

Maximum A-Posteriori Adaptive Masking for Clutter Suppression in Automatic Radar Target Recognition

Jingjing Cui, Jon Gudnason, Mike Brookes
EEE Department, Imperial College London, London, UK
Email: {Jingjing.Cui, Jon.Gudnason, Mike.Brookes}@imperial.ac.uk

Abstract—To suppress the effect of clutter in automatic target recognition we propose a maximum a-posteriori adaptive mask that isolates the target scattering centres in the data. This improves the quality of the construction of the target feature vector which results in an improved classification performance. We evaluate the technique on a closed-set identification task using targets from the MSTAR database and, on a 10-target forced-choice experiment, obtain a recognition error rate of only 1.2%. We also conduct an open set identification experiment and show that the technique can achieve an equal error rate of 1.82%.

I. INTRODUCTION

The automatic detection and classification of targets from their radar signatures is an important and difficult problem that has attracted considerable research effort. In previous papers [1], [2] we have presented novel methods for extracting features from high range resolution (HRR) radar data that characterize both the target scattering centres and the shape of the radar shadow. These features have been used with a Hidden Markov model (HMM) that represents variations to target orientation. The resulting classifier is able to recognize stationary targets with excellent results. Because the identification of the target scattering centres is performed using the super-resolution MUSIC algorithm, the azimuth resolution obtained is independent of the azimuth aperture of the data acquisition system and good recognition performance has been obtained even with small azimuth apertures.

The problem of recognizing ground targets is made more difficult by the presence of clutter which can contain isolated peaks that are comparable in magnitude to the returns from target scattering centres. Many alternative techniques have been proposed to improve the discrimination between target and clutter returns. In [3], [4], the mean power of the clutter returns in the vicinity of the target is estimated and this is used to determine an adaptive detection threshold that is chosen to maintain a constant false alarm rate. In [5], [6], a target-free portion of data is used to estimate an autoregressive model of the clutter. This is then used to whiten the clutter in the target region thereby improving the signal-to-noise ratio. In [7], a wavelet transform is adaptively selected to achieve the greatest concentration of image energy into a relatively small number of coefficients. To remove clutter, a threshold is applied to the observation data in this domain which is then transformed back into the SAR image domain for recognition.




Fig. 1. (a) A sequence of 100 HRR profiles from a T72 tank in the MSTAR database and (b) the resultant SAR image.

A wavelet-based approach to removing clutter is also used in [8], but in this case it is applied to the HRR profiles and a translation-invariant transform is used. In a completely different approach, [9], [10] attempt to distinguish between target and clutter returns by estimating the fractal dimension of the returned signal. Both authors report significant differences and improved discrimination.

The MUSIC-based algorithm used to identify the target scattering centres in [1], [2] is inherently robust to clutter noise. In the work described previously, we have nonetheless eliminated gross errors by rejecting any centres that lie outside a fixed-size mask around the target. Because the mask size is fixed, it is sometimes the case that scattering centres arising from nearby clutter are wrongly included in the feature extraction process. Conversely, a large target may occasionally extend beyond the mask resulting in the exclusion of valid scattering centres. In this paper we introduce an adaptive mask whose size adjusts to the observed target and provides the improved recognition performance.

This paper begins, in Section II, with a brief description of the feature extraction process. In Section III, we give details of how the size and position of the adaptive target mask is determined and how the target features are derived. In Section IV, we then give a brief overview of the hidden Markov model architecture for which a more detailed explanation is presented in [1], [2]. In Section V we present evaluation results for the system that incorporates adaptive masks. We conclude the paper in Section VI with a brief summary.




Fig. 2. Feature extraction process. The three parallel paths are respectively for shadow feature extraction (upper path), target feature extraction (centre path) and target mask identification (lower path).

II. RECOGNITION FEATURE SET

The features that we use for target recognition are derived from the sequence of complex-valued HRR profiles, $x(n, k)$, obtained by applying a discrete Fourier transform (DFT) to the windowed phase history radar returns. Here n is the profile index and k is the range-bin index covering the region of interest. Fig. 1(a) shows a typical plot of $|x(n, k)|$ and Fig. 1(b) shows the SAR image that results from windowing $x(n, k)$ and taking the DFT with respect to n .

Visible in the image are the target itself, with signal levels well above the clutter noise level and also a well defined shadow region with very low signal levels. For each value of the profile index n , we obtain a feature vector, $\mathbf{u}(n)$, that characterises the positions and intensities of the scattering centres within the target and another, $\mathbf{w}(n)$, that characterises the shape of the shadow area. Both these feature vectors are derived from $2P + 1$ consecutive profiles centred on profile n . We therefore define the data matrix $x_n(p, k) = x(n + p, k)$ where $p \in \{-P, \dots, P\}$.

The stages of feature extraction are illustrated in Fig. 2 where it can be seen that the observed HRR profiles are processed in three parallel paths: shadow feature extraction (upper path), target feature extraction (centre path) and target mask identification (lower path). The method of forming target and shadow features has been described in detail in [1], [2]. In this paper, we concentrate on the target mask identification procedure which forms the lower path in Fig. 2.

III. TARGET MASK IDENTIFICATION AND FEATURE EXTRACTION

The identification of the adaptive target mask is based solely on the observed data and forms part of the feature extraction process that converts a sequence of HRR profiles into a feature vector.

A. Region of Interest

Fig. 3 shows (a) a sequence of 100 HRR profiles from a T-72 tank from the MSTAR database and (b) the resultant SAR image. The solid lines in the SAR image define the $8\text{m} \times 8\text{m}$ square region that is the initial region of interest and it can be seen that the target lies within this region but is offset slightly to the left. Using the procedure described below, we determine the rectangular target mask that is shown on the image using dashed lines. It can be seen that this mask fits

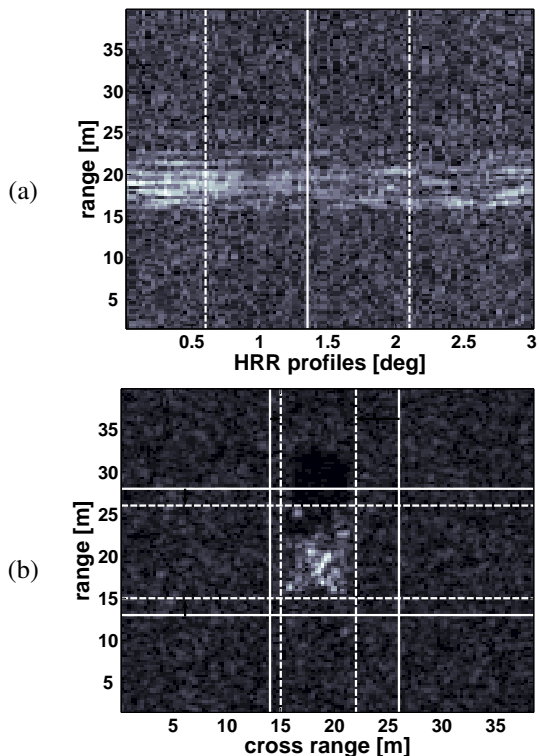


Fig. 3. (a) A sequence of 100 HRR profiles from a T72 tank in the MSTAR database and (b) the resultant SAR image. The vertical dashed lines in (a) show the sequence of HRR profiles used in forming a single feature vector. In (b), the solid lines show the initial region of interest and the dashed lines show the adaptive mask region.

snugly around the true target position and includes all its true scattering centres.

This target mask is used, as indicated in Fig. 2, to eliminate any scattering centres that may have been wrongly extracted from the clutter region. The positions of the target scattering centres that have been extracted via the central path of Fig. 2 are compared with the mask; any centres that lie outside the mask are assumed to arise from clutter and are discarded.

B. Adaptive Mask Identification

To identify the target mask shown with dashed lines in Fig. 3(b), we first form an image from the sequence of HRR profiles. As described in [2], the number of profiles used to form each feature vector is restricted to $2P + 1$ where P is a system parameter. The vertical dashed lines in Fig. 3(a) indicate the range of profiles used for a single feature vector when $P = 25$. These $2P + 1$ are used to form an image of the target whose central portion is shown in Fig. 4(a). To determine the target mask, we first segment the pixels of this image into “target” and “clutter” classes using a maximum likelihood criterion and we then use as the target mask the smallest bounding rectangle that encloses all the “target” pixels.

In the first step, we will partition the image pixels into two classes, a “target” class ξ_1 , and a “clutter” class ξ_2 . We assume that within each of these classes, the pixel values are taken

from a complex-valued gaussian distribution with zero mean and variance σ_k in both real and imaginary parts. Thus the log-likelihood of the complex-valued pixel $x(i, j)$, conditioned on the class ξ_k ($k = 1, 2$), is given by

$$\begin{aligned} L(x(i, j)|\xi_k) &= \log(\mathcal{N}(x_{\Re}(i, j); 0, \sigma_k^2)) + \\ &\quad \log(\mathcal{N}(x_{\Im}(i, j); 0, \sigma_k^2)) \\ &= -\log(2\pi\sigma_k^2) - \frac{|x(i, j)|^2}{2\sigma_k^2} \end{aligned} \quad (1)$$

where $\mathcal{N}(x; \mu, \sigma_k^2)$ is a normal distribution with mean μ and variance σ_k^2 and $x_{\Re}(i, j)$ and $x_{\Im}(i, j)$ are the real and imaginary parts respectively of the pixel $x(i, j)$.

We apply an iterative procedure to determine both the partitioning of the image pixels and the two variances, σ_k^2 . We begin by estimating initial values for σ_1^2 and σ_2^2 from the variance of small blocks of pixels in the centre of the image and outside the region of interest respectively. We next classify each of the image pixels within the region of interest based on their smoothed log-likelihood:

$$k(i, j) = \underset{k}{\operatorname{argmax}} (w(i, j) * L(x(i, j)|\xi_k)) \quad (2)$$

where $*$ denotes 2-dimensional convolution and $w(i, j)$ is the impulse response of a smoothing filter. The choice of $w(i, j)$ is a compromise between the elimination of isolated strong responses from the clutter region and the preservation of small regions that truly belong to the target. In the results described below, we have used a uniform 3×3 window for the smoothing filter; other researchers have used non-linear smoothing operations in a similar way [11].

Using the partition defined by (2), we iteratively recalculate the class variances σ_k^2 and re-evaluate $k(i, j)$ for 10 iterations or until

$$\sum_{i, j} L(x(i, j)|\xi_{k(i, j)}) \quad (3)$$

ceases to increase.

Finally, we introduce prior probabilities p_k satisfying $p_1 + p_2 = 1$ and determine their values so that the partition

$$k(i, j) = \underset{k}{\operatorname{argmax}} (w(i, j) * L(x(i, j)|\xi_k) + \log(p_k)) \quad (4)$$

maximizes the smoothed log likelihood summed over all pixels.

Fig. 4(b) shows the resultant partitioning of the pixels together with the bounding rectangle that is used as the final target mask. It is clear that the limits of the target have been correctly identified in this instance.

C. Image Reconstruction

We have described in previous papers [1], [2] how we identify target scattering centres using the MUSIC algorithm [12] based on the following data model

$$x(p, r) = \sum_{m=1}^{M_r} a_{m,r} e^{j\omega_{m,r}p} + v(p, r) \quad (5)$$




Fig. 4. (a) A SAR image formed using $2P + 1 = 51$ HRR profiles and (b) the pixels in the “target” class with the final target mask shown as a bounding rectangle.

where r is the range-bin index, M_r is the number of extracted scattering centres in the range bin, $p \in \{-P, \dots, P\}$ is the HRR profile index, $a_{m,r}$ and $\omega_{m,r}$ are the radar cross section and azimuth location of the m^{th} scattering centre and $v(p, r)$ is assumed to be white noise. These super-resolution scattering centres are low-pass filtered and sampled on a fixed grid according to

$$y(l, r) = \sum_{m=1}^M |a_{m,r}|^2 h(l - \beta\omega_{m,r}) \quad (6)$$

where the low-pass filter response is given by

$$h(l) = \frac{\sin(\pi l)}{\pi l} \quad (7)$$

and the scaling constant β is given by $\beta = 0.5\lambda(2\pi\Delta\phi\Delta r)^{-1}$ where $\lambda, \Delta\phi, \Delta r$ are respectively the wavelength, azimuth increment and cross range resolution.

Fig. 5 shows the reconstructed image both without (a) and with (b) the use of the adaptive mask. Two scattering centres that lie outside the mask are explicitly identified in Fig. 5(a) and it can be seen that the corresponding impulses are missing in (b). Also visible in this figure is the effect of the low-pass filter impulse response (7) which smears the image in cross-range. The smearing is emphasized by the log intensity scale of these images.

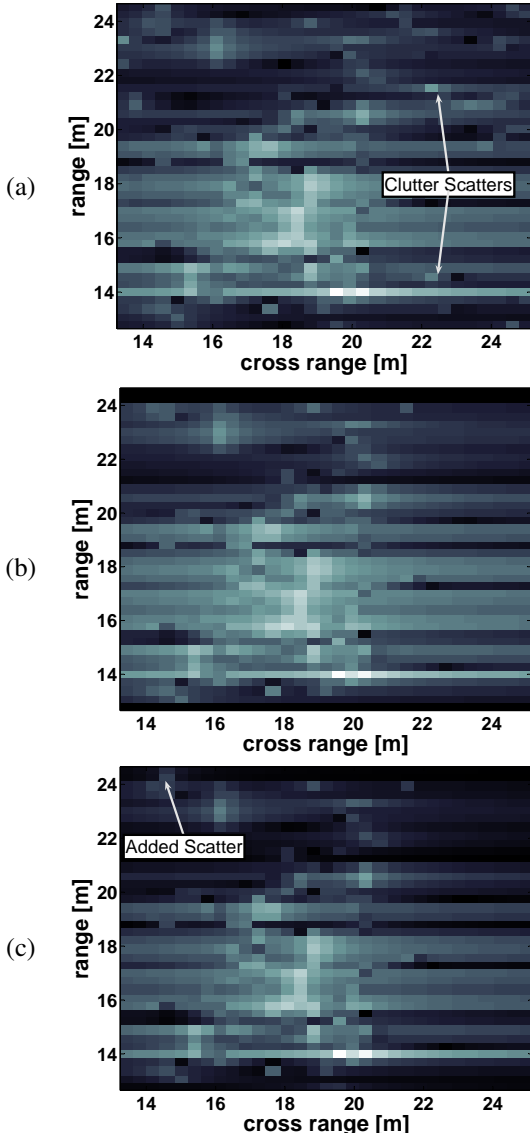


Fig. 5. Target images reconstructed from the extracted scattering centres and displayed on a logarithmic intensity scale. In (a) no target mask has been used; two of the false scattering centres are indicated with arrows. In (b) an adaptive target mask has been used; “black lines” containing no scattering centres are visible at the top and bottom of the image. In (c) at least one scattering centre is extracted for each row of the image.

D. Black Line Elimination

Following the image reconstruction process, the target feature vector is formed by taking the 2-dimensional Discrete Cosine Transform (DCT) of the log of the image in Fig. 5 and retaining only the low frequency coefficients. Unfortunately, when the target mask is applied, it is possible that some range bins contain no valid scattering centres at all. These “black lines” can occur in any of the range bins, but most commonly arise at the top and bottom of the image. Whenever this happens, all the pixels in that range bin will be identically zero and, following the log operation, these pixels will dominate the feature vector resulting in poor performance. To avoid this situation, we detect any range bin, r , that is empty of scattering

centres and, if necessary, increase M_r in (5) until at least one scattering centre is found that lies within the cross-range limits of the target mask.

Fig. 5(b) shows the reconstructed target image before this black line elimination procedure has been applied and it can be seen that there are black lines visible at the top and bottom of the image. In Fig. 5(c) additional scattering centres have been allowed and each range bin now contains at least one. The added scattering centres are normally of low intensity.

Finally, as noted above, the target feature vector is formed by taking the DCT of the clipped log of the image in Fig. 5(c) and retaining 55 low frequency coefficients. The $(0, 0)$ coefficient depends on the gain of the radar system and is therefore excluded from the feature vector.

IV. ASPECT ANGLE HMM

The architecture of our target recognition system was described in [1], [2] and only a brief description will be repeated here. The target features extracted from the radar returns will vary with target orientation and to account for this variation we use a hidden Markov model each of whose states corresponds to a range of aspect angles. We fix the number of azimuth states at 60 and initialize them to represent equal aspect ranges of 6° . Fig. 6 illustrates the HMM and shows the allowed state transitions. During the training of a particular target model, the aspect range of each state is allowed to vary freely although the total number of states remains fixed at 60. After training, some of the states correspond to azimuth ranges of only about 1° while others to substantially larger ranges. This reflects the fact that for some target aspects, the radar range profile changes very rapidly with azimuth. Within each state we represent the probability distribution of the feature vector as a mixture of multivariate Gaussians having diagonal covariances.




Fig. 6. The architecture of the Hidden Markov Model whose states represent different target orientation.

V. EVALUATION OF ADAPTIVE MASKS

A. MSTAR Database

The experimental evaluations make use of the Moving and Stationary Target Acquisition and Recognition (MSTAR) database collected by the Sandia National Laboratory using an X-band SAR sensor in 0.3m resolution spotlight mode [13], [14]. The database contains complex-valued SAR image chips of 10 confusable targets and their variants. For each target, the images cover a full azimuth range at depression angles of 17° and 15° for training and test data respectively. The SAR images have a resolution of $\Delta r = 0.3\text{m}$ in both the range (horizontal) and cross range (vertical) directions.

For our experiments, the SAR image chips were converted into a sequence of HRR profiles by taking an inverse Fourier transform, removing the zero padding and finally undoing the Taylor window in the cross-range direction [15].

B. Closed Set Results using Adaptive Masks

To evaluate the effectiveness of using adaptive masks, we compared closed-set target identification results from experiments with and without adaptive masks. We also evaluated the effect of including or omitting the black line removal procedure described in Section III-D.

Table I shows the misclassification rate of each target in the MSTAR database using target features alone (\mathbf{u}) and when combined with shadow features ($\mathbf{u} + \mathbf{w}$) as described in [1], [2]. The inclusion of the shadow features always results in a substantial improvement in performance. The first pair of columns in Table I show the results obtained for an HRR profile sequence spanning a 3° aperture angle. Each row shows the misclassification rate for one of the targets in the database while the last row shows the overall test-set misclassification rate (TMCR). The second pair of columns show the corresponding figures when the adaptive mask is used with the black line removal procedure. The final pair of columns show the results when the adaptive mask is used but no black line elimination is performed. It can be seen that the presence of black lines seriously degrades performance and that for almost all targets, the performance is worse than without the adaptive mask.

We see that the overall test set misclassification rate when using target features alone (\mathbf{u}) is reduced from 4.2% to 3.2% which corresponds to an error rate reduction of 24%. All but two of the targets show an improvement and the error rate for the most difficult target falls from 12.8% to 8%. When the target features are combined with shadow features, the overall test set misclassification rate falls from 1.3% to 1.2% which represents an improvement of about 8% on an already excellent performance.

The error rate of 1.2% obtained using a 3° aperture can be directly compared with other published results based on the MSTAR database with the same recognition task. In [3], the authors obtained error rates of 4.1% using an approach based on the SAR image and in [15] an error rate of 17.8% was obtained when performing recognition on the HRR profiles directly.

C. Open Set Experiments

The experiments presented above are based on closed-set identification where the observation is known to be of a target of interest. It is also important that a recogniser should be able to reject an out-of-set object, or non-target. Here we selected clutter regions that are centred on a distinctive radar return as impostor observations. We have designed a region of interest (ROI) detector that is able to extract plausible impostor images from the clutter observations available in the MSTAR database, shown in Fig. 7. We describe below the operation of

TABLE I
TEST SET MISCLASSIFICATION RATE FOR CLOSED-SET MSTAR TARGET IDENTIFICATION. THE TABLE DEMONSTRATES THE EFFECTIVENESS OF USING ADAPTIVE MASK AND THAT ENHANCING THE IMAGE BY COMPENSATING FOR BLACK LINES IS NECESSARY.

Target	No mask		No Black Lines		With Black Lines	
	\mathbf{u}	$\mathbf{u} + \mathbf{w}$	\mathbf{u}	$\mathbf{u} + \mathbf{w}$	\mathbf{u}	$\mathbf{u} + \mathbf{w}$
BMP2	3.6	1.0	2.6	1.6	10.8	5.7
BRDM2	12.8	3.7	8.0	4.0	18.3	9.9
BTR60	0.0	1.6	0.5	1.6	1.0	2.1
BTR70	4.1	2.1	2.6	0.5	5.2	2.6
D7	1.8	0.4	1.5	0.0	2.6	1.1
T62	6.6	0.4	5.5	1.1	10.9	2.2
T72	1.6	0.0	1.0	0.0	2.6	0.5
ZIL131	1.8	0.0	1.5	0.0	6.2	0.4
ZSU234	1.5	0.7	2.6	0.7	2.2	0.4
2S1	5.8	2.9	4.7	2.2	13.5	3.3
TMCR	4.2	1.3	3.2	1.2	7.7	2.8

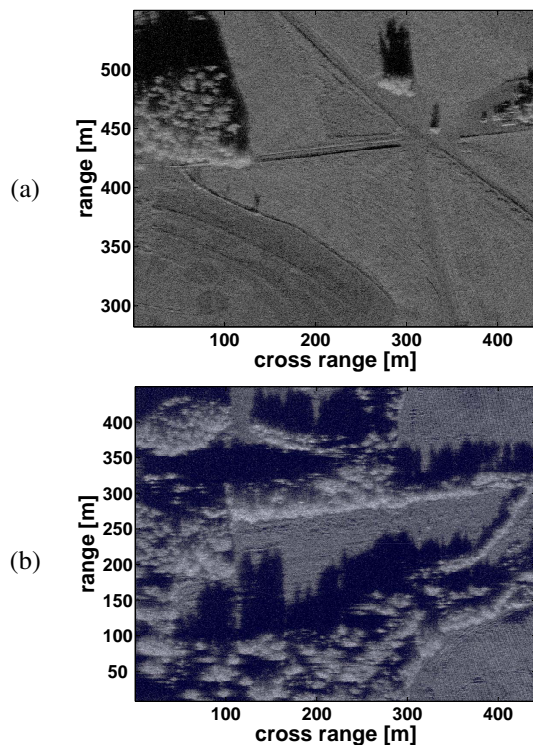


Fig. 7. Clutter images from the MSTAR database.

the ROI detector and present the open-set identification results that we have obtained.

Our region-of-interest (ROI) detector extracts small impostor regions from the $500 \text{ m} \times 500 \text{ m}$ clutter images shown in Fig. 7 that are from part of the MSTAR dataset. We scan these images to search for plausible targets in the clutter based on the power in the sub-image and the target mask that we employed. The algorithm divides each clutter image into smaller sub-images with an overlap of 15%. The power of the sub-image is computed as

$$P_{Image} = \sum_{(i,j) \in \text{Image}} |I(i,j)|^2 \quad (8)$$




Fig. 8. Impostor-targets selected from the MSTAR clutter images and used in open-set evaluations.

and similarly for the power of the target mask

$$P_{Mask} = \sum_{(i,j) \in Mask} |I(i,j)|^2 \quad (9)$$

where $I(i,j)$ is the complex SAR sub-image, $Image$ is the set of all the pixels in the image and $Mask$ is the set of all pixels in the target mask. The images with the highest ratio $P_{Mask} : P_{Image}$ were selected to be impostor-targets in the open-set identification experiments presented below. Fig. 8 shows two impostor-targets from the clutter images.

As in the closed-set experiments, a test sequence is compared against each of the 10 target models to give a list of likelihoods but now the observed sequence can be rejected as a non-target if the highest likelihood falls below a predetermined likelihood-threshold. Otherwise, the observation is accepted as a genuine target, whose identity is that of the target model that produced the highest likelihood. This procedure gives rise to three types of error that we present below: (1) false rejection, when a genuine target sequence is rejected as a non-target, (2) false acceptance, when an impostor sequence is classified as a target, and (3) misclassification, when a genuine target sequence is identified as the wrong target.

The ROC curve in Fig. 9(a) shows the false acceptance rate (FAR) as a function of the false rejection rate (FRR) using only the target features, \mathbf{u} . Each point in the plot represents the false rejection and false acceptance rates obtained with a particular threshold. Setting the threshold to a different value changes these rates. We have also, for convenience, plotted the equal error rate line for which FAR and FRR are equal. We can see the curve has an equal error rate of 2.54%. Fig. 9(b) shows the misclassification rate as a function of the false rejection rate as the threshold is varied using only the target features, \mathbf{u} . As the false rejection rate increases, the misclassification rate falls to a limiting value of about 1%. This indicates that the marginally accepted targets are those most likely to be misclassified. Fig. 10 shows the same plots as Fig. 9 but using the combined target and shadow feature set, $\mathbf{u} + \mathbf{w}$. The ROC curve has moved closer to the axes, the equal error rate has fallen to 1.82% and the limiting misclassification rate is now about 0.3%.

VI. SUMMARY AND CONCLUSIONS

This paper has described the use of an adaptive target mask that allows us to exclude spurious scattering centres arising

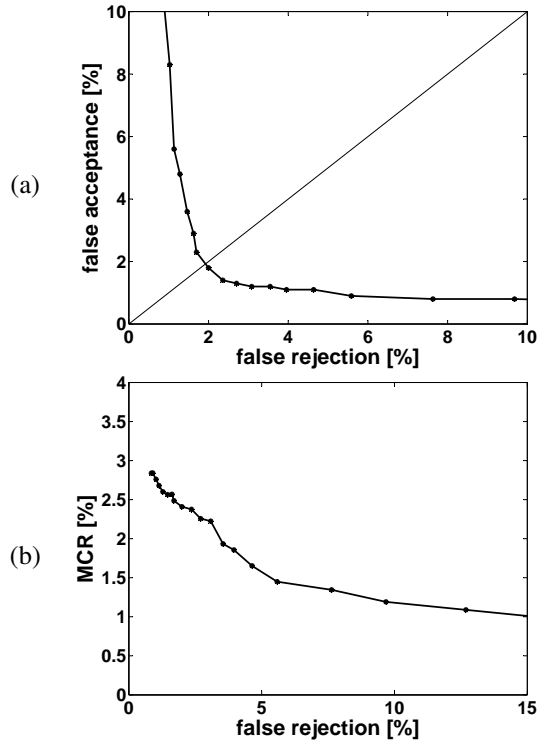


Fig. 9. Open-set evaluation results using only the target features (\mathbf{u}). The ROC curve, (a), plots the false acceptance rate versus the false rejection rate as the acceptance likelihood threshold is varied. Graph (b) plots the misclassification rate versus the false rejection rate.

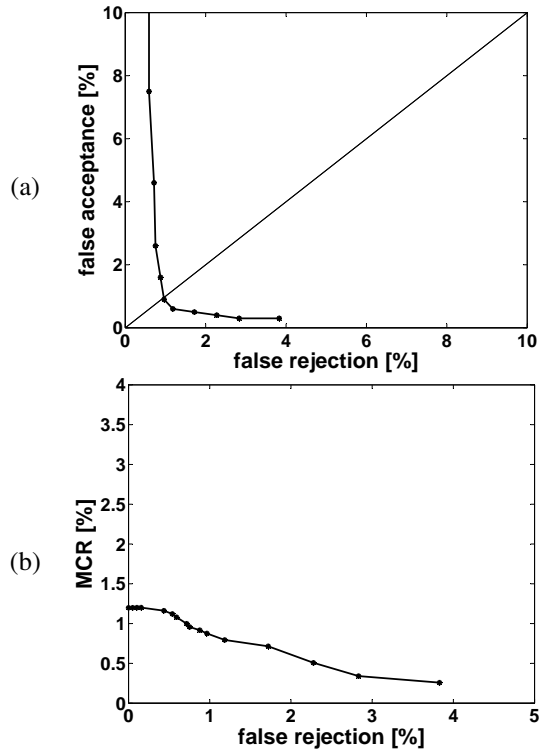


Fig. 10. Open-set evaluation results using both target and shadow features ($\mathbf{u} + \mathbf{w}$). The ROC curve, (a), plots the false acceptance rate versus the false rejection rate as the acceptance likelihood threshold is varied. Graph (b) plots the misclassification rate versus the false rejection rate.

from clutter adjacent to the target. We show that the use of this target mask gives a reduction in error rate from 4.2% to 3.2% when only the target features are used. When target features are combined with shadow features, the error rate falls more modestly from 1.3% to 1.2%. We have found that in order to obtain these improvements it is necessary to address the situation when some of the range bins contain no valid scattering centres at all and we present an effective method of doing this. We have also demonstrated that the feature sets work well on an open set identification task and achieve an equal error rate of 1.82%.

VII. ACKNOWLEDGMENT

This work was supported by the UK MoD through work funded by the Defence Technology Center for Data and Information Fusion.

REFERENCES

- [1] J. Cui, J. Gudnason, and M. Brookes, "Automatic recognition of MSTAR Targets using radar shadow and superresolution features," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2005.
- [2] ———, "Radar shadow and superresolution features for automatic recognition of MSTAR targets," in *Proc IEEE Int. Radar Conf*, May 2005, pp. 534–539.
- [3] L. M. Novak, G. J. Owirka, and A. L. Weaver, "Automatic target recognition using enhanced resolution SAR data," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 35, no. 1, pp. 157–175, 1999.
- [4] G. A. Lampropoulos, G. Gigli, L. Sevigny, A. Beaudoin, and J. Secker, "Detection of targets from electro-optical and SAR data using chaotic predictors and optimal CFAR detectors," in *IEEE International Geoscience and Remote Sensing Symposium*, vol. 1, 2002, pp. 107–109.
- [5] J. Li, G. Liu, N. Jiang, and P. Stoica, "Airborne phased array radar: clutter and jamming suppression and moving target detection and feature extraction," in *IEEE Sensor Array and Multichannel Signal Processing Workshop*, 2000, pp. 240–244.
- [6] N. Jiang and J. Li, "Multiple moving target feature extraction for airborne HRR radar," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 37, no. 4, pp. 1254–1266, 2001.
- [7] H. Deng and H. Ling, "Clutter reduction for synthetic aperture radar images using adaptive wavelet packet transform," in *IEEE International Symposium Antennas and Propagation Society*, vol. 3, 1999, pp. 1780–1783.
- [8] B. M. Huether, S. C. Gustafson, and R. P. Broussard, "Wavelet preprocessing for high range resolution radar classification," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 37, no. 4, pp. 1321–1332, 2001.
- [9] J. I. Butterfield, "Fractal interpolation of radar signatures for detecting stationary targets in ground clutter," *IEEE Antennas and Propagation Magazine*, vol. 6, no. 7, pp. 10–14, 1991.
- [10] J. Kirk, J. C., "Using the fractal dimension (FD) to discriminate between targets and clutter," in *IEEE Radar Conference*, 1993, pp. 76–78.
- [11] S. Haker, G. Sapiro, and A. Tannenbaum, "Knowledge-based segmentation of SAR data with learned priors," *IEEE Trans. on Image Processing*, vol. 9, no. 2, pp. 299–301, Feb. 2000.
- [12] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [13] T. D. Ross, S. W. Worrell, V. J. Velten, J. C. Mossing, and M. L. Bryant, "Standard SAR ATR evaluation experiments using the MSTAR public release data set," in *Proc SPIE: Algorithms for Synthetic Aperture Radar Imagery V*, vol. 3370, Apr. 1998, pp. 566–573. [Online].
- [14] T. D. Ross and J. C. Mossing, "The MSTAR Evaluation Methodology," in *Proc SPIE: Algorithms for Synthetic Aperture Radar Imagery VI*, vol. 3721, Aug. 1999, pp. 705–713.
- [15] X. Liao, P. Runkle, and L. Carin, "Identification of ground targets from sequential high-range-resolution radar signatures," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 38, no. 4, pp. 1230–1242, Oct. 2002.