

Robust UAV Search for Environments with Imprecise Probability Maps

L. F. Bertuccelli and J. P. How

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

Abstract— This paper introduces a new framework for UAV search operations and proposes a new approach to calculate the minimum number of looks needed to achieve a given level of confidence of target existence in an uncertain gridded environment. Typical search theory formulations describe the uncertainty in the environment in a probabilistic fashion, by assigning probabilities of target existence to the individual cells of the grid. While assumed to be precisely known in the search theory literature, these probabilities are often the result of prior information and intelligence, and will likely be poorly known. The approach taken in this paper models this imprecise knowledge of the prior probabilities in the individual cells using the Beta distribution and generates search actions that are robust to the uncertainty. Use of the Beta distribution leads to an analytical prediction of the number of looks in a particular cell that would be needed to achieve a specified threshold in the confidence of target existence. The analytical results are demonstrated in both an expected value setting and a framework that takes into account the variance of the posterior distribution. The effectiveness of the proposed framework is demonstrated in several numerical simulations.

I. INTRODUCTION

Future UAV missions will require increasingly higher-level planning capabilities onboard the vehicles using information acquired through sensing or communications with other assets, including other UAVs (e.g., see Refs. [1], [2], [3], [5], [9], [11], [12] and the references therein). The quality and timeliness of this information will have a direct impact on the overall mission performance. Therefore an important aspect of the overall planning capability will be the ability to predict the consequences of the UAV actions in an uncertain environment. For example, in the context of an area search problem, it is desirable to anticipate how many image frames would be required by a camera onboard a vehicle in order to classify a target as “detected” or “undetected” with uncertain prior information on the target existence. This paper presents several solutions to this key question by investigating the uncertainty associated with a typical search problem.

UAV search problems are typically formulated by gridding the environment into a set of cells (e.g., Refs. [11], [14], [15], [16]). Each of the cells contains a point estimate of the probability that a target exists in that cell. The search problems are then solved with complete knowledge of this information. In general, however, it is unlikely that full knowledge of the target information is available to the mission planner. In the so-called “fog of war” there will be uncertainty in the information, due to poor intelligence or noisy sensors. For example, an unexplored region of a

battlefield may have a high probability of target existence, but there may also be a high uncertainty in that probability. Thus, rather than using point estimates to describe the target existence probability, it is more realistic to use sets that could describe the range of possible values. In order to do this, however, a formal mathematical formulation is required that will incorporate this uncertainty into the planning problem.

This paper presents a formulation that incorporates uncertainty with the use of Beta distributions and creates robust search actions. The Beta distribution has the appealing property that it is a conjugate distribution under a Bernoulli sensor model; this is an appropriate sensor model for this type of algorithm, since it abstracts the complexity of the image processing algorithms into a binary decision: target detected, or not. As such, the Bayesian measurement update step is *exact*, which is in contrast to other density functions (e.g., truncated Gaussian) that could be used, which are not conjugate and thus would result in approximate measurement updates. The probability of target existence is treated as a random variable in the open interval from $(0, 1)$ described by a probability density instead of as a point estimate. This embeds any uncertainty in the point estimates in the planning problem. Using this density for the uncertainty in the probability leads to analytic results for the number N of “effective looks” needed to search an area. This is an important metric in mission management and design, since it describes the number of observations that must be taken *with an imperfect sensor*, subject to uncertain information. From a practical sense, it also provides an estimate for how much time a vehicle should spend taking measurements of a target. For the case of a camera, the number of looks could denote the number of frames that must be taken and stored, and if these are stored at a known image rate, λ (in frames/second), a total *Time on Target* can be calculated as N/λ .

The structure of this paper is as follows. Section II discusses the general search problem, introduces the Beta distribution and motivates its use for modeling poor knowledge in the prior probability of target existence. Section III introduces some of the definitions and nomenclature used in this paper, including probabilistic models of the measurement equation. Section IV then describes the Bayesian update, demonstrating the conjugacy property, and discusses the problem statement of obtaining probabilistic thresholds to uniquely declare the existence (or lack thereof) of a target based on the number of observations that can be made in each cell. Specifically, this uncertain search problem is

analyzed to determine the impact of uncertainty measures on N . The analytical properties are presented based on expected value and variance objectives. These results clearly show a diminishing returns effect with additional looks, which is generally modeled in classical search theory (see Ref. [15]) with the decaying exponential detection function. Section V demonstrates the benefits of the approach with several numerical simulations.

II. SEARCH PROBLEM AND BETA DISTRIBUTIONS

Search problems are generally posed by discretizing the environment in an grid of cells over a 2-dimensional space indexed by $(i, j) \in (I, J)$. Each cell is available to contain a target, but this knowledge is not known prior to the search. It is therefore estimated with a target existence probability, P_{ij} . Also known as a target occupancy probability (e.g., [11]), the target existence probability is the probability that the particular cell contains an active target. The complement of this probability, $1 - P_{ij}$, denotes the probability that a cell does not contain a target. If the target *is known* to exist in a cell, then the precise probability of target existence in the cell is $P_{ij} = 1$. Alternatively, if the target does not exist, $P_{ij} = 0$.

Given a limited number of UAVs to accomplish the search, typical allocation problems ([11], [16]) assign the vehicles to the cells that have the highest probability of target existence. This is a reasonable objective if the prior probabilities are known precisely; however, this will be difficult to achieve in real-life operations. In general, the prior probabilities come from intelligence and previous observations, and cannot be treated as completely certain. Earlier results indicating the effect of uncertainty in resource allocation problems were presented in Refs. [4], [5], and showed that performance can be lost if uncertainty is not properly accounted for in resource allocation problems. As such, a practical framework needs to be developed that appropriately takes into account the uncertainty in this information for search operations.

Various approaches have been developed in the literature to describe the uncertainty in a probability P_{ij} . The approaches in [6], [7] have considered interval probability descriptions where the probability is within a given range of values ($P_{ij} \in [P_{\min}, P_{\max}]$). This formulation is appealing because of its simplicity, but there are difficulties in propagating this range of probabilities in a Bayesian framework. An alternative method of describing the uncertainty in the probability is to use scenarios (e.g., [13]), where each of the scenarios consists of a realization of the different prior probabilities for each cell. However, these approaches may require a large number of scenarios before the entire uncertain probability is accurately modeled, and can be computationally intensive.

The uncertainty in the probability can also be described using Beta distributions (e.g., [8], [10]). The Beta distribution is a very general distribution that treats P_{ij} as a random variable, and has various appealing features: 1) The support of the Beta distribution, the random variable P_{ij} , is on the open interval $(0, 1)$; 2) The Beta distribution is *conjugate*

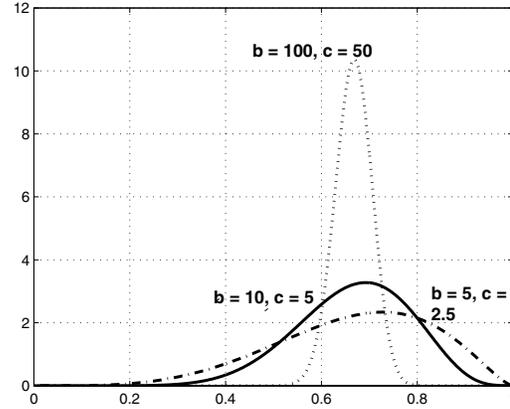


Fig. 1. Three examples of Beta distributions: $B_1(100, 50)$, $B_2(10, 5)$ and $B_3(5, 2.5)$ respectively. $B_i(b_i, c_i)$ denotes the i^{th} Beta distribution with parameters b_i and c_i . Note that the mean in all these cases is 0.67, but the distributions are clearly different.

under a Bernoulli sensor likelihood function, a simplified but common sensor model used in the UAV community.¹ The interest in using the Bernoulli model is that it abstracts away the complexities of the image processing algorithms, and results in the binary decisions that would come out of the algorithms: “yes” a target was detected, or “no” a target was not detected.

The conjugacy property ensures that the support of the posterior distribution is also on the open interval from $(0, 1)$, which is critical for propagating the distributions in a Bayesian framework. Note that since a point estimate for a target existence probability is effectively an “impulse” probability distribution, the Beta distribution can likewise be interpreted as a new form of the impulse distribution, or as a “distribution” on the distribution.

The Beta distribution is defined by

$$P(x) = \frac{\Gamma(b+c)}{\Gamma(b)\Gamma(c)} x^{b-1}(1-x)^{c-1}, \quad x \in (0, 1) \quad (1)$$

$b > 1$, $c > 1$ are weights that can be used as tuning parameters to define the initial distribution (the case of $b = 1$ and $c = 1$ denotes the uniform distribution). x varies over the continuous range from 0 to 1 and represents the support of the probability distribution. The Γ function, $\Gamma(m+1)$ is defined as

$$\Gamma(m+1) = \int_0^\infty y^m e^{-y} dy \quad (2)$$

where in the case when m is integer, this expression defines the factorial $\Gamma(m+1) = m!$.

Three examples of the Beta distribution are shown in Figure 1. While the relative weighting of the b and c values in this figure is the same ($b/(b+c) = 0.67$ for all cases), the actual values of these weighting parameters heavily influence the form of the distribution. By appropriate choice of the weighting parameters b and c , one can generate distributions

¹A prior distribution belonging to a certain class is termed *conjugate* if the posterior distribution is in the same class.

of the probabilities of target existence that capture any prior knowledge of the target existence. For example, if a mission designer was very certain that a target existed in a particular cell, a weighting of $b = 50$, $c = 1$ could be chosen, resulting in a distribution heavily skewed towards $x = 1$. If the designer were highly uncertain of the target existence, a uniform distribution could be used, by using $b = c = 1$. Choosing higher values for b and c places more weight on the prior, and the updated distribution requires a larger number of observations (compared to a prior with lower values of the shaping parameters) for the measurements to have an effect on it.

III. DEFINITIONS

As in classical search theory, the update on the probability densities is done using Bayes' Rule

$$P(x|Y) = \frac{P(Y|x) P(x)}{\int_0^1 P(Y|x)P(x)dx} \quad (3)$$

where: *i*) $P(Y|x)$ denotes the *likelihood function* which is the probability distribution of the sensor that is being used; *ii*) $P(x)$ denotes the prior distribution of the target; *iii*) the denominator serves as a normalizing factor to preserve the property of a distribution that its integral sum to 1; and, *iv*) $P(x|Y)$ is the posterior distribution, which is the updated distribution based on new measurements.

In this paper, the prior is described by a Beta distribution as in Equation 1. The likelihood distribution is given by a series of N observations that indicate whether a target is detected or not in a cell. It is a Bernoulli distribution, since the sensor returns whether a target was detected or not. In the following equation, this is represented by $Y = 1$ if a target is detected, and $Y = 0$ if a target is not detected.

$$P(Y|x) = B_n x^{\gamma_1} (1-x)^{\gamma_2} \quad (4)$$

B_n is the Binomial coefficient; γ_1 denotes the number of times a target is detected, and γ_2 indicates the number of times a target is not detected. Note that the total number of looks is given by $\gamma_1 + \gamma_2 = N$.

Finally, define a sensor accuracy parameter, ξ , which is the probability of "correct detections"; that is, the probability that if a target exists, then it will be detected by the sensor. The probability of a missed detection is given by $1 - \xi$. In general there are analogous probabilities in the case when a target does not exist, i.e., the probability of a correct rejection, χ , and a probability of false alarm, $1 - \chi$. For sake of generality, we will assume that $\xi \neq \chi$.

IV. PROBLEM STATEMENT

Having defined the notation that will be used throughout this paper, we now define the main problem statement. Consider a cell with an imprecise target existence probability, x , described by a Beta distribution. Taking a series of observations in a particular cell results in a new Beta distribution which is shifted rightward towards $x = 1$ if the noisy observations indicate a target exists, or leftward towards

$x = 0$ if the opposite is true. The key point is that this new distribution will have an updated set of statistical moments which can be used to establish new confidence values for the target existence. For example, if the expected value of the posterior distribution after a series of N observations is greater than a threshold α close to 1, we can declare that a target is in fact present in the cell. Likewise, if this expected value is less than an analogous threshold $\hat{\alpha}$ close to 0, the cell is assumed not to contain a target. Using our proposed framework, it is now possible to predict the number of looks that are required to exceed a pre-defined threshold to unambiguously conclude the presence or absence of a target for the case when the probabilities are precise.

The first result in this section is for the case when only the expected value of the distribution is used as a criterion for the declaration of target existence or absence.

Objective A: Expected Value Formulation Given *pre-defined sensor errors and a prior described by the Beta distribution, find the minimum number of looks, N , that raise the expected value of the posterior distribution, $\bar{x}(N)$, above a threshold α*

$$\{ \min N \mid \bar{x}(N) \geq \alpha \} \quad (5)$$

Result A: *The minimum number of looks, $N = \gamma_1 + \gamma_2$, required to raise the expected value of the posterior distribution to α , is given by*

$$N \geq \frac{(\alpha - 1)b + \alpha c}{\xi - \alpha} \quad (6)$$

Thus, given knowledge of the sensor error, and prior weighting given to the prior distribution, we can estimate the number of looks required to achieve an expected value for the probability that a target is actually there.

Proof: The proof consists of three steps: constructing the posterior distribution; finding the expected value of the distribution; and solving for N . The posterior distribution, $P(x|Y)$ is

$$\begin{aligned} P(x|Y) &= \frac{P(Y|x) P(x)}{\int_0^1 P(Y|x)P(x)dx} \\ &= \frac{\Gamma(b + \gamma_1 + c + \gamma_2)}{\Gamma(b + \gamma_1) \Gamma(c + \gamma_2)} x^{b+\gamma_1-1} (1-x)^{c+\gamma_2-1} \end{aligned}$$

The expected value of the posterior distribution is given by

$$\begin{aligned} \bar{x} &= \int_0^1 x P(x|Y) dx \\ &= \frac{\Gamma(b + \gamma_1 + c + \gamma_2)}{\Gamma(b + \gamma_1 + c + \gamma_2 + 1)} \frac{\Gamma(b + \gamma_1 + 1)}{\Gamma(b + \gamma_1)} \end{aligned} \quad (7)$$

Using the property of Gamma functions that

$$\Gamma(m + 1) = m \Gamma(m) \quad (8)$$

Eq. 7 simplifies to

$$\bar{x} = \frac{b + \gamma_1}{b + \gamma_1 + c + \gamma_2} \quad (9)$$

Since the probability of a correct detection is the probability of a target being detected by the sensor when there is a target present in a cell, this can be estimated over a set of measurements, N , as

$$\xi \approx \frac{\gamma_1}{N} \quad (10)$$

where γ_1 indicates the number of measurements indicating a target was detected in the cell. Hence, $\gamma_1 \approx N\xi$. This can be combined with Eq. 9 to solve for the total number of measurements (or “looks” in the cell) explicitly. Substituting this result in Eq. (9) and solving for N

$$N \geq \frac{(\alpha - 1)b + \alpha c}{\xi - \alpha} \quad \blacksquare \quad (11)$$

The key appeal of this result is that it relates the initial condition of the distribution (summarized by b and c) and sensor accuracy (ξ), to the objective α . This formulation only relies on the first moment of the distribution, since it is based on achieving a certain threshold on the expected value. It does not consider the entire distribution, which could be undesirable. For example, with the number of looks expressed in Result A, an operator may want the posterior distribution to be tightly distributed about the mean, while exceeding the threshold. This will ensure that there is a much higher “certainty” about the mean. Thus, we are also interested in a result that incorporates the tightness of the distribution, in a similar approach to that used in Ref. [4]. Hence a new formulation that includes a variance term is presented.

Objective B1: Robust Search Formulation Find the number of looks required to raise the expected value of the posterior distribution to exceed a threshold η (subject to a penalty of uncertainty, σ , and a robustness parameter μ)

$$\{ \min N \mid \bar{x}(N) - \mu\sigma(N) \geq \eta \} \quad (12)$$

Result B1: The number of looks required in a “robust” search problem is given by the solution to the following cubic equation

$$G_3 N^3 + G_2 N^2 + G_1 N + G_0 \geq 0 \quad (13)$$

where $G_i \forall i = 0, 1, 2, 3$ are constant parameters.

Proof: Since the expected value is given by

$$\bar{x} = \frac{b + \gamma_1}{b + c + \gamma_1 + \gamma_2} = \frac{b + N\xi}{b + c + N} \quad (14)$$

and the standard deviation is

$$\begin{aligned} \sigma &= \sqrt{\frac{(b + \gamma_1)(c + \gamma_2)}{(b + c + N)^2(b + c + 1 + N)}} \\ &= \frac{1}{b + c + N} \sqrt{\frac{(b + N\xi)(c + N(1 - \xi))}{b + c + 1 + N}} \end{aligned} \quad (15)$$

The optimization is to find the minimum N such that

$$\frac{b + N\xi}{b + c + N} - \mu \frac{1}{b + c + N} \sqrt{\frac{(b + N\xi)(c + N(1 - \xi))}{b + c + 1 + N}} \geq \eta \quad (16)$$

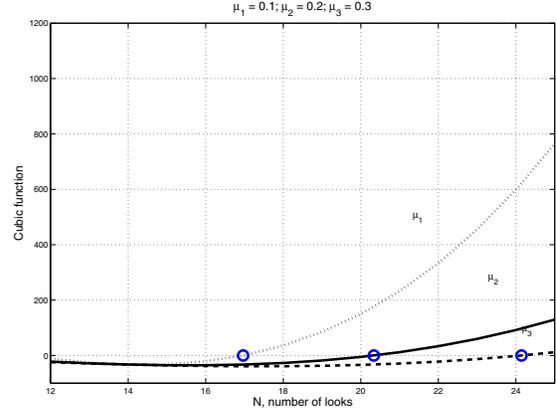


Fig. 2. A plot of the cubic function of μ . The result shows that as a greater confidence is required (from $\mu = 0.1$ to $\mu = 0.3$), the number of looks increases, from approximately 17 to 24. All cases shown were with $\xi = 0.85$, $\eta = 0.8$, and initial uniform distribution ($b = 1$, $c = 1$).

which, after some algebra, results in the cubic equation that must be solved for N

$$G_3 N^3 + G_2 N^2 + G_1 N + G_0 \geq 0 \quad (17)$$

where

$$\begin{aligned} G_3 &= (\eta - \xi)^2 \\ G_2 &= M(\eta - \xi)^2 + 2L(\eta - \xi) - \mu^2 \xi(1 - \xi) \\ G_1 &= 2LM(\eta - \xi) + L^2 - \mu^2(\xi c + b(1 - \xi)) \\ G_0 &= M L^2 - \mu^2 bc \\ L &= (b + c)\eta - b \\ M &= b + c + 1 \end{aligned} \quad \blacksquare$$

Intuitively, any form of this result where $\mu > 0$ will result in a greater number of looks in the particular region of interest than for $\mu = 0$. This effect is shown in Figure 2, where three sample cubic functions were plotted directly as a function of the number of looks N , for varying μ . It is clear from this figure that as the tuning parameter, μ , is increased (a much tighter posterior distribution is desired), the number of looks required to achieve the same threshold η increases.

Remark 1 (Results A and B): In Objectives A and B, results for the number of looks were only presented for the case when the target was present. Of course an analogous set exist that indicate the number of looks that should be made in a cell to declare that a target is *not present* in a cell. A mission designer would then use the maximum number of looks of the two formulations to unambiguously declare the existence or absence of a target.

Remark 2 (Result A): The fundamental constraint for Result A to be valid is that $N > 0$, and the non-trivial inequality that must be satisfied by the solution for N is

$$\xi > \alpha > \frac{b}{b + c}$$

It displays the interesting property that the parameter α cannot be chosen arbitrarily, and is constrained by the sensor

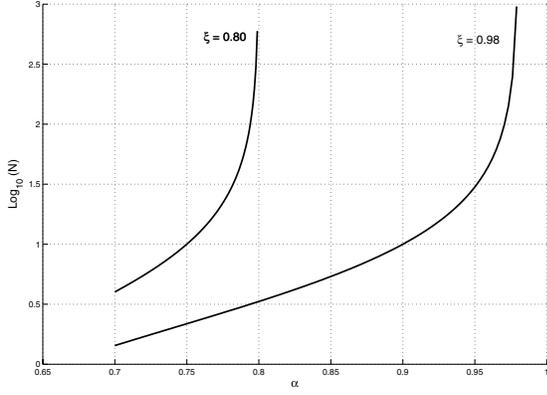


Fig. 3. The logarithm of the number of looks is plotted against α for two distinct sensor errors. Note that as $\alpha \rightarrow \xi$, the number of looks increases very rapidly, and there is a diminishing returns behavior.

accuracy ξ . This means that there is a fundamental limitation to the expected value of the posterior subject to the sensor accuracy: the chosen threshold α cannot exceed the accuracy of the sensor.

Furthermore, the number of looks can become increasingly large as the expected value of the posterior density is increased. In the case that $\alpha \rightarrow \xi$ the righthand side of Eq. (11) tends to infinity (see examples in Figure 3). This property underscores the “diminishing returns” behavior of this model. As α is increased, the number of looks rapidly increases (note that the y -axis is a logarithmic scale). One important observation from the figure is that, for lower sensor accuracies, there is a fundamental limit to the expected value of the posterior density. Further, if this value is chosen to be close to the sensor accuracy, the number of looks can become prohibitively large.

Remark 3 (Result B) The two parameters η and μ have a significant impact on the number of looks that are required to achieve a certain confidence in the target presence. There is an intrinsic interplay between the parameters μ and η in finding the optimal arrangement to suit a mission, and this will be ultimately the responsibility of the mission designer. Care must also be taken in choosing the parameters η, μ to ensure that when the cubic equation is solved, that N is real and positive.

Table I shows two cases for the uniform distribution ($b = 1, c = 1$), with a sensor accuracy of $\xi = 0.85$. Table I shows the effect of varying η for constant μ . As shown in Figure 3, when $\eta \approx \xi$, the number of looks increases significantly for a marginal increase in the threshold.

V. NUMERICAL RESULTS AND DISCUSSION

This section presents some numerical results that show the effectiveness of the expected value approach of Section IV and compares it to heuristics assuming nominal information.

In this analysis 4 targets are present in the environment; the imprecise prior information indicates that there is only a 50% chance of the targets actually existing (this is the point

TABLE I
NUMBER OF LOOKS FOR DIFFERENT η : $b = 1, c = 1, \xi = 0.85$

η	μ	N
0.70	0.01	2.7
0.75	0.01	5.1
0.80	0.01	12.2
0.84	0.01	71.1

estimate, and implies the total number of targets is unknown). The priors are described by Beta distributions, where the (b, c) values are next to the target number in Table II. Note that the cells have different (b, c) values, and will thus require a different number of looks to exceed the threshold. The purpose of the mission is to determine how many targets actually exist by assigning the UAV as indicated in Figure 4, and taking the predicted number of looks in each cell. The total number of looks will allow the mission designer to *predict* the length of the mission. In the case of the analytical result, a total Time on Target (ToT) for the i^{th} target (T_i) is calculated based on the prior information, using Eq. 11. The thresholds that are obtained with these number of looks are calculated in 1000 Monte Carlo simulations, and compared in Table II.

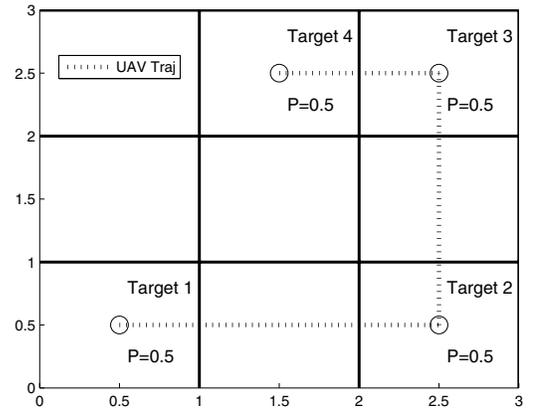


Fig. 4. Visualization of the environment discretized in 9 cells. 4 targets are present in the environment, and the prior probability for each is 0.5. The UAV has a unique number of looks that it can place on each target.

The heuristics used specify that looking in a cell for a predefined amount of time will allow the operator to exceed the threshold for which the target is unambiguously declared absent or present in the cell. Some of these heuristics have been used in the literature to model the need to look in a cell a sufficient number of times to increase the threshold in a cell ([9]). Three heuristics are compared to the analytical expressions obtained in this paper. In the first heuristic, the maximum amount of time ($T_A = \max_i(T_i), \forall i$) is taken in each cell; in the second heuristic, the minimum time ($T_B = \min_i(T_i), \forall i$) is taken in each cell; in the third heuristic, the average length of time ($T_C = \bar{T}_i, \forall i$) is taken in each cell. The key point in this experiment is that the individual

heuristics assume an equal amount of time sent in each cell, since the prior probabilities of the targets are identical.

In each simulation, samples of actual target existence of the target were generated from a Beta distribution; the individual search strategies were then implemented to investigate the level of confidence achieved for each cell. For the first heuristic, while only a total of 112 units of time are spent on target (shorter mission time), only one of the four targets is unambiguously identified. In the second heuristic, all four targets are identified, but at the expense of increasing the mission time to 280. In the third heuristic, which has an equal *total* number of looks as the analytical results, still only two of the four targets are unambiguously identified. As in the earlier heuristics, even though each target is allocated the same number of looks, target 4 requires a larger number of looks to exceed the threshold due to the large weighting on the prior.

The poor knowledge of the probabilities plays a significant factor in determining a priori the number of looks to be assigned in each cell. While the second heuristic results in a correct detection of all the targets, it does so at the expense of a mission time that is 43% longer than that predicted by the analytical expression that incorporates the uncertainty in the distribution. Also, while the third heuristic results in the same number of looks as the analytical approach, it does not unambiguously identify half of the targets. Clearly, the initial imprecision in the target probability generated search actions that either did not exceed the specified thresholds or resulted in longer mission times than necessary. The new framework compensates for the different levels of imprecision associated with each cell and generated missions that allocated a different number of looks to each prospective target. This increased flexibility in the mission design enabled the operator to complete the objectives using shorter missions.

TABLE II
TIME ON TARGET (TOT): $\alpha = 0.85, \xi = 0.95$

Target (b, c)	α_{exp}, T_A	α_{exp}, T_B	α_{exp}, T_C	α_{analy}
1 (6,6)	0.81	0.89	0.86	0.85
2 (4,4)	0.85	0.90	0.89	0.85
3 (8,8)	0.79	0.87	0.84	0.85
4 (10,10)	0.76	0.85	0.82	0.85
Total ToT	112	280	196	196

VI. CONCLUSION

This paper has presented a new framework for UAV search operations in an uncertain environment. The uncertainty in the environment was modeled as an imprecise knowledge of the prior distributions in the cells of the discretized environment. A formal mathematical framework with the use of the Beta distributions was proposed to model this lack of precision, and results that predicted the number of looks to be made in each cell were presented. These results were initially with an expected value objective, which was then extended to include a variance penalty that leads to tighter

posterior distributions. This framework was then compared in numerical simulations to a heuristic framework, and was shown to be more successful in coping with the uncertain environment. Future work will extend this approach to the case of dynamic targets and further investigate the role of uncertainty in UAV search missions.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Alighanbari, M., L. F. Bertuccelli, and J. P. How. *Filter-Embedded UAV Task Assignment Algorithms For Dynamic Environments* AIAA Guidance, Navigation and Control Conference, Aug 2004.
- [2] Bellingham, J., M. Tillerson, A. Richards, and J. P. How. *Multi-Task Allocation and Path Planning for Cooperative UAVs*. In S. Butenko, R. Murphey, and P. M. Pardalos, *Cooperative Control: Models, Applications, and Algorithms*, Pg. 23-41. Kluwer Academic Publishers, 2003.
- [3] Bellingham, J. S., M. Tillerson, M. Alighanbari and J. P. How, *Cooperative Path Planning for Multiple UAVs in Dynamic and Uncertain Environments*, Proceeding of the IEEE Conference on Decision and Control, Dec. 2002.
- [4] Bertuccelli, L. F. *Robust Planning for Heterogeneous UAVs in Uncertain Environments*, Massachusetts Institute of Technology, M. S. Thesis, 2004.
- [5] Bertuccelli, L. F. , M. Alighanbari, and J. P. How. *Robust Planning For Coupled Cooperative UAV Missions*, IEEE Conference on Decision and Control , 2004.
- [6] Boeva, V. and B. De Baets. *A New Approach to Admissible Alternatives in Interval Decision Making*, Proceedings of 2nd Intelligent Systems Conference, 2004. Col 1, Pg. 22-24, 2004.
- [7] Brown, D. E. and C. C. White. *Methods for Reasoning with Imprecise Probabilities in Intelligent Decision Systems*, Conference on Systems, Man and Cybernetics, Pg. 161-163, 1990.
- [8] Ferson, S., V. Kreinovich, L. Ginzburg, D. S. Myers, and K. Sentz. *Constructing Probability Boxes and Dempster-Shafer Structures*, Sandia National Laboratories. SAND2002-4015, 2002.
- [9] Flint, M., E. Fernandez-Gaucherand, and M. Polycarpou. *Cooperative Control for UAVs Searching Risky Environments for Targets*, IEEE Conference and Decision and Control, 2003.
- [10] Hutter, M. *Robust Estimators under the Imprecise Dirichlet Model*, Proc. 3rd International Symposium on Imprecise Probabilities and Their Applications (ISIPTA-2003) 274-289.
- [11] Jin, Y., A. Minai, and M. Polycarpou. *Cooperative Real-Time Search and Task Allocation in UAV Teams*, IEEE Conference on Decision and Control, 2003.
- [12] Jun, M. and R. D'Andrea. *Probability Map Building of Uncertain Dynamic Environments with Indistinguishable Obstacles*, IEEE American Control Conference, 2003.
- [13] Krokhmal, P., Murphey R., Pardalos, P., Uryasev, S., and G. Zrazhevsky. *Robust Decision Making: Addressing Uncertainties in Distributions*. S. Butenko et al (Eds.) Cooperative Control: Models Applications and Algorithms. Kluwer Academic Publishers, Pg. 165-185, 2003.
- [14] Polycarpou, M., Y. Yang, and K. Passino. *A Cooperative Search Framework for Distributed Agents*, IEEE International Symposium on Intelligent Control, 2001.
- [15] Stone, L. *Theory of Optimal Search*, 2nded., Military Applications Society, 2004.
- [16] Yang, Y., A. Minai, and M. Polycarpou. *Decentralized Cooperative Search by Networked UAVs in an Uncertain Environment*, American Control Conference, 2004.