

Bayesian approach to exploiting prior targeting information within a weapon seeker

Keith Copsey, Richard O. Lane, Sunil Manchanda and Andrew R. Webb

QinetiQ,
Malvern Technology Centre,
St Andrews Road,
Malvern,
Worcestershire, WR14 3PS,
United Kingdom.

K.Copsey@signal.QinetiQ.com, rlane1@qinetiq.com, smanchanda@qinetiq.com,
A.Webb@signal.QinetiQ.com

ABSTRACT

This paper considers an automatic target recognition (ATR) application in which a targeting sensor is used to guide a seeker-equipped weapon to an area containing high-value relocatable targets. The weapon seeker then needs to engage the high value targets, while minimising collateral damage. A Bayesian approach is proposed that enables the weapon seeker to exploit the targeting information before making its final decision. Specifically, the approach matches the scenes in the seeker domain with those from the targeting sensor, while taking into account uncertainty and data latency. The proposed solution utilises a Bayesian technique known as particle filtering. This paper outlines the approach, and presents results for a synthetic example. Future work will conduct a performance assessment using scenarios derived from real long-range and short-range SAR imagery.

1.0 INTRODUCTION

1.1 General

This paper considers the problem of using a targeting sensor to guide a seeker-equipped weapon to an area containing high-value relocatable targets. The aim is for the seeker-equipped autonomous weapon to exploit targeting information before making its final decision. This is illustrated in Figure 1, where targeting data at time t_0 is used to aid classification of the weapon seeker data at time t_1 . The generic problem addressed is one where two sensor images, separated in time, are available to classify relocatable targets in a scene. Issues to be addressed include:

- Uncertainty in the targeting identifications.
- Change in the target layout configuration during weapon fly-out (i.e. staleness of the targeting information).
- Differing imaging geometry between the targeting sensor and the seeker.

A particularly adverse effect of the last two items is that a target designated correctly by the targeting sensor may have both a different location and an altered signature by the time that the weapon has reached the targeted area. This will have a significant effect on the ability of the weapon to engage the pre-selected target, especially in typical scenarios where collateral damage must be minimised.




Figure 1: Exploitation of targeting information by a weapon seeker.

1.2 Bayesian approach

A Bayesian approach is proposed which uses a Particle Filter [4] to draw samples from the posterior distribution for the target locations and classes, given the information derived from the targeting sensor and weapon seeker. Since the posterior distribution contains the relevant information on the target locations and classes, the samples can be used as inputs to the final decision making process of the weapon. Successful production of this information will improve the ability of the weapon to engage the targets designated on launch of the weapon, while minimising collateral damage. The proposed approach fits within a larger framework detailed in an accompanying paper [17].

1.3 Bayesian approach

The main motivation behind a Bayesian approach [11] to the problem lies in the unique ability of Bayesian statistics to handle limited and possibly conflicting pieces of information in a fully consistent manner. In particular, Bayesian statistics provides a consistent mechanism for manipulating probabilities assigned to observed data. Further advantages to the use of Bayesian techniques include the ability to cope with additional prior information, perhaps elicited from expert knowledge, and the production of confidence intervals and other statistics for the parameters estimated.

1.4 Outline of the paper

The structure of this paper is as follows. Section 2 specifies the problem being examined. Section 3 proposes a Bayesian solution. Section 4 presents the results for a synthetic example. Conclusions and future work are given in Section 5.

2.0 PROBLEM SPECIFICATION

2.1 Introduction

In the considered scenario, an image from a targeting sensor is obtained at time t_0 . Target detection techniques are then applied to obtain a set of targeting detections. An image chip is obtained for each

detected object/target by centring an input window (of sufficient size to cover expected targets) at the location of each detection. ATR algorithms [16] are then applied to estimate the class of the object leading to each detection. If the target classes correspond to high value targets a seeker-equipped weapon is launched to engage the targets.

The weapon seeker reaches and images the highlighted area at a later time, t_1 . Similarly to the targeting sensor processing, target detection algorithms are then applied to obtain a set of seeker detections, along with associated image chips (centered on the locations of the detections). The task is to determine the locations and classes of the targets at time t_1 , utilising both the seeker information and the targeting information. The solution to this task lies in determining the posterior distribution of the locations and classes at time t_1 . For the purposes of this paper it is assumed that we are only interested in the targets detected in the targeting image.

2.2 Related work

Work by Gordon and Salmond [8] tackles the problem of matching target detections from a targeting sensor with a missile seeker, but assumes that no ID information can be inferred from the seeker or targeting measurements. Gaussian models for bulk and individual target motion during weapon fly-out were introduced, and a closed form solution was obtained. Work by the same authors on group and extended object tracking [13] hints at a non-linear approach to the same problem using particle filters [4], but does not take into account object characteristics.

The work described in this paper differs from previous approaches to matching target detections through:

- Estimation of full class probabilities, utilising the targeting and seeker sensor measurements of the objects/targets.
- Potential for modelling complicated target motion during the time-gap between the seeker measurements and the targeting measurements.
- No assumption that a target is correctly designated in the targeting image.

Work by Gordon et al [6] has developed a Bayesian approach to joint tracking and identification, which is relevant to the problem addressed within this paper. The focus within that work was ensuring efficiency for multiple sensor returns.

2.3 Targeting detections

The number of targeting sensor detections at time t_0 is denoted by N_t . The locations of the detections and associated image chips (ID sensor measurements) are denoted by l_1, \dots, l_{N_t} and r_1, \dots, r_{N_t} respectively. For notational ease we define $T_i = (l_i, r_i)$ for $i = 1, \dots, N_t$.

Assuming that there are J possible target classes, the ID sensor measurements are used to obtain J -dimensional class probability vectors ψ_i for each detection, where $\psi_{i,j}$ is the estimated probability that the i -th detection is the j -th class, for $i = 1, \dots, N_t$ and $j = 1, \dots, J$. Such class probabilities would be estimated using Bayesian ATR algorithms, or possibly via human intervention.

The measurement errors for the target locations are assigned Gaussian distributions, so that $l \sim N(x, \Sigma_l)$ where x is the actual target location, and Σ_l is the covariance matrix for the measurement errors. The covariance matrix should be determined by considering the sensor performance characteristics along with

the imaging conditions.

2.4 Seeker detections

The number of seeker detections at time t_1 is denoted by N_s . Since the targeting sensor indicates that there are N_t targets present, the threshold for detecting objects within the seeker image is assumed to be set so that $N_s \geq N_t$. Adaptation of the proposed approach to cope with $N_s < N_t$ would be trivial. The locations of these seeker detections are denoted y_1, \dots, y_{N_s} , and the associated image chips (ID sensor measurements) are z_1, \dots, z_{N_s} . For notational ease we define $D_i = (y_i, z_i)$ for $i = 1, \dots, N_s$.

An example of a DBS seeker image is provided in Figure 2. Input windows have been placed on the locations of an example set of seeker detections. The windows are colour-coded so that red indicates an actual target and green indicates the type of background clutter that might pass through the initial target detection stage. These input windows define the image chips that are extracted to provide ID sensor measurements.




Figure 2: Example of a DBS seeker image (range along the horizontal axis, cross range along the vertical axis). Red boxes highlight targets, green boxes highlight background clutter.

It is assumed that density models (conditional on class) can be estimated for the ID sensor measurements. Estimating these distributions given only limited training data for the weapon seeker is covered in the accompanying paper [17]. These distributions can be represented by $p(z | C = j)$, where z is the image chip and j is the index of the class C of the object.

In addition to probability densities for target image chips, it is assumed that a probability density has been estimated for image chips that correspond to the sort of background noise and clutter that will pass through the target detection algorithm. This density is denoted by $p(z | C = 0)$.

If the class of a target is unassigned, a mixture distribution is used for the ID sensor measurements:

$$p(z) = \pi_0 p(z | C = 0) + (1 - \pi_0) \sum_{j=1}^J \pi_j p(z | C = j) \quad (1)$$

where π_1, \dots, π_J represent the prior class probabilities excluding background clutter, and π_0 is the prior probability for background clutter. Note that the prior probability for background clutter will be related to the false alarm probability of the detection algorithm, rather than the ratio of background clutter to targets. This reflects the fact that the initial detection stage will already have eliminated most of the background noise.

The measurement errors for the object locations are assigned Gaussian distributions, so that $y \sim N(x, \Sigma_s)$ where x is the actual object location, and Σ_s is the covariance matrix for the measurement errors. As with the targeting sensor, the covariance matrix should be determined by considering the sensor performance characteristics together with the imaging conditions. Locations of any additional targets and background clutter are assumed to be distributed uniformly over the surveyed region.

3.0 BAYESIAN SOLUTION

3.1 Posterior distribution

The actual classes and locations of the targets detected by the targeting sensor at time t_0 are denoted by (c_1, \dots, c_{N_t}) and $(x_{0,1}, \dots, x_{0,N_t})$ respectively. By time t_1 the new locations are represented by $(x_{1,1}, \dots, x_{1,N_t})$. This reflects the fact that the targets may have relocated during the weapon fly-out time $t_1 - t_0$. The actual classes are of course unchanged. Using the definitions in Section 2 the posterior distribution of interest at time t_1 is:

$$p(x_{1,1}, \dots, x_{1,N_t}, c_1, \dots, c_{N_t} | T_1, \dots, T_{N_t}, D_1, \dots, D_{N_s}) \quad (2)$$

3.2 Prior Evolver

A *Prior Evolver* is used to update the information from the targeting sensor to allow for target motion during weapon fly-out. Specifically, this consists of predicting how the detections gleaned from the targeting sensor at time t_0 will have changed by the time t_1 that the weapon seeker views the targeted area. The Prior Evolver can be represented by a distribution:

$$p(x_{1,1}, \dots, x_{1,N_t} | x_{0,1}, \dots, x_{0,N_t}, c_1, \dots, c_{N_t}) \quad (3)$$

In the tracking literature [2] the Prior Evolver corresponds to the system model for the state.

The simplest non-trivial form for the Prior Evolver consists of independent Gaussian perturbations for each detection. Bulk motion of targets (perhaps reflecting the motion of a convoy) can be included easily, using a global Gaussian translation [7]. More complicated motion, incorporating knowledge of the terrain and likely target behaviour, is also possible. For example, targets can be predicted to follow a road network [1][10]. This is done by perturbing the locations using a Gaussian distribution with covariance matrix chosen so that the variance along the road is much higher than the variance orthogonal to the road. This has the effect of making the uncertainty along the road more than the uncertainty orthogonal to the road, which fits with the road motion constraints. Various rules can then be applied for targets near road junctions and for entry/exit conditions from roads. A further possibility is the construction of a potential-

field constraint, to bias predicted target motion in the direction of assumed desired target locations (such as potential hide locations) and away from impenetrable terrain (such as rivers). The tracking literature [12][15] contains examples of such an approach.

3.3 Particle Filter

Analytical calculation of the desired posterior distribution is feasible only for the simplest of measurement distributions and Prior Evolver. Thus, a Particle Filter is used to obtain the samples from the posterior distribution. The particle filter is an extension of importance sampling [5][14] to sequential sampling. For clarity, we now (briefly) describe the underlying idea behind importance sampling.

Suppose that we have a set of n independent samples, $\phi^{(1)}, \dots, \phi^{(n)}$, from a probability distribution with density function proportional to $g(\phi)$, but we are actually interested in making inference on a probability distribution with density function proportional to $f(\phi)$. If a set of unnormalised importance weights:

$$w^{(s)} = f(\phi^{(s)}) / g(\phi^{(s)}), \tag{4}$$

is defined, the expectation of a function $a(\phi)$ with respect to the distribution defined by $f(\phi)$ can be estimated by:

$$\bar{a}_f = \sum_{s=1}^n w^{(s)} a(\phi^{(s)}) / \sum_{s=1}^n w^{(s)}. \tag{5}$$

Consider now, the special case where the distribution defined by $f(\phi)$ is the posterior distribution for a likelihood function $l(x|\phi)$ and prior distribution $\pi(\phi)$, while the distribution defined by $g(\phi)$ is the prior distribution $\pi(\phi)$. Then, the importance weights defined in (4) become:

$$w^{(s)} = l(x|\phi^{(s)}). \tag{6}$$

Thus, we have a mechanism for making inference on a posterior distribution by sampling from the prior, and weighting the samples by the likelihood function. This forms the basis for the particle filter solution to our problem. For fuller details of particle filters the reader is referred to the book by Doucet et al [4].

3.4 Application of the Particle Filter

Initialisation of the filter requires a set of N_p equally weighted particles $((x_{0,1}^{(s)}, c_{1,1}^{(s)}), \dots, (x_{0,N_t}^{(s)}, c_{N_t,N_t}^{(s)}))$ from the joint posterior distribution for the classes and locations at time t_0 . The locations at time t_0 can be sampled according to the distribution for the targeting sensor measurement errors. Using definitions from Section 2.3 we set $x_{0,i} \sim N(l_i, \Sigma_i)$, for $i = 1, \dots, N_t$. The initial class samples for each detection are drawn probabilistically according to the class probability vectors ψ_i .

The Particle Filter algorithm works by passing samples through the Prior Evolver (defined in Section 3.2). Specifically, for $s = 1, \dots, N_p$ we:

- Sample $\{x_{1,1}^{(s)}, \dots, x_{1,N_t}^{(s)}\}$ from $p(x_{1,1}, \dots, x_{1,N_t} | x_{0,1}^{(s)}, \dots, x_{0,N_t}^{(s)}, c_{1,1}^{(s)}, \dots, c_{N_t,N_t}^{(s)})$, the Prior Evolver distribution (defined in Section 3.2).

- Evaluate the importance weights $w^{(s)} = p(D_1, \dots, D_{N_s} | x_{1,1}^{(s)}, \dots, x_{1,N_s}^{(s)}, c_1^{(s)}, \dots, c_{N_s}^{(s)})$ using the measurement likelihoods to be defined in Section 3.5.

The weighted samples $((x_{1,1}^{(s)}, c_1^{(s)}), \dots, (x_{1,N_s}^{(s)}, c_{N_s}^{(s)}))$ (with weights $w^{(s)}$) can then be used to approximate the required posterior distribution.

If the scenario were to be extended to a full tracking problem in which a time series of seeker images is obtained, this procedure would need to be altered to prevent degenerate weights (i.e. a few particles with very large weights, and the rest with small weights) [6].

3.5 Likelihood

Before the likelihood of the seeker measurements can be calculated, the images from the seeker and targeting sensors need to be registered. Various techniques for image registration could be used [3], such as those that extract and then match lines in the images [18]. In the example presented in this chapter the registration (and the uncertainty in this registration) is incorporated into the Prior Evolver.

Evaluation of the likelihood function is complicated by the need to associate the seeker measurements/detections with the targets. This association requires definition of the set of feasible association hypotheses θ . Each hypothesis associates a subset θ_d of seeker measurements with the detections from the targeting sensor. The cardinality of θ_d is denoted θ_{N_d} , and for each $i \in \theta_d$ we define λ_i to be the target (from the targeting sensor) to which that measurement is assigned. The remaining $\theta_{N_s} = N_s - \theta_{N_d}$ seeker measurements (defined by the indices $i \notin \theta_d$) are taken to correspond to additional targets or clutter measurements. Using the association hypotheses, the likelihood function can be expressed as:

$$p(D_1, \dots, D_{N_s} | x_{1,1}, \dots, x_{1,N_s}, c_1, \dots, c_{N_s}) = \sum_{\theta} \{ p(D_1, \dots, D_{N_s} | \theta, x_{1,1}, \dots, x_{1,N_s}, c_1, \dots, c_{N_s}) \times p(\theta | x_{1,1}, \dots, x_{1,N_s}, c_1, \dots, c_{N_s}) \} \quad (7)$$

Assuming independence between measurements/detections, the likelihood conditioned on the association hypothesis is given by:

$$p(D_1, \dots, D_{N_s} | \theta, x_{1,1}, \dots, x_{1,N_s}, c_1, \dots, c_{N_s}) = \left(\prod_{i \in \theta_d} p(y_i, z_i | x_{1,\lambda_i}, c_{\lambda_i}) \right) \times \left(\prod_{i \notin \theta_d} p(y_i, z_i) \right) \quad (8)$$

where using notation defined in Section 2.4:

$$\begin{aligned} p(y_i, z_i | x_{1,\lambda_i}, c_{\lambda_i}) &= N(y_i | x_{1,\lambda_i}, \Sigma_s) p(z_i | C_i = c_{\lambda_i}) \\ p(y_i, z_i) &= A_s^{-1} p(z_i) \end{aligned} \quad (9)$$

where A_s is the area of the surveyed region.

Following Gordon et al [7], the prior probabilities for the association hypotheses are expressed as:

$$p(\theta | x_{1,1}, \dots, x_{1,N_t}, c_1, \dots, c_{N_t}) = \text{Poisson}(\theta_{N_a} | \rho A_s) \times \text{Bin}(\theta_{N_d} | N_t, p_d) \times \left(\frac{N_t!}{(N_t - \theta_{N_d})! \theta_{N_d}!} \times \frac{N_s!}{\theta_{N_a}!} \right)^{-1} \quad (10)$$

The first term of equation (10) models the number of additional target or clutter detections in the processed seeker image by a Poisson distribution with mean ρA_s . Here, additional target refers to a target that was not detected by the targeting sensor, but which is never-the-less a proper target. A clutter detection corresponds to an object that does not belong to any of the J specific target classes. The second term of equation (10) is a Binomial distribution $\text{Bin}(\theta_{N_d}; N_t, p_d)$ for the number of detected targets θ_{N_d} . Specifically, each target from the targeting sensor is assumed to be detected in the seeker image with independent probability p_d (more complicated models would alter p_d according to the class of the target). Given θ_{N_d} , it is assumed that the allowable associations between targets and seeker measurements are equally likely. This produces the third term of equation (10). Allowable associations given θ_{N_d} are obtained by selecting the θ_{N_d} detected targets (the number of such subsets is $N_t! / ((N_t - \theta_{N_d})! \theta_{N_d}!)$), and then assigning these targets to the seeker measurements (the number of possible assignments for each subset is $N_s! / \theta_{N_a}!$).

In practical use, many of the association hypotheses will contribute only a negligible amount to the likelihood in (10), and can therefore be removed by a gating procedure based on thresholds for the class probabilities and location measurement errors. The gating procedure works by examining the associations between seeker measurements and Prior Evolver predictions. The procedure is best illustrated with an example. Suppose that it is proposed to associate a Prior Evolver prediction with class $C = j$ and location x with a seeker detection with location y and image chip z . A gating based on the class probabilities is obtained by comparing the posterior class probability based upon the seeker image chip:

$$p(C = j | z) = \frac{\pi_j p(z | C = j)}{\sum_{j'=1}^J \pi_{j'} p(z | C = j')} \quad (11)$$

with a pre-specified threshold. If the probability falls below the threshold then the association is rejected. A gating based upon the location can be obtained by using the seeker location measurement error distribution to set a threshold on the allowable distance between x and y .

3.6 Use of the particles

Numerous quantities of interest can be determined using the particles. These include such quantities as the class probabilities, mean target locations and the most likely association hypotheses for each of the seeker detections. The most likely association hypotheses would be relevant if a specific target is designated in the targeting data as being of interest. The most likely association hypothesis would then indicate which object detected by the seeker is most likely to correspond to the designated target. In this paper we concentrate on the target classes and locations.

The posterior probability that the i -th targeting detection is an object of class j is approximated by:

$$p(C_i = j | D_1, \dots, D_{N_s}, T_1, \dots, T_{N_t}) \propto \sum_{s=1}^{N_p} w^{(s)} I(c_i^{(s)} = j) , \quad (12)$$

where I is the indicator function (so $I(x = y) = 1$ if $x = y$ and 0 otherwise). The mean locations of the

targets at time t_1 can be approximated by:

$$\hat{\mu}_i = E(x_{1,i} | D_1, \dots, D_{N_t}, T_1, \dots, T_{N_t}) = \frac{1}{\sum_{s=1}^{N_p} w^{(s)}} \sum_{s=1}^{N_p} w^{(s)} x_{1,i}^{(s)} \quad (13)$$

for $i = 1, \dots, N_t$. Note, however, that estimation of the target locations by the mean of the location posterior distribution might not be appropriate. If the target locations have multi-modal distributions then the mean values might be away from the actual target locations. Ideally, we would examine the full distribution of possible target locations.

4.0 SYNTHETIC EXAMPLE

4.1 Description

The performance of the approach is illustrated with a synthetic example, in which there are three classes of target. Detections in a subset of the x-y plane, along with corresponding image chips were generated randomly to represent the targeting information after application of initial target detection algorithms. The number of targets detected was sampled from a Poisson distribution with restricted range:

$$N_t \sim \text{Poisson}(4) \times I(2 \leq N_t \leq 5) \quad (14)$$

The targeting sensor detections were restricted to the region $\{0.0 < x < 1.0, 0.0 < y < 1.0\}$. The covariance matrix of the Gaussian targeting sensor measurement error was given by:

$$\Sigma_t = \begin{pmatrix} 0.025^2 & 0.0 \\ 0.0 & 0.025^2 \end{pmatrix} \quad (15)$$

Thus the standard deviation of the targeting sensor measurement error was 0.025 along each axis. For the purposes of the documented experiments all targeting detections belonged to the set of three classes (i.e. there were no clutter objects in the targeting detections).

For demonstration purposes, the image chips were replaced by samples from 2-dimensional Gaussian distributions, whose parameters depended on the class of the target. These ID measurements could correspond to length and width, for example. Denoting the mean measurement vectors for class 1, 2 and 3 by μ_1 , μ_2 and μ_3 respectively, and the corresponding covariance matrices by Σ_1 , Σ_2 , Σ_3 we used:

$$\mu_1 = \begin{pmatrix} 1.0 \\ 1.0 \end{pmatrix}, \mu_2 = \begin{pmatrix} 2.0 \\ 0.9 \end{pmatrix}, \mu_3 = \begin{pmatrix} 1.1 \\ 2.0 \end{pmatrix}, \Sigma_1 = \begin{pmatrix} 0.43 & 0.07 \\ 0.07 & 0.43 \end{pmatrix}, \Sigma_2 = \begin{pmatrix} 0.26 & -0.02 \\ -0.02 & 0.30 \end{pmatrix}, \Sigma_3 = \begin{pmatrix} 0.27 & 0.05 \\ 0.05 & 0.41 \end{pmatrix} \quad (16)$$

The mean vector $\mu_{clutter}$ and covariance matrix $\Sigma_{clutter}$ for the clutter class were:

$$\mu_{clutter} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}, \Sigma_{clutter} = \begin{pmatrix} 4.0 & 0.0 \\ 0.0 & 4.0 \end{pmatrix} \quad (17)$$

These measurement distributions were selected to have considerable overlap between the classes. Thus, there will be a non-trivial error rate if classification is attempted using just a single image chip

measurement.

Target relocation during weapon fly-out was simulated via a global Gaussian shift of the targets, followed by independent local Gaussian perturbations. Specifically, a global shift of $N(\mu_{global}, \Sigma_{global})$ was applied equally to each target position, followed by an independent perturbation of $N(\mu_{local}, \Sigma_{local})$. The parameters were given by:

$$\mu_{global} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix}, \Sigma_{global} = \begin{pmatrix} 0.2^2 & 0.0 \\ 0.0 & 0.2^2 \end{pmatrix}, \mu_{local} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix}, \Sigma_{local} = \begin{pmatrix} 0.05^2 & 0.0 \\ 0.0 & 0.05^2 \end{pmatrix} \quad (18)$$

To simulate the seeker detections, the relocated targets have been detected with independent probabilities $p_d = 0.8$. Additionally, extra detections in line with a restricted Poisson distribution have been generated. Specifically, extra detections have been generated according to a Poisson distribution with mean $\rho = 2$ but subject to $N_t \leq N_s \leq 7$. The extra detections were restricted to the region $\{0.0 < x < 1.0, 0.0 < y < 1.0\}$, and were equally likely to be from any of the target classes and the background clutter class. The covariance matrix for the Gaussian seeker measurement error was set to:

$$\Sigma_s = \begin{pmatrix} 0.025^2 & 0.0 \\ 0.0 & 0.025^2 \end{pmatrix} \quad (19)$$

Thus the standard deviation of the seeker measurement error was 0.025 along each axis.

The image chips for the seeker detections were simulated by sampling from 2-dimensional class-conditional Gaussian measurement distributions. The same Gaussian measurement distributions as for the targeting sensor have been used. Although having the same measurement distributions for the targeting sensor and seeker may appear unrealistic, it does not bias the results, because the actual ID measurements are not compared within the Bayesian combination algorithm. Instead, only the likelihood values are combined, via Bayes' theorem.

The Prior Evolver was set to be a global Gaussian perturbation $N(\mu_{global}^{PE}, \Sigma_{global}^{PE})$, followed by independent local perturbations $N(\mu_{local}^{PE}, \Sigma_{local}^{PE})$. If the image chips are ignored, such a specification allows analytical calculations to be made, using Kalman Filters [7]. It is likely that alterations could be made so that the Kalman Filter could be used for the full problem. However to keep the approach generic (*i.e.* applicable to more complicated Prior Evolvers) the full particle filter was used. The parameters used were:

$$\mu_{global}^{PE} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix}, \Sigma_{global}^{PE} = \begin{pmatrix} 0.22^2 & 0.0 \\ 0.0 & 0.22^2 \end{pmatrix}, \mu_{local}^{PE} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix}, \Sigma_{local}^{PE} = \begin{pmatrix} 0.075^2 & 0.0 \\ 0.0 & 0.075^2 \end{pmatrix} \quad (20)$$

Note that the standard deviations of the Gaussian perturbations have been set to be slightly larger than those used to generate the synthetic data, to simulate uncertainty in our knowledge about target relocation. Furthermore, note that the relatively unrestricted target motion is actually making the problem harder, since the lack of constraints means that more particles are needed to ensure that all the possible target behaviour is accounted for.

Within the particle filter algorithm, the parameters of the location and sensor measurement distributions were all set to the same values used to generate the synthetic data. Similarly, the seeker detection probabilities $p_d = 0.8$ and the mean $\rho = 2$ of the Poisson distribution for extra target detections in the

seeker image (ignoring the restriction on the number of seeker detections) were all set to the same values used to generate the synthetic data. In real use this would not be possible, and these parameters would need to be estimated, or assigned using expert knowledge.

4.2 Monte Carlo assessment

A Monte Carlo assessment of performance has been conducted. The presented results are based on 200 random simulations, each using 5000 particles. On average there were 3.7 targets detected by the synthetic targeting sensor, and 5.0 objects detected by the seeker.

Two indicators of performance are presented. The first (Table 1) presents the average classification rate for the targets using the Bayesian combination procedure, with the classifications determined according to the maximum class probability calculated using (12). As baseline performance indicators the targeting sensor and seeker sensor classification rates are also presented (for the seeker sensor we treat missed detections as wrong classifications). Both the targeting and seeker sensor results assume that the target relocations are known (i.e. we correctly associate the detections with the actual targets). We can see that the Bayesian approach has been able to maintain the classification rate from the (idealised) targeting sensor, where-as the seeker sensor is penalised for the missed detections. Indeed, the Bayesian approach has actually been able to improve the classification performance over that of the targeting sensor only. This is presumably a result of combining the classification probabilities from the targeting sensor and the seeker.

Targeting sensor only	Seeker only	Bayesian combination
72.5%	59.0%	74.6%

Table 1: Classification rates

The second set of results (Table 2) gives an indication of the performance in determining target locations. In each case the quoted figure is the percentage of detections for which the actual target lay within a circle centred on the detection, with radius equal to three times the average sensor measurement error standard deviation. The detection locations used in the Bayesian combination results were the means of the posterior locations. The performance from the targeting sensor alone is very poor, due to the relocation of targets during weapon fly-out. The idealised seeker is penalised for its missed detections, but overall performs well since the true associations between measurements and targets have been used.

Targeting	Seeker	Bayesian	Bayesian with s.d.
8.2%	80.9%	52.8%	84.5%

Table 2: Performance estimating location

The Bayesian combination algorithm suffers from the fact that only the mean of the posterior distribution has been used, rather than the full posterior distribution. Thus, no account is being made of the estimates of the uncertainty in target location that are inherent within the posterior distribution. To show the effect of this, the percentage of detections for which the actual locations fell within a circle centred on the posterior mean, with radius equal to twice the estimated average posterior standard deviation of the location, is presented to the far right of Table 2. As can be seen by examining the two right-hand columns, taking this uncertainty into account produces much better performance. A factor of two has been used around the average standard deviation, rather than the factor of three used earlier, to penalise the potentially larger (compared to the sensor measurement errors) estimated standard deviations. Specifically, there is a trade-off between correctly estimating the possible variation in the target locations and having such a large spread in the posterior distribution locations that the particles fail to pin-point the targets

efficiently (which would negate the military utility of the algorithm).

5.0 SUMMARY AND FUTURE WORK

This paper has successfully developed a Bayesian procedure to enable exploitation of targeting information by a weapon seeker. The aim has been to use a targeting sensor to guide a seeker equipped weapon to an area containing high-value relocatable targets. The weapon seeker then needs to engage the high value targets, while minimising collateral damage. A Bayesian particle filter based solution has been developed. The procedure has been demonstrated successfully on a synthetic problem. Current work is applying the approach to more realistic scenarios. These scenarios include:

- Complicated target motion during weapon fly-out.
- Use of automated target detection algorithms applied to real data.
- Use of real data chips for the ID sensor measurements of targets.
- Use of more appropriate (and realistic) sensor measurement distributions.

Although, on the face of it, such extensions would be expected to make the problem harder, this is not necessarily the case. For example, real targets might actually be more separable than the overlapping multivariate Gaussian distributions used in the synthetic example. Furthermore, taking into account more sophisticated target motion should improve performance, since it has the effect of constraining the possible target relocations.

6.0 REFERENCES

- [1] M.S. Arulampalam, N. Gordon, M. Orton, and B. Ristic. A variable structure multiple model particle filter for GMTI tracking. *Proceedings of 5th International Conference on Information Fusion*, 927–934, 2002.
- [2] Y. Bar-Shalom and X.R. Li. *Multitarget-Multisensor Tracking*. YBS Publishing, 1995.
- [3] L.G. Brown. A survey of image registration techniques. *ACM Computing Survey*, 24(4):325–376, 1992.
- [4] A. Doucet, J.F.G. de Freitas, and N.J. Gordon. *Sequential Monte Carlo Methods in Practice*. New-York:Springer-Verlag, January 2001.
- [5] J. Geweke. Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57(6):1317–1339, November 1989.
- [6] N. Gordon, S. Maskell, and T. Kirubarajan. Efficient particle filters for joint tracking and classification. *Signal and Data Processing of Small Targets 2002*, ed O.E. Drummond, SPIE Vol 4728, 2002.
- [7] N. Gordon, D. Salmond, and D. Fisher. Bayesian target tracking after group pattern distortion. *Signal and Data Processing of Small Targets 1997*, ed O.E. Drummond, SPIE Vol 3163:238–248, 1997.
- [8] N.J. Gordon and D. Salmond. Bayesian pattern matching technique for target acquisition. *Journal of Guidance, Control, and Dynamics*, 22(1):68–76, January-February 1999.
- [9] 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, 1993.
- [10] T. Kirubarajan, Y. Bar-Shalom, K.R. Pattipati, and I. Kadar. Ground target tracking with variable structure IMM estimator. *IEEE Transactions on Aerospace and Electronic Systems*, 36(1):26–46, January 2000.
- [11] P.M. Lee. *Bayesian Statistics, An Introduction*. Arnold, London, 2nd edition, 1997.
- [12] P.O. Nougues and D.E. Brown. We know where you are going: tracking objects in terrain. *IMA Journal of Mathematics Applied in Business and Industry*, 8:39–58, 1997.
- [13] D. Salmond and N.J. Gordon. Group and extended object tracking. *Proc. of SPIE Conference on Signal and Data Processing of Small Targets 1999, Denver, Colorado*, 3809:284–296, July 1999.
- [14] A.F.M. Smith and A.E. Gelfand. Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician*, 46(2):84–88, May 1992.
- [15] C. Sodtke and J. Llinas. Terrain based tracking using position sensors. *Proceedings of 4th International Conference on Information Fusion*, 2, 2001.
- [16] A.R. Webb. *Statistical Pattern Recognition*. John Wiley & Sons, Chichester, 2nd edition. 2002.
- [17] K.D. Copley, R.O. Lane, S. Manchanda and A.R. Webb, Bayesian approach to recognising relocatable targets. *NATO RTO SET Symposium SET-080, Target Identification and Recognition using RF Systems*, Oslo, Norway, October, 2004.
- [18] M.L. Williams, R.C. Wilson, and E.R. Hancock. Multiple graph matching with Bayesian inference. *Proceedings of the International Conference - Pattern Recognition in Practice V*, 1997.

7.0 ACKNOWLEDGMENT

This research was sponsored by the UK MOD Corporate Research Programme. The authors thank Dr David Salmond, QinetiQ Farnborough, for helpful early discussions.