Assessment of a Bayesian approach to recognising relocatable targets

Richard O. Lane, Keith D. Copsey, and Andrew R. Webb
QinetiQ, Malvern Technology Centre, St. Andrews Road, Malvern,
Worcestershire, WR14 3PS, United Kingdom.
rlane1@QinetiQ.com, kkcopsey@QinetiQ.com, arwebb@QinetiQ.com

ABSTRACT

This paper considers an automatic target recognition concept in which a long-range targeting sensor is used to aid a radar seeker-equipped weapon operating in an area containing high-value relocatable targets. The weapon seeker is designed to engage the high-value targets, while minimising collateral damage. Previous work proposed a Bayesian approach that enables the weapon seeker to exploit the targeting information before making its final decision. 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, and had previously only been applied to a synthetic example. This paper summarises the approach, and presents results from an assessment using scenarios derived from an airborne data set containing short-range Doppler beam sharpened imagery. The issue of differing resolutions for the two sensors is addressed by super-resolution techniques, which are also assessed.

1. INTRODUCTION

1.1. General

This paper concerns a Bayesian framework for using target-specific information acquired by one sensor within the automatic target recognition (ATR) system for a second sensor on another platform. The particular scenario chosen for development is to enable a Doppler beam sharpened (DBS) radar seeker-equipped autonomous weapon to exploit long-range synthetic-aperture radar (SAR) targeting information. The generic problem addressed is one where two images from the same type of sensor, separated in time, are available to classify relocatable targets in a scene.

The aim of the work is to enable a pre-targeted seeker-equipped weapon to engage relocatable targets, while minimising collateral damage. The approach is required to cope with uncertainty in the targeting identifications, change in the target layout configuration during weapon fly-out, differing imaging geometry and image resolution between the targeting sensor and the seeker, and deployment of countermeasures by the targets. A particularly adverse effect of the last three points is that a target designated correctly by the targeting sensor may have both a different location and an altered signature by the time the weapon seeker 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. Overall framework

The work documented in this paper fits within the Bayesian framework detailed in [1]. The aim is to take outputs from a long range targeting sensor and a weapon seeker, and use them to estimate the posterior distribution for the target locations and classes at the time that the weapon seeker images the scene.

The main motivation behind a Bayesian approach [2, 3] 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 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 estimated. The issue of confidence is of importance in military operations because a degree of certainty is necessary before decisions can be made. The main disadvantage of Bayesian techniques

NATO RTO SET-096 specialists' meeting on the millimeterwave advanced target recognition and identification experiment (MATRIX 2005), Oberammergau, Germany, 10-12 May 2005.

is that the model used in calculations is assumed to be correct, where in practice it is not known. This may be mitigated by ensuring the model covers all likely possibilities.

The proposed Bayesian approach is a pragmatic one in which the various aspects of the problem are treated in a modular fashion, with the outcomes of each module being combined under a Bayesian probabilistic formalism. An advantage of the modular approach is that existing non-Bayesian techniques such as super-resolution, for example, can be used where appropriate. The disadvantage is that modules do not iteratively interact and make full use of all the information. It is unlikely, however, that this would make much difference in the overall performance.

1.3. Outline of paper

Two strands of research are assessed in this paper. The first of these covers a specific Bayesian approach using 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. An assessment based on millimetre wave (MMW) data collected using an airborne platform is covered in section 2.

The second strand of research covers the change in target signature when imaged by two different sensors. The particular problem studied here is where the resolution of the two sensors is different. In this case super-resolution can be used to aid the signature matching process used in the first strand of research. Initial results using MMW inverse SAR (ISAR) turntable data are presented in section 3. Overall conclusions are given in section 4.

2. POSITION MATCHING IN A RURAL SCENE

2.1. Problem Specification

In the scenario considered here, an image from a targeting sensor is obtained at time t_0. Detection algorithms are then applied to obtain a set of target locations. An image chip is obtained for each detected target by centring an input window of sufficient size to cover expected targets at the location of each detection. Target recognition procedures are then applied to estimate the class of the object leading to each detection. These procedures could take the form of ATR algorithms [5], or could be decisions made by a human operator, taking into account context in addition to the image chip. If a target class corresponds to a high-value target then a seeker equipped weapon is launched to engage this nominated target.

The weapon seeker reaches and images the highlighted area at a later time, t_1. Automated target detection algorithms are then applied to obtain a set of seeker detections, along with associated image chips centred 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.

To update the information from the targeting sensor and allow for target motion during weapon fly-out a *prior evolver* attempts to predict 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. In the tracking literature [6] 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, or an image registration error) can be included easily, using a global Gaussian translation [7]. More complicated motion, incorporating knowledge of the terrain and likely target behaviour may be modeled – targets can be predicted to follow a road network [8, 9], for example.

Precise details of the algorithm assessed here are given in [10] with a basic three-target example using synthetically generated locations and two feature measurements for each target, which could correspond to length and width, for example. The next section gives an assessment of performance using real data, a larger number of targets and more features.

2.2. Real data example

2.2.1. Introduction

This section outlines on-going work assessing the proposed Bayesian approach using scenarios derived from MMW DBS imagery. The data were collected in trials held during 2002. A short dwell time was used to mimic final target acquisition by a weapon-borne seeker. Nine military ground vehicles are present within the data, including armoured personnel carriers, air defence units and main battle tanks. These targets are embedded in a rural environment containing gravel roads, hedges and trees (both isolated and in a forest). The type and configuration of each of the nine targets is such that they can be considered to belong to different target classes. Thus, we have a nine class problem ($J = 9$).

The available data in its original form is insufficient to meet our experimental requirements completely. Although long-range spotlight SAR imagery of the targets present at Copehill Plantation is available, the target locations are the same as within the DBS imagery. Therefore, although a full experiment for utilising SAR targeting information to assist a DBS seeker could be conducted, the experiment would not cover the issue of staleness of targeting information arising as a result of target relocations during weapon fly-out. Rather than proceed with such an assessment, the conducted experiments are based on a number of scenarios simulating target relocations. Currently, these experiments use the DBS sensor data only. The main consequences of this are that the targeting and seeker images are perfectly registered, and the resolution of both test and training databases is the same.

2.2.2. Feature measurements

The amount of available DBS imagery of the military vehicles is too small to generate enough feature measurements for our experimental requirements, so the real data has been used to specify synthetic feature measurements. For both the targeting sensor and the seeker, these synthetic feature measurements have been used in place of real target image chips.

The QinetiQ SAR Machine-Aided Recognition Toolbox (SMARTBox) has been used to extract sets of features from a limited amount of DBS training data for each target class. The feature-extraction process involved automatically extracting the target from the DBS imagery, and computing a set of 13 features, including: mean and variance of the power of the pixels within the extracted target; area and perimeter of the extracted target; length and width of the extracted target; auto-correlation; shadow height and width.

These features are the same as those used by the classifiers within the SMARTBox, and are therefore assumed representative of the discriminating information present within the DBS image chips. A separate 13-dimensional Gaussian feature distribution $N(\upsilon_j, \Sigma_j)$ has been estimated for each class $j = 1, \ldots, J$ using SMARTBox features extracted from real DBS target image chips. A clutter class $N(\upsilon_{\text{clutter}}, \Sigma_{\text{clutter}})$ has also been specified, using features extracted from DBS image chips of trees.

2.2.3. Experiment description

We now describe scenarios used to assess the Bayesian framework. The composite DBS image that is being used as the targeting sensor image is displayed in Figure 1a. The image contains nine targets, whose locations are (550,650), (650,650), (350,660), (550,750) (470,670), (500,750), (390,720), (400,600) and (650,620). The pixel spacing in these images is 1m in both the x and y directions. The units for spatial co-ordinates in the remainder of this section are in metres.

It is assumed that all nine targets in the targeting sensor image are detected manually prior to weapon launch ($N_t = 9$). To reflect the errors inherent in such a procedure, the two-dimensional coordinates $l_1, \ldots, l_{N_t}$ of each targeting sensor detection are drawn from multivariate Gaussian distributions $l_i \sim N(x_{0,i}, \Sigma_t^{\text{true}})$, where $i = 1, \ldots, N_t$, $x_{0,i}, \ldots, x_{0,N_t}$ are the ground-truth locations at time t_0, and:

$$\Sigma_t^{\text{true}} = \begin{pmatrix} 5.0^2 & 0.0 \\ 0.0 & 5.0^2 \end{pmatrix} \tag{1}$$

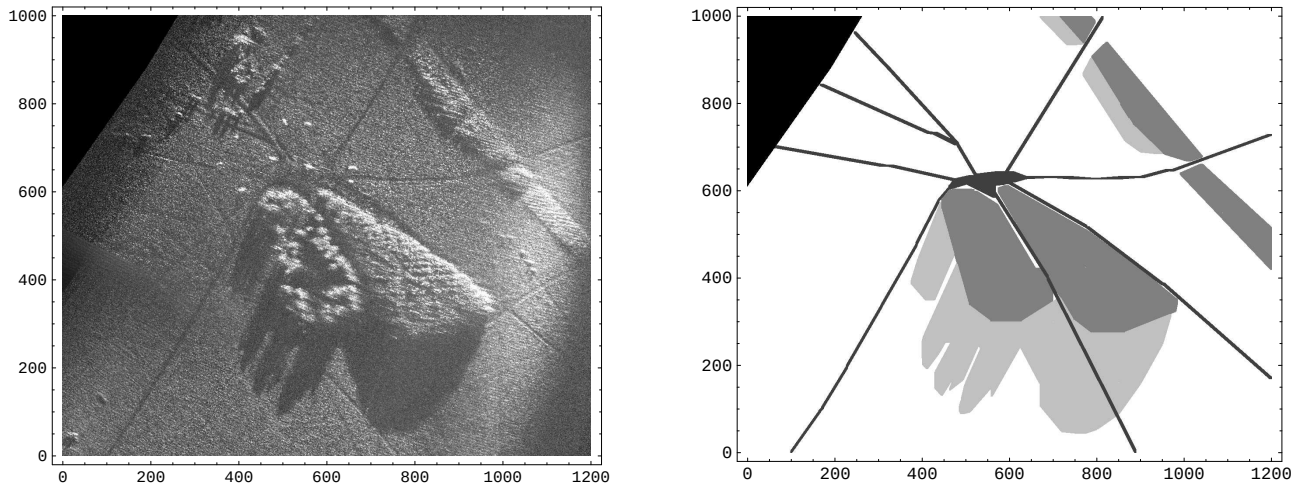

(a) Targeting sensor image

(b) Segmented image

Figure 1: Image of scene. (a) Original image. (b) Segmentation of the targeting and seeker scene by ground type (black – roads/undetermined, dark grey – forest, light grey – shadow, and white – grass).

Associated with each detected target is a 13-dimensional feature vector, simulated from the feature measurement distribution determined by the target class. Specifically, for the i-th target, of target class c_i, we sample the feature vector r_i from the distribution $N(\upsilon_{c_i}, \Sigma_{c_i})$.

Target relocation during weapon fly-out has been simulated for this dataset. To date, seven scenarios have been considered. In each case the same high-value target at (550,650) in Figure 1a has been moved to a different location, following the road network. The new target locations for each scenario are (635,700), (680,765), (650,630), (860,630), (310,780), (480,620) and (350,440). This encompasses a variety of situations including relatively small movements, a very close spacing of three targets and a wide separation between the high-value target and the remaining targets.

In addition to the pre-specified relocations, for each scenario, all the targets are subjected to an additional independent perturbation of $N(\mu_{\text{local}}, \Sigma_{\text{local}})$. The parameters are given by:

$$\mu_{\text{local}} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix} \qquad \Sigma_{\text{local}} = \begin{pmatrix} 15.0^2 & 0.0 \\ 0.0 & 15.0^2 \end{pmatrix} \tag{2}$$

To simulate the seeker detections, the high-value relocated target is detected with probability 0.99, and the remaining relocated targets are detected with independent probabilities $p_d^{\text{true}} = 0.90$. A 13-dimensional feature vector is simulated for each seeker detection, by sampling from the appropriate target class feature measurement distribution $N(\upsilon_j, \Sigma_j)$. The feature measurements for the targeting sensor and seeker are therefore drawn from the measurement distributions independently.

In addition to the seeker detections of the relocated targets, a further nine detections corresponding to clutter are generated. Five of these clutter detections are placed randomly within the forested areas of the DBS image, while four are placed randomly within the grassy areas of the image. Such a distribution of false alarms is considered typical when current target detection algorithms are applied to DBS imagery of this size in rural environments. A 13-dimensional feature vector from the clutter measurement distribution $N(\upsilon_{\text{clutter}}, \Sigma_{\text{clutter}})$ is simulated for each clutter false alarm.

To simulate errors in detecting the centre of the targets, the two-dimensional coordinates $y_1, \ldots, y_{N_s}$ of each seeker detection are drawn from multivariate Gaussian distributions with mean equal to the ground-truth

location (target or clutter) and covariance matrix set to:

$$\Sigma_s^{\text{true}} = \begin{pmatrix} 10.0^2 & 0.0 \\ 0.0 & 10.0^2 \end{pmatrix} \tag{3}$$

The variances are larger than for the targeting sensor (see (1)), to reflect the typically lower resolution of seekers.

2.3. Implementation details

Having introduced the scenarios, we now describe the implementation details for the particle filter-based Bayesian combination scheme.

One of the most important parts of the scheme is the *prior evolver*, which predicts 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. If complicated target relocations are taking place, it is vital that the prior evolver incorporates behaviour that can match those relocations. For this example, we assume that expert knowledge and intelligence information has indicated that the specific high-value target will attempt to relocate, while the other targets are likely to remain close to their original locations.

The prior evolver predictions are determined with the aid of a segmentation of the scene into roads, forest, shadow and grass. The segmented image is displayed in Figure 1b. For this example, the segmentation has been performed manually. However, techniques for automated segmentation and terrain-classification of SAR imagery exist and could be used within a full system [11].

Using this segmentation, the prior evolver predictions for the relocation of the high-value target are generated using the following procedure:

(1) With a certain probability, specify the terrain of the yet to be specified predicted location according to:

$$\text{terrain}^{\text{PE}} = \begin{cases} \text{Road} & p(\text{Road}) = 0.81 \\ \text{Forest} & p(\text{Forest}) = 0.09 \\ \text{Shadow} & p(\text{Shadow}) = 0.04 \\ \text{Grass} & p(\text{Grass}) = 0.06 \end{cases} \tag{4}$$

(2) Sample the high-value target location change $d \sim N(\mu_{\text{global}}^{\text{PE}}, \Sigma_{\text{global}}^{\text{PE}})$.

(3) Set $x' = x + d$, where x is the location of the high-value target as fed into the prior evolver (*i.e.* the location from the targeting sensor).

(4) Looking only within the terrain of type $\text{terrain}^{\text{PE}}$, find the closest location to x' in terms of Euclidean distance, and set the prediction to be that closest location.

The parameters used here are:

$$\mu_{\text{global}}^{\text{PE}} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix} \qquad \Sigma_{\text{global}}^{\text{PE}} = \begin{pmatrix} 150.0^2 & 0.0 \\ 0.0 & 150.0^2 \end{pmatrix} \tag{5}$$

These values should change with the length of the time gap $t_1 - t_0$ between collection of the targeting sensor measurements and the seeker measurements. The parameter values in (4) encode a prior belief that the target is most likely to relocate along the road network. However, the precise values have been selected arbitrarily. It is planned to develop a more sophisticated prior evolver in future work, which will predict realistically the motion of targets along the road network during the time gap $t_1 - t_0$. The current implementation merely predicts that the relocated target will, in the main, end-up on the road (*i.e.* a route is not traced out).

In addition to the prior evolver prediction for the relocation of the high-value target, independent local Gaussian perturbations $N(\mu_{\text{local}}, \Sigma_{\text{local}})$ are applied to each location. The parameters used are:

$$\mu_{\text{local}}^{\text{PE}} = \begin{pmatrix} 0.0 \\ 0.0 \end{pmatrix} \qquad \Sigma_{\text{local}}^{\text{PE}} = \begin{pmatrix} 5.0^2 & 0.0 \\ 0.0 & 5.0^2 \end{pmatrix} \tag{6}$$

The class probability vectors ψ_i for each targeting sensor detection, have been obtained using a Gaussian classifier applied to the feature vectors $r_1, \ldots, r_{N_t}$ of the targeting detections. The parameters of the Gaussian classifier are estimated using simulated training data generated from the feature measurement distributions $N(\upsilon_j, \Sigma_j)$, for target classes $j = 1, \ldots, J$. Specifically, for each target class j in the targeting sensor image, one-hundred training data examples from the class are used to estimate a mean vector $\upsilon_{t,j}$ and covariance matrix $\Sigma_{t,j}$.

In a similar manner, the seeker feature measurement likelihood distributions for the target classes are set according to $N(\upsilon_{s,j}, \Sigma_{s,j})$. The parameters $(\upsilon_{s,j}, \Sigma_{s,j})$ are estimated using training data generated from the feature measurement distributions $N(\upsilon_j, \Sigma_j)$. To reflect the difficulty of obtaining suitable training data for a seeker, only fifty training data examples from each class are used compared to one-hundred for the targeting sensor. Thus although both targeting and seeker feature distributions are derived from the same underlying distribution the estimates of the mean and covariance are different. The distribution for the clutter class was set using the same method.

The covariance matrices for targeting and seeker sensor location measurements used within the algorithm are set to:

$$\Sigma_t = \begin{pmatrix} 7.5^2 & 0.0 \\ 0.0 & 7.5^2 \end{pmatrix} \qquad \Sigma_s = \begin{pmatrix} 12.5^2 & 0.0 \\ 0.0 & 12.5^2 \end{pmatrix} \tag{7}$$

Note that the parameters of the location measurement errors Σ_t and Σ_s are slightly different to the ground-truth values specified in (1) and (3), to reflect real situations in which the detections are determined using automated target detection algorithms, rather than via a specification of a measurement error model.

2.4. Results

The procedures described above have been run five times for each of the seven scenarios (each time using different realisations for the feature measurements and ground-truth local perturbations). 10000 particles have been used for each run of the particle filter-based Bayesian combination scheme. Two performance metrics are used – one for correct classification rates and one for target location errors.

Table 1 displays the average classification rates for the targets. Two rates are provided for the Bayesian combination procedure. The first, and most appropriate, is based on classifications to the maximum class probability using all particles, and is referred to as the averaged Bayesian result. The second is based on classifying to the classes specified by the particle with the largest weight – the maximum *a posteriori* (MAP) particle. Classification rates for the targeting sensor alone and the seeker alone are presented also, with missed detections for the seeker being counted as wrong classifications. All the results are based on correct associations between the detections and the actual targets.

Table 1: Probability of correct classification.

Targeting sensor alone	Seeker sensor alone	Bayesian (average)	Bayesian (MAP)
88%	76%	90%	90%

Table 2: Average location error.

	Targeting alone	Seeker alone	Bayesian (average)	Bayesian (MAP)
Relocated	195.3	21.4 (0%)	43.5	39.2
Remaining	22.2	22.6 (11%)	21.6	20.0

The seeker classification rates are lower than those from the targeting sensor and the Bayesian combination scheme due to a combination of the missed detections, and the lower amount of training data made available to estimate the likelihoods for the seeker feature measurements. Both the averaged and MAP Bayesian combination scheme provide better classification performance than obtained from the targeting sensor alone, reflecting the fact that two sets of features measurements (targeting and seeker) are used by the Bayesian scheme.

Performance measures for estimating target locations are displayed in Table 2. The performance measures are based on a location error measure. For each simulation, the true assignment between estimated locations and ground-truth locations is used to specify the error measure. The table separately documents the average measure (over all scenarios and all runs) for the relocated high-value target and the remaining 8 targets. Since the measures are based on Euclidean distances, the lower the measure, the better the performance. Again, two figures for the Bayesian combination scheme are provided: the average over all particles and the MAP particle alone. The figure for the seeker is based on the subset of targets detected by the seeker, with percentage of targets missed displayed also in brackets. Note that this artificially biases the figures in favour of the seeker, since the targets that are missed by the seeker are likely to be the targets with the poorest performance in the Bayesian combination scheme.

As would be expected, reliance on the targeting sensor alone results in poor performance for the relocated target. The performance measures for the Bayesian combination scheme indicate that the scheme is able to handle the non-trivial target relocations this example provides. However, the seeker alone does provide better performance for the relocated target. Therefore, to improve performance the Bayesian scheme could be used to estimate which seeker detections are associated with which targeting detections. This would provide a mechanism to use the better classification rates for the Bayesian scheme in conjunction with the accuracy of the locations provided by the seeker processing.

For the targets which have only a local perturbation, the Bayesian scheme is able to combine the information from the two sensors to improve the accuracy of the target estimates. The advantage over the seeker is due to the lower accuracy of seeker measurements compared to targeting sensor measurements, combined with the fact that if a target relocation is only small, the original targeting sensor detection will still be good, even before the prior evolver is applied.

2.5. Conclusions

The Bayesian scheme has been demonstrated successfully and, based on classification performance, is shown to be superior to using either the targeting sensor measurements alone or the seeker measurements alone. The real data simulations reported here extend the synthetic problem previously addressed [10] by dealing with larger numbers of targets, more complicated target relocations during weapon fly-out, and feature measurements derived from real data.

Providing an overall quantitative measure for the performance of the Bayesian scheme has proved difficult, because the error measures for assessing target classification performance are different to those used for assessing the performance in locating the targets. Currently, the two aspects are assessed separately. Moreover, assessing the performance in locating targets is awkward, because a number of factors need to be taken into account. A solution that indicates uncertainty in the exact target locations is generally preferable to one that does not indicate the uncertainty. However, a solution that always provides vague predictions is less useful than one that accurately pin-points the targets. A further issue arises from how to penalise missed detections when assessing the performance for the seeker alone. Ideally, the proposed approach should be evaluated according to a measure of effectiveness based on the military utility of the algorithm outputs. This measure should take into account the final decision to be made by the weapon, of whether to attack a specific location, or abort in mid-air.

3. SUPER-RESOLUTION ALGORITHM COMPARISON

3.1. Problem Specification

As part of the framework assessed in the previous section it is necessary to match features of target images obtained using two different sensors. In the example above, training data was made available for both sensors. In many situations the training data available for one sensor, typically the seeker, will be limited. This does not pose a problem if the underlying distribution of the features is the same for both sensors because a common distribution can be used. However, if the resolution of the sensors is different then the feature distributions estimated using these images will also be different. One way of circumventing this problem is to use super-resolution on the target images to achieve a resolution independent of the sensor used to collect the data.

Each type of target is assumed to have a high-resolution two-dimensional plan view back-scattering coefficient. When imaged from a particular geometry this gives rise to a complex image. The image is formed by the convolution of a point spread function (PSF) and the high-resolution 2D representation of the target. This effectively places an amplitude-scaled copy of the PSF at each scatterer location. A bulk phase is present in the complex image, which is due to the distance between the radar and the target. In addition to this the various scatterers on the target will have a phase relative to each other dependent on their precise position and construction. The imaging process is described mathematically by

$$g = Tf + n, \tag{8}$$

where f is a complex vector denoting the raster-scanned high-resolution 2D target representation, T is an appropriately formatted matrix that applies the effect of the PSF, and g is a complex vector representing the resultant low-resolution image. Thermal noise in the radar receiver is modeled by n, which is a zero-mean circularly complex Gaussian with a diagonal covariance matrix N. The signal-to-noise ratio (SNR) is defined as

$$SNR_{in} = \frac{||Tf||^2}{||n||^2}. \tag{9}$$

Determination of f given g is a super-resolution problem. An overview of various algorithms for tackling this problem is now given.

3.2. Algorithm Descriptions

3.2.1. Matrix Inverse

The benchmark solution against which other super-resolution algorithms are compared is the matrix inverse. The estimate for f is given by

$$\hat{f}_{INV} = (T^H T)^{-1} T^H g. \tag{10}$$

Performance using this procedure is usually poor due to the ill-conditioning of the problem.

3.2.2. Singular Value Decomposition

In singular value decomposition (SVD) the PSF matrix T is decomposed into three matrices, $T = USV^H$, where U and V are orthogonal matrices and S is a diagonal matrix whose elements comprise the singular values. Any singular values below a threshold are set to zero and the inverse of the new matrix is calculated as $T_s^{-1} = VS^{-1}U^H$ with the inverse of the zero values also set to zero [12]. The level at which the threshold is set depends on the specific application but during initial testing of the algorithm it was found that a threshold of $s_{max}/\sqrt{SNR}$ produced good results. Thus the SVD solution is given by

$$\hat{f}_{SVD} = T_s^{-1} g. \tag{11}$$

If all the singular values are used then this is equivalent to the matrix inverse.

3.2.3. Minimum Mean Square Error

If the covariance matrix of the *a priori* distribution of f is W then the solution that minimises the mean square error is given by [13–15]

$$\hat{f}_{MMSE} = WT^H(TWT^H + N)^{-1} g. \tag{12}$$

However, in practice W is not known to a sufficiently useful accuracy so an iterative scheme that updates the W matrix is used [14]. In this scheme W is set to the identity matrix for the first iteration and in subsequent iterations is estimated from $\hat{f}$. Any element of $\hat{f}$ whose power is below a threshold is assumed to be clutter and the combined variance of all these elements is calculated and entered into W at the appropriate point. Any element i above the threshold has its variance set to $|\hat{f}_i|^2$. The procedure is terminated when $\hat{f}$ changes by less than a specified amount between iterations. An alternative algorithm is also used here where all diagonal elements of W are set to $|\hat{f}_i|^2$. The two algorithms are referred to minimum mean square error threshold (MMSE-T) and minimum mean square error all (MMSE-A), respectively.

3.2.4. MUSIC

The estimated intensity of f is given by

$$\widehat{|f|^2}_{MUSIC} = (T^H T)^{-1} T (gg^H - \lambda_{min} I) T (T^H T)^{-1} g, \tag{13}$$

where λ_{min} is the minimum eigenvalue of gg^H and I is the identity matrix. If $\lambda_{min} = 0$ then this is equivalent to the matrix inverse. In the complete MUSIC algorithm [16] there is an extra step that estimates the number of scatterers and their precise positions. Here however, in common with the algorithms above, the scene is assumed to contain a set number of scatterers that exist on a discrete grid of points.

3.2.5. Bayesian Super-resolution

The full Bayesian solution differs in one fundamental aspect from the other algorithms discussed here. Rather than estimating a single value for f the probability density for many values of f is estimated. This is given by

$$p(f, \theta|g) = \frac{p(g|f, \theta) p(f) p(\theta)}{p(g)}, \tag{14}$$

where

$$p(g|f, \theta) = \frac{1}{\det(\pi N)} \exp\left[-(g - T(\theta)f)^H N^{-1} (g - T(\theta)f)\right]. \tag{15}$$

Solution of (14) uses a Markov chain Monte Carlo (MCMC) algorithm, which produces a series of samples that characterise the probability density function. The specific algorithm used here is the Metropolis-Hastings algorithm [17]. After initialisation, at each step of the algorithm a new sample is generated according to a proposal distribution; if it is more likely (including the effect of the proposal distribution) than the current one it is always accepted but less likely ones are also accepted with a certain probability, which avoids getting trapped in local maxima. If the proposed sample is rejected then the current sample is used in the next iteration. Initial samples generated during the burn-in period depend on the starting position and must be discarded if they are not reasonable values of the distribution. The remaining samples are distributed from $p(f, \theta|g)$ as required.

This algorithm has previously been applied to synthetic one-dimensional signals to show its utility in an ATR context [18]. Results using a 2D synthetic non-coherent data example are shown in [19]. Full details of the algorithm, including a more complicated model than the one used here are given in [20].

3.3. Results

To test the super-resolution algorithms, MMW ISAR imagery has been used as the high-resolution representation of targets, f. An example of one of these images is shown in the top left of Figure 2. The image is of a T-72, with range along the horizontal axis and cross-range along the vertical axis. A point spread function representative of what may result due to simple movement of the radar platform, after the application of an autofocus algorithm, has been used. The PSF is a sinc function in the range direction and a modified sinc function [21] in the cross-range direction. The magnitude of the PSF is depicted in the top right of Figure 2. Application of the PSF to the ISAR imagery results in a lower resolution image g, shown in the bottom left of Figure 2. Noise is then added at a specified SNR.

A basic demonstration that super-resolution recovers the high-resolution target is shown in the bottom right of Figure 2, where the MMSE-T algorithm was used. The overall shape of the target and the position of the strongest scatterers have been correctly recovered. However, the magnitude of the scatterers are not always the same as the original target image. A measure of this difference – the output SNR – is defined as

$$SNR_{out} = \frac{||f||^2}{||\hat{f} - f||^2}. \tag{16}$$

A Monte Carlo assessment has been carried out to compare the various super-resolution algorithms. Each algorithm was applied using the same input image with SNR_{in} varying from 0 dB to 50 dB. Ten Monte Carlo

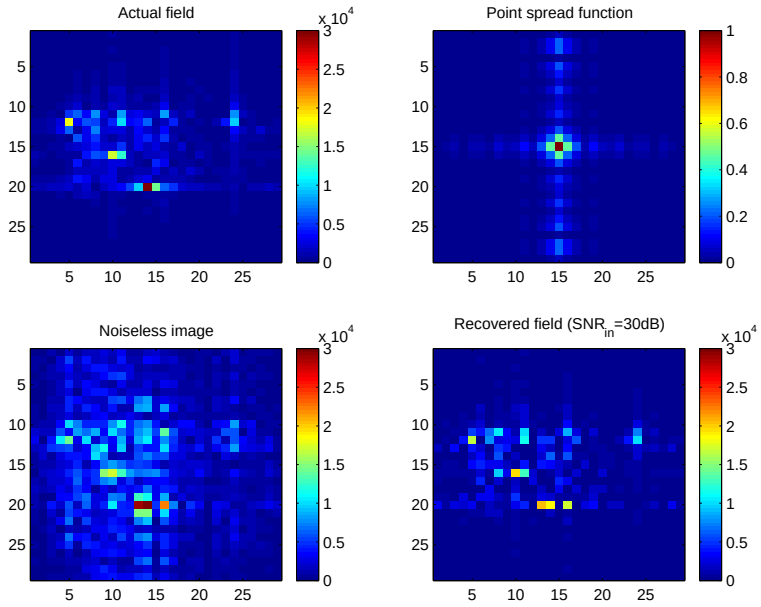

Figure 2: Super-resolution using MMW ISAR imagery

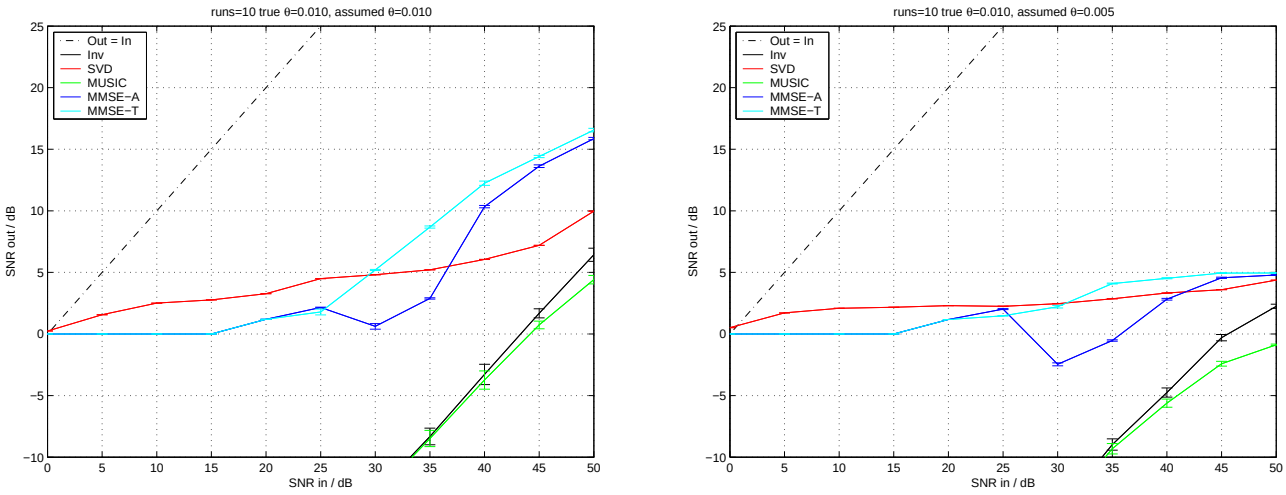

(a) Correct PSF

(b) Modified PSF

Figure 3: Super-resolution performance as a function of SNR

runs for each value of SNR_{in} were used and SNR_{out} was calculated for each run. The results of this are shown in Figure 3a, where solid lines show the mean and error bars show the standard error.

It is immediately evident from Figure 3a that the algorithms provide a trade-off between resolution and SNR. Although the resolution has improved, this is at the expense of a loss in SNR. At high input SNRs the performance of the algorithms begins to converge. However, at low SNRs the SVD and MMSE algorithms perform best. In general, the output SNR seems low due to an inadequacy of the metric as a performance measure. If a strong scatterer is present and its amplitude is correctly estimated but its position is incorrect by one pixel then it appears to be a strong noise spike, which degrades the SNR. However, by visually examining the recovered field solution in Figure 2, the major features of the target are still present when this is the case.

Therefore a feature-based ATR system should still have a reasonable performance even with a low apparent SNR.

The results displayed in Figure 3a are based on the situation where there is perfect knowledge of the PSF. In practice this is unlikely to be the case because there will be some residual error in the autofocus procedure. The effect of this has been modeled by reducing the level of defocus the algorithms assume by half, while actually applying the original level of defocus. The results of this are shown in Figure 3b. There is little change in performance at low SNRs because the noise dominates any error in the PSF. However, as the SNR increases the effect of an incorrect PSF is apparent – it puts a limit on output performance whatever the input SNR. If a parametric model of the PSF distortion is available then this performance loss could be mitigated by estimating parameters of the model [20, 21].

Results of the Bayesian algorithm are not shown in Figure 3. This is because the output of the algorithm is a distribution of estimates for f, which should include an estimate close to the true value, giving a very high output SNR. However, an advantage of the Bayesian algorithm is that it also includes other valid, whilst less likely, estimates which provide information on uncertainty. Therefore, it is difficult to provide a single performance metric for Bayesian super-resolution. The best super-resolution metric for this ATR application would be the correct classification rate of the whole ATR system, which would depend on knowledge about uncertainty in intermediate stages. The best way to apply this is a subject of ongoing research.

3.4. Conclusions

A variety of super-resolution algorithms have been demonstrated successfully using two-dimensional ISAR imagery. All algorithms improve resolution at the expense of a loss in SNR and are sensitive to accurate knowledge of the point spread function. However, the metric based on output SNR used to assess performance will not necessarily capture the effects of feature extraction when used in an ATR system and does not apply to Bayesian super-resolution. Further investigation should analyse the effect of super-resolution on target feature values. Also, since it has been shown previously that super-resolution can aid ATR [18, 22], these algorithms should be incorporated into an ATR system and assessed based on classification performance.

4. OVERALL CONCLUSIONS

A Bayesian procedure to enable exploitation of targeting information by a weapon seeker has been developed. 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 solution based on the use of a particle filter (importance sampling) has been assessed using scenarios derived from real data. Key to the approach is the specification of a prior evolver to make predictions on how the targets will relocate during weapon fly-out. These predictions should incorporate knowledge of the terrain and likely target behaviour including travelling along a road network, motion towards assumed desired target locations, such as potential hide locations, and movement away from impenetrable terrain, such as rivers. It is important that the prior evolver covers the actual target relocations.

The position matching element of this work that estimates both location and class probabilities is shown to work well in a typical rural scene where a single target has relocated and others have not moved. The classification performance of the system is better than when a single sensor is used. Future work should assess performance using other target behaviour such as convoys or a group movement towards a certain area such as a tree line.

As part of the framework assessed here it is necessary to calculate target class likelihoods. These are calculated by comparing features of a target image to a database. If the resolution of the imagery is different for different sensors then super-resolution may be used to aid the match. A comparison of super-resolution algorithms using real data has been made and shown that performance should be sufficient to aid the ATR process. Further work is required to incorporate this strand of research into the full Bayesian framework.

Acknowledgments

This work was sponsored by the UK Ministry of Defence Corporate Research Programme.

REFERENCES

1. K. D. Copsey, 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.
2. J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*, John Wiley & Sons, 1994.
3. P. M. Lee, *Bayesian Statistics, An Introduction*, Arnold, London, 2nd ed., 1997.
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. A. R. Webb, *Statistical Pattern Recognition*, John Wiley & Sons, Chichester, 2nd ed., August 2002.
6. Y. Bar-Shalom and X. R. Li, *Multitarget-Multisensor Tracking*, YBS Publishing, 1995.
7. N. Gordon, D. Salmond, and D. Fisher, "Bayesian target tracking after group pattern distortion," *Proc. of SPIE vol. 3163, Signal and Data Processing of Small Targets 1997, ed O.E. Drummond* , pp. 238–248, 1997.
8. 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* , pp. 927–934, 2002.
9. 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**, pp. 26–46, January 2000.
10. K. D. Copsey, R. O. Lane, S. Manchanda, and A. R. Webb, "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* , October 2004.
11. D. Blacknell, N. S. Arini, and I. McConnell, "High-level SAR image interpretation using contextual information," *Journal of Defence Science* **8**, May 2003.
12. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, Cambridge University Press, 2nd ed., 1992.
13. D. Blacknell and S. Quegan, "SAR super-resolution using a perturbed point spread function," *IGARSS* , pp. 2592–2595, July 1989.
14. S. P. Luttrell and C. J. Oliver, "Prior knowledge in synthetic-aperture radar processing," *J. Phys. D: Appl. Phys.* **19**, pp. 333–356, 1986.
15. L. M. Delves, G. C. Pryde, and S. P. Luttrell, "A super-resolution algorithm for SAR images," *Inverse Problems* **4**(3), pp. 681–703, 1988.
16. R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation* **34**, pp. 276–280, March 1986.
17. S. Chib and E. Greenberg, "Understanding the Metropolis-Hastings algorithm," *The American Statistician* **49**(4), pp. 327–335, 1995.
18. K. D. Copsey, R. O. Lane, and A. R. Webb, "Designing NCTR algorithms when operating sensor conditions differ from training conditions," *International conference on radar systems (Radar2004), Toulouse, France*, October 2004.
19. R. O. Lane, "Estimating radar cross section using Bayesian image restoration," *Proceedings of the London Communications Symposium* , pp. 1–4, September 2003.
20. R. O. Lane, K. D. Copsey, and A. R. Webb, "A Bayesian approach to simultaneous autofocus and super-resolution," *Proc. of SPIE vol. 5427, Algorithms for synthetic aperture radar imagery XI* , pp. 133–142, April 2004.
21. S. P. Luttrell, "A Bayesian derivation of an iterative autofocus/superresolution algorithm," *Inverse Problems* **6**(6), pp. 975–996, 1990.
22. L. M. Novak, G. J. Owirka, and A. L. Weaver, "Automatic target recognition using enhanced resolution sar data," *IEEE Transactions on Aerospace and Electronic Systems* **35**, pp. 157–175, January 1999.