

Bayesian Forecasting in Multi-vehicle Search Operations

L. F. Bertuccelli* and J. P. How†

Aerospace Controls Laboratory
Massachusetts Institute of Technology
{lucab, jhow}@mit.edu

This paper discusses robust UAV search operations in the context of decision-making under uncertainty and presents a new metric called the modified Bayes Factor (MBF) that encompasses the notion of uncertainty and can be used for robust planning. Various results are shown with the MBF embedded in our planning schemes. This paper also presents a new approach to predict the value of future information by evaluating the most likely future measurements by using Bayesian forecasting, and using these pseudomeasurements to calculate the expected MBF or reduction in variance if additional observations were available. Numerical results are shown that demonstrate the improvement that can be obtained when comparing these robust approaches to nominal ones in a poorly known environment. Key benefits are also demonstrated with the use of Bayesian forecasting in predicting the most likely locations for improvement in information.

I. Introduction

This paper presents methods for performing robust search task allocation under uncertainty, and an approach to predict the most likely future information in search tasks for autonomous agents. The objective is to assign N_V agents ($N_V \gg 1$) to regions in space to maximize the expected reward of finding N_T possibly moving targets. The number of targets is *a priori* unknown by the agents, and there is some uncertain prior information that can be used to *direct* the search. Previous research has developed techniques for modeling this uncertainty,^{6,7} but had not fully addressed the robust vehicle tasking. This paper discusses various robust tasking approaches and demonstrates their use in multi-vehicle search operations.

In the search theoretic literature, the search problem is tackled by discretizing the environment in N_{cell} cells, where each cell is described by a prior probability of containing a target.¹⁸ The agents use these probabilities to maintain an *information map*, $\mathcal{IM}_{k|k}$, and are initially assigned to cells in space based on this prior information. As they move throughout the environment, the agents update their information map by taking observations in the cells they are overflying. The observation at time k , Y_k , is typically modeled as “target detected” ($Y_k = 1$) or “no target detected”

*Research Assistant, Dept. of Aeronautics and Astronautics

†Associate Professor, Dept. of Aeronautics and Astronautics

($Y_k = 0$), and is a function of the error characteristics of the sensor and the true cell occupancy state (i.e., whether a target truly is in a cell or not). The simplest task allocation problem can thus be expressed as

$$\max_{x_k \in X_k} f(\mathcal{IM}_{k|k}, B_k, x_k) \quad (1)$$

where x denotes the binary decision variables of which agent is assigned to which region, $\mathcal{IM}_{k|k}$ denotes the information map at time k given all the measurements at time k , and B_k denotes additional parameters that may be used in the optimization, such as distance measures to targets or target velocities. Additionally, the decision variables are constrained to lie in the set \mathcal{X}_k , and typical constraints include total vehicle availability at time k . Note that, similar to the path planning problem, this is a typically a time-varying scheme, and the allocations may be re-optimized at distinct times, hence the indexing of the decision variables by time k . Typical assignments will generate a *sequence* of target lists (in this case, of cells to visit) at each time the optimization is run. These waypoints are then passed on to a lower level path planner that will generate optimal paths to those waypoints.

An example of a full mission segment is shown in Figure 1, in which 3 UAVs are exploring a poorly known environment with 8 targets. The waypoints are generated by the task assignment algorithm, and these waypoints are passed on to path planning algorithms that optimizes the trajectories to these points. The probability map given in the figure shows the nominal probabilities of target existence in each of the cells. In realistic simulations, this probability map is constantly changing depending on the sensor measurements obtained by the vehicles, and these peaks may in general diffuse in the environment over time.

Due to ambiguity or conflicting information, resource constrained search operations must be made robust to the uncertainty in the information maps. In most search-theoretic literature, the information maps are represented by probability maps, or occupancy maps.^{8,9,10,13,14,16,18,19,20,21} These probability maps are a gridded representation of the environment, and each cell is described by the probability of containing a target, P_j . These probabilities are generally treated as being perfectly known, and the tasking algorithm selects the cells with the highest likelihood of containing a target, and allocates the vehicles to those regions of space. This is done because the value of each cell is proportional to the probability of the cell containing a target. In general, however, the probabilities may not be precisely known due to sensor ambiguity or conflicting information from multiple sources. Hence, the cell's value cannot just be related to the expected value of the probability, but must also take into account higher moment information that describes the uncertainty, or ambiguity, of the probability.^{6,7} Various ways of describing this uncertainty include variance, and entropy. This paper shows how to embed uncertainty in the task allocation using techniques from robust optimization.

At each time step $k + 1, k + 2, \dots$, the measurements taken by the agents will update the information map to $\mathcal{IM}_{k+1|k+1}, \mathcal{IM}_{k+2|k+2}, \dots$; these information maps can then be used to re-allocate agents based on current information. However, it has been shown that performance can be improved by incorporating estimates of future information.^{5,15,17,21} In particular, a look-ahead heuristic can be used that forecasted future observations based on previous measurements;²¹ these

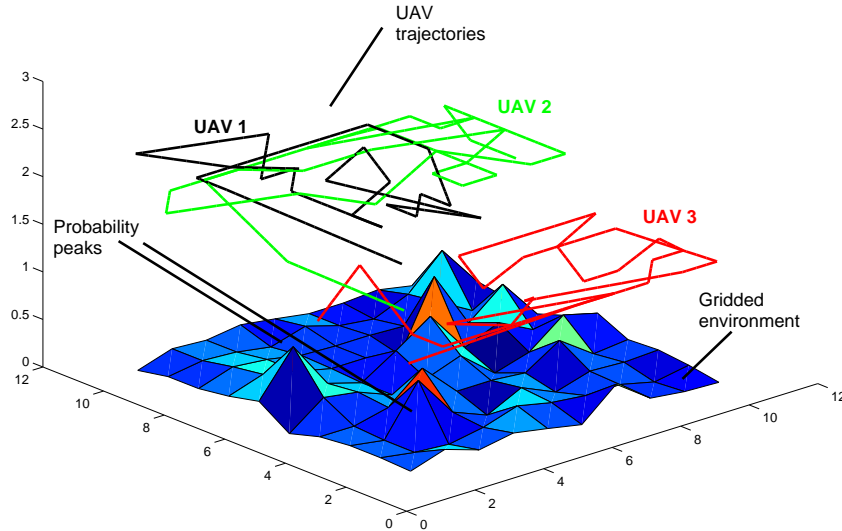


Fig. 1: Sample UAV Mission showing 8 targets found. The task allocation algorithm directed the 3 UAVs to the locations of the environment with the highest peaks and the path planner resulted in the following trajectories. The information map shows approximately 8 probability peaks (regions where the targets are most likely to exist) overlaid on the gridded environment.

future observations can then be used to predict the expected reduction in uncertainty of visiting a cell. The aforementioned results assumed that the probabilities were known precisely, while in this work we allow for uncertainty in these probabilities. Furthermore, in our numerical examples, we investigate the case of moving targets with a motion model we have discussed in a previous paper.⁷

The inference of a future measurement can be explicitly calculated using the *predictive distribution* of a future measurement, and embedding the most likely observation as a *pseudomeasurement* in calculating the expected uncertainty reduction of visiting a particular region. Equivalently, we are interested in finding the maximum likelihood measurement sequence, $\{\hat{Y}_{k+1|k}^*, \dots, \hat{Y}_{k+T|k}^*\}$, such that it maximizes the predictive distribution given all the previous measurements $\mathcal{I}_{k|k} = \{Y_1, Y_2, \dots, Y_k\}$

$$\{\hat{Y}_{k+1|k}^*, \dots, \hat{Y}_{k+T|k}^*\} = \arg \max_Y P(Y_{k+1}, \dots, Y_{k+T} | \mathcal{I}_{k|k}) \quad (2)$$

For simplicity, consider the case where the forecasting is over 1 horizon ($T = 1$). The measurement $\hat{Y}_{k+1|k}^*$ that maximizes Eq. 2 is the most likely future measurement. Treating this as a pseudomeasurement, it can then be used to evaluate the expected uncertainty reduction if the measurement were actually available. For example, the pseudomeasurement can be used to calculate the entropy of the forecasted probability density with the current pdf thereby indicating which cell has the greatest reduction in uncertainty, or the change in variance of a particular region in the environment. While solutions exist for one-step forecasts for the particular uncertainty model used in this paper, we also develop expressions for multi-step forecast.

This paper is organized as follows. Section II describes the problem statement and provides the necessary motivation for coming up with a predictive algorithm that can forecast future measurements. Section III discusses different robust allocation techniques. Section IV shows the derivation

for the 1-step and multi-step Bayesian forecasting scheme, and numerical results are shown in Section V.

II. Problem Description

The environment used in this work is a continuous space that is discretized in a finite number of uniform cells, where each cell is described by the probability of containing a target. The UAV sensors are modeled by a Bernoulli model that outputs $Y_k = 1$ if a target is detected, and $Y_k = 0$ if a target is not detected at time k ,

$$P(Y_{k+1}|p) = Mp^{\gamma_1-1}(1-p)^{\gamma_2-1} \quad (3)$$

where γ_1 (γ_2) denote the number of times a target is detected (and not detected, respectively), and M is a normalizing constant that ensures that the pdf integrates to 1. This type of model is a simplified but reasonable method to approximate the image detection process. The intricacies of the image processing are abstracted away by treating the outputs as binary measurements with an error probability that is representative of the image classification errors. These errors are the correct detection probability (ξ) and false alarm probability ($1 - \chi$), and are assumed to be known.

A. Nominal Formulations

Nominal search algorithms rely on making the UAV allocation based on the probability P_j of a target being in a cell j . A simple tasking algorithm that allocates vehicles $x \in \mathcal{X}_k$, at each time step k , to regions in space with the highest probabilities assigns a value, $V_j = C(P_j)$ of finding a target, where C is a function that may embed notions of distances between targets and UAVs in the optimization. The optimization thus takes the form

$$J_{NOM} : \max_{x \in \mathcal{X}_k} \sum_j C(P_j)x_j \quad (4)$$

B. Robust Formulations

There are cases however, when the probabilities P_j are uncertain, and can be assumed to lie in some uncertainty set \tilde{P}_j . The uncertainty in the probabilities may be due, for example, to conflicting prior information, user subjectivity, target motion, or simply poor knowledge. For example, an agent may believe from his information that a target is in cell i with probability 0.7, while another agent may believe the probability is actually 0.3, and hence there is implicit uncertainty in these distributions. In these cases, it is thus desirable to model this probability $p = P_j$ by a probability density function (pdf) that captures the uncertainty in our knowledge of the correct probability in each cell. This work uses the Beta pdf,

$$P(p) = Kp^{\zeta_1-1}(1-p)^{\zeta_2-1} \quad (5)$$

where ζ_1 and ζ_2 are previous ‘‘counts’’ that indicate how many times a target may have been detected (and not detected) in a previous reconnaissance mission with heterogeneous sensors, for example. K is a normalizing factor that ensures the pdf integrates to 1. The Beta pdf is chosen in this research for several reasons: *i*) the support of the Beta distribution is $[0, 1]$, and *ii*) the Beta distribution is conjugate with the sensor model used in this paper.^{6,7}

A straightforward technique to perform the optimization when the P_j is uncertain is to use the expected value and solve the nominal assignment problem

$$J_{EXP} : \left\{ \max_{x \in \mathcal{X}_k} \sum_j C(\bar{P}_j) x_j \right\} \quad (6)$$

Since the uncertainty (described by an uncertainty model \tilde{P}_j) can negatively impact the performance of the optimization, the goal of the robust optimizations are to hedge against this uncertainty. Hence, the robust decision problem can be formulated as

$$J_{ROB} : \left\{ \min_{P_j \in \tilde{P}_j} \max_{x \in \mathcal{X}_k} \sum_j C(P_j) x_j \right\} \quad (7)$$

In this type of optimization, cells are selected whose worst-case probability values (in the set \tilde{P}_j) will not affect the performance of the overall search operation. Rather than looking at the nominal value, the choice of the set \tilde{P}_j , or realizations within that set, can be varied to add higher moment information to the optimization. Typically, optimizations that only look at the worst-case tend to be overly conservative, as the worst case is not very likely to occur; hence, techniques that can *tune* the level of conservatism to reduce the degree of conservatism easily are highly desirable.

A tunable approach that is applicable in this framework was introduced in some of our earlier work.⁵ We formulated the robust optimization by effectively penalizing the nominal score, \bar{P}_j by a fraction of the standard deviation, σ_j . A tuning parameter α is used to embed higher or lower degrees of uncertainty, and the overall problem can be cast as

$$J_{ROB,\alpha} : \left\{ \max_{x \in \mathcal{X}_k} \sum_j (\bar{P}_j - \alpha \sigma_j) x_j \right\} \quad (8)$$

The degree of conservatism of the solution relies heavily on the choice of α . For an environment where a high degree of uncertainty exists, α can be chosen to be approximately 2-3, to capture the greatest deviations in the target probabilities. Note that for the Beta distribution as defined in Eq. 5, the mean and variance can be evaluated directly,

$$\begin{aligned} \bar{P}_j &= \frac{\zeta_1}{\zeta_1 + \zeta_2} \\ \sigma_j^2 &= \frac{\zeta_1 \zeta_2}{(\zeta_1 + \zeta_2)^2 (\zeta_1 + \zeta_2 + 1)} \end{aligned} \quad (9)$$

Other approaches suggested in literature^{3,4} place a budget on the total uncertainty, $U_j = U(\tilde{P}_j)$, in the problem. This can be accomplished, for example, by selecting cells with the highest expected probabilities and adding a constraint, Γ , on the total change in entropy desired

$$J_{ROB,Budget} : \left\{ \max_{x \in \mathcal{X}_k} \sum_j C\bar{P}_j x_j \mid \sum_j U_j x_j \leq \Gamma \right\} \quad (10)$$

The key requirement for this approach is that the uncertainty U_j must be evaluated and we can show that for the Beta distribution such an expression can be expressed in closed form.

Alternative metrics that capture both expected value and uncertainty, like the Modified Bayes Factor (introduced in the next section) can also be used to achieve the objective of assigning the vehicles robustly to targets in the environment. For example, by introducing a new metric $J_j = J(\tilde{P})$, then an alternative optimization is one such that

$$J_M : \left\{ \max_{x \in \mathcal{X}_k} \sum_j J_j x_j \right\} \quad (11)$$

An advantage of finding such a metric is that it could reduce the amount of tuning parameters in the constrained optimization of Eq. 10. Furthermore, a single metric that captures both expected value and uncertainty has the more appealing property that it may be simpler to use real-life operations. Such a metric is the topic of the next section.

III. Uncertainty Measures

A. Modified Bayes Factor

The Modified Bayes Factor is a modification of the conventional Bayes Factor (see for example, Ref.¹¹). The Bayes Factor is used to compare posterior to prior odds of a pdf for two different hypotheses,

$$BF = \frac{P(y|H_1)}{P(y|H_0)} = \frac{\int_p P(p|H_1)P(y|p, H_1)dp}{\int_p P(p|H_0)P(y|p, H_1)dp} \quad (12)$$

The two hypotheses are H_0 and H_1 are the hypotheses representing that the target is not in the cell, or is in the cell, respectively. $P(p|H_i)$ is the prior pdf for the probability of target existence under each hypothesis, while $P(y|p, H_i)$ is the sensor likelihood.

Unlike the generalized likelihood ratio test which maximizes the likelihood, the Bayes factor smooths over the uncertain parameter p . It is appropriate for this research since p is treated as a random variable indicating the probability of target existence. However, rather than integrating over the entire support of the parameter p , we modify the Bayes Factor and come up with a Modified Bayes Factor (MBF), which integrates over particular regions P_i of the support,

$$MBF = \frac{\int_{p \in P_1} P(p|H_1)P(y|p, H_1)dp}{\int_{p \in P_0} P(p|H_0)P(y|p, H_1)dp} \quad (13)$$

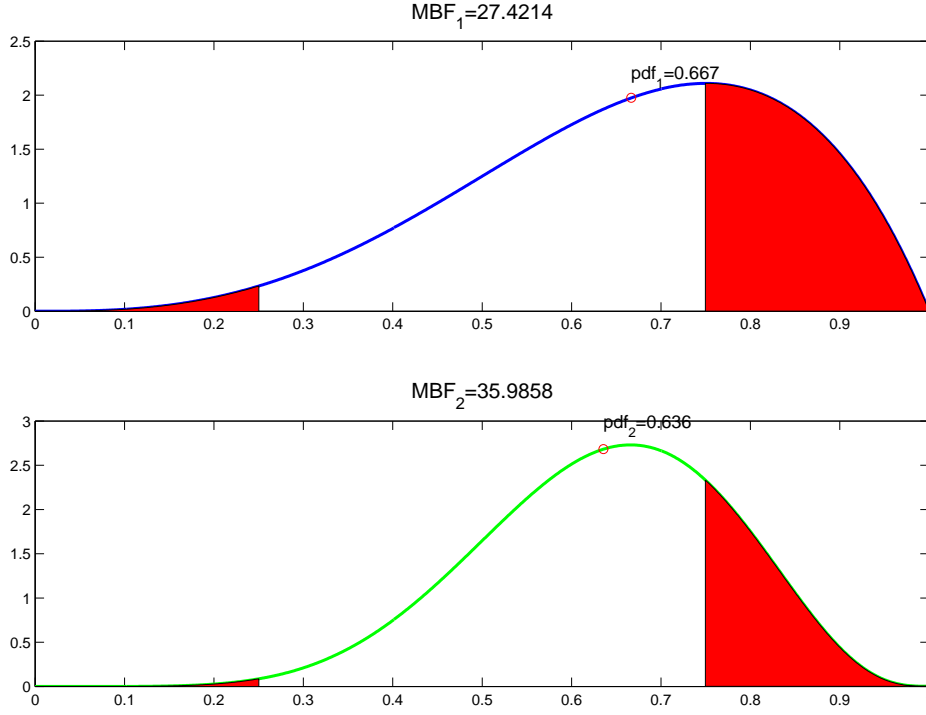


Fig. 2: Figures depicting two different Beta distributions with different prior counts. Both distributions have similar mean value, but clearly different mass distributions. For this case, the regions of Eq. 13 are $P_0 = [0, 0.25]$ and $P_1 = [0.75, 1]$.

By substituting in for the individual pdfs, $P(p|H_1)$ is the prior probability of Eq. 5; $P(y|p, H_1)$ is the likelihood model of Eq. 3. By substituting in these quantities, we obtain

$$MBF = \frac{\int_{p \in P_1} p^{\alpha_1 - 1} (1 - p)^{\alpha_2 - 1} dp}{\int_{p \in P_0} p^{\beta_1 - 1} (1 - p)^{\beta_2 - 1} dp} \quad (14)$$

Intuitively, the MBF is integrating the pdf's mass over different regions of the support. Integrating over the P_0 region in effect integrates over the region where a target should not exist, while integrating over P_1 integrates over the region where a target should exist. The quotient of the two quantities is a ratio of the masses of the two regions. If there is more evidence that a target exists (for example, by having more measurements indicating a target is present), the MBF will be very large, while if more evidence is that a target does not exist, the MBF will be very small. A convenient way to use the MBF is by using its *logarithm*, namely

$$LMBF = \log_{10} MBF = \log_{10} \left(\frac{\int_{p \in P_1} p^{\alpha_1 - 1} (1 - p)^{\alpha_2 - 1} dp}{\int_{p \in P_0} p^{\beta_1 - 1} (1 - p)^{\beta_2 - 1} dp} \right) \quad (15)$$

This is convenient because the value of a cell whose MBF is less than one (meaning that there is less evidence of a target being there), will have negative value. Since our objective is to allocate vehicles to regions that can increase the objective of the optimization, the regions with negative value will not be visited since there is evidence against a target being there.

Both the numerator and denominator of Eq. 14 cannot be evaluated in closed form, since they are not integrating over the entire support $p \in [0, 1]$, but rather over only part of the support $p \in \mathcal{P}_0$. However, such integrals can be computed by using the incomplete Beta function.¹

A brief example with the MBF will show its use for UAV planning in a simple problem. Take the choice between visiting two cells, where the first cell is described by a Beta(5, 3.1) and the second by a Beta(2, 1.2) distribution (in this case it is assumed that a target truly exists in the cell, but a poor sensor is used thereby creating seemingly ambiguous observations). In the case of allocating a single UAV to visit each cell, cell 1 would be preferred in the case of nominal planning since it has a higher expected value (0.63 to 0.62). However, the second cell has a higher total number of observations, and should be chosen since there are fewer conflicting measurements (and thus lower variance). The MBF calculation shows that in this case, cell 2 would indeed be picked since $MBF_2 > MBF_1$.

B. MBF vs. nominal

We compare the results obtained with the MBF and compared to the nominal objective of Eq. 4. There were 16 total static targets in an environment discretized by 25 cells, and the number of UAVs was varied from 10 to 20 in intervals of 2. The prior counts (b, c) for the Beta distribution in each cell were varied randomly in 500 Monte Carlo randomized runs, consistent with whether a target was present or not in that cell. For instance, if a target truly was present in a cell, a sample prior count (b, c) for that cell could have been (4, 3), while if no target was present a sample prior count was (3, 5). The goal of these simulations was to evaluate the MBF in situations with little and poor information, and hence the prior counts were randomized between 1 and 5.

The MBF was evaluated against a nominal scheme, which only takes into account the expected value $b/(b+c)$. In the nominal scheme, the UAVs were assigned to the regions with highest expected value (i.e., the optimization of Eq. 4), while in the MBF approach, the UAVs were assigned to cells with the highest Modified Bayes Factor (calculated using Eq. 13).

The UAVs were assigned in one cycle, and at each run the UAVs were assigned each to one cell. At each run, the total number of cells that the UAVs were assigned that actually contained a target were tabulated. The following quantities were evaluated to compare the two different approaches (where T_i is the total number of targets found in run i),

$$\begin{aligned}\bar{N} &= \frac{1}{N_{TOT}} \sum_{i=1}^{N_{TOT}} T_i \\ \sigma_N^2 &= \frac{1}{(N_{TOT} - 1)} \sum_{i=1}^{N_{TOT}} (T_i - \bar{N})^2 \\ N_{\min} &= \min_i T_i\end{aligned}\tag{16}$$

The results are shown in Table 1. The first column shows the number of UAVs used, the second column shows the average values, and the third and fourth columns show the standard deviation and worst-case number of targets found. The nominal results are in the left column, while the

Table 1: Comparison of MBF to Nominal

Num Vehicles (MBF)	\bar{N}	σ_N	N_{\min}	Num Vehicles (Nom)	\bar{N}	σ_N	N_{\min}
12	8.8	1.5	7	12	7.8	1.6	5
14	11.0	1.3	9	14	9.1	1.4	7
16	12.9	1.0	11	16	10.8	1.8	8
18	14.1	1.3	12	18	12.6	1.3	10
20	14.5	1.0	11	20	13.9	0.9	11

MBF results are in the right column. Note that by using the MBF, two more targets are found on average, and the worst-case for each scenario is improved by finding two more targets than nominal. These improvements occur because the MBF captures both the expected value and variance of the underlying Beta pdfs, and assigns the vehicles to the cells that have the greatest expected value and lower uncertainty. The nominal assignment only assigns the vehicles to the cells with the highest expected value, which may not be a robust decision in the absence of little and/or conflicting information.

IV. Predictive Distribution

While the optimizations in the previous sections can be shown to result in more robust decisions, recent UAV search results²¹ have shown that making predictions of future measurements and evaluating the expected reduction in uncertainty may have beneficial impacts on the search operation. Thus, the second main result of this paper investigates the use of the predictive distribution in estimating the *most likely* future measurements given all previous measurements obtained in a cell. In particular, the one-step forecasting for pdfs involving the Beta distribution are shown and extended to a multi-step forecasting in the next section, where the probability of obtaining a string of $N \gg 1$ future measurements will be developed.

A. Single Step Prediction

The single step predictive distribution allows the calculation of the most likely measurement that can be obtained, given a string of previous measurements. It is of key importance that this single step distribution can be calculated in closed form for the Beta distribution. This results is shown next.

Proposition 3 *The 1-step predictive distribution of a Beta distribution, $P(Y_{N+1} = y_j | Y_1, Y_2, \dots, Y_N)$ is proportional to the expected value of the distribution at time N , $P(Y_N | Y_1, Y_2, \dots, Y_{N-1})$.*

Proof: Define the following quantities

$$\begin{aligned}
 P(p) &= L_p p^{\alpha_0 - 1} (1 - p)^{\beta_0 - 1} && \text{Prior Distribution} \\
 P(Y_1, Y_2, \dots, Y_N | p) &= K_p p^{\alpha_1} (1 - p)^{\beta_1} && \text{Measurement Likelihood}
 \end{aligned}$$

The predictive distribution describes the probability of obtaining measurement $Y_{N+1} = y_j$ given all the previous measurements, where $y_j = \{0, 1\}$,

$$P(Y_{N+1} = y_j | Y_1, Y_2, \dots, Y_N) \quad (17)$$

This is equal to the quotient of the joint pdfs

$$P(Y_{N+1} = y_j | Y_1, Y_2, \dots, Y_N) = \frac{P(Y_1, Y_2, \dots, Y_N, Y_{N+1} = y_j)}{P(Y_1, Y_2, \dots, Y_N)} \quad (18)$$

So find the individual densities individually by marginalizing out the parameters p

$$P(Y_1, Y_2, \dots, Y_N) = \int_p P(Y_1, Y_2, \dots, Y_N | p) P(p) dp \quad (19)$$

The details of this rely on the following fact of the Beta distribution

$$\frac{\Gamma(b_1)\Gamma(b_2)}{\Gamma(b_1 + b_2)} = \int p^{b_1-1}(1-p)^{b_2-1} dp$$

Then, the marginalization process can be simplified to

$$P(Y_1, Y_2, \dots, Y_N) = K_p L_p \frac{\Gamma(\alpha_0 + \alpha_1)\Gamma(\beta_0 + \beta_1)}{\Gamma(\alpha_0 + \alpha_1 + \beta_0 + \beta_1)} \quad (20)$$

To integrate the term with the new measurement

$$P(Y_1, Y_2, \dots, Y_{N+1} = y_j) = \int_p P(Y_1, Y_2, \dots, Y_{N+1} = y_j | p) P(p) dp \quad (21)$$

The same mathematical machinery follows from above identically except for the term where the measurement will be forecasted.

$$P(Y_1, Y_2, \dots, Y_{N+1} = y_j) = M_p L_p \frac{\Gamma(\alpha_0 + \alpha_1 + q)\Gamma(\beta_0 + \beta_1 + r)}{\Gamma(\alpha_0 + \alpha_1 + \beta_0 + \beta_1 + q + r)} \quad (22)$$

where $q + r = 1$ implying that one measurement is forecasted, and M_p is the normalizer for the distribution. Substituting for these equations (see Eq. 18), the simplified quotient is

$$P(Y_{N+1} = y_j | Y_1, \dots, Y_N) = \frac{\Gamma(\alpha_0 + \alpha_1 + q)}{\Gamma(\alpha_0 + \alpha_1)} \frac{\Gamma(\beta_0 + \beta_1 + r)}{\Gamma(\beta_0 + \beta_1)} \frac{M_p}{K_p (\alpha_0 + \alpha_1 + \beta_0 + \beta_1)}$$

which simplifies to

$$P(Y_{N+1} = 1 | Y_1, \dots, Y_N) \propto \frac{\alpha_0 + \alpha_1}{(\alpha_0 + \alpha_1 + \beta_0 + \beta_1)}$$

$$P(Y_{N+1} = 0 | Y_1, \dots, Y_N) \propto \frac{\beta_0 + \beta_1}{(\alpha_0 + \alpha_1 + \beta_0 + \beta_1)}$$

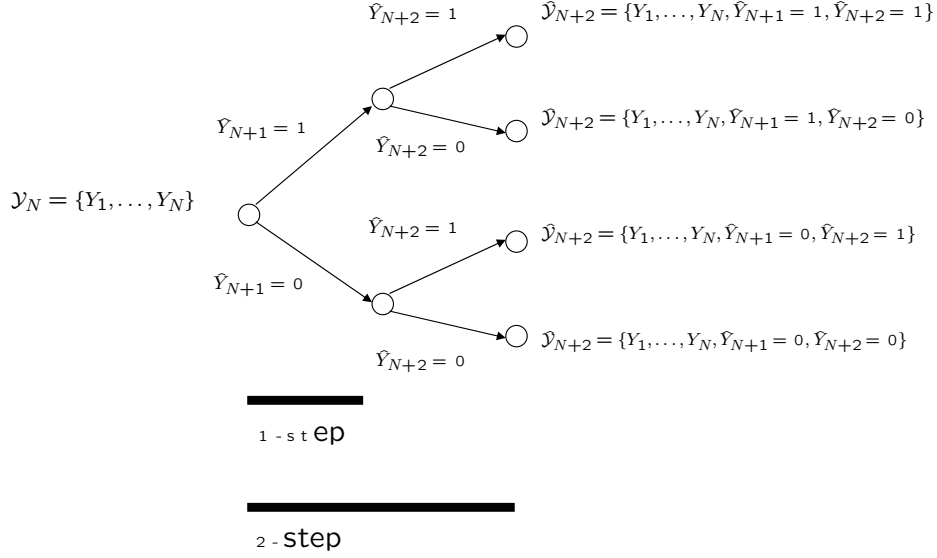


Fig. 3: Forecasting of measurements in the future.

B. Multi-step Predictive Distribution

While a 1-step predictive distribution is useful, a multi-step forecast will use more than one predicted observation to calculate the reduction in uncertainty. In a sense, this allows the prediction to be less greedy, and look further in the future. Hence, we can extend the results from the previous section to the case of multiple step forecast.

Proposition 4 *The T -step predictive distribution of a Beta distribution,*

$$P(Y_{N+1}, \dots, Y_{N+T} | Y_1, Y_2, \dots, Y_N) \quad (23)$$

can be expressed in closed form.

Proof: Define the following quantities

$$\begin{aligned} P(p) &= K_0 p^{\alpha_0 - 1} (1 - p)^{\beta_0 - 1} \\ P(Y_1, Y_2, \dots, Y_N | p) &= K_1 p^{\alpha_1} (1 - p)^{\beta_1} \\ P(Y_1, Y_2, \dots, Y_N, Y_{N+1}, \dots, Y_{N+T} | p) &= K_2 p^{\alpha_1 + q} (1 - p)^{\beta_1 + r} \end{aligned}$$

Here q and r are defined as the total number of times the future measurements will indicate that a target will be detected (and not detected, respectively)

$$q = \sum_{Y_i=1, i \geq N+1} Y_i, \quad r = \sum_{Y_i=0, i \geq N+1} Y_i \quad (24)$$

Note that since

$$P(Y_1, Y_2, \dots, Y_N) = \int_0^1 P(Y_1, Y_2, \dots, Y_N|p)P(p)dp$$

then it follows that

$$\begin{aligned} P(Y_1, Y_2, \dots, Y_N) &= \int_0^1 K_1 p^{\alpha_1} (1-p)^{\beta_1} K_0 p^{\alpha_0-1} (1-p)^{\beta_0-1} dp \\ &= K_1 K_0 \frac{\Gamma(\alpha_0 + \alpha_1) \Gamma(\beta_0 + \beta_1)}{\Gamma(\alpha_0 + \beta_0 + \alpha_1 + \beta_1)} \end{aligned} \quad (25)$$

Likewise, the marginal distribution for all the following $N + 1, N + 2, \dots, N + T$ measurement sequence will be given by

$$\begin{aligned} P(Y_1, Y_2, \dots, Y_N, Y_{N+1}, \dots, Y_{N+T}) &= \int_0^1 P(Y_1, Y_2, \dots, Y_N, Y_{N+1}, \dots, Y_{N+T}|p)P(p)dp \\ &= K_2 K_0 \frac{\Gamma(\alpha_0 + \alpha_1 + q) \Gamma(\beta_0 + \beta_1 + r)}{\Gamma(\alpha_0 + \beta_0 + \alpha_1 + \beta_1 + q + r)} \end{aligned} \quad (26)$$

Now, since the predictive distribution can be expressed

$$P(Y_{N+1}, \dots, Y_{N+T}|Y_1, Y_2, \dots, Y_N) = \frac{P(Y_1, Y_2, \dots, Y_N, Y_{N+1}, \dots, Y_{N+T})}{P(Y_1, Y_2, \dots, Y_N)} \quad (27)$$

we can take the quotient directly, thereby obtaining

$$\begin{aligned} P(Y_{N+1}, \dots, Y_{N+T}|Y_1, Y_2, \dots, Y_N) &= \\ \frac{K_2}{K_1} \frac{\Gamma(\alpha_0 + \alpha_1 + q)}{\Gamma(\alpha_0 + \alpha_1)} \frac{\Gamma(\beta_0 + \beta_1 + r)}{\Gamma(\beta_0 + \beta_1)} \frac{\Gamma(\alpha_0 + \alpha_1 + \beta_0 + \beta_1)}{\Gamma(\alpha_0 + \alpha_1 + \beta_0 + \beta_1 + q + r)} \end{aligned} \quad (28)$$

From our earlier definitions, we know that $K_1 = \frac{\Gamma(\alpha_1 + \beta_1 + 2)}{\Gamma(\alpha_1 + 1)\Gamma(\beta_1 + 1)}$ and $K_2 = \frac{\Gamma(\alpha_1 + \beta_1 + q + r + 2)}{\Gamma(\alpha_1 + q + 1)\Gamma(\beta_1 + r + 1)}$. By introducing the following notation

$$\begin{aligned} M_1 &\doteq \alpha_1 + \beta_1 \\ M_2 &\doteq \alpha_0 + \alpha_1 \\ M_3 &\doteq \beta_0 + \beta_1 \\ M_4 &\doteq M_2 + M_3 \\ F(b, r) &\doteq \frac{\Gamma(b)}{\Gamma(b + r)} \end{aligned}$$

and using a well known fact that $\Gamma(n + q) = n(n + 1) \dots (n + q - 1)\Gamma(n)$, we can thus obtain the final form for the predictive distribution

$$P(Y_{N+1}, \dots, Y_{N+T}|Y_1, Y_2, \dots, Y_N) = \frac{F(\alpha_1, q)F(M_2, q)F(M_3, r)}{F(M_1, q + r)F(\beta_1, r)F(M_4, q + r)} \quad (29)$$

Table 2: Multi-step forecasting example ($T = 2$)

q	r	Measurements
2	0	$\{Y_{N+1} = 1, Y_{N+2} = 1\}$
1	1	$\{Y_{N+1} = 1, Y_{N+2} = 0\}$
1	1	$\{Y_{N+1} = 0, Y_{N+2} = 1\}$
0	2	$\{Y_{N+1} = 0, Y_{N+2} = 0\}$

C. Remarks

This result explicitly evaluates the most likely measurement sequence obtained in the future; these measurements can in turn be used to predict the future reduction in uncertainty. Note that this result takes into account the prior distribution (by the shape parameters (α_0, β_0) which act as “prior counts”) and the actual measurements obtained (by the parameters (α_1, β_1) which serve as “measurement counts”).

The main tuning parameter is the forecast horizon, T , which dictates how far in the future to forecast. To find the maximum likelihood measurement sequence, the algorithm must try all possible combinations of future measurement sequences (see the case for the forecast horizon $T = 2$ in Table 2), and find the measurement sequence that maximizes the likelihood. Recall that q indicates all the times that a target is detected, and r denotes the times that it is not detected. Hence, these combinations of observations must be tried out, and the sequence that has the maximum probability is used as the expected pseudomeasurements in evaluating the uncertainty measures derived earlier. Note that since the measurements sequence will be a binary string, there is a finite number of them, though the combinations of possible sequences will grow larger as T increases.

V. Numerical simulations

This section describes the numerical simulations that were used to test the developed algorithms. There are a total of N_V vehicles that must be allocated to search an environment with a total of N_T targets. The continuous environment is discretized in uniform grid of dimension $M \times N$, with a total of $N_{cell} = MN$ cells. The objective of these simulations is to compare the relevant statistics of the total number of targets found (mean, maximum, and minimum) using the predictive approaches and non-predictive approaches. In the first section, the objective is to maximize the forecasted MBF and compare the number of targets found with the MBF without the use of forecasts. In the second section, target motion will be embedded, and we will compare the use of the predictive distribution robust objective of Eq. 8, and otherwise nominal prediction approaches.

A. Comparison with the Modified Bayes Factor

Each cell is initialized with a Beta distribution, with shape parameters b_1 and b_2 that represented the uncertainty in the probability in that cell. Each cell was initialized with a set of shape parameters.

Algorithm 1 Multi-Step Forecast Algorithm

```
for all  $i = 1, \dots, N_{cell}$  do
   $P_i \propto p^{a_i-1}(1-p)^{b_i-1}$  [Initialize cells]
  Initialize  $T$  [Prediction horizon]
end for
for all  $k = 0, \dots, T_{max}$  do
   $\{\hat{Y}_{k+1|k}^*, \dots, \hat{Y}_{k+T|k}^*\} = \arg \max_Y P(Y_{k+1}, \dots, Y_{k+T} | Y_1, Y_2, \dots, Y_k)$ 
   $\widehat{MBF}_{k+T|k} = g(MBF_{k|k}, \{\hat{Y}_{k+1|k}^*, \dots, \hat{Y}_{k+T|k}^*\})$ 
  if Replan then
     $\max_{x_k \in X_k} f(\widehat{MBF}_{k+T|k}, B_k, x_k)$ 
    Implement  $x_k$ 
  end if
end for
```

The key point here is that each cell's pdf on the probability of target existing is distinct and will have different moments. For example, cell 1 may have a distribution with a higher expected value and higher variance, but cell 10 may have a slightly lower probability of containing a target with lower variance. The simulations shown here compares the predictive algorithm (Algorithm 1) to a non-predictive algorithm. These simulations did not involve replanning, but future simulations will investigate the full multi-step algorithm. In this problem, the predictive algorithm was used to forecast ahead various time steps, evaluate the highest likelihood measurements sequence, and calculate the uncertainty if those observations were actually realized. The non-predictive algorithm did not forecast ahead and only used the current information to evaluate the uncertainty.

Target truths (that is, assigning a target to a random cell) were generated in Monte Carlo runs. Each of the cells was initialized with Beta shape parameters that indicated some prior knowledge corrupted by some random noise. This was done by choosing the shape parameters with a target being present or not in a cell, and adding random noise to the parameters. The simulations were parameterized both by forecast horizon (T) and total number of UAVs used. Each data point corresponds to the 300 Monte Carlo realizations.

Two measurements per cell that represented typical UAV updates on the Beta distributions were simulated, and these observations were used to update the priors. The forecasted measurements were then evaluated based on the updated priors and the actual observations, and these forecasted measurements were then used to assign the UAVs to the cells based on the predicted MBF, as shown in Algorithm 1.

These simulations show a “one-shot assignment”, that is, UAVs are assigned in one mission ($T_{max} = 1$) and they have to visit the cells with the greatest MBF. For the purposes of these simpler simulations, a target is declared found if a UAV is assigned to a cell with a target.

Numerical results associated with this figure are shown in Table 3. In contrast to the nominal planning scheme showed in the earlier sections, using the MBF and the forecasted MBF, allows a mission planner to select targets that have the lowest ambiguity in the information. These planning schemes typically select cells that have the sharpest pdf on the probability of target existence. The forecasted MBF (left column) on average finds more 1-2 targets than the un-forecasted MBF (right

Table 3: Comparison of MBF Forecasts

Num Vehicles (MBF, 1 step)	\bar{N}	σ_N	N_{\min}	Num Vehicles (MBF)	\bar{N}	σ_N	N_{\min}
12	7.5	2.7	4	12	9.9	3.3	5
14	9.1	3.3	4	14	9.9	3.8	6
16	10.6	3.5	6	16	10.3	3.6	7
18	12.8	4	9	18	11.1	2.8	9
20	14.7	3.5	11	20	12.5	2.8	10
22	16.4	2.4	13	22	14.3	3.9	11
24	17.6	1.5	14	24	15.9	2.3	12

column); furthermore, the minimum number of targets found by the forecasted MBF is higher than the unforecasted MBF. This can be attributed to the use of the predictive measurement that is used to calculate the MBF with the pseudomeasurement. By tasking vehicles with the most likely measurement, the algorithm can use the expected information to make more judicious decisions.

B. Predictive distribution with Target Motion

In this section, the effect of target motion is analyzed with extensive Monte Carlo simulations. Problems with target motion can particularly benefit from the predictive distribution. As the vehicles move in the environment, the propagation step effectively diffuses the target location to the adjacent cells, and the new Beta shape parameters for each cell will change as a result of this step. The predictive approach can nonetheless be used to evaluate the cells that are most likely to return a positive measurements, and use this measurement to evaluate the reduction in uncertainty.

Some of our work using the uncertain probabilities in the context of moving targets has addressed the issue of target motion in the search problem,⁷ and the following set of numerical experiments will address the use of the predictive distribution in these models.

1. Target Motion Model

The target motion is modeled with an environmental transition matrix Q whose $Q(i, j)$ entry is the probability that a target will transition to an adjacent cell i at time $k + 1$ subject to the target having been in cell j at time k :

$$Q(i, j) = \text{Prob}[i_{k+1}|j_k] \quad (30)$$

The transition matrix, $Q \in \mathcal{R}^{N \times N}$, is a stochastic matrix where each column sum is constrained to unity

$$Q = \begin{bmatrix} Q(1,1) & Q(1,2) & \dots & Q(1,N) \\ Q(2,1) & Q(2,2) & \dots & Q(2,N) \\ \dots & & & \\ Q(N,1) & Q(N,2) & \dots & Q(N,N) \end{bmatrix} \quad (31)$$

The diagonal elements, $Q(j, j)$ are related to the mean sojourn time,² or the expected amount of time that the target will remain in a particular cell. The mean sojourn time is varied in these

simulations by weighting each of the diagonal elements of the transition matrix with a scaling factor $\lambda > 1$. The matrix is re-normalized to ensure its stochastic property is satisfied.

The Beta parameters of each of the cells are initialized to (ζ_1, ζ_2) , and the uncertain probabilities are propagated forward using the motion model Q . The propagated distribution for each of the cells is approximated as a new Beta distribution with parameters (ζ'_1, ζ'_2) . The predictive distribution is evaluated, and the most likely measurement $Y_{k+1}^* \in \{0, 1\}$ for each of the cells is obtained. This expected measurements is used to evaluate the new mean and variance of Eq. 8, that is

$$\begin{aligned}\bar{P}_j &= \frac{\zeta'_1 + Y_{k+1}^*}{\zeta'_1 + \zeta'_2 + 1} \\ \sigma_j^2 &= \frac{(\zeta'_1 + Y_{k+1}^*)(\zeta'_2 + (1 - Y_{k+1}^*))}{(\zeta'_1 + \zeta'_2 + 1)^2(\zeta'_1 + \zeta'_2 + 2)}\end{aligned}\quad (32)$$

2. Simulation

The simulations were performed by using different number of targets (N_T) and varying weightings on the mean sojourn times, λ . The simulation parameters are

- $N_{cell} = 100$
- N_T varied from 3 to 40
- λ varied from 3 to 20

For each N_T and λ in the simulations, target locations were randomly initialized over 200 scenarios. The targets were propagated forward one time step using the motion model Q , and the UAVs were allocated to the cells that maximized the nominal allocation of Eq. 4 and the robust allocation of Eq. 8. While we experimented with various choices of α , the result for $\alpha = 3$ are presented here.

We then compared the expected number of targets found by the nominal and robust allocations, shown in Fig. 4. The x -axis shows the changing target number, while the y -axis shows the variation in λ . The z -axis gives the difference of total number of targets found (in percentage). The results shows that the robust allocation approach finds more targets on average with the prediction than the nominal approach. This is because the nominal prediction disregards the uncertainty, while the robust approach includes it in the planning. In particular, for 20 targets there is a 11% performance improvement (meaning the robust allocation finds 2 more targets than the nominal approach), while for 40 targets there is a 10% performance improvement (meaning that the robust allocation finds 4 more targets on average). This is consistent with our earlier results on planning with Uncertain Probability Maps.⁷

By using the predictive distribution in the measurements, the benefit of using the prediction with more agile targets is clear from this figure. As the targets are moving around in the environment, the counts (or shape parameters) of the Beta distribution are diffused to the adjacent cells. Hence, the cells that are most likely of returning a positive measurement if a UAV were assigned to them are constantly changing. A non-predictive objective does not take this into account however. By using the predictive approach, as the target number increases, the robust optimization finds 4 more targets on average than the nominal.

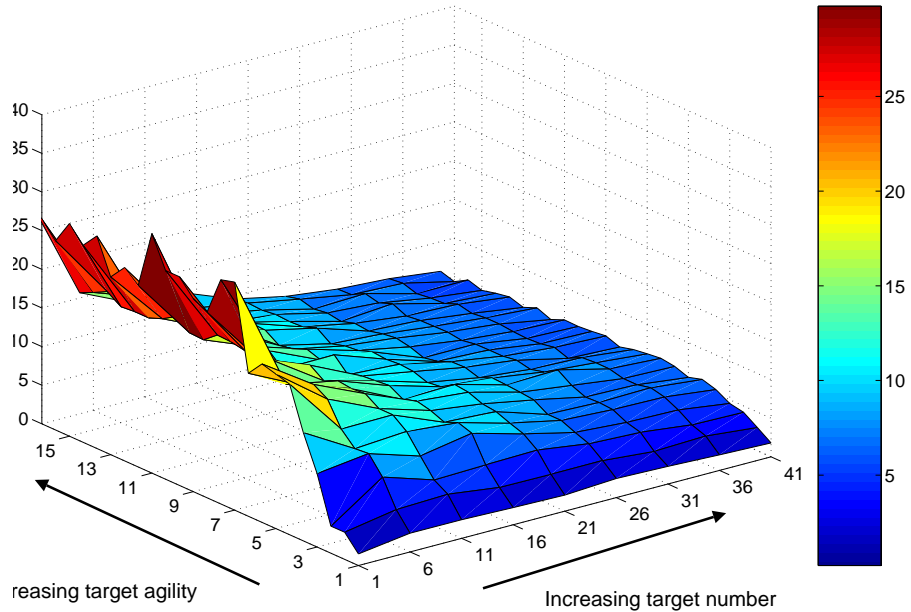


Fig. 4: Improvement in number of targets found of robust with forecast to nominal, $\alpha = 3$

VI. Conclusion and Future Work

This paper presented a new approach to UAV search task allocation using Bayesian forecasting. Key results were: the Modified Bayes Factor, which can be used in UAV planning to allocate UAVs to regions in space taking into account the uncertainty in the probability of target existence; and the design of a multi-step predictive forecasting algorithm that estimates the maximum likelihood *future* measurements (Algorithm 1).

An immediate goal is to investigate the effect of selecting a good forecast horizon. Currently, we have been using 1- or 2-step forecast, and have been able to empirically show that the search can benefit from such a predictive inference. However, as this forecast is extended, it is possible that the benefits of the predictive inference may not be as obvious, as the uncertainty in the forecast will also increase. Understanding the tradeoffs of choosing a sufficiently beneficial forecast is still a topic of research.

This work can be extended to much more general UAV problems. First, being able to predict future measurements will increase our ability to task *heterogeneous* types of vehicles, that is UAVs with different capabilities. UAVs that must do search missions along with rescue vehicles, for example, operate much more efficiently in a coupled manner, and using forecasted measurements in these coupled operations has the potential of improving overall mission performance.

Acknowledgments

The first author was supported in part by a National Defense, Science, and Engineering Graduate Fellowship, a Draper Fellowship, and AFOSR grant FA9550-04-1-0458.

References

- [1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. 9th Edition, Dover, New York, 1972.
- [2] Behrends, E. "Introduction to Markov Chains: With Special Emphasis on Rapid Mixing." Vieweg: Baun-schweigh, 2000.
- [3] D. Bertsimas and M. Sim, "Robust Discrete Optimization and Network Flows," submitted to *Operations Research Letters*, 2002.
- [4] D. Bertsimas and M. Sim, "Price of Robustness," submitted to *Mathematical Programming*, 2002.
- [5] Bertuccelli, L. F. , M. Alighabari, and J. P. How. *Robust Planning For Coupled Cooperative UAV Missions*, IEEE Conference on Decision and Control, 2004.
- [6] Bertuccelli, L. F. and J. P. How. *Robust UAV Search for Environments with Imprecise Probability Maps*, IEEE Conference on Decision and Control, 2005.
- [7] Bertuccelli, L. F. and J. P. How. *Search for Dynamic Targets with Uncertain Probability Maps*, IEEE American Control Conference, 2006.
- [8] Bourgault, F., T. Furukawa, and H. F. Durrant-Whyte. *Coordinated Decentralized Search for a Lost Target In a Bayesian World*, IEEE/RSJ International Conf. on Intelligent Robots and Systems, October 2003.
- [9] Bourgault, F., T. Furukawa, and H. F. Durrant-Whyte. *Process Model, Constraints, and the Coordinated Search Strategy*, IEEE International Conf. on Robotics and Automation, April 2004.
- [10] Bourgault, F., T. Furukawa, and H. F. Durrant-Whyte. *Decentralized Bayesian Negotiation for Cooperative Search*, IEEE/RSJ International Conf. on Intelligent Robots and Systems, October 2004.
- [11] Carlin, B. P. and T. A. Lewis. *Bayes and Empirical Bayes Methods for Data Analysis, 2nd Ed.* Chapman and Hall, 2000.
- [12] Gradshteyn, I. S., and I. M. Ryzhik. *Table of Integrals, Series, and Products. 4th Ed.* Academic Press; New York, 1965.
- [13] Jun, M., A. I. Chaudhry, and R. D'Andrea. *The Navigation of Autonomous Vehicles in Uncertain Dynamic Environments: A Case Study*, IEEE Conference on Decision and Control, 2002.
- [14] Jun, M. and R. D'Andrea. *Probability Map Building of Uncertain Dynamic Environments with Indistinguishable Obstacles*, IEEE American Control Conference, 2003.
- [15] Kreucher, C., K. Kastella, and A. O. Hero. *Sensor Management Using Relevance Feedback Learning*, IEEE Trans. on Signal Processing: Special Issue on Machine Learning, 2003.
- [16] Pollock, S. M. *A Simple Model of Search for a Moving Target*, Operations Research, Vol. 18(5), Pg. 883-903.
- [17] Popova, E. and D. Morton. *Adaptive Stochastic Manpower Scheduling*, Winter Simulation Conference, 2003.
- [18] Stone, L. *Theory of Optimal Search, 2nd ed.*, Military Applications Society, 2004.
- [19] Wong, E-M., F. Bourgault, and T. Furukawa. *Multi-vehicle Bayesian Search for Multiple Lost Targets*, IEEE International Conf. on Robotics and Automation, April 2005.
- [20] Yang, Y., M. Polycarpou, A. A. Minai. *Opportunistically cooperative neural learning in mobile agents*, Neural Networks: IJCNN 2002.
- [21] Yang, Y., M. Polycarpou, A. A. Minai. *Decentralized Cooperative Search in UAV's Using Opportunistic Learning*, AIAA GNC, 2002.
- [22] Yang, Y., A. Minai, and M. Polycarpou. *Evidential Map-Building Approaches for Multi-UAV Cooperative Search*, American Control Conference, 2005.