2014

Modeling and Analysis of Exhaustive Probabilistic Search

Chung, Timothy H.; Silvestrini, Rachel T.

http://hdl.handle.net/10945/46137

Modeling and Analysis of Exhaustive Probabilistic Search

Timothy H. Chung,[1] Rachel T. Silvestrini[2]

[1] *Department of Systems Engineering, Naval Postgraduate School, Monterey, California 93943*

[2] *Department of Operations Research, Naval Postgraduate School, Monterey, California 93943*

Received 14 November 2012; revised 14 January 2014; accepted 16 January 2014
DOI 10.1002/nav.21574
Published online 14 February 2014 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: This article explores a probabilistic formulation for exhaustive search of a bounded area by a single searcher for a single static target. The searcher maintains an aggregate belief of the target's presence or absence in the search area, concluding with a positive or negative search decision on crossing of decision thresholds. The measure of search performance is defined as the expected time until a search decision is made as well as the probability of the search decision being correct. The searcher gathers observations using an imperfect detector, that is, one with false positive and negative errors, and integrates them in an iterative Bayesian manner. Analytic expressions for the Bayesian update recursion of the aggregate belief are given, with theoretical results describing the role of positive and negative detections, as well as sensitivity results for the effect of the detection errors on the aggregate belief evolution. Statistical studies via design of simulation experiments provide insights into the significant search parameters, including imperfect sensor characteristics, initial belief value, search decision threshold values, and the available prior probability information. Regression analysis yields statistical models to provide prescriptive guidance on the search performance as a function of these search parameters. Naval Research Logistics 61: 164–178, 2014

Keywords: probabilistic search theory; autonomous systems; expected time to decision; design of experiments; generalized linear models

1. INTRODUCTION

Unmanned systems play an increasingly prominent role in diverse information gathering missions, with examples in civilian and defense contexts including search and rescue (SAR) and intelligence, surveillance, and reconnaissance (ISR) operations. In many cases, these new technologies for autonomy have offered revolutionary advances in areas such as sensor data fusion, optimized search planning, and distributed resource allocation. Much of the current active research consider applications that assume or require the use of sophisticated sensing, computation, coordination, and communication capabilities. However, there is still a measurable gap between the employment of these research-oriented autonomous systems and actively fielded unmanned systems, that is, those that are readily used in practice, e.g., by actual SAR operators. One reason for this unrealized potential may be due to the absence of transitional models and corresponding solutions that resemble legacy, human-derived, approaches while providing a clear view of the advantage afforded by increased autonomy. One such legacy approach is the use of exhaustive search patterns, such as sweeping or lawnmower trajectories, in many practical search contexts. The complete coverage offered by these patterns ensures that the searcher visits all locations within the designated search area and allows for analysis of search performance, e.g., expected time until detection under mild assumptions on the searcher's characteristics.

This work seeks to revisit the use of exhaustive search strategies as the basis of the search process, and leverage new probability models and statistical analysis techniques to help inform or refine concepts of operations. Not only is investigation of the numerous variables involved in searching for a stationary target possibly located within a bounded area possible but also the resulting ideal variable values can be used to guide the design and selection of an appropriate unmanned platform, sensor payload, and search parameters. The main contributions of this article are the analytic formulations and closed-form expressions describing the probabilistic exhaustive search process, and the statistical regression analysis providing a characterization of the search decision effectiveness as a function of multiple significant variables. This work focuses on studying the "evolution" of the search process, rather than on its optimization, so as to accentuate

Correspondence to: T.H. Chung (thchung@nps.edu)

the value of the probabilistic models to both search theorists and practitioners.

The organization of the article is as follows. Section 2 offers an overview of related work in the areas of exhaustive search, ranging from robotic coverage to probabilistic methods and applications. The exhaustive probabilistic search formulation is presented in Section 3, which describes the evolution of the search decision process in determining whether or not a stationary target is present in the search area. Analytic investigations of the search decision process including study of several special cases are described in Section 4. Section 5 describes the inputs and outputs of the simulation experiments, including the statistical methodology used to develop a model relating these two elements. Following Section 6 is a discussion of the statistical regression results, concluding with a summary of this articles main results and avenues for future research.

2. PREVIOUS WORK

Search theory and operations analysis have their origins in the use of applied mathematical models to aid in the hunt for submarines in World War II, though applications for physical search range from counter-mine operations to SAR of downed aircraft to looking for misplaced car keys. The seminal reports by [13, 14] outline various models governing the detection of targets, probability-based maps of possible target locations, and planning and allocation of search resources. In these and related works, exhaustive search patterns are proposed to offer guidance to commanders on how to employ their search assets. The notion of optimizing search in these early approaches occurs within the exhaustive search context, such as seeking to maximize the probability of detection of a target by optimizing the spacing between parallel tracks in the back-and-forth search pattern [23].

Similarly, one of the early applications of mobile robotic systems was the coverage problem, that is, to develop algorithmic and/or theoretical methods for enabling robots to completely cover an area. Given the breadth of relevant applications, the field of robotic coverage research is extensive and [5] provides a survey of many of these methods. To distinguish probabilistic search of an area with the coverage problem, search implies the use of probability models for the target's distribution as well as sensor models to characterize the detection of the target. Additionally, whereas the coverage task is concluded if or when the area has been completely covered, the search termination criterion includes measures such as attaining a minimum level of confidence, e.g., a sufficient probability of target detection. In this manner, classical coverage problems can be related to the physical search problem under the assumptions of uniform target distribution and perfect, finite-range detections. The work presented in this article can leverage the exhaustive or covering patterns developed by the robotic coverage community while incorporating methods for probabilistic search of a target.

More recently, there has been an abundance of theoretical operations research (see [19, 20], or [2] for a survey) and sensor-based robotics (e.g., [3, 16]) in the optimization of search trajectories; however, these computationally complex solutions have often gone unused in practice despite the potential improvements in the search performance they offer. Given the variability present in different search missions (e.g., weather, terrain), thorough studies of the significant variables and sensitivity to these search parameters, such as sensor characteristics, are necessary to help provide search practitioners with a further understanding of the underlying principles of search models before such sophisticated methods can be applied in the field.

In contrast, exhaustive search patterns, such as lawnmower and sweeping paths, offer the advantage of an extensive empirical basis as well as intuitive theoretical results. Standard procedures for executing land or maritime SAR operations describe the various conditions for using different types of exhaustive search patterns and provide tabulated "rules of thumb" parameters for guiding the search [8, 21]. More recently, [24] examine exhaustive search using Bayesian updating in the presence of ancillary dependencies, that is, detections depend on conditions such as the searcher's direction of travel. [15] investigate a threshold-based search decision process, similar to that presented in this article, but for search of individual, independent cells. As such, their treatment does not consider the dependent nature of many search contexts, that is, search for a specific single target (e.g., a downed aircraft) in one location should necessarily affect the belief of its presence elsewhere. [6] provide a probabilistic framework to define the evolution of the search decision of determining the presence or absence of a target, but focus on comparison of search strategies for improving search performance, rather than exploring its relationship and sensitivity to significant search parameters. This article presents both a formal analysis and statistical verification of the probabilistic search models specifically in the context of exhaustive search, extending the results of [6] so as to help provide a bridge between current practices in field operations and those offered in current research.

3. PROBABILISTIC MODEL FORMULATION

In search problems, as the searcher moves through a search region, it gathers information regarding the presence or absence of the target within the searched location. The aggregation of these positive or negative detections results in a search decision, where the searcher ultimately determines that the target is present at a specific location or is altogether

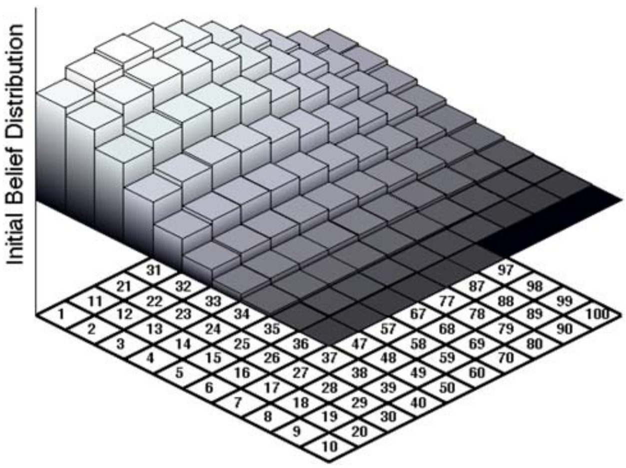

Figure 1. Illustration of a discretized environment and a model for a initial probability distribution or map, $\mathcal{M}(0)$, where the height and shading of each bar represent the belief that the target is present in that cell. The top-to-bottom, left-to-right convention for labeling cells in the grid is also shown. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

absent from the region. The probabilistic model formulation is based on this concept and discussed in this section. First the search environment and search patterns will be discussed, followed by a discussion on the observation sequence.

3.1. Search Environment and Exhaustive Search

Consider a search environment that is discretized into cells. Then, the probability that the target is present in the search area is given by the combined probability that it is present in any single cell. Let H be a Bernoulli random variable describing the state of the target's presence or absence in the search area. Then, the following expression (consistent with [6])

$$Pr[H = 1] = \sum_{c=1}^{C} Pr\,[X_c = 1] \tag{1}$$

is termed the "aggregate belief" that the target is present in the search area. The Bernoulli random variable, X_c, represents the probabilistic presence of the target in a specific cell c, such that $Pr\,[X_c = 1]$ is termed the "cell belief". The manner in which cell beliefs are dependent can be seen in (1), i.e., an increase in the cell belief in one location must decrease that of other cells. The additional task for the searcher is to localize the target to a specific cell, given it is indeed present. Figure 1 shows an example of a grid search area with $C = 100$ cells, and an illustrative initial target probability map, denoted $\mathcal{M}(0)$, where each location's height is the cell belief.

Note that the partitioning of the environment need not be spatially uniform, that is, cells can be differently sized, subject to environmental (e.g., terrain or building layout) or sensor-related (e.g., size of sensor footprint). Further, the partitioning could also incorporate nonuniform prior probabilities by constructing variably sized cells containing equal probability. The presented formulation can be applied to any exact decomposition of the search area, and the uniform grid is presented without loss of generality.

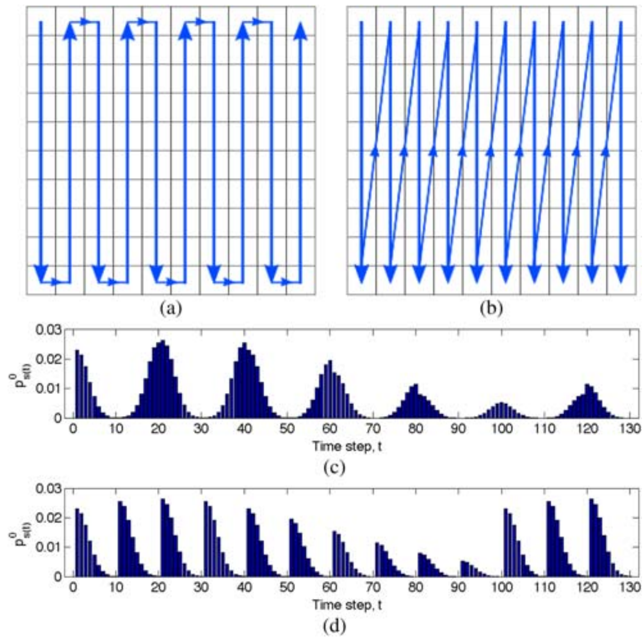

Figure 2. Illustrations of (a) lawnmower search and (b) sweeping search patterns for the example 10×10 grid. For lawnmower search, the searcher weaves back-and-forth in an alternating but contiguous path. Upon reaching the end of a complete pass (e.g., cell 90), the searcher reverses course by exactly retracing its previously traveled path. For sweeping search, the searcher executes adjacent top-to-bottom paths until reaching the end of its complete pass (at cell 100), and resumes the search again at the starting location, cell 1. The sequence of cells and their corresponding initial cell beliefs are shown for (c) lawnmower and (d) sweeping search trajectories using the initial target distribution shown in Fig. 1. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

The searcher inspects an individual cell, $s(t) \in \{1, \ldots, C\}$, in each discrete time step t, such that the searcher's trajectory from time step 1 through time step t is defined by the sequence $\mathcal{S}_t = \{s(1), \ldots, s(t)\}$. Two trajectories representing exhaustive search are considered in this article, namely the "lawnmower search" and "sweeping search" patterns. The former is also referred to as "ladder search" (in SAR contexts) [21] or a "Boustrouphedon" path in the robotic coverage literature [5], and is typically used in many real world contexts. Given the labeling convention and the known search area decomposition, the lawnmower trajectory can be predetermined and constructed prior to the start of the search process. Figure 2(a) illustrates this pattern, which is repeated in reverse on reaching the final cell in the pattern.

Note that the searcher requires a single pass to exhaustively search the area but twice the total number of cells in order to complete a repeating pattern (as the searcher must retrace its steps).

Alternatively, the sweeping search pattern offers a simpler searcher trajectory, though may not always be feasible or desirable due to the constrained motion of the searcher, while also ensuring exhaustive search of the environment. Detailed analysis of this search pattern and generated insights represent a major contribution of this work. Illustrated in Fig. 2b, the sweeping search trajectory can be constructed "a priori", with the repeating sequence of $\{1, \ldots, C\}$, restarting at the beginning of this sequence upon reaching its end. Rather than offline planning, the sweeping search trajectory can be constructed iteratively, starting at $s(1) = 1$, according to the following expression:

$$s(t+1) = \begin{cases} \text{mod}\ (t, C), & \text{mod}\ (t, C) \neq 0 \\ C, & \text{mod}\ (t, C) = 0 \end{cases}$$

where mod represents the modulus function. While sweeping may pose inefficiencies for a ground-based mobile robot, as it must transit to the top of the next column after reaching the bottom, the sweeping search pattern still plays a role in a variety of other search contexts and searcher motion capabilities, such as raster scanning of an environment with a tower-mounted pan-tilt camera, or periodic swaths taken by orbiting, satellite-based sensors, or a spiraling path of ascending flight around a mountain. The sweeping search pattern also facilitates analytic study, as described further in Section 4.

The mathematical framework governing the evolution of the belief and subsequent analysis do not depend on which search strategy, using either a lawnmower or sweeping search pattern, is selected. In fact, the proposed framework applies to any exhaustive search pattern that, like the lawnmower or sweeping search patterns, ensures that a cell in the search area is revisited only after all other cells have each been visited once. This fact can be observed by noting that, for a given initial prior distribution (e.g., Fig. 1), the two strategies can be related via a straightforward reordering of the cells. The order of cell probabilities of the visited cells is shown in Figs. 2c and 2d.

3.2. Observation Process

The searcher is endowed with an imperfect detector, which returns an observation $Y_{s(t)}$ in cell $s(t)$ at time step t. This observation can reflect a positive or negative detection, that is, $Y_{s(t)} = 1$ or $Y_{s(t)} = 0$, respectively. However, since the observations may be erroneous, false positive and false negative detections are possible and characterized by their respective error probabilities. Given the probabilistic observations, imperfect sensing is captured by likelihood functions, that is, the conditional probability of a particular observation in a cell given the true state of the cell. In other words, the detection error probabilities are given by:

$$\alpha_{s(t),t} = Pr\left[Y_{s(t)} = 1 \mid X_{s(t)} = 0\right]$$
$$\beta_{s(t),t} = Pr\left[Y_{s(t)} = 0 \mid X_{s(t)} = 1\right]$$

where $\alpha_{s(t),t}$ and $\beta_{s(t),t}$ are cell-dependent and also time-dependent false positive and false negative probabilities. These detection error probabilities depend on the selected sensor, such as an electro-optical camera often used for aerial search, and its employment, i.e., nadir pointing at fixed altitude. Nominal theoretical models for detection can be found in the search theory literature [22,23], such as the inverse cube law for visual detections, or alternatively empirical results exist, e.g., [21], in lookup tables from which the operator can select appropriate values depending on the conditions of the search mission. For succinctness in this article, we consider constant detection error probabilities and drop the subscripts. Also, it can easily be seen that $1 - \alpha - \beta \neq 0$ must be true, as otherwise, there is no benefit to taking observations in updating the prior probability. Further, without any loss of generality, we posit that $1 - \alpha - \beta > 0$, since otherwise we could reverse the detection cue (e.g., define $\bar{Y} = 1 - Y$) to arrive at the desired inequality.

3.3. Bayesian Update Expressions

Consider the probabilistic formulation presented in [6], which governs the evolution of the aggregate belief that a target is present within the search environment as a function of the gathered observations and the searcher's trajectory through the region. Each observation at time step t updates all cell beliefs (and hence the aggregate belief, $B(t)$, and the target probability map, $\mathcal{M}(t)$), as cells are assumed dependent via knowledge that only a single target may be present.

$$p_c^t \triangleq Pr\left[X_c = 1 \mid \mathcal{Y}_t\right] = \frac{Pr\left[Y_{s(t)} \mid X_c = 1\right]\ p_c^{t-1}}{Pr\left[Y_{s(t)} \mid X_{s(t)} = 1\right] p_{s(t)}^{t-1} + Pr\left[Y_{s(t)} \mid X_{s(t)} \neq 1\right]\left(1 - p_{s(t)}^{t-1}\right)}. \quad (2)$$

Define the following functions of random variables, denoted the "specificity" and "sensitivity" functions, respectively, to aid in the subsequent analysis:

$$\Phi\left(Y_{s(t)}\right) \triangleq (1-\alpha)\left(1 - Y_{s(t)}\right) + \alpha\ Y_{s(t)}$$
$$\Psi\left(Y_{s(t)}\right) \triangleq \beta\left(1 - Y_{s(t)}\right) + (1-\beta)\ Y_{s(t)}$$

and

$$\Theta_c\left(Y_{s(t)}\right) \triangleq \begin{cases} \Psi\left(Y_{s(t)}\right) & \text{if } c = s(t) \\ \Phi\left(Y_{s(t)}\right) & \text{if } c \neq s(t) \end{cases}$$

$$\Omega\left(Y_{s(t)}\right) \triangleq \left(\Psi\left(Y_{s(t)}\right) - \Phi\left(Y_{s(t)}\right)\right) = \left(2Y_{s(t)} - 1\right)(1 - \alpha - \beta)$$

Then, the above single time step recursion (2) can be expressed succinctly as

$$p_c^t = \frac{\Theta_c\left(Y_{s(t)}\right)\ p_c^{t-1}}{\Phi\left(Y_{s(t)}\right) + \Omega\left(Y_{s(t)}\right)\ p_{s(t)}^{t-1}}. \tag{3}$$

We can evaluate this recursion over multiple time steps to find a closed-form expression for a cell belief value, given the searcher's trajectory and observation histories, $\mathcal{S}_t$ and $\mathcal{Y}_t$, through time step t. Iteration and collection of terms yields:

$$p_c^t = \frac{\prod_{k=1}^{t} \Theta_c\left(Y_{s(k)}\right)\ p_c^0}{\prod_{k=1}^{t} \Phi\left(Y_{s(k)}\right) + \sum_{k=1}^{t} f\left(k, t, Y_{s(k)}\right)}, \tag{4}$$

where $f\left(k, t, Y_{s(k)}\right)$ is given by the summand

$$f\left(k, t, Y_{s(k)}\right) \triangleq \left(\prod_{l=1}^{k-1} \Theta_{s(k)}\left(Y_{s(l)}\right)\right) \Omega\left(Y_{s(k)}\right) \left(\prod_{m=k+1}^{t} \Phi\left(Y_{s(m)}\right)\right) p_{s(k)}^0$$

For the "outside" cell ($c = 0$), $\Theta_0\left(Y_{s(t)}\right) = \Phi\left(Y_{s(t)}\right)$, as that the searcher can neither observe nor visit this location. Observing that the aggregate belief (i.e., the probability that the target is in the search area) is related to this null belief (the probability that the target is not in the search area) by $p_0^t = 1 - B(t)$, we arrive at another means of examining the evolution of the aggregate belief:

$$B(t) = 1 - p_0^t = \frac{\left(1 - p_0^{t-1}\right)\ \Phi\left(Y_{s(t)}\right) + \Omega\left(Y_{s(t)}\right)\ p_{s(t)}^{t-1}}{\Phi\left(Y_{s(t)}\right) + \Omega\left(Y_{s(t)}\right)\ p_{s(t)}^{t-1}}, = \frac{\left(1 - p_0^0\right) \prod_{k=1}^{t} \Phi\left(Y_{s(k)}\right) + \sum_{k=1}^{t} f\left(k, t, Y_{s(k)}\right)}{\prod_{k=1}^{t} \Phi\left(Y_{s(k)}\right) + \sum_{k=1}^{t} f\left(k, t, Y_{s(k)}\right)}. \tag{5}$$

The above expression for the aggregate belief as a function of time step describes the sequential nature of the search process. However, the searcher is ultimately tasked to arrive at a search decision which states that the target is present in the search area, and if so, at which location, or absent altogether. This search decision reflects the search termination criterion, which is an essential component of any search process. In the presented model, the searcher can conclude that the target is present in cell c' (and thus in the search area) if the cell belief, $p_{c'}^t$, exceeds an upper search decision threshold, denoted $\overline{B}$. Conversely, the negative search decision is made when the aggregate belief falls below a lower search decision threshold, that is, when $B(t) \leq \underline{B}$. These search decision thresholds represent the levels of confidence required of the searcher when making a decision, e.g., a value of $\overline{B} = 0.95$ represents the minimum requirement that the searcher must believe the target is present in a specific cell with 95% certainty. The thresholds could, in general, be time-varying, i.e., $\overline{B}(t)$ and $\underline{B}(t)$, to reflect dynamic search confidence requirements, but are assumed constant (in time) in the presented work. An illustration of the search decision and its evolution as a function of time is shown in Fig. 3. In this manner, one can study the effect of different search parameters, e.g., the initial aggregate belief value ($B(0)$), the initial target probability distribution map ($\mathcal{M}_0$), the search decision thresholds ($\overline{B}$ and $\underline{B}$), and different sensor characteristics, such as the detection error probabilities (α, β), on the search decision performance.

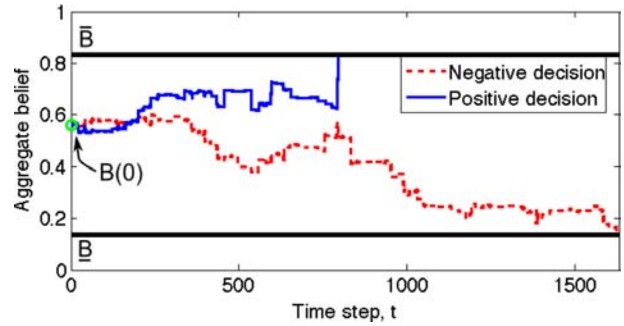

Figure 3. Illustration of the evolution of the search decision process starting at initial aggregate belief, $B(0)$. Two traces are shown, highlighting the positive or negative decisions representing the conclusions that the target is present or absent, respectively, in the search area. The crossing of the upper and lower decision thresholds, $\overline{B}$ and $\underline{B}$, terminates the search process, though the evaluation of the search decision's correctness depends on the true state (i.e., presence and/or location) of the target. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Note that erroneous search decisions are possible, that is, the searcher might make a positive decision (for example, after false positive detections) but the target is truly absent, or alternatively, the searcher incorrectly concludes the target is not present (i.e., a negative decision) but, in fact, the target is present but merely erroneously undetected. In this manner, a key performance metric of the search decision process can be the searcher's ability to make a correct search decision, e.g., as measured by the probability that the search decision is a correct (positive or negative) one.

4. THEORETICAL RESULTS

In general, computation of the aggregate belief evolution requires both the searcher's trajectory, $\mathcal{S}_t$, and its observation history, $\mathcal{Y}_t$, which, in addition to known search and sensor parameters, wholly provide the inputs for evaluation of (5). This section includes a general result for the probabilistic search model describing the effect of positive or negative detections on the aggregate belief. Further, given that exhaustive search patterns completely define the searcher's trajectory "a priori", analytic study of the search process as a function of the sequential uncertain observations is readily possible. In this manner, special cases of exhaustive search examined herein offer several insightful results in a formal context. Such insights are further compared with and validated by data from simulation experiments in Section 6.

4.1. Rate-of-Change of the Search Decision Evolution

We begin our theoretical study with the statement of the following result.

LEMMA 1: A positive detection, $Y_{s(t+1)} = 1$, at the next time step $t+1$ in any cell $s(t+1)$ always leads to an increase in the aggregate belief, $B(t)$, whereas a negative detection, $Y_{s(t+1)} = 0$, always results in a decrease in $B(t)$.

PROOF: Let $\Delta B = B(t+1) - B(t)$ represent the change in the aggregate belief in a single time step. Substitution of the one-step update Eq. (3) yields the expression:

$$\Delta B = p_0^t - p_0^{t+1} = \frac{\Omega\left(Y_{s(t+1)}\right) p_{s(t+1)}^t p_0^t}{\Phi\left(Y_{s(t+1)}\right) + \Omega\left(Y_{s(t+1)}\right) p_{s(t+1)}^t} \quad (6)$$

Then, the time rate-of-change of the aggregate belief can be seen to be positive (that is, $\Delta B > 0$) when $Y_{s(t+1)} = 1$ and negative ($\Delta B < 0$) when $Y_{s(t+1)} = 0$. □

The preceding result can offer some insights to the general behavior of the aggregate belief evolution and its connection to the search decision on termination of the search process. This relationship implies that a positive search decision, that is, concluding that the object is present in a given cell, is usually associated with the gathering of positive detections, and a negative search decision of declaring the search area clear of the object is tied to negative ones. Note, however, that there may be situations where it is possible, due to the interdependence of cell beliefs, that a negative detection in one cell can cause another cell's belief value to increase and cross the upper threshold. This possibility highlights the additional difficulty in addressing dependent cells and the richness of the presented formulation, as a positive search decision can arise from a negative detection in the above manner, which does not arise in the simpler formulations assuming independent cells found in previous works. We note that, in practice, such occurrences are rare, and the nominal insights are still useful for describing the general behavior of the aggregate belief evolution.

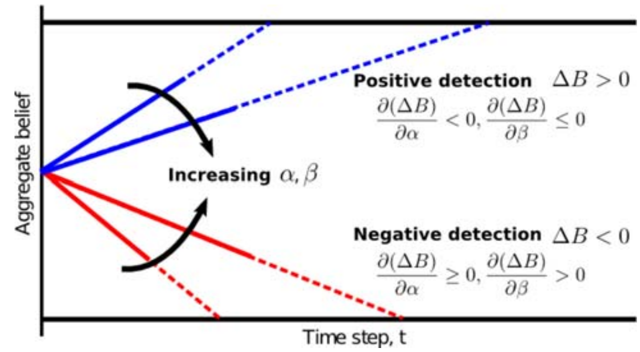

Figure 4. Effect of a positive or negative detection on the aggregate belief, and the dependence on the detection error probabilities. A positive detection increases the aggregate belief (i.e., $\Delta B > 0$), whereas a negative one decreases it ($\Delta B < 0$). Increasing the detection error probabilities causes the change in the search decision evolution to occur more slowly, such that one can conjecture that the time until a search decision is made will also increase. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

We can continue to examine the dependence of the rate-of-change of the aggregate belief on the respective detection error probabilities, that is, on α and β. The partial derivatives of ΔB with respect to α and β describe how changes to these error rates affect the effective rate of the search decision process and are given by:

$$\frac{\partial(\Delta B)}{\partial \alpha} = \frac{p_{s(t+1)}^t p_0^t \left[\beta - Y_{s(t+1)}\right]}{\left[\Phi\left(Y_{s(t+1)}\right) + \Omega\left(Y_{s(t+1)}\right) p_{s(t+1)}^t\right]^2}$$

$$\frac{\partial(\Delta B)}{\partial \beta} = \frac{p_{s(t+1)}^t p_0^t \left[(1-\alpha) - Y_{s(t+1)}\right]}{\left[\Phi\left(Y_{s(t+1)}\right) + \Omega\left(Y_{s(t+1)}\right) p_{s(t+1)}^t\right]^2}$$

For positive detections $\left(Y_{s(t+1)} = 1\right)$, the slope of the rate-of-change is seen to be negative, i.e., $\frac{\partial(\Delta B)}{\partial\alpha} < 0$ and $\frac{\partial(\Delta B)}{\partial\beta} \leq 0$, which implies that increasing α or β leads to a "decrease" in the rate-of-change ΔB, recalling that positive detections result in positively valued ΔB. Conversely, $\frac{\partial(\Delta B)}{\partial\alpha} \geq 0$ and $\frac{\partial(\Delta B)}{\partial\beta} > 0$ for negative detections $\left(\text{i.e., } Y_{s(t+1)} = 0\right)$, or in other words, increases in α or β effectively "increases" ΔB, which is to say that ΔB becomes less negative (since ΔB is negative for any negative detection). This relationship is depicted graphically in Fig. 4 for the dependence on the detection error probabilities for positive and negative detections.

As seen by the extrapolated lines, the qualitative effect of increasing error probabilities is also to increase the number of time steps needed before either a positive or negative search decision is made, i.e., the upper or lower threshold is crossed. Simulation studies discussed in detail in Section 6 corroborate this intuition on how worsening sensor characteristics correspond to diminished search performance, measured by the expected time until the search decision is made.

4.2. Special Case: Sequence of Negative Detections

Suppose the sweeping search strategy is executed, and we are interested in the sequence of negative detections, $\mathcal{Y} = \{0, \ldots, 0\}$. Then, $\Phi(0) = 1 - \alpha$ and $\Psi(0) = \beta$, and the term $\Theta_{s(k)}\left(Y_{s(l)} = 0\right)$ in (5) will occasionally be $\Psi(0)$, since cell $s(k)$ will be revisited once in each additional search pass for time step $t > C$. Substitution and manipulation of (5) leads to the following closed-form expression when $t \leq C$ (where no revisits have yet occurred, so $\Theta_{s(k)}\left(Y_{s(l)} = 0\right) = \Phi(0)$):

$$B(t) = \frac{B(0)\ (1-\alpha) - (1-\alpha-\beta)\sum_{k=1}^{t} p_{s(k)}^{0}}{(1-\alpha) - (1-\alpha-\beta)\sum_{k=1}^{t} p_{s(k)}^{0}}$$

Note that for uniform prior distribution, where $p_c^0 = \frac{B(0)}{C}$ $\forall c$, then the closed-form expression for the belief evolution for this special case reduces to previous results [6]:

$$B(t) = \frac{(C-t)\,B(0)\ (1-\alpha) + \beta\ B(0)\,t}{(C - B(0)\,t)\,(1-\alpha) + \beta\ B(0)\,t}$$

For repeated search passes, let n denote the iteration counter of the number of complete sweeps made through the environment, with $n = 0, 1, \ldots$. Substitution and manipulation of (5) yields the following expression for the belief evolution with an additional parameter, n, for $nC < t \leq (n+1)C,\ n = 0, 1, 2, \ldots$:

$$B(t) = \frac{B(0) + g(t,n)}{1 + g(t,n)}, \tag{7}$$

where

$$g(t,n) \triangleq B(0)\left[\left(\frac{\beta}{1-\alpha}\right)^{n} - 1\right] + \left[\left(\frac{\beta}{1-\alpha}\right) - 1\right]\left(\frac{\beta}{1-\alpha}\right)^{n}\left[\sum_{k=nC+1}^{t} p_{s(k)}^{0}\right].$$

The above closed-form equation for the evolution of the search process with a sequence of null detections is illustrated in Fig. 5.

4.3. Special Case: Sequence of Negative Detections with Periodic Correct Positive Detection

Whereas the previous section explored a trajectory ultimately arriving at the negative search decision (i.e., target is nowhere present in the search area), consider instead a target which is located in cell c' and causes a correct positive detection $Y_{c'} = 1$ during the searcher's visit to this cell, which

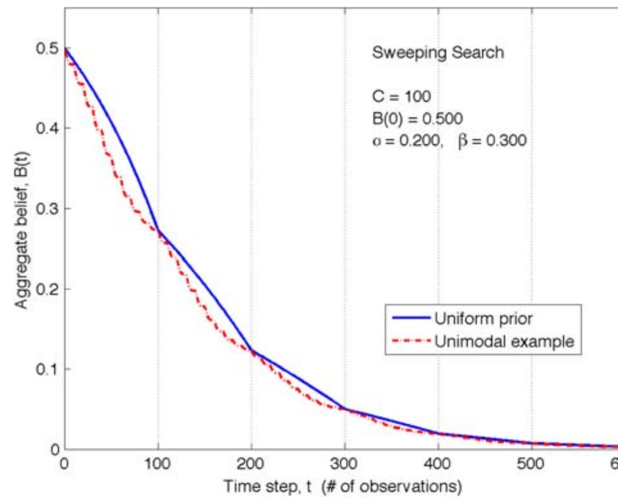

Figure 5. Evolution of the aggregate belief for a sequence of negative detections, as described by (7) for given fixed parameters. The solid line corresponds to the evolution for a uniform initial target distribution, whereas the dashed line represents the use of the uni-modal target distribution depicted in Fig. 1. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

otherwise (correctly) obtains negative detections elsewhere, i.e., $Y_{s(t)} = 0$ for $s(t) \neq c'$. Then, the history of observations is a sequence of negative detections with a periodic positive detection when the searcher inspects cell c', that is, $\mathcal{Y} = \{0, \ldots, 0, Y_{c'} = 1, 0 \ldots, 0, Y_{c'} = 1, 0, \ldots\}$.

The evolution of the aggregate belief in this case is given by the expression below:

$$B(t) = \frac{B(0) + g(t,n) + h(t,n,c')}{1 + g(t,n) + h(t,n,c')}, \tag{8}$$

where $g(t,n)$ is as defined in the previous section, and $h(t,n,c')$ is dependent on whether the searcher has yet visited cell c' in the current sweep. For $nC + 1 < t < nC + c'$, that is, the portion of the sweep prior to arriving at cell c', the term is given by:

$$h(t,n,c') \triangleq p_{c'}^{0}\left[\left(\frac{1-\beta}{\alpha}\right)^{n} - \left(\frac{\beta}{1-\alpha}\right)^{n+1}\right],$$

whereas if the searcher's sweep has already passed cell c', that is, when $nC + c' \leq t \leq (n+1)C$, we have

$$h(t,n,c') \triangleq p_{c'}^{0}\left[\left(\frac{1-\beta}{\alpha}\right)^{n+1} - \left(\frac{\beta}{1-\alpha}\right)^{n+1}\right].$$

Figure 6 illustrates the evolution of the aggregate belief for this context, which represents the idealized outcome for a target present in a fixed location (cell 42).

The theoretical results of this section show how the exhaustive search patterns facilitate deeper analysis of the probabilistic search model presented in this article. The special

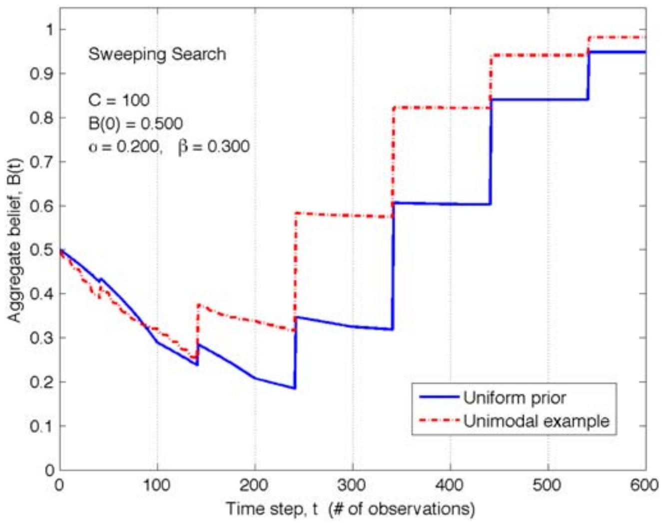

Figure 6. Evolution of the aggregate belief for a sequence of negative detections containing a single positive detection at a fixed location during each search pass. Consider a positive detection occurring in cell 42 in every sweep of the search area, with negative detections elsewhere. The aggregate belief diminishes with the exception of progressively larger increases in $B(t)$ where repeated positive detections occur, until the upper threshold is crossed. The case of a uniform prior target distribution (solid) is contrasted with the effect of using the uni-modal target distribution (dashed) shown in Fig. 1. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

cases provide further opportunity for future work in exploring the exhaustive probabilistic search model, such as examining bounds on the time until a (correct) search decision is achieved. These special cases update the probabilistic model one time step after another and are a single realization of a the simulation driven by the probabilistic models developed in Section 3 of this article.

In the next section, we present results based on an extensive set of simulations run, rather than single instances. This development of the simulation model allows the study of how various components of the model formulation influence metrics of interest, e.g., for a SAR operator. Unfortunately, this type of study on a stochastic model cannot be achieved without some sort of empirical research study, which is extremely useful in answering practical questions, such as "how does increasing false positive detection rates change the time it takes to arrive at a correct decision?" We can exploit the simulation to answer this by using Monte Carlo methods. However, the drawback in using Monte Carlo methods is that we generally fix all variables in the model at a particular (baseline) level and only study the influence of changing one variable at a time. This is referred to as a one-factor-at-a-time study.

The major disadvantage to this type of one-factor-at-a-time Monte Carlo strategy is that it fails to consider any possible interaction between the variables of interest. For example, maybe the time it takes to make a correct decision increases as the false positives decrease when a lawnmower search pattern is used, but stays relatively unchanged when an optimal search strategy is employed. This type of interaction would go unnoticed unless the variables (in this case false positive detection rate and search strategy) are varied together. In order to study how the variables in our probabilistic model influence SAR operator metrics of interest, we empirically collect results using experimental design strategies on our simulation model of the SAR process. This is described in detail in the next section.

5. SIMULATION MODEL

In Section 3, a probabilistic model formulation of the SAR problem was presented. This model can be used as a stochastic simulator. In this simulator, there are a number of variables that can be set prior to running (stepping through, by increasing the time steps) the simulation. The particular variables of interest that were identified in Section 3 are: the detection error probabilities (α and β), the aggregate belief starting point ($B(0)$, that is, the belief at time zero), and the search decision thresholds ($\underline{B}$ and $\overline{B}$). Changing any one of these variables changes the outcome of the simulation model. The goal, in this section, is to study how much of an influence each of these variables has on several different outcomes collected during the course of a single simulation, run. Experimental design is used to study these relationships.

The outcomes of the simulation model (also known as responses in the experimental design community) are presented in Section 5.1. The variables manipulated (or set) during each subsequent iteration in the simulator, and their respective setting choices and ranges are discussed in Section 5.2. The variables mentioned in the above paragraph are presented as well as two additional variables (namely, initial target probability map and search pattern). Finally, this section concludes with the experimental design set up used to collect a set of empirical results (Section 5.3). These empirical data collected by running the experimental design are then used to lead to insights on informing or refining concepts of unmanned vehicle operation in SAR missions. The statistical results and data will be presented in Section 6.

5.1. Probabilistic Search Responses

As mentioned previously, there are numerous measures of search performance involved in the analysis of probabilistic search methods. Given operational relevance of the time necessary to complete search, variants of the expected time until a search decision is reached are included in the responses of interest. The implemented simulation model is an example of a "terminating simulation", since the search

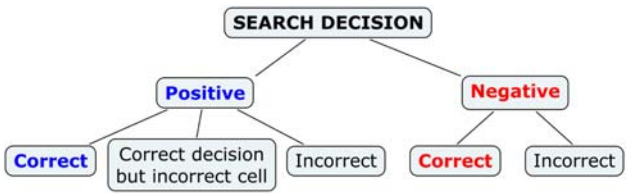

Figure 7. Taxonomy of probabilistic search decisions, where positive and negative decisions can be erroneous due to uncertainty in the search process. The bold labels represent the two measures of performance—percentage of correct positive and correct negative search decisions—highlighted in the statistical studies presented in this article. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

process terminates with the searcher's arrival at a search decision about the target's presence or absence in the search area. Since the quality of the search decision itself is also of high value in operational contexts, the percentage of correct search decisions also serves as a valid measure and, thus, provides additional responses for study. Investigation of both the classes of responses is facilitated by the probabilistic search model formulation developed in this article.

Consider refinement of the latter responses of interest, where the search decision is categorized into two types, that is, positive and negative, search decisions. Recall that a positive search decision is made when the searcher declares the target to be present in the search area (crossing the upper decision threshold) and specifies the target's likely location, whereas the latter search decision type reflects the searcher stating the target's absence from the search area (on crossing the lower decision threshold). Within these decision types, the imperfect nature of the searcher's sensor and the uncertain evolution of the aggregate belief introduce the possibility of search decision errors. While the negative search decision can be in error only when the target is indeed present, the positive search decision can be incorrect in two ways: the searcher can declare the target to be present when it is actually absent or when it is truly present but in a location other than the one specified by the searcher. Figure 7 illustrates the classification of search decision types and possible errors. In the study presented in this article, we focus on two responses—percentage of correct positive and correct negative search decisions—out of the five possible classes, as these metrics are often of primary interest when evaluating operational search methods. These response values (between 0 and 100%) are computed based on the total number of simulation replications for each set of variable values, where the decision type and error are recorded for each run.

Additionally, this research is also concerned with evaluating the search performance as measured by the expected number of time steps necessary for the searcher to reach a decision, specifically correct decisions, denoted $\mathbb{E}[TTCD]$. Deeper inspection on the expected number of time steps to make correct positive and correct negative decisions, denoted $\mathbb{E}[TTCPD]$ and $\mathbb{E}[TTCND]$, respectively, provide further insights of operational value. Observe that these average times are related by conditional expectation, conditioned on the target's presence or absence in the search area (i.e., r.v. $H = \{0, 1\}$):

$$\begin{aligned}\mathbb{E}[TTCD] &= \mathbb{E}[TTCD \mid H = 1]\, Pr[H = 1] \\ &\quad + \mathbb{E}[TTCD \mid H = 0]\, Pr[H = 0], \\ &= \mathbb{E}[TTCPD]\;\; B(0) + \mathbb{E}[TTCND]\;\; p_0^0.\end{aligned}$$

To summarize, the five responses examined in this work include percentage of (1) correct positive search decisions and (2) correct negative search decisions, as well as expected number of time steps for (3) correct search decisions, $\mathbb{E}[TTCD]$, (4) correct positive search decisions, $\mathbb{E}[TTCPD]$, and (5) correct negative search decisions, $\mathbb{E}[TTCND]$.

Table 1. Design Variables in Exhaustive Probabilistic Search.

Variable	Label	Description	Range or levels
α	x_1	False positive detection error	[0.0, 1.0]
β	x_2	False negative detection error	[0.0, 1.0]
$B(0)$	x_3	Initial aggregate belief	[0.3, 0.7]
$\overline{B}$	x_4	Upper decision threshold	[0.8, 0.95]
$\underline{B}$	x_5	Lower decision threshold	[0.05, 0.2]
$\mathcal{M}(0)$	x_6	Initial target probability map	{good, bad, none}

5.2. Simulation Variables of Interest

In practice, countless variables can affect the search mission performance; however, ones that are often incorporated in theoretical and applied probabilistic search models include quantities such as search area parameters, sensor characteristics, search patterns, and available prior information. This section defines the set of variables considered in this research to better understand how they relate both to each other and the search responses described above. In particular, we study seven variables of interest, namely the detection error probabilities (α, β), the initial aggregate belief $B(0)$, the search decision thresholds $(\overline{B}, \underline{B})$, the initial target probability map $\mathcal{M}(0)$, and the search pattern SP (i.e., lawnmower vs. sweeping pattern). Table 1 identifies labels for these variables, and their valid ranges of values as considered in the statistical simulations. The variables denoted x_1 through x_5 are termed "continuous," meaning they can take values within a continuous real-valued interval, whereas x_6 is an example of a "categorical" variable, which characterize discrete, non-numeric variables.

The categorical variable, x_6, has three levels denoted "good," "bad," and "none," representing the quality of the information available prior to the start of search. Such information is often present in search missions (e.g., last

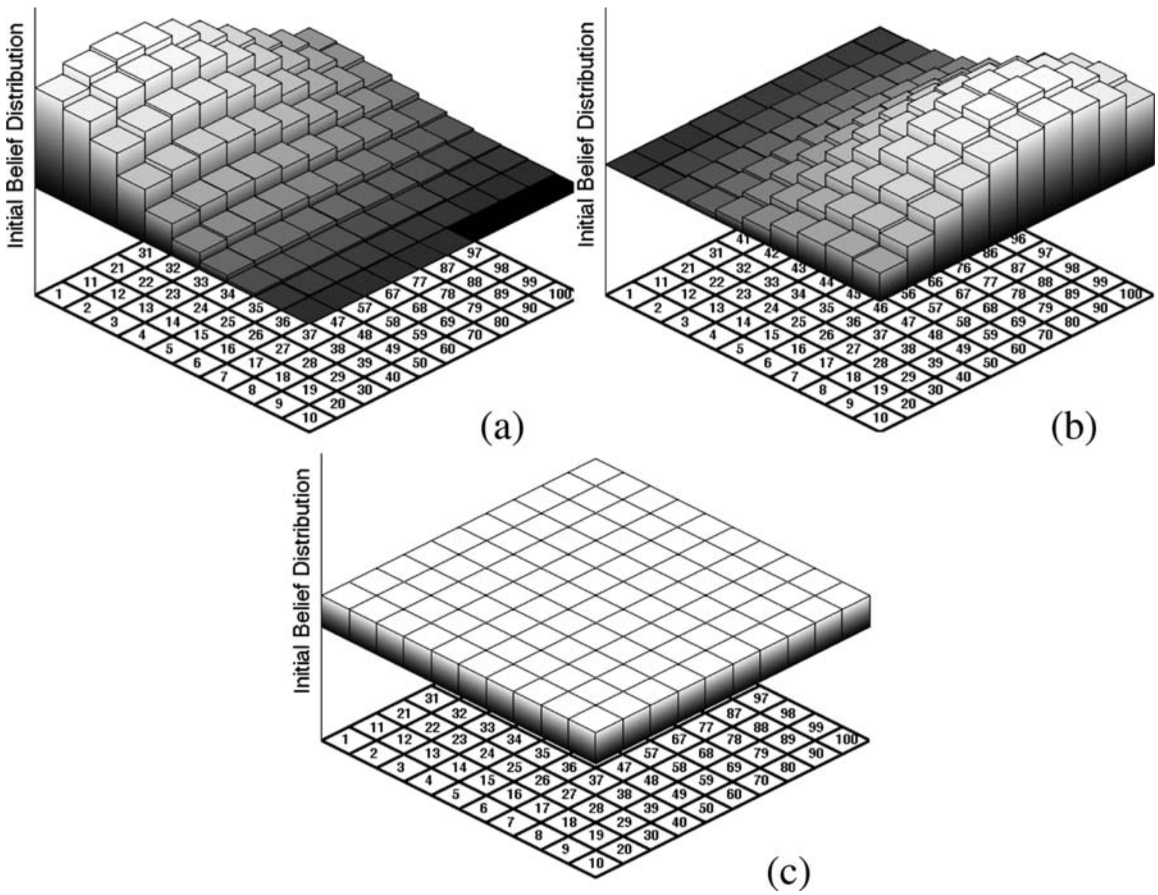

Figure 8. Three initial target probability distributions representing the categorical variable, $x_6 \in$ {good, bad, none}. For the first two levels, the searcher assumes the random target location is distributed according to (a), even though for the "bad" case the true target distribution is given by (b). This variable provides insight into the value of prior information, compared with the uninformed case given by (c) a uniform target distribution.

known location, intelligence reports, historical data), and in general, can be used to inform the search process. While exhaustive search is considered "nonadaptive," that is, the searcher's trajectory does not depend on the available and/or updated probability information, the evolution of the aggregate belief (and thus, the search decision) depends on the initial target probability map, e.g., see (4) and (5). Consider Fig. 8, which illustrates the three true probability distributions used to initialize the true target location in the simulation experiments. In both cases of "good" and "bad" levels, the searcher assumes the target location is randomly selected from the same distribution, i.e., that of Fig. 1 (equivalent to Fig. 8a). However, in the latter "bad" case, the "true" distribution of the target location is shown in Fig. 8b. This discrepancy reflects the possibility of incorporating misleading information, whether accidental or intentional (i.e., a deceptive target), in the execution of the search mission, which can potentially lead to greater errors in the final search decision. In the third information type (x_6 = none), no prior information is reflected by a uniform probability distribution, serving as a baseline for the benefit or penalty of good or bad information, respectively.

5.3. Design of Experiments

Statistical design of experiments is a methodology for studying the relationship between inputs and output(s) for a particular system of choice, whether physical or simulated. When researchers are interested in learning about or improving such a system, experimental design techniques can be us in order to develop a set of experimental runs that can efficiently and effectively study the influence of the inputs on the output(s). In general, the choice of experimental design depends on: the goal of the experiment (e.g., screening or sensitivity analysis), the number of variables to be explored, and the number of experiments that can reasonably be conducted for limited resources (e.g., time, budget). Based on the screening objective of this study, one can often expect linear regression models to be sufficient for explaining the simulation results, which makes factorial or fractional factorial designs suitable [17]. This section describes the choice of experimental design for this study, providing the foundation for the statistical results described in Section 6.

This research has two elements that prevent a stand-alone factorial design from being used, namely the use of percentage-valued (vice unbounded, real-valued) responses and a constrained variable space. Specifically, as we consider percentage metrics of correct positive and correct negative search decisions, it is necessary to use "generalized linear" modeling techniques, such as logistic regression, which requires transformation of the response data. Additionally, the variables α and β are constrained by the relationship that $1 - \alpha - \beta > 0$ (c.f. Section 3.2), which restricts the experimental design space. Also, preliminary simulation tests indicated that as $\alpha + \beta$ approaches one, the number of time steps required to make a search decision grows unbounded. To facilitate reasonable simulation execution times, the valid region explored in α, β-space was restricted further in the experimental design such that $0.8 > \alpha + \beta$.

The two unique research components, that is, responses that call for a mixture of linear and generalized linear regression models and the constrained design region, lead to the selection of a hybrid experimental design. Optimal designs, such as the D-optimal design, are well suited for screening experiments, especially in those cases where the design region is constrained, whereas space-filling designs, such as the uniform design, are well suited for experiments where multiple models are to be used to summarize the output response data.

"Optimal designs" (or alphabetic optimality designs) were first introduced by [10, 11], and [12]. Different choices of optimal experimental designs can be found in relevant textbooks (e.g., see [1]), such as the D-optimal design, including methods for their creation and application. In particular, the D-optimal design minimizes the joint confidence region of the unknown and uncertain regression model coefficients to improve the precision of their estimates. Further, D-optimal design is a suitable choice for "screening" experiments, that is, experiments seeking to identify the most influential variable parameters or "variables" in the system. Alternatively, "uniform designs," first introduced in [7], are a good choice of space-filling designs, i.e., ones which uniformly explore the entire variable space, when flexibility in fitting the output data is required, such as when the analyst wishes to remain agnostic in choice of regression model.

In this work, a hybrid design consisting of a combination of D-optimal and uniform design is used. Hybrid design approaches combining optimal with space-filling designs are discussed in [9] and shown to perform well when the form of the statistical model is not known in advance and the response may have nonlinearities present. The selected hybrid experimental design for the present study uses 16 D-optimal design points and 9 uniform design points.

Given the advantage of simulation experiments, repeated runs or "replications" for each "design point" (i.e., a given combination of variable values) can be conducted to generate a comprehensive statistical picture of the output. The hybrid design of 25 combinations in the five continuous variables was applied to each of the six combinations of levels of the two categorical variables, yielding an experimental design comprising 150 unique design points. We executed this experimental design replicating each design point 1,200 times to efficiently explore the multidimensional variable space.

6. STATISTICAL ANALYSIS AND RESULTS

Given the output of the stochastic simulation runs, summary data can be generated for the five responses of interest in preparation for statistical analysis. In order to perform statistical screening techniques, explicit regression model types need to be considered. The approach of logistic regression is well suited for the two percentage-valued responses, as it leverages an invertible mapping (via the logit function) to transform the response data such that standard multiple linear regression methods can be applied. (We assume that the underlying transformed model is intrinsically linear. Refer to [18] for a reference on logistic regression methods.) For the remaining three responses, i.e., $\mathbb{E}[TTCD]$, $\mathbb{E}[TTCPD]$, and $\mathbb{E}[TTCND]$, we can apply standard multiple linear regression techniques, which require specification of the polynomial bases for the fitting equation to capture all variables that may influence the response. We choose to seek a model for the response data which includes main effects, two-variable interactions, and squared terms (though other models are possible and subject of future analysis). This choice is described further in the following brief primer on the statistical regression approaches.

6.1. Statistical Regression Approaches

The objective of linear regression is to construct a model, as specified by its polynomial coefficient values,[1] denoted β_i, that sufficiently explains the response data. Consider an arbitrary response y, such that the general form of the multiple linear regression model (with main effects) is given by

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k + \varepsilon, \tag{9}$$

where ε is the error and x_i denotes each of the k regressors in the model. Written in matrix notation, the main effects linear model is given by $\boldsymbol{y} = \boldsymbol{X\beta} + \boldsymbol{\varepsilon}$ where $\boldsymbol{y} \in \mathbb{R}^{n\times 1}$ is

[1] Despite possible confusion with the false negative detection error probability ($\beta_{s(t),t}$), the notation for the regression coefficients, β_i, is retained to cohere with the statistics and experimental design communities.

Table 2. Significant variables for Lawnmower Search Regression Models.

	Factors														
		Main Effects							Interactions					Squared Effects	
						X_6				X_1, X_6					
Lanwnmower Response	Intercept	X_1	X_2	X_3	X_5	bad	good	none	X_1, X_5	bad	good	none	X_3, X_5	X_1^2	X_2^2
% Correct Neg	−0.147	−0.172		−0.888											
% Correct Pos	−0.917	−0.094		1.163	−0.404	−0.657	0.330	0.330		−0.436	−0.198	0.238	0.271		
E[TTCD]*	8.644	2.007	1.125		−0.298									−0.580	−1.245
E[TTCND]*	8.462	1.800	1.175	0.415	−0.319										−1.541
E[TTCPD]*	8.737	2.262	0.961	0.110					−0.278					−0.533	−1.549

The response variables pertaining to the expected time metrics (marked with asterisks) exhibited residuals possessing "heteroscedastic variance," which violates assumptions on the error, $\boldsymbol{\varepsilon}$, for using standard regression techniques. A log transformation was applied to these three responses to stabilize the residuals (see [4] for a description of standard transformations).

the response vector (using indicator variables for categorical effects), $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ represents the experimental design matrix, $\boldsymbol{\beta} \in \mathbb{R}^{p \times 1}$ are the $p = k + 1$ regression coefficients, and $\boldsymbol{\varepsilon} \in \mathbb{R}^{n \times 1}$ is the vectorized error. The above regression model can be expanded to include two-variable interactions (e.g., $x_i x_j$) and squared (e.g., x_i^2) terms resulting in an expanded matrix $\boldsymbol{X}$ [17]. Two-variable interactions are addressed by multiplying the $(i+1)^{\text{th}}$ and $(j+1)^{\text{th}}$ columns[2] corresponding to the two associated variables (x_i and x_j), and squared dependencies on a given variable are accounted for by squaring the associated main effect column. The strength of experimental design is the ability to generate a model for all such higher-order terms in an efficient manner. In other words, varying a single variable's value while holding all remaining variables at fixed levels, as often done in many research communities, fails to capture interactions and higher-order relationships, which can often reflect significant influence on the responses.

Before continuing in the regression analysis, suppose that a response variable, e.g., the percentage of correct positive decisions, represents a probability value between zero and one denoted $P(z)$ (termed the "success probability"). If the response is assumed to have the form of the logistic model given by

$$P(z) = \frac{1}{1 + e^{-z}},$$

where $z = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k$ is called the "linear predictor" (with main effects only), then the "log odds" or "logit" of $P(z)$ is the inverse mapping $[0, 1] \to \mathbb{R}$, given by

$$\operatorname{logit}(P(z)) = \log\left[\frac{P(z)}{1 - P(z)}\right] = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k,$$

which now reflects a linear relationship between the regression coefficients and the transformed response. The above model can be extended to include higher-order terms, including two-variable interactions and quadratic elements, as done for the statistical analysis presented in this article.

The linear nature now facilitates the use of standard multiple linear regression techniques. "Stepwise linear regression" can be used in order to screen for the significant variables, that is, those having a statistically significant impact on the response, and such tools are readily available in standard statistical analysis software, such as JMP®. For the presented research, the stepwise procedure was carried out until the resulting model explained 95% of the variability in the response data. Note that the fitted regression models are generally given in "coded units" [17], which represents a normalization of all continuous variables on the interval $[-1, 1]$ to allow direct comparison of sign and magnitude of the variables' effects on a unit-less scale.

6.2. Fitted Regression Models

The screening results for the simulated search using the lawnmower search pattern are summarized in Table 2, where each entry represents a fitted regression coefficient values for significant variables present in the model. As a note, the equivalence of the lawnmower and sweeping search patterns was verified through additional regression analysis, finding no statistically significant differences in their mean responses for any of the five measures of performance. The results presented herein represent lawnmower search, given its relevance to practical and commonly employed search operations.

The results of the regression analysis show that only five main effects, three two-variable interactions, and two quadratic terms are collectively present in the fitted models for the five responses of interest. The only variable exhibiting significant impact in all five models is x_1, i.e., the false positive detection probability, α, which gives merit to expending resources to minimize false alarms. In contrast, the x_4 variable

[2] Note the first column of $\boldsymbol{X}$ corresponds to the constant intercept, β_0.

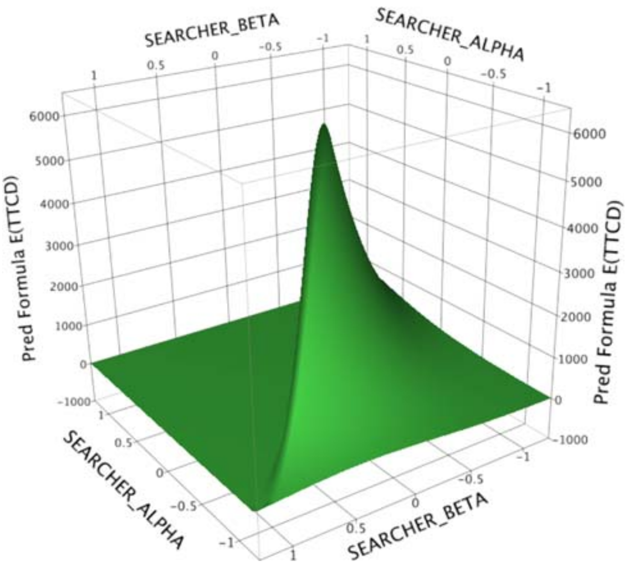

Figure 9. Illustration of the prediction Eq. (11) for $\mathbb{E}[TTCD]$ as a function of false positive (α) and false negative (β) detection error probabilities (in coded units). The remaining variable, x_5, represents the lower search decision threshold, fixed at $\underline{B} = 0.125$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

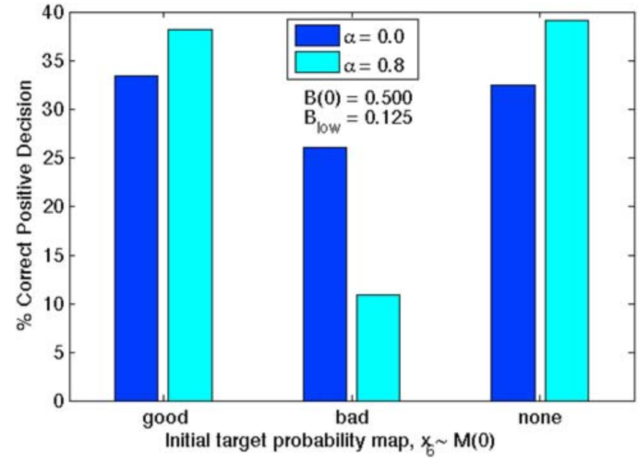

Figure 10. Interaction diagram between false positive detection probability α (x_1) and the initial target probability map, $\mathcal{M}(0)$ (x_6). The use of incorrect prior information on the distribution of target location severely diminishes the search performance as measured by the percentage of correct positive decisions. Remaining variables are held fixed at their middle values, i.e., $x_3 = 0 \to B(0) = 0.5$ and $x_5 = 0 \to \underline{B} = 0.125$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

representing the upper search decision threshold, $\overline{B}$, is absent as a significant variable in any of the responses. When considering the expected time responses, i.e., $\mathbb{E}[TTCD]$, $\mathbb{E}[TTCPD]$, and $\mathbb{E}[TTCND]$, both false detection probabilities, α and β, are most significant and manifest as such in the regression models. Observe that the direction and magnitude of their influence (i.e., the large positive values for main effect entries, x_1 and x_2) validates the conjecture from theoretical results derived in Section 4, that is, by increasing α and β, the amount of time the searcher takes to make a decision also monotonically increases, both theoretically and empirically. The regression models for these expected time responses also contain squared terms, which corresponds to the rapid increase in the number of time steps as one approaches the (α, β)-variable space constraint boundary at $\alpha + \beta = 0.8$, for example, as depicted in Fig. 9 for the $\mathbb{E}[TTCD]$ response, c.f. (11).

Only two of the regression models—the percentage of correct positive decisions and the expected time until making a correct positive decision, $\mathbb{E}[TTCPD]$—contain interaction terms. The x_1x_6 interaction, that is, between the false alarm probability, α, and the initial target probability map type ("good," "bad," or "none"), offers some interesting insights, illustrated in Fig. 10. The vertical axis is the transformed response (logit of percentage of correct positive decisions), the horizontal axis shows the three levels of the categorical variable, $\mathcal{M}(0)$, and the two lines represent the high ($x_1 = 1 \to \alpha = 0.8$) and low ($x_1 = -1 \to \alpha = 0.0$) levels of false positive detection probability, α. For a perfect sensor ($\alpha = 0$, no false positives), having good prior information slightly increases the percentage of correct positive decisions over the uninformed (i.e., no information) case, whereas both still have higher percentage of correct decisions over the misinformed (i.e., "bad" initial target probability map) case. For an extreme case of an imperfect sensor, with $\alpha = 0.8$, there is a much higher performance penalty incurred for conducting search with low fidelity prior information on target location.

Full meta-models for each of the five responses of interest can be reconstructed by the coefficient values contained in Table 2, which can then be used to inform or evaluate future search operations. Two such models are developed below to illustrate the value of these regression models and the statistical methods used to generate them. Consider the fitted regression model for the percentage of correct negative decisions, $P(CND)$, taking the coefficients from the first row of Table 2 to construct the linear regression model

$$\text{logit}\,(P(CND)) = -0.147 - 0.172x_1 - 0.888x_3$$
$$\Longrightarrow \quad P(CND) = \frac{1}{1 + e^{0.147+0.172x_1+0.888x_3}}, \qquad (10)$$

where the second transformed expression represents the fitted logistic regression model. Equation (10) can then be used to make predictions on the percentage of correct negative decisions as a function of false positive detection probabilities, α, and initial aggregate belief values, $B(0)$, which are the significant variables (x_1, x_3) in the model.

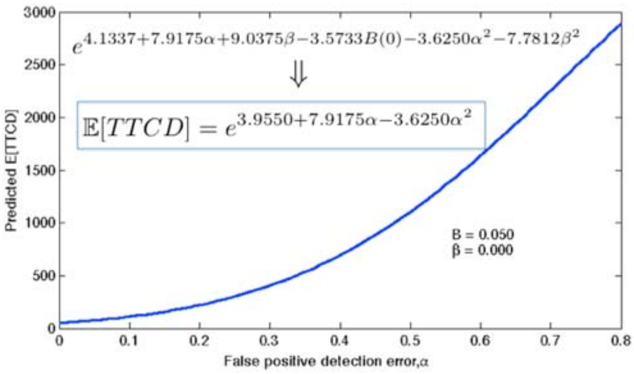

Figure 11. The "one-factor-at-a-time" prediction expression for $\mathbb{E}[TTCD]$ as a function of α (in engineering units). This function is found by examining the regression model (11) and fixing all other variables at constant values, e.g., $\beta = 0$ and $\underline{B} = 0.05$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

As another example, consider the expected number of time steps until a correct decision is made, that is, $\mathbb{E}[TTCD]$. The fitted regression model using the log transform (see note 2) of the response and its equivalent form in actual response units are given by:

$$\begin{aligned} \log(\mathbb{E}[TTCD]) &= 8.644 + 2.007x_1 + 1.125x_2 \\ &\quad - 0.268x_5 - 0.580x_1^2 - 1.245x_2^2 \\ \Longrightarrow \quad \mathbb{E}[TTCD] &= \exp(8.644 + 2.007x_1 + 1.125x_2 \\ &\quad - 0.268x_5 - 0.580x_1^2 - 1.245x_2^2). \end{aligned} \tag{11}$$

Note that the fitted regression models constructed from Table 2, such as the above examples, are still in coded or normalized units (i.e., continuous variables lie in $[-1, 1]$). To revert to the original scales or "natural or "engineering units" associated with the coded variables, the x_i's in the fitted models can be replaced by the conversion (with an example for x_3):

$$x_i = \frac{a_i - \left(\frac{a_{\text{high}}+a_{\text{low}}}{2}\right)}{\left(\frac{a_{\text{high}}-a_{\text{low}}}{2}\right)} \quad \text{Ex: } x_3 = \frac{a_3 - \left(\frac{0.7+0.3}{2}\right)}{\left(\frac{0.7-0.3}{2}\right)} = \frac{a_3 - 0.5}{0.2}$$

where a_i is the actual value in engineering units, and the a_{low} and a_{high} terms represent the low and high values, respectively, for the range of the continuous variable. This final conversion allows the analyst to examine the functional dependencies of the responses on the significant variables, and thereby predict the response values throughout the region of the variable space for which the regression models were generated. Note that these prediction equations are intended to allow interpolation in the valid variable subspace, vice as a means of extrapolating outside this regime, since prediction variance outside of the experimental design region can be shown to grow exponentially as distance from the region increases.

The prediction equations generated from the statistical analysis can also be used to study "one-factor-at-a-time" dependencies, which represent a "slice" in the multi-dimensional variable space. Figure 11 shows the predicted $\mathbb{E}[TTCD]$ as a function of α, which can be constructed by fixing other variables to constant values (e.g., $\beta = 0$ and $\underline{B} = 0.05$). The ability to generate such single-parameter dependencies demonstrates the flexibility of the prediction equations and highlights the advantage of conducting rigorous statistical experimental design and regression analysis.

7. CONCLUSIONS AND FUTURE WORKS

This article investigated the problem of probabilistically searching for a stationary target by means of exhaustive search patterns commonly used in analytic and practical search settings. In this exhaustive search context, the evolution of the search decision of determining whether the target is present (and if so, within which discrete location) or absent from the search area was examined analytically. Theoretical results derived from closed-form expressions describing the search decision process provide insight into the role of positive and negative detections and their corresponding error probabilities, and special cases of interest offer additional insights from the theoretical treatment.

Following the theoretical portion of the article, was an empirical portion, which applied experimental design and statistical analysis to studying the probabilistic based simulation model. A hybrid optimal (D-optimal) and space-filling (uniform) design was used to generate the design matrix (which specifies the simulation runs) for the study. Screening and mathematical modeling were then used to assess the statistically significant variables and generate regression equations for each of the five responses. While many of the results agreed with intuitive notions or analytic results, some unexpected insights include the lack of statistical difference between the lawnmower and sweeping search patterns. Another interesting outcome was the limited influence of the initial target probability map on the search performance measures (only affected the percentage of correct positive search decisions), though one might have expected the quality of available prior information on the target location would be pertinent in many (if not all) of the responses.

In order to investigate some aspects of this study in more detail, there are several plans for future work. A more detailed investigation of search strategy will more closely examine the relationship between the lawnmower and sweeping search patterns. It is anticipated that larger search areas and special types of target probability distributions (e.g., near the boundaries of the area versus the interior region) could prove to show a difference in the time it takes to determine a target's true location. In addition to more work on the mathematical modeling of responses, focus on extreme behaviors will be researched. While the regression equations were good at predicting responses, there was some lack of fit present when predicting cases when the expected time until decision was very large. More sophisticated statistical modeling techniques and/or data transformations will also be considered in the future studies.

REFERENCES

[1] A.C. Atkinson, A.N. Donev, and R. Tobias, Optimum experimental designs, with SAS, Oxford University Press, USA, 2007.
[2] S.J. Benkoski, M.G. Monticino, and J.R. Weisinger, A survey of the search theory literature, Nav Res Logis 38 (1991), 469–494.
[3] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte, "Optimal search for a lost target in a Bayesian world," in: Field and Service Robotics, Vol. 24, STAR Springer Tracts in Advanced Robotics, Y.S. Äôichi, A. Hajima, P. Erwin, T. Takashi, and T. Sebastian (Editors), Heidelberg, 2006, pp. 209–222.
[4] G.E.P. Box, J.S. Hunter, and W.G. Hunter, Statistics for experimenters. Design, innovation, and discovery, 2nd ed., Wiley series in probability and statistics, Wiley, Hoboken, NJ, 2005.
[5] H. Choset, Coverage for robotics—A survey of recent results, Ann Math Artif Intell 31 (2001), 113–126.
[6] T.H. Chung and J.W. Burdick, "A decision-making framework for control strategies in probabilistic search," in: Proceedings of the 2007 IEEE International Conference on Robotics and Automation, Rome, Italy, IEEE, 2007, pp. 4386–4393.
[7] K.T. Fang, The uniform design: Application of number-theoretic methods in experimental design, Acta Math Appl Sin 3 (1980), 363–372.
[8] IAMSAR International aeronautical and maritime search and rescue manual, International Maritime Organization Publishing, London, United Kingdom, 2009.
[9] R.T. Johnson, D.C. Montgomery, B. Jones, and P.A. Parker, Comparing computer experiments using high order polynomial metamodels, J Qual Technol 42 (2009), 86–102.
[10] J. Kiefer, Optimum experimental designs, J R Stat Soc Series B Stat Methodol 21 (1959), 272–319.
[11] J. Kiefer, Optimum designs in regression problems, II, Ann Math Stat 32 (1961), 298–325.
[12] J. Kiefer and J. Wolfowitz, Optimum designs in regression problems, Ann Math Stat 30 (1959), 271–294.
[13] B.O. Koopman, The theory of search, Part {II}. Target detection, Oper Res 4 (1956), 503–531.
[14] B.O. Koopman, Search and screening: General principles with historical application, 2nd ed., Pergamon Press, Elmsford, NY, 1980.
[15] M. Kress, R. Szechtman, and J.S. Jones, Efficient employment of non-reactive sensors, Mil Oper Res 13 (2008), 45–57.
[16] H. Lau, Optimal search in structured environments, Ph.D. dissertation, University of Technology, Sydney, Australia, 2007.
[17] D.C. Montgomery, Design and analysis of experiments, 7th ed., John Wiley, Hoboken, NJ, 2009.
[18] R.H. Myers, D.C. Montgomery, G.G. Vining, and T.J. Robinson, Generalized linear models with applications in engineering and the sciences, 2nd ed., John Wiley, Hoboken, NJ, 2010.
[19] H. Sato, Path optimization for single and multiple searchers: Models and algorithms, Ph.D. Dissertation, Naval Postgraduate School, Monterey, CA, 2008.
[20] L.D. Stone, J. Kettelle (Editor), Theory of optimal search, 2nd ed., Academic Press, Hanover, MD, 1989.
[21] United States National Search and Rescue, United States National Search and Rescue Supplement, Washington DC, 2000.
[22] D.H. Wagner, Naval operations analysis, United States Naval Institute, Annapolis, MD, 1999.
[23] A.R. Washburn, Search and detection, 4th ed., Topics in operations research series INFORMS, Hanover, MD, 2002.
[24] T.A. Wettergren and J.G. Baylog, Modeling sequential searches with ancillary target dependencies, Adv Decis Sci 2010 (2010), 1–27.