

Optimal Search for a Lost Target in a Bayesian World

Frédéric Bourgault¹, Tomonari Furukawa², and Hugh F. Durrant-Whyte¹

¹ ARC Centre of Excellence for Autonomous Systems (CAS)
Australian Centre for Field Robotics
The University of Sydney, Sydney, NSW 2006, Australia
{f.bourgault, hugh}@acfr.usyd.edu.au
<http://www.acfr.usyd.edu.au>

² ARC Centre of Excellence for Autonomous Systems (CAS)
School of Mechanical and Manuf. Engineering
The University of New South Wales, Sydney, NSW 2052, Australia
t.furukawa@unsw.edu.au

Abstract. This paper presents a Bayesian approach to the problem of searching for a single lost target by a single autonomous sensor platform. The target may be static or mobile but not evading. Two candidate utility functions for the control solution are highlighted, namely the Mean Time to Detection, and the Cumulative Probability of Detection. The framework is implemented for an airborne vehicle looking for both a stationary and a drifting target at sea. Simulation results for different control solutions are investigated and compared to demonstrate the effectiveness of the method.

1 Introduction

“We are sinking fast. Position ten miles south of San Remo...”¹

When rescue authorities receive a distress signal time becomes critical. The probability of survival decreases rapidly in a matter of hours when lost at sea. The prime focus of a rescue mission is to search for and find the castaways in the smallest possible amount of time. The search, based on some coarse estimate of the target location, must often be performed in low visibility conditions and despite strong winds and high seas making the location estimate even more uncertain as time goes by. Keeping these time and physical constraints in mind, what should be the optimal search strategy?

This paper presents a Bayesian framework that integrates and predicts the effects of the observations and the process model on the target distribution. The control solution formulation is then described for a single autonomous vehicle searching for a single non evading, but possibly mobile target.

The paper is organized as follows. First, the Bayesian filtering algorithm that accurately maintains and updates the target state probability density function (PDF) is described in Sect. 2. Section 3 describes the searching problem, highlights two utility candidate functions and formulates the control optimization problem. Section 4 implements the framework for an airborne search vehicle and investigates the

¹ Mayday from yacht *Winston Churchill*, April 1959 [10]

control solutions and the effectiveness of the approach for both a stationary, and a drifting target. Finally, conclusions and ongoing research directions are highlighted.

2 Bayesian Analysis

This section presents the mathematical formulation of the Bayesian analysis from which the control solutions presented in this paper are derived. The Bayesian approach is particularly suitable for combining, in a rational manner, non-linear motion models and heterogeneous non-Gaussian sensor measurements with other sources of quantitative and qualitative information [8][1].

In Bayesian analysis any quantity that is not known is modelled as a random variable. The state of knowledge about such a random variable is expressed in the form of a probability density function (PDF). Any new information in the form of a probabilistic measurement or observation is combined with prior information using the Baye's theorem in order to update the state of knowledge and form the new a posteriori PDF. That PDF forms the quantitative basis on which all control decisions or inferences are made.

In the searching problem, the unknown variable is the target state vector $\mathbf{x}^t \in \mathcal{X}^t$ which in general describes its location but could also include its attitude, velocity, etc. The analysis starts by determining the a priori PDF of \mathbf{x}^t , $p(\mathbf{x}_0^t | \mathbf{z}_0) \equiv p(\mathbf{x}_0^t)$, which combines all available information including past experience. For example, this prior PDF could be in the form of a Gaussian distribution representing the prior coarse estimate of the parameter of interest. If nothing is known about the parameter, a least informative approach is to represent this knowledge by a uniform PDF. Then, once the prior distribution has been established, the PDF of the target state at time step k , $p(\mathbf{x}_k^t | \mathbf{z}_{1:k})$, can be constructed recursively, provided the sequence $\mathbf{z}_{1:k} = \{\mathbf{z}_1, \dots, \mathbf{z}_k\}$ of all the observations made by the sensor(s) on board the search vehicle, \mathbf{z}_k being the observation (or the set of observations, if multiple sensors) made a time step k . This recursive estimation is done in two stages: prediction and update.

2.1 Prediction

A prediction stage is necessary in Bayesian analysis when the PDF of the state to be evaluated is evolving with time i.e. the target is in motion or the uncertainty about its location is increasing. Suppose we are at time step k and the latest PDF update, $p(\mathbf{x}_{k-1}^t | \mathbf{z}_{1:k-1})$ (from the the previous time step) is available. Then the predicted PDF of the target state at time step k is obtained from the following Chapman-Kolmogorov equation

$$p(\mathbf{x}_k^t | \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k^t | \mathbf{x}_{k-1}^t) p(\mathbf{x}_{k-1}^t | \mathbf{z}_{1:k-1}) d\mathbf{x}_{k-1}^t \quad (1)$$

where $p(\mathbf{x}_k^t | \mathbf{x}_{k-1}^t)$ is a probabilistic Markov motion model. If the motion model is invariant over the target states, then the above integral is simply a convolution.

Practically, this convolution is performed numerically by a discretization of the two PDF's on a grid, followed by the multiplication of their Fast Fourier Transforms (FFT)'s, followed by an inverse FFT of the produce to retrieve the result.

2.2 Update

At time step k a new observation \mathbf{z}_k becomes available and the update is performed using Bayes rule where all the observations are assumed to be independent. The update is performed simply by multiplying the prior PDF (posterior from the prediction stage) by the new conditional observation likelihood noted $p(\mathbf{z}_k|\mathbf{x}_k^t)$ as in the following

$$p(\mathbf{x}_k^t|\mathbf{z}_{1:k}) = K p(\mathbf{x}_k^t|\mathbf{z}_{1:k-1}) \cdot p(\mathbf{z}_k|\mathbf{x}_k^t) \quad (2)$$

where the normalization factor K is given by

$$K = 1 / \int [p(\mathbf{x}_k^t|\mathbf{z}_{1:k-1})p(\mathbf{z}_k|\mathbf{x}_k^t)] d\mathbf{x}_k^t \quad (3)$$

Practically, the multiplication of (2) is performed numerically by multiplying together the corresponding elements of a grid.

3 The Searching Problem

This section describes the equations for computing the probability of detection of a lost object referred to as the target. For further details on the searching problem the reader is referred to [7] and [6].

If the target detection likelihood (observation model) at time step k is given by $p(\mathbf{z}_k|\mathbf{x}_k^t)$ where $\mathbf{z}_k = D_k$ for which D_k represents a ‘‘detection’’ event at t_k , then the likelihood of ‘‘no detection’’, given a target state \mathbf{x}_k^t is given by its complement

$$p(\overline{D}_k|\mathbf{x}_k^t) = 1 - p(D_k|\mathbf{x}_k^t) \quad (4)$$

At time step k , the conditional probability that the target does not get detected during a sensor observation, $p(\overline{D}_k|\mathbf{z}_{1:k-1}) = q_k$, depends on two things: the ‘no detection’ likelihood (4), and the latest target PDF $p(\mathbf{x}_k^t|\mathbf{z}_{1:k-1})$ (from the prediction stage (1)). In fact q_k corresponds exactly to the volume under the surface formed when multiplying the two together (element-by-element for each given target state \mathbf{x}_k^t) as in the following

$$p(\overline{D}_k|\mathbf{z}_{1:k-1}) = \int p(\overline{D}_k|\mathbf{x}_k^t)p(\mathbf{x}_k^t|\mathbf{z}_{1:k-1})d\mathbf{x}_k^t = q_k \quad (5)$$

Hence q_k is given by the reduced volume (< 1) under the target state PDF after having been carved out by the ‘no detection’ likelihood in the update stage (2), but before applying the normalization factor to it. Notice that this volume is exactly the inverse of the normalization factor K (see (3) for a ‘no detection’ event ($\mathbf{z}_k = \overline{D}_k$)), so $q_k = 1/K$ and is always smaller than 1. The joint probability of failing to detect

the target in all of the steps from 1 to k , noted $Q_k = p(\overline{D}_{1:k})$, is obtained from the product of all the q_k 's as follows

$$Q_k = \prod_{i=1}^k p(\overline{D}_i | \overline{D}_{1:i-1}) = \prod_{i=1}^k q_i = Q_{k-1} q_k \quad (6)$$

where $\overline{D}_{1:i-1}$ corresponds to the set of observations $\mathbf{z}_{1:i-1}$ where all observations are equal to \overline{D} . Therefore, in k steps, the probability that the target has been detected, denoted P_k , is given by

$$P_k = 1 - Q_k \quad (7)$$

It is also possible to compute the probability that the target gets detected for the first time on time step k , denoted p_k , as follows

$$p_k = \prod_{i=1}^{k-1} p(\overline{D}_i | \overline{D}_{1:i-1}) [1 - p(\overline{D}_k | \overline{D}_{1:k-1})] = \prod_{i=1}^{k-1} q_i [1 - q_k] = Q_{k-1} [1 - q_k] \quad (8)$$

which in turn by summing over k provides a sequential method for evaluating P_k as

$$P_k = \sum_{i=1}^k p_i = P_{k-1} + p_k \quad (9)$$

For this reason P_k will be referred to as the ‘cumulative’ probability of detection at time k to distinguish it from the conditional probability of detection at time k which is equal to $1 - q_k$. Notice that as k goes to infinity, the cumulative probability of detection increases towards one. With k increasing, the added probability of detection p_k gets smaller and smaller as the conditional probability of detection $(1 - q_k)$ gets discounted by a continuously decreasing Q_{k-1} .

The mean time to detection (MTTD) is the expectation of the number of steps required to detect the target

$$E[k] = \sum_{k=1}^{\infty} k p_k = \text{MTTD} \quad (10)$$

The goal of the searching strategy could either be to maximize the chances of finding the target given a restricted amount of time by maximizing P_k over the time horizon, or to minimize the expected time to find the target by minimizing the MTTD. The difficulty though in evaluating the MTTD lies in the fact that one must in theory evaluate p_k for all k 's up to infinity.

3.1 Optimal Trajectory

Optimality is always defined in relation to an objective, or utility function [9]. For the searching problem there are two suitable candidates to evaluate a trajectory utility, namely the cumulative probability of detection P_k (9), and the MTTD (10).

For an action sequence $\mathbf{u} = \{u_1, \dots, u_{N_k}\}$ over a finite time horizon of length $T = N_k dt$, we thus have as an objective function either

$$J(\mathbf{u}, N_k) = \sum_{i=k}^{k+N_k} p_i = P_{k+N_k} - P_k \quad \text{or} \quad J(\mathbf{u}, N_k) = - \sum_{i=k}^{k+N_k} i p_i \quad (11)$$

The optimal control strategy \mathbf{u}^* is the sequence that maximizes that utility subject to the vehicle limitations $\mathbf{u}_{LB} \leq \mathbf{u} \leq \mathbf{u}_{UB}$.

$$\mathbf{u}^* = \{u_1^*, \dots, u_{N_k}^*\} = \arg \max_{\mathbf{u}} J(\mathbf{u}, N_k) \quad (12)$$

For the searching problem, because early actions strongly influence the utility of subsequent actions, the longer the time horizon, the better the computed trajectory. However, the computational cost follows the ‘‘curse of dimensionality’’ and with increasing lookahead depth the solution becomes intractable. In practice only solutions for very restricted lookahead length are possible. One way to increase the lookahead without increasing the cost of the solution too much is to use piecewise constant control sequences (see [5] and [2]) where each control parameters is maintained over a specified number of time steps. Such control solutions are said to be ‘quasi-optimal’ as they compromise the global optimality of the control solution for a lower computation cost, but nevertheless, depending on the problem at hand, often provide better trajectories than the ones computed with the same number of control parameters but with shorter time horizons.

3.2 One-step Lookahead

Planning with a time horizon of only one step is an interesting special case of the searching problem as both objective functions reduce to $J(\mathbf{u}, 1) = p_k$.


Also, because $p_k = Q_{k-1}(1 - q_k)$ (8), maximizing p_k at time step t_k is equivalent to maximizing $(1 - q_k) = p(D_k | \mathbf{z}_{1:k-1})$, the conditional probability of detecting the target (which corresponds to the volume under the surface resulting from the multiplication of the ‘detection’ likelihood with the predicted target PDF), or conversely minimizing $q_k = p(\overline{D}_k | \mathbf{z}_{1:k-1})$ (5), the conditional probability of ‘not detecting’ the target (volume under the surface resulting from multiplying the ‘no detection’ likelihood with the predicted target PDF). As will be seen in the results section 4.4, this greedy form of searching strategy provides very sensible control solutions at very low computational costs.

4 Application

The goal of the work presented in this paper is to ultimately implement and demonstrate the framework for an autonomous search on one of the ACFR’s unmanned air vehicle (UAV) as shown in Fig. 1a. ACFR has also developed a high fidelity simulator (Fig. 1b) of the UAV’s hardware, complete with different sensor models,



(a)



(b)

Fig. 1. Application: (a) one of the Brumby Mark-III uav's been developed at ACFR as part of the ANSER project. This flight vehicle has a payload capacity of up to 13.5 kg and operational speed of 50 to 100 knots; (b) display of the high fidelity simulator

on which the flight software can be tested before being implemented on board the platform almost without any modifications [3].

The rest of this section describes the implementation of the Bayesian searching framework for a single airborne vehicle searching for a single non-evading lost target that could either be stationary or mobile. However, the method is readily applicable to searching problems of all kinds, be it ground, underwater or airborne search for bushfires, lost hikers, enemy troops in the battlefield, or prospection for ore and oil, or even to search for water or evidence of life on another planet.

4.1 Problem Description

The problem chosen for the illustration of the framework involves the search by an airborne vehicle for a life-raft lost at sea. The search platform is equipped with a GPS receiver, i.e. assuming perfect localization, and a searching sensor (e.g radar, human eye, infrared or CCD camera) that can be modelled by a likelihood function (over range and bearing) hence relating the control actions to the probability of finding the target. There is one observation (full scan) made once every second. The sensor is assumed to have perfect discrimination i.e. no false target detection. However, it may fail to call a detection when the target is present i.e. miss contact.

The omnibearing sensor's maximum range (400m) is much smaller than the size of the searching area (2km x 2km). Drift current and winds (of up to 30 knots) affect the target distribution over time in a probabilistic way through the process model. The target PDF is of general form (i.e. non-Gaussian) and is evaluated and maintained on a discrete grid. As the length of the search is limited by the vehicle fuel autonomy, the utility function selected is given by (11) (left) and consists of maximizing the cumulative probability of finding the target in a fixed amount of time.

4.2 Motion Prediction

Vehicle Model The vehicle pose prediction model used for the planning purposes is the following discrete time non-linear constant velocity model

$$x_{k+1}^s = x_k^s + \frac{2V}{u_k} \sin\left(\frac{1}{2}u_k dt\right) \cos\left(\theta_k^s + \frac{1}{2}u_k dt\right) \quad (13)$$

$$y_{k+1}^s = y_k^s + \frac{2V}{u_k} \sin\left(\frac{1}{2}u_k dt\right) \sin\left(\theta_k^s + \frac{1}{2}u_k dt\right) \quad (14)$$

$$\theta_{k+1}^s = \theta_k^s + u_k dt \quad (15)$$

where the turn rate control command u_k is maintained over the time interval dt . For $u_k dt \ll 1$, i.e. turn rate close to zero, (13) and (14) reduce to

$$x_{k+1}^s = x_k^s + V dt \cos(\theta_k^s) \quad (16)$$

$$y_{k+1}^s = y_k^s + V dt \sin(\theta_k^s) \quad (17)$$

The maximum turn rate amplitude ($u_{max} = \pm 1.1607$ rad/s) corresponds to a 6g acceleration, the UAV's manoeuvre limit at $V = 50$ m/s (100 knots).

Process Model The model of the target state evolution noted $p(\mathbf{x}_k^t | \mathbf{x}_{k-1}^t)$, also called the target process, or motion model maps the probability of transition from a given previous state to \mathbf{x}_k^t , the target state at time t_k . It is defined by the target's equations of motion and the known statistics of the wind and the drift currents orientations and speeds. In this example, the life-raft is assumed to be drifting in the same direction and at a velocity proportional to the wind velocity. It was found that a joint distribution combining a Gaussian distribution for the wind direction with mean μ_θ and variance σ_θ^2 , and a Beta distribution for the velocity amplitude v where $v \in [0, v_{max}]$ as in the following expression

$$p(v, \theta) = \frac{c}{v_{max}} \left(\frac{v}{v_{max}}\right)^{a-1} \left(1 - \frac{v}{v_{max}}\right)^{b-1} \frac{1}{\sqrt{2\pi}\sigma_\theta\mu_v} e^{-\frac{(\theta-\mu_\theta)^2}{2\sigma_\theta^2}} \quad (18)$$

where the mean velocity $\mu_v = a/v_{max}(a+b)$, and a, b, c are the Beta distribution parameters, with $c = \frac{(a+b+1)!}{(a-1)!(b-1)!}$, seems to agree well in many cases with real wind data. The nice characteristics of a Beta distribution, over a Gaussian distribution for example, is that the distribution is defined only on a limited interval which is physically more realistic, and the function can also be skewed to various degrees by

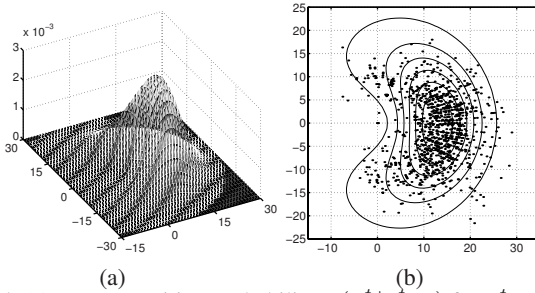


Fig. 2. Motion model: (a) target transition probability, $p(\mathbf{x}_k^t | \mathbf{x}_{k-1}^t)$ for $\mathbf{x}_{k-1}^t = [0, 0]$, and (b) its corresponding contour plot with actual wind data

adjusting the parameters a and b to match the actual data. Figure 2a shows a 3D plot of the target transition probability where $a = 4$, $b = 5$, $\sigma_\theta = \frac{\pi}{4}$ and $v_{max} = 30$ m/s. Figure 2b shows the contour plot of the function in good agreement with real wind data². For the problem described in this paper the same parameters were used except that the maximum wind velocity was set to 60 m/s giving a mean velocity of about 20 m/s (10 knots). Notice though that applying the convolution of the target prior PDF with the motion model multiple times is the same mathematically as convolving the motion model with itself multiple times and then convolving the results with the prior target PDF. The convolutions of the motion model with itself renders it more and more Gaussian like, even if the function was really far from being a Gaussian in the beginning. Therefore, for a very long searching plan, or for the case where observations only come very sporadically, a Gaussian approximation to the motion model is satisfactory.

4.3 Observation Model

The observation or sensor model is a probabilistic function representing the likelihood of the target being detected, or not ($\mathbf{z}_k = D$ or \overline{D}), conditioned on the sensor location and the state of the world.

It is not a trivial task to accurately model the sensor as many factors affect its performance: the distance to the target, the target footprint and reflectance, the transmission attenuation, and other environmental factors such as temperature, clutter and obstructions, etc.

For the purpose of this paper an active sensor model such as a downward looking millimeter wave radar was selected. It is assumed that the life raft has a radar reflector mounted on its canopy. For such a sensor, the approximate signal power, S' , received at the antennae after illumination of a target located at a distance d can be described by the following formula:

$$S' = C \frac{S A_{ant} A_t \rho}{16\pi^2 d^4} e^{-2\alpha d} \tag{19}$$

² Wind data measured at the MIT sailing pavillon on the Charles River, Cambridge, MA. Thanks to Eric Wile. <http://cbiwind.org>

where S is the emitted power, A_t and A_{ant} are the target and antennae footprints respectively, ρ is the target backscattering coefficient and α is the transmission attenuation factor which is greatly affected by the size, and density of the particles (e.g. rain) in the atmosphere. The constant C accounts for other environmental factors (e.g. background noise, temperature, etc) and could be a function of d .

If the probability of target detection is a function of the received power and the signal-to-noise ratio, then the following expression should hold true

$$\frac{P}{P_{std}} = \frac{S'}{S'_{std}} \Rightarrow P = P_{std} \frac{S'}{S'_{std}} \quad (20)$$

where by design, the reference, or ‘standard’ detection likelihood, P_{std} has a value of one (or less) for a given amount of received signal power S'_{std} evaluated at $\{d_{std}, \alpha_{std}\}$. Hence, by plugging (19) into the right side of (20), and after reduction, a closed form expression for the detection likelihood is obtained:

$$P = P_{std} \frac{d_{std}^4}{d^4} e^{-2(\alpha d - \alpha_{std} d_{std})} = p(\mathbf{z}_k = D_k | \mathbf{x}_k^t) \quad (21)$$

where the distance parameter $d (= \sqrt{h^2 + r^2})$ is a function of the vehicle altitude h and the "ground" range r ($r^2 = (x^t - x_k^s)^2 + (y^t - y_k^s)^2$) to the target. In this paper the following parameter values were used: $P_{std} = 0.8$, $d_{std} = 250$, $h = 250$, and $\alpha = \alpha_{std} = 1/250$. Figure 3 illustrates the corresponding detection likelihood and its complement for a case where the sensor is located above $x = y = 0$. Another




Fig. 3. Observation model: (a) conditional detection likelihood, $p(\mathbf{z}_k = D | \mathbf{x}_k^t)$ for $\mathbf{x}_k^s = [0, 0]$; (b) conditional detection complement likelihood (likelihood of ‘miss’), $p(\mathbf{z}_k = \bar{D} | \mathbf{x}_k^t) = 1 - p(\mathbf{z}_k = D | \mathbf{x}_k^t)$

very important parameter, not considered, that would contribute to a decrease in the detection likelihood with the ‘ground’ distance would be the height and wavelength of the seas. In this paper it is assumed that the radar reflector is always above the wave crests ensuring a direct line of sight to the emitting antennae.

Notice that in general, the detection likelihood of (21), $p(\mathbf{z}_k = D | \mathbf{x}_k^t)$, should be conditioned on the uncertain sensor state, \mathbf{x}_k^s , and written $p(\mathbf{z}_k | \mathbf{x}_k^t, \mathbf{x}_k^s)$. Hence, it should be convolved with the latest sensor state pdf, $p(\mathbf{x}_k^s | \mathbf{z}_{1:k}^s)$, to obtain $p(\mathbf{z}_k | \mathbf{x}_k^t)$ prior to using it in the update equation (2). In this paper, perfect localization is assumed so $p(\mathbf{z}_k | \mathbf{x}_k^t, \mathbf{x}_k^s) = p(\mathbf{z}_k | \mathbf{x}_k^t)$.

On a practical note when implementing an observation model, it is important that the function be smooth, and that it decreases progressively to zero without any steps. Otherwise, the objective function becomes jagged, effectively creating a multitude of local minima along the function. This quantization effect is due to the discretization of the target state PDF over the grid and has a very adverse effect for the convergence of the control optimization algorithm making it very difficult to obtain, if at all, the proper control value.

Also, because of the various assumptions made when modelling the observation likelihood, one must be aware of the possibility of discrepancies between the computed results and what would be the actual probability of detection. For computing accurately the ‘cumulative’ probability of detection (9), one would have to use an accurate observation model obtained through extensive in situ experimental testing of the search sensor. Nevertheless a theoretical model, as obtained in (21), provides a reasonable approximation for P_k , and is certainly sufficient for planning purposes, as well as for evaluating different solutions and comparing between them.

4.4 Results

For all the results presented in this section, the initial target PDF is assumed to be a symmetric Gaussian distribution centered at $x = y = 0$ with a standard deviation of 500m, and the searching vehicle is flying at an altitude of 250m, with the following initial pose $\mathbf{x}_0^s = [x_o^s = -900, y_o^s = -900, \theta_o^s = 0]$.

Stationary Target Figure 4 shows the resulting ‘greedy’ (1-step lookahead) search trajectory and the corresponding 3D views of the target PDF evolution at different stages as the search progresses from 0 to 300 seconds. Although this solution is very cheap computationally it often produces reasonable plans as it corresponds to maximizing the local payoff gradient. However because of the myopic planning, the vehicle fails to detect higher payoff values outside its sensor range and would keep spiraling further and further away from the center as can be seen on Fig. 4d. Figure 4e displays in solid line the conditional probability of detection ($1 - q_k$) obtained at every time step t_k . The dashed line represents the actual probability p_k that the target gets detected for the first time on that time step, which is the same as the solid line but discounted by Q_{k-1} . Notice the peaks in both functions as the search vehicle flyby over a mode in the PDF. Figure 4f shows the ‘cumulative’ probability P_k that the target has been detected by time step t_k . It is obtained from the integration of the payoff function (dashed line) from Fig. 4e. Another phenomenon to notice about the greedy search is the fact that because the volume under the PDF is always equal to one, as the vehicle traverses a mode of the function (e.g. when it crosses the original PDF mode for the first time (Fig. 4a), it has the effect of pushing away the probability mass hence increasing the entropy of the distribution, consequently making it harder and harder to increase the utility as time passes. The phenomenon will be referred to as the scattering effect.

Intuitively, for a given fixed trajectory length, one could imagine that instead of rushing to the PDF’s peak as in the greedy solution, the optimal strategy would be




Fig. 4. Greedy search for a static target: (a) to (d) 3D views of the vehicle trajectory and updated target PDF at time $t_k = 25, 60, 180$ and 300 respectively; (e) conditional (*solid line*) and ‘discounted’ (*dashed line*) probability of detecting the target on step k ($p(\mathbf{D}_k | \mathbf{z}_{1:k}) = 1 - q_k$, and $p_k = Q_k(1 - q_k)$); (f) accumulated probability of detection P_k




Fig. 5. Trajectory optimization: (a), (b) quasi-optimal path for a 120s search (12 control parameters maintained for 10s each) at time $t_k = 60$ and 120 respectively; (c) comparison between P_k evolutions (*top*), and control selections $u(k)$'s (*bottom*) for the ‘greedy’ solution (*solid line*), and the quasi-optimal solution (*dashed line*)

to circle around the peak but without flying over it, in such a manner as to plow the probability mass towards the peak, effectively compressing it (reducing the entropy), in order to increase the payoff of the last observations. In fact, as shown on Figs. 5a and 5b, this is exactly what happens. The piecewise constant ‘optimal’ control solution with 12 parameters, for a 120s trajectory, shows the path spiraling in instead of spiraling out. The comparison between the utility function evolutions (Fig. 5c) shows what one would anticipate. The greedy solution first gets a head start as it goes straight to the peak to finish with $P_{120} = .59$, but the ‘quasi-optimal’ solution progresses steadily to ultimately finish with $P_{120} = .77$, a 29% increase.

Drifting Target This section demonstrates the method for a drifting target with a process model as described in Sect. 4.2. The optimization technique is the same

used as for the static target, but the computational costs are increased by a few fold as the convolution operation needed for the target prediction stage is quite costly. This is also compounded by the fact that because the PDF is moving, a larger grid is necessary, making it even more costly to perform the convolution and the optimization. Nevertheless, the greedy solution is still very effective. The 3D plots of the search evolution are shown on Fig. 6. This time, though because of the lack of anticipation intrinsic to the greedy solution, it does not quite reach as high a level of cumulative detection probability (solid line on Fig. 8) as it did for the static case ($P_{300} = .71$ vs. $.91$). Figure 7 shows the results for a piecewise constant control




Fig. 6. Greedy (1s lookahead) search for a drifting target: (a) to (d) 3D views of the searching vehicle trajectory and updated target PDF at time $t_k = 60, 120, 180,$ and 300 respectively

solution with a time horizon of 30s split into three parameters and recomputed every 10s as in a feedforward control strategy. Comparing Fig. 7a with Fig. 6a really shows the positive effect of anticipation. This effect is also seen in the ultimate value of P_k ($.87$ vs. $.71$), a gain of over 22% (Fig. 8). The computational cost though is about 25 times higher.

5 Summary and Future Work

This paper introduced a general Bayesian framework for the searching problem of a single target. The approach presented explicitly considers the search vehicle kinematics, the sensor detection function, as well as the target arbitrary motion model. It was demonstrated to find efficient search plans that maximize the probability of finding the target given a fixed time limit by maintaining an accurate target location PDF of general form, and by explicitly modelling the target's process model.

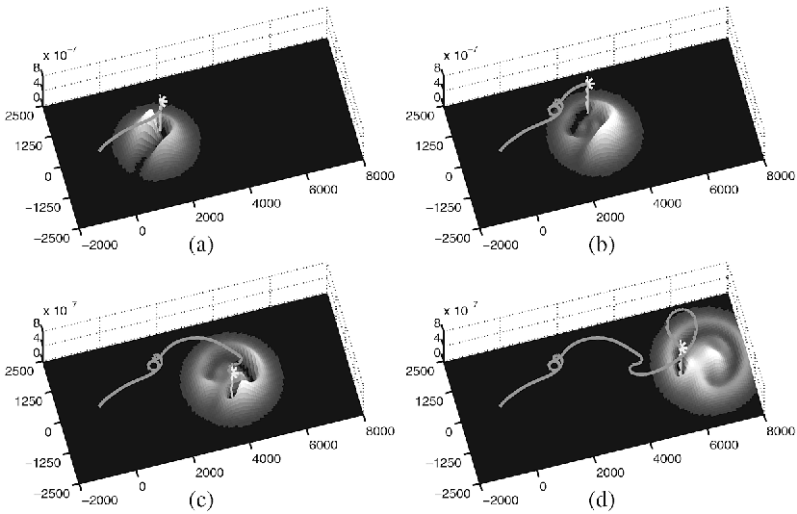


Fig. 7. Feedforward piecewise constant parametric control for a drifting target: 3D views of vehicle trajectory and updated target PDF at time $t_k = 60, 120, 180,$ and 300 respectively. The planning is done with 30s lookahead (3 control parameters each maintained for 10s) and replanned every 10s.




Fig. 8. Drifting target results: comparison between the detection probability P_k results from the greedy, 1s lookahead (*solid line*), and the piecewise constant (*dashed line*) solution from Figs. 6 and 7

The control solutions presented included the special case of a one step lookahead solution. This greedy solution demonstrated quite sensible trajectories for a very low computational cost. The ‘quasi-optimal’ solution for the static target was obtained with a piecewise constant control parametrization. It showed that the optimal solution for a given trajectory length initially delays its reward and tries to concentrate the probability mass into one location, reducing the entropy, and hence delaying the “scattering” effect, in order to reap greater benefits later in time. For the drifting target, it was shown that increasing the lookahead depth over the greedy solution improved the trajectory efficiency, at the expense of greater computational costs, by providing it with a sense of ‘anticipation’.

Having a sensor range much smaller than the searching area cause the target PDF to rapidly become very non-Gaussian even if it was originally the case. Because of the nature of the search problem it is very important to be able to keep track of the complete target distribution. Any grid based approach such as the one followed in this

paper is intrinsically subject to the curse of dimensionality, and as soon as one needs to increase the search area, the resolution of the grid, or the number of dimensions in the state-space, computational costs tend to get out of hand. As part of the ongoing research effort, techniques such as Monte Carlo methods, or particle filters [4], as well as the so called kernel methods are being investigated. A decentralized version of the Bayesian framework presented in this paper for a multiple vehicle search is also part of ongoing investigations.

Beyond the demonstration of the approach on a single and then multiple real autonomous platforms, the ultimate objective of this research is to eventually have multiple platforms participating in actual search and rescue (SAR) missions with real-time cooperative planning and fully integrated human in the loop inputs. As shown by the results presented, the technique has the potential to greatly improve on current SAR protocols, which in turn could be critical in saving human lives.

Acknowledgement

This work is partly supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government. The authors wish to thank Ali Göktoğan from ACFR, the developer of the RMUS simulator, for his assistance with the simulation implementation. Also, thanks to Erik Wile from MIT for the wind data.

References

1. J.O. Berger. *Statistical decision theory and Bayesian analysis*. Springer series in statistics. Springer-Verlag, New York, 2nd edition, 1985.
2. T. Furukawa. Time-subminimal trajectory planning for discrete nonlinear systems. *Engineering Optimization*, 34:219–243, 2002.
3. A.H. Goktogan, E. Nettleton, M. Ridley and S.Sullarieh. Real time multi-uav simulator. In *IEEE International Conference in Robotics and Automation*, Taipei, Taiwan, 2003.
4. N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings-F*, 140(2):107–113, April 1993.
5. H.J.W. Lee, K.L. Teo, V. Rehbock, and L.S. Jennings. Control parametrization enhancing technique for time optimal control problems. *Dyn. Sys. and Appl.*, 6(2):243–262, April 1997.
6. J.S. Przemieniecki. *Mathematical Methods in Defence Analyses*. AIAA Education Series. American Institute of Aeronautics and Astronautics, Inc., Washington, DC, 2nd edition, 1994.
7. L.D. Stone. *Theory of Optimal Search*, volume 118 of *Mathematics in Science and Engineering*. Academic Press, New York, 1975.
8. L.D. Stone, C.A. Barlow, and T.L. Corwin. *Bayesian Multiple Target Tracking*. Mathematics in Science and Engineering. Artech House, Boston, 1999.
9. K.L. Teo, C.J. Goh, and K.H. Wong. *A Unified Computational Approach to Optimal Control Problems*. Longman Scientific and Technical, 1991.
10. Debbie Whitmont. *An Extreme Event. The compelling, true story of the tragic 1998 Sydney-Hobart Race*. Random House Pty Ltd, Sydney, Australia, 1999.