SAR Machine – Aided Recognition Toolbox

D M Carrington

QinetiQ Malvern[†] and University College London

Abstract: A number of algorithms and techniques have been developed in QinetiQ for use in Automatic Target Recognition (ATR) using Synthetic Aperture Radar (SAR) imagery. This paper details a software system that has been designed to evaluate and demonstrate the capabilities of these tools. This system has been named the SAR Machine–Aided Recognition Toolbox (SMARTBOX). This paper gives a brief overview of the algorithms used to process the SAR imagery and some example results from experiments on real SAR imagery.

1. Introduction.

Synthetic Aperture Radar (SAR) provides the ability to obtain high – resolution radar imagery in all-weather conditions from long stand-off ranges, making it highly suited for a range of military applications such as surveillance. Analysing the imagery will present a significant workload on the image analysts, because of the volume of imagery that will be gathered in a typical mission and the need to analyse the imagery in a timely manner. This increases the risk of missing or incorrectly identifying targets of interest, usually military vehicles, that may be present in the imagery. To ensure timely and accurate analysis of the imagery, the development of a set of automatic analysis tools has been recognised as being crucial.

This paper describes the creation of an Aided Target Recognition (ATR) tool, known as SMARTBOX. This software was designed to pull together a number of algorithms to form a complete ATR processing chain. As well as being used for a number of experiments to evaluate the effectiveness of ATR, its graphical interface makes it a useful tool for demonstrating the individual elements in the ATR process.

2. The ATR Processing Chain.

The processing chain has been divided into three components:

- Pre - screening
- Target Delineation
- Target Classification

In addition, there is a training module that provides reference data for the classification component.

The processing algorithms either already existed at the start of the project or were developed in parallel with SMARTBOX. The primary function of the SMARTBOX project was to provide a framework to link these modules together in a way that a user can easily interpret the results.

There follows a brief definition of each module and a description of how it is integrated into the SMARTBOX structure.

2.1. Target Pre-Screening

The aim of the pre-screening stage is to detect pixels that belong to man-made objects such as vehicles. The resulting detections are then passed on to the delineation and classification stages. The pre-screening stage is broken down into three modules:

- Detection – identification of anomalously bright pixels in the SAR image that probably belong to man-made objects;
- Clustering – groups adjacent detections together to form regions that may represent a single target;

- Discrimination – Analyses the region represented by a cluster and either accepts it as a man-made objects or rejects it.

2.1.1 Detection

The aim of the detection module is to mark pixels that are brighter than their surroundings, indicating that they belong to a man-made metallic object. Detection is achieved by moving a sampling window, as shown in figure 1, over the image. The characteristics of the background clutter Probability Density Function (PDF) are estimated from the background ring. The likelihood of the centre test pixel arising from the PDF is calculated. The statistics of the background are used to calculate a threshold for the test pixel. The mask prevents other pixels belonging to the target from being used for the background calculation. If the test pixel exceeds this threshold it is flagged as a detection. This process is repeated for each pixel in the image.

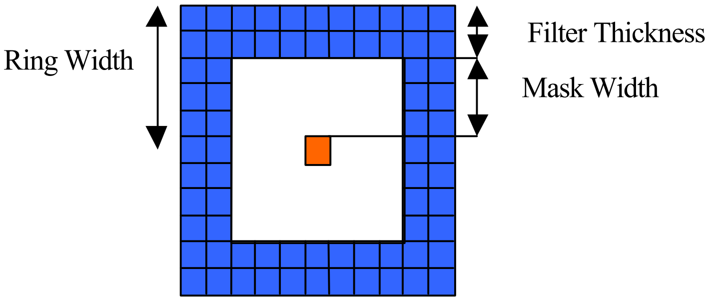

Figure 1: Window used for pixel detection

The threshold can be chosen in a way to keep the false alarm rate constant, since the background clutter can be robustly characterised. SMARTBOX allows the user to select from a variety of CFAR detectors and to change all relevant parameters. The map of detections generated by this module is then passed on to the clustering module.

2.1.2 Clustering

Clustering operates on the output of the detection algorithm. It aims to identify groups of pixels that belong to the same target vehicle.

The algorithm searches from left to right and top to bottom through the map of detections. When a detection is found, it is marked as belonging to a new cluster. The algorithm then searches for other detections within a rectangle centred on the first detection. The dimensions of the rectangle can be specified by the user. If any other detections are found within the rectangle, these are marked as belonging to the same cluster as the first pixel. The rectangle is then applied to the newly identified pixel, and the process continues until no more new pixels are found. The algorithm then starts to search for a new 'seed' pixel for the next cluster. When clustering is complete, a rectangular section is extracted from the original image, centred on each cluster. These sections, known as target chips, are then passed on to the discrimination module.

2.1.3 Discrimination

Discrimination takes the target chips generated by the clustering stage and performs a rough classification of the chip as either man-made or not man-made. The aim is to remove any detections due to bright natural objects such as trees and hedges. Discrimination uses physical aspects of the target, such as size, shape and the strength of the radar return. Constraints on these factors are set by the user, and the algorithm rejects any target which does not meet the criteria for a man-made object. Targets that do meet the criteria are passed on to the delineation stage.

2.2 Delineation

As the name suggests, the purpose of the delineation stage is to delineate the boundary between the target and background in a target chip. Once the target region has been delineated, the target shadow can also be found, which is useful for performing classification later. Active contour methods (often known as ‘snakes’) have been found to be a good solution to the delineation problem [1].

Traditional active contour methods have used pixels on the contour boundary and did not work well on SAR imagery due to the presence of speckle noise. Region – based active contour methods consider the noise statistics of regions inside and outside the active contour and have been used successfully on SAR imagery.

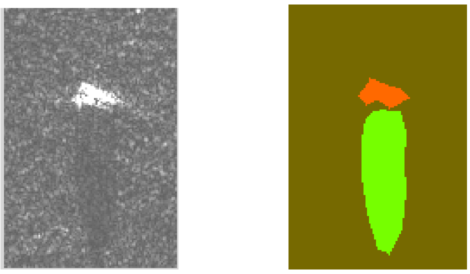

Figure 2: Target chip (left), results of target and shadow delineation (right)

The region containing the target and shadow are used to calculate features used by the Classification module.

2.3 Classification

The classification module uses a feature-based maximum likelihood classifier. This compares values of selected features (e.g. as area, variance and perimeter) with those from known data from a training set. Two possible modes of operation are available with classification: the first is fully automatic, where the classification module decides on a single vehicle class that is the best match for the given feature set. The alternative method is to present the user with the best 2 or 3 matches and allow the user to decide which class the target belongs to. In both cases, the standard maximum likelihood classifier has been modified in SMARTBOX so that targets that do not lie within the 95% confidence interval of any class are classified as unrecognisable.

3. Experimental Results and discussion

The complete processing chain was tested using 9 target vehicles, with 110 images of each vehicle. The targets were a mixture of military vehicles, civilian vehicles and inflatable decoys. The results of the SMARTBOX classification results were used to populate classification matrices showing correct and incorrect classifications.

If the fully automatic classification approach is used, classification of a target is only considered successful if the top declared classification is correct. For this 9-class problem, a correct classification rate of 52% was achieved.

Using the Aided Recognition approach, the best 2 or 3 matches are presented to the user, who can then choose which is the closest match to the target. For the purposes of this test, the classification was considered correct if the true class was represented in the 2 or 3 classes presented. If 2 classes were presented, the correct classification rate was found to be 70%. If 3 classes were presented, this rose to 82%.

The Aided Recognition test did not take into account the likelihood of the user selecting the correct class from those presented. This is a significant factor in the performance of the system, since a recent study has shown that on SAR imagery displayed as in figure 3, the human classification accuracy was 77%. Work is in progress to investigate ways of improving this performance, in particular by altering the proportion of the dynamic range of the display covering the target rather than the background. By assigning more of the dynamic range of greyscales to brighter pixels, it is hoped that more detail on the target will be visible, avoiding the saturation effects visible in figure 3. The appearance of speckle in the background will also be reduced. Initial experiments on these techniques have shown that they can increase human classification accuracy to 92%.

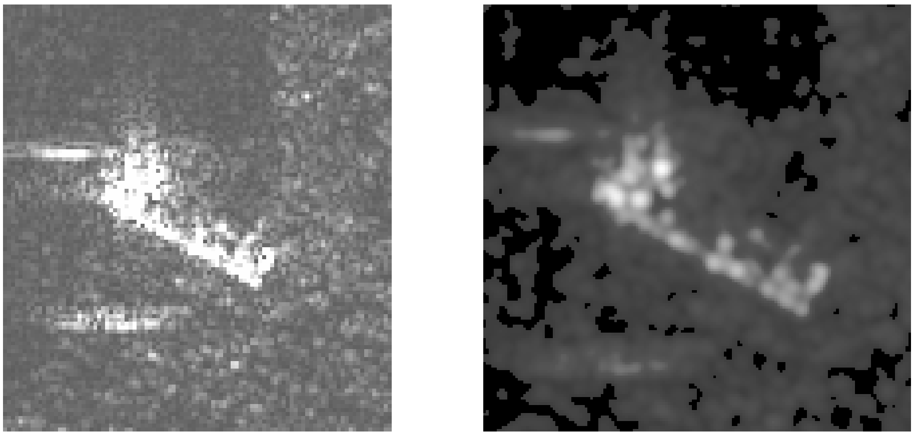

Figure 3: SAR image of a military vehicle, unmodified image (left) and modified dynamic range (right). Note the saturation of the bright regions of the target in the unmodified image and the removal of background noise from the modified image. Additional detail is also visible on the modified image, particularly towards the top left of the vehicle.

4. Conclusions

This work has produced a software tool that allows evaluation of ATR algorithms on real SAR data. This software has uses both as a demonstration tool and for gathering experimental data on the effectiveness of the algorithms used in the ATR processing chain. The graphical user interface makes it highly suited for demonstrating the utility of ATR for the pre-screening of SAR images for vehicle targets. The interactive nature of the software makes the task of customising the processing chain easy, allowing the user to modify the operation of the software to suit a particular set of data or search criteria.

The performance of the ATR algorithms has been shown to be excellent under good conditions, although it does degrade when targets are obscured by trees or are in urban terrain. This can be improved by using aided rather that automatic recognition, in which the final classification of the target is left to the operator, choosing from a list of 2 or 3 classes generated by the software. Work is in progress to evaluate and improve the performance of a human operator in identifying vehicles in SAR images.

Acknowledgements

This work was partially funded by the MoD under Surface Surveillance Technologies.

References.

[1] Pedlar DN and Blacknell D 'Target Delineation and Classification using a region-based active contour and a support vector machine classifier on SAR imagery' EUSAR Conference Proceedings, Ulm, Germany, May 2004

† Address: QinetiQ, Malvern Technology Centre, Worcs. WR14 3PS UK. Email: dmcarrington@qinetiq.com