Utility-Based Adaptive Path Planning For Subsea Search

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

***Abstract*—We present a decision-theoretic adaptive path planning approach to efficiently search for an unknown number of stationary objects distributed in a bounded search area subject to time or distance constraint. We find the optimal finite-length search trajectory by maximizing a decision-theoretic utility function, and we re-plan continually as new information is gathered. We assume that the performance of the search sensor varies with variations in the environment, and our planning approach incorporates stochastic predictions of the environment if available. Our work is inspired by subsea applications. We compare our approach with a naive mowing-the-law approach that is commonly used in subsea search. We also investigate how uncertain information affects search performance.**

I. Introduction

This paper presents a non-myopic adaptive approach to search planning that is based on maximizing the expected value of a decision-theoretic utility function. We address the case that a mobile search agent aims to locate an unknown number of stationary objects distributed in a bounded search domain. We consider that due to limited energy storage capacity of the search agent, the search task is subject to a time or distance constraint. The bounded search area is partitioned into a number of disjoint cells and we assume there is a known upper bound for the number of objects residing in each cell. Local environmental conditions in the search area are assumed to affect the sensor performance, which renders some parts of the search domain more informative than others. We show that our adaptive search strategy leads the search agent to explore more informative parts of the search area and maximizes the aggregate certainty about the number of objects.

Search theory has its roots in numerous civilian and military applications. A common task in much of the search literature is to find the most likely location of a single stationary object in some search domain 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. In a realistic search problem, there are certain limitations for successfully locating the target such as imperfect sensor readings or uncertain knowledge of the search environment. Noisy sensor observations 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 [1] and [2] consider the effect of missed detections. However, the issue of false alarms is less often addressed, see for example [3–7]. In contrast, false alarms are addressed in [8] and [9]. We build upon prior work by accounting for false alarms, missed detections and uncertainty in the environment. Our goal is to maximize the probability of detection in a limited time.

This paper extends and clarifies the probabilistic search strategy in [10]. Although our work is motivated by subsea applications, it can be applied to a variety of search problems including search and rescue applications and mine countermeasures. We maximize expected percent clearance which is the expected probability that an estimate of the number of objects at a location is correct. Our adaptive mission planning is implemented in a receding horizon strategy in which an optimal finite-length path is continually computed using real-time data as the search agent travels. Our main contribution in this study is a decision-theoretic cost function for realistic search problems that include the effects of false alarms, missed detections, uncertainty in the environment, and any prior information that might be available on the number of objects at a location. We also investigate how uncertain environment information affects search performance.

The remainder of this paper is organized as follows. In Section II, we formulate the search problem and define the observation model. In Section III, we describe our proposed cost function that is based on a zero-one utility function. In section V, we describe the limitations of search based on our cost function and offer a brief literature review on possible workarounds. Section VI shows the numerical illustration of the proposed search strategy.

II. Problem Formulation

A. Preliminaries

Given a bounded search area $S = \{s(1), s(2), \dots, s(r)\}$ partitioned into a grid with r disjoint cells, for $i \in \mathbb{N}_r^+$ where $\mathbb{N}_r^+$ denotes the set of integers from 1 to r, cell $s(i) \in S$ denotes the ith cell in the search area. We associate with each cell a random variable X which represents the number of objects in the cell. We presume $X(i)$ is independent of $X(j)$ when $i \neq j$. The searcher's objective is to estimate $X(1), \dots, X(r)$ 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 the environment in each cell and

978-0-933957-43-5 ©2015 MTS

a corresponding stochastic description of sensor performance in the environment.

We assume that the environment in each cell is from a finite set of possible environments $\mathbf{e} = \{b_1, b_2, \ldots, b_m\}$. That is, for all $s(i) \in S$, $E(i)$ is the random variable and $e(i) \in \mathbf{e}$ is the environment. 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 $\Pi(i) = [p_1(i), p_2(i), \ldots, p_m(i)]$ 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

When the searcher vehicle samples from a cell, it returns with noisy object count observation $Z = z$. Observed data z depends on true number of objects and governing environmental conditions in the cell. The value of the measurement z is the sum of the correct detections d and possible false alarms f. As the physics that generate the correct detections and false alarms are different, we assume they are probabilistically independent. We model the likelihood of observing $Z = z$ objects if the search area is governed by the environment $b_j \in \mathbf{e}$

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

where $P_D(d = l \mid x, b_j)$ is the probability that the sensor detects l objects which reside in the cell, and $P_F(f = k \mid b_j)$ is the probability that the sensor returns with k number of false alarms. Convolution of P_D and P_F follows from the assumption that the number of missed detections and false alarms are statistically independent. We model the probability of false alarms with a geometric distribution, and the probability of correct detections with a Binomial distribution,

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

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

where $0 < \alpha_j \leq 1$ denote the probability of one or more false alarms and $0 < D_j \leq 1$ 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,

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

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 those. However, we note that other expressions are also possible for the false alarm model, and that our results do not depend on this specific false alarm model except for numerical illustrations. We refer the reader to [10] for more in-depth description of how (5) is derived.

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 $Z = z$ 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 as

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

For any cell $s(i)$ with the environment probability distribution $\Pi(i) = [p_1(i), p_2(i), \ldots, p_m(i)]$, we update the probability of $x(i)$ objects located in the cell when $Z(i) = z(i)$ is observed

$$P_{\Pi(i)}\Big(x(i) \mid z(i)\Big) = \sum_{j=1}^{m} P\Big(x(i) \mid z(i),\ b_j\Big) P\Big(b_j \mid z(i)\Big) \tag{8}$$

where

$$P\Big(b_j \mid z(i)\Big) = \frac{P\Big(z(i) \mid b_j\Big) P\Big(b_j\Big)}{P\Big(z(i)\Big)} \tag{9}$$

III. Adaptive Path Planning

We assume that the goal of a search mission is to maximize *percent clearance* in a fixed amount of time. Percent clearance for each cell can be interpreted as the probability that the estimated number of objects in that cell is correct. Since percent clearance is dependent on observed data which is not available before the mission is executed, we select paths that maximize *expected* percent clearance. Using a zero-one utility function, it is possible to express percent clearance as the expected value of the utility function.

Given measured data z, we define the utility of forming the estimate $\delta(z)$ when x is the true value of parameter X

$$U\Big(x, \delta(z)\Big) = \begin{cases} 1 & \text{if } x = \delta(z) \\ 0 & \text{if } x \neq \delta(z) \end{cases} \tag{10}$$

We then compute the expected utility of visiting a cell for each environment $b_j \in \mathbf{e}$,

$$\mathbb{E}\Big[U\Big(x, \delta(z)\Big) \mid b_j\Big] = \sum_z \sum_x P(x \mid z, b_j) P(z \mid b_j) U\Big(x, \delta(z)\Big) \tag{11}$$

$$= \mathbb{E}\Big[\max_x P(X = x \mid z,\ b_j)\Big] \tag{12}$$

which matches with expected probability of correctly estimating the number of objects assuming environment b_j. Note that $\max_x P(X = x \mid z,\ b_j)$ is percent clearance after observing the data z. The zero-one utility function in (10) 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.

Measurements from different cells are independent, and thus percent clearance for a path that passes through multiple cells is simply a product of the percent clearance 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 $\gamma = \{q_1, \ldots, q_N\}$ be a path of length N where $q_i \in \mathbb{N}_r^*$ represents the cell index of the ith element of γ. 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 $\gamma_d \subset \gamma$ denote the set of distinct cells in γ. 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 $\sum_{i=1}^{M} m_i = N$. Let $\mathbf{z}(q_i) = [z_1(q_i), \ldots, z_{m_i}(q_i)]$ be the m_i independent measurements acquired at q_ith cell. The expected utility of traversing γ is

$$\mathbb{E}\Big[U_\gamma\Big(x, \delta(\mathbf{z}(\gamma))\Big)\Big] = \prod_{q_i \in \gamma_d} \mathbb{E}\Big[U\Big(x(q_i), \delta(\mathbf{z}(q_i))\Big)\Big] \times \prod_{i \in S \setminus \gamma} \max_{x(i)} P\Big(X(i) = x(i)\Big) \tag{13}$$

where $\max_{x(i)} P(X(i) = x(i))$ is the prior certainty of the ith cell and

$$\mathbb{E}\Big[U\Big(x(q_i), \delta(\mathbf{z}(q_i))\Big)\Big] = \sum_{j=1}^{m} P(E(i) = b_j) \times \mathbb{E}\Big[\max_{x(q_i)} P\Big(x(q_i) \mid \mathbf{z}(q_i),\ b_j\Big)\Big] \tag{14}$$

This shows that the expected utility of traversing a path is the expected increase in the aggregate certainty of the number of objects for each cell. Thus (13) implicitly penalizes multiple visits to a cell and encourages the searcher vehicle to explore the cells that haven't been visited yet. Let $\Gamma(k)$ denote the finite collection of N-length paths available to the searcher at time step k. Then, the optimal path is

$$\gamma^\star(k) = \arg \max_{\gamma_j \in \Gamma(k)} \mathbb{E}\Big[U_\gamma\Big(x, \delta(z(\gamma_j))\Big)\Big] \tag{15}$$

We assume the search mission is subject to a time or distance constraint. We call this constraint the *mission length*, i.e. total number of cells the vehicle can travel through during the search 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. Planning the path for a horizon smaller than the mission length may yield suboptimal trajectories, but the computational expense of adaptive planning is reduced to a manageable level for autonomous operations.

IV. Effect of Uncertainty in the Environment

In Section II-A, we define the search problem with uncertainty in the environmental conditions. We assume that the environment is from a finite set of possible environments, and the environment type affects performance of the search sensor. However, only stochastic information about the environment is available to the searcher. In this section, we analyse the effect of environment uncertainty on search performance.

We have shown that the observation model is dependent on the environment type. Therefore, uncertainty in the observation model should lead to uncertainty in the posterior distribution of the object count which will cause expected percent clearance when the environment is uncertain be different than expected percent clearance when the environment is known.

We seek a method of formally characterizing the effect of environment uncertainty on the deviation in expected percent clearance. For each cell $s(i) \in S$, let $\Pi(i) = [p_1(i), p_2(i), \ldots, p_m(i)]$ be the given environment distribution and $e(i) \in \mathbf{e}$ be the true environment. For notational convenience, let $V(\cdot)$ denote expected percent clearance under uniform prior $P(X(i) = x(i))$, e.g. $V(e(i)) = \mathbb{E}\big[\max_{x(i)} P(x(i) \mid z(i), e(i))\big]$. We define μ to be the total deviation from the expected percent clearance when true environment for each cell is known.

$$\mu = \sum_{i=1}^{r} \Big|\ V\Big(e(i)\Big) - V\Big(\Pi(i)\Big)\ \Big| \tag{16}$$

As the certainty in the environment increases, μ in (16) is reduced and consequently we expect that planned expected percent clearance before a mission better matches actual percent clearance after a mission. We say we underestimate the environment if $V(e) > V(\Pi)$, and overestimate if $V(e) < V(\Pi)$. While underestimating the environmental conditions may decrease search efficacy, overestimation may yield inaccurate decisions on object count. In subsea applications, the consequences of the latter case might be severe. In Section VI, we present the results of our Monte Carlo analysis that suggests reduced environmental uncertainty improves search performance.

In [11], the authors conduct mowing-the-law experiments first when there is uncertainty in the environment, and later when the environmental conditions are deterministically known. They show that deterministic knowledge of the environment shortens the mission duration and improves search performance.

V. Limitation of Proposed Search Strategy

In Section III, we present our adaptive path planning approach. We show that among paths available to the searcher vehicle, the path that maximizes the expected value of a utility function is selected. In this section, we seek to formally characterize some of the limitations of our proposed search strategy. Given the utility function in (10), we compute expected utility for each possible environment $b_1, \dots b_m$. This gives us the set of possible outcomes of the search. Let O denote the set of possible outcomes, we have

$$O = \left\{ \mathbb{E}\Big[\max_x P\Big(x \mid z,\ b_1\Big)\Big], \dots, \mathbb{E}\Big[\max_x P\Big(x \mid z,\ b_m\Big)\Big] \right\} \tag{17}$$

where $\mathbb{E}\Big[\max_x P\Big(x \mid z,\ b_j\Big)\Big]$ matches with the expected probability that our estimate is correct given the cell is governed by the environment b_j. Expected utility theory suggests that for a cell with environment probability distribution $\Pi = [p_1, p_2, \dots, p_m]$, we compute the expected value of these outcomes.

$$\mathbb{E}\Big[U\Big(x, \delta(z)\Big)\Big] = \sum_{j=1}^{m} P(b_j)\ \mathbb{E}\Big[\max_x P(x \mid z,\ b_j)\Big] \tag{18}$$

However, this method fails to penalize the uncertainty in the environment. We compare the utility of searching different cells based on the mean values of their corresponding probability distributions on outcomes. If a cell with uncertainty in the outcome is chosen, depending on the variation in the outcome, this may result in achieving a clearly undesirable expected percent clearance. We note that achieving a greater expected percent clearance than assumed has no benefit since it does not enhance our trust on the object count estimates. On the other hand, if the expected percent clearance turns out to be worse, this may cause over-confidence on our estimates and lead to inaccurate estimations. Thus, among different choices with the same mean value but different environmental distributions, there are likely many practical scenarios where the choice with less variation in the outcomes should be selected.

In the literature, the goal of reducing uncertainty when maximizing expected utility is sometimes referred to as ambiguity aversion [12]. More generally, ambiguity aversion addresses the principal that one should choose a known risk over unknown risks. There exist a number of well-known approaches to addressing ambiguity aversion. While some of these approaches are not suitable for the search problem, for example [13] and [14], we believe the rank dependent utility theory [15] offers a promising solution.

VI. Numerical Results

In this section, we present simulation results that show the efficacy of adaptively selecting paths that maximize percent clearance, given stochastic assumptions on the environment and sensor performance. Our contribution is a cost function for selecting paths rather than a fundamentally new class of adaptive search algorithm. Our numerical illustrations aim to evaluate search performance for simplistic scenarios that are inspired by mine-hunting missions.

We divide the bounded search area S into a grid with 10×10 non-overlapping cells. For each cell $s(i) \in S$, we assume there is $0 \le x(i) \le L$ number of objects bounded by an upper bound L. In mine hunting missions, $x(i)$ represents the number of mines residing in cell $s(i)$. We assume no prior information exists about the object count in any cell. We note that L is typically not known beforehand; however, letting L be a sufficiently large number will capture every possible scenario. In our simulations, we set $L = 2$. We assume there are three possible environments in the search area. Probability of detection, D, and probability of at least one false alarm, α, for each environment are given as follows: for the first environment, $D = 0.65$ and $\alpha = 0.4$, for the second environment, $D = 0.8$ and $\alpha = 0.3$, and for the third environment, $D = 0.9$ and $\alpha = 0.1$. We note that the amount of information about the object count in a cell increases with the probability of detection and decreases with the probability of false alarm of the environment to which the cell belongs. In this regard, the first environment is the least and the third environment is the most informative environment. In Fig. 1, we define the search area and introduce corresponding environment distributions for each cell. Note that cell colors indicate the most probable environment for each cell. For darker gray cells, the most likely environment is the first environment, for medium gray cells the second environment and for lighter gray cells the third environment.

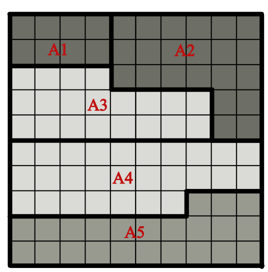

A1	$\rightarrow$	$[0.9, 0.05, 0.05]$
A2	$\rightarrow$	$[0.85, 0.1, 0.05]$
A3	$\rightarrow$	$[0.1, 0.1, 0.8]$
A4	$\rightarrow$	$[0, 0.1, 0.9]$
A5	$\rightarrow$	$[0.2, 0.1, 0.7]$

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

In subsea applications, 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 the 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. The vehicle passes through cells that are outside of the search area when transitioning between rows. Passing cells that are

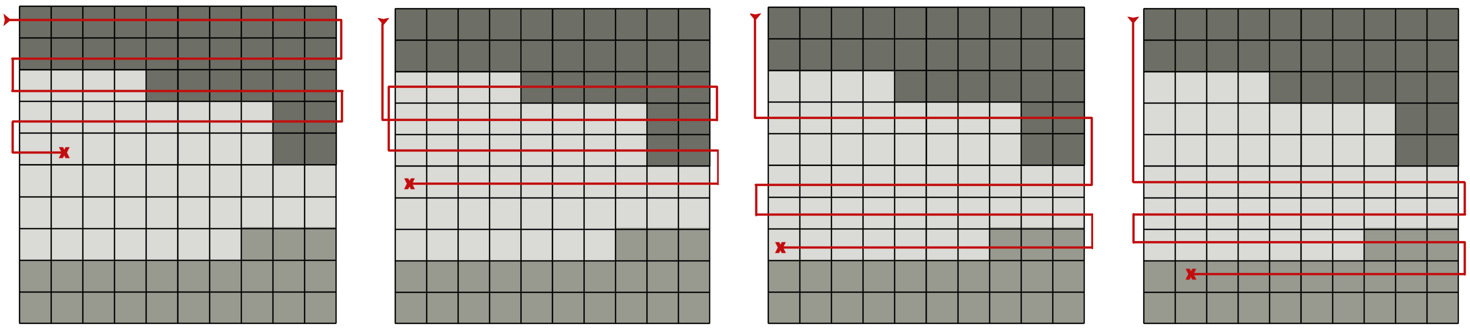

(a) (b) (c) (d)

Fig. 2: Results for mowing-the-law search and adaptive search for different path lengths. (a) planned path with mowing-the-law search (b) and adaptive search when path length is 7 (c) path length is $10, 15$ (d) path length is $20, 25$.

outside the search area require time but do not reduce percent clearance since no measurements are acquired. A possible search trajectory is defined as a set of connected cells. Given a mission length for the search mission, the search trajectory contains grid cells and possible transition cells where the total number of these cells matches with the mission length. For each simulation, we assume the mission length is 50.

We assume that the mission is so long that an optimal solution cannot be computed in near real-time using reasonable computational resources. Therefore we adopt a receding horizon strategy. An optimal N-length path is computed where N is much shorter than the total mission length. The searcher vehicle travels to the first cell on the path, samples the cell, updates percent clearance for the cell, and then re-computes an optimal N-length path using whatever new information is available.

We conduct eight different numerical tests. We first employ a naive mowing-the-law approach as a baseline against which to assess the efficacy of the adaptive path planning strategy. Fig.2a shows the search trajectory when mowing-the-law approach is employed. The Red line represents the searcher vehicle path. The red cross is the end cell of the path. The searcher vehicle travels back and forth until the mission length is met. For the other tests, we employ the adaptive search strategy for varying path lengths N and record the computational time spent for planning the path. Fig.2b through Fig.2d show optimal N-length receding horizon search trajectories when the path length N is 7 (Fig.2b), when N is 10 or 15 (Fig.2c), and when N is 20 or 25 (Fig.2d). Corresponding expected percent clearance and computational time for each test are given in Table I. To avoid the issue of arithmetic underflow, we compute negative logarithm of expected percent clearance and sum over each cell. Note that negative logarithm of expected percent clearance increases as expected utility decreases, i.e. smaller negative log of expected percent clearance implies better search performance. The results show that the adaptive search strategy yields better expected percent clearance when compared to mowing-the-law approach. It is also seen that search performance doesn't necessarily improve when path length increases. For example, expected percent clearance when path length is 10 is better than expected percent clearance when path length is 20. Fig. 2d shows that when the path length is 20, the sixth and seventh rows, the most informative parts of the search area, are within the planning horizon. However, after traveling through these rows, the vehicle is stuck at row 9 which is the less informative part of the search area. On the other hand, Fig. 2c shows that for a 10-length path, the sixth and seventh rows are not within the planning horizon and the vehicle travels through the forth row, instead. This prevents the vehicle to travel through row 9. This phenomena suggests that it is important for the vehicle to detect when undesirable results might occur. The results also show that computational time grows exponentially with the path horizon and for some large path horizon autonomous search might be practically infeasible.

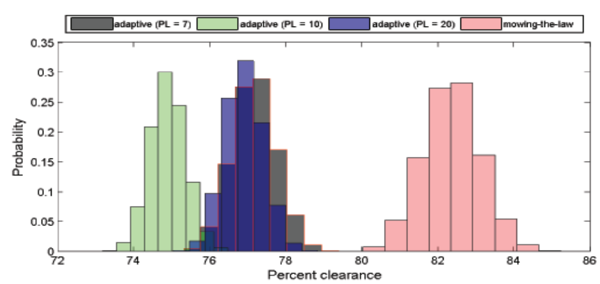

Fig. 3: Frequency of percent clearance for mowing-the-law and adaptive search for different path lengths (PL)

TABLE I: Search Performance

Path length	Expected percent clearance	Computational time
Mowing-the-law	82.2	–
Adaptive ($N = 5$)	82.2	0.76 sec
Adaptive ($N = 7$)	77.4	1.65 sec
Adaptive ($N = 10$)	74.9	6.1 sec
Adaptive ($N = 15$)	74.9	29.8 sec
Adaptive ($N = 20$)	76.7	216.6 sec
Adaptive ($N = 25$)	76.7	779 sec
Adaptive ($N = 50$)	74.9	$>$ 2 hour

We conduct Monte Carlo simulations to observe variation in percent clearance for each search trajectory given in Fig. 2. We iterated for 10000 times to eliminate the effects due to random nature of observations and plotted the results in Fig. 3. The results show that the mean of the random observations for each search trajectory matches the expected percent clearances given in Table I. Note that when the mean for a search trajectory increases, the variation in percent clearance also increases. This is because when the vehicle samples from less informative parts of the search area the observations have more variation, which yields more variation in percent clearance.

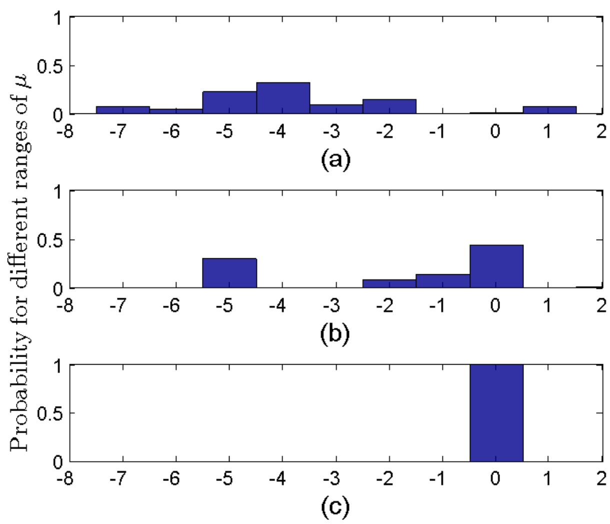

Fig. 4: Frequency of deviation from actual expected percent clearance for different ranges of μ. (a) when $16 \leq \mu \leq 18$, (b) when $10 \leq \mu \leq 12$, (c) when $3 \leq \mu \leq 5$

In order to simulate the effect of environment uncertainty on search performance, we conduct Monte Carlo simulations where the environment in each cell is randomly drawn from the environment distributions in Fig. 1. We analyse how expected percent clearance after a search mission changes when μ in (16) decreases. We define *actual* expected percent clearance to be expected percent clearance that we in fact achieve instead of expected percent clearance assumed to be achieved when environment is uncertain. Fig. 4 shows that as μ decreases, the difference between actual expected percent clearance and expected percent clearance when environment is uncertain decreases. This shows that when μ is closer to zero, our search results provide more reliable information about the object count.

VII. Conclusions

In this paper, we present a decision-theoretic adaptive path planning approach to search for an unknown number of stationary objects distributed in a bounded search area subject to time or distance constraint. Our observation model accounts for possible false alarms, missed detections and environment uncertainty in the search area. We investigate the effect of environment uncertainty on search performance and present a new measure μ for the uncertainty in the environment. We conduct Monte Carlo analyses that showed that decreasing the environment uncertainty improves search results.

References

[1] 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.

[2] 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.

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

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

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

[6] 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.

[7] 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.

[8] 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.

[9] K. E. Wilson, R. Szechtman, and M. P. Atkinson, "A sequential perspective on searching for static targets," *European Journal of Operational Research*, vol. 215, no. 1, pp. 218–226, 2011.

[10] A. Shende, M. J. Bays, and D. J. Stilwell, "Toward a mission value for subsea search with bottom-type variability," in *Oceans, 2012*. IEEE, 2012, pp. 1–5.

[11] 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.

[12] D. Ellsberg, "[risk, ambiguity, and the savage axioms]: Reply," *The Quarterly Journal of Economics*, pp. 336–342, 1963.

[13] I. Gilboa and D. Schmeidler, "Maxmin expected utility with non-unique prior," *Journal of mathematical economics*, vol. 18, no. 2, pp. 141–153, 1989.

[14] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica: Journal of the Econometric Society*, pp. 263–291, 1979.

[15] J. Quiggin, "A theory of anticipated utility," *Journal of Economic Behavior & Organization*, vol. 3, no. 4, pp. 323–343, 1982.