

method of Lagrangian multipliers

$$\tilde{\mathbf{c}} = \frac{\mathbf{Q}^{-1}\mathbf{e}}{\mathbf{e}^T\mathbf{Q}^{-1}\mathbf{e}}. \quad (14)$$

ACKNOWLEDGMENTS

The authors thank the Systems Surveillance Division at DSTO for allowing use of the Jindalee array, and acknowledge J. Ayliffe for assisting in operating the receiving system, P. Bawden and B. White for help in replaying the data, and M. Peake for useful discussions.

G. A. FABRIZIO
Surveillance Systems Division
Defence Science and Technology Organisation
P.O. Box 1500
Salisbury
South Australia 5108
Australia
E-mail: (joe.fabrizio@dsto.defence.gov.au)

and
Dept. of Electrical and Electronic
Engineering
The University of Adelaide
South Australia, 5005

D. A. GRAY
Cooperative Research Centre for Sensor Signal
and Information Processing
SPRI Building, Warrendi Rd.
Mawson Lakes
South Australia, 5095
Australia
E-mail: (dgray@cssip.edu.au)

and
Dept. of Electrical and Electronic
Engineering
The University of Adelaide
South Australia, 5005

M. D. TURLEY
Surveillance Systems Division
Defence Science and Technology Organisation
P.O. Box 1500
Salisbury
South Australia, 5108
Australia
E-mail: (mike.turley@dsto.defence.gov.au)

REFERENCES

- [1] Farina, A. (1992)
Antenna-Based Signal Processing Techniques for Radar Systems.
Norwood, MA: Artech House, 1992.
- [2] Compton, R. T. (1988)
Adaptive Antennas—Concepts and Performance.
Englewood Cliffs, NJ: Prentice Hall, 1988.
- [3] Monzingo, R. A., and Miller, T. W. (1980)
Introduction to Adaptive Arrays.
New York: Wiley, 1980.
- [4] Johnson, J. R., Fenn, A. J., Aumann, H. M., and Willwerth, F. G. (1991)
An experimental adaptive nulling receiver utilising the sample matrix inversion algorithm with channel equalisation.
IEEE Transactions on Microwave Theory and Techniques, **39**, 5 (May 1991), 798–808.
- [5] Abramovich, Y. I., Kachur, V., and Struchev, V. (1984)
Methods of digital channel corrections in multichannel radar receivers.
Radio Engineering and Electronic Physics, **29**, 9 (1984), 62–68.
- [6] Lewis, B. L., Kretschmer, F. F., and Shelton, W. W. (1986)
Aspects of Radar Signal Processing.
Norwood, MA: Artech House, 1986.
- [7] Fabrizio, G. A., Abramovich, Y. I., Anderson, S. J., Gray, D. A., and Turley, M. D. (1998)
Adaptive cancellation of nonstationary interference in HF antenna arrays.
IEEE Proceedings—Radar, Sonar and Navigation, **145**, 1 (Feb. 1998), 19–26.
- [8] Fabrizio, G. A. (2000)
Space-time characterisation and adaptive processing of ionospherically-propagated HF signals.
Ph.D. dissertation, Adelaide University, Australia, July 2000.
- [9] Compton, R. T. (1982)
The effect of random steering vector errors in the Applebaum adaptive array.
IEEE Transactions on Aerospace and Electronic Systems, **AES-18**, 5 (Sept. 1982), 392–400.
- [10] Mallet, J. D., Reed, I. S., and Brennan, L. E. (1974)
Rapid convergence rate in adaptive arrays.
IEEE Transactions on Aerospace and Electronic Systems, **10**, 6 (1974), 853–863.
- [11] Albert, A. (1972)
Regression and the Moore–Penrose Pseudoinverse.
New York: Academic Press, 1972.

Recognizing Occluded Objects in SAR Images

Recognizing occluded vehicle targets in synthetic aperture radar (SAR) images is addressed. Recognition algorithms, based on local features, are presented that successfully recognize highly occluded objects in both XPATCH synthetic SAR signatures and real SAR images of actual vehicles from the MSTAR data. Extensive experimental results are presented for a basic recognition algorithm, using SAR scattering center relative locations as features with the XPATCH data and for an improved

Manuscript received October 4, 1999; revised May 8, 2000; released for publication November 17, 2000.

IEEE Log No. T-AES/37/1/02942.

Refereeing of this contribution was handled by L. M. Kaplan.

This work was supported in part by Grants F49620-97-1-0184 and F33615-99-C-1440. The contents and information do not reflect the policies and positions of the U.S. Government.

0018-9251/01/\$10.00 © 2001 IEEE

algorithm, using scatterer locations and magnitudes with the real SAR targets in the MSTAR data. The results show the effect of occlusion on recognition performance in terms of probability of correct identification (PCI), receiver operating characteristic (ROC) curves, and confusion matrices.

I. INTRODUCTION

The focus of this work is recognizing occluded objects, starting with occluded versions of synthetic aperture radar (SAR) image chips of various target vehicles and ending with the identification of the specific vehicle type (e.g., a SCUD missile launcher). We use both synthetic target images, generated by the XPATCH radar signature prediction code [1] and real SAR images of actual vehicles from the MSTAR (public) targets data set [15]. We present a basic recognition system for the XPATCH data, using SAR scattering center relative locations as features, and an improved recognition system, using scatterer locations and magnitudes, that achieves improved performance with the real SAR targets in the MSTAR data. These local features are the key to successfully recognizing highly occluded objects. The recognition approach uses standard models of the objects (at 1° azimuth increments) to recognize occluded versions of the same objects in standard and articulated configurations. (Articulated objects have at least one major part that can move relative to the rest of the object, such as the turret rotation of a tank or a SCUD missile erected on the launcher vehicle.) The recognition process is an efficient search for positive evidence, that uses table look-ups based on information in the occluded test image to generate votes for the appropriate object (and azimuth pose). The recognition system performance is characterized in terms of the effect of occlusion on such performance measures as probability of correct identification (PCI), confusion matrices, and ROC (receiver operating characteristic) curves.

General reviews of automatic target recognition concepts and technologies can be found in [3, 8]. The detection theory [4, 5], pattern recognition [13, 16, 17] and neural network [7, 18] approaches to SAR recognition all tend to use global features that are optimized for standard, nonarticulated, nonoccluded configurations. Approaches that rely on global features are not appropriate for recognizing occluded (or articulated) objects because occlusion (or articulation) changes global features like the object outline and major axis [20]. Some of the SAR recognition techniques, e.g., MINACE filters [6], PERFORM [10], mean squared error template matching [14] and invariant histograms [9] have reported limited test results for small amounts of occlusion, typically 25 percent or less. In addition, the developers of the MSTAR search engine reported [19] using a shadow inferencing technique to hypothesize

targets with up to 30 percent occlusion in the cross-range direction.

In contrast, we present an approach to SAR target recognition, specifically designed to accommodate articulated and occluded targets, that achieves excellent recognition results for highly occluded data with over 50 percent target occlusion. This paper is both an extension of our earlier work [11] on occluded and articulated XPATCH objects (e.g. adding performance analysis, predictions and occlusion from an unknown second object) and a significant advancement that addresses the development of new algorithms needed and the extensive experimental results obtained with the real MSTAR data. In addition to locations of scattering centers as features, this algorithm takes into consideration location uncertainty, magnitude of scattering centers as features and its uncertainty. Furthermore by comparing predicted and actual performance, it shows that these constraints are effective for target recognition. A variety of test data (that is distinct from the training/model data) is used to perform experiments. This includes 1) simulated controlled occlusion, 2) random positional and magnitude noise as additional scattering centers, 3) articulated occluded objects, 4) occlusion from an unmodeled second object, and 5) allowing for confusers and a reject or unknown class.

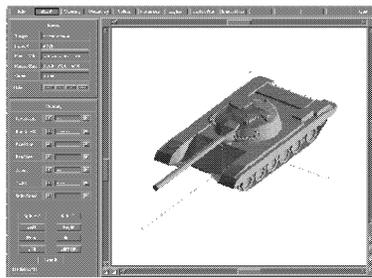
II. SAR TARGET CHARACTERISTICS

A. XPATCH Objects, SAR Images and Scattering Centers

The XPATCH radar signature prediction code [1] is used to generate 6 in resolution target chips at 360 azimuth angles (at a 15° depression angle) from CAD models of three tanks (T72, T80, and M1a1) and a SCUD missile launcher (with numbers of model surface facets ranging from 5345 to 32954). The scattering center locations are determined by finding local eight-neighbor maxima in the radar image. Examples of the object CAD model, SAR image and scattering center locations (as black squares) are shown in Fig. 1 for the T72 tank at 30° azimuth. (Fig. 1 is not to scale and the image is displayed at 8 intensity levels, the scattering center map at 256 levels.)

B. MSTAR Objects, SAR Images and Scattering Centers

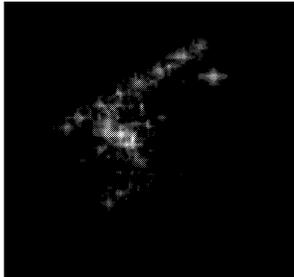
Objects from the MSTAR public data [15] used in this paper include: BMP2 armored personnel carrier (APC) serial number (#) c21, BTR70 APC #c71, T72 tanks #132 and #a64 (#a64 in the articulation experiments), ZSU23/4 anti-aircraft gun #d08 and BRDM2 APC #e71. Regions of interest (ROI) are found in the MSTAR SAR target chips by reducing speckle noise using the Crimmins algorithm in



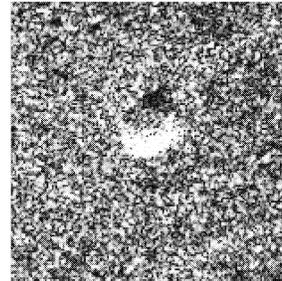
(a)



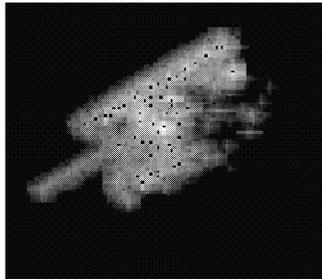
(a)



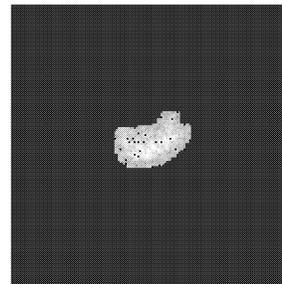
(b)



(b)



(c)



(c)

Fig. 1. Examples of T72 tank CAD model, XPATCH SAR image and scattering center locations superimposed on image for 30° azimuth.

Fig. 2. Example MSTAR target photo, SAR image, ROI (with peaks) for ZSU 23/4 #d08.

Khoros [12], thresholding at the mean plus two standard deviations, dilating to fill small gaps between regions, eroding to have one large ROI and little regions, discarding the small regions with a size filter and dilating to expand the extracted ROI. All the procedural parameters in the ROI extraction algorithm are fixed for the entire MSTAR data used to demonstrate results in this work. The scattering centers are extracted from the SAR magnitude data (within the boundary contour of the ROI) by finding local eight-neighbor maxima. Example photograph, SAR image and the ROI, with the locations of the scattering centers superimposed, are shown in Fig. 2 for ZSU23/4 #d08. The MSTAR data is all at 1 ft resolution and the data is at 15° depression angle (unless otherwise noted).

C. Target Occlusion

There are no real SAR data with occluded objects available to the general public (limited data on vehicles in revetments [14] and partially hidden

behind walls [19] has been reported to exist, but it has not yet been released for unrestricted use). In addition, there is no standard, accepted method for characterizing or simulating occluded targets. Typically occlusion occurs when a tank backs up into a tree line, for example, so that the back end is covered by trees and only the front portion of the tank is visible to the radar. Thus, the bright target becomes a much smaller sized object to the ATR. In addition, the tree tops can produce bright peaks that are of similar strength to target peaks at many azimuths.

The occluded test data here is simulated by starting with a given number of the strongest scattering centers and then removing the appropriate number of scattering centers encountered in order, starting in one of four perpendicular directions d_i (where d_1 and d_3 are the cross range directions, along and opposite the flight path respectively, and d_2 and d_4 are the up range and down range directions). This simulates the spatial correlation of the missing features in occluded objects. Then the same number of scattering centers (with random magnitudes, within the overall range of

target scatterers) are added back at *random locations* within the original bounding box of the chip. This synthetic noise keeps the number of scatterers constant and acts as a surrogate for some potential occluding object.

Our approach, using simulated occlusion provides an enormous amount of test data with varying known amounts of occlusion and additional random noise for carefully controlled experiments, discussed in Section IV for XPATCH data and Section V for MSTAR data. In the cases that are not forced recognition, the test data includes unmodeled confuser vehicles, where the correct system response should be the unknown category. In the occluded articulated object experiments the test data has missing features and additional noise due to articulation as well as the simulated occlusion. XPATCH is used to generate (simulated SAR) test data with actual occlusion and interference from a second object for the experiments reported in Section IVD.

III. SAR RECOGNITION ALGORITHM

The basic SAR recognition algorithm is an off-line model construction process and a similar on-line recognition process. The approach is designed for SAR and is specifically intended to accommodate recognition of occluded and articulated objects. Standard nonarticulated, nonoccluded models of the objects are used to recognize these same objects in nonstandard, articulated and occluded configurations in the presence of noise. The models are a look up table and the recognition process is an efficient search for *positive evidence*, using relative locations of the scattering centers in the test image to access the look-up table and generate votes for the appropriate object (and azimuth pose).

The relative locations and magnitudes of the N strongest SAR scattering centers (local maxima in the radar return signal) are used as characteristic features (where N , the number of scattering centers used, is a design parameter). Because of the specular radar reflections in SAR images, a significant number of features do not typically persist over a few degrees of rotation (in either the XPATCH data [11] or the real MSTAR data [2]). Consequently, we model each object at 1° azimuth increments. Any local reference point, such as a scattering center location, can be chosen as a basis point (or origin) to establish a reference coordinate system for building a model of an object at a specific azimuth angle pose. The relative distance and direction of other scattering centers can be expressed in radar range and cross-range coordinates and naturally tessellated into integer buckets that correspond to the radar range/cross-range bins. For ideal data, picking the location of the strongest scattering center as the basis point is sufficient. However, for potentially corrupted data where any scattering center could be spurious or

missing (due to the effects of noise, target articulation, occlusion, nonstandard target configurations, etc.), we use all N strongest scattering centers in turn as basis points to ensure that a valid basis point is obtained. Thus, to handle occlusion and articulation, the size of the look-up table models (and also the number of relative distances that are considered in the test image during recognition) are increased from N to $N(N-1)/2$. The models are constructed using the relative positions of the scattering centers in the range and cross-range directions as the initial indices to a look-up table of labels that give the associated target type, target pose, basis point range and cross-range positions and the magnitudes of the two scatterers. Since the relative distances are not unique, there can be several of these labels (with different target, pose, etc. values) at each look-up table entry. The outline of the model construction algorithm for an object at a specific azimuth is as follows.

- 1) Obtain location and magnitude of the strongest n scatterers.

- 2) For each scatterer i ($1 \leq i \leq n$, where $i \neq$ origin scatterer), calculate relative range (dR_i), cross-range (dC_i) location from the origin and at model look up table location (dR_i, dC_i) append to list an entry with: object name, azimuth angle, origin range, origin cross-range, origin magnitude, and scatterer i magnitude.

- 3) Repeat Step 2 using other scatterers as the origin.

The recognition process uses the relative locations of the N strongest scattering centers in the test image to access the look-up table and generate votes for the appropriate object, azimuth, range, and cross-range translation. Constraints are applied to limit the allowable *percent difference* in the magnitudes of the data and model scattering centers (difference normalized by model magnitude, all expressed as scaled radar cross-section) to $\pm 9\%$, based on measured probability mass functions of scatterer magnitude invariance with target configuration variants and articulations [2]. (Given the MSTAR targets are centered in the chips, a ± 5 pixel limit on allowable translations is imposed for computational efficiency.) To accommodate some uncertainty in the scattering center locations, the eight-neighbors of the nominal range and cross-range relative location are also probed and the translation results are accumulated for a 3×3 neighborhood in the translation subspace. The recognition process is repeated with different scattering centers as basis points, providing multiple looks at the model database to handle spurious scatterers that arise due to articulation, occlusion, or configuration differences. The recognition algorithm actually makes a total of $9N(N-1)/2$ queries of the look-up table to accumulate evidence for the appropriate target type, azimuth angle, and translation. The models (labels with object, azimuth, etc.)

associated with a specific look-up table entry are the real model and other models that happen by coincidence, to have a scatterer pair with the same (range, cross-range) relative distance. The constraints on magnitude differences filter out many of these false matches. In addition, while these collisions may occur at one relative location, the same random object-azimuth pair doesn't keep showing up at other relative locations with appropriate scatterer magnitudes and mapping to a consistent 3×3 neighborhood in translation space, while the correct object does. The outline of the recognition algorithm is as follows.

- 1) Obtain the location and magnitude of the strongest n scatterers in the test SAR image.
- 2) For each scatterer i ($1 \leq i \leq n$, where $i \neq$ origin scatterer), calculate the relative range (dR_i) and cross-range (dC_i) location from the origin, look up list of entries at (dR_i, dC_i) in the model table, and for each entry in the list, if the translation of the origin and percent magnitude changes of both the scatterer and the origin (test data versus model entry) are all within limits, then increment the vote count of the appropriate object name, azimuth angle and (range, cross-range) translation combination.
- 3) Repeat Step 2 using the 3×3 neighborhood about each nominal (dR_i, dC_i) .
- 4) Repeat Step 2 and Step 3 using other scatterers as the origin.
- 5) Tally the votes for each object, azimuth, translation combination by summing votes in a 3×3 neighborhood in translation space.

The basic decision rule used in the recognition is to select the object-azimuth pair (and associated best translation) with the highest accumulated vote total. To handle identification with unknown objects, we introduce a criteria for the quality of the recognition result that the votes for the potential winning object exceed some minimum threshold v_{\min} . By varying the decision rule threshold we obtain a form of ROC curve with $PCI = P\{\text{decide correct object} \mid \text{object is true}\}$, versus probability of false alarm, $P_f = \{ \text{decide any object} \mid \text{unknown is true} \}$. We call the algorithm a 6D recognition algorithm since, in effect, we use the range and cross-range positions and the magnitudes of pairs of scattering centers. (When using 40 scatterers, this 6D algorithm takes an average of 2.5 s to process a test chip on a Sun Ultra2 without any optimizations.)

IV. XPATCH RESULTS

Results with XPATCH data are based on using a 2D recognition algorithm that is an earlier, simpler version of the 6D algorithm described above. The 2D algorithm [11] uses only the relative range and cross-range distances; it does not compute the translation; it only considers the exact scatterer

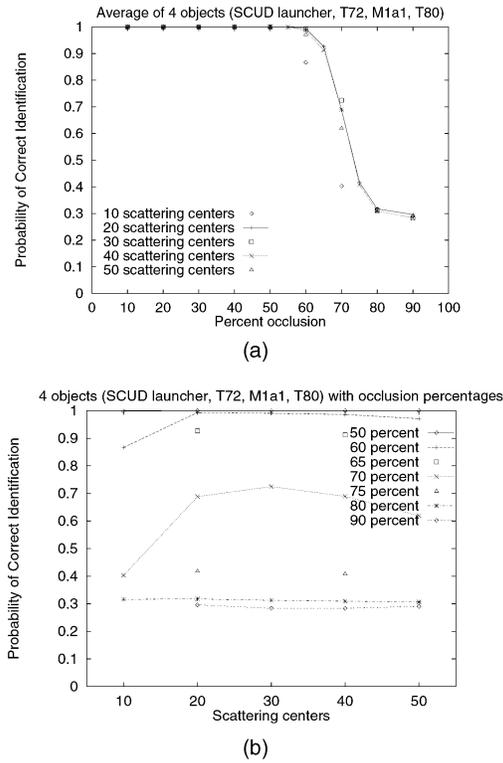


Fig. 3. Effect of XPATCH occlusion and number of scatterers on recognition rate. (a) Effect of occlusion. (b) Effect of number of scatterers.

location; and it does not use the magnitude information.

A. XPATCH Forced Recognition Performance

The performance of the 2D recognition algorithm with nonarticulated objects that are occluded (using the approach outlined in Section IIC) is shown in Fig. 3(a) in terms of the PCI as a function of percent occlusion with the number of scattering centers used as a parameter. The results of 288,000 test cases are shown as 50 points (10 to 50 scatterers in 10 scatterer steps for 10 to 90% occlusion in 10% steps, plus 20 and 40 scatterers at 55, 65 and 75% occlusion, minus 10 scatterers at 90% occlusion is 50 points), where each point for a specific percent occlusion and number of scattering centers is the average PCI for all 4 occlusion directions, the 4 objects and the 360 azimuths (5760 test cases). The overall 2D recognition algorithm performance is almost perfect for up to 60% occlusion (for 40 scatterers with 60% occlusion the PCI is 0.986). By 80 to 90% occlusion, the results are not much better than the 0.25 PCI one would expect by chance from the 4 possible objects. These performance results are replotted as Fig. 3(b) to illustrate the effect of the number of scattering centers used on the recognition rate for the highly occluded cases. This indicates that optimal performance is in the range of 20 to 40 scattering centers.

B. Performance Analysis

The performance of the 2D recognition algorithm can be expressed in terms of the number of votes received for the true case (the actual object, azimuth used in the test instance) and the highest number of votes received for any random case. The true case will have n matching points (scattering centers with the same range, cross-range location) of M points considered (where $n \leq M$ due to occlusion). The number of votes for the true case, V_t , is given by

$$V_t = \frac{n(n-1)}{2} + p \cdot \left[\frac{M(M-1) - n(n-1)}{2} \right] \quad (1)$$

where p is the probability that an object, azimuth instance has an entry at a random location in the relative distance look-up table. The first term in (1) reflects distances with both end points (scattering centers) matching, while in the remaining term, one or both of the end points are not matching.

The average number of votes received for the true test object, the predicted number of votes from the first term of (1), the average votes for the highest ranking other object and the average votes for the incorrect winning object when recognition fails are shown in Fig. 4(a) as a function of the percent valid (or unoccluded) data for a 4-object case (SCUD launcher, T72, M1a1, and T80) with 40 scattering centers. Each point in Fig. 4(a) for the test object and highest other object is the average of 5760 test cases (4 objects, 360 azimuths, 4 occlusion directions). The true object receives more votes than predicted by the first term of (1) because of the random contributions from the nonmatching points. These contributions are predicted by the second term of (1) to range from 11.3 votes at 10 percent valid down to 2.2 votes at 90 percent valid. These predictions are much lower than the actual results, because the relative distance table density distribution is very nonuniform with short distances being much more common than longer distances.

The average votes received for the true object is shown in Fig. 4(b) for the four test objects, and the average of the highest votes for another object is shown in Fig. 4(c) with the true object as a parameter. Note that the two larger objects (in our case the SCUD launcher and T72 tank) have fewer than average votes (see Fig. 4(b)) and result in the highest other object receiving fewer than average votes (Fig. 4(c)). In generating the occluded test data, the occluded points are placed back within the original bounding box of the object SAR image, so the occluded test cases for the larger objects are relatively sparse and have more long distances. Because the larger objects have fewer of the more common short distances, they have a smaller contribution from random coincidences and the number of votes is thus smaller and closer to the number predicted solely

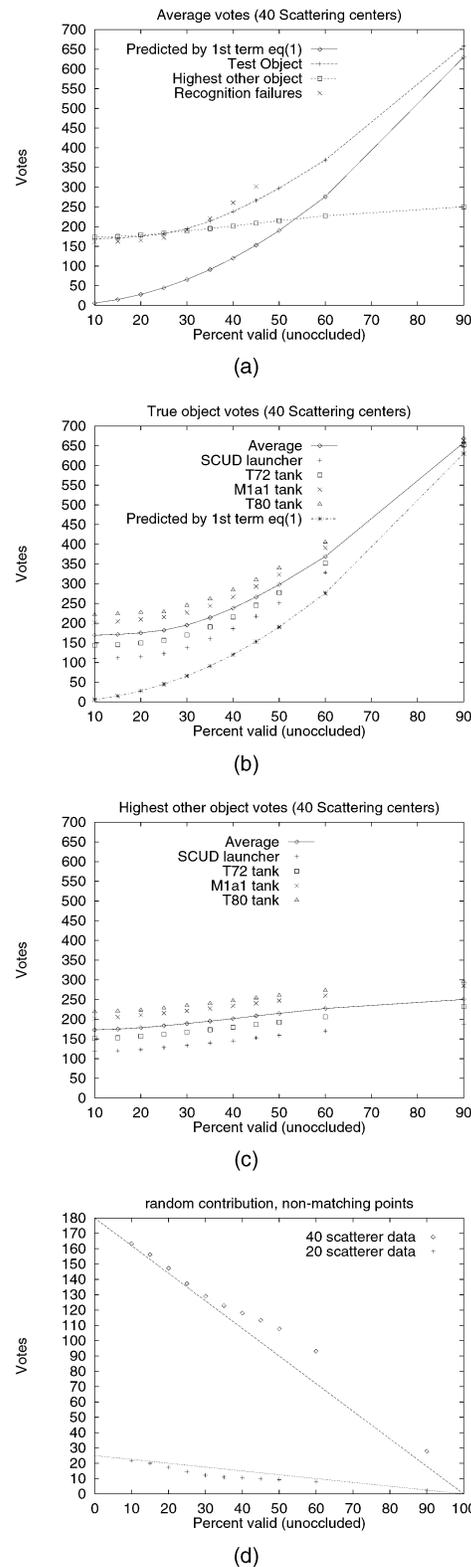


Fig. 4. Occluded performance analysis for XPATCH data. (a) Average votes. (b) True object votes. (c) Highest other object votes. (d) Random contribution to winning total votes.

from the valid points (the first term of (1)). Similarly, because of the relatively fewer smaller distances with the larger object test cases, the smaller other objects do not collect many votes.

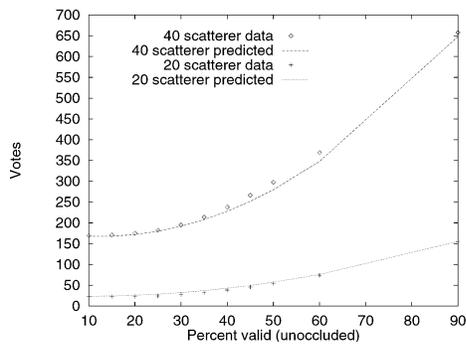


Fig. 5. Performance prediction by (2).

If we call the random contribution of the nonmatching points V_n , then (1) for the number of votes in the true case can be written as $V_t = n(n-1)/2 + V_n$. In Fig. 4(a), V_n is the Test object curve minus the Predicted curve. Fig. 4(d) shows V_n as a function of percent of valid points n/M for $M = 20$ and 40 scattering centers. We can approximate V_n , as shown in Fig. 4(d), by a linearly decreasing function of n with a random contribution $V_n = Q$ when $n = 0$ (no matching points) and $V_n = 0$ when $n = M$ (all scattering centers match). Thus, we obtain $V_n = Q(1 - n/M)$. If we assume that Q is given by a simple polynomial of the form $Q = aM^b$, based on Fig. 4(d) with $Q(40) = 180$ and $Q(20) = 25$, we derive the coefficient values $a = 4.9 \times 10^{-3}$ and $b = 2.85$. Thus we obtain

$$V_t = n(n-1)/2 + aM^b(1 - n/M). \quad (2)$$

Using (2) to predict the total votes, the predicted and actual votes for the 2D recognition algorithm with 20 and 40 scattering center models are shown in Fig. 5.

C. XPATCH Decision Rule, ROC Curve and Unknown Object Results

The decision rule used to determine the recognition algorithm result is a design parameter that can be varied to optimize the algorithm performance. Using a vote ratio decision rule (i.e., the ratio of votes for the potential winning object v_1 to the votes for the second place different object v_2 must be greater than some minimum ratio r), results are obtained for a 4-object case (SCUD launcher, T72, M1a1, and T80) with test data that are occluded versions of those four objects and a similarly occluded FRED tank as an unknown object. Fig. 6 gives the ROC curves (PCI versus P_f) of the 2D recognition algorithm with 40 scattering centers for 50%–70% occlusion using vote ratio decision rule. (The upper-right point on each ROC curve, maximum PCI, and maximum P_f , is $r = 1.0$ which corresponds to the forced recognition case).

Fig. 7(a), which plots P_f versus r for 40 scatterers with 50%–70% occlusion, shows that for a given vote ratio, the false alarm rate is largely independent of the

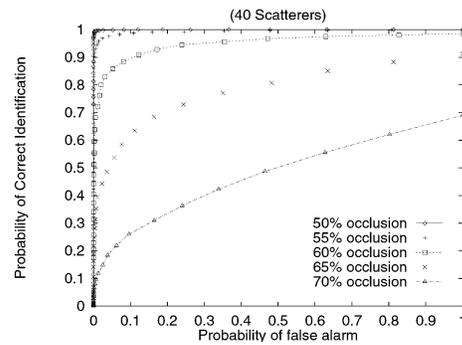
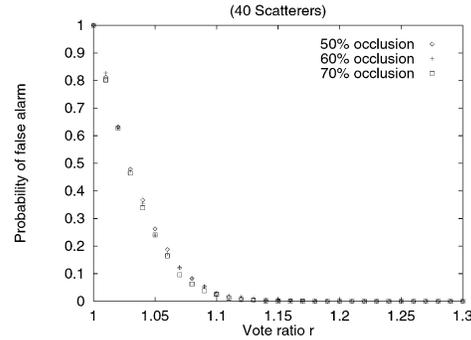
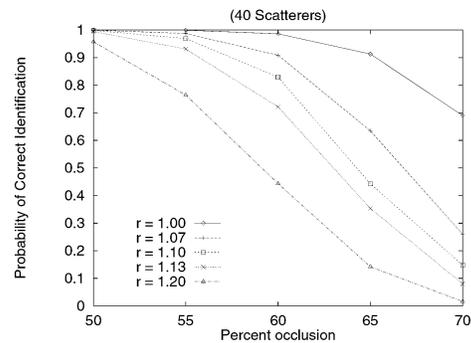


Fig. 6. ROC.



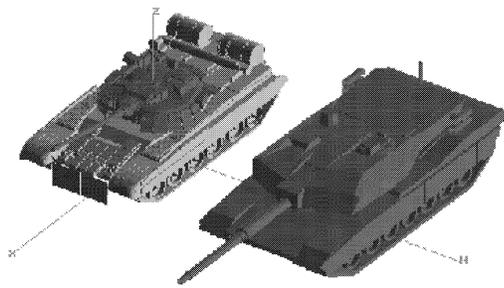
(a)



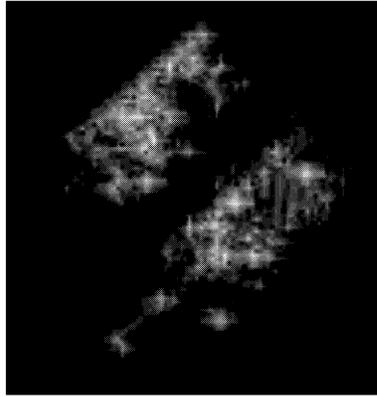
(b)

Fig. 7. Effect of vote ratio. (a) False alarm rate. (b) Recognition rate.

occlusion of the unknown object. Thus, the desired maximum false alarm rate (such as 10%) dictates a minimum vote ratio (at least 1.08 for 40 scatterers). Higher vote ratios, such as 1.10 for 40 scatterers, reduce the false alarm rate to about 3%. However, increasing the vote ratio r reduces the recognition rate, PCI, for highly occluded objects. Fig. 7(b) shows how the recognition rate degrades more quickly with occlusion for higher values of vote ratio. Thus, the optimum vote ratio is the smallest value that will meet the desired false alarm rate. For 60% occlusion with $r = 1.1$ (40 scatterers) the PCI is 0.829 and P_f is 0.031, with only one target misidentification (and all other target misses are classified as unknown). Previously, for the same (60% occlusion, 40 scatterer) case, but with forced recognition, $r = 1.0$, there are 80 misidentifications for a PCI of 0.986. This is typical because as vote ratio increases, not only do some of the weaker identifications move into the unknown



(a)



(b)

Fig. 8. Example of CAD models and XPATCH SAR image (40° azimuth) for T80 (left) and FRED (right) tanks together.

column, but also many of the misidentifications become unknowns.

D. Occlusion and Noise from an Unknown Second XPATCH Object

Another method for generating occluded target data with additional noise is to introduce a competing second object that is unknown. The CAD model of the FRED tank is positioned parallel to the T80 tank (arbitrarily 83 in apart) and 360 6 in resolution SAR images are generated with XPATCH (in steps of 1° azimuth). An example of the geometry and a SAR image are shown in Fig. 8. In this case FRED is not another target, recognizing closely spaced targets is a different research problem, rather FRED is used to provide a measurable “occlusion” environment. The close presence of FRED provides blockage, strong spurious scatterers and interaction effects. Recognition results (for 40 scattering centers and a vote ratio of 1.1) are shown as a confusion matrix in Table I with an overall PCI of 97.8%. The PCI is 99.4% for cases with the T80 in front (0°–179° azimuth) and 96.1% for cases with the FRED tank in front (180°–359° azimuth). Since the T80 CAD model location is fixed, the strongest 40 scattering center locations from each image with both (T80 and FRED) tanks together can be compared with the strongest 40 scatterers from the image at the same azimuth with the T80 tank alone

TABLE I
Confusion Matrix for XPATCH T80 and FRED Tanks Together

Test Targets	Identification Results				
	SCUD	T72	M1a1	T80	Unknown
T80 and FRED together	0	0	1	352	7

Note: 40 scatterers, $r = 1.1$.

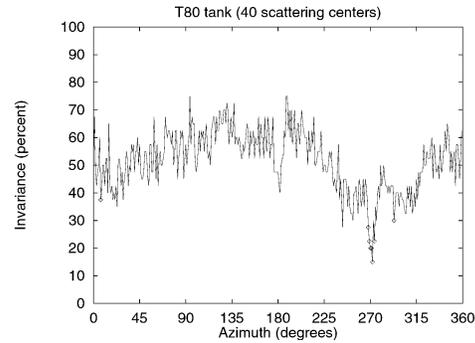


Fig. 9. XPATCH T80 tank recognition failure plot (◊) on invariance curve for T80 and FRED tanks together.

to determine the percent invariance at that azimuth. Fig. 9 shows this T80 invariance (with 40 scatterers) as a function of azimuth and plots, as diamonds (◊), the recognition failures (7 where the presence of the FRED tank causes the answer to be unknown and 1 misidentification). From Fig. 9 it can be seen that most of the recognition failures occur near 270° azimuth, where the FRED tank is broadside to the radar and directly in front of the T80 tank.

E. Occluded Articulated XPATCH Objects

The occluded articulated data is produced in the same manner as the nonarticulated occluded data (see Section IIC). The same tanks are used, but with a 90° turret rotation and the missile is erect on the SCUD launcher. There are 9 occluded articulated data sets (for 20 scattering centers with 10 to 90% occlusion), each with 5760 test cases as before, for a total of 51,840 test cases. Fig. 10 shows the average and individual test object performance of the 2D recognition algorithm (using 20 scattering centers) as a function of percent occlusion with 4 different articulated objects. The results of 51,840 test cases are shown, for example, the overall performance for these articulated objects with 30% occlusion is a 0.698 PCI.

V. MSTAR RESULTS

A. MSTAR Forced Recognition Performance

The performance of the 6D recognition algorithm with MSTAR objects that are occluded (as outlined

TABLE II
Forced Recognition Confusion Matrix for 70% Occluded MSTAR Objects

70% Occluded Test Targets:	Identification Results				Pose Accuracy			
	BMP	BTR	T72	ZSU	BMP	BTR	T72	ZSU
BMP	769	1	2	0	768c, 768e	773c, 770e	767c, 765e	1084c, 1072e
BTR	1	774	1	0				
T72	3	0	773	0				
ZSU	1	1	0	1094				

Note: 40 scatterers, e = exact pose, c = pose within $\pm 5^\circ$.

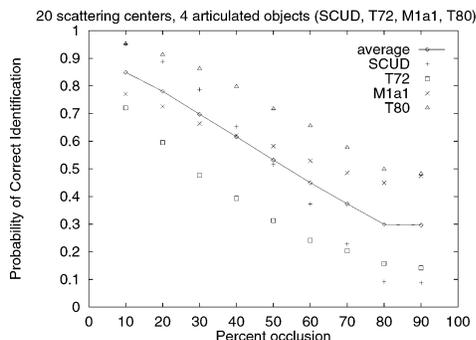
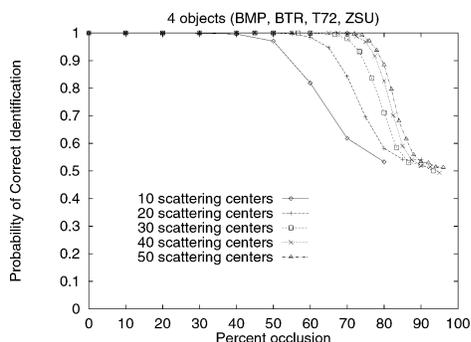
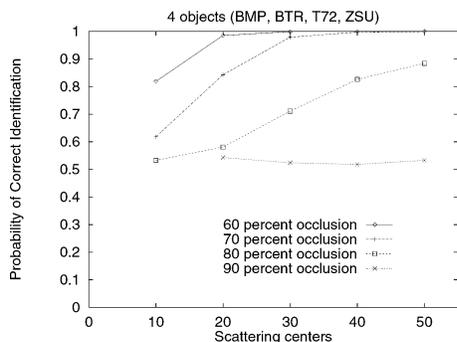


Fig. 10. Effect of occlusion on articulated object recognition rate using XPATCH data.



(a)



(b)

Fig. 11. Effect of occlusion and number of scatterers on recognition rate using MSTAR data. (a) Effect of occlusion. (b) Effect of number of scatterers.

in Section IIC) is shown in Fig. 11 in terms of PCI as a function of percent occlusion with the number of scattering centers used as a parameter. Each point

for a specific number of scattering centers and percent occlusion is the average PCI for all four occlusion directions, the 4 objects (BMP, BTR, T72, and ZSU) and the number of available test azimuths. We defined the available test azimuths as azimuths that had at least the number of scattering centers used present in the data, thus we avoid introducing an uncontrolled variable: the number of scattering centers actually available for some instance of an object at a specific azimuth orientation. (In practice, if some target aspect did not have the appropriate number of scattering centers, the performance would degrade as if the missing scatterers were occluded.) The forced recognition results for the MSTAR data in Fig. 11 are comparable to the XPATCH results of Fig. 3. For the 6D algorithm and the MSTAR data the breakpoint is at 60–75% occlusion for 20 scatterers or more, compared with 55–60% for the 2D algorithm with XPATCH data. In addition the 6D algorithm optimizes at a higher number of scattering centers used than the 2D version. At very high occlusion levels one would expect that the recognition results with four objects would approach 25%, due to chance. With the MSTAR data we achieve 50% recognition, because the ZSU is almost never confused with the other vehicles, so the three remaining vehicles at a little over 33% and the ZSU at over 90% yields an overall rate of about 50% recognition.

Typical forced recognition results for 40 scattering centers and 70% occlusion are shown as a confusion matrix in the left half of Table II. With 3410 correct identifications in 3420 trials, the overall PCI is 0.9971. The right half of Table II shows the pose accuracy results, where 99.18% of the time the pose is correct within $\pm 5^\circ$, in 98.68% of the cases the pose is exactly correct.

B. MSTAR Performance Analysis

For the 6D recognition algorithm, the additional constraints on the range and cross-range translations and the percent magnitude changes for the scatterers significantly reduce the number of random matches. In the case of occlusion, for the true case (the actual object, azimuth used in the test instance) there are n valid scatterers of M scatterers used (where $n \leq M$ due to occlusion) and, neglecting any random contribution

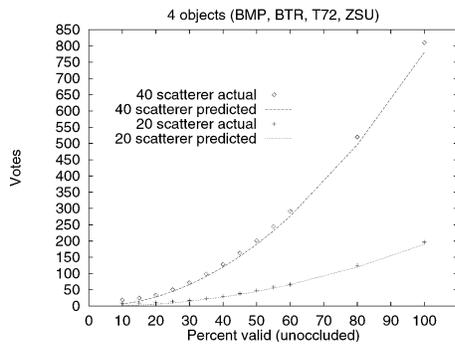


Fig. 12. MSTAR occluded performance prediction.

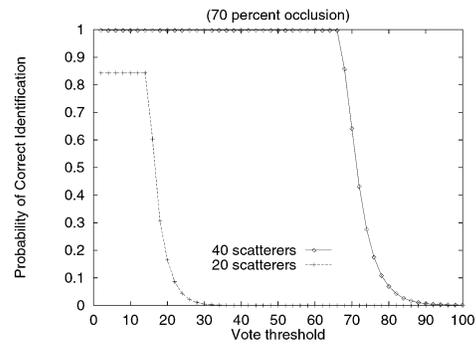
of nonmatching points, the number of votes for the true case V is given by $V = n(n - 1)/2$. Fig. 12 shows that the actual number of votes received lies just slightly above the prediction curve and that the random contributions of nonmatching points are negligible. These results for the 6D algorithm are in contrast to the results for the earlier 2D algorithm, Fig. 4(b), where there was a significant random contribution from the nonmatching points. Thus, the additional features and constraints used in the 6D recognition are quite effective.

C. MSTAR Decision Rule, ROC Curve and Unknown Object Results

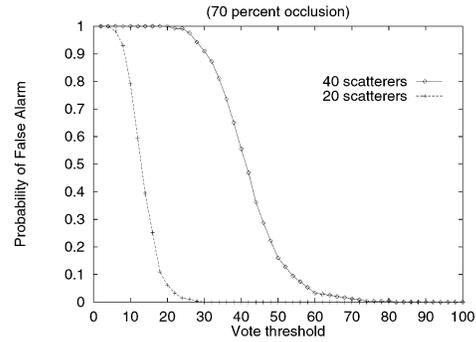
The vote ratio decision rule was not used with the MSTAR data, because the BRDM2 confuser vehicle would consistently be identified as a T72 with a high vote ratio, even though the total number of votes was quite low (compared with what a T72 test case would get). So, using a vote threshold decision rule (i.e., the votes for the potential winning object exceed some threshold, v_{\min}), recognition results were obtained with the MSTAR data for occluded versions of the BMP, BTR, T72, and ZSU test vehicles as well as a similarly occluded BRDM2 confuser vehicle. Figs. 13(a), 13(b), and 13(c) give the PCI, probability of false alarm, and probability of miss, $P_m = P\{\text{decide unknown} \mid \text{object is true}\}$, respectively, as a function of v_{\min} for 20 and 40 scatterers and 70% occlusion. The resulting ROC curves for 20 and 40 scatterers and 70% occlusion are shown in Fig. 13(d).

Fig. 14 gives the ROC curves for 40 scatterers with 65%–80% occlusion and is comparable to Fig. 6. An illustrative confusion matrix for 70% occlusion and $v_{\min} = 65$ (40 scatterers) is shown in Table III. The overall PCI is 0.997 and the P_f is 0.025. The results in Table III are significantly better than the comparable 2D algorithm XPATCH results in Section IVC (which are: PCI = 0.829, $P_f = 0.031$ at 60% occlusion).

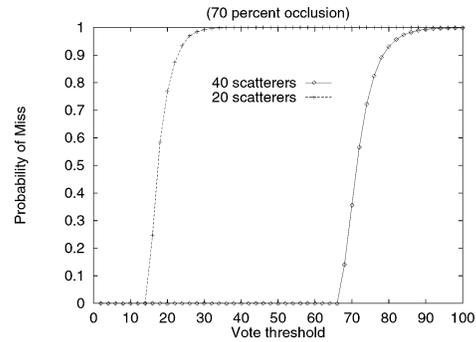
Fig. 15 compares the mean and standard deviation of the votes generated by the test objects with the votes generated by the BRDM2 confuser vehicle for 40 scatterers as a function of the percentage



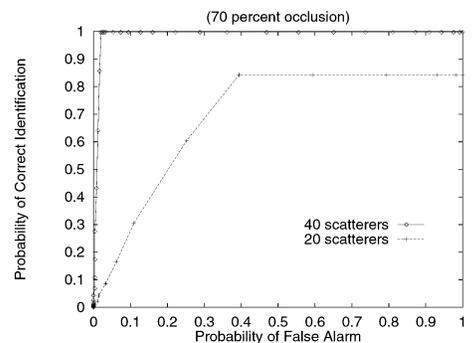
(a)



(b)



(c)



(d)

Fig. 13. MSTAR vote threshold and ROC (70% occlusion). (a) Probability of correct identification. (b) Probability of false alarm. (c) Probability of miss. (d) ROC.

of valid (unoccluded) scatterers. This shows that, with 40 scatterers, for above 30% valid data (or less than 70% occlusion) the occluded BRDM2 is not in competition with the actual object. However, while

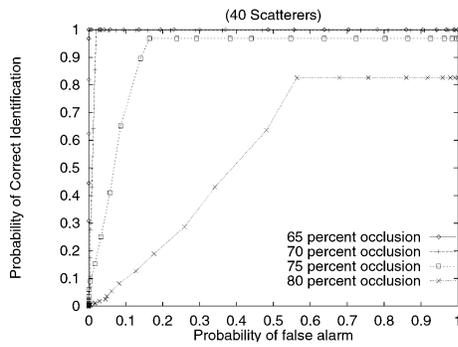


Fig. 14. Effect of occlusion on MSTAR ROC (40 scatterers).

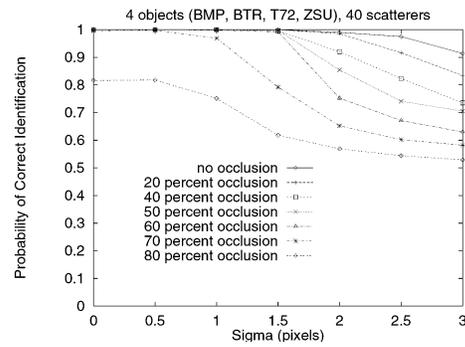


Fig. 16. Effect of positional noise on MSTAR occluded object recognition.

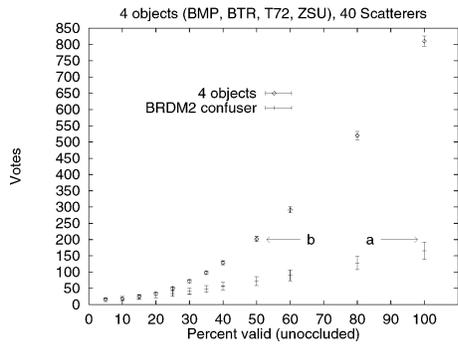


Fig. 15. MSTAR occluded performance with unknown object.

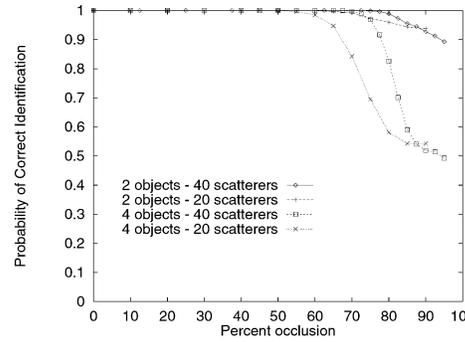


Fig. 17. Effect of scaling on MSTAR occluded object recognition.

TABLE III

Typical Confusion Matrix for 70% Occluded MSTAR Objects

70% Occluded Test Targets:	Identification Results				
	BMP	BTR	T72	ZSU	Unknown
BMP	769	1	2	0	0
BTR	1	774	1	0	0
T72	3	0	773	0	0
ZSU	1	1	0	1094	0
BRDM2	2	3	0	1	237

the target may be occluded, the confuser vehicle may not necessarily be occluded in the practical case. Hence to cope with *unoccluded confusers*, one would need to set a threshold of about 200 votes for a valid identification (labeled a in Fig. 15), which would then limit the ability to recognize targets to about 50% target occlusion (b in Fig. 15).

D. MSTAR Occlusion and Positional Noise

Test data with positional noise are generated by adding Gaussian noise with zero mean and standard deviation sigma (in units of 1 ft resolution pixels) to the range and cross-range locations of the scattering centers. The overall recognition performance for four objects, using 40 scatterers with varying amounts of occlusion, is shown in Fig. 16 as a function of positional noise. Fig. 16 confirms that the objective of the 6D recognition algorithm to accommodate a

one pixel uncertainty in scattering center location has been achieved for up to 70 percent occlusion.

E. MSTAR Occlusion and Scaling

The previous MSTAR occlusion experiments all involved recognizing four objects (BMP, BTR, T72, and ZSU). Because there are only two articulated objects available (T72 and ZSU), it is useful to establish the effect of scaling the forced recognition problem from four occluded objects to two occluded objects prior to investigating the effect of occlusion on articulated objects. Fig. 17 illustrates the effect of scaling on occluded object recognition. The 20 and 40 scatterer curves for four objects are the same data as previously shown in Fig. 11(a). As one would expect, the results for the two-object case are better than the four-object case: the break point is less pronounced and the PCI is higher at very high occlusion levels.

F. Occluded Articulated MSTAR Objects

In the articulated object experiments the models are nonarticulated versions of T72 #a64 and ZSU23/4 #d08 and the test data are the articulated versions of these same serial number objects (with the turret rotated to 315°) that are occluded in the same manner as before. The MSTAR articulated data is at 30 deg depression angle. Fig. 18 shows the effect of occlusion on recognition of these MSTAR articulated

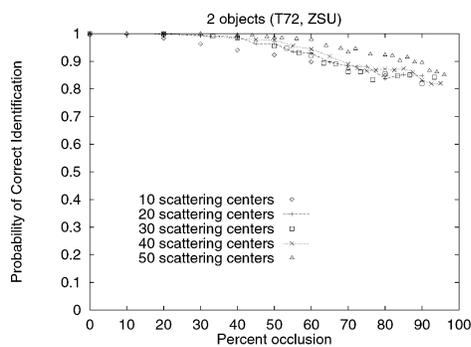


Fig. 18. Effect of occlusion on MSTAR articulated object recognition.

objects for various numbers of scattering centers used. The 6D algorithm results on MSTAR data in Fig. 18 are less sensitive to occlusion than the 2D algorithm results on XPATCH data, shown in Fig. 10, although the excellent results at higher occlusion levels are due to the fact that only two articulated objects are available in the MSTAR data.

VI. CONCLUSIONS

Local features provide a successful approach to recognizing highly occluded objects (with 50% or more occlusion). The basic 2D algorithm, while sufficient for the simulated XPATCH SAR data, had a significant level of potential false matches that are nearly eliminated by the introduction of additional features and constraints in the 6D algorithm. The techniques introduced in the 6D algorithm used both location and magnitude of scattering centers as features and successfully accommodated uncertainty in these features. The possibility of an *unoccluded confuser* vehicle is an important practical limiting factor on the performance that can be achieved in recognizing highly occluded vehicle targets (e.g., 50% occlusion versus over 70% with occluded confusers). These algorithms achieve excellent occluded object recognition results for simulated and real SAR data. In addition to the significant results for occlusion alone, the local features approach also produced good results for the combined effects of occlusion and articulation.

GRINNELL JONES, III
BIR BHANU
 Center for Research in Intelligent Systems
 University of California at Riverside
 B232 Bourns Hall
 Riverside, CA 92521-0425

REFERENCES

[1] Andersch, D., Lee, S., Ling, H., and Yu, C. (1994) XPATCH: A high frequency electromagnetic scattering prediction code using shooting and bouncing ray. In *Proceedings of Ground Target Modeling and Validation Conference*, Aug. 1994, 498–507.

[2] Bhanu, B., and Jones, G. (1999) Recognizing MSTAR target variants and articulations. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery VI*, **3721** (Apr. 1999), 507–519.

[3] Bhanu, B., Dudgeon, D., Zelnio, E., Rosenfeld, A., Casasent, D., and Reed, I. (1997) Introduction to the special issue on automatic target detection and recognition. *IEEE Transactions on Image Processing*, **6**, 1 (Jan. 1997), 1–6.

[4] Carlson, D., Kumar, B., Mitchell, R., and Hoffelder, M. (1997) Optimal trade-off distance classifier correlation filters (OTDCCFs) for synthetic aperture radar automatic target recognition. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery IV*, **3070** (Apr. 1997), 110–120.

[5] Casasent, D., and Shenoy, R. (1997) Synthetic aperture radar detection and clutter rejection MINACE filters. *Pattern Recognition*, **30**, 1 (Jan. 1997), 151–162.

[6] Casasent, D., and Ashizawa, S. (1997) Synthetic aperture radar detection, recognition and clutter rejection with new minimum noise and correlation energy filters. *Optical Engineering*, **36** (Oct. 1997), 2729–2738.

[7] Casasent, D., and Shenoy, R. (1997) Feature space trajectory for distorted-object classification and pose estimation in SAR. *Optical Engineering*, **36** (Oct. 1997), 2719–2728.

[8] Dudgeon, D., and Lacoss, R. (1993) An overview of automatic target recognition. *The Lincoln Laboratory Journal*, **6**, 1 (1993), 3–9.

[9] Ikeuchi, K., Shakunga, T., Wheeler, M., and Yamazaki, T. (1996) Invariant histograms and deformable template matching for SAR target recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 1996, 100–105.

[10] Jachimczyk, W., and Cyganski, D. (1997) Enhancements of pose-tagged partial evidence fusion SAR ATR. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery IV*, **3070** (Apr. 1997), 334–345.

[11] Jones, G., and Bhanu, B. (1999) Recognition of articulated and occluded objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**, 7 (July 1999), 603–613.

[12] *Khoros Pro v2.2 User's Guide* (1998) Addison Wesley Longman Inc., 1998.

[13] Meth, R., and Chellappa, R. (1996) Automatic classification of targets in synthetic aperture radar imagery using topographic features. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery III*, **2757** (Apr. 1996), 186–193.

[14] Mossing, J., and Ross, T. (1998) An evaluation of SAR ATR algorithm performance sensitivity to MSTAR extended operating conditions. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery V*, **3370** (Apr. 1998), 554–565.

[15] Ross, T., Worrell, S., Velten, V., Mossing, J., and Bryant, M. (1998) Standard SAR ATR evaluation experiments using the MSTAR public release data set. *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery V*, **3370** (Apr. 1998), 566–573.

- [16] Ryan, T., and Egaas, B. (1996)
SAR target indexing with hierarchical distance transforms.
SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery III, **2757** (Apr. 1996), 243–252.
- [17] Verly, J., Delanoy, R., and Lazott, C. (1993)
Principles and evaluation of an automatic target recognition system for synthetic aperture radar imagery based on the use of functional templates.
SPIE Proceedings: Automatic Object Recognition III, **1960** (Apr. 1993), 57–71.
- [18] Waxman, A., Seibert, M., Bernardon, A., and Fay, D. (1993)
Neural systems for automatic target learning and recognition.
The Lincoln Laboratory Journal, **6**, 1 (1993), 77–116.
- [19] Wissinger, J., Ristroph, R., Diemunsch, J., Severson, W., and Freudenthal, E. (1999)
MSTAR’s extensible search engine and inferencing toolkit.
SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery VI, **3721** (Apr. 1999), 554–570.
- [20] Yi, J. H., Bhanu, B., and Li, M. (1996)
Target indexing in SAR images using scattering centers and the Hausdorff distance.
Pattern Recognition Letters, **17** (1996), 1191–1198.

Errata: Fast Converging Adaptive Processor for a Structured Covariance Matrix¹

1) The title of this contribution should read as above; inadvertently **for** was printed as **or** in the original.

2) On page 1116: In the last line of the first paragraph in the second column, which now reads “nonsingular **SMC** [12].” should read “nonsingular **SCM** [12].”

3) On page 1123: The first line of V. SUMMARY should read “A fast converging adaptive processor has been —.”

The Editors apologize for these errors and any inconveniences.

¹ Steiner, M., and Gerlach, K., *IEEE Transactions on Aerospace and Electronic Systems*, **36**, 4 (Oct. 2000), 1115–1126.

Manuscript received December 18, 2000.

IEEE Log No. T-AES/37/1/02943.

Author’s address: Naval Research Laboratories, Code 5341, 4555 Overlook Ave. SW, Washington, DC 20325-5320.

0018-9251/01/\$10.00 2001 IEEE

U.S. Government work not protected by U.S. copyright.