Auton Robot (2008) 24: 387–399
DOI 10.1007/s10514-007-9081-4

Dynamic space reconfiguration for Bayesian search and tracking with moving targets

Benjamin Lavis · Tomonari Furukawa · Hugh F. Durrant Whyte

Received: 18 April 2007 / Accepted: 17 December 2007 / Published online: 12 January 2008

Abstract This paper presents a technique for dynamically reconfiguring search spaces in order to enable Bayesian autonomous search and tracking missions with moving targets. In particular, marine search and rescue scenarios are considered, highlighting the need for space reconfiguration in situations where moving targets are involved. The proposed technique improves the search space configuration by maintaining the validity of the recursive Bayesian estimation. The advantage of the technique is that autonomous search and tracking can be performed indefinitely, without loss of information. Numerical results first show the effectiveness of the technique with a single search vehicle and a single moving target. The efficacy of the approach for coordinated autonomous search and tracking is shown through simulation, incorporating multiple search vehicles and multiple targets. The examples also highlight the added benefit to human mission planners resulting from the technique's simplification of the search space allocation task.

Keywords Recursive Bayesian estimation · Search and tracking

1 Introduction

Picture a common marine search and rescue scenario involving the crew of a ship in distress and with failing navigation systems, dispatching a final mayday signal and boarding life rafts. The crew can then only wait for discovery and rescue, all the while being buffeted by strong winds and waves, and drifting due to ocean currents. In such a case the objective for authorities is to rescue the crew before their survival expectancy vanishes, requiring an efficient search for the life rafts and the prompt and safe rescue of the victims (IMO/ICAO 2006).

The use of a team of fast autonomous unmanned aerial vehicles (UAVs) together with a team of autonomous helicopters bearing rescue crews has been presented as an implementable robotics approach to marine search and rescue (SAR) (Wong et al. 2005). In such an application the use of multiple autonomous agents can produce a significant reduction in search times. If the UAVs have the additional ability to track the found life rafts as they drift, providing high powered lighting and/or gathering environmental information, then the safety of the rescue operation may be improved. As a result, search and tracking can be referred to as the two indispensable control objectives of autonomous vehicles for successful marine SAR missions.

The majority of the work on search and tracking to date has been carried out in the independent fields of either search or tracking. Many of the fundamental issues of search theory were first posed by B. O. Koopman and colleagues in the Antisubmarine Warfare Operations Research Group (ASWORG) during World War II (Dobbie 1968). Since the search problem is primarily concerned with the area to be searched, initial studies simplified the search problem to an area coverage problem. The introduction of the probability of detection along with advances in computational hardware led to more optimal allocation of search effort (Stone 1977, 1989a, 1989b). Later years saw the implementation of recursive Bayesian estimation (RBE) for manned SAR and anti-submarine search operations (Richardson et al. 2003).

B. Lavis (✉) · T. Furukawa
The School of Mechanical and Manufacturing Engineering,
The University of New South Wales, Sydney 2052, Australia
e-mail: benjamin.lavis@student.unsw.edu.au

H. F. Durrant Whyte
The Australian Centre for Field Robotics,
The University of Sydney, Sydney 2006, Australia

RBE techniques recursively update and predict the probability density function (PDF) of the target's state with respect to time. More recently, techniques have been formulated for decentralized search using multiple vehicles (Bourgault et al. 2003) and optimal autonomous search using the grid-based method for RBE (Bourgault et al. 2004).

On the other hand, target tracking, which initially consisted of simple feedback motion tracking, has evolved with the development of a variety of RBE techniques such as the Kalman filter (KF) (Salmond 2001), the extended Kalman filter (EKF) (Nordsjo 2005), sequential Monte Carlo (SMC) methods (Hue et al. 2002), sequential quasi-Monte Carlo (SQMC) methods (Julier et al. 2000; Guo and Wang 2006), and their variants (Schumitsch et al. 2005; Kräußling et al. 2006). These techniques for tracking seek computational efficiency in representing the sharp and often near-Gaussian PDF of an observable target, with little thought about representing the boundary of the target search space, which is an important consideration for search missions. While the KF and EKF represent the target PDF with a mean and a covariance matrix, the SMC and SQMC methods represent it with a set of particles, which moves freely with a resampling technique such as sequential importance sampling (Siegmund 1976; Doucet and Godsill 2000).

Recently, as a result of a general interest in multi-objective missions, in particular amongst the UAV community, several approaches to unified search and tracking have appeared (Butenko et al. 2004). Whilst the majority of these approaches have consisted of collections of independent search and tracking techniques, the authors proposed a unified *search-and-tracking* (SAT) approach within the RBE framework (Furukawa et al. 2006). This approach used the grid-based method, and subsequently the element-based method (Furukawa et al. 2007), for RBE, as these methods accurately represent the search space with a small number of sample points, unlike Monte Carlo based techniques, and are thus beneficial to SAT. However the maintenance of a large search space, necessary to include all the possible states of the moving targets, yields an excessively large amount of computational effort, without guaranteeing the inclusion of moving targets in the search space in a longer time horizon.

Meanwhile, reconfiguration of the search space boundary has been studied under the name of frontier based exploration (Yamauchi 1997), which identifies unknown areas in the entire space and allows robots to explore them. Frontier based exploration is a long established and popular approach to exploration (Thrun et al. 2005, Chap. 17), which has been extended to encompass exploration with multiple agents (Yamauchi 1998; Burgard et al. 2000; Tovar et al. 2006). It is generally associated with occupancy grid maps, but as a concept can be applied to other useful representations such as topological maps (Choset and Nagatani 2001) or manifolds (Howard et al. 2006), provided the representation allows for a distinction between known areas and unknown areas (Burgard et al. 2005). The applications of frontier based exploration, however, have been limited to information gathering tasks, where the coverage of unexplored areas and the continuous expansion of the explored areas are the primary concerns.

This paper presents a technique that performs dynamic reconfiguration of search spaces in order to estimate the state of moving targets during SAT operations. The proposed technique considers the reachable states of the moving targets and reconfigures the search space in order to maintain the validity of the PDF describing them, and to remove redundancy in the representation. The technique may be applied using any RBE technique that represents a bounded search space, including the grid-based and element-based methods. As such the technique is valid for all types of PDFs, including multimodal and non-Gaussian PDFs. Furthermore, the technique is formulated such that the established objective function for coordinated Bayesian SAT requires no alteration.

This paper is organized as follows. Section 2 covers the unified approach to SAT using RBE. Numerical methods for RBE are presented in Sect. 3. Section 4 defines the proposed space reconfiguration technique for SAT with moving targets. Section 5 presents a number of numerical examples which demonstrate the efficacy of the proposed technique. Conclusions and future considerations are discussed in the final section.

2 Recursive Bayesian search and tracking

Recursive Bayesian estimation is a basis for estimating nonlinear non-Gaussian models. This section presents the vehicle and observation models, and the fundamental steps required for RBE in coordinated SAT.

2.1 Target and sensor platform models

In order to successfully apply RBE for autonomous SAT it is essential to have accurate models of both the target of the search and rescue mission (the lost entity) and the autonomous sensor platforms conducting the operation. The motion of a target, t, is given by the following discrete time equation

$$\mathbf{x}_{k+1}^t = \mathbf{f}^t(\mathbf{x}_k^t, \mathbf{u}_k^t, \mathbf{w}_k^t) \tag{1}$$

where $\mathbf{x}_k^t \in \mathcal{X}^t$ is the state of the target at time k, $\mathbf{u}_k^t \in \mathcal{U}^t$ describes the target's control inputs and $\mathbf{w}_k^t \in \mathcal{W}^t$ represents the 'system noise', which includes environmental influences on the target. In general, the state of the target describes its

two-dimensional location, but may also include other variables such as velocity.

An autonomous sensor platform, s, used to perform the search and rescue operation is assumed to accurately know its own state via global sensors such as a combination of GPS, a compass and an inertial measurement unit (IMU). The motion of the search vehicle with state $\mathbf{x}_k^s \in \mathcal{X}^s$ and control inputs $\mathbf{u}_k^s \in \mathcal{U}^s$ is therefore given by

$$\mathbf{x}_{k+1}^s = \mathbf{f}^s(\mathbf{x}_k^s, \mathbf{u}_k^s). \tag{2}$$

The observation made from the vehicle s at time k is ${}^s\mathbf{z}_k \in \mathcal{X}^t$ and is subject to observation noise, $\mathbf{v}_k^s$. The observation model is given by the formula

$${}^s\mathbf{z}_k = \mathbf{h}^s(\mathbf{x}_k^t, \mathbf{x}_k^s, \mathbf{v}_k^s). \tag{3}$$

Note that if the autonomous search vehicle only carries a single sensor for observation, then the terms 'sensor platform' and 'autonomous search vehicle' may be used interchangeably. It will be assumed that this is the case for the remainder of this paper.

2.2 Recursive Bayesian estimation

In general, RBE is concerned with maintaining two PDFs, known as the posterior distribution and the prediction. Two distinct stages are used: the update stage and the prediction stage. The update stage calculates the posterior distribution of the current state given a prior estimation of the state (based on the sequence of previous observations) and a new observation at the present time. The prediction stage calculates the PDF of the next state using the posterior density of the current state and a probabilistic Markov motion model.

Before these stages are discussed in greater detail, a number of terms must first be defined. The PDF of a continuous random variable X in n_x dimensional Euclidean space is $p(\mathbf{x})$, where $p(\mathbf{x}) \geq 0, \forall \mathbf{x} \in \mathbb{R}^{n_x}$, and satisfies the following two conditions

$$\Pr(\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{n_x}) = \int_{\mathcal{X}} p(\mathbf{x}) d\mathbf{x} \tag{4}$$

and

$$\Pr(\mathbf{x} \in \mathbb{R}^{n_x}) = \int p(\mathbf{x}) d\mathbf{x} \equiv 1. \tag{5}$$

The sequence of states of a sensor platform s from time step 1 to k is given by the term

$$\tilde{\mathbf{x}}_{1:k}^s \equiv \{\tilde{\mathbf{x}}_i^s | \forall i \in \{1, \ldots, k\}\}. \tag{6}$$

The tilde is used here to indicate an instance $\tilde{(\cdot)}$, of a variable $(\cdot)$. Additionally, let the sequence of observations made by sensor platform s from time step 1 to step k be given by

$${}^s\tilde{\mathbf{z}}_{1:k} \equiv \{{}^s\tilde{\mathbf{z}}_i | \forall i \in \{1, \ldots, k\}\}. \tag{7}$$

It is possible to estimate the posterior distribution of the target at any time step k, $p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s)$, given a prior distribution of the target, $p(\tilde{\mathbf{x}}_0^t)$, and the sequences of sensor platform states and observations, $\tilde{\mathbf{x}}_{1:k}^s$ and ${}^s\tilde{\mathbf{z}}_{1:k}$. This is achieved via the recursive application of the two stages of RBE, update and prediction, described below.

2.2.1 Update

The update stage considers a new observation ${}^s\mathbf{z}_k$ in light of the corresponding estimation of state based on the sequence of previous observations, $p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k-1}, \tilde{\mathbf{x}}_{1:k-1}^s)$, and calculates the posterior density $p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s)$. The update equation is given by

$$p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s) = \frac{p({}^s\tilde{\mathbf{z}}_k | \mathbf{x}_k^t, \tilde{\mathbf{x}}_k^s) p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k-1}, \tilde{\mathbf{x}}_{1:k-1}^s)}{\int_{\mathcal{X}^t} p({}^s\tilde{\mathbf{z}}_k | \mathbf{x}_k^t, \tilde{\mathbf{x}}_k^s) p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k-1}, \tilde{\mathbf{x}}_{1:k-1}^s) d\mathbf{x}_k^t} \tag{8}$$

where $p({}^s\tilde{\mathbf{z}}_k | \mathbf{x}_k^t, \tilde{\mathbf{x}}_k^s)$ is the observation likelihood given knowledge of the current target state. However, given that in SAT problems the target state is generally unknown, the observation likelihood $l(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_k, \tilde{\mathbf{x}}_k^s)$ is used instead. $l(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_k, \tilde{\mathbf{x}}_k^s)$ may be defined with reference to the probability of detection, $0 < P_d(\mathbf{x}_k^t | \mathbf{x}_k^s) \leq 1$. The "detectable region" ${}^s\mathcal{X}_d^t$ of the sensor platform s may be defined as

$${}^s\mathcal{X}_d^t = \{\mathbf{x}_k^t | \epsilon < P_d(\mathbf{x}_k^t | \tilde{\mathbf{x}}_k^s) \leq 1\} \tag{9}$$

where ϵ is a positive threshold value which determines a detection event. Therefore the likelihoods for SAT may be expressed as

$$l(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_k, \tilde{\mathbf{x}}_k^s) = \begin{cases} 1 - P_d(\mathbf{x}_k^t | \tilde{\mathbf{x}}_k^s) & \nexists {}^s\tilde{\mathbf{z}}_k^t \in {}^s\mathcal{X}_d^t, \\ p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_k^t, \tilde{\mathbf{x}}_k^s) & \exists {}^s\tilde{\mathbf{z}}_k^t \in {}^s\mathcal{X}_d^t \end{cases} \tag{10}$$

where search occurs when no observation exists within the detectable region, and tracking takes place if the observation is within the observable region. It is also important to note that when $k = 1$ the update is carried out using $p(\mathbf{x}_k^t | {}^s\tilde{\mathbf{z}}_{1:k-1}, \tilde{\mathbf{x}}_{1:k-1}^s) = p(\tilde{\mathbf{x}}_0^t)$.

2.2.2 Prediction

The prediction step applies the Total Probability Theorem to the target's probabilistic Markov model, $p(\mathbf{x}_{k+1}^t | \mathbf{x}_k^t)$ defined by (1), and the posterior distribution (8), in order to compute

the PDF for the next time step, $p(\mathbf{x}^t_{k+1}|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k})$. The prediction is performed using the Chapman-Kolmogorov equation

$$p(\mathbf{x}^t_{k+1}|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) = \int_{\mathcal{X}^t} p(\mathbf{x}^t_{k+1}|\mathbf{x}^t_k)\, p(\mathbf{x}^t_k|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) d\mathbf{x}^t_k. \tag{11}$$

2.3 Coordinated control for search and tracking

Sensor data fusion is one form of coordination for multiple sensor platforms. The conditional independence of observations from n_s sensors leads to the multiple-sensor observation likelihood

$$l(\mathbf{x}^{t_j}_k|{}^s\tilde{\mathbf{z}}_k, \tilde{\mathbf{x}}^s_k) = \prod_{i=1}^{n_s} l(\mathbf{x}^{t_j}_k|{}^{s_i}\tilde{\mathbf{z}}^{t_j}_k, \tilde{\mathbf{x}}^{s_i}_k) \tag{12}$$

where $\tilde{\mathbf{x}}^s_k = \{\tilde{\mathbf{x}}^{s_i}_k | \forall i \in \{1, \ldots, n_s\}\}$ and ${}^s\tilde{\mathbf{z}}^{t_j}_k = \{{}^{s_i}\tilde{\mathbf{z}}^{t_j}_k | \forall i \in \{1, \ldots, n_s\}\}$ represent the states of the n_s platforms at time k, and their corresponding observations of target t_j, respectively, and $l(\mathbf{x}^{t_j}_k|{}^{s_i}\tilde{\mathbf{z}}^{t_j}_k, \tilde{\mathbf{x}}^{s_i}_k)$ is the observation likelihood for sensor platform s_i given the observation ${}^{s_i}\tilde{\mathbf{z}}^{t_j}_k$. For fully connected sensors with lossless and delay free communication, each platform i can receive the individual likelihoods $l(\mathbf{x}^{t_j}_k|{}^{s_q}\tilde{\mathbf{z}}_k, \tilde{\mathbf{x}}^{s_q}_k)$, $\forall q \neq i$, and decentrally construct (12). Substitution of (12) into (8) in the place of $p({}^s\tilde{\mathbf{z}}_k|\mathbf{x}^t_k, \tilde{\mathbf{x}}^s_k)$, gives the update equation for multiple sensors. Prediction using (11) completes the general decentralized data fusion approach to RBE as presented by Bourgault et al. (2003).

3 Numerical methods for recursive Bayesian estimation

The implementation of the two stages of RBE, described in Sect. 2.2, essentially requires the evaluation of a function at an arbitrary point in the target space, and the integration of a function over the target space. This section details two numerical methods for RBE which are suitable for SAT. Both the grid-based method and the element-based method may be used to create a continuous representation of the target's PDF and both maintain the accuracy of the representation in areas of low probability density, which is important during the search phase.

3.1 The grid-based method

Target space representation Consider a two dimensional target space such that $\mathbf{x}^t = [x^t, y^t] \in \mathcal{X}^t$. The grid-based method can represent this space continuously by considering a set of grid cells (rather than grid points). An approximate target space $\mathcal{X}^g$, consisting of n_g grid cells, is described by

$$\mathcal{X}^g \equiv \{\mathcal{X}^g_1, \ldots, \mathcal{X}^g_{n_g}\} \approx \mathcal{X}^t, \tag{13}$$

where $\bigcup_{i=1}^{n_g} \mathcal{X}^g_i = \mathcal{X}^g$ and $\bigcap_{i=1}^{n_g} \mathcal{X}^g_i = \emptyset$. Furthermore, the center and area of each grid cell are denoted $\bar{\mathbf{x}}^g_i$ and A^g_i, $\forall i \in \{1, \ldots, n_g\}$, respectively.

Function evaluation and integration Given an approximate function $f^g : \mathcal{X}^g \to R$, the value of the function at a given point in the approximate search space, $\tilde{\mathbf{x}}^t \in \mathcal{X}^g$, can be computed using

$$f(\tilde{\mathbf{x}}^t) \approx f^g(\tilde{\mathbf{x}}^t) \equiv \sum_{i=1}^{n_g} f(\bar{\mathbf{x}}^g_i)\delta_i(\tilde{\mathbf{x}}^t - \bar{\mathbf{x}}^g_i), \tag{14}$$

where $\delta_i(\cdot)$ is the indicator function, defined as

$$\delta_i(\tilde{\mathbf{x}}^t - \bar{\mathbf{x}}^g_i) = \begin{cases} 1 & \tilde{\mathbf{x}}^t \in \mathcal{X}^g_i, \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

The integral of the function may also be computed using the approximate function, as follows:

$$I = \int_{\mathcal{X}^t} f(\mathbf{x}^t) d\mathbf{x}^t \approx \int_{\mathcal{X}^g} f(\mathbf{x}^t) d\mathbf{x}^t = \sum_{i=1}^{n_g} f(\bar{\mathbf{x}}^g_i) A^g_i. \tag{16}$$

3.2 The element-based method

Target space representation The element-based method continuously approximates the target space and PDF using irregularly shaped elements described by shape functions. Generally the target space is first defined by a number of nodes which are then connected so as to create elements. For two-dimensional search spaces the simplest such elements are linear triangular elements generated via Delaunay triangulation. However elements need not be limited by shape or linearity; triangular or quadrilateral elements and higher-order elements with more nodes are all possible, as shown in Figs. 1 and 2.

Let an approximate target space $\mathcal{X}^e$, consisting of n_e elements, be described by

$$\mathcal{X}^e \equiv \{\mathcal{X}^e_1, \ldots, \mathcal{X}^e_{n_e}\} \approx \mathcal{X}^t, \tag{17}$$

where $\bigcup_{i=1}^{n_e} \mathcal{X}^e_i = \mathcal{X}^e$ and $\bigcap_{i=1}^{n_e} \mathcal{X}^e_i = \emptyset$. As such, any point in the target space, $\tilde{\mathbf{x}}^t \in \mathcal{X}^t$ may be located in one of the

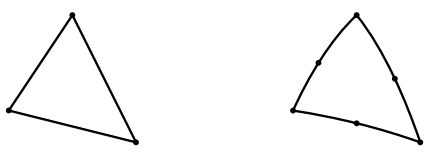

(a) 3 nodes (b) 6 nodes

Fig. 1 Triangular element types

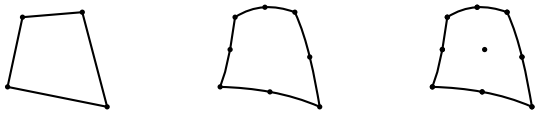

(a) 4 nodes (b) 8 nodes (c) 9 nodes

Fig. 2 Quadrilateral element types

elements. For a point in the ith element, $\tilde{\mathbf{x}}^t \in \mathcal{X}_i^e$, the point may be expressed in terms of the n_v nodes of the element. For nodes, $\check{\mathbf{x}}_{ij}^e = [\check{x}_{ij}^e, \check{y}_{ij}^e]^T, \forall j \in \{1, \ldots, n_v\}$, $\tilde{\mathbf{x}}^t$ may be expressed as

$$\begin{aligned} x^t &= \varphi_x(\xi, \eta) \equiv \sum_{j=1}^{n_v} \check{x}_{ij}^e N_j(\xi, \eta), \\ y^t &= \varphi_y(\xi, \eta) \equiv \sum_{j=1}^{n_v} \check{y}_{ij}^e N_j(\xi, \eta) \end{aligned} \tag{18}$$

where $N_j(\xi, \eta)$ is the shape function, which must satisfy

$$0 < N_j(\xi, \eta) < 1, \quad \sum_{j=1}^{n_v} N_j(\xi, \eta) = 1 \tag{19}$$

and $\xi \in \Xi = [\xi_{\min}, \xi_{\max})$ and $\eta \in H = [\eta_{\min}, \eta_{\max})$ are known as the natural coordinates.

The shape function and the ranges of the natural coordinates vary according to the type of element. In general the shape function takes the form

$$N_j(\xi, \eta) = \sum_{k=1}^{n_v} a_{jk} b_k(\xi, \eta) \tag{20}$$

where a_{jk} is a coefficient determined by the constraints (19) and $b_k(\xi, \eta)$ is the basis function of monomials in the natural coordinates. $b_k(\xi, \eta)$ may be determined using the binomial theorem:

$$\begin{aligned} &b_1(\xi, \eta) = 1, \\ &b_2(\xi, \eta) = \xi, b_3(\xi, \eta) = \eta, \\ &b_4(\xi, \eta) = \xi\eta, \\ &b_5(\xi, \eta) = \xi^2, b_6(\xi, \eta) = \eta^2, \\ &b_7(\xi, \eta) = \xi^2\eta, b_8(\xi, \eta) = \xi\eta^2, \\ &b_9(\xi, \eta) = \xi^2\eta^2, \\ &\ldots \end{aligned} \tag{21}$$

For triangular elements the ranges of the natural coordinates are $[\xi_{\min}, \xi_{\max}] = [0, 1]$ and $[\eta_{\min}, \eta_{\max}] = [0, 1 - \xi]$, and for quadrilateral elements the ranges of the natural coordinates are $[\xi_{\min}, \xi_{\max}] = [-1, 1]$ and $[\eta_{\min}, \eta_{\max}] = [-1, 1]$.

Function evaluation and integration Using the element-based method the evaluation of a function at a point and the integration of a function may be carried out by considering the natural coordinates. For a point in the ith element, $\tilde{\mathbf{x}}^t \in \mathcal{X}_i^e$, the natural coordinates of the point in the element may be determined using

$$\left[\tilde{\xi}, \tilde{\eta}\right]^T = \varphi^{-1}(\tilde{\mathbf{x}}^t) \tag{22}$$

where φ^{-1} is the inverse of the set of functions $\varphi = \{\varphi_x, \varphi_y\}$. The function value at $\tilde{\mathbf{x}}^t$ is then given by

$$f(\tilde{\mathbf{x}}^t) \approx f^e(\tilde{\mathbf{x}}^t) \equiv \sum_{j=1}^{n_v} f(\check{\mathbf{x}}_{ij}^e) N_j(\tilde{\xi}, \tilde{\eta}). \tag{23}$$

Integration is performed with respect to the natural coordinates according to the transformation

$$d\mathbf{x}^t = \det \mathbf{J}(\xi, \eta) d\xi d\eta \tag{24}$$

where $\mathbf{J}$ is the Jacobian matrix

$$\mathbf{J}(\xi, \eta) = \begin{bmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial y}{\partial \xi} \\ \frac{\partial x}{\partial \eta} & \frac{\partial y}{\partial \eta} \end{bmatrix}. \tag{25}$$

Integration over the target space is given by

$$I = \int_{\mathcal{X}^t} f(\mathbf{x}^t) d\mathbf{x}^t \approx \sum_{i=1}^{n_e} I_i^e \tag{26}$$

where I_i^e, the integral over an element, is

$$I_i^e = \int_{\mathcal{X}_i^e} f^e(\mathbf{x}^t) d\mathbf{x}^t = \int_{\Xi, H} f^e(\varphi(\xi, \eta)) \det \mathbf{J} d\xi d\eta. \tag{27}$$

Note that this integral is only analytically derivable for triangular elements with three nodes. In general Gauss integration may be used to numerically calculate the integral over each element.

4 Dynamic reconfiguration of search boundaries

In this section the proposed technique for dynamic search space reconfiguration is introduced. The requirements for a comprehensive approach to SAT are considered first, followed by a detailed description of the technique itself, and how it satisfies those requirements.

4.1 Requirements for comprehensive search and tracking

The previous sections have shown that the implementation of RBE for SAT requires numerical integration in order

to compute the posterior and prediction PDFs. The limitation of existing techniques is that they only consider static search spaces. With static search spaces a problem arises where it is possible for the target to move beyond the designated search space, that is the probability that the target is not within the search space at the next time step becomes non-zero: $\Pr(\mathbf{x}^t_{k+1} \notin \mathcal{X}^t) > 0$. In such cases the prediction $p(\mathbf{x}^t_{k+1}|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k})$, is no longer a valid PDF since

$$\begin{aligned}&\int_{\mathcal{X}^t} p(\mathbf{x}^t_{k+1}|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) d\mathbf{x}^t_{k+1}\\ &\quad = \Pr(\mathbf{x}^t_{k+1} \in \mathcal{X}^t)\\ &\quad = 1 - \Pr(\mathbf{x}^t_{k+1} \notin \mathcal{X}^t) < 1\end{aligned} \tag{28}$$

which violates condition (5), since the prediction PDF is only defined for $\mathcal{X}^t$. Therefore, the PDF is insufficient for maintaining an accurate estimate of the target's state at time step $k+1$, and a new search space $\mathcal{X}^t_{k+1}$ must be defined such that $\Pr(\mathbf{x}^t_{k+1} \in \mathcal{X}^t_{k+1}) = 1$. A comprehensive approach to SAT must therefore be able to recognize the potential for the target to move beyond the search space before it occurs and reconfigure the search space accordingly, such that both the accuracy of the estimate and the previously gathered information, are maintained.

4.2 Forward reachable sets

The state space that a target can reach within a certain time frame is called the target's reachable set. Of particular interest is the target's *forward reachable set* (FRS) which is the state space reachable forwards in time given an initial target state. The forward reachable set for target t from time $k+1$ to $k+n_k$, is denoted $\mathcal{A}^{\tilde{\mathbf{x}}^t}_{k+1:k+n_k}$. There are a number of methods for determining a vehicle's FRS, the most common being kinematic analysis. It is also possible to describe a target's FRS in terms of its probabilistic motion model, (1). The FRS evaluated at a point $\tilde{\mathbf{x}}^t \in \mathcal{X}^t_k$ may therefore be expressed as

$$\mathcal{A}^{\tilde{\mathbf{x}}^t}_{k+1:k+n_k} = \{\mathbf{x}^t_{k+1:k+n_k} | p(\mathbf{x}^t_{k+1:k+n_k}|\tilde{\mathbf{x}}^t_k) > 0\}. \tag{29}$$

4.3 Search space expansion based on reachable set analysis

In its usual sense, frontier based exploration involves a robot or robots moving to a frontier in order to gain some additional information about the environment, where a frontier is considered to consist of those areas of the search space about which some information is known and which are adjacent to space about which no information is known.

For SAT a frontier, $\mathcal{F}_k$, is defined as the set of points from which the target is able to, in the next time step, reach beyond the boundaries of the current search space:

$$\mathcal{F}_k = \{\mathbf{x}^t_k | \mathcal{A}^{\mathbf{x}^t_k}_{k+1:k+1} \not\subseteq \mathcal{X}^t_k\}. \tag{30}$$

For any frontier node, $\tilde{\mathbf{x}}^t_k \in \mathcal{F}_k$, if the value of the posterior distribution is greater than zero,

$$p(\tilde{\mathbf{x}}^t_k|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) > 0 \tag{31}$$

then there is some probability that the target will not be within the search space in the next time step. The direct result of this is that application of the prediction equation (11) will always result in an invalid PDF, as described by (28). In such cases, a new search space ${}^f\mathcal{X}^t_k$ should be defined, based on the present search space $\mathcal{X}^t_k$ and the target's FRS evaluated at those frontier nodes described by (31),

$${}^f\mathcal{X}^t_k \equiv \mathcal{X}^t_k \cup \mathcal{A}^{\tilde{\mathbf{x}}^{t+}_k \in \mathcal{F}_k}_{k+1:k+1} \tag{32}$$

where $\mathbf{x}^{t+} \in \{\mathbf{x}^t | p(\mathbf{x}^t|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) > 0\}$.

This new search space may be approximated, using either the grid-based or the element-based method, as follows:

$${}^f\mathcal{X}^{g/e} \equiv \{\mathcal{X}^{g/e}, {}^f\mathcal{X}^{g/e}_1, \ldots, {}^f\mathcal{X}^{g/e}_{n_{g/e}}\} \approx {}^f\mathcal{X}^t_k \tag{33}$$

where $\mathcal{X}^{g/e}$ is the grid or element based approximation of $\mathcal{X}^t$, and $n_{g/e}$ is the number of additional grid cells or elements required to approximate the new search space, such that $\mathcal{X}^{g/e} \bigcup_{i=1}^{n_{g/e}} {}^f\mathcal{X}^{g/e}_i = {}^f\mathcal{X}^{g/e}$ and $\mathcal{X}^{g/e} \bigcap_{i=1}^{n_{g/e}} {}^f\mathcal{X}^{g/e}_i = \emptyset$.

Before this new search space can be incorporated into the recursive Bayesian estimation, the posterior distribution must also be correctly defined over it. However this is made straightforward by observing that at time k, $\Pr(\mathbf{x}^t_k \notin \mathcal{X}^t_k) = 0$. Therefore put

$$p(\mathbf{x}^t_k|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) = \begin{cases} p(\mathbf{x}^t_k|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) & \text{if } \mathbf{x}^t_k \in \mathcal{X}^t_k,\\ 0 & \text{otherwise}\end{cases} \tag{34}$$

for $\mathbf{x}^t_k \in {}^f\mathcal{X}^t_k$. Using this newly defined search space and posterior distribution function the prediction stage may now be performed without loss of information using

$$\begin{aligned}&p(\mathbf{x}^t_{k+1}|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k})\\ &\quad = \int_{{}^f\mathcal{X}^t_k} p(\mathbf{x}^t_{k+1}|\mathbf{x}^t_k) p(\mathbf{x}^t_k|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}^s_{1:k}) d\mathbf{x}^t_k.\end{aligned} \tag{35}$$

The update stage should then be applied using $\mathcal{X}^t_{k+1} = {}^f\mathcal{X}^t_k$. It should also be noted that as there is no loss of information, the technique requires no change to the objective function for coordinated autonomous SAT, established in the literature.

4.4 Search space reduction

Whilst FRS based space expansion is theoretically capable of maintaining a complete representation of the target PDF for all time steps $k > 0$, the reality of limited computational capacity means that the unlimited expansion of search

spaces is not practically possible. If, however, those areas of the search space which provide no significant information are removed, then it may be possible to indefinitely maintain the target PDF using dynamic search spaces. It is logical to consider using an approach that is well suited to the method presented for FRS based expansion. Those subsets of the search space whose integrals are very small, effectively offer no information. Therefore, for some significance threshold, $0 < \varepsilon \ll 1$, and a grid cell or element, $\tilde{\mathcal{X}}_i^{g/e} \subset {}^f\mathcal{X}_k^t$, where

$$\Pr(\mathbf{x}_k^t \in \tilde{\mathcal{X}}_i^{g/e}) = \int_{\tilde{\mathcal{X}}_i^{g/e}} p(\mathbf{x}_k^t|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s) d\mathbf{x}_k^t < \varepsilon \tag{36}$$

then

$$\int_{{}^f\mathcal{X}_k^t} p(\mathbf{x}_k^t|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s) d\mathbf{x}_k^t - \int_{{}^f\mathcal{X}_k^t \setminus \tilde{\mathcal{X}}_i^{g/e}} p(\mathbf{x}_k^t|{}^s\tilde{\mathbf{z}}_{1:k}, \tilde{\mathbf{x}}_{1:k}^s) d\mathbf{x}_k^t < \varepsilon. \tag{37}$$

Furthermore, if this subspace is not *reachable* by the target in the next time step, that is if

$$\mathcal{A}_{k+1:k+1}^{\mathbf{x}_k^{t+}} \cap \tilde{\mathcal{X}}_i^{g/e} = \emptyset \tag{38}$$

then it may safely be eliminated from the search space without loss of information. So, if there are ${}^r n_{g/e}$ grid cells or elements $\tilde{\mathcal{X}}_i^{g/e}$ which satisfy conditions (36) and (38), the removable subspace ${}^r\mathcal{X}_k^t$ may be given by

$${}^r\mathcal{X}_k^t = \bigcup_{i=1}^{{}^r n_{g/e}} \tilde{\mathcal{X}}_i^{g/e}. \tag{39}$$

The update stage may then be applied using

$$\mathcal{X}_{k+1}^t = {}^f\mathcal{X}_k^t \setminus {}^r\mathcal{X}_k^t. \tag{40}$$

5 Numerical examples

This section presents three scenarios in order to demonstrate search space reconfiguration for autonomous SAT. The first two scenarios consider a single target and a single sensor in order to highlight the advantages of dynamic search spaces over static search spaces for SAT missions in general. The third scenario considers two lost targets and the use of three UAVs, to demonstrate the ease with which the approach extends to the coordinated SAT problem.

The targets in each of the scenarios share a common model. They move in the horizontal plane according to

$$\begin{aligned} x_{k+1}^{t_j} &= x_k^{t_j} + \Delta t \cdot v_k^{t_j} \cos\gamma_k^{t_j}, \\ y_{k+1}^{t_j} &= y_k^{t_j} + \Delta t \cdot v_k^{t_j} \sin\gamma_k^{t_j} \end{aligned} \tag{41}$$

Table 1 Vehicle model control limits

	Sensor platform	Target
Maximum speed [m/s]	20	10
Minimum speed [m/s]	10	0
Maximum turn rate [deg/s]	3	60

where $v_k^{t_j}$ and $\gamma_k^{t_j}$ are, respectively, the speed and direction of the target's motion due to external disturbances such as wind and current, subject to Gaussian noise. A time increment, $\Delta t = 60$ s, was used in each example. The sensors also move in the horizontal plane, and share the following model

$$\begin{aligned} x_{k+1}^{s_i} &= x_k^{s_i} + \Delta t \cdot v_k^{s_i} \cos(\theta_k^{s_i} + \gamma_k^{s_i}), \\ y_{k+1}^{s_i} &= y_k^{s_i} + \Delta t \cdot v_k^{s_i} \sin(\theta_k^{s_i} + \gamma_k^{s_i}), \\ \theta_{k+1}^{s_i} &= \theta_k^{s_i} + \gamma_k^{s_i} \end{aligned} \tag{42}$$

where $v_k^{s_i}$ is the speed of the sensor platform and $\gamma_k^{s_i}$ is the angle through which the platform turns. Table 1 shows the sensor platform and target control limits used during each of the simulations.

The probability of detection was given by a zero mean Gaussian distribution with constant covariance. The value of ϵ was set such that a detection occurs if the sensor comes within 1 km of the target. The observation likelihood was given by a zero mean Gaussian distribution with covariance proportional to the distance between the sensor platform and the target. A five step lookahead with receding horizon was used for control optimization.

In each of the simulations the sensors are indicated with a cyan circular marker and the targets with a yellow triangular marker. The search space is shown as a colored surface, representing the posterior distribution, against a grey background.

5.1 Single UAV, single target—static search spaces

Let us first consider two scenarios involving SAT over static search spaces. In both examples a single UAV undertakes a SAT mission for a single lost target over a 20 km × 20 km search space. The search space was represented using linear triangular elements according to the element-based method. The sensor's initial position is the same in both examples, $\tilde{\mathbf{x}}^{s_1} = [-10, 0]^T$, as is the initial distribution, which is Gaussian and centered at $[-1, 0]^T$, the only difference is the target's actual initial position. In the first example the target is located at $[-2.5, 3]^T$, whilst in the second example it is located at $[2, -1]^T$. Simulation results using these conditions are shown in Figs. 3 and 4 respectively.

It is clear from Figs. 3 and 4 that in both scenarios the sensor failed to track the target for the duration of the two hour mission. Figure 5 shows results useful in the analysis

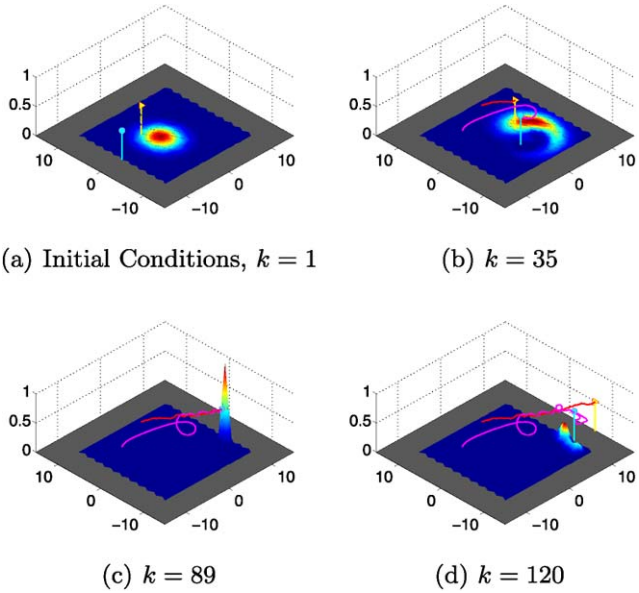

(a) Initial Conditions, $k = 1$ (b) $k = 35$

(c) $k = 89$ (d) $k = 120$

Fig. 3 Single sensor, single target SAT using a static search space—scenario 1

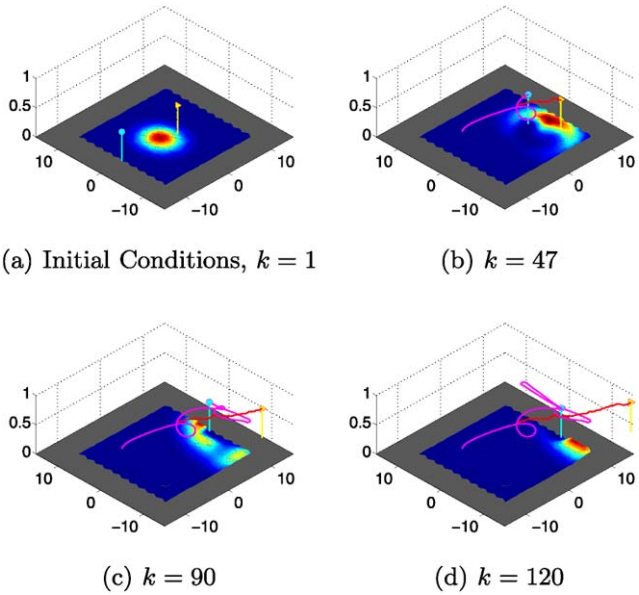

(a) Initial Conditions, $k = 1$ (b) $k = 47$

(c) $k = 90$ (d) $k = 120$

Fig. 4 Single sensor, single target sat using a static search space—scenario 2

of these failures. It can be seen that in the first scenario the search component was successfully completed, as the sensor moved to within detection range of the target at $k = 42$, resulting in a dramatic decrease in the information entropy. Even though the information entropy, which may be used as a measure of uncertainty in the sensor's belief, stayed relatively low for the remainder of the mission, the target drifted beyond the sensor's range at $k = 93$ and was not detected again. This is because the limitations of the static search space prevented the sensor from accurately estimating the

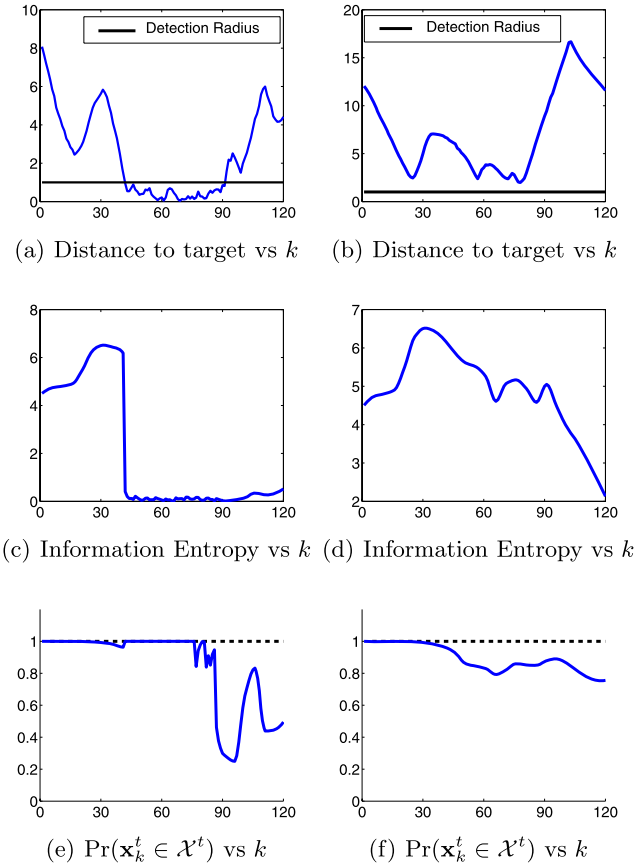

(a) Distance to target vs k (b) Distance to target vs k

(c) Information Entropy vs k (d) Information Entropy vs k

(e) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k (f) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k

Fig. 5 [L-R] scenario 1, scenario 2: single sensor, single target with a static search space

target's state. This can be seen in Fig. 5(c), which shows that the probability that the target was within the search space at $k = 93$ is less than 0.3.

The failure of the sensor in the second scenario is such that even though it came within three kilometers of the target on four separate occasions, see Fig. 5(b), it was unable to successfully locate the target before the target drifted beyond the search space. Again the information entropy was an insufficient measure of uncertainty, as it dropped dramatically despite the probability that the target was in the search space dropping too.

These two scenarios serve to highlight the limitations of static search spaces, in that even in the event of a successful search, it was possible for the sensor to permanently lose the target. In the second scenario the poor definition of the initial search space, had even worse results, in that the target remained lost and the search effort was completely unsuccessful.

5.2 Single UAV, single target—with search space expansion

The same scenarios were again simulated, this time with the inclusion of the FRS based space expansion method. Fig-

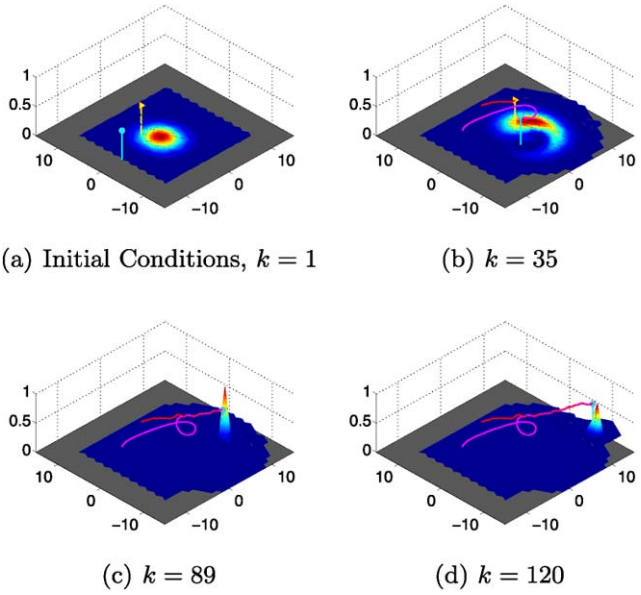

(a) Initial Conditions, $k = 1$ (b) $k = 35$

(c) $k = 89$ (d) $k = 120$

Fig. 6 Single sensor, single target SAT with FRS based space expansion—scenario 1

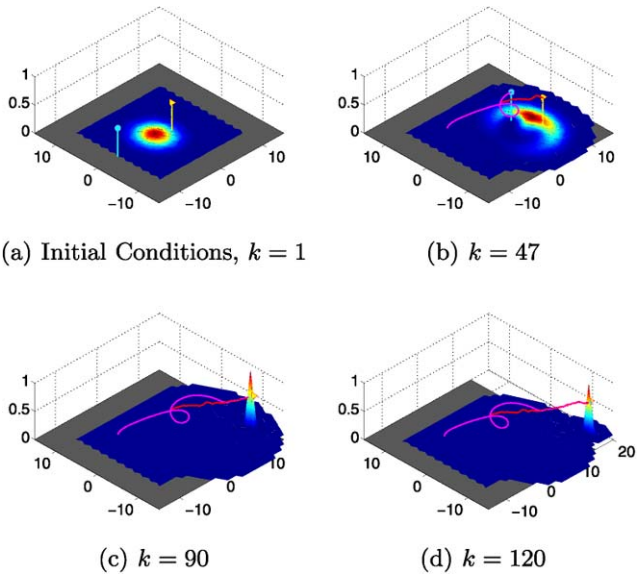

(a) Initial Conditions, $k = 1$ (b) $k = 47$

(c) $k = 90$ (d) $k = 120$

Fig. 7 Single sensor, single target SAT with FRS based space expansion—scenario 2

ures 6 and 7 show the simulations results based on Scenario 1 and 2, respectively.

These results and the results shown in Fig. 8 clearly show that in both cases the sensors were capable of not only finding the lost target, but also of keeping it within sensor range and tracking it for the remaining duration of the mission. In both cases detection of the target caused a sharp decrease in the information entropy, which stayed appropriately low as the probability that the target was in the search

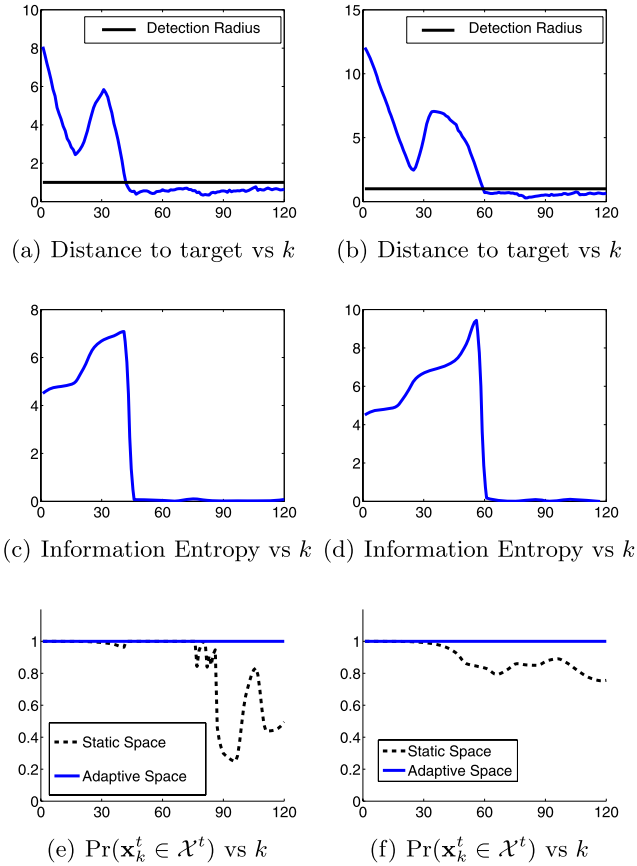

(a) Distance to target vs k (b) Distance to target vs k

(c) Information Entropy vs k (d) Information Entropy vs k

(e) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k (f) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k

Fig. 8 [L-R] scenario 1, scenario 2: single sensor, single target with space expansion

space remained at one, for the entire duration of both simulations.

These results clearly highlight the benefit of FRS based space expansion for SAT. In the first scenario the sensor detected the target at $k = 42$, just as in the previous example, however since the probability that the target was within the search space never dropped below one, the sensor was able to maintain an accurate belief of the target's state. The target was thereby successfully tracked from the time of detection to mission termination at $k = 120$, well beyond the failure point of $k = 93$ in the static search space example. In the second scenario the improvement in PDF representation due to FRS based expansion was such that the sensor successfully located the lost target at time step $k = 60$, despite the target going undetected in the previous example when the same initial search space was held static. Despite these improvements, it is important to note the large areas of extremely low probability density in Figs. 6(c) & (d) and Figs. 7(c) & (d). Evaluation of the update and prediction values for these areas is a computational burden which may be avoided using the search space reduction technique proposed in Sect. 4.4.

5.3 Single UAV, single target—with search space expansion and reduction

The scenarios from the previous example were again simulated, this time incorporating the search space reduction technique introduced in Sect. 4.4. The simulation results based on Scenario 1 and 2, are shown in Figs. 9 and 10, respectively. Again the sensor was successful in its search for, and subsequent tracking of, the lost target in both scenarios. The similarities between the results shown in Figs. 11(a)–(d) and the results of the previous example, Figs. 8(a)–(d), show that the proposed technique ensures that there is no loss of information as a result of search space reduction. Therefore the probability that the target was in the search space remained at a value of one for the duration of both simulations (Figs. 11(e) & (f)) and, importantly, the times required for the sensor to find the target in each scenario matched the search times required in the previous examples involving expansion only.

The advantage of the search space reduction technique may be seen in Figs. 11(g) and (h), which compare the area of the search spaces defined using static search spaces, the proposed FRS based expansion method, and the proposed search space reduction technique, for the two SAT scenarios, respectively. Whilst in the first scenario the FRS based expansion method allowed the target to be tracked for the duration of the mission, the search space grew by over 45% from 400 km^2 to 581.5 km^2, a figure that was increasing at the time of mission termination. Implementation of the search space reduction technique however, saw the search space grow to a maximum of only 455 km^2 during the search phase, and a reduction to an average of 37 km^2 during the tracking phase. At the conclusion of the mission, $k = 120$, this represents more than a 93% reduction of the search space generated using expansion alone.

The second scenario saw the search space grow by 83% from 400 km^2 to 732 km^2 using the expansion method alone. Implementation of the search space reduction technique resulted in a maximum search area of 535 km^2 during the search phase and the area again averaged 37 km^2 during the tracking phase, less than 6% of the search space at $k = 120$ generated using the expansion method alone.

5.4 Coordinated search and tracking—multiple targets

In this example three UAVs cooperatively search for and track two lost targets. The targets used were the same as the two targets used in the scenarios presented in the previous examples. Two sensors were introduced in addition to the sensor used in the previous examples. The initial locations of the new sensors were $\tilde{\mathbf{x}}^{s_2} = [-10, -2.5]^T$ and $\tilde{\mathbf{x}}^{s_3} = [-10, -5]^T$, such that the initial distance between each of the additional sensors and each of the targets was greater than the initial distance between the original sensor and the two targets. The simulation results are shown in Fig. 12.

These results show the generalization of the proposed technique to SAT problems involving multiple search vehicles and multiple targets. Furthermore, Figs. 12(e) and (f) highlight the advantage of cooperative SAT in reducing the time required for searching. Target 2, the target used in the second scenario in the previous examples and which went undetected when a static search space was used, was the

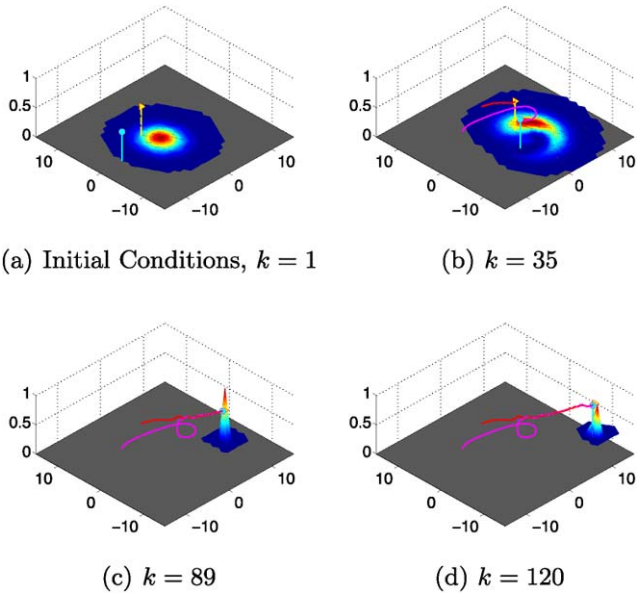

(a) Initial Conditions, $k = 1$ (b) $k = 35$

(c) $k = 89$ (d) $k = 120$

Fig. 9 Single sensor, single target SAT with FRS based expansion and search space reduction—scenario 1

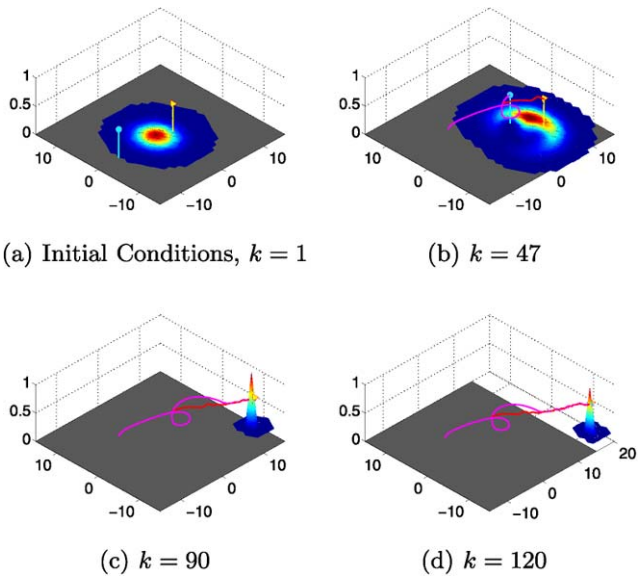

(a) Initial Conditions, $k = 1$ (b) $k = 47$

(c) $k = 90$ (d) $k = 120$

Fig. 10 Single sensor, single target SAT with FRS based expansion and search space reduction—scenario 2

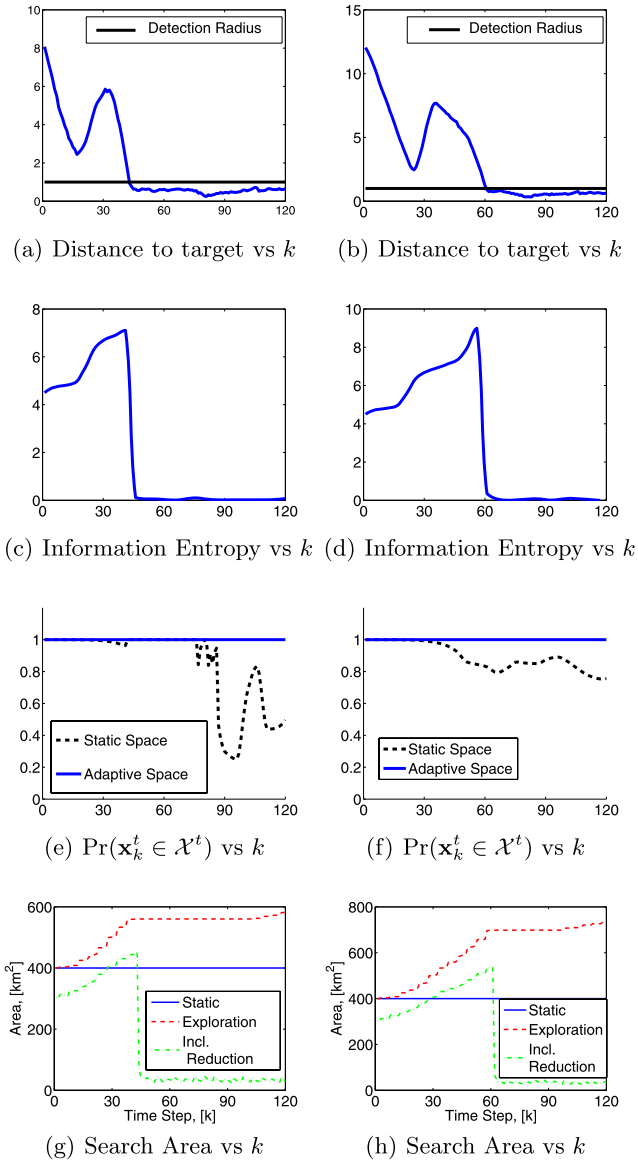

(a) Distance to target vs k (b) Distance to target vs k

(c) Information Entropy vs k (d) Information Entropy vs k

(e) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k (f) $\Pr(\mathbf{x}_k^t \in \mathcal{X}^t)$ vs k

(g) Search Area vs k (h) Search Area vs k

Fig. 11 [L-R] scenario 1, scenario 2: single sensor, single target with search space expansion and reduction

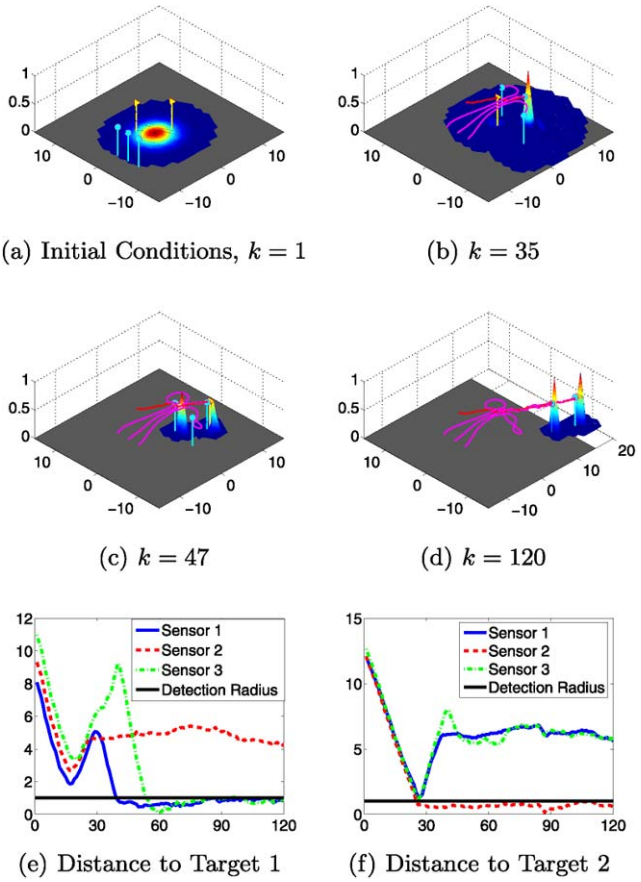

(a) Initial Conditions, $k = 1$ (b) $k = 35$

(c) $k = 47$ (d) $k = 120$

(e) Distance to Target 1 (f) Distance to Target 2

Fig. 12 Three sensors, two targets using search space expansion and reduction

first of the targets to be located. It was detected by Sensor 2 at $k = 26$, less than half the time required by the lone sensor in the previous examples, despite Sensor 2 being farther from the target than Sensor 1 at the start of the mission. Following the detection of Target 2, Sensor 2 commenced tracking, whilst Sensors 1 and 3 continued searching for Target 1. The lone sensor in the previous examples required 41 minutes ($k = 42$) to locate Target 1, however when cooperative search was used, Target 1 was located at $k = 40$. Therefore despite the presence of two targets, the use of three sensor platforms reduced the time required to locate each of the targets, and dramatically reduced the overall required search time. Once each target had been located, they were successfully tracked for the remainder of the mission.

6 Conclusions and future work

A technique which dynamically reconfigures the search space of moving targets during Bayesian SAT operations was presented. The technique is independent of the method of search space representation, may be decentrally applied and has been easily extended to scenarios involving multiple sensors and multiple targets. It was demonstrated through a number of examples, that during SAT, a sensor with search space expansion capabilities has an improved likelihood of succeeding in its mission compared with a sensor without such capabilities, especially over extended time periods. The proposed technique included a method for removing the redundant areas of the search space to improve the computational efficiency when using dynamic search spaces. A further advantage is a reduction of the burden placed on human mission planners during the initial stages of SAR operations, as search spaces need only to be defined on the basis of the targets' initial distributions, rather than on an estimation of the mission duration and all possible target locations in that time frame. These advantages are significant in scenarios such as marine SAR, where the consequences of mission failure may be dire.

The compatibility of the proposed technique with element-based representations presents an opportunity for future investigation of approaches from finite element theory such as remeshing and adaptive meshing. Approaches such as these might be used to improve the quality of the PDF representation and therefore improve the overall performance of the SAT operation. Furthermore, a hardware-in-the-loop simulator and hardware demonstrators are planned developments for this project. Four model helicopters are currently being equipped with autopilots and additional sensors, with the aim of demonstrating the proposed technique in the field.

References

Bourgault, F., Furukawa, T., & Durrant-Whyte, H. F. (2003). Coordinated decentralized search for a lost target in a Bayesian world. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 48–53).

Bourgault, F., Furukawa, T., & Durrant-Whyte, H. F. (2004). Process model, constraints and the coordinated search strategy. In *IEEE international conference on intelligent robots automation* (Vol. 5, pp. 5256–5261).

Burgard, W., Moors, M., Fox, D., Simmons, R. G., & Thrun, S. (2000). Collaborative multi-robot exploration. In *IEEE international conference on intelligent robots automation* (Vol. 1, pp. 476–481).

Burgard, W., Moors, M., Stachniss, C., & Schneider, F. E. (2005). Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, *21*(3), 376–386.

Butenko, S., Murphey, R., & Pardalos, P. M. (Eds.) (2004). *Recent developments in cooperative control and optimization*. Boston: Kluwer Academic.

Choset, H., & Nagatani, K. (2001). Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, *17*(2), 125–137.

Dobbie, J. M. (1968). A survey of search theory. *Operations Research*, *16*(3), 525–537.

Doucet, A., Godsill, S., & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, *10*(3), 197–208.

Furukawa, T., Bourgault, F., Lavis, B., & Durrant-Whyte, H. F. (2006). Recursive Bayesian search-and-tracking using coordinated UAVs for lost targets. In *Proceedings of IEEE international conference on robotics and automation* (pp. 2521–2526). Orlando, Florida.

Furukawa, T., Durrant-Whyte, H. F., & Lavis, B. (2007). The element-based method—theory and its application to Bayesian search and tracking. In *Proceedings of IEEE/RSJ international conference intelligent robots and systems* (pp. 2807–2812). San Diego, California.

Guo, D., & Wang, X. (2006). Quasi-Monte Carlo filtering in nonlinear dynamic systems. *IEEE Transactions on Signal Processing*, *54*(6), 2087–2098.

Howard, A., Sukhatme, G. S., & Matarić, M. J. (2006). Multirobot simultaneous localization and mapping using manifold representations. In *Proceedings of IEEE* (Vol. 94, pp. 1360–1369).

Hue, C., Le Cadre, J.-P., & Pérez, P. (2002). Sequential Monte Carlo methods for multiple target tracking and data fusion. *IEEE Transactions on Signal Processing*, *50*(2), 309–325.

IMO/ICAO (2006). *International aeronautical and maritime search and rescue manual: Vol. II. Mission coordination* (consolodated edn.). London & Montreal: International Maritime Orginization/International Civil Aviation Organization.

Julier, S., Uhlmann, J., & Durrant-Whyte, H. F. (2000). A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control*, *45*(3), 477–482.

Kräußling, A., Schneider, F. E., & Sehestedt, S. (2006). Tracking multiple objects using the Viterbi algorithm. In *Proceedings of international conference informatics in control, automation and robotics ICINCO-RA* (pp. 18–25), Setubal, Portugal.

Nordsjo, A. E. (2005). Target tracking based on Kalman-type filters combined with recursive estimation of model disturbances. In *Proceedings of IEEE international radar conference* (pp. 115–120).

Richardson, H. R., Stone, L. D., Monach, W. R., & Discenza, J. H. (2003). *Early maritime applications of particle filtering. SPIE* (Vol. 5204, pp. 165–174).

Salmond, D. (2001). Target tracking: introduction and Kalman tracking filters. In *Proceedings of IEE target tracking: algorithms and applications*. Vol. Workshop (Vol. 2, pp. 1/1–1/16) (Ref. No. 2001/174).

Schumitsch, B., Thrun, S., Bradski, G., & Olukotun, K. (2005). The information-form data association filter. In *Proceedings of conference on neural information processing systems (NIPS)*, Cambridge: MIT Press.

Siegmund, D. (1976). Importance sampling in the Monte Carlo study of sequential tests. *The Annals of Statistics*, *4*(4), 673–684.

Stone, L. D. (1977). Search theory: a mathematical theory for finding lost objects. *Mathematics Magazine*, *50*(5), 248–256.

Stone, L. D. (1989a). *Theory of optimal search*. Arlington: ORSA Books.

Stone, L. D. (1989b). What's happened in search theory since the 1975 Lanchester prize? *Operations Research*, *37*(3), 501–506.

Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic robotics*. Cambridge: MIT Press.

Tovar, B., Muñoz-Gómez, L., Murrieta-Cid, R., Alencastre-Miranda, M., Monroy, R., & Hutchinson, S. (2006). Planning exploration strategies for simultaneous localization and mapping. *Robotics and Autonomous Systems*, *54*(4), 314–331.

Wong, E.-M., Bourgault, F., & Furukawa, T. (2005). Multi-vehicle Bayesian search for multiple lost targets. In *Proceedings of IEEE international conference robotics and automation* (pp. 3180–3185).

Yamauchi, B. (1997). A frontier-based approach for autonomous exploration. In *Computational intelligence in robotics and automation* (pp. 146–151).

Yamauchi, B. (1998). Frontier-based exploration using multiple robots. In *Proceedings of 2nd international conference autonomous agents* (pp. 47–53).

Benjamin Lavis received the B.E. and University Medal in Mechatronic Engineering from the University of New South Wales, Australia, in 2004. He has been a Ph.D. student at the Computational Mechanics and Robotics Laboratory at the University of New South Wales since 2005. His research interests include multiple robot cooperation and probabilistic estimation for search and rescue.

Tomonari Furukawa received the B.Eng. in Mechanical Engineering from Waseda University, Japan, in 1990. He received a M.Eng. (Research) degree in Mechanical Engineering from the University of Sydney, Australia, in 1993 and a Ph.D. degree in Quantum Engineering and Systems Sciences from the University of Tokyo, Japan, in 1996. He was an Assistant Professor (1995–1997) and Lecturer (1997–2000) at the University of Tokyo, and Research Fellow (2000–2002) at the University of Sydney before joining the University of New South Wales, Australia, where he is currently a Senior Lecturer. His research interests include optimization, cooperative control and computational mechanics.

Hugh F. Durrant Whyte received the B.Sc. in Nuclear Engineering rom the University of London, United Kingdom, in 1983, and the MSE and Ph.D. degrees, both in systems engineering, from the University of Pennsylvania, in 1985 and 1986, respectively. From 1987–1995, he was a senior lecturer in engineering science, at the University of Oxford, and a fellow of Oriel College Oxford. From 1995–2002 he was a professor of mechatronic engineering at the University of Sydney. In 2002, he was awarded an Australian Research Council (ARC) Federation Fellowship. He also now leads the ARC Centre of Excellence for Autonomous Systems. His research focuses on autonomous vehicle navigation and decentralized data fusion.