vehicle when (a) proposed approach is employed (b) entropy change maximization method is employed, and (c) the mowing-the-lawn approach is employed. Figures (e-f): characterization vehicle's trajectories for the case the sensors operate on separate vehicles when the characterization locations are selected (e) by our proposed approach, (f) by entropy change maximization method. Figure (d) shows the search vehicle's trajectory when no environment information acquired.

generate the true environment e from the environmental distributions in Fig. 1 and the true number of objects x from a uniform distribution. Assuming that a cell can be visited by a vehicle at most k times, we randomly generate the set of search measurements z and the set of environmental measurements y from the sensor models $P(z \mid x, e)$ and $P(y \mid e)$ given true environment e and true number of objects x. When a vehicle visits a location, it acquires randomly generated observation(s). For each test, we compute the anticipated search performance and the actual search performance. Note that the actual search performance can be computed since the true environment is assumed to be known. We then compute the error in search performance which is the difference between the anticipated search performance and the actual search performance. We show that the error in search performance is reduced when our proposed approach is employed.

Both sensors operating simultaneously on a single vehicle

Fig. 3 shows the results after 10000 iterations for the case both sensors operate on the same vehicle. Fig. 3a on the left is the percent of occurrences of the error in search performance, and Fig. 3b on the right is the percent of occurrences of the actual search performance. The subplots from top to bottom are the results when 1) our proposed approach is employed, 2) the entropy change maximization method is employed, 3) the mowing-the-lawn approach is employed, and 4) environmental information is not available so that the vehicle acquires only the search measurements. The average value of results for each test is also shown in the plots. For convenience, the displayed results are the negative log of the computed search performance (estimation accuracy). Thus, smaller values imply better search performance. The simulations show that

- The proposed approach yields smaller error in search performance compared to the entropy change maximization and mowing-the-lawn. In addition, the actual search performance when using our approach is no worse than the actual search performance when using the other methods.
- When environmental information is acquired, the error in search performance is significantly smaller. Hence, a benefit of characterizing the environment is to better anticipate the true search performance.
- Average error for the mowing-the-lawn approach is smaller than the average error for entropy change maximization method. This is because mowing-the-lawn approach visits A1 that has no uncertainty in the envi-

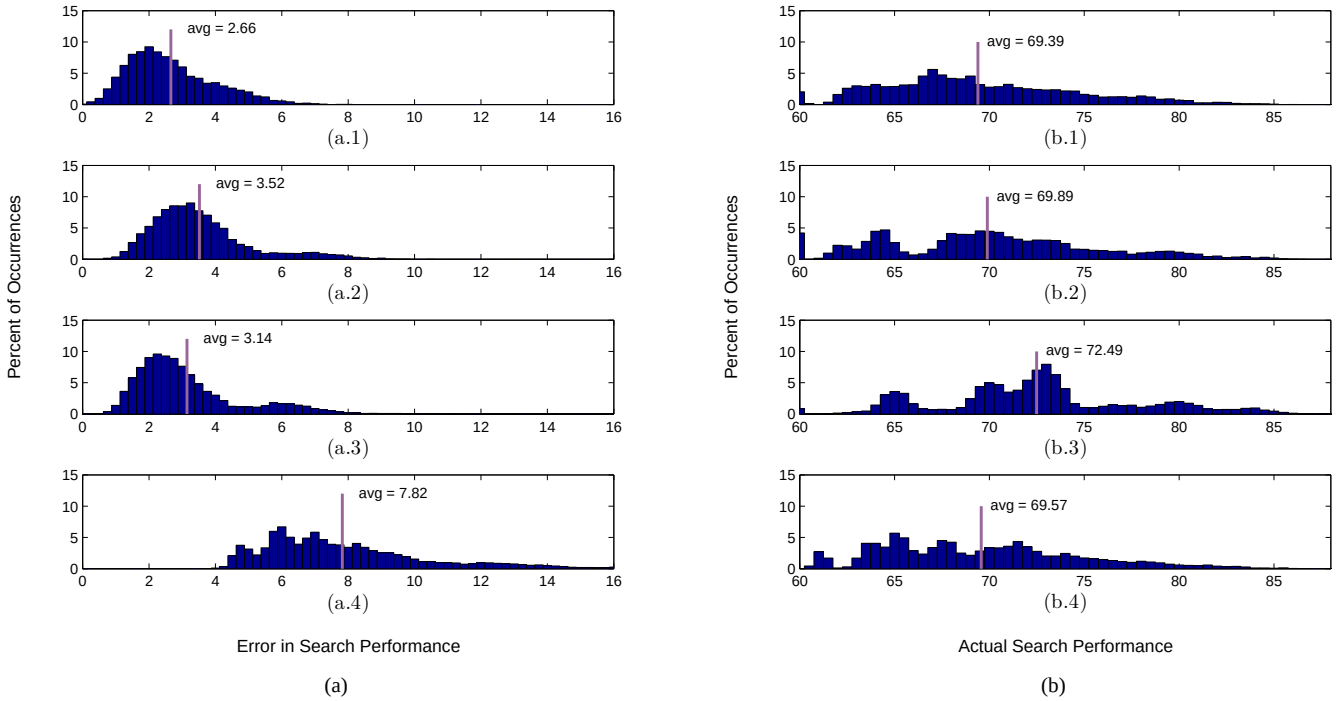

Figure 3: Percent of occurrences for (a) error in search performance and (b) actual search performance when both sensors operate on *the same vehicle*. From top to bottom, (a.1) and (b.1) correspond to the proposed approach, (a.2) and (b.2) correspond to the entropy change maximization method, (a.3) and (b.3) correspond to the mowing-the-lawn approach, and (a.4) and (b.4) correspond to the case where environment information is not available. Note that the horizontal axes is the negative log of the results. Thus, smaller values imply better search performance.

ronment while the entropy change maximization method visits A4 where the environmental uncertainty is greatest. However, as the environment in A1 is the least informative, the average actual search performance for mowing-the-lawn approach is the worst among all methods.

Each sensor on separate vehicles

The results when search and environmental characterization tasks are performed on separate vehicles are plotted in Fig. 4. Again, the left plot is the percent of occurrences of the error in search performance, and the right plot is the percent of occurrences of actual search performance. The subplots from top to bottom are the results when 1) the locations that yield the greatest reduction of uncertainty in search performance are characterized, 2) the locations that maximize the entropy change are characterized, and 3) there is no environmental characterization and the search vehicle plans its path by using the prior environmental distributions. We note that Fig. 4a.3 and Fig. 4b.3 are the same plots given in Fig. 3a.4 and Fig. 3b.4, and we show them here for convenience of comparison. It is seen that

- The average error is smaller when environmental characterization is performed at the locations selected by our proposed approach. This is expected since our approach directly penalizes the variation from the true search performance.
- The average error when the sensors are on different vehicles is higher than when both sensors operate on the same vehicle since the search vehicle may search the locations that are not characterized. On the other hand, this results in average actual search performance to be better since the search vehicle can skip the locations that are characterized and found to be uninteresting for search.

The results of Monte Carlo simulations show that our proposed approaches to select the characterization locations outperform the other strategies that frequently exist in the literature. We note that the case where the characterization vehicle and the search vehicle operates simultaneously is a subtle modification of the case characterization precedes search that we illustrate here, and the corresponding path would be the same path that is shown in Fig. 2e. Note that for each characterization location in Fig. 2e, depending on the acquired environmental information, the search vehicle either does not sample from that location or visits that location after it is characterized. Hence, the expected gain of characterizing these locations will be the same regardless of whether characterization precedes search or both vehicles perform simultaneously.

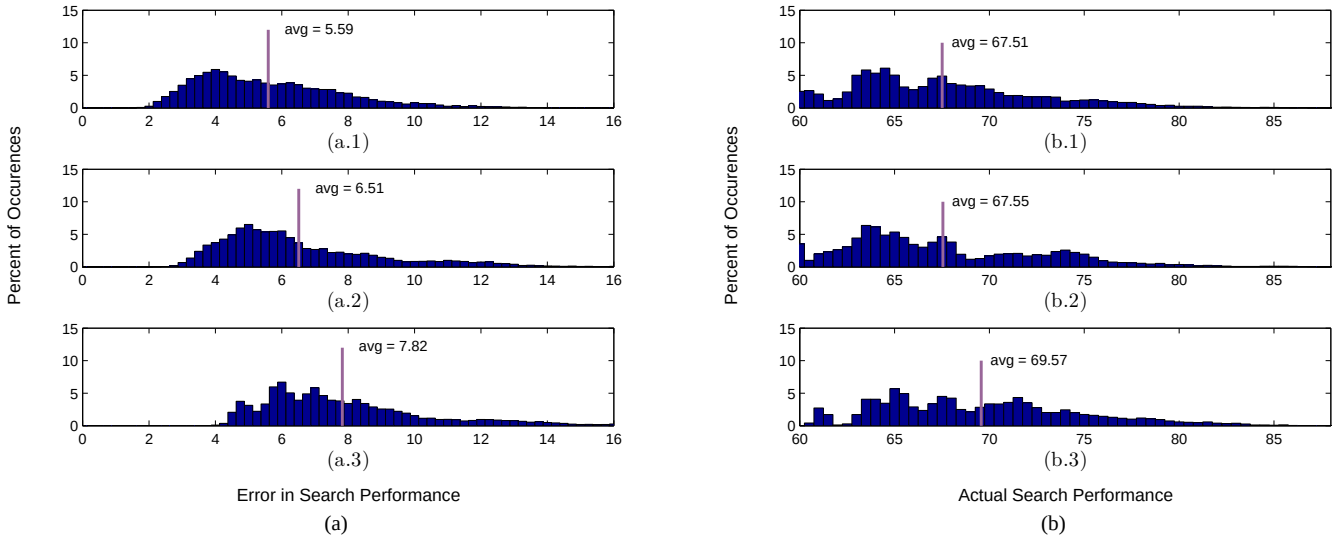

Figure 4: Percent of occurrences for (a) error in search performance and (b) actual search performance when search and characterization are performed on *separate vehicles*. From top to bottom, (a.1) and (b.1) correspond to the proposed approach, (a.2) and (b.2) correspond to the entropy change maximization method, and (a.3) and (b.3) correspond to the case where environment information is not available. Note that the horizontal axes is the negative log of the results. Thus, smaller values imply better search performance.

VIII. Conclusions

In this paper, we address the case where environmental information can be acquired to improve the performance of a search mission. We consider different scenarios where the search sensor and the environmental characterization sensor can be placed on the same vehicle or on separate vehicles. For each scenario, we derive a decision-theoretic cost function to compute the locations where environmental information should be acquired. We show that when the search sensor and the environmental characterization sensor are placed on separate vehicles, environmental information should be acquired at the locations where the greatest reduction of the uncertainty in anticipated estimation accuracy will occur. For the case where the search sensor and the environmental characterization sensor are placed on the same vehicle, we show that the expected estimation accuracy should be maximized. The results of the numerical illustrations show that for each scenario, our proposed approaches yield smaller error in search performance.

Acknowledgements

The authors gratefully acknowledge the support of the Office of Naval Research via grants N00014-12-1-0055 and N00014-16-1-2092. The assistance provided by Dr. Hongxiao Zhu (Department of Statistics, Virginia Tech) is greatly appreciated.

References

[1] C. Papadimitriou, J. L. Beck, and S.-K. Au, "Entropy-based optimal sensor location for structural model updating," *Journal of Vibration and Control*, vol. 6, no. 5, pp. 781–800, 2000.

[2] M. C. Coleman and D. E. Block, "Nonlinear experimental design using bayesian regularized neural networks," *AIChE journal*, vol. 53, no. 6, pp. 1496–1509, 2007.

[3] A. Elfes, "Dynamic control of robot perception using multi-property inference grids," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*. IEEE, 1992, pp. 2561–2567.

[4] H. Choset, "Coverage for robotics–a survey of recent results," *Annals of mathematics and artificial intelligence*, vol. 31, no. 1-4, pp. 113–126, 2001.

[5] J. Berger, N. Lo, and M. Barkaoui, "Static target search path planning optimization with heterogeneous agents," *Annals of Operations Research*, pp. 1–18, 2016.

[6] T. H. Chung and J. W. Burdick, "A decision-making framework for control strategies in probabilistic search," in *Robotics and Automation, 2007 IEEE International Conference on*. IEEE, 2007, pp. 4386–4393.

[7] H. Sato and J. O. Royset, "Path optimization for the resource-constrained searcher," *Naval Research Logistics (NRL)*, vol. 57, no. 5, pp. 422–440, 2010.

[8] S. J. Benkoski, M. G. Monticino, and J. R. Weisinger, "A survey of the search theory literature," *Naval Research Logistics (NRL)*, vol. 38, no. 4, pp. 469–494, 1991.

[9] T. H. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics," *Autonomous robots*, vol. 31, no. 4, pp. 299–316, 2011.

[10] M. C. Chew Jr, "A sequential search procedure," *The Annals of Mathematical Statistics*, pp. 494–502, 1967.

[11] J. B. Kadane, "Discrete search and the neyman-pearson lemma," *Journal of Mathematical Analysis and Applications*, vol. 22, no. 1, pp. 156–171, 1968.

[12] N.-O. Song and D. Teneketzis, "Discrete search with multiple sensors," *Mathematical Methods of Operations Research*, vol. 60, no. 1, pp. 1–13, 2004.

[13] J. Tisdale, Z. Kim, and J. K. Hedrick, "Autonomous UAV path planning and estimation," *Robotics & Automation Magazine, IEEE*, vol. 16, no. 2, pp. 35–42, 2009.

[14] G. Hollinger, S. Yerramalli, S. Singh, U. Mitra, G. S. Sukhatme *et al.*, "Distributed data fusion for multirobot search," *Robotics, IEEE Transactions on*, vol. 31, no. 1, pp. 55–66, 2015.

[15] T. H. Chung and J. W. Burdick, "Analysis of search decision making

using probabilistic search strategies," *Robotics, IEEE Transactions on*, vol. 28, no. 1, pp. 132–144, 2012.
[16] B. Kriheli, E. Levner, and A. Spivak, "Optimal search for hidden targets by unmanned aerial vehicles under imperfect inspections," *American Journal of Operations Research*, vol. 6, no. 02, p. 153, 2016.
[17] P. Elmore, W. E. Avera, M. M. Harris, K. M. Duvieilh *et al.*, "Environmental measurements derived from tactical mine-hunting sonar data," in *OCEANS 2007-Europe*. IEEE, 2007, pp. 1–5.
[18] A. Zare and J. T. Cobb, "Sand ripple characterization using an extended synthetic aperture sonar model and MCMC sampling methods," in *IEEE/MTS OCEANS*, San Diego, CA, USA, 2013, pp. 1–7.
[19] K. Takahashi, J. Igel, and H. Preetz, "Clutter modeling for ground-penetrating radar measurements in heterogeneous soils," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 4, pp. 739–747, 2011.
[20] K. Takahashi, H. Preetz, and J. Igel, "Soil properties and performance of landmine detection by metal detector and ground-penetrating radarsoil characterisation and its verification by a field test," *Journal of Applied Geophysics*, vol. 73, no. 4, pp. 368–377, 2011.
[21] P. D. Gader, M. Mystkowski, and Y. Zhao, "Landmine detection with ground penetrating radar using hidden markov models," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 39, no. 6, pp. 1231–1244, 2001.
[22] H. Yetkin, C. Lutz, and D. J. Stilwell, "Utility-based adaptive path planning for subsea search," in *Proc. IEEE/MTS OCEANS*, Washington, DC, USA, 2015.
[23] M. Harris, W. Avera, C. Steed, J. Sample, L. D. Bibee, D. Morgerson, J. Hammack, and M. Null, "Aqs-20 through-the-sensor (tts) performance assessment," in *OCEANS, 2005. Proceedings of MTS/IEEE*. IEEE, 2005, pp. 460–465.
[24] V. Myers and D. P. Williams, "A pomdp for multi-view target classification with an autonomous underwater vehicle," in *OCEANS 2010*. IEEE, 2010, pp. 1–5.
[25] Y. Zhang, A. B. Baggeroer, and J. G. Bellingham, "Spectral-feature classification of oceanographic processes using an autonomous underwater vehicle," *Oceanic Engineering, IEEE Journal of*, vol. 26, no. 4, pp. 726–741, 2001.
[26] S. Jaramillo and G. Pawlak, "Auv-based bed roughness mapping over a tropical reef," *Coral Reefs*, vol. 30, no. 1, pp. 11–23, 2011.
[27] J. O. Berger, *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media, 2013.