



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2008

Multi-agent Probabilistic Search in a Sequential Decision-theoretic Framework

Chung, Timothy H.; Burdick, Joel W.

Timothy H. Chung and Joel W. Burdick, "Multi-agent Probabilistic Search in a Sequential Decision-theoretic Framework," 2008 IEEE International Conference on Robotics and Automation, Pasadena, California. 2008.

<http://hdl.handle.net/10945/36552>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Multi-agent Probabilistic Search in a Sequential Decision-theoretic Framework

Timothy H. Chung

Department of Operations Research
Naval Postgraduate School
Monterey, California, USA
thchung@nps.edu

Joel W. Burdick

Department of Mechanical Engineering
California Institute of Technology
Pasadena, California, USA
jwb@robotics.caltech.edu

Abstract—Consider the task of searching a region for the presence or absence of a target using a team of multiple searchers. This paper formulates this search problem as a sequential probabilistic decision, which enables analysis and design of efficient and robust search control strategies. Imperfect detections of the target’s possible locations are made by each search agent and shared with teammates. This information is used to update the evolving decision variable which represents the belief that the target is present in the region. The sequential decision-theoretic formulation presented in this paper provides an analytic framework to evaluate team search systems, as it includes a performance metric (time until decision), a measure of uncertainty (decision confidence thresholds) and imperfect information gathering (detection error). Strategies for cooperative search are evaluated in this context, and comparisons between homogeneous and hybrid search strategies are investigated in numerical studies.

I. INTRODUCTION

As autonomous robots are increasingly employed for intelligent information-gathering tasks, decision-making capabilities in these systems become more relevant and essential for efficient completion of assigned tasks as well as potential interaction with their human counterparts. One such task is probabilistic search of a region for the presence or absence of a target. There is renewed interest in probabilistic search due to advances in autonomy that enable robots to realize various new search-related objectives. Search applications include traditional scenarios such as search-and-rescue missions and surveillance for intruder detection, but also extend to other arenas such as modeling of animal foraging or hunting behavior and search-driven visual attention.

The classical theory of search was introduced by Koopman [1] for the search of naval vessels (e.g., submarines) during World War II, and furthered by many others [2], [3], [4] in the context of operations research. The advent of ubiquitous sensing and inexpensive computing capabilities is facilitating a natural progression from human-guided search to autonomous search using intelligent mobile robots. Searching problems also arise in other research communities, including signal detection [5], machine learning [6], and human visual attention [7] research. Furthermore, teams of mobile search agents endowed with communication and coordination capabilities can cooperatively conduct autonomous search, taking advantage of the benefits of multi-robot systems, such as increased reliability and performance in the collective objective [8], [9], [10], [11].

However, much of previous research literature on search falls short in addressing various components of the probabilistic search task. In particular, previous methods to optimize search plans using only prior information, such as those found in classical search theory, lack the benefit of feedback, which is essential when dealing with dynamic environments (e.g., moving targets). Also, notions of the time or effort necessary to complete the search task are largely missing in robotic search literature, even though such metrics are relevant for many situations such as rescue or recognition tasks. Additionally, most prior works employ a limited model of uncertainty in observations, where only missed detections of the target are considered, neglecting possibly incorrect positive measurements. Finally, search termination criteria are often fixed arbitrarily without consideration for how they are related to other factors such as detection characteristics and search time.

In contrast, the work presented in this paper incorporates search control strategies which update the agents’ search plans iteratively (as in [12]) in order to adapt to possible changes in the search environment. Further, by casting the search task as a sequential decision, the time till a decision is made regarding the presence or absence of the target becomes meaningful and serves as a performance metric when comparing various search methodologies. Also, a general model for imperfect observations is employed, such that both missed detections and false alarms are easily accounted for. An additional strength of the decision-theoretic framework is the relationship that exists between the time till decision metric, the search termination thresholds (which correspond to confidence measures), and the amount of uncertainty in the information-gathering (i.e., detection error rates). These advantages to the approach taken in this paper help facilitate the analysis and design of efficient and practical search control strategies.

The main contributions of this work are the decision-theoretic formulation of a multi-agent probabilistic search problem, including an analytic expression governing the evolution of the decision. Additionally, the advantages of this framework for notions of heterogeneous search tasks for teams of intelligent agents are also developed. This paper is organized as follows, beginning with a formal statement and construction of the considered search problem in Section II. The evolution of the search decision process is developed and presented in Section III, followed by investigation of the

role of team search control strategies via numerical studies in Section IV. Section V summarizes the results and discusses future avenues of research.

II. PROBLEM DESCRIPTION AND FORMULATION

Consider a bounded search region, \mathcal{A} , which is discretized into $|\mathcal{A}|$ cells. Such cellular partitions of environments may arise naturally due to the structure of the search region, for example, different rooms within a building, or may represent abstract decompositions of the space due to sensing constraints, such as visibility or finite detection radius. In this paper, let \mathcal{A} be a rectangular environment, such that its discrete representation is given by a regular lattice or grid. Note, however, that the approach in this paper applies more generally to any bounded region and for any discrete partitioning of that region.

Denote the target of interest as x , which is assumed to be present in \mathcal{A} with probability, $0 < \delta \leq 1$. The parameter δ is the aggregate or cumulative probability of the target distribution, and represents the searchers' initial "guess" as to the target's presence in the search region prior to commencing the search. This target probability distribution is initially prescribed by prior information, if any is available, on the target's whereabouts. For example, imperfect knowledge of the target's preferences for specific locations or perhaps the physical presence of obstacles or terrain features in the search region can be used to shape the prior target probability distribution. As depicted in Fig. 1,

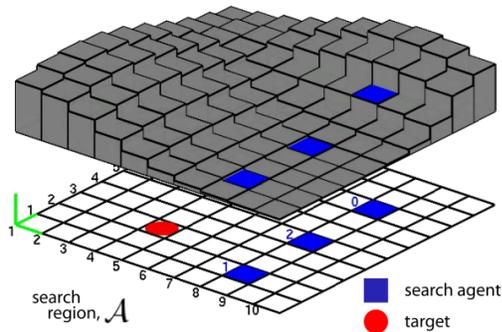


Fig. 1. Illustration of the rectangular search region \mathcal{A} with multiple searchers (blue squares) and a target (red disc). The discrete surface, depicted by the collection of gray bars, represents the assumed prior target probability distribution.

the search region is given by a discrete grid, in which the target may or may not be present. In the numerical studies presented in this paper, the target is assumed to be present and located at position $(i, j) = (4, 3)$. Three searchers are located initially at $s_0 = (8, 8)$, $s_1 = (8, 2)$, and $s_2 = (8, 5)$, equipped with identical detectors with false alarm and missed detection error rates given by, respectively, $\alpha_i = 0.20$ and $\beta_i = 0.10$, for $i = 1, 2, 3$. Their initial aggregate probability value is given by $\delta = 0.50$, which corresponds to a 50% belief that the target is initially within the search region.

The multi-agent search objective posed in this paper is to utilize a team of M search agents, s_i , for $i = 1, \dots, M$, to determine the presence of the target x in search region \mathcal{A} .

This objective can be characterized by a binary hypothesis test – the *search decision* – between hypotheses $H = 1$ and $H = 0$ (i.e., the target's presence or absence in \mathcal{A} , respectively). The searchers' search trajectories through the region are to be designed to reduce the time required to reach a decision. The search objective described above is an example of an *information-gathering task* [12], [6], in that information regarding the presence (and location) of the target is accumulated as search agents maneuver through the search region. The motion of each search agent is prescribed by its respective dynamics model, such that agents are able to either move to another adjacent cell or remain in the current location at each time step.

The searcher uses its sensor to detect the presence or absence of the target within the cell it is investigating. Let $d_i^{(a)}(t) \in \{0, 1\}$ denote the binary detection measurements by search agent i in the a^{th} cell at discrete time t . The team's task is confounded by the fact that observations are imperfect, that is, the agents' detectors may incorrectly register the presence or absence of a target within the observed cell, called *false alarm* and *missed detection* errors, respectively [13].

In much of the autonomous search literature, however, false alarm errors are generally neglected, as the search task in these cases is considered complete upon the first positive detection of the target [14], [8], [2]. Clearly, these approaches would prematurely terminate the search in more realistic scenarios where incorrect positive detections are possible. In contrast, the decision-theoretic framework presented in this paper neatly and easily addresses errors of both kind, capturing the i^{th} search agent's probability of false alarms and missed detections by α_i and β_i , respectively:

$$Pr[d_i^{(a)} | x] : \begin{cases} Pr[d_i^{(a)} = 0 | x = a] &= \beta_i \\ Pr[d_i^{(a)} = 1 | x = a] &= 1 - \beta_i \\ Pr[d_i^{(a)} = 0 | x \neq a] &= 1 - \alpha_i \\ Pr[d_i^{(a)} = 1 | x \neq a] &= \alpha_i. \end{cases}$$

The time index, t , is dropped for succinctness when no ambiguity exists.

Additional notation used in this paper is as follows. The history of detections obtained by the i^{th} search agent from the start of the search till time t is given by $D_i(1:t) = \{d_i(1), \dots, d_i(t)\}$. Denote the set of all M agents' detection histories as $\mathcal{D}(1:t)$, and let $\mathcal{D}(t)$ be the set of all M detections from all search agents at time t , i.e., $\mathcal{D}(t) = \{d_1(t), \dots, d_M(t)\}$.

A. Multi-agent Bayesian Filtering

Probabilistic methods are most appropriate for evaluating the evolution of most information-gathering tasks, as they account for the probabilistic representations of imperfect information. However, traditional tools for recursive estimation of states, such as the Kalman filter, are ill-suited for scenarios requiring the propagation and update of non-Gaussian distributions. Alternatively, the generalized Bayesian filter [15], [16] employs the Chapman-Kolmogorov

and Bayes relations as “prediction” and “correction” steps that recursively incorporate arbitrary distributions.

Fusion of multiple observations in the context of Bayesian filtering has been widely used in a variety of multi-robot applications [17], [10]. The combination of detections from multiple searchers in a Bayesian filter is briefly reviewed in the following section. Additionally, a closed-form expression governing the evolution of the decision is derived, which is attainable due to the structure of the search as a decision. This analytic formula, which aids in the fundamental analysis and efficient implementation of the proposed search strategies, is one of the main contributions of this paper.

Prediction of Target Dynamics: The discrete analog to the Chapman-Kolmogorov equation represents the propagation of the belief probability distribution to account for the target model dynamics, which has the effect of *increasing* the uncertainty in the target’s location and presence in the search region.

$$Pr[x(t) = a | \mathcal{D}(1 : t - 1)] = \sum_m p^{(m,a)} Pr[x(t-1) = m | \mathcal{D}(1 : t-1)],$$

where $p^{(m,a)} \triangleq Pr[x(t) = a | x(t-1) = m]$ represents the discrete state transition probability of the target moving to cell a given that it was previously in cell m . The transition probability matrix given by $P \triangleq [p^{(m,a)}]$ reflects the discrete structure of the region and captures the uncertainty that is present in the environment and the target motions. Such a characterization of the target motion, such as a random walk or even more complex dynamics [18], can often be constructed in a straightforward manner, and in this way, the target’s discrete trajectory through the search region can be seamlessly accounted for in the evolution of the belief distribution.

Update for Multiple Observations: At each time step, the team of searching agents provides M detection measurements, $\mathcal{D}(t)$, to be aggregated and used to update the a^{th} cell belief value. Bayes’ rule can be applied to these conditionally independent observations, such that

$$\begin{aligned} Pr[x(t) = a | \mathcal{D}(1 : t)] &= Pr[x(t) = a | \mathcal{D}(1 : t - 1), \mathcal{D}(t)] \\ &= \frac{Pr[\mathcal{D}(t) | x(t) = a, \mathcal{D}(1 : t - 1)]}{Pr[\mathcal{D}(t) | \mathcal{D}(1 : t - 1)]} Pr[x(t) = a | \mathcal{D}(1 : t - 1)], \end{aligned}$$

where $Pr[x(t) = a | \mathcal{D}(1 : t - 1)]$ is determined from the prediction step, $Pr[\mathcal{D}(t) | \mathcal{D}(1 : t - 1)]$ is the normalizing marginalization of the observations, and

$$\begin{aligned} Pr[\mathcal{D}(t) | x(t) = a, \mathcal{D}(1 : t - 1)] &= Pr[d_1(t), \dots, d_M(t) | x(t) = a, \mathcal{D}(1 : t - 1)] \\ &= \prod_{i=1}^M Pr[d_i(t) | x(t) = a]. \end{aligned}$$

Note that by considering the search in a decision-theoretic context, the likelihood function, $Pr[d_i(t) | x(t) = a]$, represents the general detection model described previously, easily capturing both missed detections and false alarms.

The independence of detections taken in different cells holds true given that observing in one location does not affect another search agent’s observation somewhere else in the search region. Further, simultaneous detection measurements (that is, within the same cell) are assumed to be probabilistically independent in that each observation does not interfere with the other sensors’ respective sensor characteristics (i.e., detection error rates), which is reasonable for practical detection modalities.

III. EVOLUTION OF THE SEARCH DECISION

Given the Bayesian filtering of detection measurements described in the previous section, the search can be readily defined as a decision process. As more information is gathered (in the form of positive and negative detections), the *belief* that the target is present in the search region also evolves with each discrete observation. This sequential notion of decision-making stems from the seminal work by [19], [20] in formulating the Sequential Probability Ratio Test (SPRT), where information is accumulated over time until one of two decision boundaries (corresponding to the acceptance or rejection of the hypothesis) is reached. The discrete observation structure of the SPRT lends itself to the discrete time search application treated in this paper. Note, however, that this is not a limitation, as the mathematically equivalent, continuous analog of the SPRT, called the *drift diffusion model* for decision making [21], can be applied for continuous search problems, and is the subject of future investigation.

The belief function, $\mathbf{B}(t)$, is given by the cumulative mass function of the target distribution, representing the aggregate probability of the target’s presence within the search region:

$$\mathbf{B}(t) \triangleq Pr[H = 1 | \mathcal{D}(1 : t)] = \sum_{a=1}^{|\mathcal{A}|} Pr[x(t) = a | \mathcal{D}(1 : t)].$$

The individual cell belief probabilities are computed using the propagation-update rules of the Bayesian filter described previously, and in practical implementations, the update step is often the computationally limiting component of such filters [10], [17]. A closed-form expression for the belief evolution is presented below.

Analytic expression for the belief evolution: Based on the approach in [12], the formulation for multiple search agents is a straightforward extension of the analysis for single observations by substituting a vector of conditionally independent observations. The aggregate belief function requires the marginalization of the observations, given by:

$$\begin{aligned} Pr[\mathcal{D}(t) | \mathcal{D}(1 : t - 1)] &= \sum_{H=\{0,1\}} Pr[\mathcal{D}(t) | H, \mathcal{D}(1 : t - 1)] Pr[H | \mathcal{D}(1 : t - 1)]. \end{aligned}$$

Once again employing the conditional independence of the M detections, the above expression can be examined by

noting that

$$\begin{aligned} Pr[\mathcal{D}(t)|H=1, \mathcal{D}(1:t-1)] \\ &= Pr[d_1(t), \dots, d_M(t)|H=1, \mathcal{D}(1:t-1)] \\ &= \prod_{i=1}^M Pr[d_i(t)|H=1, \mathcal{D}(1:t-1)], \end{aligned}$$

where the likelihood of an observation given the affirmative hypothesis is

$$\begin{aligned} Pr[d_i(t)|H=1, \mathcal{D}(1:t-1)] \\ &= \frac{\sum_{a'=1}^{|A|} Pr[d_i^{(a)}|x=a', \mathcal{D}(1:t-1)] Pr[x=a'|\mathcal{D}(1:t-1)]}{\sum_{a'=1}^{|A|} Pr[x=a'|\mathcal{D}(1:t-1)]}, \end{aligned}$$

and conversely for the null hypothesis that

$$\begin{aligned} Pr[\mathcal{D}(t)|H=0, \mathcal{D}(1:t-1)] \\ &= Pr[d_1(t), \dots, d_M(t)|H=0, \mathcal{D}(1:t-1)] \\ &\Rightarrow \prod_{i=1}^M Pr[d_i(t)|x \neq a', \mathcal{D}(1:t-1)], \forall a'. \end{aligned}$$

Following [12] and utilizing the above expansions, the modified expression for updating individual cell belief probabilities is:

$$\begin{aligned} Pr[x(t) = a|\mathcal{D}(1:t)] = \\ \frac{Pr[\mathcal{D}(t)|x=a, \mathcal{D}(1:t-1)] Pr[x=a|\mathcal{D}(1:t-1)]}{\left[\frac{Pr[\mathcal{D}(t)|x=a', \mathcal{D}(1:t-1)] Pr[x=a'|\mathcal{D}(1:t-1)]}{+ Pr[\mathcal{D}(t)|H=0, \mathcal{D}(1:t-1)](1-Pr[x=a'|\mathcal{D}(1:t-1)])} \right]}. \end{aligned}$$

Further manipulation leads to the final expression for the aggregate belief evolution:

$$\mathbf{B}(t) = \frac{\Delta(t)}{\Delta(t) + (\alpha)^l(1-\alpha)^{Mt-l}(1-\delta)}, \quad (1)$$

where l is the number of positive ($d=1$) detection observations over all M agents up till time t and

$$\begin{aligned} \Delta(t) \triangleq &\prod_{j=2}^t Pr[\mathcal{D}(j)|H=1, \mathcal{D}(1:j-1)] \\ &\times Pr[\mathcal{D}(1)|H=1] Pr[H=1]. \end{aligned}$$

The intuition behind expression (1) is simply that the belief evolves according to the probability of receiving correct positive detections divided by the total probability of total (correct plus false alarm) positive detections.

Benefit of multi-agents, even without feedback: The use of multiple searching agents is beneficial for search, even in the absence of feedback about the current belief distribution [22]. This fact is simply due to the increased amount of information that is accumulated at each time step, in the form of either multiple simultaneous observations in a single cell or the greater size of the searched area.

The improvement due to multi-agent search in the context of the search decision is evident by observing the reduction in the decision time with increasing number of searchers, as illustrated in Fig. 2 for open-loop strategies such as

random walking search and parallel sweeping search. In these nonadaptive settings, the search performance is tightly related to the re-visitation rate, or the number of times the target cell is revisited, during the course of the search. In this manner, the decision *rate*, that is, how quickly the decision evolves, is an interesting aspect of this search-as-a-decision problem to be investigated in future efforts.

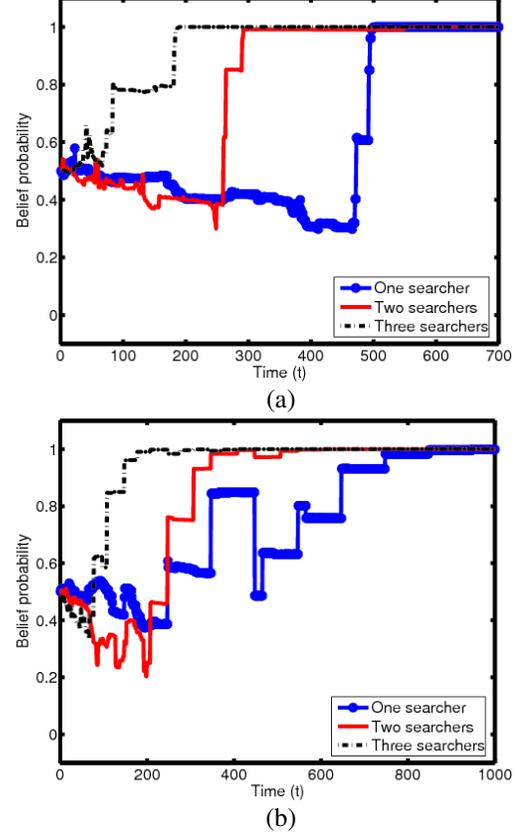


Fig. 2. Comparison in illustrative simulation trials of search decision performance improvement for increasing number of search agents employing (a) random walking, or (b) parallel sweeping search. Simulations were conducted using the setup illustrated in Fig. 1.

IV. MULTI-AGENT SEARCH STRATEGIES UNDER FEEDBACK

The use of feedback about the current belief distribution to generate updated search trajectories greatly improves the performance of the search task [12], and is known as adaptive or semiadaptive search [2]. The optimization of the search on a grid is provably complex [23], and the computational cost of approximations must be considered in their evaluations.

Implementations of (semi)adaptive strategies intuitively will outperform those of naïve ones, such as the random walk or parallel searches illustrated previously, as they employ updated information in computation of search paths. The use of multiple agents amplifies this advantage in reducing the time until the search decision is made as compared to single agent search [12]. The effect of increasing numbers of more simultaneous searchers is shown in Fig. 3.

Hybrid Search Teams: An additional advantage of search with multiple agents is that the search task can be dis-

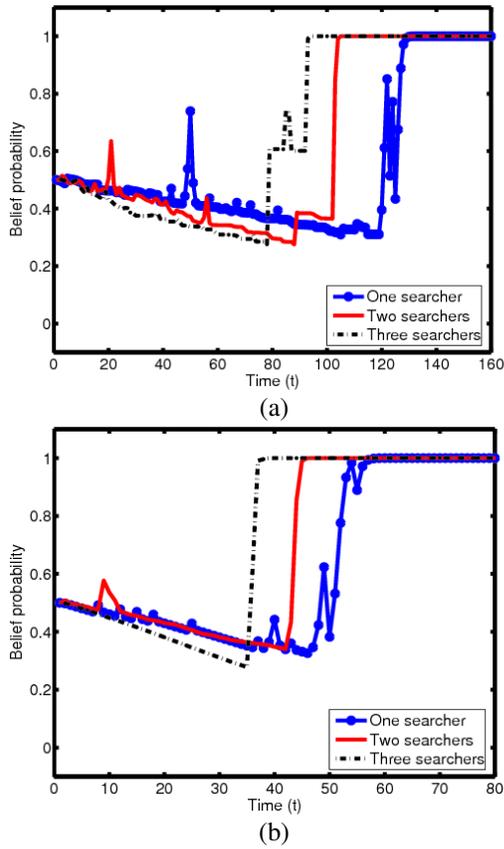


Fig. 3. Performance in illustrative simulation trials of (a) “*Drosophila*-inspired” and (b) “saccadic” search strategies using multiple search agents. Increasing the number of searchers reduces the time till a decision is made.

tributed heterogeneously, that is, different agents can employ different strategies for generating their respective search plans. A wide range of applications call for deployments of hybrid teams with varying detection and motion capabilities. For example, fast-moving unmanned aerial vehicles, which typically yield low to mid-resolution information, can be combined with slower but higher resolution, ground-based counterparts to search for, inspect possible hits, and eliminate identified hostile targets. In the decision-theoretic framework developed thus far, the overall search decision task remains the same (i.e., “Is the target present?”), and the evolution of the belief also easily incorporates detections from differing information-gathering modalities. Note that the heterogeneity may be manifest at the sensing level with varying detection error rates (e.g., different sensor suites onboard different robots) or at the strategic level with task- or resource-dependent assignment of search trajectories (e.g., fast-moving UAV’s handle broad search while slower UGV’s perform inspections). Hybrid search in the latter case is demonstrated via the following example.

Consider the application example of a building under surveillance by a distributed sensor network, where each cell or room in the complex has a camera. A security guard is able to choose which camera is turned on and displayed on his screen, while another guard patrols on foot to areas of potential intrusion. One agent using the “saccadic” strategy,

where observations are taken in the cell containing maximal belief, can be combined with another search agent employing the “*Drosophila*-inspired” search strategy from [12] (in which a searcher only updates its search plan periodically) to model the search of the informed walking guard. The simulation for this example hybrid search team is shown in Fig 4. The two search agents depicted in 4 conduct a

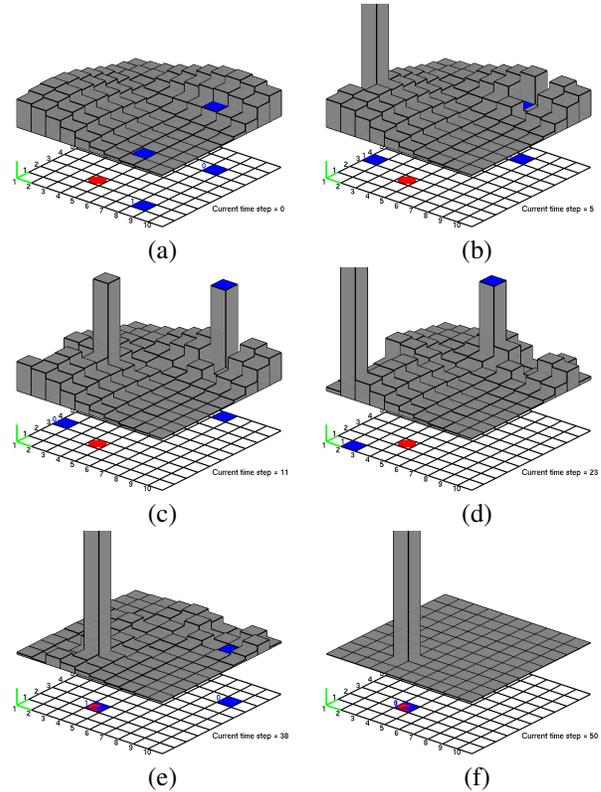


Fig. 4. Simulation snapshots of a hybrid search team’s path through the search region. Searchers 0 and 1 employ the *Drosophila*-inspired and “saccadic” search strategies, respectively. (a)-(d) Search agent 1 is able to jump to cells with high belief, while search agent 0 maneuvers through the region. (e) Searcher 1 has identified the cell likely to contain the target, and (f) search agent 0 converges on this cell upon updating its search plan.

distributed search of the region, in that their observations are used jointly to update the target belief distribution. Search agent 1, not limited by motion constraints, ensures that disjoint sections of the search region are observed, whereas its teammate augments the search by exploring cells along its search trajectory. This combination of search methodologies provides a balance between the narrow but fast search of the “saccadic” strategy and the slower yet more exploratory nature of the “*Drosophila*-inspired” one.

An illustration of this combined benefit can be seen in the belief evolution for the hybrid team, shown in Fig. 5, against the performance of homogeneous search teams using one of either the “saccadic” or the “*Drosophila*-inspired” strategies. Given identical simulation settings (i.e., initial starting locations and agent detection error rates), the two-agent hybrid search strategy outperforms the comparable teams employing only a single strategy. This (perhaps surprising) result is caused by the fact that, in the latter cases of homogeneous teams, the search agents begin to examine

the same cells at the same time. Although this may improve the combined “virtual” detector at this location, the added benefit of observing different cells is lost in these cases.

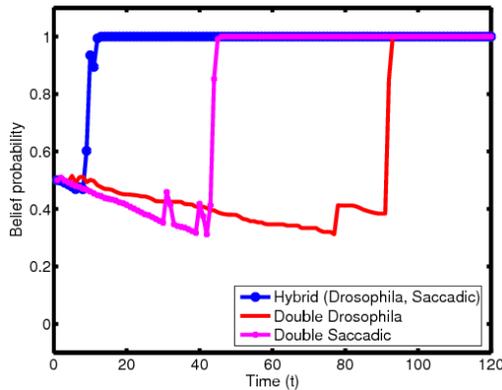


Fig. 5. Comparisons from illustrative simulation trials of the performance of the hybrid search using saccadic and *Drosophila*-inspired strategies, compared to the respective single strategy search teams.

V. CONCLUSIONS AND FUTURE WORK

This paper investigated the problem of searching for a dynamic target in a given search region using multiple search agents. Extending the decision-theoretic framework of [12] to incorporate a team of searchers, the advantage of distributed search was demonstrated by reducing the time until a decision regarding the target’s presence or absence is made. Further, the ability to employ hybrid search strategies, where search agents use different methods for generating their search plans, highlighted another benefit of the formulation presented herein. Such hybrid searches can be useful for heterogeneous teams, including human-in-the-loop applications.

Avenues for further research are abundant, including extension to multi-target scenarios, such as in [10]. In this case, more sophisticated target tracking [24] and task assignment [25] methodologies may be required. Also, for situations where a specific target is of interest among many false targets, an additional recognition or identification phase is necessary to distinguish between desired and decoy objects [26], [27]. This feature might be particularly relevant for security or surveillance applications in practical settings (such as identification of a person in a crowded airport). The sequential decision-theoretic framework presented in this paper offers the possibility to address these *two stage* decision problems by introducing intermediate decision thresholds corresponding to the *explore* and the *exploit* search modes.

Another possibly illuminating direction is to employ graph-theoretic representations of the search region, such that connections between the search as a decision, as posed in this paper, and search on graphs can be made. The wealth of available tools and relevant issues, such as re-contamination of previously searched cells [28], may provide a way, coupled with the decision-theoretic framework, to generalize the discrete search problem described in this paper.

REFERENCES

- [1] B. O. Koopman, “The theory of search, part III. the optimum distribution of searching effort,” *Operations Research*, vol. 5, pp. 613–626, 1957.
- [2] L. D. Stone, *Theory of Optimal Search*, 2nd ed., J. Kettelle, Ed. Academic Press, 1989.
- [3] A. R. Washburn, *Search and Detection*, 4th ed., ser. Topics in Operations Research Series. INFORMS, 2002.
- [4] S. J. Benkoski, M. G. Monticino, and J. R. Weisinger, “A survey of the search theory literature,” *Naval Research Logistics*, vol. 38, pp. 469–494, 1991.
- [5] S. Alhakeem and P. Varshney, “Decentralized bayesian detection with feedback,” *IEEE Trans. on Systems, Man and Cybernetics, Part A*, vol. 26, no. 4, pp. 503–513, July 1996.
- [6] M. Strens, “Learning Multi-Agent Search Strategies,” *AISB Journal*, vol. 1, no. 4, pp. 7–11, 2003.
- [7] J. Najemnik and W. S. Geisler, “Optimal eye movement strategies in visual search,” *Nature*, vol. 434, pp. 387–391, 2005.
- [8] N.-O. Song and D. Teneketzis, “Discrete search with multiple sensors,” *Mathematical Methods of Operations Research*, vol. 60, no. 1, pp. 1–13, September 2004.
- [9] T. Furukawa, F. Bourgault, B. Lavis, and H. Durrant-Whyte, “Recursive Bayesian Search-and-Tracking Using Coordinated UAVs for Lost Targets,” in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, 2006, pp. 2521–2526.
- [10] E. M. Wong, F. Bourgault, and T. Furukawa, “Multi-vehicle bayesian search for multiple lost targets,” in *Proc. of Intl. Conf. on Robotics and Automation*, April 2005, pp. 3169–3174.
- [11] R. Vidal, O. Shakernia, H. Kim, D. Shim, and S. Sastry, “Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation,” *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 662–669, Oct 2002.
- [12] T. H. Chung and J. W. Burdick, “A decision-making framework for control strategies in probabilistic search,” in *Proc. of IEEE Intl. Conference on Robotics and Automation*, 2007.
- [13] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. John Wiley & Sons, Inc., 2001.
- [14] J. Kadane, “Optimal Whereabouts Search,” *Operations Research*, vol. 19, no. 4, pp. 894–904, 1971.
- [15] S. Thrun, W. Burgard, , and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [16] A. Jazwinski, *Stochastic processes and filtering theory*. Academic Press New York, 1970.
- [17] M. Rosencrantz, G. Gordon, and S. Thrun, “Decentralized sensor fusion with distributed particle filters,” in *Proc. Conf. on Uncertainty in AI*, 2002.
- [18] F. Bourgault, T. Furukawa, and H. Durrant-Whyte, “Process model, constraints, and the coordinated search strategy,” in *Robotics and Automation, 2004. Proceedings*, vol. 5, 2004, pp. 5256–5261.
- [19] A. Wald, *Sequential Analysis*. New York, NY: John Wiley and Sons, Inc., 1947.
- [20] A. Wald and J. Wolfowitz, “Bayes solutions of sequential decision problems,” *PNAS*, vol. 35, pp. 99–102, 1949.
- [21] R. Ratcliff, “A theory of memory retrieval,” *Psychological Review*, vol. 85, no. 2, pp. 59–108, 1978.
- [22] P. Vincent and I. Rubin, “A framework and analysis for cooperative search using UAV swarms,” in *Proceedings of the 2004 ACM symposium on Applied computing*. ACM Press NY, NY, 2004, pp. 79–86.
- [23] K. E. Trummel and J. R. Weisinger, “The complexity of the optimal searcher path problem,” *Operations Research*, vol. 34, no. 2, pp. 324–327, Mar. - Apr. 1986.
- [24] D. B. Reid, “An algorithm for tracking multiple targets,” *IEEE Tran. on Automatic Control*, vol. 24, no. 6, pp. 843– 854, Dec 1979.
- [25] V. Isler, S. Khanna, J. Spletzer, and C. Taylor, “Target tracking with distributed sensors: The focus of attention problem,” *Computer Vision and Image Understanding Journal*, vol. Special Issue on Attention and Performance in Computer Vision, p. accepted, 2005, accepted.
- [26] R. Hohzaki, “Discrete search allocation game with false contacts,” *Naval Research Logistics*, vol. 54, no. 1, pp. 46–58, 2006.
- [27] D. Kalbaugh, “Optimal Search Density for a Stationary Target among Stationary False Targets,” *Operations Research*, vol. 41, no. 2, pp. 310–318, 1993.
- [28] A. LaPaugh, “Recontamination does not help to search a graph,” *Journal of the ACM (JACM)*, vol. 40, no. 2, pp. 224–245, 1993.