International Journal of
IMAGE AND GRAPHICS

Volume 4 • Number 1 • January 2004

Multiple Look Angle SAR Recognition

B. Bhanu and G. Jones III

World Scientific
NEW JERSEY • LONDON • SINGAPORE • SHANGHAI • HONG KONG • TAIPEI • CHENNAI

International Journal of Image and Graphics
Vol. 4, No. 1 (2004) 85–98

World Scientific
www.worldscientific.com

MULTIPLE LOOK ANGLE SAR RECOGNITION

BIR BHANU* and GRINNELL JONES III†
Center for Research in Intelligent Systems
University of California, Riverside, California 92521, USA
**bhanu@cris.ucr.edu*
†*grinnell@cris.ucr.edu*

Received 30 May 2002
Revised 4 March 2003

The focus of this paper is optimizing the recognition of vehicles in Synthetic Aperture Radar (SAR) imagery by exploiting the azimuthal variance of scatterers using multiple SAR recognizers at different look angles. The variance of SAR scattering center locations with target azimuth leads to recognition system results at different azimuths that are independent, even for small azimuth deltas. Extensive experimental recognition results are presented in terms of receiver operating characteristic (ROC) curves to show the effects of multiple look angles on recognition performance for MSTAR vehicle targets with configuration variants, articulation, and occlusion.

Keywords: Automatic target recognition; synthetic aperture radar; multiple recognizers; recognizing occluded objects; articulated object recognition; recognizing configuration variants.

1. Introduction

In this paper we are concerned with optimizing the recognition of vehicles in Synthetic Aperture Radar (SAR) imagery by using multiple SAR recognizers at different look angles. Methods to acquire multi-look SAR data would currently involve using multiple platforms or multiple passes, however, one can also envision the future potential of multiple sensors on a single platform. In this research, the recognition system starts with real SAR chips (at one foot resolution) of actual military vehicles from the MSTAR public data[1] and ends with the identification of a specific vehicle type (e.g. a T72 tank). A major challenge is that the vehicles can have significant external configuration variants (fuel barrels, searchlights, etc.), they can be in articulated configurations (such as a tank with its turret rotated), or they can be partially occluded. SAR recognition results for our basic approach are compared (in Ref. 2) to other different approaches using real SAR images from the MSTAR public data. This paper extends our previous work on SAR recognition[2–5] to multiple recognizers at different look angles. It takes a principled approach to multiple look angle SAR recognition and answers two fundamental questions:

(1) what should the next view angle be? and (2) how to integrate the results? We know of no prior published work on SAR recognition at multiple look angles.

The key contributions of this paper are:

- Uses the fundamental azimuthal variance of SAR scatterer locations as the basis for a principled approach to multiple look angle SAR recognition.
- Demonstrates that SAR recognition results at different azimuths are *independent*, even for small azimuths, such as one degree.
- Shows that decision level fusion of two observations at different look angles can substantially increase SAR recognition performance.
- Shows that these results apply to target configuration variants, articulated objects and occluded objects.

The remainder of this paper is organized as follows: the next section demonstrates the azimuthal variance of SAR scattering center locations; Sec. 3 gives a description of the SAR recognition system; Sec. 4 demonstrates that multiple look angle SAR recognition results are independent; Sec. 5 gives the experimental results for multiple look angle SAR recognition with configuration variant, articulated and occluded objects; and finally Sec. 6 has the conclusions and suggests future work.

2. Azimuthal Variance of SAR Scatterers

The typical detailed edge and straight line features of man-made objects in the visual realm do not have good counterparts in SAR images of sub-components of vehicle-sized objects at one foot resolution. However, there are a wealth of peaks in the radar return corresponding to scattering centers. The relative locations of SAR scattering centers, determined from local peaks in the radar, are related to the aspect and physical geometry of the object, independent of translation and serve as distinguishing features. In addition to the scatterer locations, the magnitudes of the peaks are also features that we use for recognition.

Figure 1 shows an example of a typical T72 tank SAR target image (at 66° azimuth) and the associated region of interest (ROI) with the locations of the SAR scattering centers superimposed. The typical rigid body rotational transformations for viewing objects in the visual world do not apply much for the specular radar reflections of SAR images, where we use the locations of SAR scattering centers (determined from local peaks in the radar return) as features. This is because, as we have previously shown, a *significant* number of features *do not* typically persist over a few degrees of rotation.[2]

Because the radar depression angle is generally known, the significant unknown target rotation is (360°) in azimuth. Azimuth persistence or invariance can be expressed in terms of the percentage of scattering center locations that are unchanged over a certain span of azimuth angles. It can be measured (for some base azimuth θ_0) by rotating the pixel locations of the scattering centers from an image at azimuth θ_0 by an angle $\Delta\theta$ and comparing the resulting range and cross-range locations

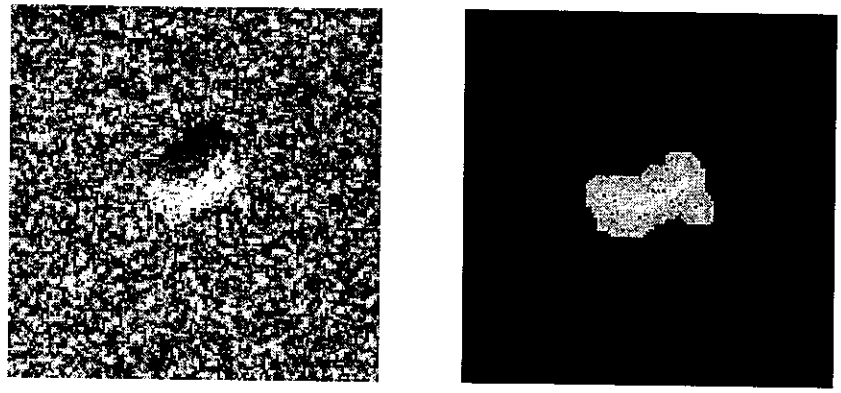

(a) T72: image. (b) T72: ROI.

Fig. 1. Example MSTAR SAR image and ROI (with peaks) for a T72 tank at 66° azimuth.

with the scatterer locations from an image of the same object at azimuth $\theta_0 + \Delta\theta$. More precisely, because the images are in the radar slant plane, we actually project from the slant plane to the ground plane, rotate in the ground plane and project back to the slant plane. Because the objects in the chips are not registered, we calculate the azimuth invariance as the maximum number of corresponding scattering centers (whose locations match within a given tolerance) for the optimum integer pixel translation. This method of registration by finding the translation that yields the maximum number of correspondences has the limitation that for very small or no actual invariance it may find some false correspondences and report a slightly higher invariance than in fact exists. To determine scattering center locations that persist over a span of angles, there is an additional constraint that for a matching scattering center to "persist" at the kth span $\Delta\theta_k$, it must have been a persistent scattering center at all smaller spans $\Delta\theta_j$, where $0 \leq j < k$. Averaging the results of these persistent scattering center locations over 360 base azimuths gives the mean azimuth invariance of the object.

Figure 2 shows an example of the mean scatterer location invariance (for the 40 strongest scatterers) as a function of azimuth angle span using T72 tank serial number (#)132, with various definitions of persistence. In the "exact match" cases the centroid of the rotated scatterer pixel from the image at θ_0 azimuth is within the pixel boundaries of a corresponding scatterer in the image at $\theta_0 + \Delta\theta$. In the "within 1 pixel" cases, the scatterer location is allowed to move into one of the eight adjacent pixel locations. Note that for a 1° azimuth span, although only 20% of the scatterer locations are invariant for an "exact match", 65% of the scatterer locations are invariant "within 1 pixel". The cases labeled "persists" in Fig. 2 enforce the constraint that the scatterer exist for the entire span of angles and very few scatterers continuously persist for even 5°. In the upper two cases (not labeled "persists") scintillation is allowed and the location invariance declines slowly with

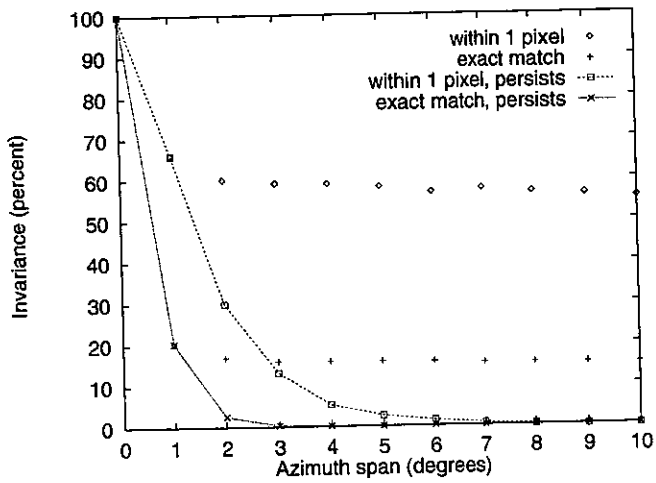

Fig. 2. Scatterer location persistence, T72 #132.

azimuth span. The "within 1 pixel" results (which allow scintillation) are consistent with the one foot ISAR results of Dudgeon,[6] whose definition of persistence allowed scintillation. Because of the higher scatterer location invariance with 1° azimuth span, in our research we use 360 azimuth models for each target, in contrast to others who have used 5°,[7] 10°,[8] and 12 models.[9]

The fact that the SAR scatterer locations *do not* persist over a span of even a few degrees, demonstrated in Fig. 2, strongly indicates that observations at different azimuth angles are independent. Thus, what had previously been viewed as a "problem" for modeling, now presents a significant opportunity for integrating the results of SAR observations at multiple look angles.

3. SAR Recognition System

Establishing an appropriate local coordinate reference frame is critical to reliable recognition, based on locations of features in SAR images, of objects that can be in articulated configurations or be partially occluded. The object articulation and occlusion problems require the use of local features; global features and measures do not work, because the center of mass and principal axes of the object change with articulation or occlusion.[2] In a SAR image the radar range and cross-range directions are known and choosing any local reference point, such as a scattering center location, establishes a reference coordinate system. The relative distance and direction of the other scattering centers can be expressed in radar range and cross-range coordinates, and naturally tesselated into integer buckets that correspond to the radar range/cross-range bins. The scale is determined by the bin size, which is a function of the frequency of the radar (one foot resolution and X-band in this case). The recognition system takes advantage of this natural system for SAR,

where selecting a single basis point performs the translational transformation and fixes the coordinate system to a "local" origin.

The SAR recognition system uses standard non-articulated models of the objects (at 1° azimuth increments) to recognize the same objects in non-standard, articulated and occluded configurations. The relative locations and magnitudes of the N strongest SAR scattering centers (local maxima in the radar return signal) are used as characteristic features. Using a technique similar to geometric hashing,[10] the relative positions of pairs of scattering centers in the range (R) and cross-range (C) directions are the initial indices to a look-up table of labels that give the associated target type/pose and the remaining features: range and cross-range position of the "origin" and the magnitudes of the two scatterers. (The "origin", O, is the strongest of a pair of scatterers, the other is a "point", P.) Any local reference point, such as a scattering center location, can be chosen as an origin to establish a reference coordinate system for building a model of an object at a specific azimuth angle pose. For ideal data, picking the location of the strongest scattering center as the origin 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, non-standard target configurations, etc.), we use all N strongest scattering centers in turn as the origin to ensure that a valid local reference point is obtained. Thus, to handle articulation and occlusion, 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 recognition process is an efficient search for positive evidence, using relative locations of scattering centers 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 to $\pm L\%$. (The design parameters N and L are optimized, based on experiments, to produce the best forced recognition results). 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 space. The process is repeated with different scattering centers as reference points, providing multiple queries of 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. In contrast to many other model-based approaches (e.g. Refs. 1 and 8), this recognition process is not exhaustively searching all the models; instead we are doing table look-ups based on the relative distances between the strongest scatterers in the test image to accumulate evidence for the recognition result. The efficiency of this approach is indicated by the typical one second time to process a single test image with 38 scatterers (using a Sun Ultra 2 workstation with no attempts to optimize the code).

To handle identification with "unknown" objects, we introduce a criteria for the quality of the recognition result (e.g. the votes for the potential winning object exceed some threshold, D). By varying the decision rule parameter (typically from 50 to 400 votes in 10 vote increments) we obtain a form of Receiver Operating Characteristic (ROC) curve with probability of correct recognition (PCR) versus probability of false alarm (PFA).

More formally, a radar image of object c at azimuth pose a consists of N (or more) scatterers, each scatterer k with a magnitude S_k and range and cross-range locations R_k and C_k, which (for consistency) are ordered by decreasing magnitude such that $S_k \geq S_{k+1}$ where $k = 1, \ldots, N$. A *model* M of object c at azimuth a is given by:

$$M(c,a) = \{V_1(c,a), V_2(c,a), \ldots, V_{N(N-1)/2}(c,a)\}\,, \tag{1}$$

which is comprised of the set of all pairwise *observations*, V_i:

$$V_i(c,a) = \{f_1, f_2, \ldots, f_6\}_i\,, \tag{2}$$

where $i = 1, 2, \ldots, N(N-1)/2$, $f_1 = R_P - R_O$, $f_2 = C_P - C_O$, $f_3 = R_O$, $f_4 = C_O$, $f_5 = S_O$, $f_6 = S_P$, and with the individual scatterers in each pair labeled O and P so that $S_O \geq S_P$ for consistency (see Fig. 3).

We define a *match*, H, as:

$$H(V_i, V_j) = \begin{cases} 1 & \text{if } |(f_b)_i - (f_b)_j| \leq \delta_b\,,\ \forall b = 1, \ldots, 6\,, \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where the match constraints are $\delta_1 = \delta_2 = 0$ pixels, $\delta_3 = \delta_4 = 5$ pixels (translation) and $\delta_5 = \delta_6 = L$ percent. A subscript t applied to a match denotes that the match, H_t, is associated with the relative translation $t(R,C) = (\Delta f_3, \Delta f_4)$ of the stronger scatterers in the two observations.

The recognition result, T, for some test image (with a test class, x, and test azimuth, y, to be determined) is a maximal match that is greater than a threshold, D, given by:

$$T = \begin{cases} [c,a], \text{ if } \arg\max\limits_{c,a,t} \left(\sum\limits_{l=1}^{9} \sum\limits_{k=1}^{N(N-1)/2} \sum\limits_{n=1}^{9} H_t^l(V_k^n(x,y), V_m(c,a)) \right) > D \\ \text{“unknown”, otherwise}\,, \end{cases} \tag{4}$$

where $V_m \in M(c,a)\ \forall m$ such that $|(f_1)_{V_k^n} - (f_1)_{V_m}| = 0$ and $|(f_2)_{V_k^n} - (f_2)_{V_m}| = 0$ (note: this formulation avoids exhaustive search of all the models and can be implemented as a look-up table). The nine observations (denoted by the superscript n in V_k^n) are made to account for location uncertainty by taking the 3×3 neighbors about the nominal values for the relative locations f_1 and f_2 of scatterer pair k in the test image. Similarly, the nine matches (denoted by the superscript l in H_t^l) are computed at the 3×3 neighbors located ± 1 pixel about the resulting nominal value for translation, $t(R,C)$, of the scatterers in the test image from the model.

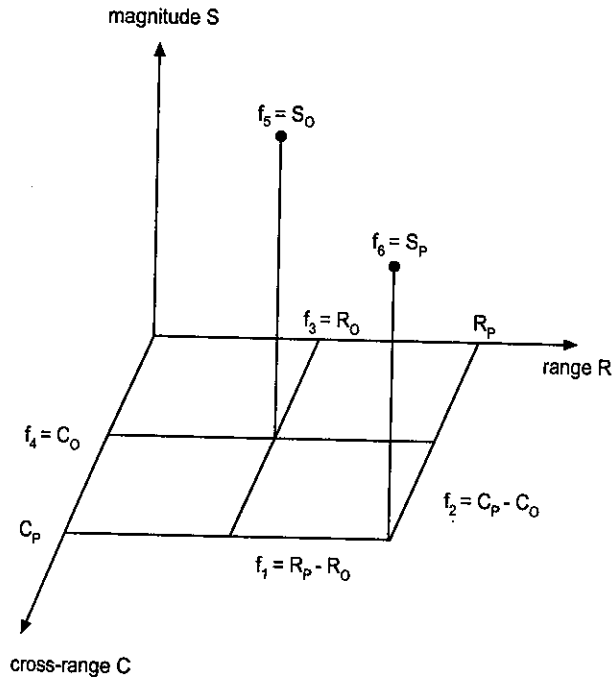

Fig. 3. Observation for a pair of scatterers O and P.

4. Independence of Multiple Look Angle SAR Recognizers

The azimuthal variance of SAR scatterer locations, previously demonstrated in Fig. 2, strongly indicates that observations at different azimuth angles are independent. Given the probability of one SAR recognizer failing, $F1$, where $F1 = 1 - PCR$; then if two recognizers are independent, the probability that both recognizers are wrong, $F2$, is simply $F2 = F1^2$. In order to obtain the most failures, we pick the object configuration variant case, which we had previously determined[2] to be the most difficult case for our SAR recognition approach (compared to the depression angle change and object articulation cases).

In the configuration variant experiments a single configuration of the T72 (#132) and BMP2 (#C21) vehicles are used as the models and the test data are two other variants of each vehicle type (T72 #812, #s7 and BMP2 #9563, #9566), all at 15° depression angle. Comparing the model with the test data at the same azimuth for different configurations of the same type of vehicle, typically less than 20 percent of the scatterer locations are invariant with configuration differences for an exact match of scatterer locations, while about 60 percent are quasi-invariant for a location match within a 1 pixel (8 neighbor) uncertainty region.[2]

Recognition results are obtained using the optimum parameters of $N = 38$ scattering centers and a percent magnitude change, L, of less than $\pm 11\%$. Figure 4 shows the probability that two recognizers are both wrong for forced recognition of configuration variants as a function of the difference in azimuth angle of the

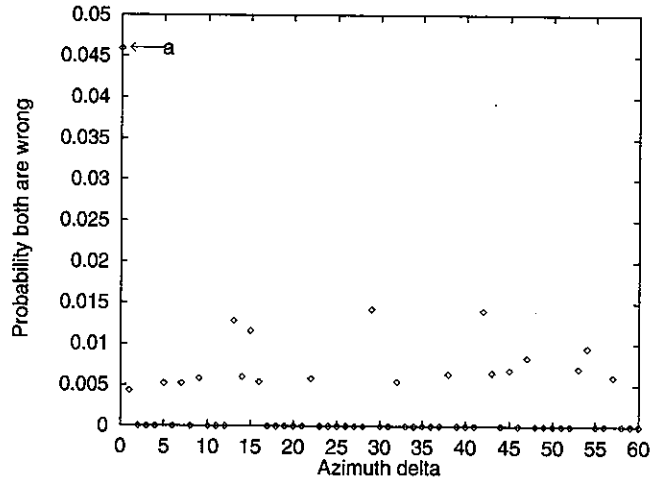

Fig. 4. Probability that two recognizers are both wrong.

object from the two recognizers. (While Fig. 4 emphasizes the small angles by only showing up to ± 60 degrees, the results out to ± 180 degrees are similar.) The single recognizer result, 19 failures of 414 trials, is an F1 failure rate of 0.0459, which is plotted for reference as point 'a' in Fig. 4. For an F1 of 0.0459 the predicted value of F2 is 0.0021, which is very close to the overall experimental average F2 of 0.0025. The other interesting observation from Fig. 4 is that the results are independent of the angle difference. Although Fig. 2 shows that there is some correlation in scatterer locations for a one degree angle change (about 20 percent for an exact match), Fig. 4 shows that the results for two recognizers are independent even for small angle differences like one degree. This demonstration that multiple look angle SAR recognition results are independent, even for small angles, provides the scientific basis for both measuring and improving the quality of recognition results.

5. Multiple Look Angle SAR Recognition Results

5.1. *Configuration variants*

In contrast to the forced recognition case described in the previous section, in these configuration variant experiments the BTR70 armored personnel carrier (#c71) is used as an unmodeled confuser vehicle to test the recognition system and the vote threshold parameter is used to generate "unknown" results. The other test conditions and parameters are the same as the forced recognition case (most significantly, the models are one configuration of the T72 and BMP2 vehicles and the test vehicles are two different configurations).

Figure 5 shows the effect of multiple look angle recognizers on the probability of false alarm using the BTR70 as a confuser. (The BTR70 is a more difficult case than other confusers such as the BRDM2 armored personnel carrier or the ZSU

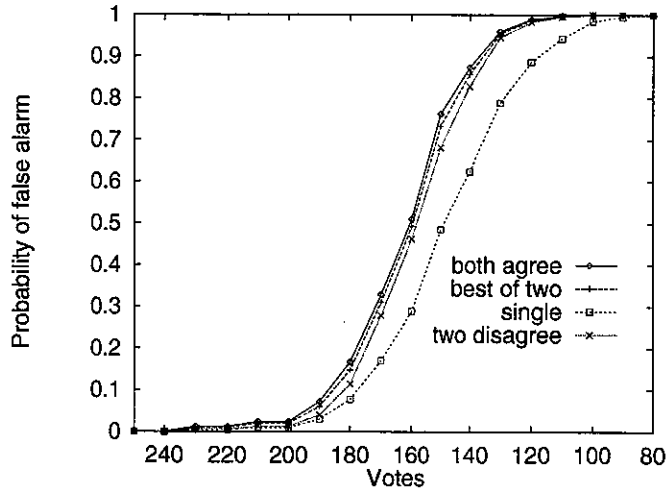

Fig. 5. Effect of multi-look on probability of false alarm.

23/4 anti-aircraft gun.[2]) In the cases with two recognizers, the decision rule is that if either gives results above the vote threshold, the result is declared a target (which, for these BTR70 confusers, would be a false alarm). Thus, with this "target bias" decision rule for the multiple recognizer cases have higher false alarms than a single recognizer. It is important to note that the penalty in increased false alarms is small for the left tail of the curve. Figure 5 also shows that the false alarm rates are similar for all the two look angle recognizer cases and that agreement on the "target" is basically irrelevant for a false alarm.

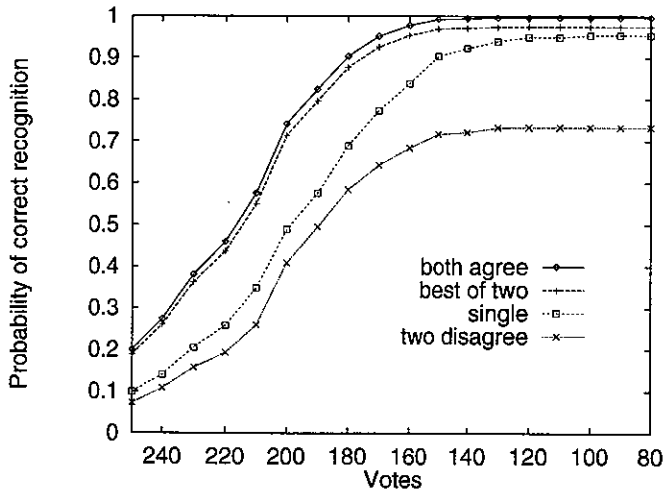

Fig. 6. Effect of multi-look on probability of correct recognition.

Figure 6 shows the effect of multiple look angle recognizers on the probability of correct recognition for the configuration variant case where the test data are different configurations of the T72 and BMP2. The top curve shows the results for the 91.7 percent of the time when two recognizers at different look angles agree on the result. The bottom curve is for the remaining 8.3 percent of the time when the two recognizers disagree and the answer that gets the most votes is chosen. The second curve, labeled "best of two", uses a decision level fusion rule that simply picks the target based on which of the two recognizers got the most votes. (This case is also the weighted average of the agree and disagree cases.) In Fig. 6 the probability of correct recognition decreases as the vote threshold increases (to the left in Fig. 6), because the higher threshold causes more targets to be classified as "unknown". The recognition results for using the best of two recognizers at different look angles are substantially better than the results for a single recognizer. This is basically the result of fewer "misses", where a target object is classified as an "unknown"; because there are two opportunities to get above the vote threshold and declare a "target".

Figure 7 combines the results of Figs. 5 and 6 and shows the effect of using multiple look angle recognizers on the Receiver Operating Characteristic curve for the configuration variant cases. These recognition results, using the best of two recognizers at different look angles, are substantially better than the results for a single recognizer. For example, at a 0.10 PFA the PCR for the best of two look angles is 0.8324, compared to 0.7091 for a single recognizer. The performance improvement is because the cost in increased false alarms (in Fig. 5) is low compared to the benefits in increased recognition (in Fig. 6), due to fewer targets being classified as "unknown".

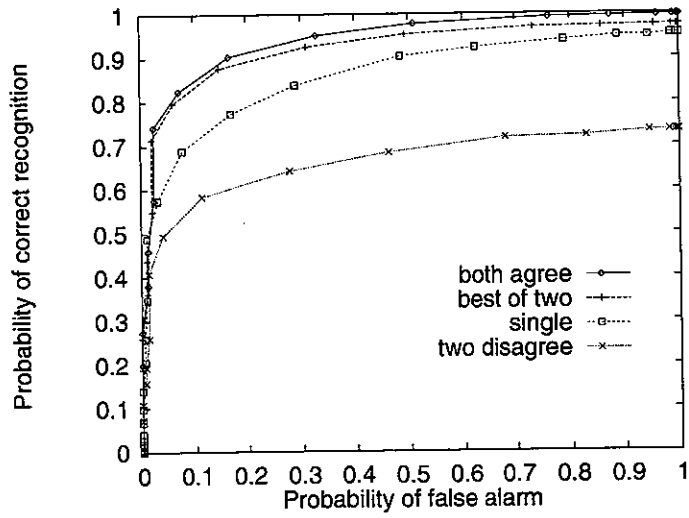

Fig. 7. Effect of multi-look on configuration variant ROC curve.

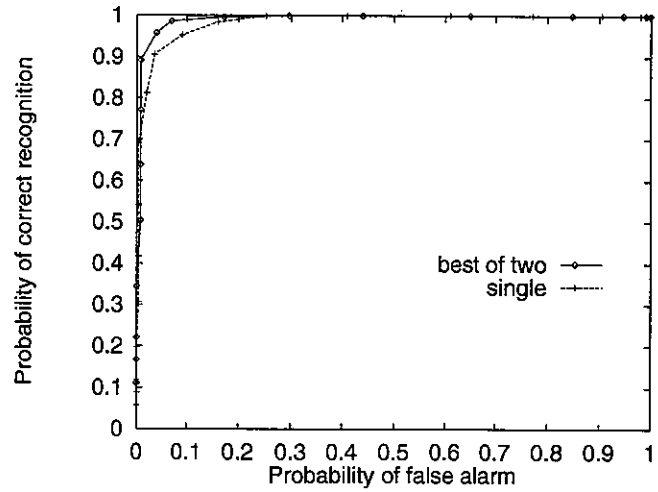

Fig. 8. Effect of multi-look on articulated object ROC curve.

5.2. *Articulation and occlusion*

In the articulated object experiments the models are non-articulated versions of the T72 tank #a64 and the ZSU 23/4 anti-aircraft gun #d08 (with the gun turret straight forward). The test data are articulated versions of these same serial number objects (with the turret rotated) and the BRDM2 armored personnel carrier #e71 is a confuser vehicle (all at 30° depression angle). The results of applying the same techniques (and all the same recognition system parameters) in these articulated object experiments are shown as ROC curves in Fig. 8. Again the results for using two recognizers at different look angles and picking the answer with the largest number of votes are better than the single recognizer results.

The occlusion experiments use four models: T72 tank #132, BMP2 APC #C21, BTR70 APC #c71 and ZSU23/4 gun #d08 and the unmodeled confuser vehicle is BRDM2 APC (#e71) (all at 15° depression angle). Since there is no real SAR data with occluded objects available to the general public, the occluded test data in this paper is simulated by starting with a given number of the strongest scattering centers in target chips of these same four objects and then removing the appropriate number of scattering centers encountered in order from 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). Then the same number of scattering centers (with random magnitudes) are added back at *random locations* within the original bounding box of the chip. This is the same technique used in Ref. 3; it keeps the number of scatterers constant and acts as a surrogate for some potential occluding object. In some of our previous work on occluded objects,[3] the confuser vehicle was occluded. However, while the target may be occluded, the confuser vehicle may not necessarily be occluded in the

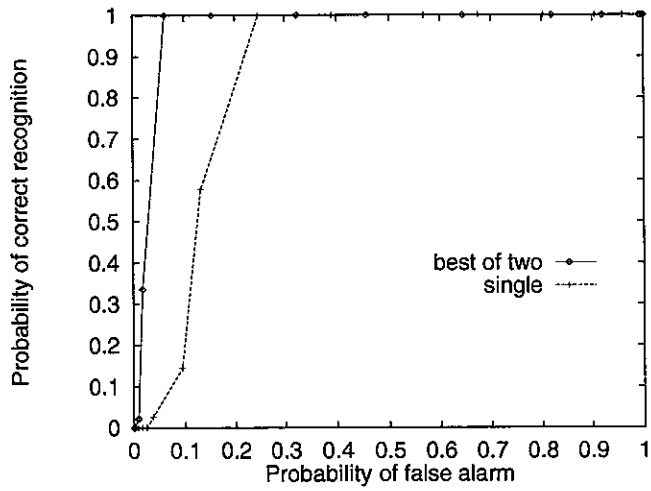

Fig. 9. Effect of multi-look on 50% occluded object ROC curve.

practical case. Hence, in this research the BRDM2 APC (#e71) is an *unoccluded confuser* vehicle, which is a substantially more difficult case.

Figure 9 shows the effect of multiple look angle recognizers on the probability of correct recognition for the case of 50% occluded targets and an unoccluded confuser. The same techniques (and all the same recognition system parameters) used in the prior configuration variant and articulation experiments are applied to these occluded object experiments. Here again, using two recognizers at different look angles and a decision level fusion rule of picking the answer with the largest number of votes gives better results than a single recognizer.

6. Conclusions and Future Work

The fundamental azimuthal variance of SAR scatterer locations can be successfully used as the basis for a principled and effective multiple look angle SAR recognition approach. The experiments demonstrate that SAR recognition results at different azimuths are *independent*, even for small azimuths, such as one degree. In addition, using decision level fusion of two observations at different look angles can substantially increase SAR recognition performance for target configuration variants. These techniques can also be successfully applied to recognition of articulated objects and occluded objects.

This paper considers only two look angles, but a similar approach could be used for more look angles. In addition, more formal methods, like Dempster-Shafer and Bayesian approaches, can be used for decision level fusion. Our approach can be applied to resolve ambiguities with multiple looks from the same or different sensor platforms.

Acknowledgments

This work was supported by DARPA/AFOSR grant F49620-97-1-0184 and AFOSR grant F49620-02-1-0315; the contents and information do not necessarily reflect the position or policy of the US Government.

References

1. T. Ross, S. Worrell, V. Velten, J. Mossing, and M. Bryant, "Standard SAR ATR evaluation experiments using the MSTAR public release data set," in *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery V*, **3370**, 566–573 (April 1998).
2. B. Bhanu and G. Jones, "Recognizing target variations and articulations in synthetic aperture radar images," *Optical Engineering* **39**(3), 712–723 (March 2000).
3. G. Jones and B. Bhanu, "Recognizing occluded objects in SAR images," *IEEE Transactions on Aerospace and Electronic Systems* **37**(1), 316–328 (January 2001).
4. G. Jones and B. Bhanu, "Recognizing articulated targets in SAR images," *Pattern Recognition* **34**(2), 469–485 (February 2001).
5. G. Jones and B. Bhanu, "Recognition of articulated and occluded objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**(7), 603–613 (July 1999),
6. D. Dudgeon, R. Lacoss, C. Lazott, and J. Verly, "Use of persistent scatterers for model-based recognition," in *SPIE Proceedings: Algorithms for Synthetic Aperture Radar Imagery* **2230**, 356–368 (April 1994).
7. L. Novak, S. Halversen, G. Owirka, and M. Hiett, "Effects of polarization and resolution on SAR ATR," *IEEE Transactions on Aerospace and Electronic Systems* **33**(1), 102–115 (January 1997).
8. K. Ikeuchi, T. Shakunga, M. Wheeler, and T. Yamazaki, "Invariant histograms and deformable template matching for SAR target recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 100–105 (June 1996).
9. J. Verly, R. Delanoy, and C. Lazott, "Principles and evaluation of an automatic target recognition system for synthetic aperture radar imagery based on the use of functional templates," in *SPIE Proceedings: Automatic Target Recognition III* **1960**, 57–71 (April 1993).
10. Y. Lamden and H. Wolfson, "Geometric hashing: A general and efficient model-based recognition scheme," in *Proc. International Conference on Computer Vision* pp. 238–249 (December 1988).

Bir Bhanu received the SM and EE degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, the PhD degree in electrical engineering from the Image Processing Institute, University of Southern California, Los Angeles, and the MBA degree from the University of California, Irvine.

Currently Professor Bhanu is the Director of the Center for Research in Intelligent Systems (CRIS) at the University of California, Riverside where he has been a Professor and Director of the Visualization and Intelligent Systems Laboratory (VISLab) since 1991. Previously, he was a Senior Honeywell Fellow at Honeywell Inc. in Minneapolis, MN. He has been on the

faculty of the Department of Computer Science at the University of Utah, Salt Lake City, UT, and has worked at Ford Aerospace and Communications Corporation, CA, INRIA-France and IBM San Jose Research Laboratory, CA. He has been the principal investigator of various programs for DARPA, NASA, NSF, AFOSR, ARO and other agencies and industries in the areas of learning and vision, image understanding, pattern recognition, target recognition, navigation, image databases, and machine vision applications. He is the co-author of books on Computational Learning for Adaptive Computer Vision (Forthcoming), Genetic Learning for Adaptive Image Segmentation (Kluwer 1994), and Qualitative Motion Understanding (Kluwer 1992). He has received two outstanding paper awards from the Pattern Recognition Society and has received industrial awards for technical excellence, outstanding contributions and team efforts. He has been the guest editor of several IEEE transactions and other journals, and on the editorial board of various journals. He holds eleven US and international patents and over 200 reviewed technical publications in the areas of his interest. He has been the General Chair for IEEE Workshops on Applications of Computer Vision, the Chair for the DARPA Image Understanding Workshop, the General Chair for the IEEE Conference on Computer Vision and Pattern Recognition, and Program Chair for the IEEE Workshops on Computer Vision Beyond the Visible Spectrum. Professor Bhanu is a Fellow of IEEE, AAAS, IAPR and SPIE, and a member of the ACM, and AAAI.

Grinnell Jones III was a National Merit Scholar who received the BS degree in mechanical engineering from the Massachusetts Institute of Technology in 1966, MS in aerospace engineering (with distinction) from the Air Force Institute of Technology in 1971, and MS in computer science from the University of California, Riverside in 1997.

After a 25-year career in development engineering, missile operations and acquisition management with the US Air Force, Lt. Col. (USAF Retired) Jones has been conducting computer science research developing systems and algorithms for automatic target recognition using synthetic aperture radar imagery, for the past seven years with the University of California, Riverside. His research interests include object recognition, computer vision, machine learning, image and video databases, and systems applications.