Effects of Image Quality on SAR Target Recognition

Leslie M. Novak
Scientific Systems Company, Inc.
500 West Cummings Park, Suite 3000
Woburn, MA 01801
UNITED STATES OF AMERICA

lnovak@ssci.com novakl@charter.net

ABSTRACT

Target recognition systems using Synthetic Aperture Radar (SAR) data require well-focused target imagery to achieve high probability of correct classification. Techniques for improving the image quality of complex SAR imagery are investigated. The application of phase gradient re-focusing of target imagery having cross-range smearing is shown to significantly improve the target recognition performance of a model-based ATR system. The application of High Definition Imaging (HDI) is also shown to enhance and improve the image quality and resolution of SAR target imagery -- the improvement in target recognition performance of a template-based ATR system using HDI-processed SAR imagery is quantified. Interruptions in SAR phase history data are shown to significantly degrade SAR image quality; CS-based image formation (Basis Pursuit Denoising) is demonstrated to mitigate the effects of data interruptions, providing complex SAR imagery with excellent image quality.

1. INTRODUCTION

SAR image quality has a significant effect on the performance of SAR automatic target recognition systems. Template-based classifiers and model-based classifiers both require well-focused imagery in order to accurately match an observed target image to a database of stored templates or features such as peak-scatterer locations, etc. High-resolution SAR requires accurate motion compensation in order to form well-focused target images, and errors in motion compensation can yield images with poor image quality, such as excessive cross-range image smearing or blurring. Section 2 of this paper demonstrates the use of phase gradient processing to refocus target imagery degraded by cross-range smearing. ATR performance of a model-based classifier is investigated; the probability of correct classification (Pcc) is compared using target imagery having significant cross-range smearing versus target imagery that has been refocused using phase gradient algorithm (PGA) processing; it is demonstrated that model-based ATR performance is improved considerably by using PGA-processing prior to passing the target imagery to the ATR.

Section 3 of this paper applies to high-resolution SAR imagery already having good image quality, including imagery that has been well-focused using PGA-processing; the technique investigated improves image quality by enhancing the resolution of well-focused complex SAR imagery. A 10-target, template-based classifier is described and classifier performance is presented using SAR imagery having 0.3m x 0.3m, 0.5m x 0.5m, and 1.0m x 1.0m resolutions; classifier performance is presented in terms of confusion matrices and probability of correct classification (Pcc). Next, enhanced resolution imagery is formed from the original 0.3m x 0.3m and 1.0m x 1.0m data using Lincoln Laboratory's High Definition Imaging (HDI) algorithm -- this processing improves (approximately) the resolution of the data to 0.15m x 0.15m and 0.5m x 0.5m, respectively; and the image background speckle noise is reduced. The improvement in the performance of the template-based classifier due to using HDI-processed data is quantified.

In Section 4 we investigate the effects of data gaps (interruptions) in SAR phase history data; we show examples of SAR imagery formed using phase history data containing various types of gapping patterns -- and we demonstrate that SAR image artifacts induced by interruptions in SAR phase history data are mitigated using Compressed Sensing image formation, resulting in excellent quality complex SAR imagery.

Section 5 of the paper summarizes the results and conclusions of these studies. Section 6 provides the references used in this research study.

2. IMPROVING ATR PERFORMANCE VIA PGA IMAGE QUALITY ENHANCEMENT

This section presents an example of an ATR performance study using high-resolution SAR imagery gathered by the Lynx SAR. In this example the imagery was gathered at a nominal resolution of 0.15m by 0.15m in spotlight mode; a contiguous sequence of seven SAR images were used in this study. Figure 1 shows these seven images. Visually, these SAR images appear to have very good image quality (IQ), however, our target recognition studies show that SAR image #1 has the best image quality and SAR image #7 has the poorest image quality.

A side-by-side comparison of SAR Image#1 with SAR Image #7 is shown in Figure 2. Comparing the bright scatterer located on the uppermost target shows that Image #7 has significant cross-range blurring, most likely due to uncompensated platform motion. Our analyses will show that the average cross-range scatterer width = 17 pixels for the scatterers in Image #1, whereas the average cross-range scatterer width = 36 pixels for the scatterers in Image #7. These "image quality feature" values were calculated from the image data during our target recognition studies. As we will show, the importance of this observation is directly related to the performance of the ATR system. A model-based target recognition system was used to classify the individual targets in each image. The classifier was designed to recognize 20+ military targets. The target array in these studies, as shown in Figures 1and 2, contained twelve military targets, and seven of these targets were contained in the classifier's set of 20+ targets.

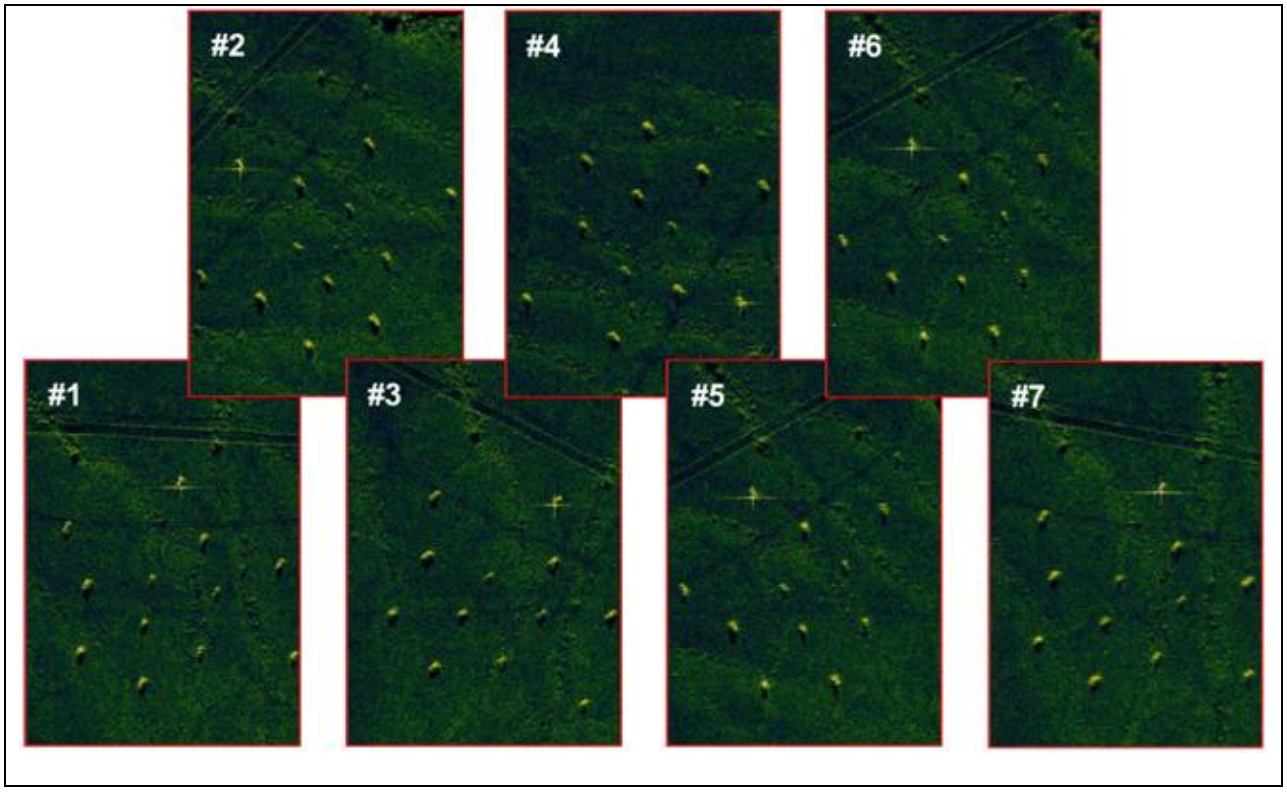

Figure 1: Sequence of seven SAR images.

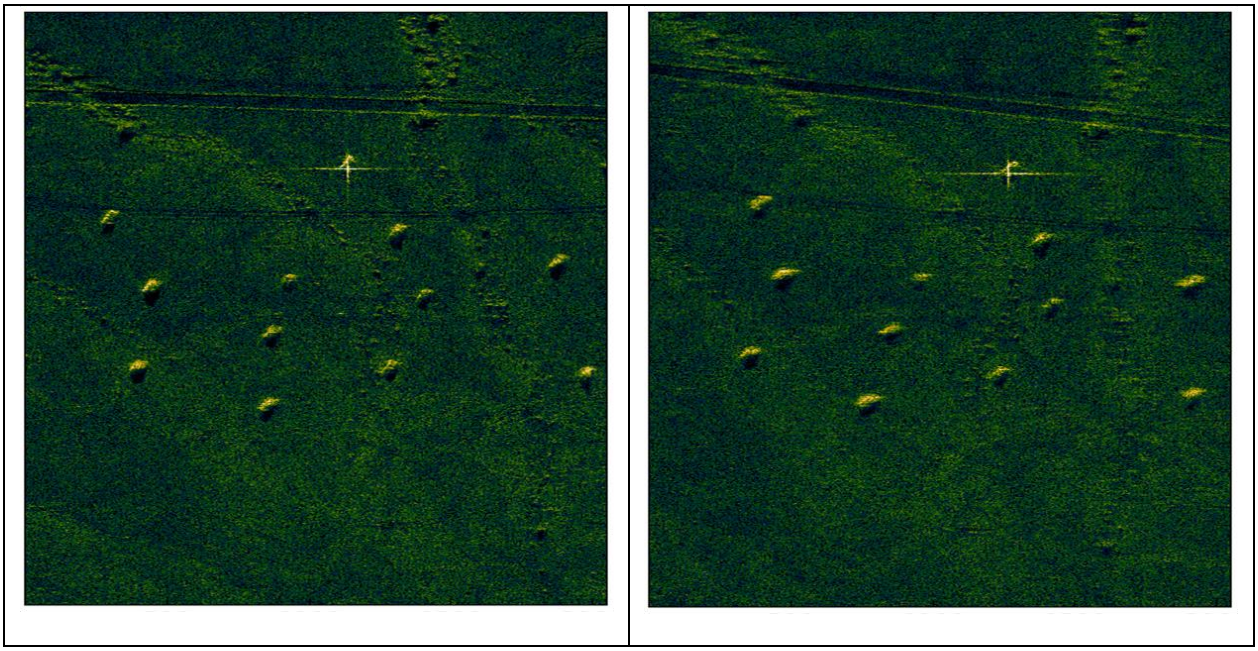

Figure 2: Image #1 (left), Image #7 (right).

A CFAR detector was used to detect the targets located in these SAR images. Each of the seven targets contained in the classifier's 20+ target set were presented to the model-based classifier; thus, a total of 49 target images were input to the classifier. Table 1 presents the target recognition results obtained for Image #1 (column 2) versus the results obtained for Image #7 (column 3). As the table shows, each of the targets contained in Image #1 were correctly classified (Pcc = 7/7). Four of the targets from Image #7 were incorrectly classified (highlighted in RED); thus, for this image, Pcc = 3/7.

Next, the targets from Image #7 were refocused using Phase-Gradient processing (see Reference [1]). Each target's brightest scatterers were CFAR detected and aligned as required by the PGA algorithm and averaged -- an average cross-range scatter width was calculated from the average of the brightest scatterers. Table 1 summarizes the cross-range scatter widths obtained for each target, and also an average width obtained for Image #1 and Image #7.

As stated previously, Image #1 has an average cross-range width = 17 pixels and Image #7 has an average cross-range width = 36 pixels. As shown in Table 1, column 4 tabulates the "classifier calls" and the average cross-range widths obtained after applying 3 iterations of phase gradient focusing to Image #7. After (PGA3) focusing, each of the targets in Image #7 were correctly classified (Pcc = 7/7) -- and the average cross-range width was reduced to 14 pixels.

TARGET TRUTH	IMAGE #1 (ORIGINAL)	IMAGE #7 (ORIGINAL)	IMAGE #7 (AFTER PGA3)
T72	"T72" WIDTH = 23	"BRDM2" WIDTH = 45	"T72" WIDTH = 7
2S1	"2S1", WIDTH = 19	"BMP2", WIDTH = 51	"2S1" WIDTH = 19
M60	"M60" WIDTH = 21	"BRDM2" WIDTH = 47	"M60" WIDTH = 21
M2	"M2" WIDTH = 11	"M2" WIDTH = 21	"M2" WIDTH = 17
BMP2	"BMP2" WIDTH = 19	"M113" WIDTH = 37	"BMP2" WIDTH = 11
BRDM2	"BRDM2" WIDTH = 11	"BRDM2" WIDTH = 21	"BRDM2" WIDTH = 15
M113	"M113" WIDTH = 15	"M113" WIDTH = 27	"M113" WIDTH = 9
	PCC= 7/7, WIDTH (AVG.) = 17	PCC= 3/7 WIDTH (AVG.) = 36	PCC= 7/7, WIDTH (AVG.) = 14

Table 1: Classifier performance, Image #1 vs. Image #7

A comparison of SAR target images from Image #7 is presented in Figure 3. The left target image shows significant cross-range image blurring; the right target image is the same target after reprocessing the complex image data using 3 iterations of the PGA algorithm. Table 2 summarizes the model-based classifier's performance for each of the seven target images processed.

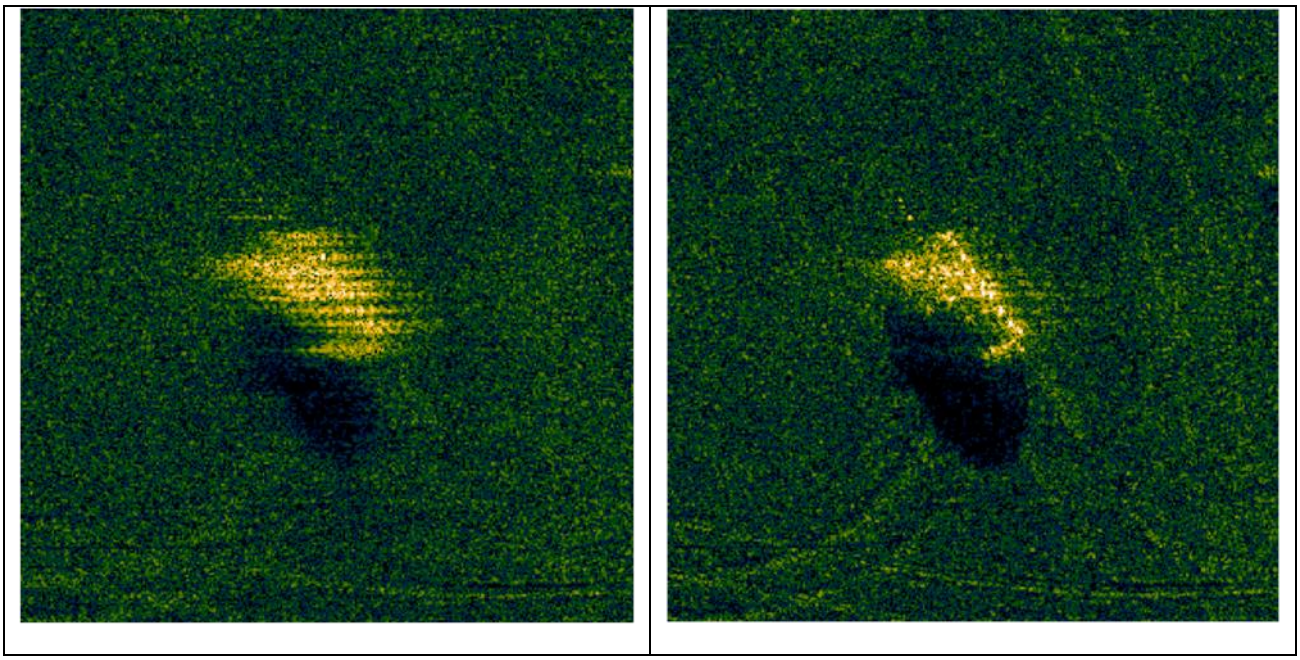

Figure 3: Target extracted from Image #7; original (left), after PGA3 (right).

Table 2: Summary of Model-based Classifier performance.

IMAGE NUMBER	AVG. WIDTH (ORIGINAL)	AVG. WIDTH (AFTER PGA3)	PCC (ORIGINAL)	PCC (AFTER PGA3)
#1	17.0	11.5	7/7	7/7
#2	21.8	13.8	7/7	7/7
#3	23.2	16.1	6/7	7/7
#4	24.1	11.8	6/7	7/7
#5	26.4	16.0	5/7	5/7
#6	31.5	12.7	6/7	6/7
#7	35.5	14.1	3/7	7/7
Averages	**25.6**	**13.7**	**40/49**	**46/49**

An alternative image-focusing algorithm based upon minimizing the image entropy is described in Reference [2]; however, results presented in Figure 4 indicate that minimum entropy focusing requires using at least 10 iterations of the entropy minimization algorithm in [2]. The entropy minimum is achieved using only 3 iterations of the PGA algorithm. Thus, PGA processing seems to be the preferred SAR image focusing technique.

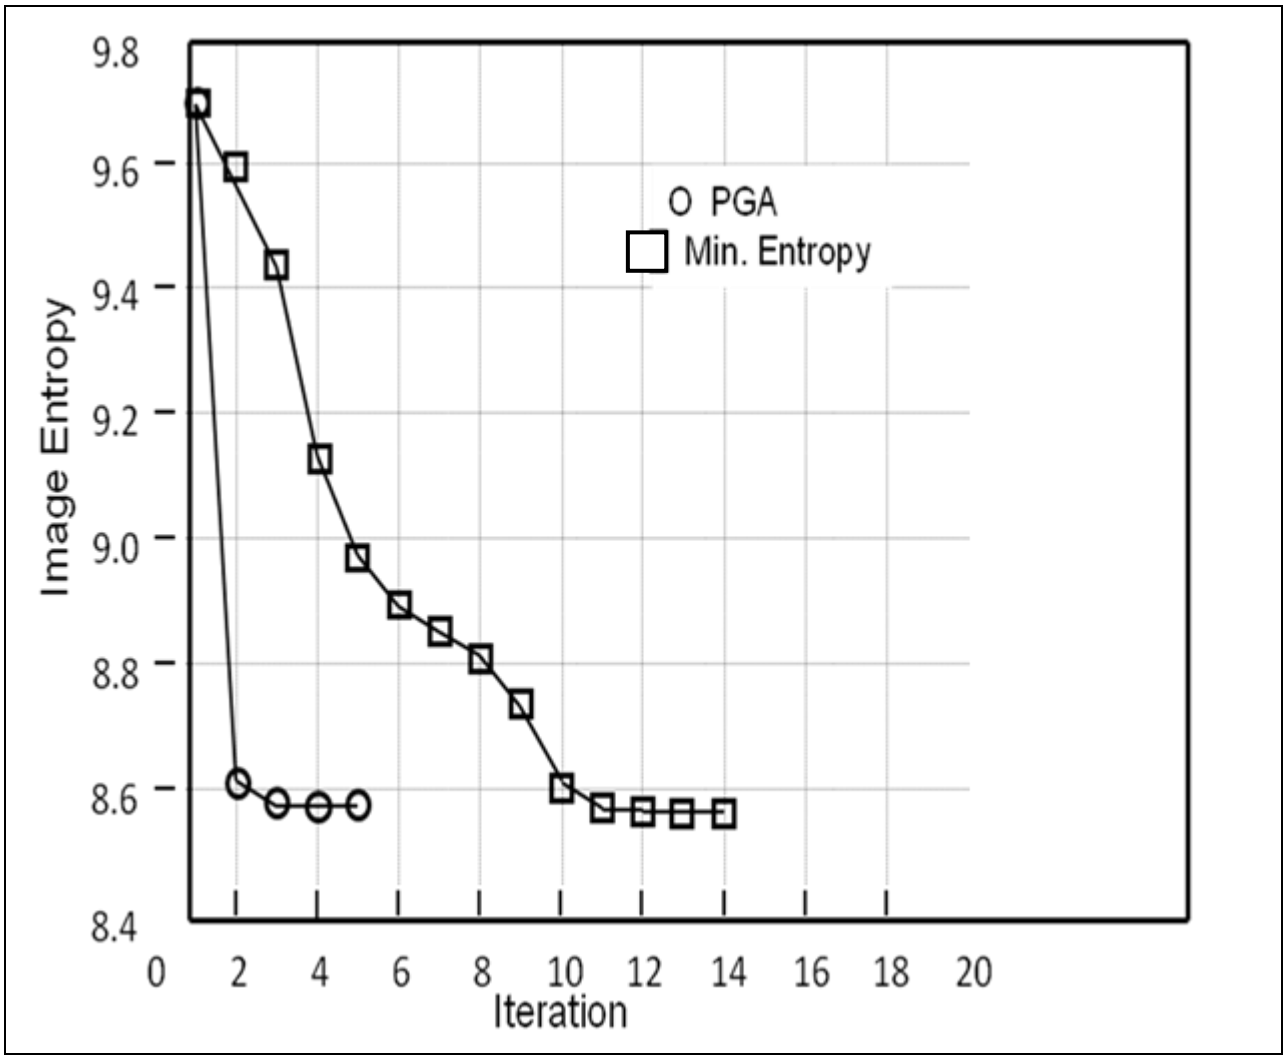

Figure 4: Image Entropy comparison, PGA vs. Min. Entropy algorithm [2].

3. IMPROVING ATR PERFORMANCE VIA HIGH-DEFINITION IMAGE PROCESSING

This section presents an approach that has been shown to improve the ATR performance of a template-based classifier [3] using complex SAR imagery that has been resolution-enhanced using Lincoln Laboratory's High Definition Image (HDI) Processing [4]. The SAR imagery used in these studies was gathered in the fall of 1995 at the Redstone Arsenal in Huntsville, AL by the Sandia X-band (9.6 GHz) HH-polarization SAR. The data comprise a large set of military targets imaged over 360 deg of aspect. In these studies the recognition performance of a template-based mean-square-error (MSE) classifier was evaluated using imagery of 18 distinct targets contained in the data set. The target set shown in Figure 5 includes three versions each of the BMP2 armored personnel carrier, the M2 armored personnel carrier, and the T72 main battle tank. The T72 tanks contain significant differences from tank to tank; T72#2 has barrels mounted on the rear of the target; T72#3 does not have skirts along the side of the target. The BMP2 and M2 armored personnel carriers have minor differences in target-to-target configuration. We trained a 10-target classifier and then evaluated the ability of the classifier to recognize and classify all 18 targets shown in Figure 5. The initial evaluations used non-HDI-processed data to establish a baseline with which the performance using HDI-processed data could be compared. The improvement in classifier performance using HDI-processed

data was then evaluated. Performance results are presented in terms of classifier confusion matrices which show the number of correct and incorrect classifications achieved; the confusion matrices are summarized in terms of a probability of correct classification (Pcc) metric. We constructed 72 classifier templates per target, covering approximately 360 deg of aspect per target; the total number of classifier templates was 720. The classifier was initially tested using the training data images as test inputs, providing a sanity check on the algorithm code.

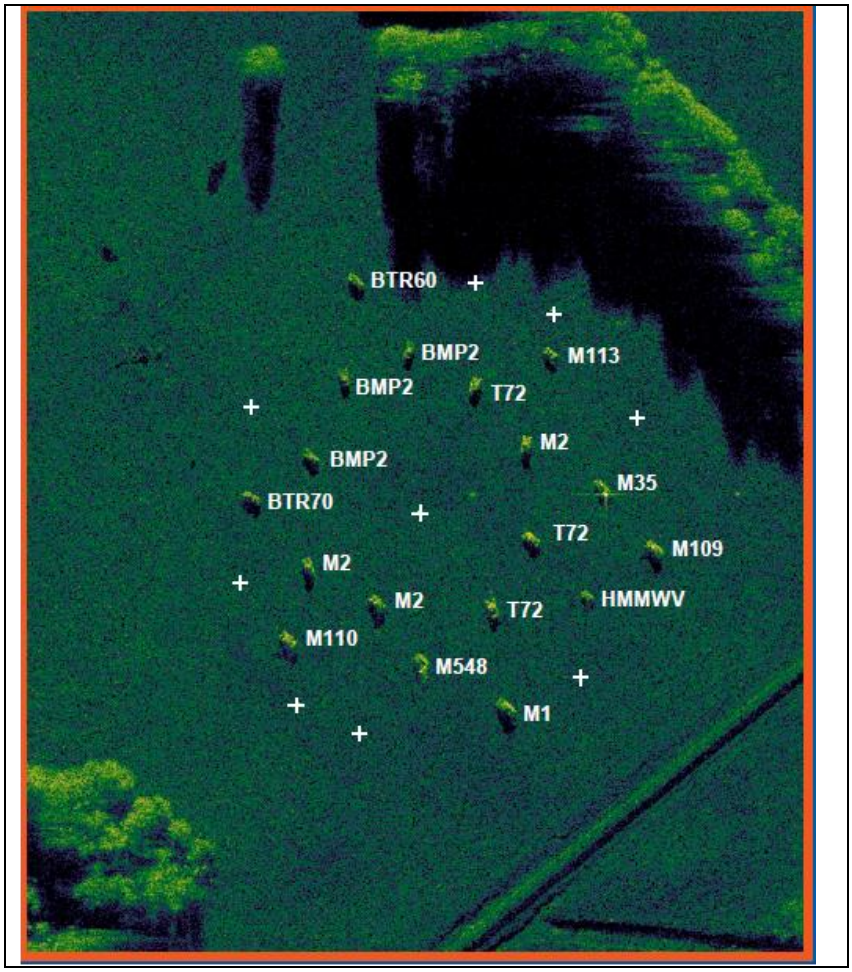

Figure 5: SAR Image of target array; plus signs (+) show corner reflector locations.

Table 3 is the classifier confusion matrix for the 0.3 m x 0.3 m resolution data. When the classifier was tested using the training data, perfect classifier performance was achieved. When the classifier was tested using the independent test data, nearly perfect classifier performance was achieved (Pcc = 93.9 %). Note, however, that the performance for T72#2, which contained extra barrels on the rear of the tank, resulted in 39 images out of the 255 total declared unknown. The performance for T72#3 (which did not have skirts along the sides of the target) was nearly perfect; only 4 images out of the 251 total were declared unknown. At this resolution, the classifier rejected a large number of confuser vehicles (438 images out of the total of 499).

Table 3: Classifier performance (0.3m x 0.3m); Pcc = 93.9 % [5].

	Number of Targets Classified as										
	BMP2	BTR60	BTR70	M109	M110	M113	M1	M2	M548	T72	Unknown
BMP2#1	255										
BMP2#2	251							1			3
BMP2#3	251							1		2	2
BTR60		256									
BTR70			256								
M109				256							
M110					256						
M113						256					
M1							256				
M2#1								256			
M2#2								251			4
M2#3								252			3
M548									256		
T72#1										256	
T72#2										216	39
T72#3										247	4
HMMWV	12	6				37		2			187
M35									4		251

Independent Test Data Confuser Vehicles (Not in Training Set)

Table 4 shows the classifier confusion matrix for 0.5 m *x* 0.5 m resolution data. The probability of correct classification for these resolution data (calculated using only the independent test vehicles and the confuser vehicles) is 84.1%. At this resolution, the M35 truck was misclassified only 13 times out of the 255 total M35 test images. The HMMWV, however, was misclassified most of the time (only 61 HMMWV images were declared unknown).

Table 4: Classifier performance (0.5m x 0.5m); Pcc = 84.1% [5].

	Number of Targets Classified as										
	BMP2	BTR60	BTR70	M109	M110	M113	M1	M2	M548	T72	Unknown
BMP2#1	256										
BMP2#2	235	1						8		7	3
BMP2#3	237	1				1		11		4	3
BTR60		255									1
BTR70			256								
M109				256							
M110					256						1
M113						255					
M1							256				
M2#1								256			
M2#2				4	5			233		4	9
M2#3				2	5	1		239		4	5
M548									254		2
T72#1										256	
T72#2					3		8	3		217	24
T72#3				1			2	1		241	6
HMMWV	29	8	1	9	3	115		14	3	1	61
M35	1			3	3	1			5		242

Independent Test Data Confuser Vehicles (Not in Training Set)

Table 5 shows the classifier confusion matrix for the 1.0m x 1.0m resolution data. For these specific targets at this resolution, we observe a very large degradation in classifier performance; the probability of correct classification degraded to 45.4%. Note, however, that nearly perfect classifier performance was achieved when the classifier was tested using the training data; this result shows the importance of testing classifiers using independent target test data.

Table 5: Classifier performance (1.0m x 1.0m); Pcc = 45.4% [5].

	Number of Targets Classifed as										
	BMP2	BTR60	BTR70	M109	M110	M113	M1	M2	M548	T72	Unknown
BMP2#1	255							1			
BMP2#2	129	15	21	8	1	39	3	15		18	6
BMP2#3	135	7	10	5	2	43		25		18	11
BTR60		255				1					
BTR70			256								
M109				256							
M110					256						
M113	1					253					2
M1				1		1	253				
M2#1						1		255			
M2#2	14	3	2	37	14	36	5	121		16	7
M2#3	21	6		18	6	47	9	118		27	4
M548									256		
T72#1						1				255	
T72#2	7	12	3	19	9	23	17	32	1	94	38
T72#3	19	8	6	23	5	12	14	19		127	18
HMMWV	48	7	4	4		145		13	1		22
M35	8			12	9	29	4	4	13	1	175

Independent Test Data Confuser Vehicles (Not in Training Set)

Table 6 shows the classifier confusion matrix for HDI-processed 0.3 m x 0.3 m resolution data (after HDI processing, the resolution of the data is approximately 0.15 m x 0.15 m). Comparing the results of Table 6 with the results of Table 3 shows somewhat-improved classifier performance; the probability of correct classification using HDI-processed data has increased to 96.4%, an improvement of 2.5% over the conventionally processed data -- and with HDI-processed 0.3 m x 0.3 m data, the classifier rejected a larger number of confuser vehicles (471 images out of the total 499).

Table 6: Classifier performance (0.3m x 0.3m + HDI); Pcc = 96.4% [5].

	Number of Targets Classified as										
	BMP2	BTR60	BTR70	M109	M110	M113	M1	M2	M548	T72	Unknown
BMP2#1	256										
BMP2#2	252										3
BMP2#3	255							1			
BTR60		256									
BTR70			256								
M109				256							
M110					256						
M113						256					
M1							255				
M2#1								256			
M2#2								251	2		2
M2#3								248	3		3
M548									256		
T72#1										256	
T72#2										232	23
T72#3										245	6
HMMWV	4	1				13		3			223
M35								1	6		248

Independent Test Data Confuser Vehicles (Not in Training Set)

Table 7 shows the classifier confusion matrix for HDI-processed 1.0 m x 1.0 m resolution data (after HDI processing, the resolution of the data is approximately 0.5 m x 0.5 m). Comparing the results of Table 7 with the results of Table 5 shows a dramatic improvement in classifier performance. The probability of correct classification using HDI-processed data has increased by approximately 30% over that achieved with conventionally processed 1.0 m x 1.0 m resolution data; the probability of correct classification has increased from 45.4% to73.4%. With HDI-processed 1.0m x 1.0m data, the number of rejected confuser vehicles increased from 197 images to 321 out of a total of 499 images. Although HDI processing of 1.0m x 1.0m data has resulted in a significant increase in the probability of correct classification (Pcc = 73.4%), performance using conventionally processed 0.5m x 0.5m resolution data gave somewhat better probability of correct classification (Pcc = 84.1 %).

Table 7: Classifier performance (1.0m x 1.0m +HDI); Pcc = 73.4% [5].

	Number of Targets Classified as										
	BMP2	BTR60	BTR70	M109	M110	M113	M1	M2	M548	T72	Unknown
BMP2#1	256										
BMP2#2	202	4	2			1		17		10	19
BMP2#3	195	2	4	2		4		22		17	10
BTR60		254									2
BTR70			256								
M109				256							
M110					256						
M113						255					1
M1							255				
M2#1								256			
M2#2	4		2	13	11	2	1	188		24	10
M2#3	8	1	2	16	5	3	1	180		27	13
M548									256		
T72#1										256	
T72#2	3			5		1	7	10		190	39
T72#3	5	1		3	4		3	10		211	14
HMMWV	25	8	8	1		101		8			93
M35	1			3	1	6			14		230

Independent Test Data Confuser Vehicles (Not in Training Set)

Figure 6 presents a side-by-side comparison of M35 Truck images; the right image was formed using conventional 2D FFT SAR processing; the left image is the corresponding HDI-processed image. This image comparison validates that HDI-processing does result in improved ATR performance. Visually, the left figure shows more clearly focused target scatterers, resulting in improved recognition of the target.

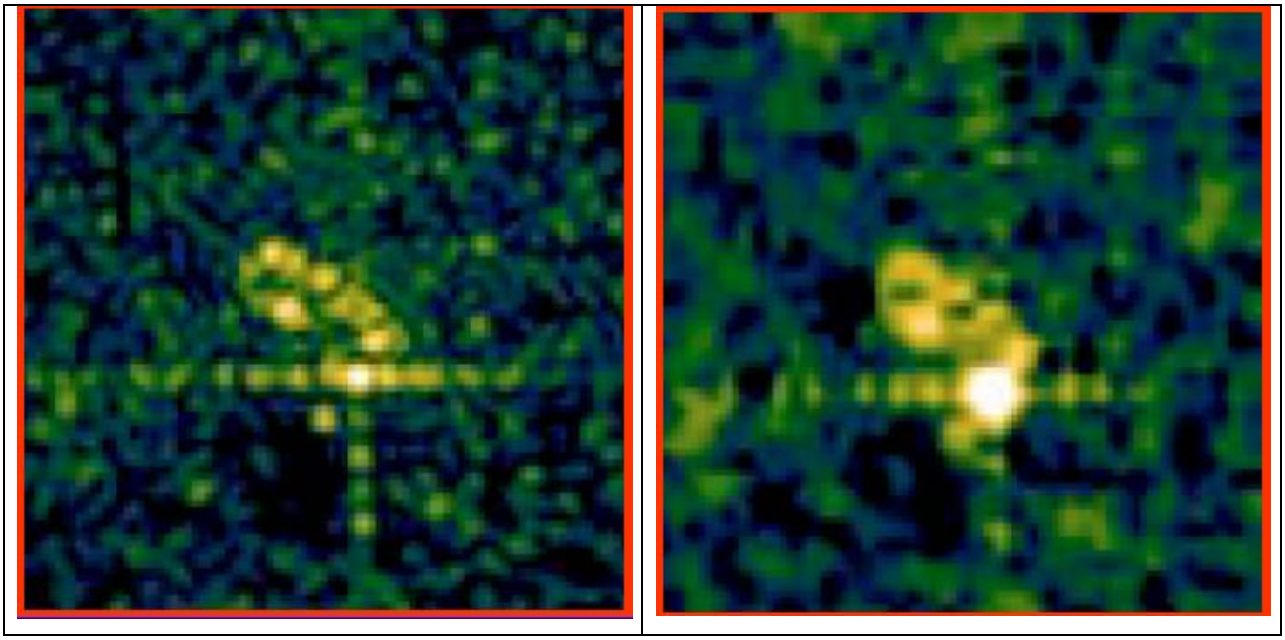

Figure 6: SAR images of an M35 Truck; Left Image, 1.0m x 1.0m HDI-Processed; Right imagine, 1.0m x 1.0m 2D FFT Processed.

Figure 7 presents a bar chart of the 10-target classifier probability of correct classification (Pcc) versus SAR image resolution. The corresponding classifier confusion matrices for the results presented in Figure 7 are presented above in Tables 3 through 7.

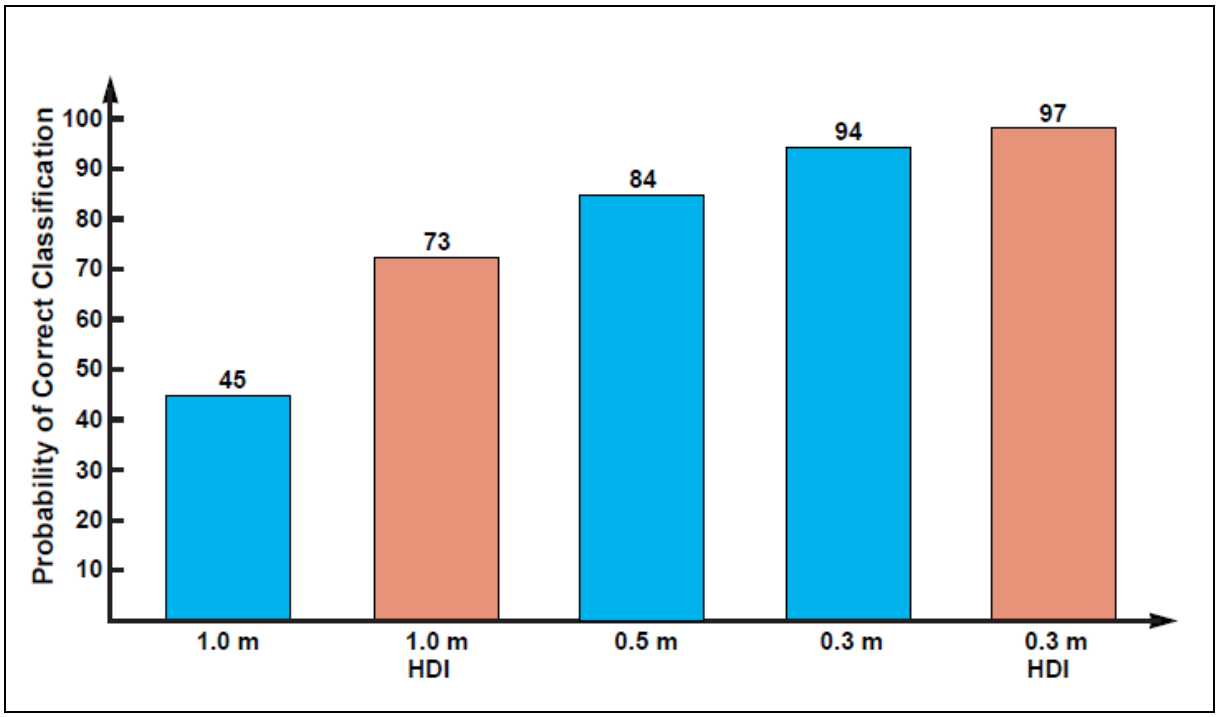

Figure 7: Summary of template-based classifier performance vs. resolution [3].

4. RECONSTRUCTION OF INTERRUPTED SAR IMAGERY

Recent advances in radar technology allow modern day radar systems to employ active array antennas with transmit/receive modules at each antenna element. Active array technology equips the airborne radar with an agile-beam, allowing the radar to accommodate multiple operational modes such as search, track, target recognition, synthetic aperture (SAR/ISAR) imaging, etc. These multiple modes may compete for radar resources, leading to an interrupted phase history data collection used in the formation of a high-resolution SAR image. Typical coherent integration times may extend to several tens of seconds, thus, potentially tying the radar resources to the single mode of operation for a long uninterrupted period of time. Depending on an aircraft's tactical situation, other high-priority modes may need to be serviced in a timely manner, and thus, it may be necessary to interrupt the SAR data collection randomly or periodically. Such interruptions leave data gaps in the coherent SAR phase history which can significantly degrade the resulting SAR image quality.

This section of the paper presents a summary of our most recent investigations into the effects on SAR image quality of various types of interruptions in the SAR phase-history data collection. We have focused on the use of compressive sensing approaches that might yield reconstructed SAR imagery having minimal loss in image quality. Specifically, we present the results of an image reconstruction study using a basis pursuit denoising (BPDN) algorithm. Our studies demonstrate that excellent SAR image quality can be achieved using phase-history data that has been degraded by various types of interruptions. These studies used very good image quality data gathered by the Sandia X-Band SAR for the MSTAR ATR program. We show that BPDN image formation of interrupted phase-history data yields reconstructed SAR imagery having excellent image quality.

The CS-based image reconstruction approach we use, referred to as "Basis Pursuit denoising", was developed by S. Chen in Reference [6]. Our use of BPDN for reconstruction of SAR imagery from gapped phase-history data is described as follows. We start with a complex SAR image having good image quality and we generate a "pseudo" phase-history data array by taking the 2-D Fourier Transform of the SAR image. We then emulate the gapped phase-history data by appropriately removing range profiles from the phase-history, thereby mimicking unobserved aspects.

We model the gapped phase-history as:

$$\underline{y} = RF\,\underline{x}, \tag{1}$$

where $\underline{y} \in C^m$ is the interrupted SAR phase-history and $\underline{x} \in C^{n^2}$ is an $n\,x\,n$ spatial reflectivity image, both lexicographically ordered as vectors. The 2-D Fourier Transform operator is denoted by $F : C^{n^2} \to C^{n^2}$, and the aspect selection operator is denoted by $R : C^{n^2} \to C^m$.

For image reconstruction we use basis pursuit denoising, defined as follows:

$$\hat{\underline{x}} = \arg_{\underline{x}} \min \|\underline{x}\|_1 \quad s.t. \left\|\underline{y} - RF\,\underline{x}\right\|_2 \le \sigma, \tag{2}$$

where the sparsity of $\underline{x}$ is enforced on the magnitude of the spatial reflectivity, i.e.

$$\|\underline{x}\|_1 = \sum_i \sqrt{(R_e[x_i])^2 + (I_m[x_i])^2} \tag{3}$$

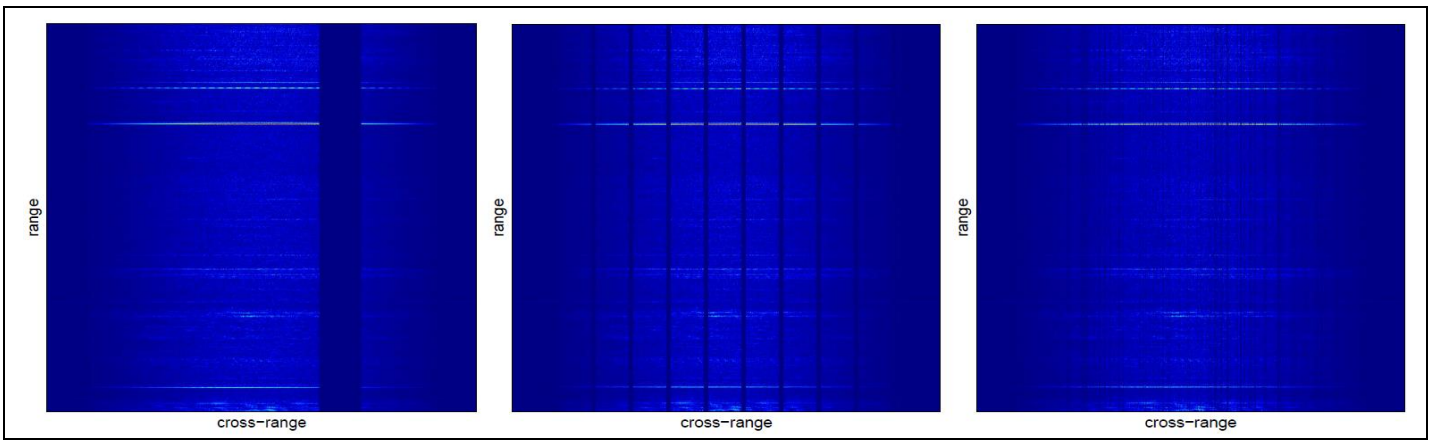

Figure 8: Range profiles of phase-history data filled with zeros during interrupts; total interrupt width is 100 in 1024 Profiles; (left) single gap; (center) periodic interrupts (10 gaps, each of width 10; (right) random gaps.

Figure 8 shows examples of the SAR phase-history data after being interrupted by several types of gapping patterns; the figure shows examples of a single contiguous gap, a periodic gapping pattern, and a random gapping pattern; each of these interruptions is comprised of 100 total unit-width interruptions.

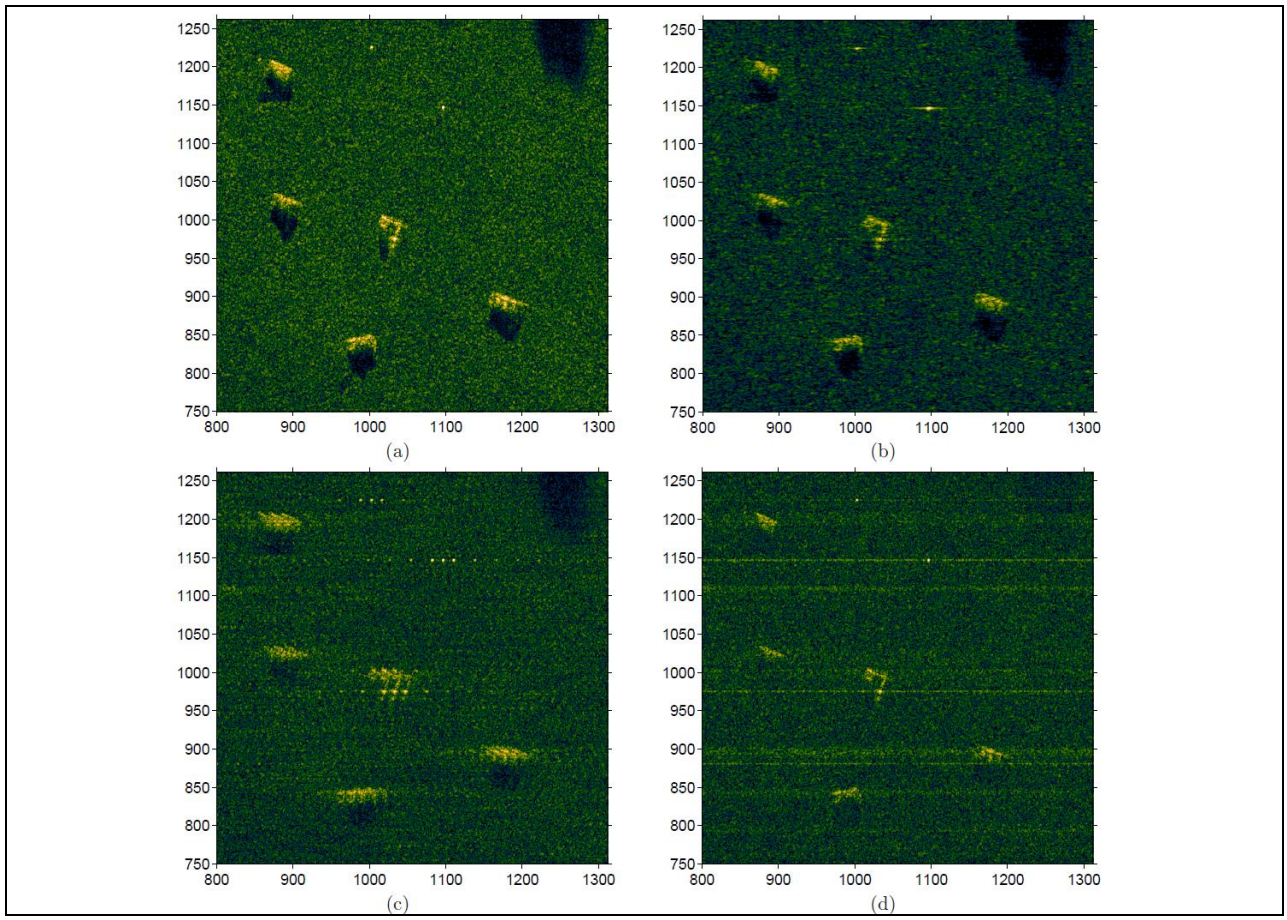

Figure 9: Matched Filter image reconstruction; (a) uninterrupted phase-history; (b) single contiguous Interruption; (c) periodic Interruptions; (d) random Interruptions.

Figure 9 shows the SAR images obtained using Matched Filter (Taylor weighted, 2D-FFT) image formation processing of the original uninterrupted phase history data (a), the single contiguous interruption (b), the periodic interruptions (c), and the random interruptions (d); (interrupt patterns shown in Figure 8). Note that the original uninterrupted SAR image contains a set of 5-military targets located in a homogeneous clutter background (a grassy field). Also visible in the image are a two good quality corner reflectors (metal spheres) – note the sharpness of these scatterers. The image quality of the images formed using matched filter processing of interrupted SAR phase history data shown in (c) and (d) of Figure 9 is significantly degraded compared to the original SAR image shown in (a) of Figure 9. For example, each of the targets in (c) of Figure 9 appears significantly blurred – and reliable recognition of the targets in this image by either an automatic algorithm or an image analyst is most likely not possible.

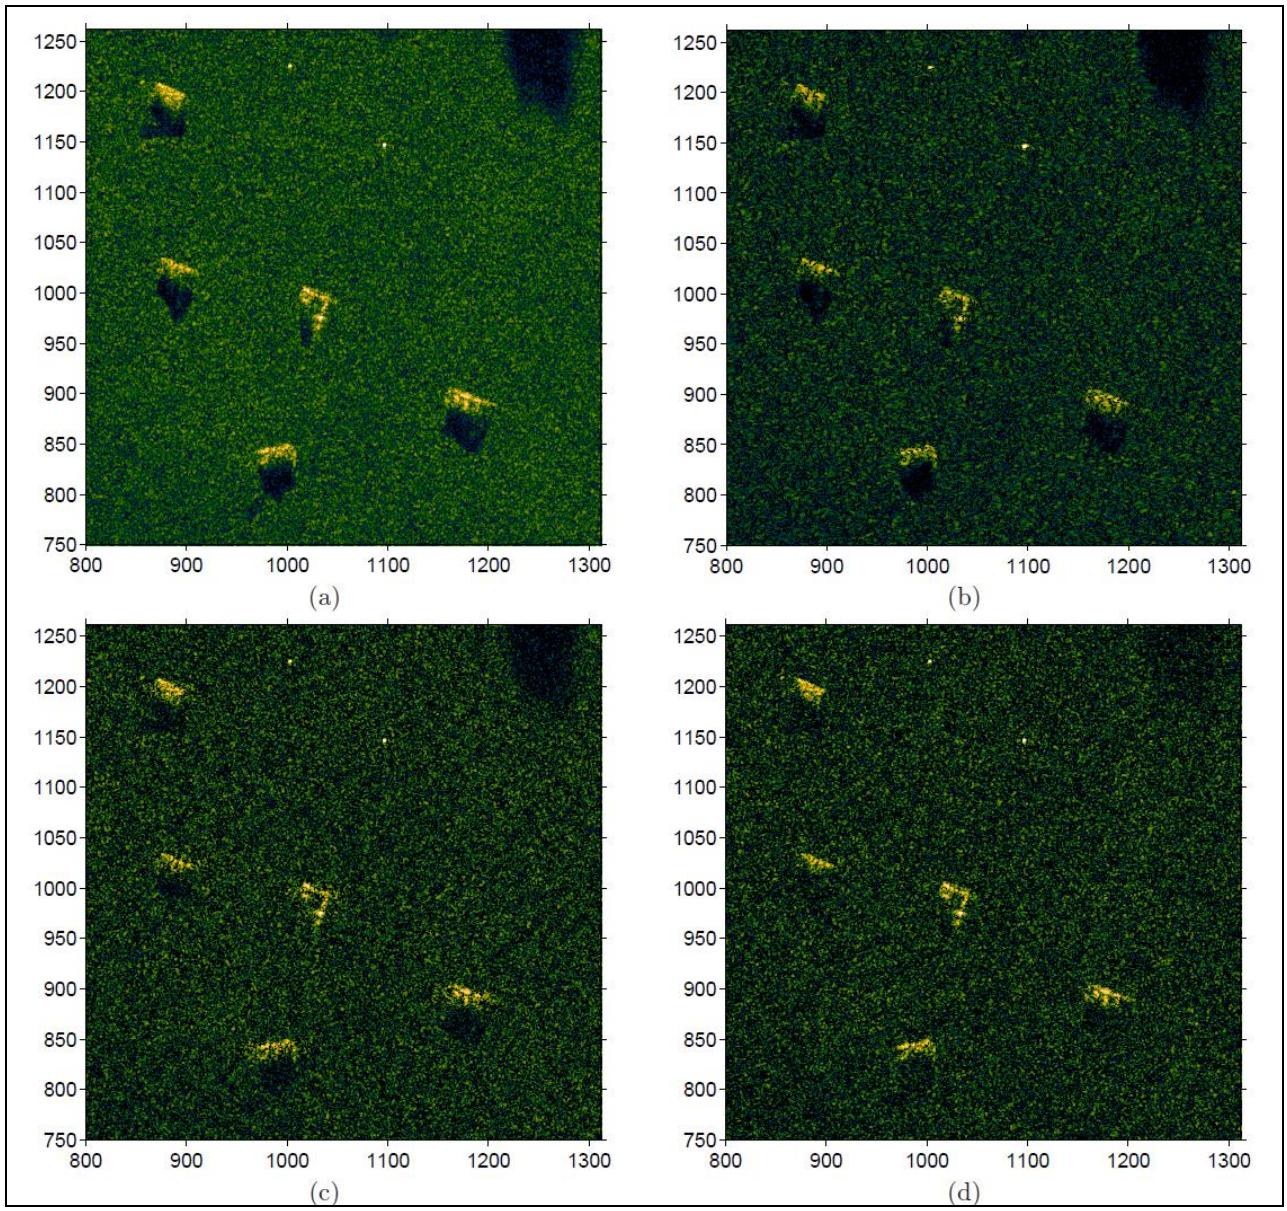

Figure 10: BPDN image reconstruction: (a) uninterrupted phase-history; (b) single contiguous interruption; (c) periodic interruptions; (d) random interruptions.

Figure 10 shows the SAR images obtained using BPDN image formation processing of the phase-history data arrays displayed in Figure 8. Note that the SAR image obtained from the BPDN-processed uninterrupted phase history data (Figure 10(a)) is essentially identical to the original uninterrupted SAR image shown in Figure 9(a) -- and the corresponding BPDN reconstructed images obtained from the interrupted phase history data arrays have excellent image quality (see b, c, and d, of Figure 10). Additional details of these studies can be found in Reference [7].

5. SUMMARY AND CONCLUSIONS

Phase gradient SAR image focusing was demonstrated to provide well-focused imagery; cross-range smearing of the imagery was significantly reduced, resulting in higher probability of correct classification as demonstrated by a 20+ target model-based classifier. High Definition Imaging was demonstrated to improve the image quality of complex SAR imagery; the effective resolution of SAR imagery was shown to be increased as demonstrated by the improved Pcc achieved by a 10-target template-based classifier. 2D-FFT image formation processing of interrupted SAR phase history data was shown to yield SAR imagery containing significant artifacts and degraded image quality; CS-based image formation processing (BPDN) was shown to mitigate these image artifacts and produced complex SAR imagery having excellent image quality.

6. REFERENCES

[1] Wahl, D. et al, "Phase gradient autofocus -- a robust tool for high-resolution SAR phase correction," IEEE Transactions on Aerospace and Electronic Systems, Vol. 30, N0. 3, July 1994.

[2] Kragh, T., "Monotonic Iterative Algorithm for Minimum-Entropy Autofocus," ASAP Conference, Lincoln Laboratory, 2006.

[3] Novak, L., et al, "Classifier Performance using Enhanced Resolution SAR Data," IEE Radar Conference, Edinburgh, UK, 1997.

[4] Benitz, G., "High Definition Vector Imaging," Lincoln Laboratory Journal, Vol. 10, No. 2, 1997.

[5] Novak, L., et al, "Automatic Target Recognition Using Enhanced Resolution SAR Data," IEEE Transactions on Aerospace and Electronic Systems, Vol. 35, No. 1, January, 1999.

[6] Chen, S., "Basis Pursuit", PhD Dissertation, Stanford University, November 1995.

[7] Stojanovic, I., Karl, W. C., and L. Novak, "Reconstruction of Interrupted SAR Imagery for Persistent Surveillance Change Detection," SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery, Baltimore, MD, April 2012.