Autonomous UAV Search Planning with Possibilistic Inputs

Emily Grayson[1], Paul Elmore[2], Don Sofge[3], Fred Petry[2]
[1]Computer Science , Carnegie Mellon University, Pittsburgh PA
[2]Geospatial Science and Technology, Naval Research Laboratory, Stennis Space Center MS
[3]Information Technology, Naval Research Laboratory, Washington DC

***Abstract*— Many aspects of decision making processes for autonomous systems involve human subjective information in some form. Methods for informing decision making processes with human information are needed to inform probabilistic information used in an autonomous system. This can provide better decisions and permit a UAV to more quickly and efficiently complete tasks. Specifically we use possibility theory to represent the subjective information and apply possibilistic conditioning of the probability distribution. A simulation platform was developed to evaluate approaches to using possibilistic inputs and showed that is was feasible to make effective usage of such information.**

Keywords— possibility distribution, possibilistic conditioning, Bayesian updating, autonomous systems

I. INTRODUCTION

Low-cost aerial vehicles equipped with on-board cameras may be used to autonomously track objects on the ground, and thereby provide a valuable capability for operations such as search-and-rescue [1] and aerial surveillance [2]. Tracking objects on the ground using multiple aerial vehicles requires that the vehicles be routed in such a way that information gain increases as the search continues. Ideally, we want to find search paths that maximize information gain.

This paper will focus specifically on path planning for autonomous systems, in particular deciding where to move in an environment in order to most quickly complete a task, given human sources of information. A prime example of such a task, where time is a factor and human information is often provided, is a search and rescue mission. Such a mission will require the missing individuals be found as quickly as possible, and sources of human information may include where they were last seen, local inhabitants' insights into the area and so forth. This information, provided by human sources, is called subjective or epistemic information.

Here we will focus on use of a possibility approach to represent such information. Possibility measures vary from probability measures in that a possibility describes a degree of belief and makes it feasible to represent incomplete knowledge. that is generalized from Bayesian methods An approach for informing a Bayesian a priori probability distribution with possibilistic information will be used in this paper. We also include an outline of the approach taken to experimentally test these methods, a description of the simulation and planner developed and discussion of results.

II. BACKGROUND

A. *Autonomous Vehicle Search and Tracking*

Recent literature on tracking multiple targets suggests approaches to addressing the planning problem through use of Bayesian inference and information theory. In [3]–[7] Bayesian inference and maximum likelihood methods are combined for multi-target tracking, while [8]–[10] describe methods for moving the mobile platforms. In fact, most of these methods utilize only a single mobile sensor platform.

Optimal sensor selection and placement is explored from the perspective of design optimization in [11]-[13], but these methods do not scale well as the number of sensors/platforms increases due to the complexity of calculating mutual information amongst the sensors and targets. This planning process must also be robust to limitations in the autonomous vehicles' ability to communicate due to bandwidth limitations or interference.

Physics-inspired swarm-based control is proposed in [14]-[15]. In this control paradigm a heat-map is imposed over the search space, and the heat is adjusted such that the vehicles follow a gradient along the heat map. The "temperature" of the surface is adjusted such that the vehicles are driven to gather new information, thereby increasing information gain. Achieving the appropriate behaviors, however, requires careful tuning of multiple parameters, resulting in a heuristic method that may not guarantee convergence to a near-optimal search policy.

In [16], Sydney implements a distributed motion planner for multiple vehicles tracking multiple moving targets through use of a Bayesian Likelihood Ratio Tracker (B-LRT). The algorithm implements the Dynamic Data-Driven Application Systems (DDDAS) [17] paradigm in which sensor measurements are used to guide subsequent data collection. Each mobile platform is assumed to cover a finite area, and its on-board sensors provide a binary "yes" or "no" reading of whether (or not) a target is present within its search area. The B-LRT assimilates the data from all of the platforms into a map, which represents a probability density function (PDF) of the likelihood that targets are present over each particular point within the search domain. Mutual information between the sensor and target states is maximized by having the vehicles following the gradient of the information surface.

The behaviors approximate the motions of molecules in one of three states of matter: solid, liquid, or gas. As compared with the methods described in [14]-[15] which required parameter tuning to produce the desired behaviors, in this method the behaviors are emergent and require fewer tuning parameters.

B. Possibilistic Conditioning

To formalize the problem, let V be a discrete variable taking values in a space X that has both aleatory and epistemic sources of uncertainty [18]. Let there be a probability distribution P: $X \rightarrow [0, 1]$ such that $p_i \in [0,1]$, : $\sum_{i=1}^{n} p_i = 1$ that models the aleatory uncertainty. Then the epistemic uncertainty can be modeled by a possibility distribution [19] such that $\Pi : X \rightarrow [0, 1]$, where $\pi(x_i)$ gives the possibility that x_i is the value of V, i = 1,2,...n. A usual requirement here is the normality condition, $\max_x [\pi (x)] = 1$, that is at least one element in X must be fully possible. Abbreviating our notation so that $p_i = p(x_i)$, etc. and $\pi_i = \pi(x_i)$, etc., we have $P = \{p_1, p_2, ...p_n\}$ and $\Pi = \{\pi_1 , \pi_2 ,...., \pi_n\}$.

In possibilistic conditioning, a function f dependent on both P and Π is used to find a new conditioned probability distribution such that [20]

$$f(P,\Pi) \Rightarrow \text{new } \hat{P}$$

where $\hat{P} = \{ \hat{p}_1, \hat{p}_2, ..., \hat{p}_n\}$ with $\hat{p}$

$$\hat{p}_i = \frac{p_i \pi_i}{K} \qquad \text{where} \quad K = \sum_{i=1}^{n} p_i \pi_i \tag{1}$$

A strength of this approach using conditioned probability is that it also captures Zadeh's concept of consistency between the possibility and the original probability distribution. Consistency provides an intuition of concurrence between the possibility and probability distributions being aggregated. In Eq (1), K is identical to Zadeh's possibility-probability consistency measure [19], $C_Z (P,\Pi)$; i.e. $C_Z (P, \Pi) = K$.

For the purposes of this paper, this conditioning was done at some time step t during the experiment. An example of the use of this method would be if, during a search and rescue mission, some new information was discovered by a human source who believed the missing person who more likely to be farther North and less likely to be in the East. This information would be represented by a possibility distribution, where the North would be assigned a value of 1 in the distribution and the East would be assigned values of or close to 0. Then the prior distribution over the area being covered in the search would be conditioned with this new possibility information.

Because of the subjective and potentially less reliable nature of the possibility information distribution more weight can be given to the probability distribution than the possibility distribution in the simulation. Another point of importance is how to utilize Zadeh consistency, as the greater the consistency, or the lower the conflict, between the possibility and probability distribution, the more informative the newly conditioned posterior will be; and the lower the consistency, or higher the conflict, between the possibility and probability distributions, the less informative the posterior will be.

C. Weighted Aggregation

The simplest method used in this paper, the weighted aggregation method, allows one to combine multiple sources of information with varying levels of reliability or credibility. So if a problem required combining information, for example, from multiple sensors of different types, where each sensor had a different accuracy, this could be done easily using this method; higher weights given to the more accurate or reliable sensors.

For the present, it is assumed only two sensors are used and each sensor is equally reliable. That is, here we used an equi-weighted approach, essentially taking the average of two sensor values.

$$\hat{p} = w_1 p_{1k} + w_2 p_{2k} = 1/2(p_{1k} + p_{2k})$$

III. SIMULATION

A. Experimental Design

In this section we describe how these methods were implemented, what was being tested, and potential results. Specifically we describe the simulation platform and planner, which determined where the simulated UAV would move, and the method used to update the a priori information as the UAV continuously sensed the space searching for the target.

Our basic evaluative criterion for these simulations was that using these methods for informing the prior would assist a UAV in detecting the target more rapidly. The specific criterion was represented by the number of iterations, i.e. how many times the UAV (or UAVs in the case of a multi agent simulation) moved. Iterations are a more accurate representation of success than real time because the simulation will run more quickly or slowly depending on the number of UAVs in simulation, size of the grid, or, to a lesser extent, the method being used to inform the prior; the number of iterations moves at a constant pace regardless of these or any other factors, and is therefore a better method for accurately comparing results. Note that if the probability and possibility distributions are in conflict, the UAV should in general take more iterations than the control to detect the target, but if they are not in conflict then there may be less time than the control to detect the target.

We begin with a previous simulation developed for persistent coverage [21] extended to target detection, moving from a risked based to certainty based planner. This approach adapts how the UAV senses the environment, adding methods for updating the grid based on new sensor information, and adding functions for defining or informing the prior with additional information. Included are variables such as which method is used, how large and how many grid cells there are, number of UAVs, and number of targets.

B. Planner

The planner used is a reactive planner, meaning it does not matter where in the environment the UAV is, but what immediately surrounds the UAV. The main benefit of a reactive planner is that it is highly scalable. Regardless of the size of the environment, obstacles placed within the environment, how much or little the UAV knows about its surroundings, the planner will be able to direct the UAV throughout the space.

The main algorithm for the planner involves cycling through 4 steps: evaluate, move, sense and update. First, the planner looks to see in which direction certainty is highest by evaluating the certainty at discretized segments along the sensor footprint. The sensor footprint is elliptical and positioned in line with but in front of the UAV. Once the direction of highest certainty is determined the UAV moves in this direction. The UAV then takes a sensor measurement in its new location. Because this is a simulated experiment, we will know in which grid cell the target is located. So, if the target is present the sensor will return a value of 0.9 and if no target is present it will return a value of 0.1, the 0.9 and 0.1 representing an imperfect sensor. These values may be changed depending on the quality of the sensor in use. Given this new sensor measurement, the prior certainty values in the grid are updated using a Bayesian update. Once this is done, the process will repeat until the cell where the target is located reaches a threshold value for detection, or the simulation goes on for too long and reaches a maximum number of iterations.

C. Update Method

The update method is a Bayesian Filter. Some common examples of Bayesian filters include a Kalman Filter, Particle Filter, and Multiple hypothesis tracking, however here a simplified filter was used [22]. The initial sensor measurement z at cell k is 0 or 1:

$$z_k \in \{0,1\}$$

The probability of detecting a target given prior knowledge of the current grid cell is between 0 and 1 exclusive and depends on quality of the sensor. In this case

$$p(z_k \mid \Theta_k) \in \{0.1, 0.9\}$$

The prior is updated by multiplying the new probability of detection with the prior at cell k and then normalizes this value.

$$update : \left\{ \begin{array}{l} p(\theta_k \mid z_k) = \dfrac{p(z_k \mid \theta_k)\, p(\theta_k \mid z_{k-1})}{p(z_k \mid z_{k-1})} \\ p(z_k \mid z_{k-1}) = \sum_{\Omega} p(z_k \mid \theta_k)\, p(\theta_k \mid z_{k-1}) \end{array} \right\}$$

At the cell k where the most recent sensor measurement occurred:

Target was Detected: Yes- $p(z_k \mid \Theta_k) = 0.9$
No - $p(z_k \mid \Theta_k) = 0.1$

At all other locations k' > k

Target will be Detected: Yes- $p(z_{k'} \mid \Theta_{k'}) = 0.9$
No - $p(z_{k'} \mid \Theta_{k'}) = 0.1$

Benefits of this basic approach include that it is a simple way to represent an imperfect sensor and that it allows for the inclusion of multiple sensors in future work. However, the main drawback of this update method is that it is computationally inefficient, as grid size increases complexity grows exponentially. At the moment, the simulation will run relatively quickly for a grid of up to 12 by 12 cells, but beyond this the simulation becomes inoperable. For the basic testing, we did not use a grid larger than 10 by 10.

IV. Simulation Results

Four sets of simulations are compared against a benchmark simulation. Table I provides a description of these computations. In each simulation, possibilistic conditioning and update occur at time step $t = 100$. The benchmark has no possibilistic conditioning, using only Bayesian updates alone. All simulations are repeated 100 times.

TABLE I. Description of Simulation

Run #	Description
1	"Random-1": The possibility of 1 is placed randomly at one location at t=100; possibility = 0 everywhere else.
2	"Target-1": The possibility of 1 placed at the true target location at t=100; possibility = 0 everywhere else. Possibilistic update at t=100.
3	"Random-1 & prior": Possibility of 1 placed randomly at t=100; every other location has possibility = prior probability.
4	"Target-1 & prior": Possibility of 1 at the true target location at t=100; every other location has possibility = prior probability.
Benchmark	"Uniform" The initial probability is set at uniform across the search space, and there is no possibilistic update during the simulation. This run serves as the benchmark.

Once the simulations are complete, histograms of the number of steps needed to find the target are examined. Figure 1 shows the histogram from Run #2 as an example. The simplest fit to each histogram from the simulations is an exponential probability density function, given by

$$p(t) = \mu^{-1} \exp(-t/\mu).$$

In Eq. (1), μ is mean. The mean and standard deviation, σ, are equivalent for an exponential probability distribution. (Numerically, $\mu > \sigma$ of update steps needed to find the target, but were the same to within 12%, 9%, 8%, 3% and 4% of the mean for Run #1 to #4, respectively, and 9% for the Benchmark run). Column 1 of Table II provides statistics from each simulation. Also listed are the upper and lower bounds for a 90th percentile confidence interval of μ, given by $\{\underline{\mu}, \overline{\mu}\}$. The maximum likelihood estimation algorithm (as implemented in the Matlab Statistics Toolbox estimate) outputs $\{\underline{\mu}, \overline{\mu}\}$.

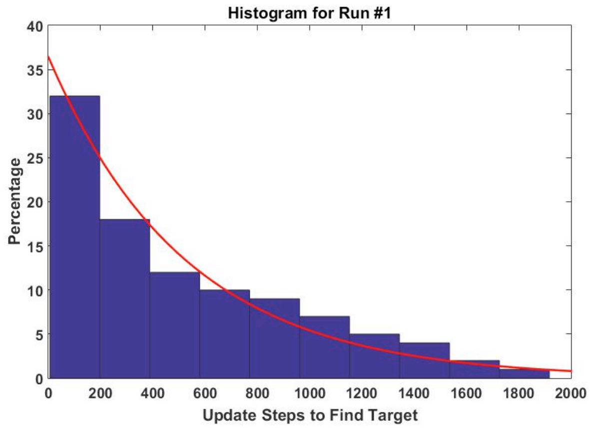

Fig. 1. Histogram of update steps for Run #1 in Tables I and II. The curve plots Eq. (1) using the computed mean.

TABLE II. MEAN, MEDIAN AND 90TH PRECENTILE TO FIND TARGET IN SIMULATIONS, LOWER AND UPPER BOUNDS ARE IN BRACKETS

Run #	Mean $\{\underline{\mu},\overline{\mu}\}$	$t_{50},\{\underline{t}_{50},\overline{t}_{50}\}$	$t_{90},\{\underline{t}_{90},\overline{t}_{90}\}$
1	524, {520, 533}	365, {361, 370}	1209, {1197, 1228}
2	500, {495, 508}	347, {344, 353}	1152, {1141, 1170}
3	432, {428, 438}	300, {297, 305}	995, {985, 1011}
4	383, {380, 389}	267, {264, 271}	884, {876, 898}
Benchmark	512, {507, 520}	356, {352, 361}	1179, {1168, 1198}

From Eq. (1) and the estimate for μ , the normalized CDF is computed over the time-steps vector of 1 to 2,001 seconds. From the CDF, the median, t_{50} , and 90th quantile, t_{90} , are found. Columns 2 and 3 in Table II provide these results. To obtain the lower and upper bounds for t_{50} and t_{90} ; $\{\underline{t}_{50},\overline{t}_{50}\}$ and $\{\underline{t}_{90},\overline{t}_{90}\}$, respectively; the CDF and quantile calculations are repeated using $\{\underline{\mu},\overline{\mu}\}$ in Eq. (1). These two curves are the upper and lower bounds for the CDF. The median and 90th quantiles from the lower bound CDF give $\underline{t}_{50}$ and $\underline{t}_{90}$. Likewise, the upper bound CDF give $\overline{t}_{50}$ and $\overline{t}_{90}$.

A. Discussion

It is expected that the better the possibilistic conditioning "helped" the target search, the lower the mean, median, and 90th quantile. In the simulations, good inputs corresponded to a possibility of 1 set at the target. To get the best reduction in time, however, the possibility elsewhere at $t = 100$ also has to be set to the prior distribution. Maximum use of prior knowledge at that time seems to have the greatest impact to the targeting algorithm.

Use of possibility of 0 in conditioning for the non-target locations actually is detrimental. The conditioning operation sets these probabilities to zero, which is actually incorrect for the sensor used and the Bayesian targeting results at that time step of the simulation. Doing so takes the Bayesian algorithm further away or "offtrack" from a convergence state at t=100, resulting in longer search iterations.

Hence for this application, it is proposed that the possibilistic conditioning needs to be revised such that

$$\tilde{p}_i = \max\left[\frac{p_i \pi_i}{K}, p_i\right].$$

A strength of the possibilistic conditioning is the relaxation in the normalization condition compared to probabilities. Although the possibility of 1 is set at the true target location, the end user has the freedom to set possibility > 0 at the other locations. If one were to use probability, then probability = 1 at the target would mean all other probabilities would be zero, which is not correct and throws the Bayesian update offtrack, creating longer convergence times.

V. CONCLUSIONS AND FUTURE WORK

The focus of this work was developing a simulation platform on which to test the effectiveness of methods for informing a priori information in a search and rescue scenario. This involved developing a reactive path planner which moved the UAV towards the area of highest certainty within the sensor footprint, taking sensor measurements along the way until a detection threshold is reached and the target is detected. Results so far show that some of the methods of informing a prior information are more effective than the uninformative prior of maximum entropy.

Some planned experiments include using a static versus a dynamic target or a single target versus multiple targets. Other measures of how informative the prior is, such as the Shannon entropy or Gini index [23], can be used to evaluate whether a more informative a priori with regards to these methods results in quicker detection of the target. Also varying the method used to update the a priori with the newest sensor measurement and determining the effect of the updating method on the number of iterations can be considered. Finally we are preparing to run experimental tests using physical UAVs in the NRL Laboratory for Autonomous Systems Research (LASR)'s high bay, in order to compare results of a physical UAV to the simulated results.

ACKNOWLEDGMENT

We would like to thank the Naval Research Laboratory's Base Program, Program Element No. 0602435N for sponsoring this research. .

REFERENCES

[1] H.A.F. Almurib, P.T.Nathan, and T.N.Kumar. "Control and path planning of quadrotor aerial vehicles for search and rescue." In Proceedings of SICE Annual Conference (SICE), pages 700–705, Sept 2011.

[2] H. C. Yang, R. AbouSleiman, B. Sababha, E. Gjoni, D. Korff, and O.

Rawashdeh. "Implementation of an autonomous surveillance quadrotor system." In Proceedings of AIAA Infotech, April 2009.

[3] W.R.Blanding, P.K.Willett, and Y.Bar-Shalom. "Multiple target tracking using maximum likelihood probabilistic data association." In 2007 IEEE Aerospace Conference, pages 1–12, March 2007.

[4] P. Willett, Y Ruan, and R. Streit. "Pmht: problems and some solutions." IEEE Transactions on Aerospace and Electronic Systems, 38(3):738–754, Jul 2002.

[5] L. D. Stone, C. A. Barlow, and T. L. Corwin. Bayesian Multiple Target Tracking. Artech House, 1999.

[6] D.E.Clark and J.Bell. "Bayesian multiple target tracking in forward scan sonar images using the phd filter." IEEE Proceedings of Radar, Sonar and Navigation, 152(5):327–334, October 2005.

[7] C-C. Huang and S-J Wang. "A Bayesian hierarchical framework for multitarget labeling and correspondence with ghost suppression over multicamera surveillance system." Automation Science and Engineering, IEEE Transactions on, 9(1):16–30, Jan 2012.

[8] M. A. Peot, T. W. Altshuler, A. Breiholz, R. A. Bueker, K. W. Fertig, A. T. Hawkins, and S. Reddy. "Planning sensing actions for UAVs in urban domains." volume 5986, pages 59860J–59860J–10, 2005.

[9] F. Rafi, S. Khan, K. Shafiq, and M. Shah. "Autonomous target following by unmanned aerial vehicles." volume 6230 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2006.

[10] J. Kim and Y. Kim. "Moving ground target tracking in dense obstacle areas using UAVs." In Proceedings of the 17th IFAC World Congress, 2008.

[11] V. Isler and M. Magdon-Ismail. "Sensor selection in arbitrary dimensions." Automation Science and Engineering, IEEE Transactions on, 5(4):651– 660, Oct 2008.

[12] G. M. Hoffmann and C.J. Tomlin. "Mobile sensor network control using mutual information and particle filters." Automatic Control, IEEE Transactions on, 55(1):32–47, Jan 2010.

[13] Chris M Kreucher, Keith D Kastella, and Alfred O Hero III. "Information-based sensor management for multi-target tracking." In Optical Science and Technology, SPIE's 48th Annual Meeting, pages 480–489. International Society for Optics and Photonics, 2003.

[14] E. Cuevas, A. Echavarria, and M. A. Ramirez-Ortegon. "An optimization algorithm inspired by the states of matter that improves the balance between exploration and exploitation." Applied Intelligence, 40(2):256– 272, 2014.

[15] W. M. Spears, D. F. Spears, J. C. Hamann, and R. Heil. "Distributed, physics-based control of swarms of vehicles. "Autonomous Robots, 17(2-3):137–162, 2004.

[16] N. Sydney and D. Sofge, "Distributed information-theoretic target detection using physics-inspired motion coordination." Proceedings of the 8th International Symposium on Resilient Control Systems, IEEE, 2015.

[17] F. Darema. "Dynamic data driven applications systems: A new paradigm for application simulations and measurements. "In Computational Science-ICCS, 2004.

[18] S. Parsons S (2001) Qualitative methods for reasoning under uncertainty. MIT Press, Cambridge MA

[19] L. Zadeh (1978) "Fuzzy sets as a basis for a theory of possibility". Fuzzy Sets and Systems 1: 3-28.

[20] R. Yager (2012) "Conditional approach to possibility-probability fusion." IEEE Trans Fuzzy Syst 20:46–56.

[21] Wallar, A., Plaku, E., & Sofge, D. "A planner for autonomous risk-sensitive coverage (PARCov) by a team of unmanned aerial vehicles." 2014 IEEE Symposium on Swarm Intelligence.

[22] Fox, D., J. Hightower, Lin Liao, D. Schulz, and G. Borriello. "Bayesian filtering for location estimation." IEEE Pervasive Computing: 24-33.

[23] P.Elmore, F. Petry, and R. Yager. "Comparative Measures of Aggregated Uncertainty Representations." Journal of Ambient Intelligence and Humanized Computing 5.6 (2014): 809-19.