

Bayesian Approach to Recognising Relocatable Targets

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 presents a framework for target acquisition. The targets of interest are relocatable ground vehicles imaged at time $t=t_0$ by a long range targeting sensor and then at a later time $t=t_1$ by a weapon platform. The framework must handle several key issues: changes in scene (vehicle movement between t_0 and t_1); incorporation of domain knowledge (terrain and vehicle type); image registration errors; differences in viewing angle; uncertainty in vehicle type, and location. A modular approach is presented in which the key quantities of interest are probability density functions. There are many technical issues that must be addressed and two in particular are highlighted: the development of generalisation procedures between sensors that enable training data gathered with one sensor to be used to classify data obtained from a different RF sensor (specifically, a procedure to enable ISAR data to be exploited); and the development of techniques that use prior knowledge from a targeting sensor to aid a weapon seeker (the use of targeting information to support acquisition). A Bayesian methodology is adopted and the research is set in the target acquisition context.

1.0 INTRODUCTION

1.1 Target Acquisition

This paper is concerned with the classification of relocatable ground vehicles in weapon seeker data. The overall aim is to develop techniques for target detection and acquisition using data from sensors that differ from the targeting sensors. A Bayesian framework is presented that allows target-specific information acquired by additional (different) sensors, together with domain knowledge of terrain and target properties, to be exploited within the automatic target recognition (ATR) framework for the weapon seeker sensor. There are two sources of additional sensor information that influence the weapon seeker classifier.

- 1) Targeting data. This is illustrated in Figure 1. A scene of interest is observed at time t_0 by a long range targeting sensor (e.g. SAR). At a later time, t_1 , the (evolved) scene is imaged by a weapon seeker. The targeting sensor provides prior knowledge as to the nature and location of target data that may be exploited in the weapon seeker algorithm.
- 2) Classifier training data. One of the main aims is to enable objects imaged by a weapon's seeker to be classified using ATR systems trained on more readily available ground-based sensor data. The exemplar application used here is to use ATR systems trained on readily available Inverse

Paper presented at the RTO SET Symposium on "Target Identification and Recognition Using RF Systems", held in Oslo, Norway, 11-13 October 2004, and published in RTO-MP-SET-080.

Bayesian Approach to Recognising Relocatable Targets

Synthetic Aperture Radar (ISAR) data to classify objects imaged by a Doppler Beam Sharpened (DBS) radar seeker. This is a non-trivial problem since key differences between the measurements from different platforms arise from differences in sensor technology, spatial resolution, polarisation, frequency, imaging geometry and target motion.

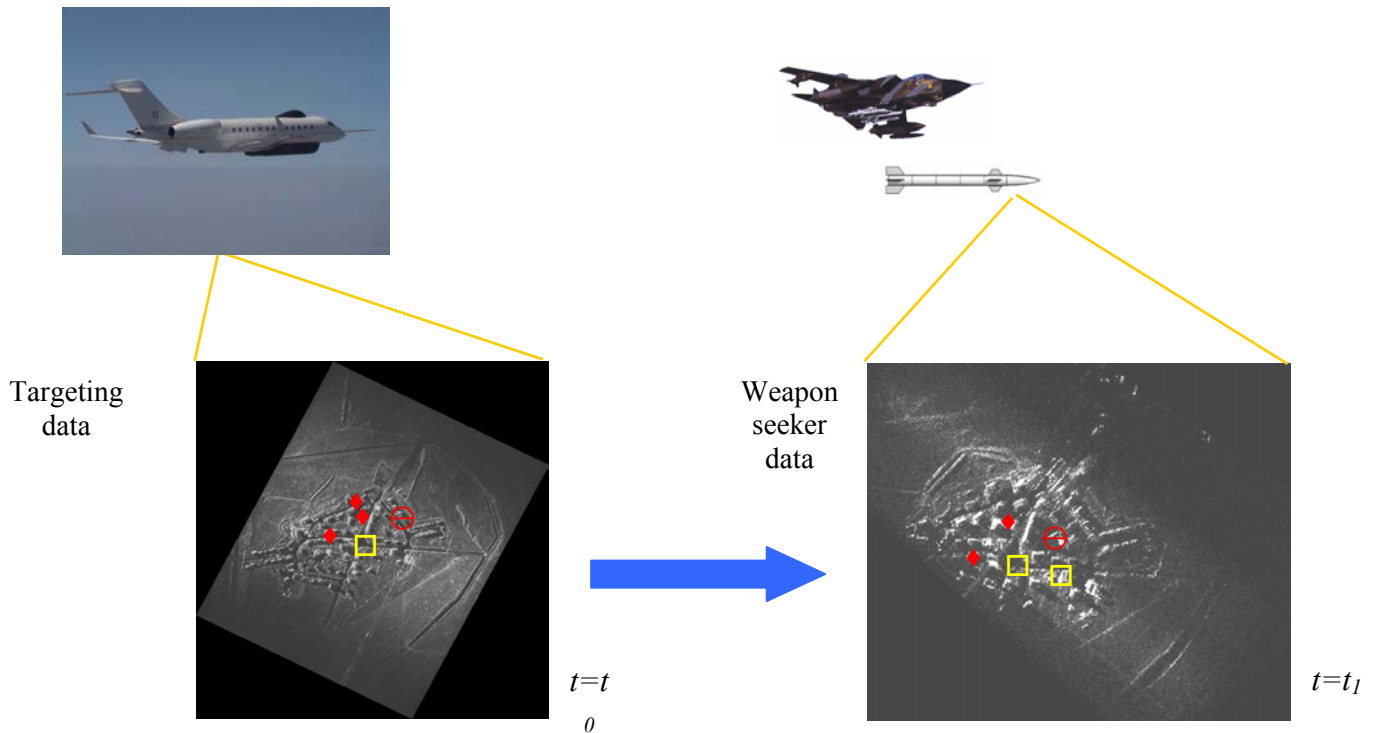


Figure 1: Exploitation of targeting information by a weapon seeker

Among the many issues that must be addressed are:

- 1) Differing imaging geometry between the targeting sensor and the seeker.
- 2) Change in the configuration of targets during weapon fly-out (i.e. staleness of the targeting information).
- 3) Deployment of countermeasures by the targets subsequent to imaging by the targeting sensor.
- 4) The difficulty (and expense) of obtaining sufficient training data for a weapon seeker ATR system.
- 5) Uncertainty in target positions and type.
- 6) Image registration errors.

A particularly adverse effect of items two and three 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.

1.2 Aim and Outline of paper

The aim of this paper is to describe a framework for target acquisition that addresses the above issues. Two specific technical issues are highlighted:

- 1) The development of generalisation procedures between sensors that enable training data gathered with one sensor to be used to classify data obtained from a different RF sensor.
- 2) The development of techniques that use prior knowledge from a targeting sensor to aid a weapon seeker.

A Bayesian methodology is adopted. The main motivation behind a Bayesian approach is the ability of Bayesian statistics to handle limited and possibly conflicting pieces of information in a fully consistent manner. Further generic arguments in favour 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 of interest.

Section 2.0 describes the target acquisition framework, with Section 2.1 summarising the technical issues that such a framework must handle and Section 2.2 presenting a framework. Section 3.0 describes the approaches to the two problems above where data from additional sensors is used in the weapon seeker classifier. Finally, we conclude with a summary of the approach.

2.0 A FRAMEWORK FOR TARGET ACQUISITION

2.1 Technical Issues

The following subsections summarise some of the main technical issues that a target acquisition framework must address.

2.1.1 Uncertainty in Vehicle Type

There is uncertainty in vehicle type both within the targeting data and the weapon seeker data. This may be reduced by using contextual information or prior knowledge, but the decision making process must be able to 'fuse' both sources of information concerning target type.

2.1.2 Vehicle Movement

The target acquisition procedures must be designed to cope with possible changes in target configuration during weapon fly-out, and possible distortion of signatures due to the deployment of countermeasures. Information that may be exploited includes terrain information, vehicle properties and intelligence information.

2.1.3 Different Sensors

In many applications of pattern classification, including target recognition, the operating conditions for the classifier differ from those used to gather data for training the classifier [6]. Thus, the training conditions are not representative of the expected operating conditions. These differences can be due to a number of factors. The specific aspect that this programme has addressed is the difference of sensors between training and operating conditions. A Bayesian inverse imaging procedure has been developed which allows the seeker data to be classified using ATR systems trained on more readily available (and cheaper) data from a second sensor (in the exemplar application, an ISAR processor). The sensors differ in range and cross-range resolution. By utilising larger amounts of training data covering more varied extended

operating conditions, this procedure is likely to lead to an improved autonomous classification ability for the weapon seeker, that should enable the seeker to identify and react to changes in the configuration of targets during weapon fly-out.

If the ATR system for the sensor providing the training data has some degree of robustness to countermeasures, the proposed seeker ATR system will inherit this robustness, provided that the effects of the countermeasures are similar for the two sensors. However, the caveat has to be added that design of an ATR system that is robust to countermeasures is a current research area in itself.

Proposed approaches to target classifier design currently under investigation include Gaussian mixture models [3], [5], non-linear dimensionality reduction techniques [9], unsupervised "symmetry-preserving" neural network techniques [11], [12], and unsupervised encoder networks.

2.1.4 Platform Motion

Uncertainties in platform motion lead to additional differences between training and operating conditions. The inverse imaging procedure above must take account of this through simultaneous auto-focus/super-resolution (see Section 3.2.2).

2.1.5 Different Imaging Geometries

The imaging geometries between the targeting sensor and the weapon seeker will differ. The scene will be observed from different elevations and directions of view. Thus, radar shadow will be different between the two measured scenes leading to a possible source of error in the image registration process.

2.1.6 Registration

In order to fuse the targeting predictions (estimates of target location and type) with the weapon seeker predictions, it is necessary to have a model for the image registration errors. These errors will depend on such factors as the position of the targets in the field of view, the imaging geometries, sensor resolution and target movement.

2.1.7 Vehicle Classes

One of the problems with ATR is the definition of the classes. Vehicles of the same basic type are used for different military purposes and therefore the importance of classifying a vehicle correctly depends on its role. It is not the decision, but the expected utility of the decision that matters classifier design. Therefore, costs of misclassification should be taken into account in the decision making process. The usual criterion of error rate as a means of assessing a classifier is deficient in that it treats all misclassifications equally.

Conversely, vehicles of different type (or the same type with different equipment fits) are used for the same military function leading to radar returns that can vary significantly within the class. Designing classifiers that are robust to intra-class variability can be an important problem to address.

2.1.8 Contextual information

In addition to the information on target type that may be gained from radar measurements of the target, further clues as to target type may result from contextual information and domain knowledge. The type of contextual information that could be incorporated could include the proximity of other potential targets or the type of terrain in which the vehicle is operating. The technical issue to address here is the specification of the domain knowledge and the description of a framework that handles such knowledge, together with target measurements, in a consistent manner.

2.2 Framework

The overall framework proposed by this programme is outlined in Figure 2. A pragmatic Bayesian approach has been proposed in which the various aspects of the problem are treated in a modular fashion, with the outcomes of each module being descriptions of probability distributions that can be combined under a Bayesian probabilistic formalism. An advantage of the modular approach is that existing techniques (such as those developed for registering images from different sensors) can be used where appropriate. In particular, the green boxes in Figure 2 relate to processes where we have relied on current state-of-the-art techniques with minimal customisation. Since many of these processes require further research to provide satisfactory solutions, the formalism is designed in such a manner that new techniques can be incorporated as they reach maturity.

We now consider the main modules in Figure 2 in turn.

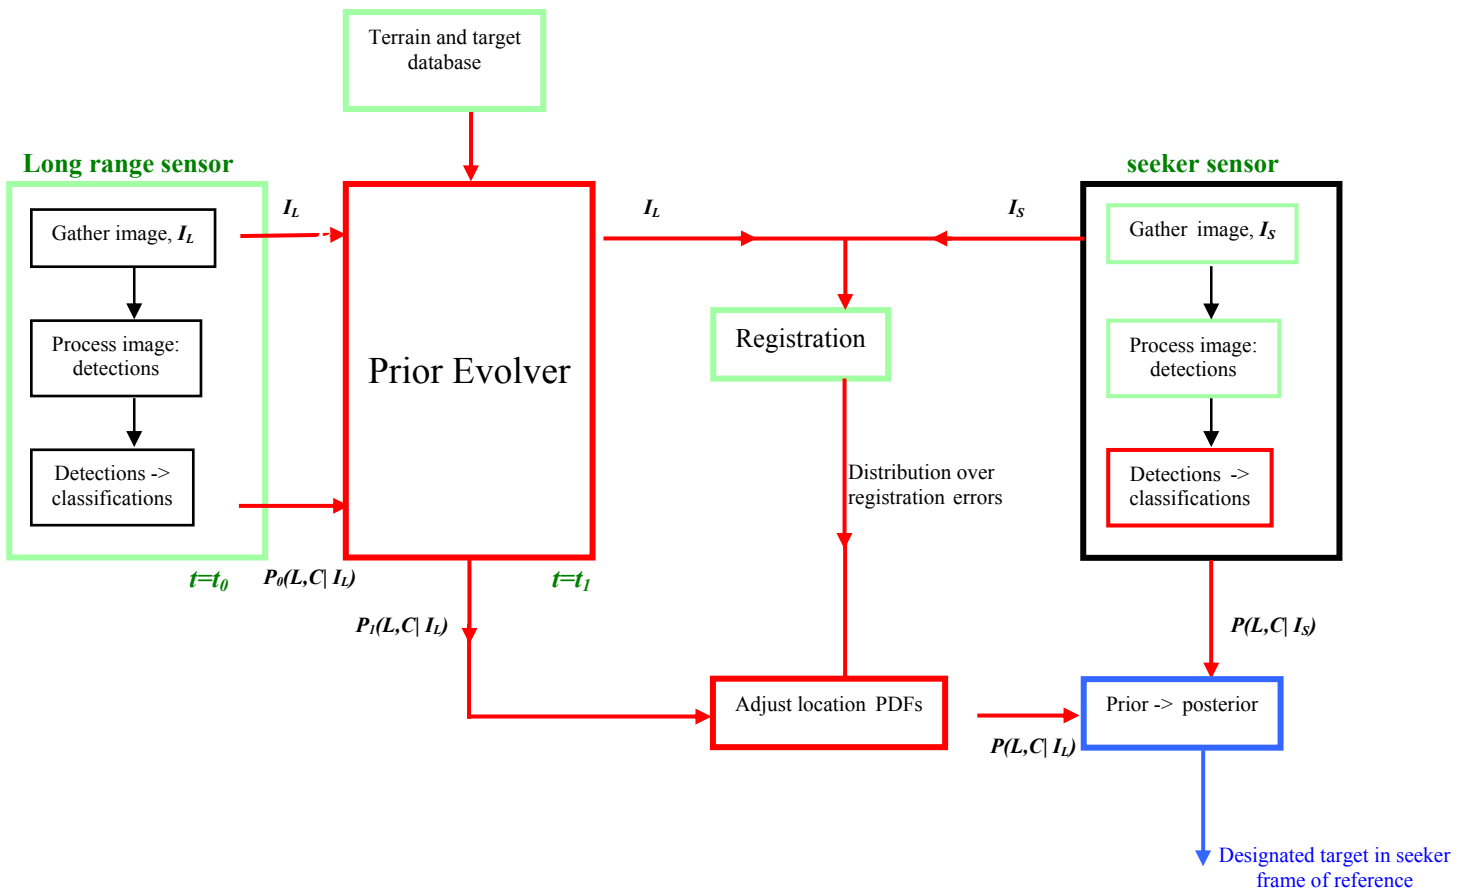


Figure 2: Framework for the exploitation of targeting information by a weapon seeker

2.2.1 Long-range Sensor

The long-range sensor, depicted as the green box on the left-hand side of Figure 2, captures image data, I_L , at time $t=t_0$ and processes the data to produce target detections and identifications. This processing may be performed manually, or by some semi-automatic process. The output of this process is an estimate of the

number of targets present, with their locations and types. This is represented by the probability density function, $p_0(L, C | I_L)$ where L represents the vector of target locations and C represents the vector of target classes. This captures the uncertainty in the quantities L and C at time $t=t_0$.

2.2.2 Prior Evolver

The prior evolver is used to update the information from the targeting sensor to allow for target motion during weapon fly-out. This consists of predicting how the detections gleaned from the targeting sensor at $t=t_0$ will have changed by the time $t=t_1$ that the weapon seeker views the targeted area. This evolution of the probability density function can be modelled to include bulk motion of targets (perhaps reflecting the motion of a convoy) and more complicated behaviour incorporating knowledge of the terrain and likely target behaviour. The result of this process is the probability density function, $p_1(L, C | I_L)$, which is the original distribution evolved to time $t=t_1$. Contextual information and domain-specific knowledge can be incorporated within this density [4].

2.2.3 Seeker Sensor

The seeker sensor box performs a similar operation to the long-range sensor. Image data, I_S , is gathered using the seeker sensor, this time at time $t=t_1$, and processed to produce $p(L, C | I_S)$, the distribution over locations and classes of target given the seeker data. Note that in this case, the dimensionality of L and C may differ from the long-range sensor case since the seeker detection process may estimate a different number of targets to be present. Also, the locations are measured in the seeker frame of reference.

The on-board classifier that produces the classifications is designed using training data. An approach that uses ISAR data to train a classifier that is applied to DBS images is outlined in Section 3.2.2.

2.2.4 Registration

Registration of the targeting sensor image, I_L , to the frame of reference of the seeker produces a translation, with associate errors, expressed as a probability density function.

2.2.5 Adjust Location PDFs

The registration information enables the updated targeting information (expressed in the form of probability density $p_1(L, C | I_L)$) to be adjusted to the seeker frame of reference. This is denoted in the figure by the probability density function, $p(L, C | I_L)$.

2.2.6 Prior -> Posterior

The incorporation of uncertain targeting information, expressed in terms of the probability density function $p_1(L, C | I_L)$, with uncertain seeker detections, described by $p_1(L, C | I_S)$, occurs in the (blue) box labelled "Prior -> posterior", and is outlined in Section 3.2. The output of this procedure is the updated information (in terms of target class probabilities) for the objects within the seeker frame of reference. This information will improve the ability of the weapon to engage the targets designated on launch of the weapon while minimising collateral damage.

3.0 INCORPORATING ADDITIONAL SENSOR INFORMATION

3.1 Introduction

The previous section described a framework for target acquisition. A modular approach was proposed in which the key quantity of interest was a probability density function. Manipulation of these pdfs by each

module resulted in a description of the designated target of interest in the seeker frame of reference. There are two main areas where information from different sensors has to be combined or fused. One is where the prior targeting information from the long-range sensor is combined with information from the seeker (the ‘prior->posterior’ box in Figure 2). This is reported in detail in [1]. The second area is where training data from an ISAR sensor is used to train a classifier applied to DBS imagery (the ‘Detections->classifications’ box in the seeker sensor). This is reported in [2] and [8]. The approaches are summarised below.

3.2 Exploiting Targeting Information

3.2.1 Targeting detections

The number of targeting sensor detections at time t_0 is denoted by N_t . The estimated 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 could be estimated using a standard ATR system or possibly via human intervention.

The measurement errors for the target locations are assigned Gaussian distributions, so that $l \sim N(x, \Sigma_t)$ where x is the actual target location, and Σ_t 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.

3.2.2 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 $i = 1, \dots, N_s$.

It is assumed that density models (conditional on class) can be estimated for the ID measurements. Estimating these distributions given only limited training data for the weapon seeker is the subject of current research. 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. Ideally, each density estimate should incorporate the uncertainty in the centre of each detected object. Mixture model densities meet many of the requirements for these class-conditional densities.

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.

3.2.3 Bayesian solution

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. The posterior distribution of interest at time t_1 is:

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

the distribution of the locations and types of the targets given seeker and targeting data. In [1], a Bayesian technique based on particle filtering [7] is used to obtain samples from the posterior distribution.

3.3 Generalising Classifiers

We now turn to the problem of designing a classifier for the seeker (DBS) data. We denote measurements in the training (ISAR) conditions by the variable x and measurements in the operating (DBS) conditions by the variable z . We suppose that we have a training set $D = \{x_i, i = 1, \dots, N\}$ of samples gathered under the training conditions. This training set is used to design a Bayesian classifier (*i.e.* a classifier based on probability distributions for the sensor measurements) [10], which outputs posterior class probabilities for each ISAR measurement of an object to be classified. The posterior class probabilities estimated by the classifier for an ISAR measurement x are denoted by $p(C = j | x, D)$, for $j = 1, \dots, J$ where J is the number of target classes.

In accordance with our operational scenario, during operational use we only have access to a DBS measurement z (rather than the ISAR measurement x) so cannot use the Bayesian classifier directly. To proceed we require a model $p(x | z)$ for the relationship between an operational sensor measurement and a training sensor measurement. Then, we can consider the expectation of the posterior class probabilities given the operational sensor measurement:

$$E[p(C = j | x, D) | z] = \int p(C = j | x, D) p(x | z) dx \tag{3}$$

In this manner, given a model for the conditional density $p(x | z)$ we can use the training sensor classifier to classify an operational sensor measurement.




Figure 3: Imaging framework

We assume a radar sensor measurement process as shown in Figure 3.

For the specified sensor measurement processes, the conditional density $p(x | z)$ can be expressed as:

$$p(x | z) = \int d\sigma p(x, \sigma | z) = \int d\sigma p(x | \sigma) p(\sigma | z) \tag{4}$$

The term $p(x | \sigma)$ represents the forward sensor measurement process for the training sensor, *i.e.* generation of a training sensor image from an underlying cross section σ (temporarily suppressing the defocus parameters within our notation). The term $p(\sigma | z)$ is the restored cross section given the operational sensor measurement. Determination of $p(\sigma | z)$ corresponds to a super-resolution problem.

Thus, Equation (4) states that to find the distribution of ISAR measurements that correspond to a DBS image, z , we find the distribution of underlying cross sections, σ , that give rise to z , and then pass these through an ISAR imaging model. The resulting conditional distribution, $p(x | z)$, is then substituted into (3) to classify z . Since (3) cannot be evaluated analytically for the majority of sensor models, samples are drawn from the relevant posterior distributions and inference is based on those samples. In [8], a Bayesian approach to simultaneous auto-focus/super-resolution is described. In [2] results of this classification procedure on synthetic data are presented.

4.0 SUMMARY

This paper has described a framework for target acquisition that has the following features

- 1) It presents a modular approach in which the key quantities of interest passed between modules are probability density functions;
- 2) It handles uncertainty in target locations and classes;
- 3) Prior targeting information is combined with seeker data in a consistent manner;
- 4) A classifier trained on ISAR data may be used to provides estimates of target type for DBS data using a Bayesian auto-focus/super-resolution approach;

- 5) Knowledge of the properties of vehicle and terrain type may be included in a 'prior evolver';
- 6) Registration errors are treated probabilistically.

The approaches to two particular aspects of the target detection problem have been outlined, namely

- 1) the combination of prior targeting data with seeker detections to produce posterior estimates of target locations and classes;
- 2) the use of a previously trained classifier on data recorded from a different sensor.

5.0 REFERENCES

- [1] Copsey, K.D., Lane, R.O., Manchanda, S. and Webb, A.R. (2004). Bayesian Approach to exploiting prior targeting information within a weapon seeker. NATO RTO SET Symposium SET-080. *Target Identification and Recognition Using RF Systems*. Oslo, Norway.
- [2] Copsey, K.D., Lane, R.O. and Webb, A.R. (2004). Designing NCTR algorithms when operating sensor conditions differ from training conditions. *Radar 2004*. International Conference on Radar Systems, Toulouse, France.
- [3] Copsey, K.D. and Webb, A.R. (2000). Bayesian approach to mixture models for discrimination. *Advances in Pattern Recognition*, F.J. Ferri, J.M. Inesta, A. Amin and P. Pudil (eds). Springer Lecture Notes in Computer Science 1876, 491-500.
- [4] Copsey, K.D. and Webb, A.R. (2002). Bayesian networks for incorporation of contextual information in target recognition systems. In T. Caelli, A. Amin, R.P.W. Duin, M. Kamel and D. de Ridder (editors) *Structural, Syntactic and Statistical Pattern Recognition*, Proceedings of the Joint IAPR International Workshops SSPR2002 and SPR2002, Windsor, Canada, August 2002. Lecture Notes in Computer Science 1876, 709-717, Springer.
- [5] Copsey, K.D. and Webb, A.R. (2003). Bayesian gamma mixture model approach to radar target recognition *IEEE Transactions on Aerospace and Electronic Systems*. 39(4), 1201- 1217.
- [6] Copsey, K.D. and Webb, A.R. (2004). Classifier design for population and sensor drift. *Structural, Syntactic and Statistical Pattern Recognition*, Proceedings of the Joint IAPR International Workshops SSPR2004 and SPR2004, Lisbon, Portugal, August 2004. Lecture Notes in Computer Science, Springer.
- [7] Doucet, A., de Freitas, J.F.G. and Gordon, N.J. (2001). *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag.
- [8] Lane, R.O., Copsey, K.D. and Webb, A.R. (2004). A Bayesian approach to simultaneous autofocus and super-resolution. Proc. SPIE 5427. *Algorithms for Synthetic Aperture Radar Imagery XI*. E.G. Zelnio and F.D. Garber (eds). Orlando, FL, USA.
- [9] Maskall, G.T. and Webb, A.R. (2002). Nonlinear feature extraction for MMW image classification: a supervised approach. Proc. SPIE 4726, 353-363. *Automatic Target Recognition XII*. Firooz A. Sadjadi (ed).
- [10] Webb, A.R. (2002). *Statistical Pattern Recognition*. Second edition. John Wiley and Sons, Chichester.

- [11] Webber, C.J. (2000). Self-organisation of symmetry networks: transformation invariance from the symmetry-breaking mechanism. *Neural Computation*, 12(3), 565-596.
- [12] Webber, C.J. (2001). Predictions of the spontaneous symmetry-breaking theory for visual code completeness and spatial scaling in single-cell learning rules. *Neural Computation*, 13(5), 1023-1043.

6.0 ACKNOWLEDGEMENT

This research was sponsored by the UK MOD Corporate Research Programme.

