



A decision-theoretic approach to acquire environmental information for improved subsea search performance

Harun Yetkin^{1,*}, Collin Lutz², Daniel J. Stilwell

Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24060, USA

ARTICLE INFO

MSC:
90B40
68T37
68T40
93C85

Keywords:

Search theory
Path planning
Environmental information
Subsea search
Autonomous underwater vehicles

ABSTRACT

This study addresses subsea search applications where an autonomous underwater vehicle (AUV) is tasked with finding the target density in a given search region within finite time. We assume that AUV is equipped with a side-scan sonar sensor that detects the targets at the sampled location. We consider that sensor performance is dependent on local environmental conditions (e.g., clutter density, sediment type) that vary throughout the search region, and we presume that environmental conditions are unknown or partially known. Due to uncertain and varying environmental conditions, resulting search performance is also uncertain and it varies by location. This paper specifically considers the cases where environmental information can be acquired either by a separate vehicle or by the same vehicle that performs the search task. Our main contribution is to formally derive a decision-theoretic cost function to compute the locations where the environmental information should be acquired so that the performance of the search task can be improved. For the cases where computing the optimal locations to sample the environment is computationally expensive, we offer an approximation approach that yields provable near-optimal paths. We show that our decision-theoretic cost function outperforms the information-maximization approach, which is often employed in similar applications.

1. Introduction

We address subsea search applications where an AUV is to find an unknown number of objects in a bounded search environment and within limited time. We assume that the environment affects sensor performance, and local environmental conditions vary by location. That is, in some locations the search sensor is expected to perform better than in other locations. Yetkin et al. (2015) shows how guidance algorithms for search missions can incorporate stochastic knowledge of the environment to improve search performance. In this study, we specifically consider the cases where environmental information can be acquired as part of the overall search task. The principal contributions of this study show where environmental information should be acquired in order to improve overall search performance. We address two cases: (1) environmental characterization is performed prior to a search mission by a separate asset than the search vehicle, and (2) environmental characterization is performed at the same time as search by the same vehicle that performs the search task. In this study, we extend our findings in our prior work (Yetkin et al., 2015, 2016).

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 reduction in entropy of the distribution that describe the environment, which is often employed in similar applications (see, for example, Coleman and Block, 2007; Elfes, 1992; Papadimitriou et al., 2000). In contrast, we show that reducing uncertainty in the environment is not the best approach. 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

* Corresponding author.

E-mail addresses: hyetkin@bartin.edu.tr (H. Yetkin), collin@vt.edu (C. Lutz), stilwell@vt.edu (D.J. Stilwell).

¹ Presently with the Department of Mechatronics Engineering, Bartın University, Turkey

² Presently with Charles Schwab, Lone Tree, CO 80124, USA

the estimate for the number of objects in the environment is correct. Computing the optimal locations to acquire the environment information can be computationally expensive when the planning horizon is large. To address the computational challenge of our approach, we show an approximation approach that yields provably near-optimal paths.

The remainder of this paper is organized as follows. An overview of search theory and the benefit of acquiring environmental information in some search missions is provided in Section 2. In Section 3, we formulate the search problem and define the observation model. In Section 4, we define the objective function that maximizes the estimation accuracy. In Section 5 and Section 6, we describe our proposed cost function to compute the locations where environmental characterization should be performed. Section 7 provides the numerical results that illustrate our approach.

2. Related work

2.1. Search theory

Search theory is concerned with finding an optimal allocation of available search effort to locate a lost or hidden target, such that a reward specified as a measure of search effectiveness is maximized. Bernard Koopman offered one of the first systematic approaches to the search problem (Koopman, 1957). His work laid the foundations of the problem and since then the search problem received a lot of attention, mainly from the operations research community. Some notable examples include Chung and Burdick (2012a), Kadane (1971), Kimeldorf and Smith (1979), Kress et al. (2008), Richardson (1988) and Stone et al. (1972).

In a real-world search application, sensor measurements can be noisy due to false negatives, i.e. failing to detect a target that is present, and false positives, i.e. falsely detecting a target when the target is not present at the location. Local environmental conditions can also affect the number of false positives and false negatives the sensor observes, and deterministic knowledge on environmental conditions is often unavailable. However, existing literature in search theory rarely addresses the requirements of real-world search applications. For instance, the existence of false positives is mostly ignored in existing works. Exceptions include Chung and Burdick (2012b), Dobbie (1973), Kress et al. (2008), Kriheli et al. (2016), Pollock (1964) and Stone et al. (1972), where the effect of environmental conditions on sensor performance is not accounted for. To best of our knowledge, there is no work in the literature that addresses both the existence of false positives and multiple targets. Our work builds upon prior work by accounting for multiple targets, false alarms, non-zero cost of moving to another search location, and uncertainty in the environment. To best of our knowledge, this is the first study to account for all these factors together. We do not address coverage problems as in Choset (2001) or exhaustive search. Rather, we consider applications where there is a time or distance constraint, and we seek solutions where we achieve the best search performance within a time or distance constraint. 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.

2.2. Effect of environment information

We consider that the local environmental conditions affect the sensor performance, and that the environment varies throughout the search region. Hence, the sensor performs better in some locations than in others. The environment at a location may not be known with certainty and we may only have a probabilistic knowledge on the environment. When possible, acquiring information about the local environmental conditions can significantly improve the results of a

follow-on search mission. The question is, where to sample the environment so that the most improvement in search performance can be obtained. This is the fundamental question we seek to answer in this study, and we show that a new approach is needed to address the cases where the limited planning horizon of the mission does not allow exhaustive search of the environment.

The effect of the environment on search performance is well-known. De Guenin is the first to account for the effect of environmental conditions on search performance (De Guenin, 1961). He illustrated the intuitive effect of environmental conditions on detection probability with a simple example of visual search where the detection probability depends upon the atmospheric conditions that may not be uniform and thus, for the same search effort, the detection probability will be higher in clear areas than in overcast areas. 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 (Elmore et al., 2007; Zare and Cobb, 2013). Similarly, in terrestrial applications using ground penetrating radar, background clutter and soil properties have significant effect on search results (Gader et al., 2001; Takahashi et al., 2011a,b). For the cases where the search environment is unknown or partially known, reducing the uncertainty in the environment can greatly improve search performance. For the particular trial in Harris et al. (2005), 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. In our prior work (Yetkin et al., 2015), we also show that uncertainty in the environment leads to uncertainty in search performance which can result in terminating the search mission too early, and thus, obtaining poor search performance.

We note that the existing literature on robotic exploration (see Bourgault et al., 2002; Carrillo et al., 2015; Makarenko et al., 2002; Yamauchi, 1997, among many examples) provides little insight to the applications we address. Robotic exploration addresses the challenge of building a map. In contrast, we seek to characterize a subset of the environment with respect to search sensor performance for the goal of improving search effectiveness.

3. 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*. In this section, we provide the notation and the formal definition of the search problem. A list of variables is provided in Table C.1 in Appendix.

3.1. Preliminaries

A search grid $\mathcal{G} \subset \mathbb{R}^2$ is partitioned into K disjoint cells and each cell in the search grid is associated with random variables X and E . Random variable X represents the number of objects, and random variable E represents the environmental conditions in the cell. We consider that the number of objects and the environmental conditions in each is independent throughout the search grid. That is, for cells i and j , X_i is independent of X_j and E_i is independent of E_j . The objective of the search mission is to estimate X_1, \dots, X_K by using a sensor to detect objects in each cell. We assume that deterministic knowledge of the environmental conditions is unavailable, but the finite set of environments w_1, w_2, \dots, w_m in the search grid is known and a probabilistic knowledge on these environments is available for each cell. For instance, the environment probability distribution for

cell i is expressed $[p_1(i), p_2(i), \dots, p_m(i)]$ where $p_j(i) = P(E_i = w_j)$ is the probability that the actual environment in cell i is w_j . We note that in practice clustering of different types of environmental conditions can be carried out by using a previously acquired environment dataset in the search domain (see, for example, [McMahon et al., 2017](#)).

3.2. Sequential Bayesian update for the search vehicle

When the search vehicle samples from a location, it observes the number of objects in that location, and an imperfect measurement $z \in Z$ of the actual number of objects in the location is acquired. We denote by Z both the set of possible search measurements (i.e. $z \in Z$) and the random variable associated with a search measurement in a cell (i.e. $Z_i = z_i$). We note that the acquired measurement z includes both false positives and false negatives. The sensor model that we choose for the numerical illustrations in Section 7 is

$$P(z | x, w_j) = \sum_{k=0}^{\min(x,z)} \binom{x}{k} D_j^k (1 - D_j)^{x-k} (1 - F_j) F_j^{z-k} \quad (1)$$

which is the likelihood of observing z objects when x is the true number of objects at the location and the environment is w_j . In (1), $0 \leq F_j < 1$ denotes the probability of one or more false alarms, and $0 < D_j \leq 1$ denotes the probability of detection. Note that both F_j and D_j are assumed to vary as functions of the environment type w_j . The details of the search sensor model is described in prior work ([Shende et al., 2012](#); [Yetkin et al., 2015](#)) and omitted here for brevity. We note that other expressions for the sensor model are also possible, and our results do not depend on this specific sensor model except for numerical illustrations. Upon acquiring a search measurement z , we use Bayesian update law to update the distribution $P(X | z, w_j)$.

$$P(X = x | z, w_j) = \frac{P(z | x, w_j) P(X = x)}{\sum_x P(z | X = x, w_j) P(X = x)} \quad (2)$$

3.3. Sequential Bayesian update for the environmental characterization vehicle

When the environment characterization vehicle samples from a location, it observes the environmental conditions in that location, and an imperfect measurement $y \in Y$ of the true environment at the location is acquired. We denote by Y both the set of possible environment measurements (i.e. $y \in Y$) and the random variable associated with an environment measurement at a cell (i.e. $Y_i = y_i$). We assume the likelihood of observing an environment measurement y when the true environment is w_j is known. We use Bayesian update law to compute the posterior distribution of the environment after observing y ,

$$P(E = w_j | y \in Y) = \frac{P(y | E = w_j) P(E = w_j)}{\sum_j P(y | E = w_j) P(E = w_j)} \quad (3)$$

3.4. 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 probability distribution of the number of objects unconditioned on the environment

$$P(X = x | z, y) = \sum_j P(X = x | z, w_j) P(E = w_j | y) \quad (4)$$

where the posterior distributions $P(X = x | z, w_j)$ and $P(E = w_j | y \in Y)$ follow from (2) and (3), respectively.

4. 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 the 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 ([Yetkin et al., 2015](#)).

The search objective is to accurately estimate the number of objects in each cell. However, due to imperfect sensor measurements, our estimate of the number of objects in a cell can be different than the true number of objects. Hence, we define that the goal of searching a location is to maximize the accuracy of our estimate of the number of objects at that location. That is, we seek to maximize the probability that our estimate of the number of objects after searching a location will be correct. We use a zero-one utility function to compute the value of forming an estimate $\delta_X(z)$ of the true number of objects x after acquiring search measurement z .

$$U(x, \delta_X(z)) = \begin{cases} 1 & \text{if } x = \delta_X(z) \\ 0 & \text{if } x \neq \delta_X(z) \end{cases} \quad (5)$$

We believe the zero-one utility function in (5) adequately assesses the value of forming an estimate for the applications where an incorrect estimate may have severe consequences (e.g. subsea mine-hunting missions). For the applications where incorrect estimates are also valued, one can use a linear loss function to assess the value to forming an estimate of the number of objects (see, for example, [McMahon et al., 2017](#)). Given that the acquired measurement is z and the environment is w_j , Bayes estimator is

$$\delta_X^*(z) = \arg \max_{\delta_X(z)} \mathbb{E} [U(x, \delta_X(z)) | z, w_j] \quad (6)$$

that maximizes the expected values of the zero-one utility function in (5). The value of searching a location is

$$\mathbb{E} [U(x, \delta_X^*(z))] = \sum_{j=1}^m P(E = w_j) \mathbb{E} [U(x, \delta_X^*(z)) | w_j] \quad (7)$$

and we call this the *anticipated* estimation accuracy.

Let $\gamma = \{q_1, \dots, q_N\}$ be a candidate search path that traverses the cells $q_1, \dots, q_N \in \mathcal{G}$, and let \tilde{C}_γ be the budget constraint on the search mission due to limited time/distance the vehicle can traverse. We define $C(\gamma)$ to denote the cost for traversing a path γ . Note that when the traversal cost for moving from one location to another is unity, $C(\gamma)$ is simply the number of cells traversed by γ .

When the vehicle makes multiple visits to a cell, we acquire a set of independent search measurements. We denote by z both a single search measurement and a set of search measurements when a cell is visited multiple times by γ . Then, the expected utility of traversing γ is

$$\mathbb{E} [U(x, \delta_X^*(z_\gamma))] = \prod_{q_i \in \gamma} \mathbb{E} [U(x_{q_i}, \delta_X^*(z_{q_i}))] \times \prod_{i \in \mathcal{G} \setminus \gamma} \max_{x_i} P(X_i = x_i) \quad (8)$$

where $\mathcal{G} \setminus \gamma$ denotes the remaining cells in the search grid that are not traversed by γ , and $\max_{x_i} P(X_i = x_i)$ is the certainty in the number of objects in cell i prior to acquiring new measurements. Let Ω_γ denote the finite collection of feasible search paths. Then, the optimal search path is

$$\gamma^* = \arg \max_{\gamma \in \Omega_\gamma} \mathbb{E} [U(x, \delta_X^*(z_\gamma))] \quad (9)$$

subject to

$$C(\gamma) \leq \tilde{C}_\gamma \quad (10)$$

5. 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 the case where environmental characterization is performed prior to search, and we assume that both environmental characterization and search cannot be performed exhaustively due to limited resources.

5.1. Entropy change maximization

When environment information can be acquired only in some locations due to limited resources, the question is to determine where to optimally sample the environment. One approach that is often used in similar applications is to maximize the change in entropy due to acquired environment measurements (Coleman and Block, 2007; Elfes, 1992; Papadimitriou et al., 2000). We briefly describe a typical entropy approach for selecting where to sample the environment so that we can compare it to our proposed approach.

Let $H(E = w_j)$ denote the prior entropy of the probability distribution $P(E = w_j)$ and let $H(E = w_j | y \in Y)$ be the posterior entropy after acquiring the environment measurement y . The expected amount of change in the entropy for a future environment measurement y can be computed by

$$J(E) = H(E = w_j) - \sum_y H(E = w_j | y \in Y)$$

Let η be a candidate path for the environment characterization vehicle, Ω_η be the finite collection of feasible paths, and \bar{C}_η be the budget constraint on the environment characterization mission. Then, the best path to characterize the environment based on the entropy change maximization method is

$$\eta^* = \arg \max_{\eta \in \Omega_\eta} J_\eta(E) \quad (11)$$

subject to

$$C(\eta) \leq \bar{C}_\eta$$

However, we note that the purpose of environmental characterization in this study is not to explore the environment, but to improve the performance of a follow-on search mission. We show in Section 7 via numerical studies that maximizing change in entropy does not maximize the performance of follow-on search missions.

5.2. Environmental loss function

In order to improve the performance of a follow-on search mission by exploring the environment, we first address how the uncertainty in the environment affects the search performance. When we compute the value of searching a location in (7), we average the search value over all possible environments with respect to the prior environment probability distribution, which can result in a different search value than the true search value given true environment at the location. That is, uncertainty in the environment results in deviations from actual search results and degrades search performance after a mission.

Let e be the true environment and $V(w_j)$ denote the value of searching a location with environment w_j

$$V(w_j) = \mathbb{E} \left[U(x, \delta_X^*(z)) | E = w_j \right] \quad (12)$$

In order to penalize the deviations from actual search performance, we define a linear loss function

$$L_V(e, \delta_E(y)) = \begin{cases} c_1 (V(e) - V(\delta_E(y))) & \text{if } \delta_E(y) \leq e \\ c_2 (V(\delta_E(y)) - V(e)) & \text{if } \delta_E(y) > e \end{cases} \quad (13)$$

where $c_1, c_2 > 0$ are the relative costs of over and underestimating search performance, and $\delta_E(y)$ is an estimate of the environment after environment measurement y is acquired. The notation $<$ and $>$ represent a preference ordering among the environments. When $V(w_i) < V(w_j)$ for environments w_i and w_j , we say $w_i < w_j$, and when $V(w_i) > V(w_j)$, we say $w_i > w_j$. Throughout this paper, we consider a *distinct* and *ordered* set of environments so that $w_1 < w_2 < \dots < w_m$.

When our environment estimate is worse than the true environment ($\delta_E(y) < e$), too much search effort can be allocated to the same location. Since the available search effort we can apply to search the region is limited, allocating unnecessary search effort to the same location may degrade the overall search performance. On the other hand, when our environment estimate is greater than the true environment ($\delta_E(y) > e$), we exaggerate the value of searching a location, which may result in poorly allocated search effort and exaggerated search performance after a mission. In the numerical illustrations in Section 7, we choose $c_1 < c_2$ so that overestimation is more penalized than underestimation.

After observing environment measurement y at a location, we compute the posterior probability distribution of the environments $P(w_j | y)$ in (3). We then compute the expected loss of forming environment estimate $\delta_E(y)$ with respect to the posterior distribution $P(w_j | y)$

$$\mathbb{E} \left[L_V(e, \delta_E(y)) | Y = y \right] = \sum_{j=1}^m P(w_j | y) L_V(w_j, \delta_E(y)) \quad (14)$$

and, we choose the Bayes estimator $\delta_E^*(y)$ that minimizes the expected loss in (14).

$$\delta_E^*(y) = \arg \min_{\delta_E(y)} \mathbb{E} \left[L_V(e, \delta_E(y)) | y \right] \quad (15)$$

Let $F_{E|y}(w_n)$ be such that

$$\begin{aligned} F_{E|y}(w_n) &:= P(E \leq w_n | y) \\ &= \sum_{j=1}^n P(E = w_j | y) \end{aligned} \quad (16)$$

for $n \in \{1, 2, \dots, m\}$ where (16) follows since $w_1 < w_2 < \dots < w_m$. Then, the Bayes estimator (15) with respect to the linear loss function in (13) is

$$\delta_E^*(y) = w_{l+1} \quad (17)$$

where

$$l = \arg \max_n \left(n \in \{1, 2, \dots, m-1\} : F_{E|y}(w_n) \leq \frac{c_1}{c_1 + c_2} \right) \quad (18)$$

We note that the proof for (17) can be found in any standard book on statistical decisions (see, for example, Berger, 2013).

5.3. Path planning

A benefit of environmental surveys is to reduce the error in anticipated search performance due to uncertainty in the environment. When a location is not visited by the search vehicle during a search mission, acquiring an environment measurement at that location will not affect search performance. From the loss function in (13), computing the estimate (15) of the environment after acquiring environment measurement y yields the conditional expected loss

$$\mathbb{E} \left[L_V(e, \delta_E^*(y)) | y \right] = \sum_{j=1}^m P(E = w_j | y) L_V(w_j, \delta_E^*(y)) \quad (19)$$

that quantifies the amount of uncertainty in anticipated estimation accuracy after acquiring y . Informally speaking, the prior loss before acquiring an environment measurement represents the prior uncertainty, and the conditional expected loss in (19) represents the posterior uncertainty in search performance. For notational convenience, we define $\mathcal{R}(y)$ to denote the reduction of uncertainty in anticipated estimation accuracy due to environment measurement y

$$\mathcal{R}(y) = \mathbb{E} \left[L_V(e, \delta_E^*) \right] - \mathbb{E} \left[L_V(e, \delta_E^*(y)) | y \in Y \right]$$

Then, the gain of acquiring an environment measurement y_i in cell i is the reduction of uncertainty in anticipated estimation accuracy given that cell i is visited by the search vehicle

$$G(y_i) = \mathbf{I}(i, \gamma^*(y_i)) \mathcal{R}(y_i) \quad (20)$$

where the notation $\gamma^*(y)$ denote the best path for the search vehicle when the probability distribution of the environment is updated with the acquired measurement y , and the indicator function $\mathbf{I}(i, \gamma(y_i)) : y_i \rightarrow [c', 1]$ is defined

$$\mathbf{I}(i, \gamma(y_i)) = \begin{cases} 1 & i \in \gamma(y_i) \\ c' & i \notin \gamma(y_i) \end{cases} \quad (21)$$

where $0 \leq c' < 1$ is a parameter to determine the relative gain of sampling the environment at locations that will not be searched. Without loss of generality, we consider that $c' = 0$.

Let $\eta = \{q_1, q_2, \dots, q_M\}$ be a candidate path for the environment characterization vehicle and recall that Ω_η is the finite collection of feasible characterization paths. We denote by y both a single environment measurement and a set of independent environment measurements when a cell is visited multiple times by η . Then, the expected characterization gain of traversing η is

$$\mathbb{E}[G_\eta] = \sum_{y_\eta} P(Y_\eta = y_\eta) \sum_{q_i \in \eta} \mathbf{I}(i, \gamma^*(y_\eta)) \mathcal{R}(y_{q_i}) \quad (22)$$

and the optimal path is

$$\eta^* = \arg \max_{\eta \in \Omega_\eta} \mathbb{E}[G_\eta] \quad (23)$$

subject to

$$C(\eta) \leq \bar{C}_\eta \quad (24)$$

5.4. Approximating the characterization gain of a path

Computing the optimal path for the environment characterization vehicle in (23) can be computationally very expensive. This is mainly due to large computational requirements of computing the optimal search path for each set of environment measurements along a candidate path η . Thus, we also propose an approximate method that reduces the computational complexity of the solution in (23). Our approximate solution yields provably near-optimal paths.

Let $\bar{\Omega}_\eta$ be the set of environment characterization paths such that $C(\eta) \leq \bar{C}_\eta$, let $|\cdot|$ denote the size of a set (or an array), and let \mathcal{S} denote the computational complexity of computing the optimal search paths. Then, the solution in (23) has a computational complexity of $\mathcal{O}(|\bar{\Omega}_\eta| m^{|\mathcal{S}|})$ where m is the number of environments in the search domain. This shows that the exponential increase in the computational complexity is dominated by the large planning horizon for the characterization vehicle. One approach to reduce this computational complexity is to use a receding horizon strategy where we compute the paths for a shorter horizon. While receding horizon approach may require less computational power compared to computing the paths for the entire planning horizon, it may still be infeasible unless the considered planning horizon is sufficiently small (in which case the resulting performance will be poor). Instead, our approach to reduce the complexity of the solution in (23) aims to approximate the characterization gain of traversing a path.

We start with re-arranging the terms in (22) by partitioning a path η into two parts; a cell $q_i \in \eta$ and the other cells in the path

$$\begin{aligned} \mathbb{E}[G_\eta] &= \sum_{y_\eta} P(Y_\eta = y_\eta) \sum_{q_i \in \eta} \mathbf{I}(q_i, \gamma^*(y_\eta)) \mathcal{R}(y_{q_i}) \\ &= \sum_{q_i \in \eta} \sum_{y_{q_i}} \mathcal{R}(y_{q_i}) \sum_{y_{\eta \setminus q_i}} \mathbf{I}(q_i, \gamma^*(y_\eta)) P(Y_\eta = y_\eta) \end{aligned} \quad (25)$$

where $\eta \setminus q_i$ denotes the set of cells in η except cell q_i . The terms up to the third summation in (25) denote the characterization gain of

acquiring the environment measurement y from cell $q_i \in \eta$, and the other terms that start with the third summation denote how likely it is that sampling cell q_i will improve the performance of a follow-on search mission. That is, it represents the chances that cell q_i will be visited during a follow-on search mission based on the environment measurements that we may acquire along the path η . Note that since $\eta = \{q_1, q_2, \dots, q_M\}$, the joint probability $P(Y_\eta = y_\eta)$ in (25) can be expressed

$$P(Y_\eta = y_\eta) = P(Y_{q_1} = y_{q_1}) \times \dots \times P(Y_{q_M} = y_{q_M}) \quad (26)$$

Thus, we can re-write (25)

$$\mathbb{E}[G_\eta] = \sum_{q_i \in \eta} \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \bar{\mathbf{P}}_{\eta \setminus q_i} \quad (27)$$

where $\bar{\mathbf{P}}_{\eta \setminus q_i}$ is the total probability of every possible sets of the environment measurements $y_{q_1}, y_{q_2}, \dots, y_{q_{i-1}}, y_{q_{i+1}}, \dots, y_{q_M}$ such that the optimal search path associated with the updated environment distributions visits cell q_i

$$\bar{\mathbf{P}}_{\eta \setminus q_i} = \sum_{y_{\eta \setminus q_i}} \mathbf{I}(q_i, \gamma^*(y_\eta)) P(Y_{\eta \setminus q_i} = y_{\eta \setminus q_i}) \quad (28)$$

$$= \sum_{y_{\eta \setminus q_i} : q_i \in \gamma^*(y_\eta)} P(Y_{\eta \setminus q_i} = y_{\eta \setminus q_i}) \quad (29)$$

Indeed, the computational cost of (27) is dominated by $\bar{\mathbf{P}}_{\eta \setminus q_i}$ in (29). Thus, we use a sample-based method to compute an empirical estimate of $\bar{\mathbf{P}}_{\eta \setminus q_i}$, which results in a significant speed-up in computing the characterization path. For each cell $q_i \in \eta$ and for every environment measurement $y_{q_i} \in Y$, we perform \bar{N} trials where, in each trial, we randomly sample an environment measurement from the probability distribution $P(Y_{q_i} = y_{q_i})$ for all $q_j \in \eta$ such that $j \neq i$. Then, with the updated environment distributions we compute the optimal search path and see whether the corresponding path visits cell q_i , or not. We simply count the number of times cell q_i is being visited by the resulting search path out of \bar{N} trials, and we denote this number by k . Since this is repeated for every other cell in the path, we may sample the same environment measurement y_{q_i} from cell q_i during a trial of another cell. Let $\bar{N}_{y_{q_i}}$ be the number of times y_{q_i} is sampled in cell q_i during the trials of the other cells in path η and let $k_{y_{q_i}}$ be the number of times cell q_i is contained in the corresponding search path out of these $\bar{N}_{y_{q_i}}$ trials. Then, the empirical estimate of $\bar{\mathbf{P}}_{\eta \setminus q_i}$ is

$$\hat{\mathbf{P}}_{\eta \setminus q_i} = \frac{k + k_{y_{q_i}}}{\bar{N} + \bar{N}_{y_{q_i}}} \quad (30)$$

Obtaining a close estimate of $\bar{\mathbf{P}}_{\eta \setminus q_i}$ is important to compute a near-optimal characterization path. We show that a bound on the distance between $\bar{\mathbf{P}}_{\eta \setminus q_i}$ and $\hat{\mathbf{P}}_{\eta \setminus q_i}$ can be computed. We first note that after each trial for a cell, that cell is either contained in the follow-on search path, or it is not contained. Thus, we can cast each trial as a Bernoulli trial where the result of the trial is either 1 if the cell is contained in the search path, or it is 0 if the cell is not contained. Then, we use Hoeffding's inequality Hoeffding (1963) to obtain a probabilistic bound on the difference between $\bar{\mathbf{P}}_{\eta \setminus q_i}$ and $\hat{\mathbf{P}}_{\eta \setminus q_i}$

$$P(|\bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i}| < \epsilon \bar{N}) \geq 1 - 2 \exp^{-2\epsilon^2 \bar{N}} \quad (31)$$

where $\bar{N} = \bar{N} + \bar{N}_{y_{q_i}}$ and $\epsilon > 0$. Replacing $\bar{\mathbf{P}}_{\eta \setminus q_i}$ with its estimate $\hat{\mathbf{P}}_{\eta \setminus q_i}$ in (27) approximates the characterization gain of a path. We denote the approximate characterization gain of a path η by $\mathbb{E}[\hat{G}_\eta]$

$$\mathbb{E}[\hat{G}_\eta] = \sum_{q_i \in \eta} \sum_{y_{q_i}} \frac{k + k_{y_{q_i}}}{\bar{N} + \bar{N}_{y_{q_i}}} P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \quad (32)$$

Finally, we select the path that maximizes the approximate characterization gain in (32) subject to the budget constraint in (24).

$$\hat{\eta} = \arg \max_{\eta \in \bar{\Omega}_\eta} \mathbb{E}[\hat{G}_\eta] \quad (33)$$

Now, we define the following theorem and the corollary, where we first bound the difference between the characterization gain and the approximate characterization gain for a path, and we then bound the difference between the optimal characterization gain and the approximately optimal characterization gain. Proofs for both the theorem and the corollary are provided in [Appendix](#).

Theorem 1. *The probability of the difference between the characterization gain (27) and its estimate (32) for a path η satisfying a specific bound is expressed*

$$P\left(\left|\mathbb{E}[G_\eta] - \mathbb{E}[\hat{G}_\eta]\right| < \epsilon \bar{N}|\eta|\right) \geq 1 - 2 \exp^{-2\epsilon^2 \bar{N}} \quad (34)$$

Corollary 1. *For the optimal characterization path η^* in (23) and the approximate path $\hat{\eta}$ in (33), the difference in expected characterization gain satisfies*

$$P\left(\mathbb{E}[G_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] < 2\epsilon \bar{N}|\eta|\right) \geq 1 - 2 \exp^{-2\epsilon^2 \bar{N}} \quad (35)$$

The proposed approximation approach yields a provably near-optimal path for the environment characterization vehicle. The computational complexity of the solution is reduced from $\mathcal{O}(|\bar{\mathcal{Q}}_\eta| m^{|\eta|} \mathcal{S})$ to $\mathcal{O}(|\bar{\mathcal{Q}}_\eta| |\eta| m \bar{N} \mathcal{S})$. In general, choosing a larger value for \bar{N} is likely to reduce the approximation error. However, our preliminary results show that a small value for \bar{N} is often sufficient to obtain a close approximation.

5.5. Approximating the characterization gain of a cell

In this section, we introduce an alternative approach to approximate the optimal characterization gain. Our alternative approach assumes that the search paths are composed of sequences of parallel straight lines. This assumption arises often in subsea applications that rely on side-scan imaging sonar (see, for example, [Hayes and Gough, 2009](#); [Houston et al., 2002](#)). Unlike the proposed approach in Section 5.4, our approach in this section can only apply to certain classes of mapping problems.

Suppose the search area \mathcal{G} consists of a set of parallel straight lines (a line is either a row or a column) l_1, l_2, \dots, l_{n_l} . When each line corresponds to a row, $n_l = n_r$, and when each line corresponds to a column, $n_l = n_c$. Suppose that the j th line traverses the cells $q_{j1}, q_{j2}, \dots, q_{jk}$. Thus, when the vehicle traverses line l_j , it samples from cells $q_{j1}, q_{j2}, \dots, q_{jk}$. Let η be a candidate path for the environment characterization vehicle that consists of lines l_1, l_2, \dots, l_k , and consider that cell $i \in \eta$ is contained in l_j . We claim that the value of characterizing cell i along path η can be closely approximated by the value of characterizing cell i along line l_j . That is

$$\mathbb{E}[G_{q_i \in \eta}] = \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \bar{\mathbf{P}}_{\eta \setminus q_i} \quad (36)$$

$$\approx \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \bar{\mathbf{P}}_{l_j \setminus q_i} \quad (37)$$

where

$$\bar{\mathbf{P}}_{l_j \setminus q_i} = \sum_{y_{l_j \setminus q_i}: q_i \in \gamma^*(y_{l_j \setminus q_i})} P(Y_{l_j \setminus q_i} = y_{l_j \setminus q_i}) \quad (38)$$

Due to (37), we can compute the value of characterizing a particular cell by only looking at the cells in the associated line (a row or a column). By doing so, we can compute the characterization gain for each cell individually. Then, the characterization gain of a path is simply the summation of the characterization gain of each cell in that path. We note that we do not have a formal guarantee of how closely (37) approximates (36). However, our initial tests as well as the intuition suggest that (36) can be well-approximated by (37).

Since computing $\bar{\mathbf{P}}_{l_j \setminus q_i}$ in (38) can be very expensive when the length of a line is large, we instead compute an empirical estimate

of $\bar{\mathbf{P}}_{l_j \setminus q_i}$ as described in Section 5.4. For each cell in the search grid, we perform \bar{N} trials to compute an empirical estimate of $\bar{\mathbf{P}}_{l_j \setminus q_i}$, where in each trial we sample environment measurements from the remaining cells in line l_j . After approximating the characterization gain for each cell individually, we apply an exact branch-and-bound method similar to our prior work in ([McMahon et al., 2017](#)) to compute the near-optimal path for the characterization vehicle. Suppose that the complexity of computing the characterization path when cell-wise characterization gains are known is similar to that of the search path. Then, this approximation approach yields a $\mathcal{O}(rm\bar{N} + 1)\mathcal{S}$ complexity of computing the near-optimal characterization paths.

6. 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 4 where the search vehicle aims to maximize the estimation accuracy with only the search measurements, we now acquire both a search measurement z and an environmental measurement y when the vehicle visits a location. Thus, the path strategy in Section 4 that do not address the acquisition of environmental measurements do not apply to this case. The results of this section partially follow from prior work ([Yetkin et al., 2016](#)) and presented here for completeness.

Let $\{w_1, w_2, \dots, w_m\}$ be a set of environments, and let $W(w_j)$ denote estimation accuracy conditioned on the environment w_j after acquiring search measurement z .

$$W(w_j) = \max_x P(X = x | z, w_j) \quad (39)$$

Note $W(w_j)$ in (39) is the accuracy of the estimate of the number of objects at a location while $V(w_j)$ in (12) is the expected accuracy when a measurement z has not yet been acquired. Suppose the environments w_1, \dots, w_m are distinct and ordered with respect to the estimation accuracy $W(w_j)$. That is, $W(w_1) < W(w_2) < \dots < W(w_m)$ implies $w_1 < w_2 < \dots < w_m$. Let e be the true environment in a cell, and let $\delta_E(y)$ be an estimate of the environment based on the environment measurement y . When the true environment is e , the loss due to the estimate $\delta_E(y)$ is defined

$$L_W(e, \delta_E(y)) = \begin{cases} c_1(W(\delta_E(y)) - W(e)) & \text{if } \delta_E(y) > e \\ c_2(W(e) - W(\delta_E(y))) & \text{if } \delta_E(y) \leq e \end{cases} \quad (40)$$

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 $\delta_E(y)$, and the corresponding Bayes estimator $\delta_E^*(y)$ are

$$\mathbb{E}[L_W(e, \delta_E(y)) | z, y] = \sum_{j=1}^m P(w_j | y) L_W(w_j, \delta_E(y)) \quad (41)$$

$$\delta_E^*(y) = \arg \min_{\delta_E(y)} \mathbb{E}[L_W(e, \delta_E(y)) | z, y] \quad (42)$$

Given search measurement z and the estimate $\delta_E^*(y)$ after acquiring environment measurement y , the probability that the estimate of the number of objects at a location is correct is computed from

$$\mathbb{E}[U(x, \delta_X(z)) | z, \delta_E^*(y)] = \max_x P(X = x | z, \delta_E^*(y)) \quad (43)$$

In order to assess the benefit of visiting a location, we compute the *estimated* estimation accuracy in (43) for each possible set of observations $z \in \mathcal{Z}$, $y \in \mathcal{Y}$. Then, the expected estimation accuracy before visiting a location can be computed

$$\mathbb{E}[W(\delta_E(y))] = \sum_z \sum_y P(z, y) \max_x P(X = x | z, \delta_E^*(y)) \quad (44)$$

where

$$P(z, y) = \sum_x \sum_{w_j} P(z | x, w_j) P(y | w_j) P(x) P(w_j) \quad (45)$$

We again consider the candidate search path γ , the finite collection of feasible search paths Ω_γ , and the budget constraint \bar{C}_γ on the vehicle. Let y_{q_i} be the set of independent environment measurements acquired at q_i th cell. The expected estimation accuracy for traversing γ is

$$\mathbb{E}\left[W(\delta_E(y_\gamma))\right] = \prod_{q_i \in \gamma} \mathbb{E}\left[W(\delta_E(y_{q_i}))\right] \times \prod_{i \in \mathcal{S}_\gamma} \max_{x_i} P(x_i) \quad (46)$$

and the optimal path is

$$\gamma^* = \arg \max_{\gamma \in \Omega_\gamma} \mathbb{E}\left[W(\delta_E(y_\gamma))\right] \quad (47)$$

subject to

$$C(\gamma) \leq \bar{C}_\gamma \quad (48)$$

7. Numerical results

This section presents simulation results and elaborates on the performance of the proposed search and environmental characterization strategies. Our numerical illustrations are inspired by subsea mine-hunting missions. However, we note that our proposed approaches also apply to other subsea search applications. 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 operates on separate vehicles, our proposed approach maximizes the reduction of uncertainty in search performance (22). Thus, our approach should, on average, display less anticipated estimation accuracy error than other approaches.

The search region \mathcal{G} is partitioned into a grid with 10×10 non-intersecting cells. It is assumed that there is no prior information on the number of objects for any cell, but an upper bound $L > 0$ on the number of objects that can reside in a cell is known. That is, for cell i ,

$$x_i \leq L \quad \text{and} \quad P(x_i = 0) = P(x_i = 1) = \dots = P(x_i = L) \quad (49)$$

with known upper bound L . There are three candidate environments in the search region and the distinct and ordered set of environments is $w_1 < w_2 < w_3$. The likelihood of observing a particular search measurement follows from (1). 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 w_1 , $D = 0.8$ and $\alpha = 0.3$ for environment w_2 , and $D = 0.95$ and $\alpha = 0.05$ for environment w_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 w_1 is the least and environment w_3 is the most informative. We consider that the sensor model for environment characterization is

$$a_{ij} = P(Y = w_i | E = w_j) \quad \text{for all } i, j \in \{1, 2, 3\} \quad (50)$$

where a_{ii} is the probability of observing the true environment w_i . For the numerical illustrations, we use the characterization sensor model with $a_{11} = 0.9$, $a_{22} = 0.92$, $a_{33} = 0.94$. That is, for example, there is 0.9 probability of acquiring environment measurement w_1 when w_1 is the true environment at the location. The noisy environment observations are due to nonzero probabilities of observing environment w_i when true environment is w_j , denoted by a_{ij} for $i \neq j$. We assume the probability of acquiring incorrect environment measurement is the same for all possible environments other than the true environment. For example, when the true environment at a location is w_1 , since $a_{11} = 0.9$, the probability of acquiring environment measurement w_2 and probability of acquiring environment measurement w_3 are $a_{21} = a_{31} = 0.05$.

We consider that the search vehicle is equipped with a side-scan sonar as is typically the case in subsea search applications. The quality of the data side-scan sonar observes can be very poor when the vehicle turns. Thus, in order to account for this limitation of the sensor on

the vehicle's motion, we consider that the vehicle travels in parallel straight lines only, and it transits between different lines only outside of the search region. We implicitly assume that the cells are sufficiently large so that any sequence of cells that correspond to "go-straight", "turn-right", or "turn-left" corresponds to a connected and dynamically feasible path. Passing through cells that are outside the search region requires time but does not improve search performance since no measurements are acquired. For the numerical illustrations, we consider a unit cost for moving from a cell to an adjacent cell. Thus, the total cost of a mission is the length of the planning horizon, which we refer to as the mission length.

7.1. Numerical illustrations

Fig. 1 shows a search area that is partitioned into regions A1 through A5. For each region, the corresponding probability distribution $\Pi = [p_1, p_2, p_3]$ is given, where p_j is the probability that the environment is w_j . For example, for the cells labeled A2, there is 0.15 probability that the environment is w_1 , 0.2 probability that the environment is w_2 , and 0.65 probability that the environment is w_3 . 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. We consider that the mission length is 60 for the search and search/ environmental characterization vehicles and 35 for the characterization vehicle. We selected the search path length to be 60 steps and the environmental characterization length to be 35 steps because for this specific illustrative problem these values best illustrated the advantage of our rigorously derived approach over a more naive approach based on entropy maximization. For other values applied to this specific problem, the advantage is smaller. We adopt a best-first branch-and-bound approach to compute the optimal paths (Wah and Yu, 1985).

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 (9), and the objective of the characterization vehicle is to maximize the expected gain of characterization in (22). However, due to large computational requirements of computing (22), we instead approximate the solution of (22) as described in Section 5.5. When both sensors operate on the same vehicle, the objective of the vehicle is to maximize expected estimation accuracy in (46).

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 approximate approach in Section 5.5 with the entropy change maximization method described in Section 5.1. Fig. 2(e) shows the trajectory for the environmental characterization vehicle when using our proposed approximate approach in Section 5.5, and Fig. 2(f) shows the trajectory when using the entropy change approach in Section 5.1. When using the proposed approximate approach, we choose \bar{N} – the number of trials to approximate $\bar{\mathbf{P}}$ in (29) – to be 100. We note that choosing $\bar{N} = 100$ suffices to compute a good approximation of the term $\bar{\mathbf{P}}$ and it yields computationally feasible results. Neither environmental characterization path visits A1 because the environments in those locations are completely known. We note that the environmental characterization path in Fig. 2(e) that

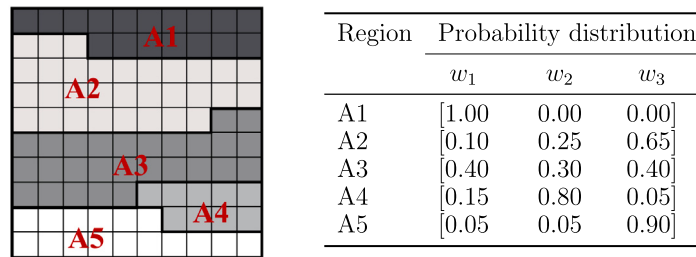


Fig. 1. Search area and cell-wise environment distributions.

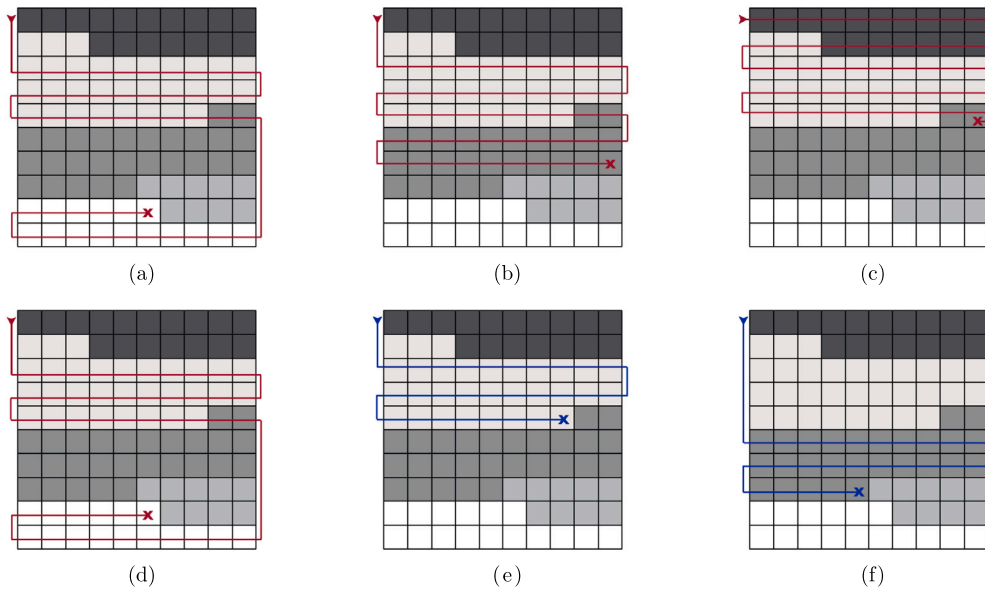


Fig. 2. Optimal trajectories for search and characterization. Figures (a–c): trajectories for the case both sensors operate on the same vehicle when (a) proposed approach is employed (b) entropy change maximization method is employed, and (c) the mowing-the-lawn approach is employed. Figures (e–f): characterization 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.

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 (47) with entropy change maximization method and with a mowing-the-lawn approach. The latter arises often in subsea applications such as mine-hunting. We note that the entropy change maximization method described in Section 5.1 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 (11) as

$$\gamma^* = \arg \max_{\gamma \in \Omega_\gamma} J_\gamma(X) + \beta J_\gamma(E) \quad (51)$$

where $J(X)$ denotes the entropy change in $P(X)$, the number of objects, and β is the relative weight of the entropy change in $P(E)$ compared to the entropy change in $P(X)$. Since the objective is to reduce the uncertainty in the number of objects, we choose $0 < \beta < 1$. Fig. 2(c) 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. 2(a) shows the trajectory for the proposed approach and Fig. 2(b) 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. 2(d). We note that, due to sonar sensor, generated paths in Fig. 2 are parallel straight lines only. However, our results apply to other domains where there is no such constraint on vehicle motion.

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 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 $P(z | x, e)$ and $P(y | e)$ 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 significantly reduced when our proposed approach is employed.

Both sensors operating simultaneously on a single vehicle

Fig. 3 shows the results after 10 000 iterations for the case both sensors operate on the same vehicle. Fig. 3(a) on the left is the percent

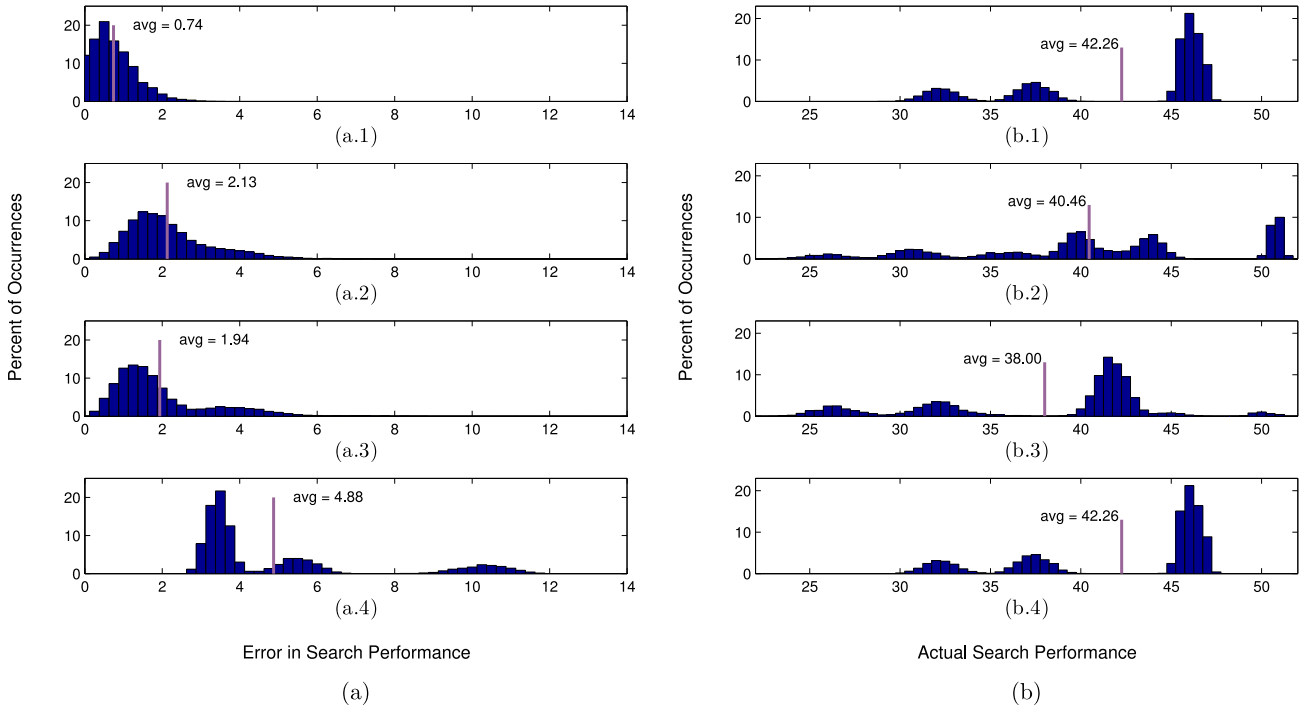


Fig. 3. Percent of occurrences for (a) error in search performance and (b) actual search performance when both sensors operate on *the same vehicle*. From top to bottom, (a.1) and (b.1) correspond to the proposed approach, (a.2) and (b.2) correspond to the entropy change maximization method, (a.3) and (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 search performance which is the achieved risk reduction after a mission. Smaller values for (a) imply less error in search performance and larger values for (b) imply better search performance.

of occurrences of the error in search performance, and Fig. 3(b) on the right is the percent of occurrences of the actual search performance. For convenience, we compute the actual search performance after traversing an optimal path γ and acquiring the search measurements z_γ as

$$-\left(\log\left(\prod_{i \in \mathcal{G}} \max_{x_i} P(x_i)\right) - \log\left(\prod_{i \in \gamma} \max_{x_i} P(x_i | z_i, e_i) \times \prod_{i \in \mathcal{G} \setminus \gamma} \max_{x_i} P(x_i)\right)\right) \quad (52)$$

where e_i is the actual environment in cell i . That is, the actual search performance is the difference between the prior certainty in the number of objects before acquiring any measurement and the posterior certainty in the number of objects after acquiring the search measurements along the path. Loosely speaking, the actual search performance plotted in Fig. 3(b) represents the amount of information we acquire on the number of objects after traversing the corresponding optimal search path. Thus, smaller values for Fig. 3(a) imply less error in search performance and larger values for Fig. 3(b) imply better search performance. The displayed results are the negative log of the computed 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. The simulations show that

- The proposed approach yields smaller error in search performance compared to the entropy change maximization and mowing-the-lawn. With respect to the case where there is no environment information (in Fig. 3(a).4), our proposed approach achieves 85% error reduction while entropy change maximization achieves 56% and mowing-the-lawn achieves 60%, on average. In addition, the actual search performance when using our approach is no worse than the actual search performance when using the other methods.

- Fig. 3(a).1 shows that in many of the iterations, the error in search performance is very close to zero. This implies that, in these trials, we correctly estimate the environmental conditions in each visited cell. We note that this is also due to the sensor model we choose in (50) for environment characterization.
- The 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 environment while the entropy change maximization method visits A3 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.
- Note that Fig. 3(b).1 and Fig. 3(b).4 are identical. This is because the search locations selected by the proposed approach (in Fig. 2(a)) are the same locations selected when environment information is not available (in Fig. 2(d)) for given search area characteristics. However, the anticipated search performances for these two cases are different. Indeed, comparing Fig. 3(a).1 with Fig. 3(a).4 shows that the anticipated search performance when environment information is available is significantly more accurate than the anticipated search performance when there is no environment information. Hence, a benefit of characterizing the environment is to better anticipate the true search performance.

Each sensor on separate vehicles

The results when search and environmental characterization tasks are performed on separate vehicles are shown 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

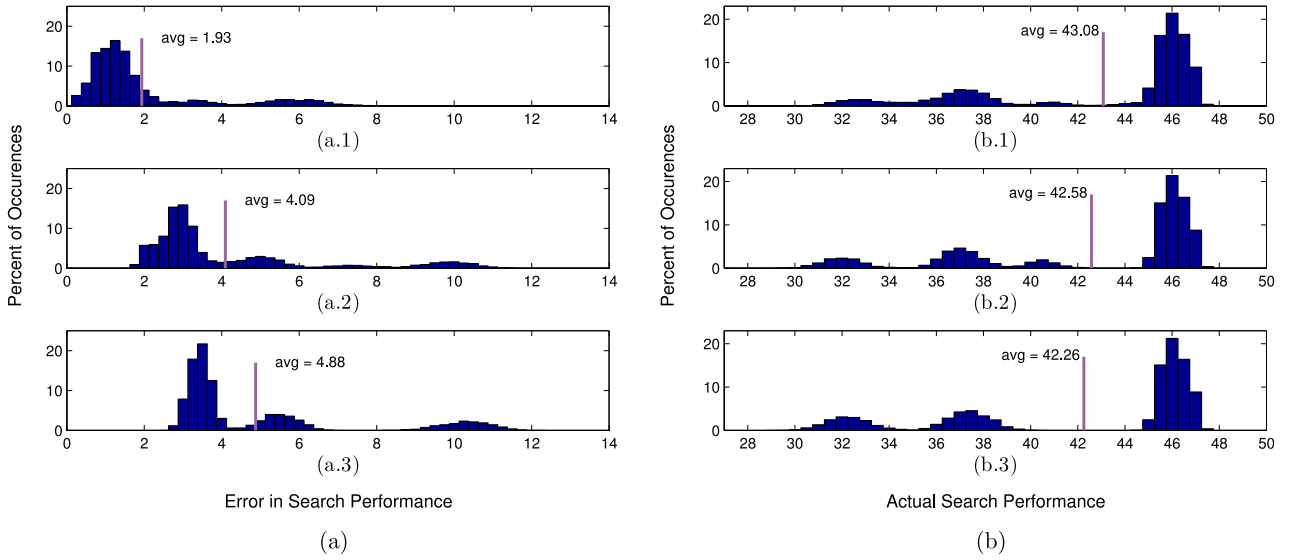


Fig. 4. Percent of occurrences for (a) error in search performance and (b) actual search performance when search and characterization are performed on *separate vehicles*. From top to bottom, (a.1) and (b.1) correspond to the proposed approach, (a.2) and (b.2) correspond to the entropy change maximization method, and (a.3) and (b.3) correspond to the case where environment information is not available. Note that the horizontal axis is the negative log of the search performance which is the achieved risk reduction after a mission. Smaller values for (a) imply less error in search performance and larger values for (b) imply better search performance.

characterization and the search vehicle plans its path by using the prior environmental distributions. We note that Fig. 4(a).3 and Fig. 4(b).3 are the same plots given in Fig. 3(a).4 and Fig. 3(b).4, and we show them here for convenience of comparison. It is seen that

- The average error is significantly smaller when environmental characterization is performed at the locations selected by our proposed approach. With respect to the case where there is no environment information (in Fig. 4(a).3), our proposed approach achieves 60% error reduction while entropy change maximization achieves only 16%, on average. Note that the distribution in Fig. 4(a).2 is very similar to the distribution in Fig. 4(a).3. This is because entropy change maximization fails to improve the performance of a follow-on search mission since it leads the vehicle to explore the parts of the search area that are less likely to be searched.
- 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.

8. 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 AUV or on separate AUVs. 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 AUVs, 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 AUV, 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.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Harun Yetkin: Conceptualization, Methodology, Software, Formal analysis, Visualization, Writing - original draft. **Collin Lutz:** Conceptualization, Methodology, Software, Formal analysis. **Daniel J. Stilwell:** Conceptualization, Methodology, Formal analysis, Supervision, Funding acquisition, Writing - review & editing.

Acknowledgments

The authors gratefully acknowledge the support of the Office of Naval Research, USA 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.

Appendix A. Proof of Theorem 1

First, observe that when $c_1, c_2 \leq 1$ in (13)

$$\begin{aligned} \mathcal{R}(y) &= \mathbb{E} \left[L_V(e, \delta_E^*) \right] - \mathbb{E} \left[L_V(e, \delta_E^*(y)) \mid y \in Y \right] \\ &\leq \mathbb{E} \left[L_V(e, \delta_E^*) \right] \\ &\leq \max_{w_i, w_j} (V(w_i) - V(w_j)) \\ &\leq 1 \end{aligned}$$

for all $y \in Y$.

The difference between $\mathbb{E}[G_\eta]$ and $\mathbb{E}[\hat{G}_\eta]$ is

$$\left| \mathbb{E}[G_\eta] - \mathbb{E}[\hat{G}_\eta] \right|$$

Table C.1
A list of variables.

Variable	Description
X	Random variable denoting the number of targets in a cell
E	Random variable denoting environmental conditions in a cell
Z	The set of search measurements and also random variable denoting a search measurement
Y	The set of environment measurements and also random variable denoting an environment measurement
x_i	True number of objects in cell i
e_i	True environment in cell i
z_i	Search measurement(s) acquired from cell i
y_i	Environment measurement(s) acquired from cell i
$\hat{\delta}_x(\cdot)$	An estimate of the number of objects
$\hat{\delta}_E(\cdot)$	An estimate of the environment
\mathcal{G}	Search grid
n_r, n_c	Number of rows and columns, respectively
w_j	Environment type j
γ	A candidate path for the search vehicle
η	A candidate path for the characterization vehicle
Ω	The set of candidate paths
\hat{C}_γ	Budget constraint on search mission
\hat{C}_η	Budget constraint on environment characterization mission
$\mathcal{C}(\cdot)$	Cost of traversing a path

$$\begin{aligned}
&= \left| \sum_{q_i \in \eta} \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \bar{\mathbf{P}}_{\eta \setminus q_i} \right. \\
&\quad \left. - \sum_{q_i \in \eta} \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \hat{\mathbf{P}}_{\eta \setminus q_i} \right| \\
&= \left| \sum_{q_i \in \eta} \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \left(\bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right) \right| \\
&\leq \sum_{q_i \in \eta} \sum_{y_{q_i}} \left(P(Y_{q_i} = y_{q_i}) \mathcal{R}(y_{q_i}) \right) \times \left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right| \\
&\leq \sum_{q_i \in \eta} \sum_{y_{q_i}} P(Y_{q_i} = y_{q_i}) \times \left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right| \\
&= \sum_{q_i \in \eta} \left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right| \\
&= |\eta| \left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right|
\end{aligned}$$

Due to the bound on the difference between $\bar{\mathbf{P}}_{\eta \setminus q_i}$ and $\hat{\mathbf{P}}_{\eta \setminus q_i}$ in (29), it follows that

$$\begin{aligned}
P\left(\left| \mathbb{E}[G_\eta] - \mathbb{E}[\hat{G}_\eta] \right| < \epsilon \bar{\mathbf{N}} |\eta| \right) &= P\left(|\eta| \left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right| < \epsilon \bar{\mathbf{N}} |\eta| \right) \\
&= P\left(\left| \bar{\mathbf{P}}_{\eta \setminus q_i} - \hat{\mathbf{P}}_{\eta \setminus q_i} \right| < \epsilon \bar{\mathbf{N}} \right) \\
&\geq 1 - 2 \exp^{-2\epsilon^2 \bar{\mathbf{N}}}
\end{aligned}$$

Appendix B. Proof of Corollary 1

By Theorem 1, we obtain the bounds for η^* in (23) and for $\hat{\eta}$ in (33)

$$\begin{aligned}
P\left(\left| \mathbb{E}[G_{\eta^*}] - \mathbb{E}[\hat{G}_{\eta^*}] \right| < \epsilon \bar{\mathbf{N}} |\eta^*| \right) &\geq 1 - 2 \exp^{-2\epsilon^2 \bar{\mathbf{N}}} \\
P\left(\left| \mathbb{E}[G_{\hat{\eta}}] - \mathbb{E}[\hat{G}_{\hat{\eta}}] \right| < \epsilon \bar{\mathbf{N}} |\hat{\eta}| \right) &\geq 1 - 2 \exp^{-2\epsilon^2 \bar{\mathbf{N}}}
\end{aligned}$$

Note that since $\hat{\eta}$ maximizes (33), $\mathbb{E}[\hat{G}_{\hat{\eta}}] \geq \mathbb{E}[\hat{G}_{\eta^*}]$. Hence,

$$\begin{aligned}
\mathbb{E}[G_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] &= \mathbb{E}[G_{\eta^*}] - \mathbb{E}[\hat{G}_{\eta^*}] + \mathbb{E}[\hat{G}_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] \\
&\leq \left(\mathbb{E}[G_{\eta^*}] - \mathbb{E}[\hat{G}_{\eta^*}] \right) + \left(\mathbb{E}[\hat{G}_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] \right) \\
&\leq \left| \mathbb{E}[G_{\eta^*}] - \mathbb{E}[\hat{G}_{\eta^*}] \right| + \left| \mathbb{E}[\hat{G}_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] \right|
\end{aligned}$$

Then, assuming $|\eta^*| = |\hat{\eta}| = |\eta|$, error in approximation of the optimal characterization gain is

$$\begin{aligned}
P\left(\left| \mathbb{E}[G_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] \right| < 2\epsilon \bar{\mathbf{N}} |\eta| \right) \\
&\geq P\left(\left| \mathbb{E}[G_{\eta^*}] - \mathbb{E}[\hat{G}_{\eta^*}] \right| + \left| \mathbb{E}[\hat{G}_{\eta^*}] - \mathbb{E}[G_{\hat{\eta}}] \right| < 2\epsilon \bar{\mathbf{N}} |\eta| \right) \\
&\geq 1 - 2 \exp^{-2\epsilon^2 \bar{\mathbf{N}}}
\end{aligned}$$

Appendix C. Table showing the list of variables

See Table C.1.

References

- Berger, J.O., 2013. *Statistical Decision Theory and Bayesian Analysis*, second ed. Springer-Verlag, New York, pp. 161–162.
- Bourgault, F., Makarenko, A.A., Williams, S.B., Grocholsky, B., Durrant-Whyte, H.F., 2002. Information based adaptive robotic exploration. In: *Proceedings of the International Conference on Intelligent Robots and Systems*. IEEE, pp. 540–545. <http://dx.doi.org/10.1109/IRDS.2002.1041446>.
- Carrillo, H., Dames, P., Kumar, V., Castellanos, J.A., 2015. Autonomous robotic exploration using occupancy grid maps and graph SLAM based on Shannon and Rényi entropy. In: *Proceedings of the International Conference on Robotics and Automation*. IEEE, pp. 487–494. <http://dx.doi.org/10.1109/ICRA.2015.7139224>.
- Choset, H., 2001. Coverage for robotics—A survey of recent results. *Ann. Math. Artif. Intell.* 31 (1–4), 113–126.
- Chung, T.H., Burdick, J.W., 2012a. Analysis of search decision making using probabilistic search strategies. *IEEE Trans. Robot.* 28 (1), 132–144.
- Chung, T.H., Burdick, J.W., 2012b. Analysis of search decision making using probabilistic search strategies. *IEEE Trans. Robot.* 28 (1), 132–144.
- Coleman, M.C., Block, D.E., 2007. Nonlinear experimental design using Bayesian regularized neural networks. *AIChE J.* 53 (6), 1496–1509.
- De Guenin, J., 1961. Optimum distribution of effort: an extension of the Koopman basic theory. *Oper. Res.* 9 (1), 1–7.
- Dobbie, J.M., 1973. Some search problems with false contacts. *Oper. Res.* 21 (4), 907–925.
- Elfes, A., 1992. Dynamic control of robot perception using multi-property inference grids. In: *Proceedings of the International Conference on Robotics and Automation*. IEEE, pp. 2561–2567.
- Elmore, P., Avera, W.E., Harris, M.M., Duvielh, K.M., Environmental measurements derived from tactical mine-hunting sonar data, in: *Proc. IEEE/MTS OCEANS*, 2007.
- Gader, P.D., Mystkowski, M., Zhao, Y., 2001. Landmine detection with ground penetrating radar using hidden Markov models. *IEEE Trans. Geosci. Remote Sens.* 39 (6), 1231–1244.
- Harris, M., Avera, W., Steed, C., Sample, J., Bibee, L.D., Morgerson, D., Hammack, J., Null, M., AQS-20 through-the-sensor (TTS) performance assessment, in: *Proc. IEEE/MTS OCEANS*, Washington, DC, USA, 2005, pp. 460–465.
- Hayes, M.P., Gough, P.T., 2009. Synthetic aperture sonar: A review of current status. *IEEE J. Ocean. Eng.* 34 (3), 207–224.
- Hoeffding, W., 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58 (301), 13–30.
- Houston, B.H., Bucaro, J.A., Yoder, T., Kraus, L., Tressler, J., Fernandez, J., Montgomery, T., Howarth, T., Broadband low frequency sonar for non-imaging based identification, in: *Proc. IEEE/MTS OCEANS*, 2002, pp. 383–387.
- Kadane, J.B., 1971. Optimal whereabouts search. *Oper. Res.* 19 (4), 894–904.
- Kimeldorf, G., Smith, F.H., 1979. Binomial searching for a random number of multinomially hidden objects. *Manage. Sci.* 25 (11), 1115–1126.
- Koopman, B.O., 1957. The theory of search: III. The optimum distribution of searching effort. *Oper. Res.* 5 (5), 613–626.
- Kress, M., Lin, K.Y., Szechtman, R., 2008. Optimal discrete search with imperfect specificity. *Math. Methods Oper. Res.* 68 (3), 539–549.
- Kriehel, B., Levner, E., Spivak, A., 2016. Optimal search for hidden targets by unmanned aerial vehicles under imperfect inspections. *Am. J. Oper. Res.* 6 (2), 153.
- Makarenko, A.A., Williams, S.B., Bourgault, F., Durrant-Whyte, H.F., 2002. An experiment in integrated exploration. In: *Proceedings of the International Conference on Intelligent Robots and Systems*. IEEE, pp. 534–539. <http://dx.doi.org/10.1109/IRDS.2002.1041445>.
- McMahon, J., Yetkin, H., Wolek, A., Waters, Z., Stilwell, D.J., 2017. Towards real-time search planning in subsea environments. In: *Proceedings of the International Conference on Intelligent Robots and Systems*. IEEE.
- Papadimitriou, C., Beck, J.L., Au, S.-K., 2000. Entropy-based optimal sensor location for structural model updating. *J. Vib. Control* 6 (5), 781–800.
- Pollock, S.M., 1964. Sequential search and detection. Technical Report, Massachusetts Institute of Technology, Cambridge Operations Research Center.
- Richardson, H.R., 1988. Search theory. *Encycl. Stat. Sci.*
- Shende, A., Bays, M.J., Stilwell, D.J., Toward a mission value for subsea search with bottom-type variability, in: *Proc. IEEE/MTS OCEANS*, Virginia Beach, VA, USA, 2012.
- Stone, L.D., Stanishine, J.A., Persinger, C.A., 1972. Optimal search in the presence of Poisson-distributed false targets. *SIAM J. Appl. Math.* 23 (1), 6–27.
- Takahashi, K., Igel, J., Preetz, H., 2011a. Clutter modeling for ground-penetrating radar measurements in heterogeneous soils. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 4 (4), 739–747.

- Takahashi, K., Preetz, H., Igel, J., 2011b. Soil properties and performance of landmine detection by metal detector and ground-penetrating radar—soil characterisation and its verification by a field test. *J. Appl. Geophys.* 73 (4), 368–377.
- Wah, B.W., Yu, C.F., 1985. Stochastic modeling of branch-and-bound algorithms with best-first search. *IEEE Trans. Softw. Eng.* SE-11 (9), 922–934.
- Yamauchi, B., 1997. A frontier-based approach for autonomous exploration. In: Proceedings of International Symposium on Computational Intelligence in Robotics and Automation. IEEE, pp. 146–151. <http://dx.doi.org/10.1109/CIRA.1997.613851>.
- Yetkin, H., Lutz, C., Stilwell, D.J., Utility-based adaptive path planning for subsea search, in: OCEANS 2015 - MTS/IEEE Washington, 2015, pp. 1–6.
- Yetkin, H., Lutz, C., Stilwell, D.J., Acquiring environmental information yields better anticipated search performance, in: OCEANS 2016 - MTS/IEEE Monterey, 2016, pp. 1–6.
- Zare, A., Cobb, J.T., 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.