Robust vehicle Detection in Synthetic Aperture Radar Imagery

By Hill, R. D.[(1)], Moate, C. P.[(2)], Leonard, T. P.[(2)]

(1) Igence Radar Ltd, Wyche Innovation Centre, Walwyn Road, Upper Colwall, Malvern, Worcs WR13 6PL

(2) QinetiQ, Malvern Technology Centre, St Andrews Road, Malvern WR14 3PS

Abstract

This paper describes the development and analysis of a pre-screening algorithm that has been demonstrated to provide robust detection performance for military vehicles deployed in close proximity and in complex clutter. The key elements of the processing chain are a Constant False Alarm Rate (CFAR) detector based on the K-distribution and a novel edge-based vehicle detector. The algorithm is demonstrated using simulations and real Synthetic Aperture Radar (SAR) imagery which show that it is capable of detecting closely spaced vehicles. The resulting scheme is computationally efficient and suitable for the first stage of an Automatic Target Recognition system designed to process large quantities of image data.

1. Introduction

There is a military requirement to detect and recognise relocatable and mobile surface targets at range and in all weathers. Tactical and strategic Synthetic Aperture Radar (SAR) systems can provide imagery of sufficient quality to potentially allow the identification of military targets. Robust Automatic Target Recognition (ATR) techniques are required to support both tactical and strategic applications. In the case of an airborne targeting platform, automation is required to minimise operator workload, whereas in the strategic application, automation is necessary to combat data deluge. ATR is the process by which targets in an image are automatically detected and classified using computer algorithms. Generally, ATR schemes implement a processing pipeline that consists of pre-screening followed by a more computationally expensive classification stage, which encompasses target recognition and identification. The pre-screening stage aims to flag up potential man-made targets in the image and as such implements target detection.

This paper describes the development and analysis of a pre-screening algorithm that has been demonstrated to provide robust detection performance for military vehicles deployed in close proximity and in complex clutter. The pre-screening algorithm consists of three main stages, a single pixel detection stage, detection filtering and clustering and an edge based vehicle detector. The single pixel detection stage is based on Constant False Alarm Rate (CFAR) processing where the detection threshold is calculated using the assumption that the clutter follows a K-distribution. Parameter estimation in standard implementations of K-CFAR is based on average image intensities within a CFAR mask surrounding a test pixel, this approach fails when vehicles are deployed in close proximity or in urban areas and a significant fraction of the mask pixels are contaminated by foreground clutter. This problem is tackled by using robust statistics to estimate the parameters of the K-distribution and we show how this modifies the calculation of the detection threshold. The performance of the new single pixel detector is compared with the standard implementation by applying both detectors to SAR images containing vehicles in close proximity. The results show a significant improvement in the probability of detection when both detectors are operated at the same false alarm rate. A novel edge-based vehicle detector is described which uses an exhaustive search to fit vehicle sized rectangles to the output of the single-pixel detector. Unlike many snake-based algorithms this approach allows a local, deterministic and exhaustive search to be carried out without recourse to random search strategies such as simulated annealing. This computationally efficient approach provides the capability to pre-screen large SAR images in realistic time-scales.

2. Algorithm Overview

An overview of the vehicle detection algorithm is provided in FIGURE 1; the key elements of the processing chain are a single pixel detection stage which consists of a CFAR algorithm based on the assumption of K-distributed clutter, an m from n filter designed to remove isolated single pixel detections and a vehicle detector that exploits edge information to find the boundary of a target return. The target dimensions are then compared with input size criteria to decide whether to reject it or flag it as a detection.

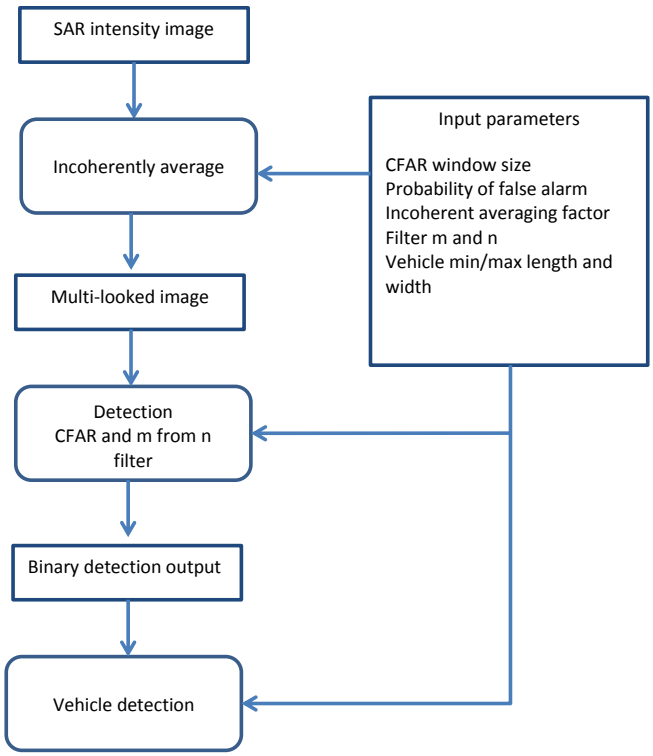

FIGURE 1. Vehicle detection algorithm

The key stages of the vehicle detection algorithm are described in more detail in the following sections.

3. K Distribution Clutter Model

The target detection scheme is based on the assumption that the SAR image clutter intensity may be described by the K distribution. The K distribution model for sea clutter was introduced by Jakeman & Pusey (1976), who suggested analogies between scattering at microwave and optical wavelengths to motivate its use. Ward (1981) then showed how the K-distribution emerged from the detailed analysis of high-resolution sea clutter. This work demonstrated the usefulness of the compound representation of the clutter process which provided a method for the systematic analysis of the effect of thermal noise and spatial and temporal correlations and their impact on radar performance. Oliver (1998) demonstrated the applicability of the K-distribution to synthetic aperture radar imagery where it has provided the basis for a wide range of image exploitation techniques including detection and segmentation.

The detection threshold is calculated assuming that the clutter is distributed according to a multi-looked K-distribution. The compound form of the K distribution (for a single look with no thermal noise) of SAR image intensities z is obtained from a speckle process with a gamma distributed local power x:

$$P_{speckle}(z|x) = \frac{1}{x}\exp\left(-\frac{z}{x}\right)$$
$$P(x) = \frac{b^{\nu}}{\Gamma(\nu)}x^{\nu-1}e^{-bx} \tag{3.1}$$

The K-distribution is then obtained by integrating the speckle process over the local power to give:

$$P(z) = \frac{b^{\nu}}{\Gamma(\nu)}\int_0^{\infty} x^{\nu-1}e^{-bx}\,\frac{1}{x}e^{-\frac{z}{x}}dx$$
$$= \frac{2b^{\left(\frac{\nu+1}{2}\right)}z^{\left(\frac{\nu-1}{2}\right)}}{\Gamma(\nu)}K_{\nu-1}\left(2\sqrt{bz}\right) \tag{3.2}$$

where an integral representation of the K-Bessel function has been identified in the second equation of (3.2). The derivation of properties of the distribution, such as the moments, invariably relies on the compound form.

The multi-looked K-distribution is obtained by incoherently averaging N clutter pixels with the same local power, which is assumed to be slowly varying across the image. The resulting multi-looked K distribution is given by Oliver and Quegan (1998) as

$$\begin{aligned} P_N(z) &= \frac{b^\nu}{\Gamma(\nu)}\int_0^\infty x^{\nu-1}e^{-bx}\,\frac{N^N z^{N-1}}{x^N\Gamma(N)}e^{-\frac{Nz}{x}}dx \\ &= \frac{2(bNz)^{\left(\frac{\nu+N}{2}\right)}}{\Gamma(\nu)\Gamma(N)z}K_{N-\nu}\left(2\sqrt{bNz}\right) \end{aligned} \tag{3.3}$$

The moments of the multi-looked K-distribution can be calculated straightforwardly from the compound form to obtain:

$$\left\langle z^r \right\rangle = \frac{\Gamma(N+r)\Gamma(\nu+r)}{\Gamma(N)\Gamma(\nu)N^r b^r} \tag{3.4}$$

The mean of the distribution is therefore given by $\mu = \frac{\nu}{b}$. In the case where the distribution parameters are known *a priori*, a detection threshold can be calculated from the required probability of false alarm using

$$pfa(T) = \frac{b^\nu}{\Gamma(\nu)\Gamma(N)}\int_0^\infty \Gamma\left(N, \frac{NT}{x}\right)x^{\nu-1}e^{-bx} \tag{3.5}$$

where $\Gamma(N,z)$ is the incomplete Gamma function.

In practice, the distribution parameters are not known and must be estimated from the SAR image. In previous work, the background level was estimated from the mean intensity within the CFAR mask ($\hat{\mu}_B$) which has well known limitations in the presence of outliers. A range of approximate estimators for the K-distribution shape parameter (ν) are described by Oliver (1994), Blacknell (1994) and Blacknell and Tough (2001). A typical example is the U estimator given by

$$U = \langle \ln z \rangle - \ln\langle z \rangle \tag{3.6}$$

The shape parameter estimate is then obtained from the solution of

$$U = \psi(\nu) + \psi(N) - \ln\nu - \ln N \tag{3.7}$$

where ψ denotes the digamma function. The threshold for unit mean, $\hat{T}_1$, is then obtained from equation (3.5) by setting b equal to the shape parameter estimate. This is implemented using a pre-calculated look-up table to avoid the necessity for repeated numerical solutions of the equation. The detection threshold is then given by

$$T_D = \hat{\mu}_B \hat{T}_1 \tag{3.8}$$

In this work the estimation scheme has been modified to avoid bias in situations where vehicles are in close proximity to each other or to buildings. The estimate of the mean background intensity is now calculated from the median which provides robustness to outliers and allows a larger CFAR estimation window to be used. The sample median is scaled by a factor, $\alpha(\hat{\nu})$, to provide an estimate of the mean image intensity. The multiplier $\alpha(\hat{\nu})$ is calculated as the reciprocal of the median of a multi-looked K distribution with unit mean. The shape parameter estimation scheme has also been modified to avoid averaging data within the estimation window. The new scheme consists of calculating the ratio of the 70th and 50th percentiles of the intensity data within the CFAR window and interpolating a look-up table to determine the shape parameter. The percentile ratio is calculated numerically using equation (3.5) with the results shown in FIGURE 2. The resulting estimates are spatially averaged to reduce the CFAR loss associated with parameter estimation.

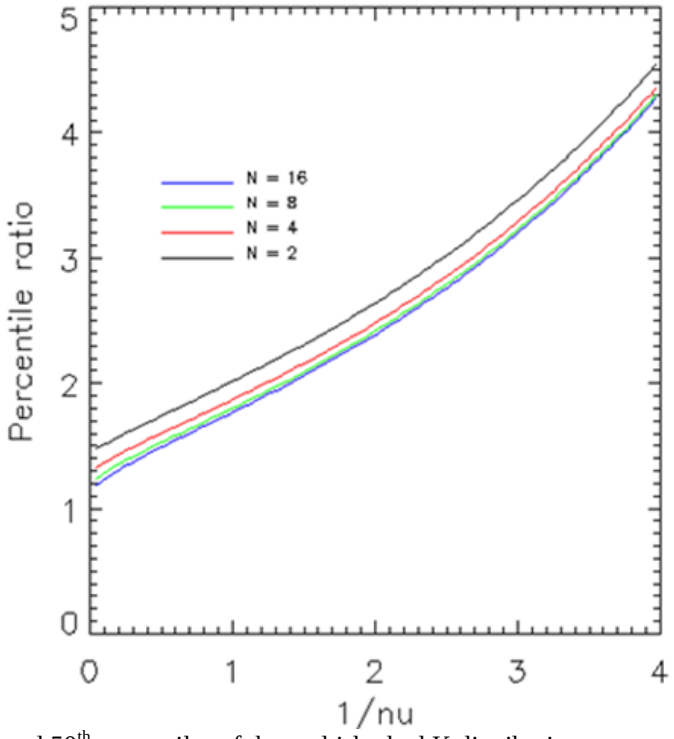

FIGURE 2. Ratio of the 70th and 50th percentiles of the multi-looked K distribution

4. Detection Filtering

The K-CFAR detection stage is followed by a simple filter designed to remove the many single isolated detections associated with the image speckle. An m out of n filter is applied in a window centred on each candidate detection; the size of the window is set to the length of a typical target. The n parameter is defined by assuming that at least half of the target lies in the window and that half of these pixels are detected. The probability of false alarm following the filter can be related to the single pixel probability of false alarm at the filter input through

$$P_{FA} = \sum_{i=m}^{i=n} \frac{n!}{(n-i)!i!} p^i (1-p)^{n-i} \tag{4.1}$$

This relationship can be used to set the probability of false alarm of the detector to optimise the performance of the vehicle detector.

5. Detection Simulations

FIGURE 3 compares the detection scheme based on the original estimators of the distribution parameters (using the mean to estimate the background and the U statistic to estimate the shape parameter) with the new scheme which uses a median background estimation and percentiles to estimate the shape parameter. The left hand image shows simulated data where targets are modelled as K-distributed clutter with a greater mean level than the background. The centre image shows the results obtained using the original detector and the right hand image those from the new detector. In the original detector, the effect of closely spaced targets is to contaminate the background estimation window and thereby raise the detection threshold and reduce the probability of detection. The results obtained using the percentile based detector suggest that this effect has been reduced and the detection of closely spaced targets is significantly improved.

FIGURE 3. Comparison of alternative parameter estimation schemes. Left: input image, Centre: Average, Right: Median

A more quantitative comparison is shown in FIGURE 4 which shows the fraction of target pixels detected plotted against the measured probability of false alarm in a background region for the scenario shown in FIGURE 3. The results for the two detection schemes are compared with those achieved using an ideal fixed threshold based on the known clutter parameters used in the simulation. The percentile based scheme is found to be a significant improvement on the original detector, however the detection performance is less than that which can be achieved with the ideal fixed threshold owing to the CFAR loss associated with parameter estimation.

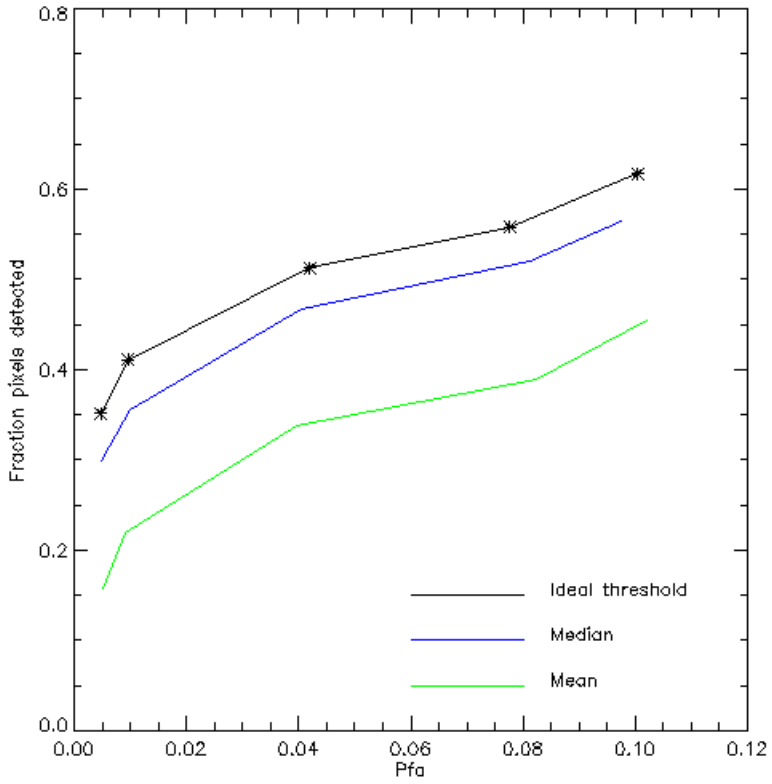

FIGURE 4. Detection performance compared with the constant *a priori* threshold

6. Edge Based Vehicle Detection

When the detection algorithm is applied to SAR imagery it produces a large number of detections from which objects of interest must be identified. The next stage of processing is delineation which involves finding the boundary of the target return in the image. This can be a crucial stage in the processing, as it enables the feature based detection of objects of interest by a down-stream Automatic Target Recognition algorithm.

Chesnaud *et al* (1999) and Pedlar (2004) showed that an area-based approach to active contours (or snakes) can find target boundaries in SAR imagery, where a single target is well separated from other targets. This technique computes statistical measures for pixels inside and outside the contour, evolving the contour to maximise a probability measure. Moate and Denton (2006) developed this approach for the more difficult case of multiple vehicles in close proximity. This problem is difficult because, with multiple vehicles, there is the potential for a contour delineating one vehicle to grab pixels from another vehicle. FIGURE 5 illustrates this problem. The figure shows how two contours are initially seeded (manually in this case) for the two targets. These are then evolved by randomly varying their properties and accepting the change if it improves the probabilistic fit criterion. Due to the random starting orientation and evolution of the contours, sometimes they correctly delineate the vehicles and sometimes they only find a local minimum. Once a contour has grabbed pixels from another vehicle, evolution to the global minimum in which both targets are correctly delineated is then only possible by coordinated movement of the contours.

An underlying problem of this approach is that the trial movements of different contours are effectively coupled, much increasing the size of the search space over which the minimisation occurs during fitting. Our previous work used two search heuristics specifically to overcome this problem – simulated annealing (allowing the contours to evolve to positions which are a worse fit, but in a limited manner) and multiple hypothesis delineation (starting the delineation in multiple different states to cover more of the minimisation space).

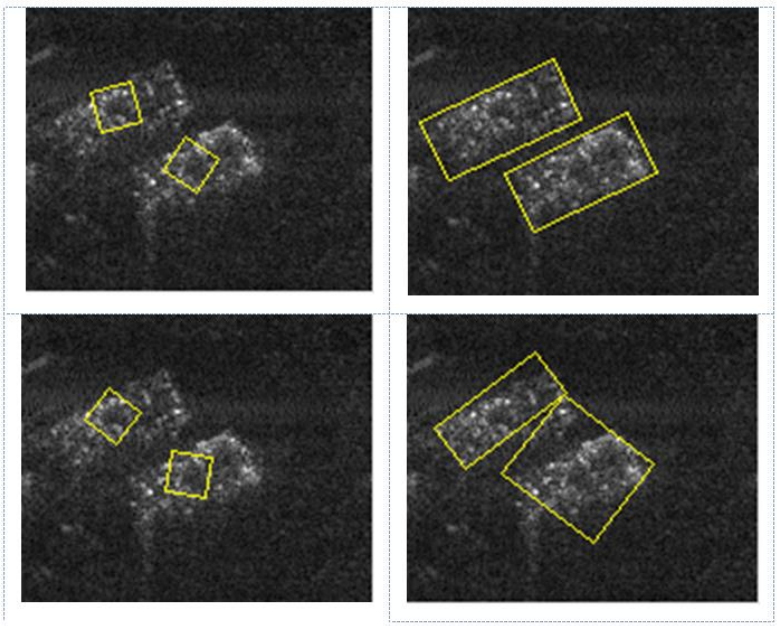

FIGURE 5. The grab problem: Left hand side: initial contour positions. Right hand side: final contour positions. Top: successful delineation fitting targets. Bottom: Failed delineation stuck in local minimum.

The current work has improved on this previous approach by removing the need for all the search heuristics (random evolution, simulated annealing and multiple hypothesis delineation) which each require parameters that can need tuning. It has also removed the sensitivity of the previous approach to the seeding positions used for the contours. A detailed comparison of the two approaches is provided in Moate & Leonard (2015) which provides a wider range of examples and timing benchmarks. The new technique is still based around the active contour approach. The key change from earlier work is that this new approach searches for strong edges in the image. The approaches are compared and contrasted in FIGURE 6.

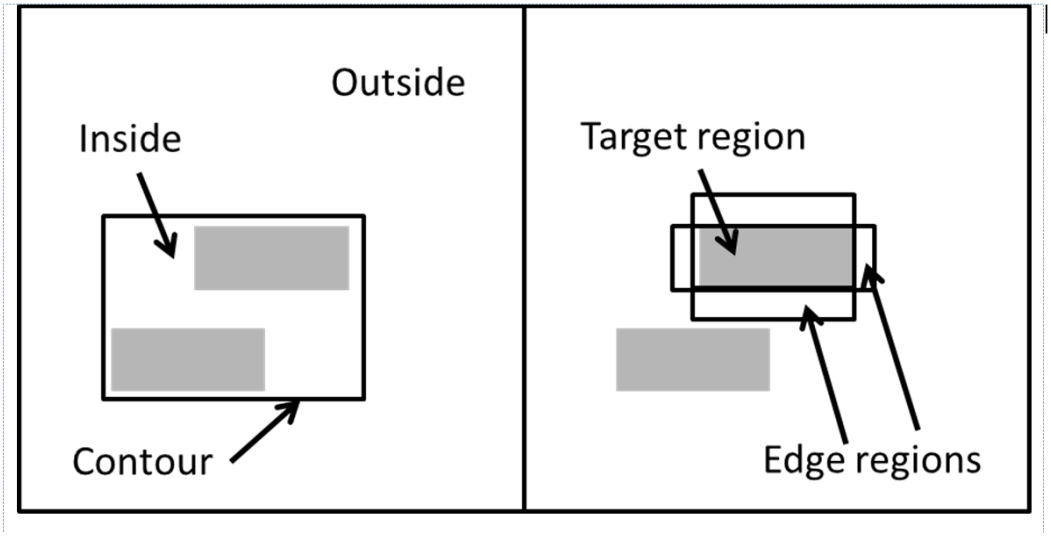

FIGURE 6. Left: the area-based approach to active contour delineation. Right: the edge-based approach

In the left hand image we show how the area-based contour divides the image into regions – inside and outside – which are assumed to have different mean intensities. Where there are additional bright returns in the background (such as fences, walls, etc, or a second vehicle as shown) there is a strong incentive for the area-based contour to grow larger and encompass these bright returns, separating them from the darker background. The right hand image shows our new approach based on edge detection. The central rectangle (target region) aims to delineate the vehicle while the surrounding 4 rectangles (edge regions) measure the local background return. An edge metric is defined as:

$$p = \sqrt[4]{\prod_{i=1}^{4} r_i} \tag{5.1}$$

where the edge ratio for each of the four edges is calculated to be:

$$r_i = \begin{cases} 1 - \dfrac{\bar{e}_i}{\bar{t}} & \bar{e}_i < \bar{t} \\ 0 & \bar{e}_i \geq \bar{t} \end{cases} \tag{5.2}$$

with the mean image return in the target region given by $\bar{t}$ and the mean in each of the edge regions by $\bar{e}_i$. The edge-based contour is therefore evolved so as to maximise this metric. The edge regions have a thickness of either two or three pixels depending on the image resolution; there is an obvious trade-off between the need to find thin edges and having enough pixels to be confident an edge is present but this has not been investigated in any detail.

The difference between the two approaches is in how the two algorithms treat the other bright returns in the image whilst trying to delineate a single target. The original area-based approach uses all of the image pixels to compute an overall likelihood, meaning that the presence of any extra bright returns in the background region is seen as unlikely. This leads to the delineation trying to encapsulate both targets within a single contour. The new approach only uses those image pixels within the target and edge regions. The additional target returns are therefore not of interest, although depending on how the implementation is constrained, the algorithm may fit to either return.

A key advantage of the new technique is therefore its ability to fit a single target at a time, rather than requiring multiple contours to be seeded and evolve in parallel. Because a single contour is fitted at a time the search space is much smaller – multiple targets do not become coupled. Therefore, with the new algorithm we simply search all possible states for the contour (width, length, orientation, and position) within a local area about the detected pixels. This removes the need to use heuristics such as simulated annealing and multiple hypothesis delineation and also removes the persistent concern that the algorithm may become stuck in local minima despite best efforts to avoid this. The exhaustive search over width, length orientation and position requires the comparison of approximately 30 million states which is readily achievable with a modern multi-core computer. The key point is that in the case of area based snakes the size of the search space may be raised to the power of n when n rectangles are coupled and the use of edge information avoids this combinatorial explosion.

7. Application to SAR imagery

Example results are presented showing the detection of a set of vehicles deployed in the open. The SAR imagery was collected using QinetiQ's Enhanced Surveillance Radar (ESR) shown in FIGURE 7. The radar is fully programmable with an operating frequency of 9.7 GHz and an instantaneous bandwidth of up to 500 MHz. Spotlight SAR imagery was produced using QinetiQ's 3rd generation image formation processor incorporating state of the art motion compensation and autofocus algorithms described in Jahangir et al 2008. FIGURE 8 shows an example of detection of vehicles in close proximity. The left hand image shows a spotlight SAR image containing a number of vehicles, the right hand image shows the set of vehicle detections overlaid on the scene. The two adjacent vehicles toward the bottom of the scene have been successfully detected and accurately delineated. A number of false alarms can be seen in the top right of the image indicating the requirement for additional processing to distinguish vehicles from clutter.

FIGURE 7. QinetiQ's Enhanced Surveillance Radar carried by BAC One-Eleven 500 ZH763

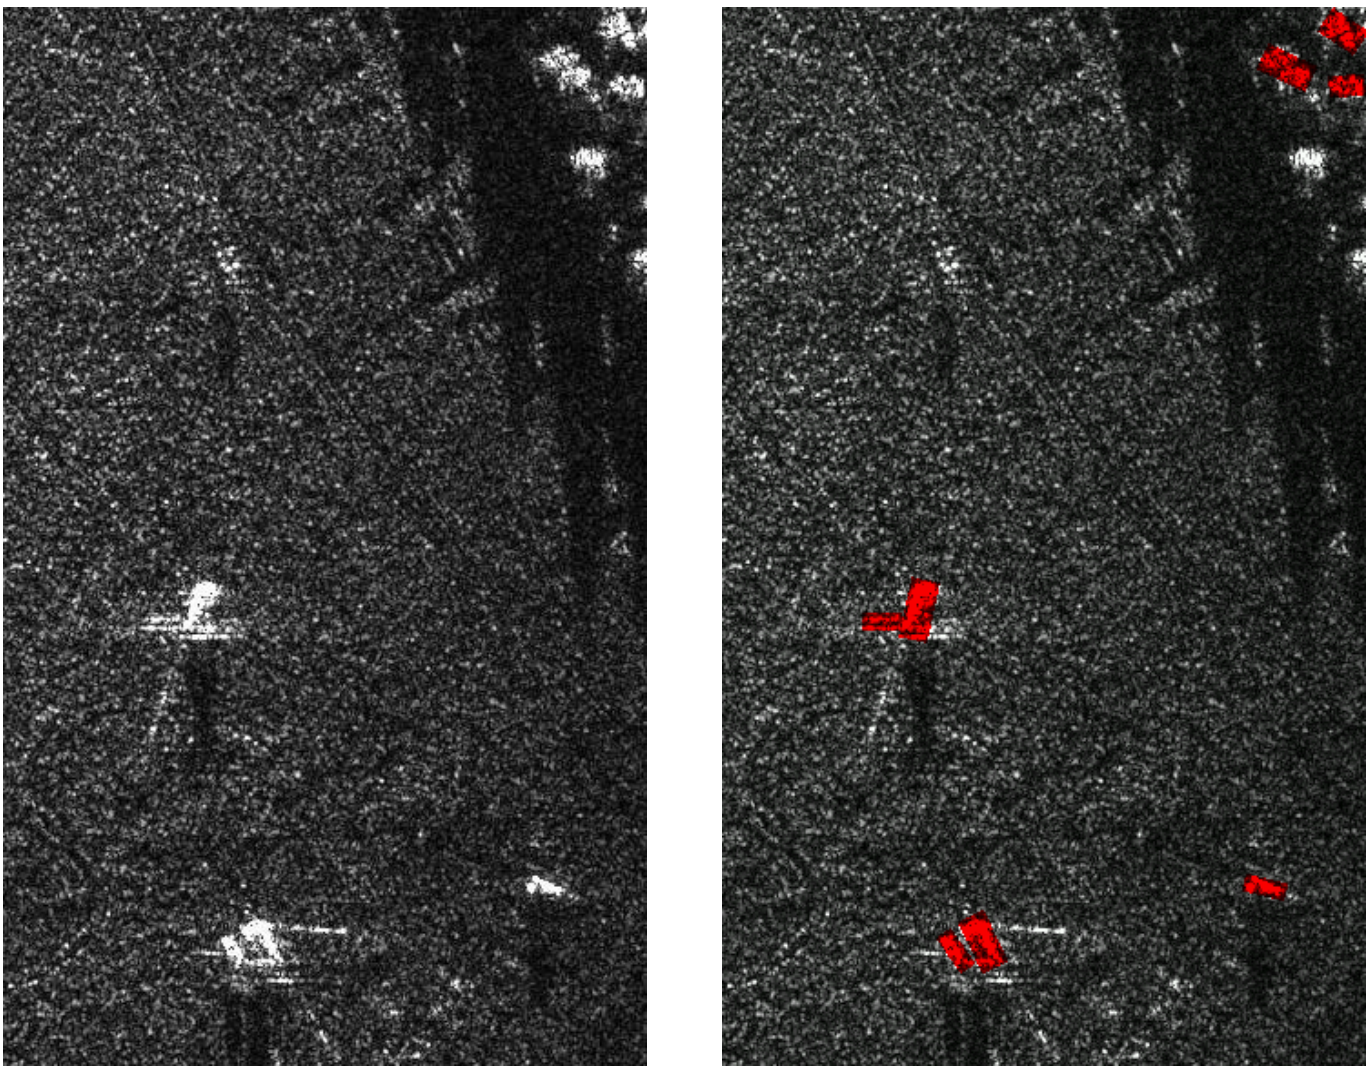

FIGURE 8. Vehicle detection results. Left: Original SAR image, right: SAR image with vehicle detections overlaid

8. Conclusions

This paper describes the development and analysis of a pre-screening algorithm that has been demonstrated to provide robust detection performance for military vehicles deployed in close proximity and in complex clutter. Target detection is based on the use of a K-CFAR scheme with parameter estimation based on percentiles, which improves the ability to detect targets in close proximity. Results from simulations illustrate the improvements over earlier work where the model parameters were estimated using the mean of the image intensity and log intensity. A novel edge based vehicle detector is described which uses an exhaustive search to fit vehicle sized rectangles and overcomes many of the problems associated with the random search strategies frequently used for snake-based delineation. The resulting scheme is computationally efficient and suitable for the first stage of an Automatic Target Recognition system designed to process large quantities of image data.

9. Acknowledgements

This work was undertaken with funding from the UK MOD via the Defence Science and Technology Laboratory (DSTL), contract number DSTLX-1000087464 (TIN 3-26). The authors would like to thank Prof David Blacknell for his support and Alan Blake for support with SAR image processing.

REFERENCES

BLACKNELL, D. 1994. Comparison of parameter estimators for the K-distribution. *EE Proc Radar, Sonar, Navig*, **141**, No 1

BLACKNELL, D. & TOUGH, R. J. A. 2001. Parameter estimation for the K-distribution based on [zlogz] *IEE Proc Radar, Sonar, Navig*, **148**, No 6

CHESNAUD, C. REFREGIER, P. & BOULET, V. 1999. Statistical Region Snake-Based Segmentation Adapted to Different Physical Noise Models, *IEEE Trans. On Pattern Analysis and Machine Intelligence*, **Vol. 21,** No. 11

JAHANGIR, M, COE, D., BLAKE, A. P., KEALEY, P. G. AND MOATE C. P. 2008 PodSAR: A versatile real-time SAR GMTI surveillance and targeting system, IEEE Radar Conference

JAKEMAN, E. & PUSEY, P. N. 1976 A model for non-Rayleigh sea echo. *IEEE Trans. Antennas Propag*, **AP-24**, 808

MOATE, C. P. & DENTON, J. 2006. SAR image delineation of multiple targets in close proximity. Proc. Of EUSAR 2006, Dresden, Germany, 16th-18th May 2006

MOATE, C. P. & LEONARD, T. P. 2015. Edge-based delineation of SAR imagery applied to multiple vehicles in close proximity. *NATO SET-228 Specialists Meeting*

OLIVER, C. J. 1994. Optimum texture estimators for SAR clutter. *J. Phys. D, Appl. Phys.*, **26**, 1824-1835

OLIVER, C. & QUEGAN, S. 1998 Understanding Synthetic Aperture Radar Images. *Artech House, Boston and London*

PEDLAR, D. 2004. Target delineation and classification using a region-based active-contour methodology. EUSAR 2004, Ulm, Germany, 25-27th May 2004

WARD, K. D. 1981. Compound representation of high-resolution sea clutter. *Electron. Lett.* **17(16)**, 561-563, 1981

WARD, K. D., BAKER, C. J. and WATTS, S. 1990. Maritime surveillance radar Part 1: radar scattering from the ocean surface *IEE Proceedings*, **137**, Part F, No 2

WARD, K. D., TOUGH, R. J. A. & WATTS, S. 2006. Sea Clutter: Scattering, the K Distribution and Radar Performance. *Institute of Engineering and Technology*