

ATR of Ground Targets: Fundamentals and Key Challenges

Professor David Blacknell

Dstl
Porton Down, Salisbury
Wiltshire SP4 0JQ
UNITED KINGDOM

dblacknell@dstl.gov.uk

Dr. Luc Vignaud

ONERA
Chemin de la Hunière
F-91761 Palaiseau Cedex
FRANCE

Luc.Vignaud@onera.fr

ABSTRACT

Since the advent of radar as sensor for military remote sensing, the potential benefits in terms of reduced operator workload that automated cueing could bring have been clear. In particular, as radar technology progressed to provide an imaging capability, the potential to go beyond simple detection and provide a level of target recognition has become apparent. This lecture provides an introduction to the fundamentals of ground target recognition using radar. In particular, automatic target recognition (ATR) based on ground target images provided by synthetic aperture radar (SAR) is considered.

1.0 INTRODUCTION

1.1 Overview

Since the advent of radar as sensor for military remote sensing, the potential benefits in terms of reduced operator workload that automated cueing could bring have been clear. In particular, as radar technology progressed to provide an imaging capability, the potential to go beyond simple detection and provide a level of target recognition has become apparent.




Figure 1.1: The origins of ground imaging radar.

One of the earliest ground imaging radars was the H2S system developed by the Telecommunications Research Establishment (TRE) in Malvern, UK during the Second World War. One appealing possible explanation for the name of this system was that the Government scientific advisor, Lord Cherwell, repeatedly declared that “it stinks” (a typical British expression of disapproval) when told of delays to the programme which had resulted from a misunderstanding between him and the developers. As a result, the developers gave the project the codename H2S, i.e. the chemical symbol for hydrogen sulphide, which of course “stinks” with a rotten egg smell. Figure 1.1 shows a typical image from the H2S with a map for comparison. It is fair to say that the resolution of this system is quite coarse by today’s standards, but recognition of landmasses is clearly possible.




Figure 1.2: Image of tank “graveyard” (courtesy of QinetiQ).

In modern radar imaging systems, much better resolutions are now achievable. Figure 1.2 shows an example of a synthetic aperture radar (SAR) image at 15cm resolution (courtesy of QinetiQ). The highlighted area is a tank “graveyard” where disused tanks have been abandoned. SAR images are not intuitively interpretable by humans who are used to seeing optical images but it is possible to see the periodic bright returns along the side of the uppermost vehicle which result from the wheels along the side of the vehicle. This illustrates a number of points. It is apparent that some level of vehicle recognition should be possible at such resolutions but the target characteristics which are most evident in radar imagery may not correspond to the most recognisable optical characteristics. It is therefore important to understand the characteristic “features” of the image. It should also be noted that, because radar imagery is not intuitively interpretable, the role of automated algorithms is even more important to act as an aid for operators.

Inspired by the level of information available in high resolution SAR imagery, much research has been undertaken over the past 20 years into automatic recognition of ground targets in SAR. The aim of this lecture is to provide an overview of the fundamentals of ground target recognition. The lecture begins with a discussion of the phenomenology of SAR images, i.e. what are the distinctive characteristics of objects when imaged using this radar technique. It should be emphasised that radar images are very different to electro-optic images as a result of the way that the radar signal interacts with the scene and the way in

which the returned radar signal is processed and so understanding the phenomenology is a very important part of target recognition. The lecture then proceeds to discuss the principles of target detection and recognition in SAR imagery and how these need to be integrated into an end-to-end system to provide a full ATR capability. It will be seen that ATR relies upon having databases of example imagery of the targets of interest. Given the huge degree of variability intrinsic in radar imagery, it is generally impractical to populate such databases entirely with real imagery and so imagery obtained from radar scattering prediction tools applied to target models must also be used. The topic of training databases and target modelling is thus the next topic that is considered. A crucial aspect of any ATR system intended for military purposes is the ability to assess how well it will perform in given circumstances. Thus the subject of ATR performance assessment is an essential component of any discussion of ground target ATR and forms the last major subject of this lecture. Finally however, ATR is considered in the context that there is a continuum of problems to be solved of varying degrees of difficulty from very constrained scenarios to a completely general recognition system. All points within this continuum provide important military capability and systems that provide a level of radar ATR are already in service and helping NATO activities. ATR is a solvable problem and this is evident in operational systems. However, the requirement is also to provide greater capability by pushing the technology further along the difficult axis. To do this, a number of challenges must be addressed. This lecture thus concludes with a discussion of the current challenges facing ATR developers which must be overcome to achieve the future advanced capability that will allow NATO to most effectively fulfil its global role.

No references have been included in the text as the student is encouraged to go out and explore the concepts introduced using all the internet tools that are now available. However, a set of references has been included at the end which could act as the starting point for this exploration.

1.2 Definitions and Acronyms

The use of automatic techniques to classify radar data gives rise to different acronyms depending on the particular radar domain involved. It is usually known as Automatic Target Recognition (ATR) when dealing with air-to-ground activities which mostly use Synthetic Aperture Radar (SAR) imaging whilst it is usually known as Non-Cooperative Target Recognition (NCTR) for ground-to-air or air-to-air activities which mostly use High Resolution Range (HRR) profiles, Jet Engine Modulation (JEM) and Inverse SAR (ISAR) imaging.

What is meant by the word recognition? Care is needed in answering this as the ATR ‘vocabulary’ is still evolving. Indeed, only few terms have been standardized by NATO but, somewhat confusingly, of these some words actually have more than one formal NATO definition.

Taking the word “identification” as an example, the NATO AAP-6 Glossary of terms and definitions says that identification is the separation of friend and foe. However, in most modern conflicts a third class has to be added to this dual separation of the world to take into account the “neutral” targets that exist independently of the “classical” enemies. Moreover this third class tends to be the focus of most actual identification efforts to avoid collateral damage. It is clear that this identification process will depend on the people involved (countries, coalition forces) and on the context (in both space and time): it should also be taken into account that a civilian “neutral” vehicle may be easily turned into an enemy weapon. Currently, “identification” relies mostly on human interpreters or transponders like the Identification Friend Foe (IFF) system. It is somewhat difficult to characterize this definition purely in terms of scientific criteria and thus very difficult to automate.

In contrast, the word “recognition” as defined by the NATO AAP-6 Glossary of Terms and Definitions is a little more precise. The process is decomposed into a kind of classification “tree” in which the targets are categorized into more and more precise sub-classes as progress is made through the tree structure. Five major classification steps are then described:

- Detection, separating targets from other objects in the scene
- Classification, giving the target a meta-class such as aircraft, wheeled vehicle etc
- Recognition, specifying the class of the targets such as fighter aircraft, truck etc
- Identification, giving the sub-class of the target such as MIG29 fighter aircraft, T72 tank etc
- Characterization, taking into account the class variants such as MIG29 PL, T72 tank without fuel barrels etc
- Fingerprinting, leading to an even more precise technical analysis such as MIG29 PL with reconnaissance pod.

It can be seen that the boundaries between these decomposition steps cannot be clearly fixed for all problems and targets. Moreover, these definitions lead to the word “classification” being reserved solely to describe the process of meta-class separation whilst it is more often used by scientists to describe the whole process of assigning objects to categories irrespective of the status of those categories. This breakdown of definitions is even more obvious with the word “identification” which has been seen to take two different meanings within a single official glossary.

The main outcome of this discussion is thus to stress the need for a precise problem formulation and description of the operational conditions applying to the particular ATR problem under consideration.

2.0 SAR PHENOMENOLOGY

In order for an ATR system to make best use of all the information contained in a SAR image, it is essential that the ATR system formulation must incorporate precise knowledge of radar imaging phenomenology. In particular, it must be taken into account that the radar is moving relative to the target which itself may also move or change with time. Consequently, 2D SAR imaging is essentially a 2 step process:

1. High Resolution Range (HRR) profiles are acquired over time: HRR are instantaneous 1D projection on the Radar Line Of Sight (RLOS) of the whole scene observed over a large frequency bandwidth
2. 2D image formation: HRRs are integrated over time and focused which means that the potential target is observed over a time/angular domain.

As a result, the SAR image is focused on a projection plane (perpendicular to the apparent target/radar rotation axis), which implies:

- A sensitivity to the direction of illumination (shape, shadows etc)
- A sensitivity to the 3D geometry of target & ground (positions, overlays, masking etc)
- A sensitivity to the possible target motion or mobile parts (wheels, tracks, rotating parts like propeller, blades or antennas etc)

Looking at targets over some angular domain, it can be clearly seen that that the elementary scatterers from which the target signatures are composed are highly dependent on the angular directivity of the radar observation, e.g. Figure 2.1. The aspect angle dependency can be so strong that a target may look completely different when seen from directions separated by only few degrees apart. From an ATR point of view this means that it may be necessary to consider the target images taken at different observation angles as many “different” classes (each with the same “tag”).

On the other hand, a 360° azimuth integration of all individual aspect angles (see Figure 2.2) gives a clear view of the target that leads to its visual recognition. However, care must be taken as this type of acquisition may not be feasible except for turntable data or persistent surveillance of a target on a clear background.




Figure 2.1: Scatterer variability seen in sequence of SAR images taken every 1°.




Figure 2.2: Incoherent combination of target images taken every degree over 360° - the colour indicates the direction of illumination of the main contributing radar energy (courtesy of QinetiQ).

Whilst there is some frequency dependence in SAR imagery as seen in Figure 2.3 which shows an example of false colour image containing three different frequency bands, the target signature dependency on the radar frequency is probably less important than the angular one. However, within the identification

process, it will be necessary to use reference data taken around the same frequencies as the test data or the signatures may look too different.




Figure 2.3: Colour coding of SAR image combination with red at 8.82GHz, green at 9.37GHz and blue at 10GHz.





Figure 2.4: Sensitivity to geometry: shadows and specular returns (left) and multipath off ladders (right).

Sensitivity to the 3D target geometry is another key point: it is possible to see through very simple examples how the radar energy interacts with the target and with the ground via multiple bounces resulting in layover effects, inner shadowing and multi-path interaction of close objects. These effects, combined with the specular nature of the radar reflection, produce target signatures that “look very fuzzy” when compared to their optical equivalent: this may explain the relative difficulty to train human interpreters to

work on radar images. However, these effects may give geometrical “fingerprints” that are well suited for automated ID. Figure 2.4 shows examples of shadows and specular reflections from planes and multipath effects from ladders on the side of oil tanks.

In summary, when considering the SAR phenomenology that should be included in an ATR system, it is important to avoid making assumptions based on experience with optical systems. SAR images have particular sensitivities to imaging geometry, radar parameters and radar scattering mechanisms which need to be taken into account when attempting to characterise target classes.

3.0 TARGET DETECTION & RECOGNITION

3.1 Pre-Screening Stage

The first stage in the target recognition process is to automatically detect potential targets in the scene which can then be passed on up the processing chain for further analysis. This task in itself can contain a number of stages in which candidate detections are identified and then filtered to reject those that do not meet the criteria for being a potential target. For this reason, this stage of the processing chain is often called “pre-screening”.

The various stages in a possible approach to target pre-screening are shown in Figure 3.1. Given a SAR scene, the first stage in the process is to perform a single pixel detection which flags up pixels which are anomalously bright in comparison to their neighbouring background pixels. A mask is placed around the pixel under test to exclude any pixels which might also belong to the target and hence bias the calculation of the background statistics. An outer ring of pixels is then used to calculate the background statistics, typically the mean and standard deviation. If the pixel under test exceeds the background mean by more than a given number of standard deviations then a detection is declared. This type of approach belongs to the class of Constant False Alarm Rate (CFAR) techniques of which there are many variants to cope with problems such as the ring being used to estimate the background containing non-background objects. However a discussion of such variations is beyond the scope of this lecture.

A number of single pixel detections will be obtained in this way but there is no understanding of which detections may belong to the same potential target. For this reason, a clustering procedure is used as the next stage. One way to do this is to start with a detection and to add other detections to the same cluster provided that they are not too far from any detections already contained in the cluster. The definition of “too far” will depend on the resolution of the system and the type of targets anticipated to be in the scene. This is the approach illustrated in Figure 3.1. Alternative approaches are possible such as those based on the morphological operations.

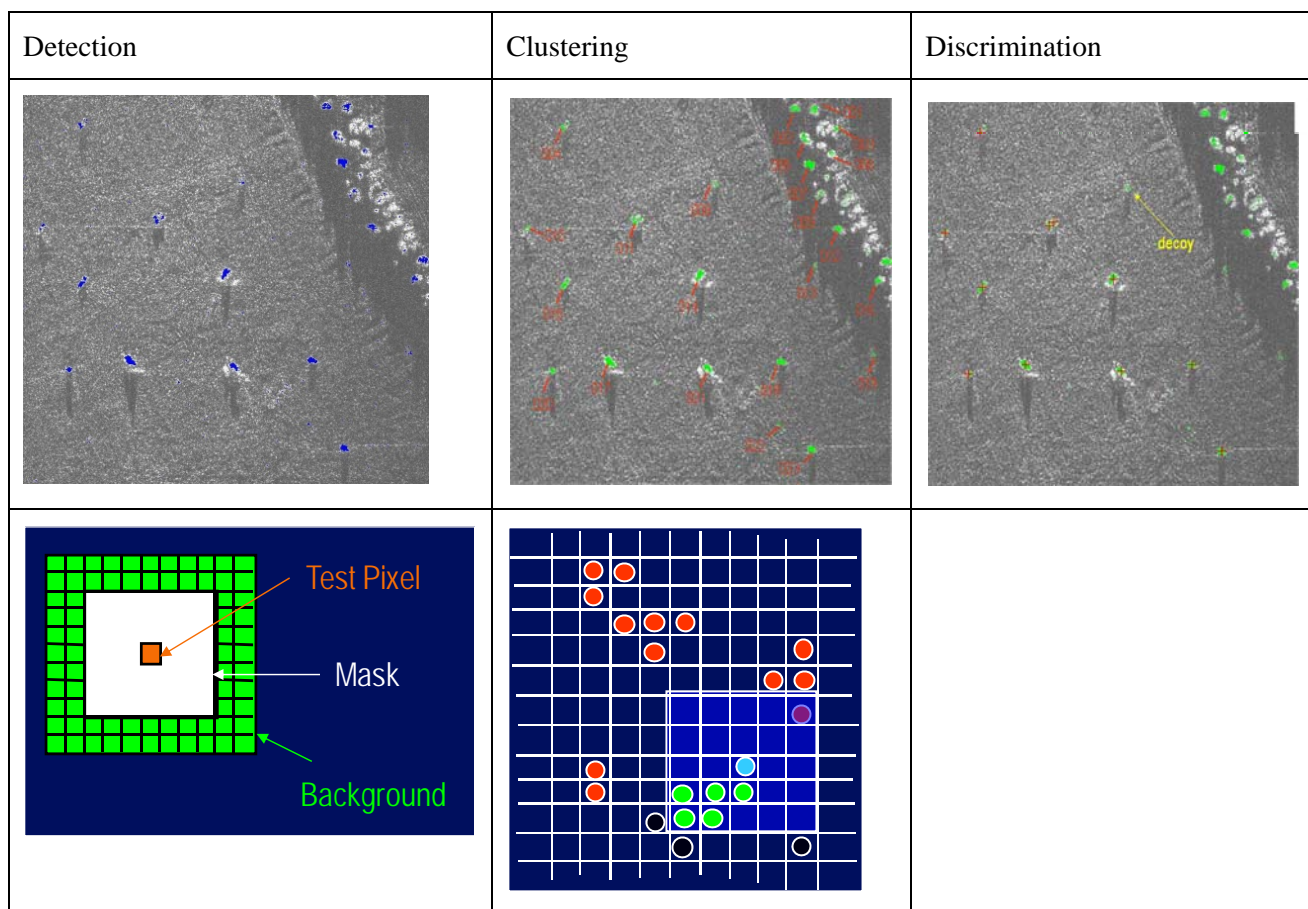


Figure 3.1: Illustration of the detection, clustering and discrimination stages of a possible approach to target pre-screening (courtesy of QinetiQ).

The final pre-screening stage is to examine each cluster and measure some simple discriminant values such as size and power. It is possible then to reject a number of candidate targets as being more likely to be discrete clutter objects such as trees rather than man-made objects. In Figure 3.1, a number of clusters along the vegetation boundary as well as a decoy target have been rejected by this process.

3.2 Classification: Template-Matching

Once the candidate targets have been identified by the pre-screening process then the classification process can begin. One conceptually simple approach is to compare the object under test with example images of the various possible targets known to the system. This is illustrated in Figure 3.2. As discussed with regard to phenomenology, radar images are very variable as a function of imaging geometry. So for each possible target class, the database must contain example images of that target at all possible geometries. Figure 3.2 illustrates a target database containing target images over 360° of aspect angle variation although in general elevation angle and many other degrees of freedom would need to be taken into account.

A measure of similarity is required to perform the comparison between the object under test and the images in the database. A natural measure is the correlation coefficient between the two images which has a maximum value of unity for two identical images. Given this maximum, it is reasonable to set a threshold correlation value such that if this threshold is not exceeded then a target classification is not made and the object is declared to be unknown. The option of making an “unknown” declaration is very important and will be discussed further with regard to performance assessment.

Whilst conceptually simple, the problem with this template-matching approach to classification is that the required databases can be huge when there are many target classes and many potential degrees-of-freedom. Thus template-matching has an important role for classification problems that are relatively constrained but an alternative approach is required for less constrained problems.




Figure 3.2: Template matching involves comparing the object under test with a number of examples of possible targets (courtesy of QinetiQ).

3.3 Classification: Feature-Based

Feature-based classification provides an alternative to template-matching which solves the issue of a requirement for huge databases of imagery by representing target classes in terms of measured features that are intended to characterise the unique properties of the target class. Before any features can be measured, it is important to establish accurately which pixels belong to the target and which to the surrounding clutter. One way of doing this is to use an active contour or snake algorithm. An initial contour is placed around a point defining the position of the target (e.g. the mean position of the detections comprising the cluster). This is shown as the approximately circular inner red circle in the image on the left in Figure 3.3 which is defined by a number of node positions. This contour is iteratively adapted by randomly moving the positions of the nodes. An objective function is defined that measures how well the statistics inside and outside of the contour match the assumed statistics for target and background. An annealing approach is taken such that a change is accepted if the objective function increases but a change is also accepted with some probability if the objective function decreases. This probability decreases as the number of iterations increases. The aim of this process is to avoid the iterative procedure becoming stuck in local maxima rather than finding the global maximum. Once the process has converged, an outline such as the outer red line shown in Figure 3.3 will be obtained.

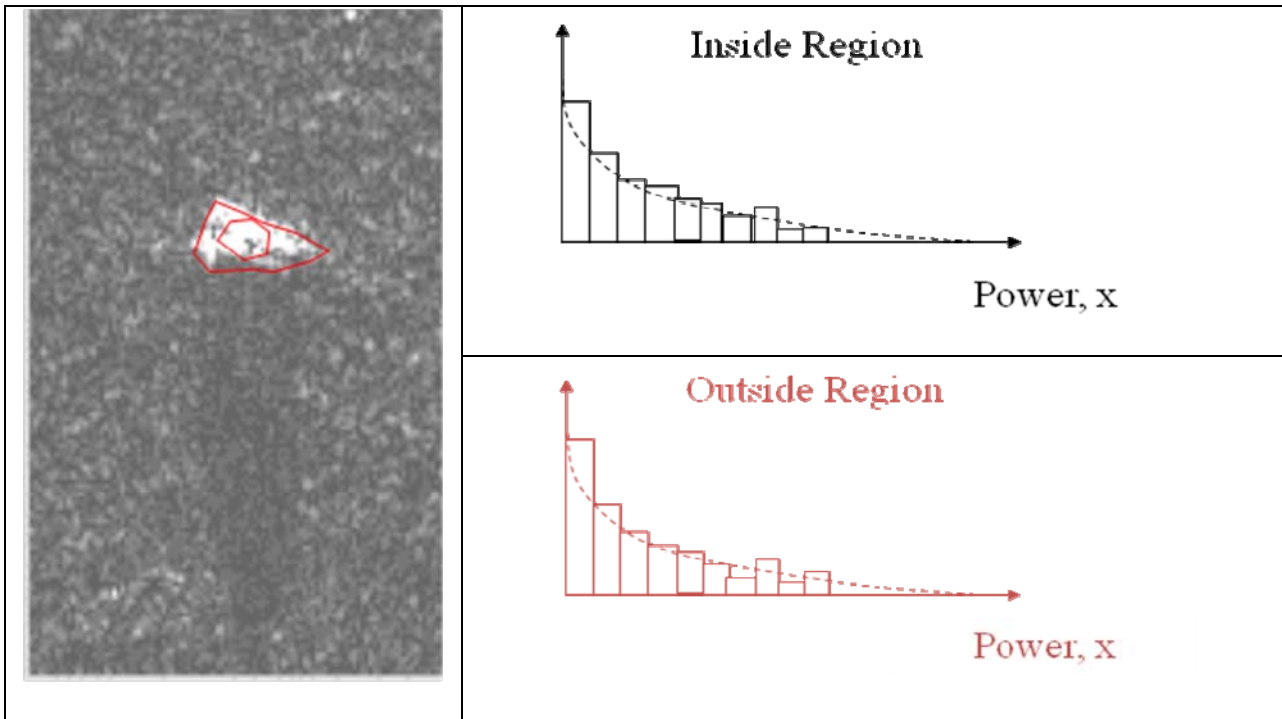


Figure 3.3: An active contour starts with an initial circular contour which is adapted iteratively until the statistics inside and outside of the region match a given model (courtesy of QinetiQ).

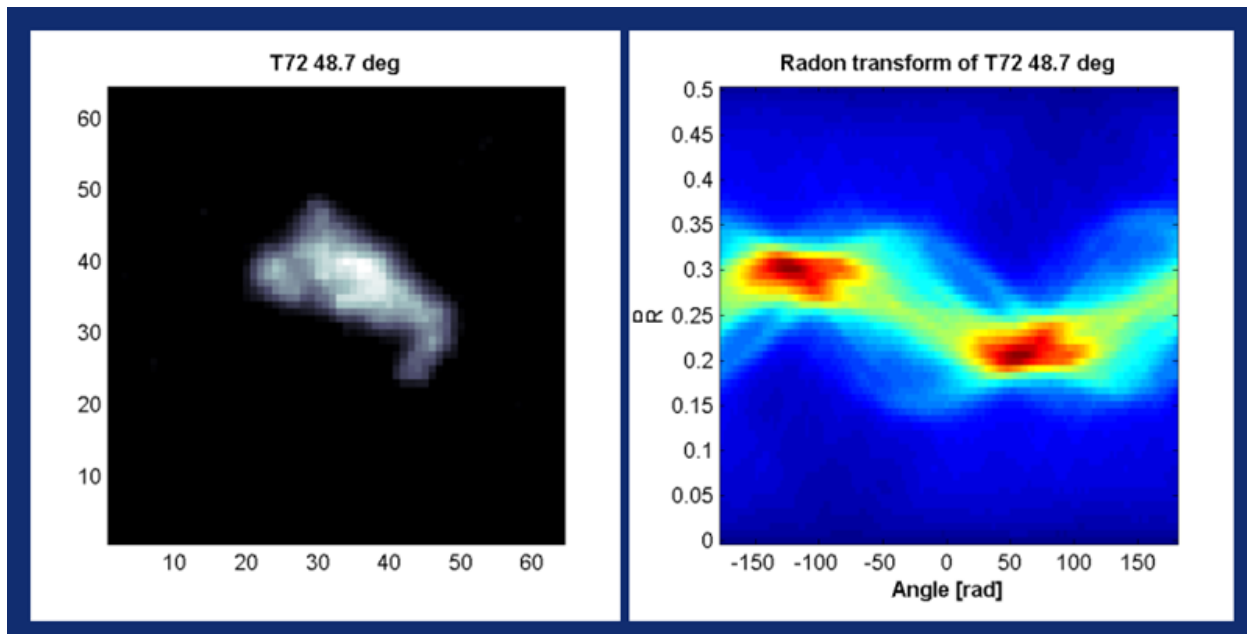


Figure 3.4: The Hough transform maps lines in the image to points in the transform domain so that bright lines show up as peaks (courtesy of FGAN).

Prior to measuring features, it can also be convenient to estimate the pose of the target, i.e. its angle with respect to the axes of the image. One way of achieving this is to use the Radon transform. The Radon transform defines lines in the image in terms of a distance from the origin and an angle with respect to the x-axis. The pixel values are summed along these lines and the result placed in the transform domain at the

corresponding (angle, distance) co-ordinates. Lines which contain many bright pixels are seen as peaks in the Radon transform domain so that identifying the brightest peak will identify the brightest line in the image. There is a 180° ambiguity but this is not of consequence. Having determined the pose of the target in this way it is then possible to measure the length and width, for example, in a consistent way.

The choice of features used to represent the target classes is the key to classification performance. There are many features which have arisen from the general pattern recognition literature including geometric features such as target dimensions, moments of inertia, fractal measures and Fourier coefficients as well as radiometric features such as mean, standard deviation, spatial correlation measures and the proportion of energy contained in the brightest pixels. A reasonable level of classification can be obtained using such features but, as will be discussed later, tuning features to the known specific properties of SAR images should be the aim of an ATR system design.

The enormous variability of target appearance as seen by radar is a key challenge. Figure 3.5 shows an example of the same target seen at different aspect angles in a SAR image. The appearance of targets can vary so much that a sensible approach is to essentially treat intervals of aspect angle as a different class.




Figure 3.5: Ground target image aspect angle variation and air target range profile variation (courtesy of QinetiQ).

Once a set of features has been established, these features are measured for examples of the different target classes, i.e. the training set. The feature values for an object under test are also measured and compared with those of the known target classes. If there is a sufficiently good match between the test and training values for a particular class then a classification is declared. This basic principle is illustrated in Figure 3.6 for a two class problem where two features are being used. The red crosses mark the positions of feature values obtained from the red class training examples and the yellow crosses the same for the yellow class training examples. A decision boundary must be drawn such that a test example with features that lie on one side will be declared red and on the other will be declared blue.




Figure 3.6: Illustration of decision boundary for two class problem (courtesy of QinetiQ).

There are many algorithms available for determining the decision boundary. In Figure 3.6, contours have been shown which arise from approximating the distributions of the feature values by two-dimensional Gaussian distributions. It is then straightforward to set up the decision via statistical considerations and this is known as Bayesian classification. Other techniques include nearest neighbour, linear discriminant analysis, neural networks and support vector machines although it is beyond the scope of this lecture to go into these methods in detail.

3.4 End-to-End ATR Processing Chain

One demonstration of a SAR ATR system has been given QinetiQ in their SAR Machine-Aided Recognition Toolbox (SMARTbox) as illustrated in Figure 3.7. This implements the entire processing chain from detection of potential targets to recognition. The idea for this demonstration is to have a man-in-the-loop. The SAR image is first processed for detection of potential man-made objects which are indicated on the SAR image by cross-hairs. The operator can then click on one of the detections for further processing. A feature-based classification is performed using features derived from both the target itself and the shadow which are both automatically delineated. Possible target classifications are then presented to the operator with an associated probability or confidence level. The operator is then able to request a prediction of the target image that would be expected to be seen for the declared target type. A slider bar allows the operator to view this prediction from a number of different aspect angles to assess the validity of the declaration. In this way, the tool acts very much as an automated operator aid. In this way it reduces operator workload but any final decision still requires a man-in-the-loop which is an important consideration given rules-of-engagement.

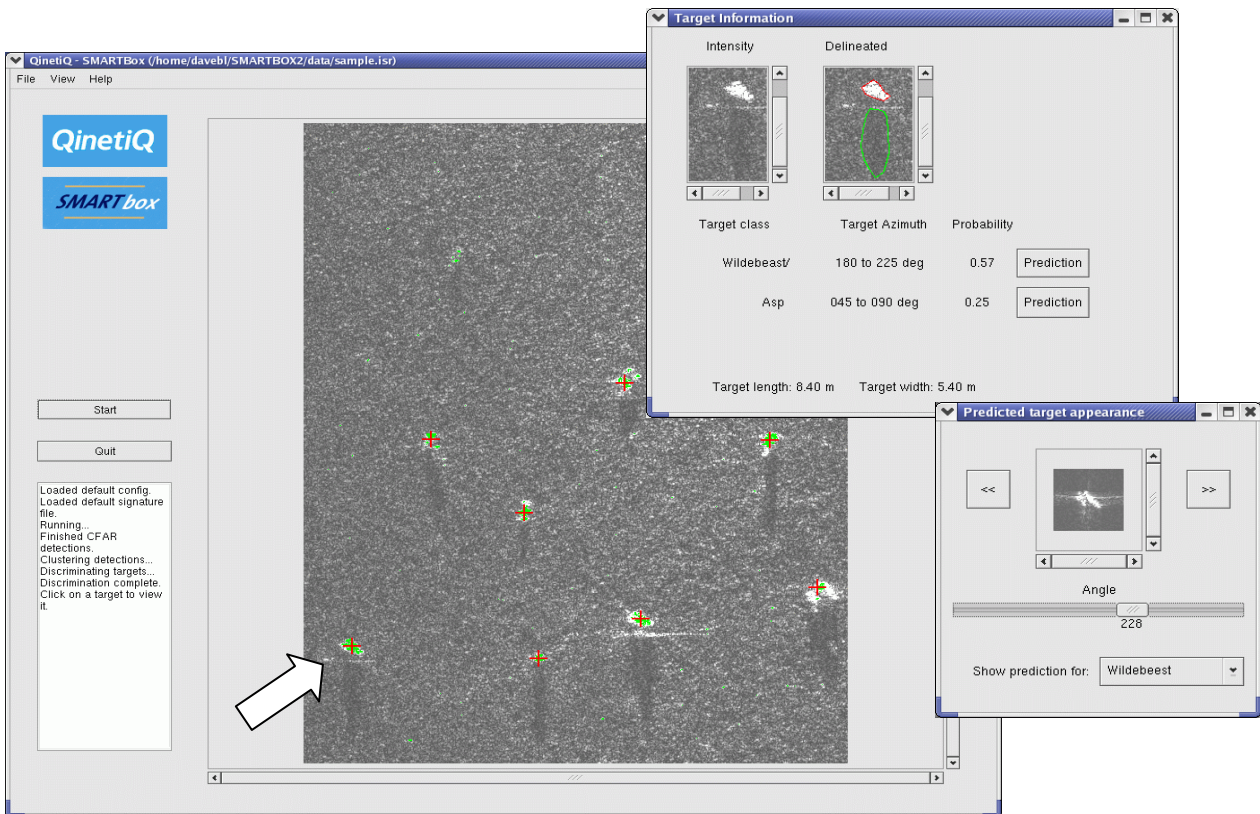


Figure 3.7: Feature-based ATR demonstration (courtesy of QinetiQ).

This is an effective demonstration of an ATR system and there are many other such demonstrations available from research groups around the world and in the literature.

3.5 Choice of Features

Features are measures of target characteristics which hopefully provide some degree of discrimination between target classes. A large number of possible target features have been proposed in the literature many of which arise from the general pattern recognition literature. As a result, measures such as length, width, compactness, elongation, pixel statistics, rank-fill ratio (concentration of energy into a small number of pixels), fractal dimension etc have been used to varying degrees in proposed ATR schemes.

Such pattern recognition features have value but in terms of finding features which provide the greatest robustness to target variability, it may be argued that features which relate specifically to the underlying physical structure of the target are likely to be the most robust as it will be the same underlying structure that is present whatever the imaging geometry or level of obscuration. Figure 3.8 shows a composite SAR image (this is an incoherent combination of images taken over a full 360° of aspect angle) of a vehicle where persistent or dominant scattering events have been associated with real structures on the vehicle.




Figure 3.8: Associating characteristic dominant scatterers with vehicle structure (courtesy of QinetiQ).

It is such scattering events which may hold the key to robust feature-based classification. However, on the basis of a single channel, single aspect image of a target there may not be sufficient information. However, polarimetric techniques allow more sophisticated characterisation of scatterers whilst multiple aspect collects or non-straight trajectories allow 3D information to be obtained. The future availability of more advanced collection modes such as these will thus open up additional avenues for defining robust features.





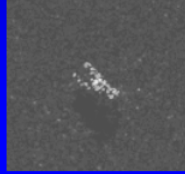
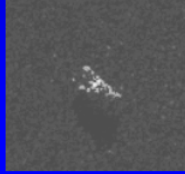







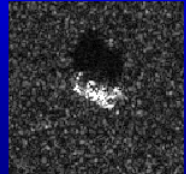
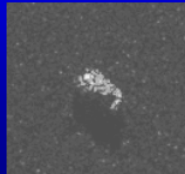
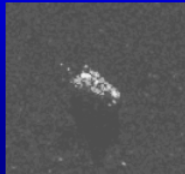
4.0 DATABASES & MODELLING

Building an ATR reference data-bank for radar presents greater challenges than, for example, building an EO database when it may be sufficient to use a handheld camera or search over the internet for a picture of a particular vehicle. To obtain representative SAR imagery of targets in the field, multiple flights are needed over an area where targets are deployed to get enough data for multiple angle acquisition. The experiment will also be more representative if it is possible to provide variants of same targets on the field. The first large data collection was done in the 90's by DARPA and is known as the MSTAR database. A small part of this database was then released to public at the beginning of the 2000's and has motivated a substantial body of SAR ATR research which has been reported in the literature. Figure 4.1 provides examples of targets and imagery from the MSTAR dataset.

However, when taking into account the totality of real world variability, it is clear that it is not feasible to collect real data that would cover all the operating conditions that may be encountered in modern conflicts. Thus it is essential that reference databases are populated to a large extent by other means such as turntable measurements or modeling. Figure 4.2 shows an example of a radar image of a tank obtained from a turntable. This provides a very controlled situation for image formation in which the radar is stationary and relative motion is introduced by rotating the turntable. Image formation is then performed using inverse SAR techniques. It is important to be aware that such imagery may not be entirely representative of target imagery taken in the field. In particular, the background clutter against which the target appears and hence ground / target interactions may not be representative.

MSTAR data set

- approx. 200-250 images per type or aspect angle (128x128 complex)
- depression angle training 17°, test 15°
- confusion vehicles (not considered yet)

	BMP2-9563	BMP2-9566	BMP2-C21	
				
				
D7	BTR70-C71	T72-132	T72-812	T72-S7
				
				

14 February 2003

Figure 4.1: The MSTAR public data set collected under a DARPA program.




Figure 4.2: An example of a radar image of a tank obtained from turntable measurements.

An alternative, as mentioned earlier, is to use modeling techniques which predict the radar signature from a CAD model of the target and thus allow images to be simulated. There exist various simulation tools

ATR of Ground Targets: Fundamentals and Key Challenges

from sophisticated electromagnetic codes that solve Maxwell equation, compute surface currents and so on, to more simplistic ones that will only generate the outline of the target's shadow. The compromise is often between accuracy and computation speed. "Exact" codes may take days to compute a single aspect angle image and so the simulation tool really has to be fitted to the ATR philosophy. This trade-off between accuracy and speed is illustrated in Figure 4.3.

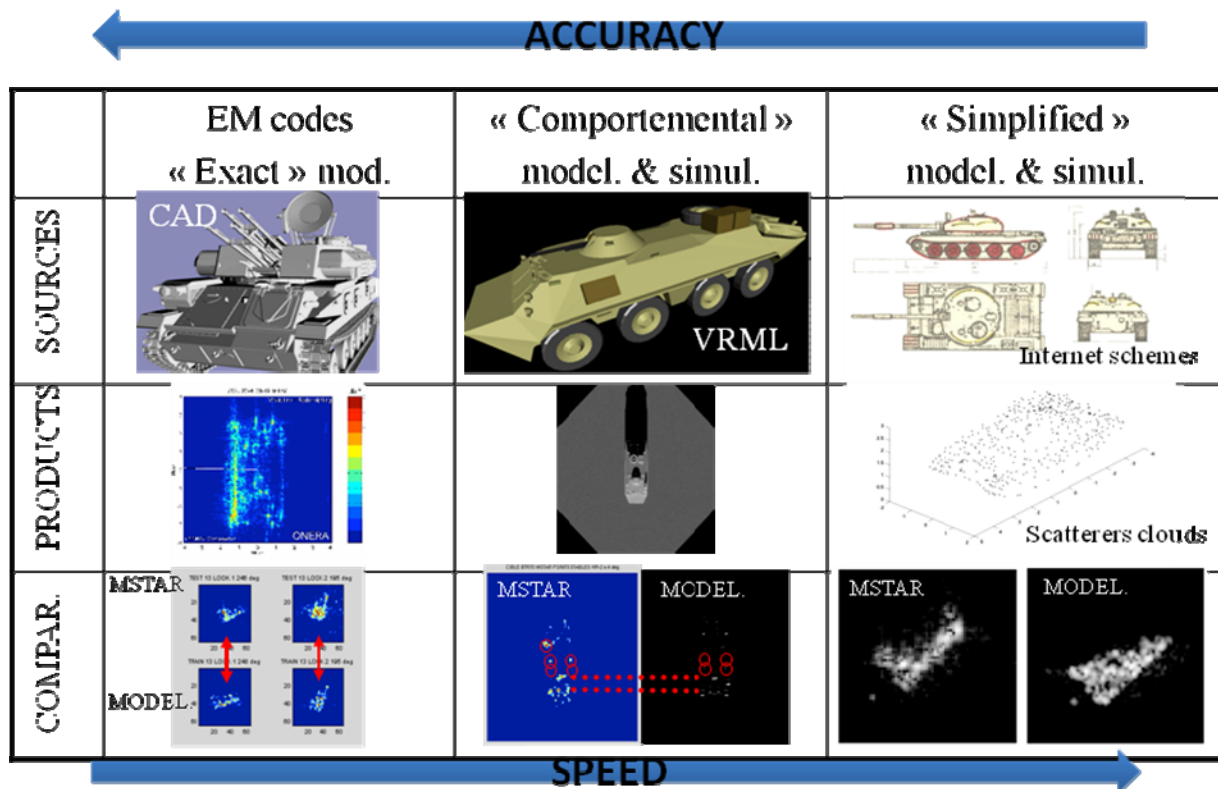


Figure 4.3: Trade-off between accuracy and speed for SAR image simulation.

A key consideration is how the accuracy of the CAD model used affects the accuracy of the simulation and hence the performance of the ATR system. If the ATR performance is critically dependent on having a very exact CAD model representation then this is unlikely to be a robust solution since the actual targets in the field may easily vary from the CAD model representation. Also there is a question as to how the CAD models are to be obtained. If the vehicle is available then it may be possible to use laser scanning techniques to obtain an accurate CAD model. However, in some situations the vehicle will not necessarily be freely available to measure and so it will be necessary to produce a CAD model from possibly a limited number of photographs of the target. A continuing challenge is to understand the impact of CAD model fidelity on ATR performance and how this relates to the accuracy of CAD models that can be obtained from, for example, photographs.

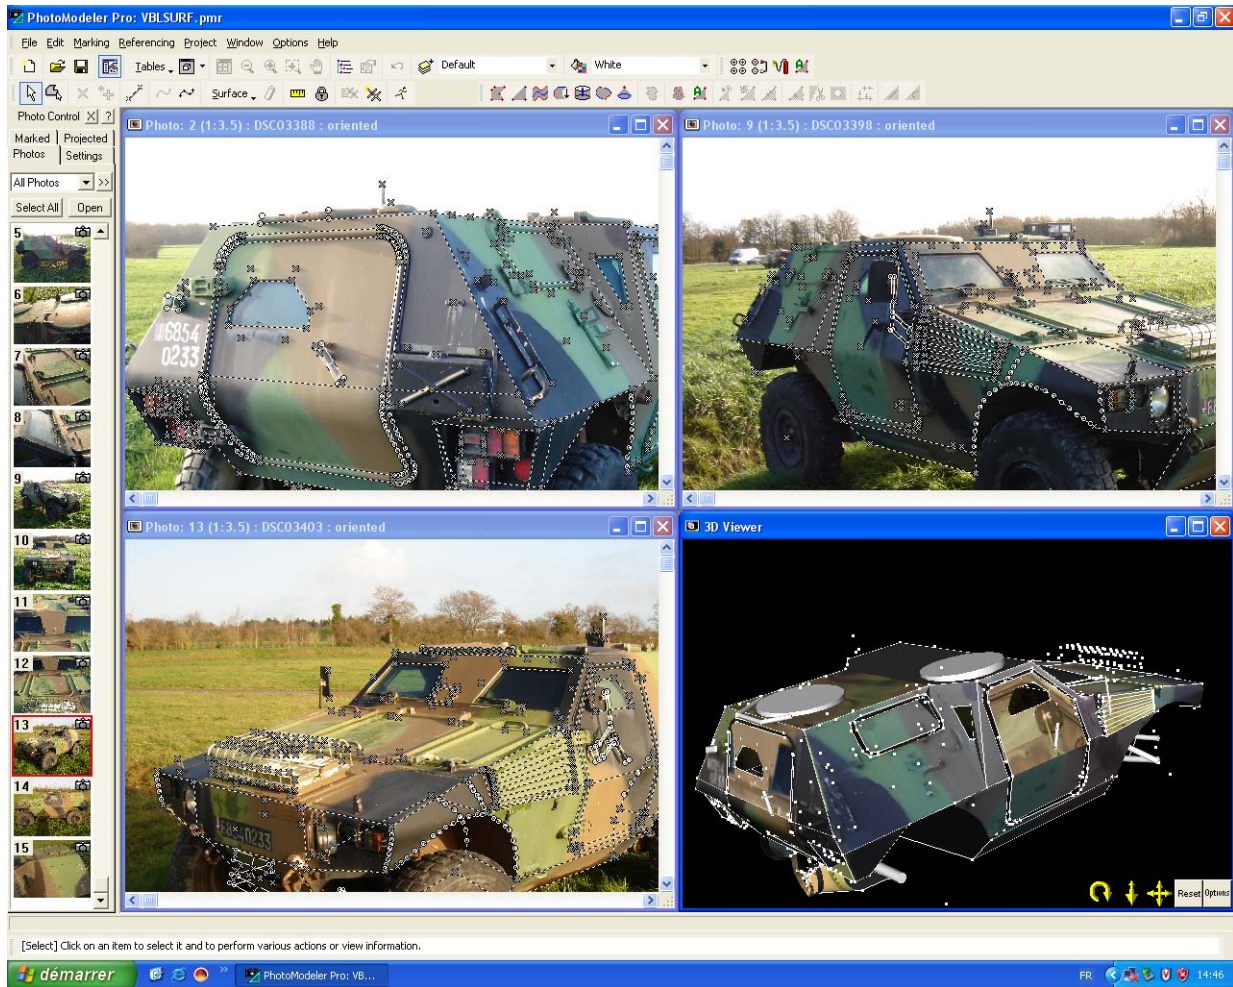


Figure 4.4: Construction of CAD models from photographs.

In summary, for most ATR applications some form database of imagery simulated from CAD models is likely to be necessary. However, it needs to be taken into account that the target signature is sensitive to geometry and radar parameters, the targets may have many variants such as articulations or extra equipment attached and that targets are not always available for detailed analysis. There is a need to build CAD models either via techniques such as laser-scanning or from photographs but a key question is how the accuracy and complexity of the model in terms of number of facets, parts, articulations and material properties affects simulation accuracy and hence ATR performance. There is thus a trade-off between accuracy and computation speed which needs to be taken into account when designing an SATR system.

5.0 PERFORMANCE ASSESSMENT & ATR THEORY

5.1 Receiver Operating Characteristic (ROC) Curves

As has been seen previously, the end-to-end ATR processing chain contains a pre-screening stage and a classification stage. The pre-screening stage is essentially a two-class problem which aims to identify targets and reject clutter. For this type of problem, the receiver operating characteristic (ROC) curve is a convenient measure of performance.

Figure 5.1 illustrates the underlying principle. It is assumed that some discriminant value is measured to determine whether an object is target or clutter. For many examples of targets this discriminant value will have a probability density function (PDF) and similarly for clutter as shown. A threshold is used to decide whether an object is target (threshold exceeded) or clutter (threshold not exceeded). Given a threshold, it is possible to calculate the probability of detection (PD) by integrating the target PDF from the threshold to infinity and the probability of false alarm by integrating the clutter PDF from the threshold to infinity as shown. As the threshold is varied, the PD and PFA will vary and this is shown for two values of the threshold.

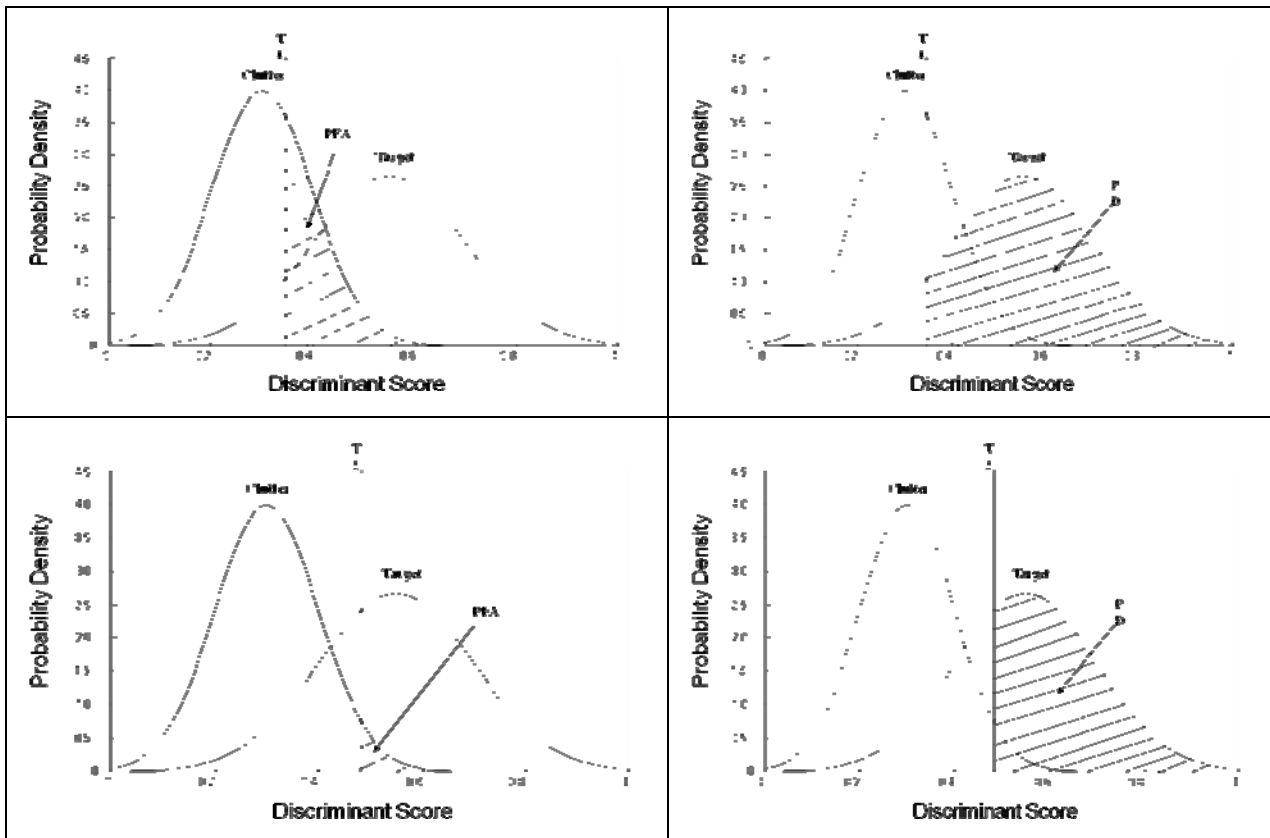


Figure 5.1: The effect on probabilities of detection and false alarm of altering the threshold.

As the threshold is varied, the resulting values of PD and PFA can be plotted against each other as shown in Figure 5.2. The ROC curve is thus the locus of PD versus PFA as defined by the threshold parameter. However, care must be taken as the same threshold will not correspond to the same PD or PFA for different missions and so maintaining equivalent performance in different circumstances can be challenging.

Figure 5.3 provides an interesting example of the use of ROC curves to show the performance of a target detector. In this example, the targets were either without optical camouflage netting (bare) or with it (cam). Also they were either in the open (open) or in light vegetation (scrub). In Figure 5.3 it can be seen that better detection performance is obtained when the target is in the open (top two curves) than when in scrub (bottom two curves). However, in both open and scrub the detection performance is better when optical camouflage netting is used than when it is not. The reason for this has not been determined conclusively but it is speculated that the netting was wet and it was a windy day. Hence in the resulting radar imagery, the netting caused a bright and smeared return which was more easily detected. The use of ROC curves was thus important in understanding what was going on in this example.




Figure 5.2: ROC curves are the locus of PD versus PFA as the threshold parameter is varied. Care must be taken as the same threshold will not correspond to the same point for different missions.




Figure 5.3: An example of the use of ROC curves.

5.2 Confusion Matrices

At the classification stage of the ATR chain, there will normally be multiple possible classes to which a target may belong and so the concept of a ROC curve is less useful. In this case, the performance of the system can be more fully understood through use of a confusion matrix. Figure 5.4 shows a basic example of a confusion matrix for a four class problem. In this example there are two classes of friendly vehicles (blue) and two classes of enemy vehicles (red). The numerical entries in the main body show the proportion of targets of a given true class that have been declared as a given declaration class. The probability of correct classification (PCC) is the probability that a target of a given class is declared as that class and this is shown in the right hand column. However, of more relevance to the war-fighter is the probability of correct label (PCL) which is the probability that a target that is declared to be of a given class is actually a member of that class. The PCL is shown in the bottom row and clearly tells a different story. For example, targets from class Red2 are declared as such 95% of the time. However, if a target is declared as Red2, it is actually only a member of that class on 83% of the time.

Truth	ATR System Output				PCC
	Declaration				
	Blue1	Blue2	Red1	Red2	
Blue1	0.78	0.03	0.06	0.13	0.78
Blue2	0.01	0.95	0.02	0.02	0.95
Red1	0.10	0.03	0.83	0.04	0.83
Red2	0.02	0.01	0.03	0.95	0.95
PCL	0.86	0.93	0.88	0.83	

Figure 5.4: A basic confusion matrix for the performance of a four class classifier.

This distinction between PCC and PCL becomes more evident when the more realistic situation is considered in which the targets under test contain vehicles of two class types that are not known to the classifier, i.e. confuser classes. Typical results in this case are shown in Figure 5.5. It can be seen that the PCC is unaltered in this case but the PCL is dramatically affected. Now a target declared as Red2 will only actually belong to the Red2 class 57% of the time and the effect is similar for all classes.

Truth	ATR System Output				PCC
	Declaration				
	Blue1	Blue2	Red1	Red2	
Blue1	0.78	0.03	0.06	0.13	0.78
Blue2	0.01	0.95	0.02	0.02	0.95
Red1	0.10	0.03	0.83	0.04	0.83
Red2	0.02	0.01	0.03	0.95	0.95
Conf1	0.24	0.28	0.22	0.27	N/A
Conf2	0.28	0.16	0.31	0.25	N/A
PCL	0.55	0.65	0.56	0.57	

Figure 5.5: A confusion matrix for the performance of a four class classifier including confuser targets.

This illustrates a problem with the classifier whose performance is being assessed by these confusion matrices. Basically it is being forced to classify every target under test as one of the four classes it knows about even though it is being exposed to targets outside its database as is virtually inevitable in ground target recognition situations. It is thus important, as has been mentioned earlier, to include an unknown class in the classification procedure. Figure 5.6 then shows an example output from the classifier. It can be

seen that the PCC values are reduced as some of the previous correct declarations were not actually sufficiently confident and have now been declared as unknown. However, the PCL is now significantly improved over the case when a forced decision was made. This emphasises the importance of including an unknown class.

Truth	ATR System Output					PCC
	Declaration					
	Blue1	Blue2	Red1	Red2	Unknown	
Blue1	0.70	0.01	0.00	0.12	0.17	0.70
Blue2	0.00	0.95	0.01	0.00	0.04	0.95
Red1	0.09	0.01	0.79	0.02	0.09	0.79
Red2	0.01	0.01	0.02	0.95	0.02	0.95
Conf1	0.04	0.01	0.00	0.03	0.93	N/A
Conf2	0.04	0.00	0.07	0.04	0.85	N/A
PCL	0.80	0.96	0.89	0.82	N/A	

Figure 5.6: A confusion matrix for the performance of a four class classifier including confuser targets and an unknown class.

Another factor that needs to be taken into account is that some declaration errors can have more serious consequences than others. For example, if Blue is declared but the target is Red then an enemy attack may not be averted whilst if Red is declared but the target is Blue then friendly fire may occur. Thus it is important to include the probability of critical error (PCE) as has been done in Figure 5.7 to understand the significance of a false declaration.

Truth	ATR System Output					PCC
	Declaration					
	Blue1	Blue2	Red1	Red2	Unknown	
Blue1	0.70	0.01	0.00	0.12	0.17	0.70
Blue2	0.00	0.95	0.01	0.00	0.04	0.95
Red1	0.09	0.01	0.79	0.02	0.09	0.79
Red2	0.01	0.01	0.02	0.95	0.02	0.95
Conf1	0.04	0.01	0.00	0.03	0.93	N/A
Conf2	0.04	0.00	0.07	0.04	0.85	N/A
PCL	0.80	0.96	0.89	0.82	N/A	
PCE	0.11	0.02	0.09	0.16	N/A	

Figure 5.7: A confusion matrix for the performance of a four class classifier including confuser targets, an unknown class and showing probability of critical error.

Yet another factor that needs to be considered is the “order of battle”, i.e. the probable number of units present. In this example it is assumed that there are 10 times as many Red units as Blue units and confusers. The confusion matrix declarations in Figure 5.8 have been adjusted to give the number of declarations if there were 1000 of each Red unit and 100 of each other unit. This does not change the PCC but has a drastic impact on the PCL and PCE. In this situation, the war-fighter would have a lot more confidence acting on a Red declaration but there would be little confidence in a Blue declaration.

Truth	ATR System Output					PCC
	Declaration					
	Blue1	Blue2	Red1	Red2	Unknown	
Blue1	70	1	0	12	17	0.70
Blue2	0	95	1	0	4	0.95
Red1	90	10	790	20	90	0.79
Red2	10	10	20	950	20	0.95
Conf1	4	1	0	3	93	N/A
Conf2	4	0	7	4	85	N/A
PCL	0.40	0.81	0.97	0.96	N/A	
PCE	0.57	0.17	0.01	0.02	N/A	

Figure 5.8: A confusion matrix for the performance of a four class classifier including confuser targets, an unknown class and showing probability of critical error.

Thus it has been seen that the confusion matrix provides a powerful means of representing classifier performance but it is important to take all factors into account when interpreting the results. A few key considerations have been introduced here but this has by no means been exhaustive.

5.3 Operational Assessment

There are many, many sets of ATR performance results published in the literature in the form of confusion matrices. Almost every paper presents a new technique and shows that it gives better performance than existing techniques. How can this be? Figure 5.9 illustrates a few of the ways in which reported performance results may not truly reflect the operational performance of the system. This is a significant issue and emphasises the need for a strategy for ATR performance assessment at a national or international level in which some independent body maintains a set of test data which is used only for performance assessment and rigorously dictates the assessment procedure. Without this, most reported ATR performance figures are meaningless.

Performance assessment should, of course, reflect the performance that will be achieved operationally. However, a major issue for operational assessment is that algorithm development and performance assessment take place in circumstances that are not necessarily representative of the real operational environment. This is illustrated in Figure 5.10. In this illustration, the reference conditions relate to data that was gathered to support algorithm development, the model-based conditions relate to data that has been used to populate databases, the test conditions relate to data collected to support algorithm performance assessment and the operational conditions relate to data to which the system is to be applied operationally. As has been illustrated conceptually, the intersection of all these conditions is very small and, in particular, there is a large portion of operational space which is not represented. The challenge for the future is to develop assessment methodologies which provide the procurer of such systems greater confidence that they will indeed perform as intended when used operationally.

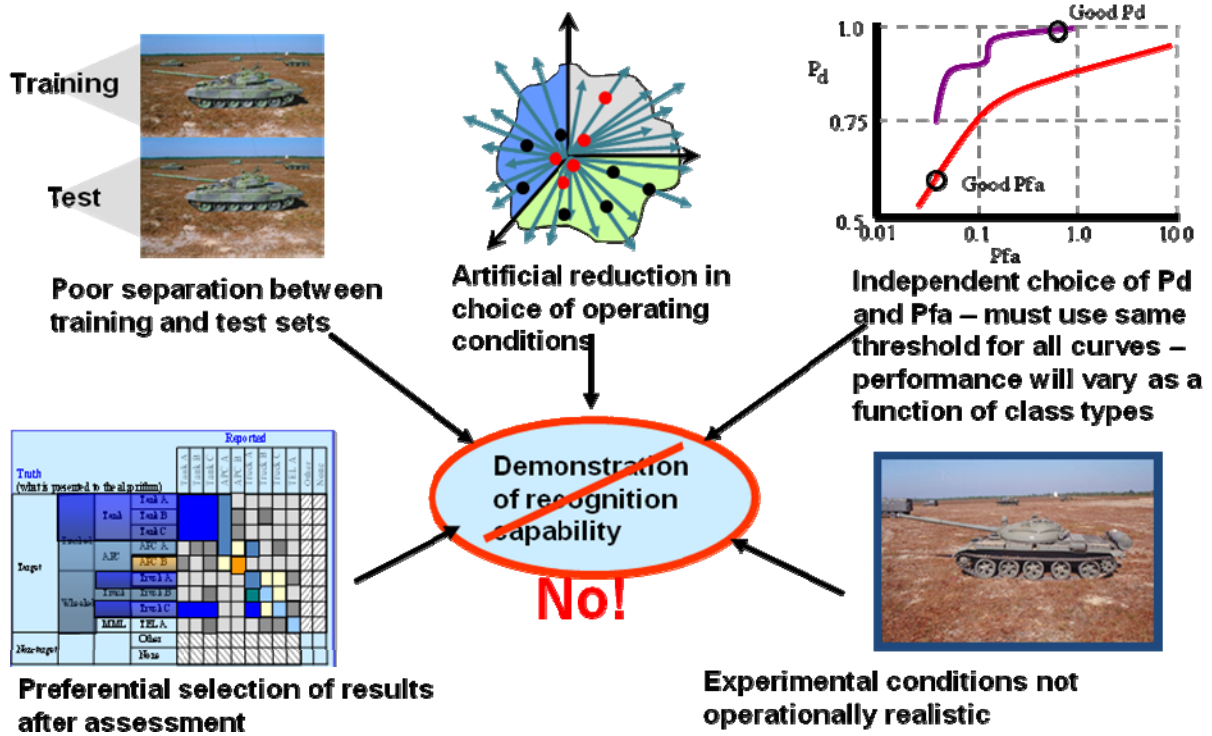


Figure 5.9: Some common mistakes made when assessing performance (courtesy of SET-111).

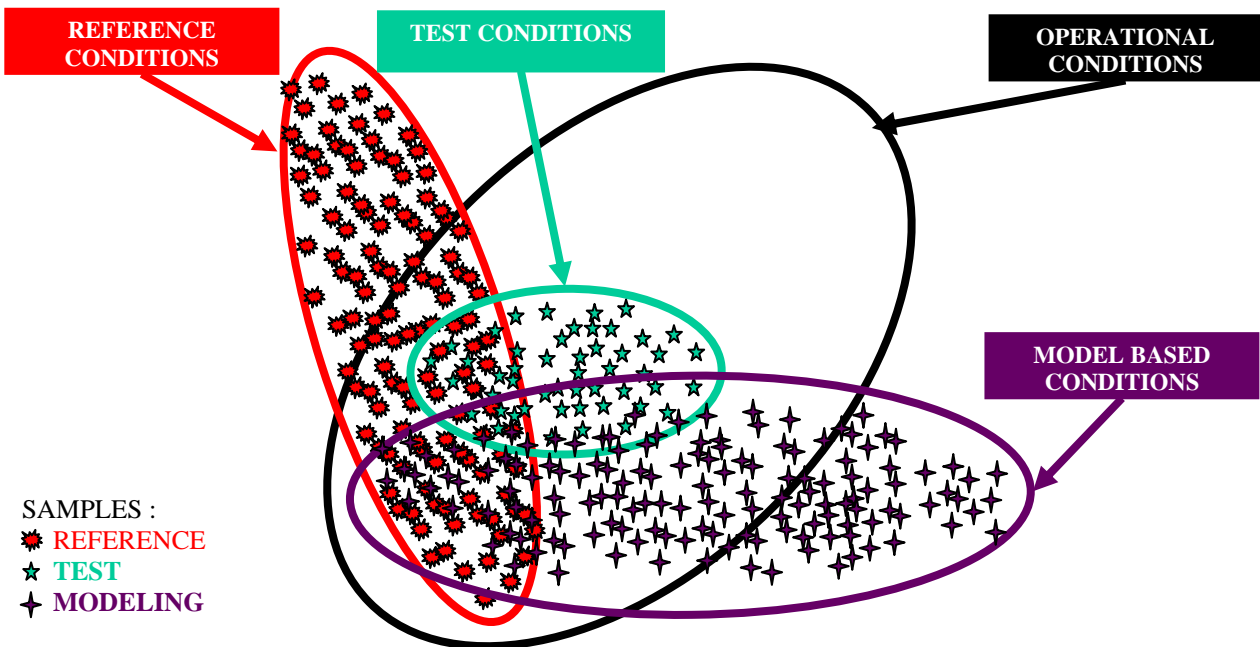


Figure 5.10: Illustration of performance assessment issues (courtesy of SET-111).

6.0 KEY CHALLENGES

In this section, some of the challenges are discussed which must be addressed to achieve the future advanced radar ATR capability that will allow NATO to most effectively fulfil its global role.

There are many who are sceptical about radar ATR ever providing operational capability. However the fact is, of course, that some degree of operational radar ATR capability has already been achieved. A fact that is not always appreciated is that a requirement for ATR does not represent a single problem but actually represents a continuum of problems of varying degrees of difficulty from very constrained scenarios to a completely general recognition system. All points within this continuum provide important military capability and systems that provide a level of radar ATR are already in service and helping NATO activities.

Figure 6.1 gives a simplified (and UK-centric) illustration of the development of radar ATR. From the mid-80's there was research into image formation and optimum detectors which were subsequently incorporated into land-based and maritime systems. Subsequent research in the 90's resulted in operational capability such as the man-portable MSTAR system which classifies between tracked & wheeled vehicles and personnel on the basis of Doppler and the Brimstone missile which classifies between valid and non-valid targets on the basis of an automated algorithm. Ongoing research is aiming to push further the difficulty and complexity level of ATR algorithms for insertion into future platforms. The point is that ATR is a solvable problem which is evident from the existing operational systems. However, the requirement is to provide greater capability by pushing the technology further along the difficult and complexity axis. To do this, a number of challenges must be addressed as will be discussed next.

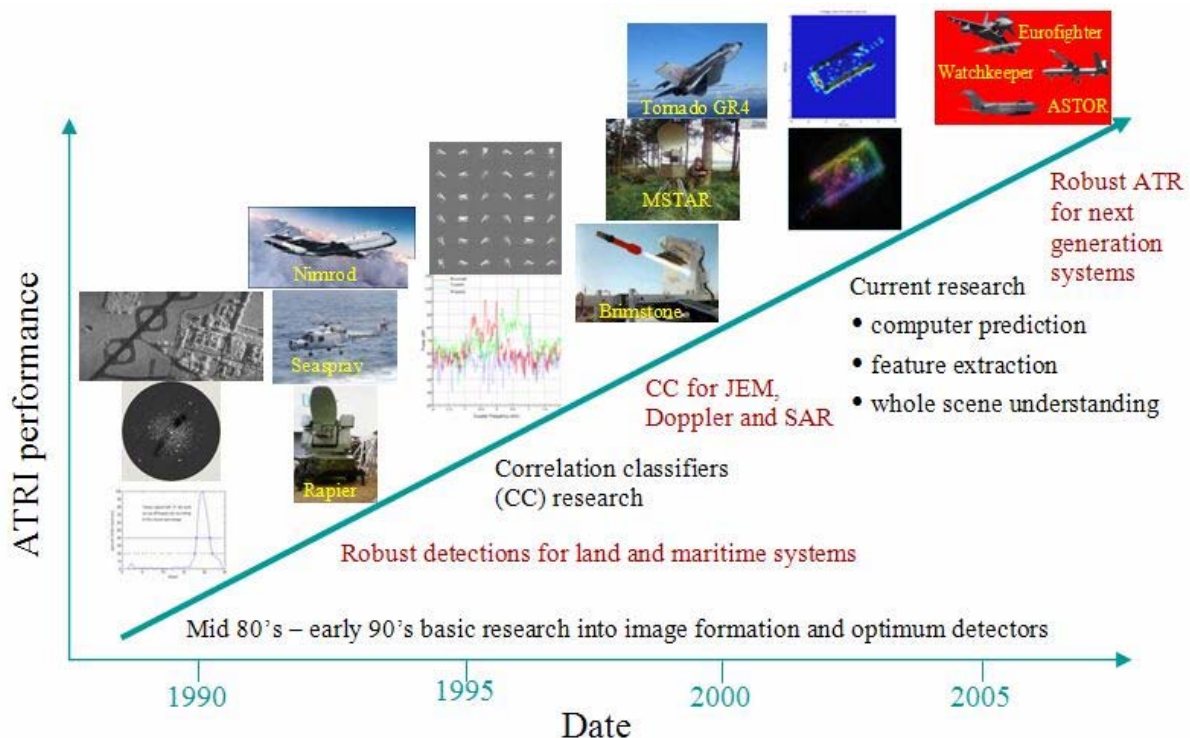


Figure 6.1: Radar ATR: Past and Future – a UK view (courtesy of QinetiQ).

6.1 Complex Clutter Environments

As has been previously mentioned, progress on NCI / ATR development has proceeded to some extent through constraining the problem. For example, for the ground target case, much development has taken place assuming that the targets are seen against an essentially featureless background. This assumption is valid for some operational scenarios and provides military capability. However, in progressing towards further level of complexity and utility, a key consideration is the situation when the targets are located in more complex clutter environments, e.g. heavy scrubland or urban areas.

Some progress can be made by the use of advanced radar modes. Figure 6.2 shows a single channel SAR image (left) of an area of scrubland containing some vehicles. These are difficult to see at first glance and would be challenging to an operator having to examine huge amounts of SAR imagery. However, the use of polarimetric decompositions enhances the contrast between the targets and the background making them more visible and hence more easily detected. Fundamentally, hard targets contain a lot of dihedral scattering events and so an odd / even bounce decomposition will highlight these against the mixed polarisation of the scrubland.

Moving towards the challenge of more complex clutter and the urban environment, approaches based on image segmentation to delineate and remove the unwanted clutter objects are one way forward. Alternatively, collateral information in the form of maps or other sensor information can provide a template of the buildings in the scene which can be masked out so that the search for targets is confined to only those regions where valid targets can be found. This is an ongoing challenge and leads on to the idea of using contextual information.




Figure 6.2: Use of polarimetry to enhance contrast of target signatures against scrubland background (courtesy of QinetiQ).

6.2 Contextual Information

It has been standard approach to ATR to extract an image chip containing the target and perform the classification purely on the basis of this chip. However, certainly in the ground target case, there is a

substantial amount of contextual information that an operator would use in making a target declaration which is ignored in this “tunnel-vision” approach. For example, the nearby lines of communication, the type of terrain and the military doctrine which govern deployment of targets will influence the probability of a target of a particular type being present.

To illustrate this, Figure 6.3 shows the process involved in achieving enhanced detection using such contextual information. There are three types of contextual information. Firstly it is known that some areas of terrain are easier for vehicles to move over than others. Secondly, military doctrine dictates that vehicles park up close to hedge boundaries and tree lines rather than in the open. Thirdly, military doctrine dictates that vehicles travel in groups rather than individually. On the basis of the first two of these, the thresholds of a detector are adjusted as illustrated in the centre image so that regions where it is more likely that targets are present are examined more closely (whiter). Once an initial set of detections has been obtained, an iterative process is followed in which the thresholds are adjusted in the locality of existing detections to accept more detections in these regions on the basis that there are likely to be more targets there. This simulated example showed that significant gains in detection performance could be achievable using such contextual information.

More generally, the challenge of contextual information is to bring together collateral information in the form of maps and other geospatial products, other sources of imagery and rules of military behaviour to provide not only an enhanced classification scheme but also an enhanced understanding of the entire scene from an operational perspective.




Figure 6.3: Illustration of enhanced detection based on use of contextual information (courtesy of QinetiQ).

6.3 Performance Assessment & Prediction

A crucial aspect of achieving an operational ATR capability is the process of validating performance against the required specification. This is an area with multiple challenges. Firstly, being able to articulate the required operational performance in terms that can be meaningfully assessed is not straightforward. Typically the metrics used in discussions of performance in the literature will not ideally match the key performance metrics for operations.




Figure 6.4: Conceptual illustration of output from ATR theory (courtesy of AFRL).

A second issue with performance assessment is that test data will not always be available. It may be that it is prohibitively expensive to gather sufficient data to populate the region of operational conditions. Alternatively, in the design stage it may be wished to understand the trade-off in radar parameter choices in terms of eventual ATR performance prior to building any demonstrator system. The resulting challenge is then to develop a theory of ATR that allows prediction of performance as a function of the variable that define the radar system, the target set, the operational scenario, the environmental conditions etc. This is illustrated in Figure 6.4 which shows conceptually how performance may be plotted against a number of variables as a result of such a theory to allow the interaction of the various defining factors to be explored. It is fair to say that this is a huge challenge although research is ongoing in some areas.

7.0 CONCLUSIONS

An overview of the fundamentals of ground target recognition using SAR has been given. It has been argued that radar ATR is not a single problem that can or cannot be solvable but that it is a continuum of problems of varying degrees of difficulty and complexity all of which provide useful military capability. This was essentially the viewpoint articulated by the NATO SET111 Task Group on ground target recognition at the conclusion of that activity as illustrated in Figure 7.1. Operational radar ATR systems already exist and the overarching challenge is to push forward the solution space to achieve successful operation in more difficult and complex circumstances. To achieve this, a number of key specific challenges have been identified and discussed and to a large extent it is these challenges which current researchers are tackling. Hopefully this lecture series will inspire those attending to contribute to the effort to solve these challenges and provide NATO with enhanced military capability.

<i>Scene :</i>	
Nb. targets	0 1 ... 10 ... unknown
Nb. variants	1 ... known ... estimated ... unknown
Clutter	uniform (field) ... dense & complex (urban, forest)
<i>Performance :</i>	
Level	detection ... classification (rough ... fine) ... identification ... technical analysis
Nb.Target/class	1 ... 10 ... all targets on scene
Time to	days ... hours ... minutes ... real time
Error (PFA)	10% ... 1% ... 5e-3%
Success (PD)	50% ... 80% ... 90% ... ~100%
Cost	« Free » ... radar cost ... multi radars &/ collections
<i>Reference :</i>	
Ref. Data	Photos ... expert ... 3D CAD (rough ... fine) ISAR/turntable SAR
Radar(s) mode	monolook ... multi-looks ... multi-aspects ...
Side orders	bandwidth ... channels (freq., polar., InSAR ...) ... multi sensors

Figure 7.1: Complexity of challenges (green easier, red harder) as assessed by SET-111.

8.0 REFERENCES

- [1] T D Ross, J J Bradley, L J Hudson, M P O'Connor "SAR ATR- So What's The Problem? An MSTAR Perspective" SPIE Algorithms for SAR VI vol 3721 pp662-672, 1999
- [2] E G Zelnio "Advanced decision-making systems in future avionics : automatic target recognition example" Aerospace conference Proc IEEE Vol 1 1998
- [3] M J Coulombe, T Horgan, J Waldman, J Neilson, S Carter and W Nixon "A 160 GHz polarimetric compact range for scale model RCS measurement" Proc Antenna measurements and techniques association pg239 October 1996
- [4] R H Giles, W T Kersey, A J Gatesman, M J Coulombe, M S McFarlin, R Finley and W E Nixon "A study of the X-band radar signature characteristics for main battle tanks in operational environments" Proc SPIE, vol 4718, pp336-343, April 2002
- [5] R H Giles, H J Nielson, D M Healy Jr, T Grayson, R Williams, and W E Nixon "Acquisition and analysis of X-band moving target signature data using a 160 GHz compact range" Proc SPIE, vol 4379, ATR XI, pp289-299 April 2001
- [6] S D Turner "RESPECT: Rapid electromagnetic scattering predictor for extremely complex targets" IEE Proceedings, Vol 137, part F, No 4, August 1990
- [7] D Andre, D Blacknell, and J Hare "Generation and validation of a simulated radar ground target database" SET080 NATO conference, 2004
- [8] R O Duda, P E Hart, D G Stork "Pattern Classification" Second Edition, John Wiley and Sons, 2001

- [9] T D Ross “*Confidence intervals for ATR performance metrics*”
- [10] J A Richards, W J Bow, and B K Bray “*An informative confidence metric for ATR*” Algorithms for SAR X, Proc SPIE Vol 5095 pp336-348, 2003
- [11] J A Ratches, C P Walters, R G Buser, and B D Guenther “*Aided and automatic target recognition based upon sensory inputs from image forming systems*” IEEE Trans. Pattern Analysis and Machine Intelligence Vol 19 No 9, 1997
- [12] R Hummel “*Model-based ATR using synthetic aperture radar*” IEEE Radar Conference, pp856-886, 2000
- [13] E R Keydel and S W Lee “*Signature prediction for model-based Automatic Target Recognition*” Algorithms for Synthetic Aperture Radar III, Proc SPIE 2757, pp306-317, 1996
- [14] J Wissinger, R Washburn, D Morgan, C Chong, N Friedland, A Nowicki, R Fung “*Search algorithms for model-based SAR ATR*” Algorithms for Synthetic Aperture Radar III, Proc SPIE 2757, pp279-293, 1996
- [15] E Keydel, W Williams, R Sieron, V Rajlich, and S Stanhope “*Reasoning support and uncertainty prediction in model-based vision SAR ATR*” Algorithms for Synthetic Aperture Radar VI, Proc SPIE 3721, pp620-631, 1999
- [16] G J Ettinger, G A Klanderma, W M Wells, and W E L Grimson “*A probabilistic optimisation approach to SAR feature matching*” Algorithms for Synthetic Aperture Radar III, Proc SPIE 2757, pp318-329, 1996
- [17] T D Ross and J C Mossing “*The MSTAR evaluation methodology*” Algorithms for SAR VI vol 3721 pp705-713, 1999
- [18] V Velten, T Ross, J Mossing, S Worrell, and M Bryant “*Standard SAR ATR evaluation experiments using the MSTAR public release data set*” Algorithms for Synthetic Aperture Radar V, Proc SPIE 3370,566-573 (1998)
- [19] J C Mossing and T D Ross “*An evaluation of SAR ATR algorithm performance sensitivity to MSTAR extended operating conditions*” Algorithms for Synthetic Aperture Radar V, Proc SPIE 3370, 554-565, 1998
- [20] L M Novak, G J Owirka, W S Brower, and A L Weaver “*The automatic target recognition system in SAIP*” The Lincoln Laboratory Journal, Volume 10, number 2, 1997
- [21] L M Novak, G J Owirka, and A L Weaver “*Automatic target recognition using enhanced resolution SAR data*” IEEE Trans AES Vol 35 No 1, 1999
- [22] G J Owirka, S M Verbout and L M Novak “*Template-based SAR ATR performance using different image enhancement techniques*” Algorithms for SAR VI, SPIE proceedings v3721, p302-319, 1999
- [23] L M Novak, G J Owirka, and W S Brower “*An efficient multi-target SAR ATR algorithm*” Asilomar 32, Vol 1, p3-13, 1998
- [24] L M Novak, G J Owirka, and W S Brower “*Performance of 10 and 20 target MSE classifier*” IEEE Trans AES Vol 36 No 4 2000

- [25] L M Novak, S D Halverson, G J Owirka, and M Hiatt “*Effects of polarisation and resolution on SAR ATR*” IEEE Trans AES, Vol 33, pp102-116, 1997
- [26] M Cetin, W C Karl, and D A Castanon “*Feature enhancement and ATR performance using nonquadratic optimization-based SAR imaging*” IEEE Trans AES Vol 39 Np4 2003
- [27] A Mahalanobis, A V Forman, N Day, M Bower, and R Cherry “*Multi-class SAR ATR using shift-invariant correlation filters*” Pattern Recognition Vol 27 No 4 pp619-626 1994
- [28] J A O’Sullivan, M D DeVore, V Kedia, M I Miller “*SAR ATR performance using a conditionally Gaussian model*” IEEE Trans AES Vol 37 No 1 2001
- [29] J H Yi, B Bhanu, and M Li “*Target indexing in SAR images using scattering centers and the Hausdorff distance*” Patt Rec Letts 17 pp1191-1198 1996
- [30] G Jones III and B Bhanu “*Quasi-invariants for recognition of articulated and non-standard objects in SAR images*” IEEE workshop on computer vision beyond the visible spectrum, Fort Collins, CO, pp88-97, 1999
- [31] B Bhanu, Y Lin, G Jones, J Peng “*Adaptive target recognition*” Machine Vision and Applications Vol 11(6) pp289-299, 2000
- [32] B Bhanu and G Jones III “*Object recognition results using MSTAR synthetic aperture radar data*” Proc IEEE workshop on computer vision beyond the visible spectrum, Hilton Head, SC, pp55-62, 2000
- [33] B Bhanu and G Jones III “*Recognizing target variants and articulations in synthetic aperture radar images*” Optical Engineering Vol 39 No 3 2000
- [34] G Jones III and B Bhanu “*Recognizing articulated objects in SAR images*” Pattern Recognition 34, pp469-485, 2001
- [35] G Jones III and B Bhanu “*Recognizing occluded objects in SAR images*” IEEE Trans AES Vol 37 No1 2001
- [36] J S Ahn and B Bhanu “*Model-based recognition of articulated objects*” Patt Rec Letts 23 pp1019-1029 2002
- [37] B Bhanu and G Jones III “*Increasing the discrimination of synthetic aperture radar recognition models*” Optical Engineering Vol 41 No 12 2002
- [38] L C Potter and R L Moses “*Attributed scattering centers for SAR ATR*” IEEE Trans. Image Processing 6 No 4 1997
- [39] H J Wolfson and I Rigoutsos “*Geometric hashing : an overview*” IEEE Computational science and engineering V4 pp10-21 1997
- [40] B Bhanu and Y Lin “*Stochastic models for recognition of occluded targets*” Pattern Recognition 36 pp2855-2873, 2003
- [41] C Nilubol, R M Mersereau, and M J T Smith “*A SAR target classifier using radon transforms and hidden markov models*” Digital Signal Processing 12, pp274-283, 2002

- [42] Q Zhao, J C Principe, V L Brennan, D Xu, and Z Wang “*Synthetic aperture radar automatic target recognition with three strategies of learning and representation*” Opt Eng 39(5) 1230-1244, 2000
- [43] Q Zhao and J C Principe “*Support vector machines for synthetic aperture radar automatic target recognition*” IEEE AES 37(2), 2001
- [44] M J Gerry, L C Potter, I J Gupta “*A parametric model for synthetic aperture radar measurements*” IEEE Trans. Antennas and Propagation Vol 47 No 7 July 1999
- [45] H Chiang, R L Moses, and L C Potter “*Model-based classification of radar images*” IEEE Trans Information Theory Vol 46 No 5 2000
- [46] H Chiang, R L Moses, and L C Potter “*Model-based Bayesian feature matching with application to synthetic aperture radar target recognition*” Pattern Recognition 34 pp1539-1553 2001
- [47] B Ulug, S C Ahalt, and R A Mitchell “*Efficient ATR using compression*” IEEE Trans AES Vol 33 No 4 1997
- [48] C Daniel, A Mahalanobis, and R Goodman “*Object recognition in subband transform-compressed images by use of correlation filters*” Applied Optics, Vol 32 No 32 2003
- [49] S Suvorova and J Schroeder “*Automated target recognition using the Karhunen-Loeve transform with invariance*” DSP 12 pp295-306, 2002
- [50] D Casasent and R Shenoy “*Feature space trajectory for distorted-object classification and pose estimation in synthetic aperture radar*” Opt Eng 36(10) pp2719-2728 1997
- [51] R G Congalton “*A review of assessing the accuracy of classifications of remotely sensed data*” Remote Sens. Environ. 37 pp35-46 1991
- [52] S V Stehman “*Selecting and interpreting measures of thematic classification accuracy*” Remote Sens. Environ. 62 pp77-89 1997
- [53] M Z Brown “*Analysis of multiple-view Bayesian classification for SAR ATR*” Algorithms for Synthetic Aperture Radar X Proc SPIE 5095 pp 265-274 2003
- [54] E Krogager “*New decomposition of the radar scattering matrix*” Electronics Letters, Vol 26, pp1525-1527, 1990
- [55] L M Novak, G J Orwirka, and C M Netishen “*Performance of a high-resolution polarimetric SAR automatic target recognition system*” The Lincoln Laboratory Journal, Vol 6, pp11-24, 1993
- [56] R D Chaney, M C Burl, L M Novak “*On the performance of polarimetric target detection algorithms*” IEEE International radar conference, pp520-525, 1990
- [57] L M Novak, M C Burl, W W Irving “*Optimal polarimetric processing for enhanced target detection*” IEEE Trans AES, Vol 29, pp234-244,1993
- [58] F Sadjadi “*Improved target classification using optimum polarimetric SAR signatures*” IEEE Trans AES, Vol 38, pp38-49, 2002
- [59] P Pudil, F J Ferri, J Novovicova, and J Kittler “*Floating search methods for feature selection with nonmonotonic criterion functions*” Proc of the 12th International Conference on IAPR, Vol 2, IEEE, 1994

- [60] A Jain and D Zongker “*Feature selection: evaluation, application and small sample performance*” IEEE Trans PAMI, Vol 19, 1997
- [61] T H Pink and U Ramanathan “*Intelligent selection of features for optimal feature based classification*” Proc IEEE IGARSS, pp3012-3014, 2000
- [62] K Torkkola “*Feature extraction by non-parametric mutual information maximization*” Journal of Machine Learning Research 3, pp1415-1438, 2003
- [63] D Blacknell “*Statistical target behaviour in SAR images*” IEE Proc Radar, Sonar, and Navigation, 147, issue 3, pp143-148, 2000
- [64] D S Rosario “*Managing within-class target variability in SAR imagery with a target decomposition model*” IEEE Int conf image processing Vol 3, pp935-938, 1996
- [65] R Meth and R Chellappa “*Stability and sensitivity of topographic features for SAR target characterization*” IEEE Int conf image processing, Vol 3, pp464-467, 1997
- [66] R Meth and R Chellappa “*Target indexing in synthetic aperture radar imagery using topographic features*” IEEE ICASSP-96, Vol 4, pp2152-2155, 1996
- [67] D E Dudgeon “*ATR performance modelling and estimation*” Digital Signal Processing 10, pp269-285, 2000
- [68] K Fukunaga and D L Kessell “*Nonparametric Bayes error estimation using unclassified samples*” IEEE Trans Information Theory IT-19, No4, 1973
- [69] K Fukunaga and R R Hayes “*Estimation of classifier performance*” IEEE Trans Pattern Analysis and Machine Intelligence Vol 11, No 10, 1989
- [70] A J Viterbi, and J K Omura “*Principles of digital communication and coding*” McGraw-Hill 1979
- [71] E Zelnio and F Garber “*A characterization of ATR performance evaluation*” Signal Processing, Sensor Fusion, and Target Recognition V, pp2-13, Proc SPIE 2755, 1996
- [72] F Garber and E Zelnio “*On some simple estimates of ATR performance and initial comparisons for a small data set*” SPIE Algorithms for Synthetic Aperture Radar IV, Proc SPIE 3070, pp150-161, 1997
- [73] F Kanaya and K Nakagawa “*On the practical implication of mutual information for statistical decision making*” IEEE Trans Information Theory, Vol 37, No 4, pp1151-1156, 1991
- [74] S D Briles “*Information-theoretic performance bounding of Bayesian identifiers*” Automatic Object Recognition III, Proc SPIE Vol 1960, 1993
- [75] S D Morgera and M R Soleymani “*A structural review of pattern recognition from the point of view of rate-distortion theory*” in Pattern Recognition and Artificial Intelligence Ed E S Gelsema and L N Kanal Elsevier Science Publishers, 1988
- [76] L L Horowitz and G F Brendel “*Fundamental SAR ATR performance predictions for design tradeoffs*” SPIE Algorithms for Synthetic Aperture Radar IV, Vol 3070, 1997
- [77] Z Knowles and D Parker “*A Monte Carlo simulation based approach to a priori performance prediction for target detection and recognition in cluttered synthetic aperture radar imagery*”

- [78] A M Horne “*An information theory for the prediction of SAR target classification performance*” Algorithms for Synthetic Aperture Radar VIII, Proc SPIE Vol 4382, 2001
- [79] A Jain, P Moulin, M I Miller, K Ramachandran “*Information theoretic bounds on target recognition performance based on degraded image data*” IEEE Trans Pattern Analysis and Machine Intelligence, Vol 24, No 9, 2002
- [80] M Boshra and B Bhanu “*Predicting the performance of object recognition*” IEEE Trans Pattern Analysis and Machine Intelligence Vol 22 No 9, 2000
- [81] M Boshra and B Bhanu “*Predicting an upper bound on SAR ATR performance*” IEEE Trans AES Vol 37 No 3 2001
- [82] A C Williams and B Clark “*Evaluation of SAR ATR*” SPIE signal processing, sensor fusion, and target recognition Proc SPIE 2755, pp36-45, 1996
- [83] M I Miller, U Grenander, J A O’Sullivan, and D L Snyder *Automatic target recognition organised via jump-diffusion algorithms*” IEEE Trans Image Processing Vol 6 No 1, 1997
- [84] U Grenander, M I Miller, and A Srivastava “*Hilbert-Schmidt lower bounds for estimators on matrix Lie groups for ATR*” IEEE Trans Patt Anal and Mach Int Vol 20, No 8, 1998
- [85] U Grenander, A Srivastava, and M I Miller “*Asymptotic performance analysis of Bayesian target recognition*” IEEE Trans Information Theory Vol 46, No 4, 2000
- [86] M L Cooper and M I Miller “*Information measures for object recognition accommodating signature variability*” IEEE Trans Information Theory Vol 46 No 5, 2000
- [87] M D DeVore “*Recognition performance from synthetic aperture radar imagery subject to resource constraints*” Phd Thesis, Washington University, 2001, www.people.virginia.edu/~md9c/
- [88] M D Devore, J A O’Sullivan, R D Chamberlain, M A Franklin “*Relationships between computational system performance and recognition system performance*” Automatic Target Recognition XI, Proc SPIE 4379, pp355-363, 2001
- [89] M D DeVore and J A O’Sullivan “*Performance complexity study of several approaches to automatic target recognition from SAR images*” IEEE Trans AES Vol 38, No 2, pp632-648, 2002
- [90] M D DeVore, R Chamberlain, G Engel, J O’Sullivan, and M Franklin “*Tradeoffs between quality of results and resource consumption in a recognition system*” Proc IEEE Conference on application specific systems, architectures and processors, pp391-402, 2002

