

Robust Search for Structured Object Placement using Unmanned Vehicles

John G. Baylog and Thomas A. Wettergren
Naval Undersea Warfare Center
1176 Howell Street
Newport, RI 02841
baylogjg@npt.nuwc.navy.mil
t.a.wettergren@ieee.org

John C. Hyland and Cheryl M. Smith
Naval Surface Warfare Center
110 Vernon Avenue
Panama City, FL 32407
john.hyland@navy.mil
cheryl.m.smith1@navy.mil

Abstract—We develop a robust method for planning an undersea search by unmanned vehicles; the goal of which is to find groups of objects that may be placed randomly or arranged in predefined patterns on the ocean floor. This approach revolves around the computation of target pattern priors that describe the underlying geometrical structure of the expected group of objects. These priors are updated based on limited search observations, followed by a search performance assessment in the remaining region utilizing these observed target pattern distributions. We outline the derivation of this approach to assessing the performance of unmanned searches. The algorithms performance is tested in a simulation example to illustrate the improvement in search evaluation accuracy provided by the consideration of object structure. We conclude with a discussion of how this search planning strategy can be employed in realistic scenarios, searching for objects that are expected to appear in more general patterns.

I. INTRODUCTION

The development of methods for evaluating the expected performance of a search for hidden objects has a long history, going back to the early work of Koopman [1] and others in the 1940s. That early work focused on problems of finding single randomly placed objects in a region using random search patterns. Over time, the scope of problems studied advanced to examining structured ladder search patterns [2] and optimal spatial allocations of search effort [3], yet the complexity of the analysis has continually been placed in the search allocation. The object of search has traditionally been a single randomly placed object. However, in many undersea applications, searches are conducted for groups of objects that occur in patterns on the ocean floor, not just single independently placed random objects. These patterns occur either due to the logistical demands of deployment (for man-made objects), the tactical utility of patterns (for minefields), or simply are the result of bathymetric features in a highly non-homogeneous environment. Thus, to accurately characterize searches for these objects, it is important to properly account for the geometric structure in the placement of the group of objects.

Once a capability exists to accurately and robustly evaluate expected search performance, it is possible to develop improved planning aids. In particular, the search evaluation component of a search planning aid is often the fundamental

component of technology that separates methods that lead to poor searches from methods that lead to great searches. Given parameterizations of search performance, various optimization strategies may be utilized to identify the parametric combinations that yield the best performance. These parameterizations, in turn, provide a representation of the desired search strategy.

When unmanned vehicles are used as the platforms performing the search, there is additional complexity in the search planning process. In particular, large search missions are often covered by a group of vehicles making multiple individual sorties. While this may create a logistical challenge, it enables improved search through the use of early searches to inform the remaining searchers of the nature of any observed structure in the search objects pattern. We thus concern ourselves with the problem of using the result of a partial search to improve our understanding of the underlying structure of object placement in the ocean. Furthermore, new search evaluation methods are created to utilize this object structure to more accurately represent the expected search performance. Given this search performance capability, the development of adaptive methods for re-planning future vehicle paths based on early search results is realizable.

In this paper we develop the methodology for assessing the performance of these searches for structured object placement. In the next section, we outline the components of this assessment strategy and provide an overall framework for the search evaluation process. In section III, the development of probability priors representing the object pattern structure is presented along with an algorithm for updating these priors based on results of a partial search. Then, in section IV, we present a method for evaluating the performance of a search for objects placed according to this probability description of the object placement structure. Finally, we conclude the paper with a simulation example of a search for undersea mines that are expected to occur in line patterns. This example illustrates the performance of both the prior probability calculations as well as the search effectiveness prediction. In the context of the example, we then explain how these methods may be applied to planning these complex multi-vehicle search operations.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Robust Search for Structured Object Placement using Unmanned Vehicles				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Undersea Warfare Center,1176 Howell Street,Newport,RI,02841				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002176. Presented at the MTS/IEEE Oceans 2008 Conference and Exhibition held in Quebec City, Canada on 15-18 September 2008.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

II. FRAMEWORK FOR SEARCH PLANNING, EVALUATION, AND ADAPTATION

When planning searches for unmanned vehicles, search theoretic methods use performance estimates of the search sensors to determine search trajectories and any required overlap necessary to achieve the desired search effectiveness. However, existing search evaluation methodologies are often based on singular object placement paradigms and uniform placement distributions. As such, when employed for multiple object searches, the anticipated performance can be misleading and lead to less efficient search trajectories for the autonomous platform in the given period of time. When trying to utilize scarce resources in a cost-efficient manner, such inefficiencies should necessarily be examined in the context of the overall system design. By considering object placement structure in the specification of the placement density prior, we have developed an effective method to evaluate search effectiveness when such placement strategies are employed. This methodology is used to more effectively plan exploratory searches over regions of interest when seeking objects aligned in such patterns.

The overall search planning, updating, evaluation, and re-planning methodology is outlined in Fig. 1. In this process, the search problem is described by the two components on the left side of the figure. These two components enter into the *search evaluation* process inside of the upper feedback loop, corresponding to the *plan optimization* process. At each iteration within the plan optimization process, the *vehicle path* is updated and the evaluation is re-processed. Given an optimal plan, we begin execution of the search. As some data is returned from early sorties (represented as the *execute partial search* block), the priors that describe the search object parameterization are updated in the *update priors* process. This lower feedback process provides the ability to re-plan searches based on these new priors, leading to a new iteration of the upper feedback process of plan optimization. This process repeats until the search effort is completed (or at least until the last sortie is planned). In the remainder of this paper, we describe the details of the two critical blocks that are shaded in Fig. 1. In the next section, we describe the *update priors* block, and in section IV the *search evaluation* block is detailed. The *execute partial search* block is not discussed, since it merely represents a physical process (albeit one we simulate in the examples), and the *plan optimization* function is beyond the limited scope of this paper, its development is the subject of ongoing research.

III. SEARCH AREA CHARACTERIZATION AND PRIOR FORMULATION

To illustrate the procedure for updating object pattern structure priors, we focus on the example problem of finding undersea mines in a region that is (potentially) anticipated to contain mine lines. In this manner, we illustrate a specific type of prior associated with the combination of mine line characteristics mixed with randomly placed mine-like objects, and thus provide a context for using prior information. This is also a generically illustrative example since many structures




Fig. 1. Unmanned Vehicle Search Planning and Re-planning Process

underwater are laid in lines due to the necessity of deployment from moving crafts. In this section we outline an algorithm for utilizing prior information about the nature of mine lines to examine a group of mine-like sensor contacts to sort which are likely in lines from those that are not. We then show how this information is used to update the priors on expected search object structure for use in search evaluation.

As currently implemented, the procedure to characterize a search area in terms of likely mine lines assumes that the search sensors all have identical, constant probabilities of detection and identical, constant false alarm rates. The sensors search an area and a classifier algorithm then classifies all the resulting detections as either mine-like or non-mine-like. All mine-like contacts (*mlc*'s) are then passed on to the search area characterization algorithm which provides further sub-classification into one of three classes: *fixed* – the *mlc* fits into a mine line pattern with fixed mine spacing; *random* – the *mlc* fits into a mine line pattern with random mine spacing; or *false alarm* – the *mlc* does not fit into a mine line pattern and is subsequently not considered a mine.

A. Line Pattern Algorithm Description

Since a unique line is defined by two points, we begin by considering all possible pairings of detections that fall between prescribed minimum and maximum mine line lengths. For each candidate pair (p_i, p_j) , the algorithm creates a candidate rectangular mine lane formed by extending the (p_i, p_j) line by the predicted off-track mine placement error. For convenience, the point with the smallest x -value is labeled p_0 and the remaining point is labeled p_1 . Fig. 2 shows that only mines inside this region are considered candidate mines.

Given N individual *mlc*'s, the algorithm's uncontrolled computational burden to check for lines is $\mathcal{O}(N^2)$. Before performing more detailed and computational intensive analysis on the candidate line, the algorithm immediately applies a false alarm count test to dismiss lines that do not have high probabilities. The number of expected false alarms N_{FA} in the candidate mine lane region is

$$N_{FA} = 4\lambda d\sigma_w \quad (1)$$




Fig. 2. Candidate Mine Lane Region

where

$$\begin{aligned}\lambda &= \text{sensor false alarm rate (FA/m}^2\text{)} \\ d &= \text{length of candidate line (m)} \\ \sigma_w &= \text{off-track placement error standard deviation (m)}\end{aligned}$$

Let H_{ML} and $H_{\overline{ML}}$ denote the respective hypotheses that a mine line is present or not. Let the variable C represent the number of targets in the mine lane. We assume that the false alarm count, given by the number of false alarms k , behaves as a fixed rate Poisson process. Therefore, under the null hypothesis $H_{\overline{ML}}$, the probability of observing at least C false alarms is the Poisson probability, with mean N_{FA} , of obtaining a value of at least C . Let CT denote this probability such that $CT = \Pr(k \geq C | H_{\overline{ML}})$.

Only candidate lines that meet or exceed the $\alpha = 0.05$ level of significance pass this preliminary screening. Furthermore, we assume that when a mine line is present

$$\Pr(\text{observed targets} \geq C | H_{ML}) = 1.0 \quad (2)$$

such that the observed count along the line is at least as big as the actual number of targets along the mine lane.

In high false alarm environments, the odds of randomly choosing the pair (p_i, p_j) to both be mines is approximately 2 to 98. In order to be conservative, we then take the prior probability that (p_i, p_j) correctly defines a mine line as 0.05. Therefore, the probability of obtaining an observed number of mlc 's X on a line that is as large (or larger) than C is

$$\begin{aligned}\Pr(X \geq C) &= \Pr(X \geq C | H_{ML}) \Pr(H_{ML}) \\ &\quad + \Pr(X \geq C | H_{\overline{ML}}) \Pr(H_{\overline{ML}}) \\ &\approx 1.0 \times 0.05 + CT \times 0.95\end{aligned} \quad (3)$$

Then, the posterior probabilities are

$$\begin{aligned}\Pr(H_{\overline{ML}} | C) &= \Pr(C, H_{\overline{ML}}) / \Pr(C) \\ &\approx \frac{CT \times 0.95}{1.0 \times 0.05 + CT \times 0.95}\end{aligned} \quad (4)$$

and

$$\Pr(H_{ML} | C) = 1.0 - \Pr(H_{\overline{ML}} | C). \quad (5)$$

This surprising simple calculation provides an adequate measure for the likelihood of the presence of a mine line. If the posterior probability of a mine line (from equation (5)) exceeds a pre-determined threshold p_{min} , we assume a line is present and proceed to further analyze its properties to determine whether the mine line has fixed or random mine spacing.

Under the assumption that p_0 and p_1 are both mines and part of a fixed spacing pattern, then the distance between p_0 and p_1 must be nearly an integer multiple of the fixed spacing. We assume that the miner used due diligence in placing the mines. That is, the mine spacing pattern is within the suggested values for the type of mine and water depth, and the miner followed the pattern to within navigational capabilities. With these assumptions we can set a minimum and maximum spacing and estimate the placement error. Using this minimum and maximum along with the known distance between p_0 and p_1 , we have a finite list of possible mine spacings. For each possible spacing we use the estimated placement error to define rectangular cells likely to contain mines as shown in Fig. 3.




Fig. 3. Candidate Mine Cell Regions

Under the assumption that each cell contains a mine, the sequence of detections become a sequence of Bernoulli trials and the number of detections in cells is binomial with probability of success equal to the sensor probability of detection P_D . Also, all targets in the mine lane area but outside of the cells are assumed to be false alarms, so the count of targets outside cells should be Poisson random variable with parameter λ times the area. Applying these deductions, one can write the exact probability of obtaining the observed count. Unfortunately, though technically correct, this calculation is highly sensitive to errors in the sensor's estimated probability of detection P_D , and produces values too small for accurate calculations of the posterior probabilities.

To avoid these issues, we concentrate not on the targets in the cells, but rather on the number of targets that fail to be covered by the cells. If too many targets lie outside the cells, we clearly have the wrong cell spacing. For the conditional probability of obtaining as many as the observed number of targets outside the cells in the presence of mine line with u spacing, we again use the Poisson distribution, this time with

mean equal to the area outside the cells times the sensor false target rate λ .

The binomial nature of the target detections is not completely ignored. Using the normal approximation to a binomial random variable, we include a loose binomial qualifying test [4] for fixed spaced lines, requiring that the count be greater than the expected mean less one standard deviation. If the candidate line fails the binomial qualifying test, the posterior probability for the fixed-spaced line is set to zero. Otherwise it is calculated from the possible spacings.

The lack of an adequate prior for the true value of u makes the calculation of the overall probability of a fixed spaced mine line difficult. Initial attempts to assign equal likelihood to all reasonable values for u again produced unmanageable calculations. In practice, a few related spacings produce conditional probabilities several orders of magnitude larger than the rest. However, among these large probabilities, one u typically stands out clearly as the “best” choice. Without a good source of prior information, our solution is to take the largest candidate u that maximizes the probability of the observed count and assume that it is the only possible u . If the candidate line with the chosen u satisfies all qualifying tests, the prior probabilities for all other u ’s are set to zero. The posterior probability for the fixed-spaced mine line is then calculated using only the conditional probability of the fixed-spaced line given the observations.

If the candidate line fails some criterion for fixed-spaced lines, the posterior probability for a randomly-spaced line is calculated as the difference between the fixed line posterior and the total posterior probability for a mine line. Otherwise, the probability for a random line is set to zero. This bias toward fixed spaced line is necessary since all lines meet the random line criteria.

To insure that spurious random points are not attached to the candidate line a final requirement is that the maximum distance between any two consecutive points along the line is no greater than that reasonably expected from random chance. In the case of the fixed spacing, the number of filled and empty cells are known and the likelihood of the observed maximum gap is a combinatorics problem. For the random case, the maximum allowed gap is the 20% cutoff for the first order statistic for a sequence of exponential waiting times with expectation based on the assumed false target rate.

If the candidate line meets all the above requirements, its posterior probabilities are compared with all those from other choices for endpoints, p_0 and p_1 . The best line is then accepted, and all targets suspected of being mines within the line are eliminated from further consideration. For a random line, all targets within the error bound around the line (mine lane) are marked as mines. For a fixed spacing line, only targets in cells are marked as mines. Thus the set of detections is slowly purged of targets likely to be members of mine lines.

Once a line is established some effort is made to insure that it completely spans its true mine line. This is done by attempting to extend the successful line in both directions. This is based strictly on the false target rate. Extending the line

from either end we consider the distance from the endpoint to the next target within range of the extended line. If the distance to that target is too small, that target is likely part of the line. Here, “too small” is defined as the 5% cutoff for the exponential waiting time to the next false target. The line is extended as long as points within range remain. These newly added points are then purged from the detection list, and the extended line is saved. Then the process begins again, repeating until no more lines pass the criteria.

B. Updating Priors Using Line Information

Given an area plot of the *mle*’s that result from a partial search, human observers can typically detect most of the mine lines. However, discerning mine spacing and the beginning and ending points of the mine line point is more difficult. The target detection analysis described in section III-A can detect mine lines not obvious to a human observer and accurately estimate their mine spacings. However, the real value of such an algorithm is for an automated system to evaluate the current state of a mine hunting mission and to use the collected information to speculate about mine locations in the unsearched regions. When no knowledge about mine locations is known, it is typical to assume that the mine spatial distribution is uniform. Although this approach minimizes the variance of the search performance when a uniform effort is applied, the search results are not optimal if the underlying mine spatial distribution is not uniform. More accurate knowledge about the distribution will improve search results by enabling search optimization strategies to plan a more efficient search of the unsearched regions.

The calculation to update the object structure priors begins by computing the observed mine density. This is done by dividing the total number of mines (*random* plus *fixed*) by the size of the area searched. This observed density is then applied in a uniform fashion across the entire unsearched areas and is identical to the typical uniform spatial distribution assumption used when no knowledge is available. However, it is reasonable to assume that the density estimate is more accurate than in the no-knowledge case. Later, if a *random* or *fixed* line extends into the unsearched areas, the *random* or *fixed* line density information overwrites the previous uniform density information.

Often target lines extend to the edge of the searched region. Clearly these lines likely extend into the unsearched area. Given the length of the line within the searched area, and using a model for the length of a mine line, we can predict the target density along the line in the unsearched region. For each line we test the distance between its endpoints and the unsearched regions. If this distance is less than the distance between targets allowed for a line extension, we allow the line to extend into the unsearched region. As the line extends further and further we have less confidence about its presence. To model this we assume that the length of a mine line is normally distributed with known mean and variance. Then at each point along the line we can calculate the probability that

the line extends that far given that its length is at least that seen in the searched area.

For randomly placed lines we simply extend the mine lane into the unsearched area. We assume that objects are placed uniformly within that region for each line length. We use the density of *mlc*'s (given by D) in the portion of the line in the searched area as an estimate of the entire line's density. Let $N(x)$ represent the cumulative Gaussian distribution with the appropriate mean and variance for mine line lengths evaluated at $x \in \mathbb{R}$. Then for a point at length y from the end of the original line whose length was x , we have

$$E\{\text{Density at } y\} = \frac{D(1 - N(x + y))}{1 - N(x)} \quad (6)$$

For fixed spacing lines, we calculate the projected positions of cells in the unsearched area. To account for placement error, we place a Gaussian error mask over the predicted location and scale by the line length factor as in the random case. The result is an array of target density predictions for each point in the unsearched area.

IV. EVALUATION OF THE SEARCH PLAN

The search evaluation capability discussed here employs a grid based design for partitioning the search space into an exclusive and exhaustive set of evaluation cells. Unlike typical approaches found in [5], [3] where region size is fairly large compared to sensor characteristics, we use a cell resolution that is large enough to preserve the mutual independence of detection events across cells yet small enough to articulate spatial variability in performance due to environmental influences and multiple pass operations of sensor platforms operating over the same search space. The approach has lead us to a multiple map representation of sensor coverage, contact likelihood and pertinent representation of placement structure. These details are discussed below.

A. Search Grid Specification

Let $S \subseteq \mathbb{R}^2$ represent a physical region in geographic coordinates over which the search is to be conducted. For simplicity of illustration, a square grid space is constructed with a rectangular *tessellation* and presented in Fig. 4. The geographic representation of S is shown on the left of the figure where an arbitrary search path is illustrated along with contours indicative of placement strategy. Probability calculations are performed over the grid cells centered about the points

$$G = \{ \langle x, y \rangle \in S \mid x = \pm(k - \frac{1}{2})\Delta X, \\ y = \pm(m - \frac{1}{2})\Delta Y, (k, m) \in \{n\}_1^{N_{max}} \} \quad (7)$$

defining the set of geographic coordinates $\mathbf{x} = \langle x, y \rangle$ in a square of side length $2N_{max}\Delta X$ centered about $\langle x, y \rangle = \langle 0, 0 \rangle$ with $\Delta X = \Delta Y$ to provide identical spacing in either coordinate direction. Hence, $4N_{max}^2$ grid points are used to construct the global representation of the search space S .




Fig. 4. Grid and Subgrid Depiction

Probabilities are updated sequentially at a time sampling occurring at regular interval over the search period

$$T = \{t_i \mid t_i = t_0 + i\Delta T, i \in \{n\}_{n=0}^{I_{max}}\} \quad (8)$$

starting at time t_0 . The sampling interval ΔT is constrained to be sufficiently small relative to vehicular dynamics such that the search trajectory over this interval is approximately linear. This serves to enable performance predictions for irregular search trajectories. However, the constraint of a small and uniform ΔT can be relaxed in alternate realizations without altering the underlying likelihood structure.

Actual calculations are performed over the much smaller local grid

$$L = \{\mathbf{x} \mid \mathbf{x} \in G, |x - x_s| \leq d_x, |y - y_s| \leq d_y\} \quad (9)$$

where coordinates $\mathbf{x}_s = \langle x_s, y_s \rangle$ denote sensor position within the grid at evaluation time t_i and displacements d_x and d_y define the extent of the extracted subset assimilating the advance associated with a temporal update interval, ΔT . This local grid is illustrated in Fig. 4 as the box surrounding the terminal point of the partial search path, with a close up view shown on the lower right illustrating a circular sensor detection region and a corresponding grid resolution applied in evaluation. The circle is of radius R_D and corresponds to a detection region employing a *definitive range law* as discussed in [1].

Local grid coordinates are translated and rotated by the transformation $\mathbf{w} = \mathbf{T}(\theta_h)(\mathbf{x} - \mathbf{x}_s)$ to set the sensor as the coordinate origin and align the y-axis with the search platform heading θ_h . This is depicted in the upper right of Fig. 4. We employ the circular detection region in part to maintain a generic (i.e., non-sensor specific) approach to calculation of the set coverage achieved by the search path and in part, to facilitate the extraction of set coverage difference as the sequential performance calculation unfolds.

Let D_{adv} represent the positional advance within an update interval. The advance region is illustrated in Fig. 4 as the slither encompassing the new set of grid points that are not

in the detection region of the previous update. Note that $D_{adv} \approx S_S \Delta T$ where S_S denotes the platform search speed. For any given search platform, search path evaluation consists of aggregating probabilities over the collection of evaluation regions defined by the respective search path. This is illustrated in Fig. 5 below.




Fig. 5. Search Path Coverage Calculation

We define the contact detection region C_i as the set of grid cells that support detection within the i^{th} interval,

$$C_i = \{\mathbf{w} \mid \|\mathbf{w}\| \leq R_D \vee |w_x| \leq R_D \vee -D_{adv} \leq w_y \leq 0\} \quad (10)$$

Let the location of each cell within the global grid structure G be denoted by its indices (k, l) and assume this mapping is retained in the local grid. Let $I_P = \{I_i\}_{i=0}^{I_{max}}$ denote the sequence of sets of grid indices corresponding to the set of updates depicted in Fig. 5. That is,

$$I_i = \{(k, l) \mid \mathbf{w} \in C_i / C_{i-1}\} \quad (11)$$

This important modeling detail provides implicit acknowledgement of the persistence of object placement and the image processing associated with side scan sonar and other sensors relevant to unmanned vehicular search. With this approach, independent detection events can only be attained by the search path circling back over itself (i.e., by design) or by distinct sensor platforms performing a coordinated search. The search path for a given platform is approximated by all the cells covered along the search trajectory; that is

$$G_{sp} = \bigcup_{i=0}^{I_{max}} \left[\bigcup_{(k,l) \in I_i} G_{k,l} \right] \quad (12)$$

Let outcomes $D = \{d, \bar{d}\}$ signify detection and non-detection events. Let the probability of a detection event occurring within a given grid cell be denoted by the likelihood function $P_{D_{k,l}}$. To develop this probability, assume there exists a placement probability density function (pdf) $f(\mathbf{x})$ with support on S and a traditional detection likelihood function in terms of source-receiver geometry. Then the joint probability that an object is located within the cell and is detected within a given search pass becomes

$$\Pr(d, \mathbf{x} \in G_{k,l}; \mathbf{x}_s) = \int_{G_{k,l}} \Pr(d|\mathbf{x}; \mathbf{x}_s) \cdot f(\mathbf{x}) d\mathbf{x} \quad (13)$$

From this, the grid cell conditional detection likelihood can be developed as

$$\begin{aligned} \Pr(d|G_{k,l}; \mathbf{x}_s) &= \frac{\Pr(d, \mathbf{x} \in G_{k,l}; \mathbf{x}_s)}{\Pr(\mathbf{x} \in G_{k,l})} \\ &= \frac{\int_{G_{k,l}} \Pr(d|\mathbf{x}; \mathbf{x}_s) \cdot f(\mathbf{x}) d\mathbf{x}}{\int_{G_{k,l}} f(\mathbf{x}) d\mathbf{x}} \\ &= \int_{G_{k,l}} \Pr(d|\mathbf{x}; \mathbf{x}_s) \cdot f(\mathbf{x}|\mathbf{x} \in G_{k,l}) d\mathbf{x} \\ &\approx P_{D_{k,l}} \end{aligned} \quad (14)$$

The resulting value denotes a weighted spatial average of the detection likelihood function over the grid cell [6]. We shall assume that grid cell size is sufficiently small such that the variation in detection and placement likelihood over the grid cell is also small allowing a nominal constant value to be presumed over the cell (and permitting us to drop the notation in $\mathbf{x}_s$).

Observe that for $P_{D_{k,l}} < 1$, there is a probability of $(1 - P_{D_{k,l}})$ that detection will not occur on the first search opportunity. It may be detected on subsequent passes, however. Subject to the set difference restrictions discussed above, we model each search pass at a given position $G_{k,l}$ as an independent Bernoulli trial. Then, the waiting time $W_{k,l}$ for detection to occur within the cell follows a *geometric distribution* [7, p. 61]. That is, the probability that the first detection will occur in cell (k, l) on the n^{th} pass becomes

$$\Pr(W_{k,l} = n) = P_{D_{k,l}} \cdot (1 - P_{D_{k,l}})^{n-1} \quad (15)$$

B. Event Map Synthesis

A number of likelihood maps are synthesized on both global and local scales to facilitate probability calculation. Exactly how they are operated on depends upon the event space constructed and the performance criteria utilized. These will vary according to the search problem being addressed. We do, however, draw distinction between detection likelihood and placement models as indicated in Fig. 6. This allows us to draw inference on search effort independent of the placement density applied when appropriate.




Fig. 6. Event Map Dichotomy

A detection probability map, $M_{P_D} = \{P_{D_{k,l}}\}$, is constructed over G to encapsulate any spatial variations in detection likelihood due to the environment. Such environmental influences on detection performance include water turbidity and bottom depth and composition. Should sensor referenced dependencies need be incorporated, these variations can be constructed over L so to modify detection likelihood according to sensor proximity.

The non-detection likelihood map, $M_{\bar{P}_D}$, is also constructed over G . Whereas the M_{P_D} map is static, the $M_{\bar{P}_D}$ map is maintained sequentially with every update. Let $\delta_{k,l}(n)$ denote the first detection within grid cell (k,l) occurring over the n^{th} pass and the event $\bar{\delta}_{k,l}(n)$ denote the occurrence of no detections through the first n passes. With no previous search conducted, these probabilities are initialized at $P_{\delta_{k,l}}(0) = 0$ and $P_{\bar{\delta}_{k,l}}(0) = 1$. The multiple-pass detection recursion equations for $n = 1, \dots, N_S$ passes are given by

$$P_{\delta_{k,l}}(n) = P_{D_{k,l}}(n) \cdot P_{\bar{\delta}_{k,l}}(n-1) \quad (16a)$$

$$P_{\bar{\delta}_{k,l}}(n) = (1 - P_{D_{k,l}}(n)) \cdot P_{\bar{\delta}_{k,l}}(n-1) \quad (16b)$$

where differing sensor types can cause the $P_{D_{k,l}}$ to vary from pass to pass. It is the result of the equation (16b) recursion that is stored in the non-detection map, $M_{\bar{P}_D} = \{P_{\bar{\delta}_{k,l}}(n)\}$.

Object placement likelihood is stored in the static, globally referenced map, $M_{P_G} = \{P_{G_{k,l}}\}$. In the simplest event space construct, M_{P_G} represents a pdf over $\mathcal{S}$ for a single placement that is known to occur. For this case, the probability of successful search along the specified search path is merely the aggregated probability

$$\begin{aligned} \Pr(d | G_{sp}) &= \int_{G_{sp}} \Pr(d|\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \\ &\approx \sum_{i=0}^{I_{max}} \Delta P_i = p_{sp} \end{aligned} \quad (17)$$

$$\Delta P_i = \sum_{(k,l) \in I_i} P_{D_{k,l}} P_{\bar{\delta}_{k,l}} P_{G_{k,l}} \quad (18)$$

Here the detection recursion of equation (16a) is embedded directly in the evaluation region calculation of ΔP_i .

If N_M distinct objects are known to exist in $\mathcal{S}$ and placement within $\mathcal{S}$ is i.i.d., then from [6] the event of finding k of them along the search path on a binomial distribution

$$\Pr(E_k | G_{sp}) = \binom{N_M}{k} p_{sp}^k (1 - p_{sp})^{N_M - k} \quad (19)$$

where E_k denotes the event that $N(G_{sp}) = k$. The clearance event, E_C , where all mines are found along the search path, $N(G_{sp}) = N_M$, has probability

$$\Pr(E_C | G_{sp}) = p_{sp}^{N_M}. \quad (20)$$

The probability of finding at least one object along the path becomes

$$\Pr(E_D | G_{sp}) = 1 - (1 - p_{sp})^{N_M} \quad (21)$$

where E_D denotes an ‘‘assured detection’’ event, $N(G_{sp}) > 0$. Note that with all placements i.i.d., the calculation of search

effectiveness probability is reduced from an operation on the $(\mathcal{S} \times D)^{N_M}$ outcome space to an apparent operation on $\mathcal{S} \times D$ through the functional composition; in this case, raising operations on M_{P_G} map to an exponent.

When the number of objects placed in $\mathcal{S}$ is unknown, one is left to evaluate search effort alone or to make assumptions on placement based on intelligence estimates. Spatial Poisson processes (SPP) provide a good alternative when the expected number of objects in $\mathcal{S}$, $E\{N_M\}$, is given. The model operates on an intensity function, $\lambda(\mathbf{x})$, which need not be homogeneous over $\mathcal{S}$ [7]. The *Coloring Theorem* allows a separation of placement events into independent detection and non-detection processes with $\lambda(\mathbf{x}) = \gamma_d(\mathbf{x})\lambda(\mathbf{x}) + \gamma_{\bar{d}}(\mathbf{x})\lambda(\mathbf{x})$, subject to $\gamma_{\bar{d}}(\mathbf{x}) = 1 - \gamma_d(\mathbf{x})$. Here, $\gamma_d(\mathbf{x})$ plays the role of the detection likelihood function. The SPP does presume that object placement is i.i.d. and this allows the intensity function to be decomposed into two parts, $\lambda(\mathbf{x}) = E\{N_M\} \cdot f(\mathbf{x})$, where $\int_{\mathcal{S}} f(\mathbf{x}) d\mathbf{x} = 1$ and $f(\mathbf{x})$ serves as the placement pdf. Using this model, the probability of finding k objects along the search path becomes

$$\Pr(N(G_{sp}) = k) = \frac{\Lambda(G_{sp})^k}{k!} e^{-\Lambda(G_{sp})} \quad (22)$$

where

$$\begin{aligned} \Lambda(G_{sp}) &= \int_{G_{sp}} E\{N_M\} \gamma_d(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \\ &\approx E\{N_M\} \cdot \sum_{i=0}^{I_{max}} \Delta P_i. \end{aligned} \quad (23)$$

Note that this same result could have been obtained by applying the Poisson approximation to equation (19). The next detection probability analogous to equation (21) becomes

$$\begin{aligned} \Pr(N(G_{sp}) > 0) &= 1 - e^{-E\{N_M\} \cdot \sum_{i=0}^{I_{max}} \Delta P_i} \\ &= 1 - \prod_{i=0}^{I_{max}} e^{-E\{N_M\} \Delta P_i}. \end{aligned} \quad (24)$$

Next, we introduce field maps indicating structure in placement strategy. We define M_{FR} as a configuration map specifying the intended location of each placement relative to a common reference point. This reference points to a location in G , $\mathbf{x}_{ref} = \langle x_{ref}, y_{ref} \rangle$, but in the local map with field coordinates $\mathbf{x}_m = \langle x_m, y_m \rangle$, the reference point serves as the coordinate origin. The reference placement map, M_{DR} depicts a likelihood function indicating plausible placement locations given M_{FR} and a suitable drop uncertainty region associated with each placement. Examples of these map are provided in Fig. 7.

When the reference point is known, the M_{DR} map is readily superimposed over G for search path evaluation. However, the probability calculations change some. Assume that the map M_{DR} is comprised of N_M distinct placements and that the field is known to be totally populated and exists in $\mathcal{S}$ at the given reference point. Assume element placements are conditionally independent and let $\{f_j(\mathbf{x})\}_{j=1}^{N_M}$ denote the respective placement pdf's. Assume these regions, $R_j(\mathbf{x})$ are




Fig. 7. Field Structure Definition Maps

disjoint in M_{DR} and hence, S . Let $f_{j_{k,l}}$ represent their respective partitioning when mapped to G and assume that the same restrictions on grid resolution hold. Let $E_k = \{H_i\}$ designate the set of hypotheses on the N_M placements such that k detections occur, $J_{D_i} \subseteq \{j\}_{j=1}^{N_M}$ designate which ones are detected and $\bar{J}_{D_i}$ designate it's complement. Then, analogous to equation (19), the probability of event E_k occurring becomes

$$\Pr(E_k | G_{sp}) = \sum_{H_i \in E_k} \left[\prod_{j \in J_{D_i}} p_{sp}^j \prod_{j \in \bar{J}_{D_i}} (1 - p_{sp})^j \right] \quad (25)$$

where the component probabilities are now element specific,

$$\begin{aligned} p_{sp} &= \int_{G_{sp}} \Pr(d|\mathbf{x}) f_j(\mathbf{x}) d\mathbf{x} \\ &\approx \sum_{i=0}^{I_{max}} \sum_{(k,l) \in I_i} P_{D_{k,l}} P_{\delta_{k,l}} f_{j_{k,l}} \end{aligned} \quad (26)$$

For the clearance probability

$$\Pr(E_C | G_{sp}) = \prod_{j=1}^{N_M} p_{sp}^j \quad (27)$$

to be non-zero, all elements must be visited along the search path. Similarly, the next detection probability is given by

$$\Pr(E_D | G_{sp}) = 1 - \prod_{j=1}^{N_M} (1 - p_{sp}^j). \quad (28)$$

For disjoint regions $\{R_j(\mathbf{x})\}$, equations (27) and (28) can be calculated directly from M_{PD} and P_G at any given update t_i via region indexing

$$p_{sp}^j = \sum_{(k,l) \in R_j} (1 - P_{\delta_{k,l}}) P_{G_{k,l}} \quad (29)$$

where $P_{G_{k,l}}$ is synthesized as the aggregate of each projection $f_{j_{k,l}}$ on S and the $P_{\delta_{k,l}}$ applied is the post-recursion update.

More generally, the structure may not be fully populated and its reference location in S or even how many of such structures there are in S may not be known. One alternative to address modeling uncertainty is to represent the structure as an SPP. Here, the aggregate projection represented in the drop reference map, M_{DR} , becomes applicable for each placement and separation between individual placements is no longer




Fig. 8. Lines and Detections

ensured. However, under certain conditions of large N_M and incomplete population, the approximation is useful.

For the SPP, the M_{DR} map depicts the intensity function

$$\lambda(\mathbf{x}) = E\{N_M\} \cdot f(\mathbf{x}) = \frac{E\{N_M\}}{N_M} \sum_{j=1}^{N_M} f_j(\mathbf{x}) \quad (30)$$

where $E\{N_M\}$ represents the anticipated population level. For known reference point, the mapping to G is direct, $\mathbf{x} = \mathbf{x}_{ref} + \mathbf{x}_m$. For unknown reference location, we construct M_{DP} to represent field placement likelihood and perform a convolution over the two map to yield P_G . Then, calculation is as in equations (23) and (22).

V. VEHICLE COLLABORATION EXPERIMENTAL RESULTS

An example experiment is conducted to demonstrate the information processing capabilities of the search and optimization framework. The search is conducted in multiple stages. In the first stage, one quarter of the space S is searched and the contact data is processed. The information derived from the first search component is applied to extract estimates of placement structure and to apply them so to develop the expected placement of objects in the following search quadrant. Fig. 8 depicts as green stars the valid contact reports that occur over the entire search space S . These are comprised of the subset of uniform random and structured line placements that have been detected. The actual line structures follow generic mine lay down strategies [8] and are indicated in the figure.

The experiment consists of partitioning the search space into four quadrants, processing the contact data in the first quadrant and applying the results as priors for the remaining quadrants. The partitioning is illustrated in Fig. 9. Line placement hypotheses developed for quadrant I are extended according to the estimated parametric model developed for the respective lines.




Fig. 9. Search Quadrant Partitioning

In this example, the algorithm detected 6 out of 7 lines present. The seven lines were composed of 96 mines, 92 of which were correctly flagged as mines. The missed line was a fixed-spaced line that had too few targets within the first quadrant to be detected. No false targets were incorrectly associated with any lines.

To improve the future portions of search based on this information, we have developed a method to extract parameterizations of search object patterns from observed contacts in environments characterized by heavy clutter. After searching quadrant I, a set of mine line hypotheses was developed. These hypotheses were mapped (along with a random hypothesis) to an expected contact density in the remaining quadrants as described in section III. This algorithm provides an update to the priors that were utilized in planning the remainder of the search. Thus, the approach provides a level of adaptation for multi-vehicle searches that can be automatically employed for autonomous vehicles.

The magenta +’s in Fig. 9 mark the hypothesized mine locations from the detected fixed-spaced lines. Note that the algorithm has incorrectly identified the fixed-spaced line in the top of quadrant I as two separate line; this creates a short detected mine line which in turn causes the hypothetical extension to proceed too far into quadrant IV. The solid magenta areas in quadrants II and III represent the hypothesized mine extension for random lines. The extension in quadrant III is reasonably close the actual mine line while the extension in Quadrant II exceeds the true line’s length.

The search evaluation algorithm described in section IV is applied to each of the three remaining quadrants; this being indicative of search being conducted by three distinct unmanned vehicles. For each quadrant search, the following grid and update specification applies:

Search Box: 9260 m per side
Resolution: 5 m

Sweep time: 8 hrs 43 min
Update rate: 10 sec

The search path for each vehicle follows a typical uniform coverage ladder path (i.e., mow the lawn).

The three search quadrants have differing characteristics due to asymmetric portioning of anticipated line structure. In particular, the mine threat density maps generated for the quadrants vary in the relative contribution of random and structured line placements. These are indicated in Table I.

TABLE I
SEARCH QUADRANT BREAKDOWN

	Quadrant II	Quadrant III	Quadrant IV
True Number of Mines	66	48	0
Number of <i>mlc</i> ’s	107	87	81
Density Total	221	109	95

Results were compiled for each of the quadrant search efforts. Shown in Fig. 10 are the respective probabilities for finding $k > 0, 50, 100$, and 150 contacts developed over time as the search path is executed, as given by equation (22). Fig. 11 depicts the expected number of contacts, the actual number observed, and the subset of them that were actually mines placed within the line structure. For this experiment, the detection results produced for the original experiment were maintained. Hence, the numbers presented in Table I apply.




Fig. 10. Detection Probability for $k > 0$ (black), $k > 50$ (blue), $k > 100$ (red), and $k > 150$ (green) Contacts

The figures illustrate the impact that the spatially variable density maps have on the temporal probability assessments. As no attempt to optimize search path was made, the probability evolution takes on a stepped appearance indicative of crossing hypothesized line regions. This is especially true for

quadrant II results which have the highest concentration of lines. The number of the reported contacts depicted in Fig. 11 relative to the expected number highlights both the distinction between the two types of line hypotheses and the difficulty in establishing the priors. As a single experiment, the results are consistent with the predictions to the degree that they exhibit the anticipated trends in the data.




Fig. 11. Expected Detected Contacts (black), Actual Contacts (blue), and Actual Mineline-Related Contacts (red)

VI. CONCLUSION

We have presented a methodology and mathematical modeling approach to planning searches for multiple unmanned vehicles. The developed constructs are applicable to vehicles searching for objects that occur in uncertain structured patterns on the ocean floor. As such, the procedures developed readily accommodate both to planning search paths based on area coverage asset allocations as well as to coordinated group searches.

This approach was applied to a problem of multiple unmanned vehicles searching for mines that are likely to occur in linear patterns. In a simulation setting, we illustrated how the reported data from a first vehicle's search of one quadrant of the region can be used to infer prior parameters for the evaluation of search in the other regions. In this manner, we have illustrated a capability to interpret spatial variability in placement.

The example presented in this paper represents a single realization of a general toolset for evaluating search in the context of placement uncertainty. Future extensions of this work include the incorporation of spatially variable random placement strategies, the examination of nonlinear patterns, the use of nonhomogeneous sensing capabilities, and the use of these methods in a search plan optimization setting.

ACKNOWLEDGMENT

This work was sponsored by the Office of Naval Research, Code 321 Maritime Sensing.

REFERENCES

- [1] B. H. Koopman, *Search and Screening: General Principles with Historical Applications*, Pergamon Press, Elmsford, NY, 1980.
- [2] J. C. Hyland, C. M. Smith, and P. J. Stampnick, "Multi-dimensional MCM Tactical Analysis and Optimal Resource Allocation", CSS/TR-95/13, Naval Surface Warfare Center, Panama City, FL, February, 1993.
- [3] L. D. Stone, *Theory of Optimal Search*, 2nd Edition, ORSA Books, Arlington, VA, 1989.
- [4] W. Feller, *An Introduction to Probability and its Applications, Volume I*, 3rd Edition, John Wiley & Sons, Inc., New York, NY, 1968.
- [5] A. R. Washburn, *Search and Detection*, 4th Edition, INFORMS, Linthicum, MD, 2002.
- [6] J. G. Baylog and T. A. Wettergren, "Search Path Evaluation Incorporating Object Placement Structure", NUWC-NPT Technical Report 11,845, Naval Undersea Warfare Center Division, Newport, RI, 20 December 2007.
- [7] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*, Oxford University Press, Inc., New York, 2001.
- [8] Committee for Mine Warfare Assessment, Naval Studies Board Division on Engineering and Physical Sciences, National Research Council, *Naval Mine Warfare: Operational and Technical Challenges for Naval Forces*, National Academy Press, Washington, DC, August, 2001.