

Database Integrity Monitoring for Synthetic Vision Systems Using Machine Vision and SHADE

Eric G. Cooper, Steven D. Young
NASA Langley Research Center, Hampton Virginia 23681

ABSTRACT

In an effort to increase situational awareness, the aviation industry is investigating technologies that allow pilots to visualize what is outside of the aircraft during periods of low-visibility. One of these technologies, referred to as Synthetic Vision Systems (SVS), provides the pilot with real-time computer-generated images of obstacles, terrain features, runways, and other aircraft regardless of weather conditions. To help ensure the integrity of such systems, methods of verifying the accuracy of synthetically-derived display elements using onboard remote sensing technologies are under investigation. One such method is based on a shadow detection and extraction (SHADE) algorithm that transforms computer-generated digital elevation data into a reference domain that enables direct comparison with radar measurements. This paper describes machine vision techniques for making this comparison and discusses preliminary results from application to actual flight data.

Keywords: Synthetic vision, pattern recognition, geo-spatial database, terrain feature extraction

1. INTRODUCTION

The concept of using forward-looking X-band weather radar (WxR) for integrity monitoring of Synthetic Vision Systems (SVS) has been proposed in [1][2]. As described in [2], the general approach requires transformation of two independent sets of information into a common reference domain. As the radar scans the environment in front of the aircraft, a Shadow Detection and Extraction (SHADE) algorithm performs these transformations and produces radial shadow state vectors at the radar measurement rate.

Shadow state vectors produced by the WxR-derived thread of SHADE result from thresholding the measurements. As the radar antenna scans, radial measurements are provided consisting of M range bin values that represent the magnitude of the reflected energy observed by the radar for the corresponding range interval. For each range interval (bin), SHADE assigns a shadow state of one if the reflectivity value is greater than zero. If the reflectivity value is zero, SHADE assigns a shadow state of zero. As a result, the WxR-derived thread of SHADE produces shadow state vectors of length M at the measurement rate of the radar. By defining a radar scan/image that consists of N radial measurements, each at a different azimuth pointing direction, these vectors can be accumulated over the duration of a scan to produce WxR-derived images. These images will consist of NM pixels. Pixel resolution will be $\Delta\alpha$ -by- ΔR , where $\Delta\alpha$ is the angular measurement resolution and ΔR is the linear range resolution of the radar, respectively.

Shadow state vectors produced by the DEM-derived thread of SHADE result from modeling expected radar behavior. This thread makes use of five inputs: (1) antenna position derived from GPS, (2) aircraft attitude derived from an Inertial Reference Unit (IRU), (3) the antenna pointing direction, (4) an antenna beam pattern model, and (5) a Digital Elevation Model (DEM). Synthetic radial vectors are generated that emanate from the antenna position. Shadow state values are assigned to range bins in these vectors based on computed depression angle to the DEM and relative line-of-sight shadowing geometries. As with the WxR-derived thread, as the radar antenna scans, radial vectors of length M are produced at the measurement rate of the radar. By defining a radar scan/image that consists of N radial measurements, each at a different azimuth pointing direction, these vectors can be accumulated over the duration of a scan to produce synthetic DEM-derived images. As with the WxR-derived thread, these DEM-derived images will consist of NM pixels and pixel resolution will be $\Delta\alpha$ -by- ΔR .

SHADE is described in detail in [2]. Example images produced by the two threads of SHADE using actual flight data are shown in Figure 1 for three unique terrain environments near Juneau, Alaska (JNU), Modesto, California (MOD), and Reno, Nevada (RNO).

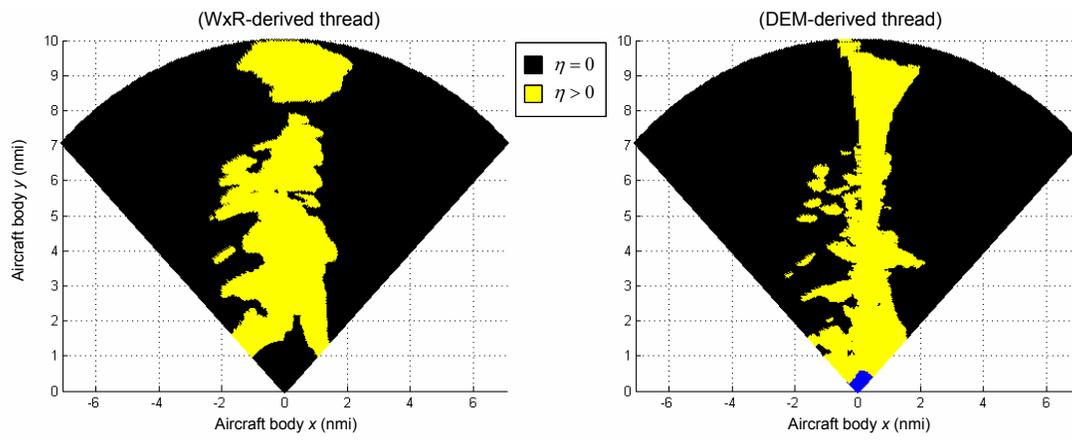


Figure 1a. SHADE images from JNU ($N: 192, M: 512, \Delta\alpha: 0.47^\circ, \Delta R: 36.2 \text{ m}$)

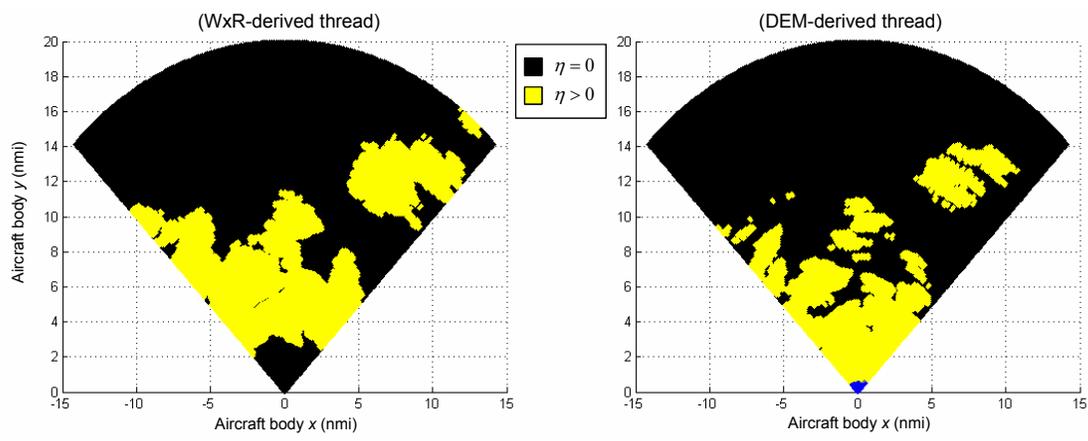


Figure 1b. SHADE images from MOD ($N: 198, M: 256, \Delta\alpha: 0.45^\circ, \Delta R: 144.8 \text{ m}$)

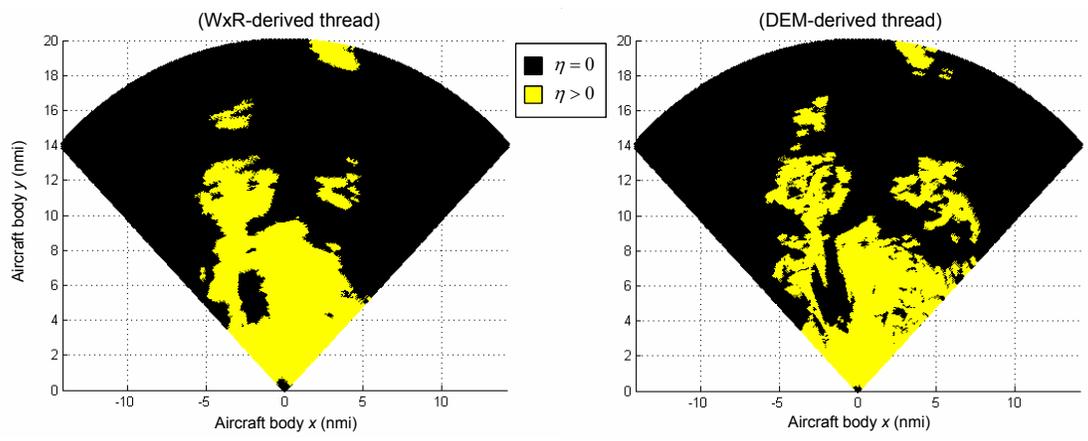


Figure 1c. SHADE images from RNO ($N: 376, M: 256, \Delta\alpha: 0.25^\circ, \Delta R: 144.8 \text{ m}$)

Given these two independent observations of the terrain environment in front of the aircraft, the monitor must compare the two during flight to establish an agreement metric. Using this metric, the monitor should be able to detect off-nominal ‘errors’ thereby increasing integrity. Specifically, this paper describes a methodology for comparing the images produced by SHADE using machine vision and pattern recognition techniques. Preliminary results from application to actual flight data are also presented.

2. APPROACH

One candidate algorithm for comparing the WxR- and DEM-derived images is shown in figure 2. In this algorithm the images are optionally preprocessed through a series of morphology operators (opening and closing) [3] to facilitate a pixel-by-pixel comparison by “filling in” gaps in the WxR image and / or filtering background clutter. After application of the morphology operators, measures of similarity are calculated and the result forwarded to a pattern classification function that establishes an integrity measure.

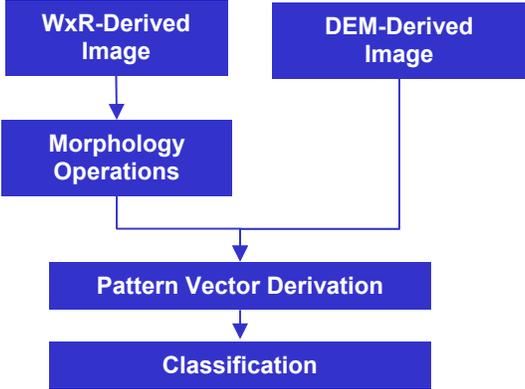


Figure 2. Algorithm for full image comparison

To facilitate machine vision analysis, the conical perspectives shown in Figure 1 are reconfigured as rectangular images. Figure 3 shows the results of this reconfiguration as applied to the RNO data shown in Figure 1c.

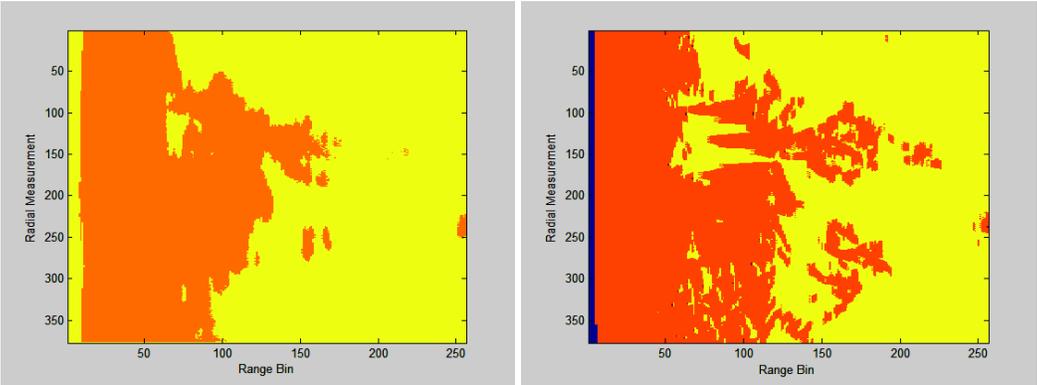


Figure 3. RNO images (a) rectangular WxR-derived scan, (b) rectangular DEM-derived scan

3. SIMILARITY MEASURES

A variety of measures can be used to assess consistency between the WxR- and DEM-derived images. One measure is the normalized correlation function [4] which provides a template-based measure as shown in equation 1.

$$r = \frac{\sum_x \sum_y [f(x,y) - \bar{f}][w(x,y) - \bar{w}]}{\sqrt{\sum_x \sum_y [f(x,y) - \bar{f}]^2 \sum_x \sum_y [w(x,y) - \bar{w}]^2}} \quad (1)$$

where $f(x,y)$ and $w(x,y)$ are the grey-level values for the pixel located at x, y for the WxR-derived and DEM-derived pixels, respectively; and \bar{f} and \bar{w} are the respective mean grey-level values.

Another measure is the mean-square error shown in equation 2.

$$mse = \frac{\sum_x \sum_y [f(x,y) - w(x,y)]^2}{n} \quad (2)$$

where $f(x,y)$ and $w(x,y)$ are the grey-level values for the pixel located at x, y for the WxR-derived and DEM-derived pixels, respectively; and n is the total number of pixels in the image.

These two measures can be complemented with a set of shape-sensitive measures based upon spatial moments, central moments, normalized central moments, and Hu invariant moments as shown in equations 3, 4, 5 and 6 [5]. Table 1 shows an algorithm for deriving the ellipse tilt angle using second moments [6].

Spatial Moment of order (p, q):

$$m_{pq} = \sum_x \sum_y x^p y^q f(x,y) \quad (3)$$

Central moment - invariant to translation:

$$u_{pq} = \sum_x \sum_y (x - \bar{x})^p (y - \bar{y})^q f(x,y) \quad (4)$$

where $\bar{x} = \frac{m_{10}}{m_{00}}$ $\bar{y} = \frac{m_{01}}{m_{00}}$

Normalized central moments - invariant to translation and scale:

$$n_{pq} = \frac{u_{pq}}{u_{00}^\gamma}, \gamma = \frac{(p+q+2)}{2} \quad (5)$$

Hu invariant moments - invariant to translation, scale, and rotation (this study only utilizes the first 3 Hu moments):

$$Hu_1 = n_{20} + n_{02} \quad (6a)$$

$$Hu_2 = (n_{20} - n_{02})^2 + 4n_{11}^2 \quad (6b)$$

$$Hu_3 = (n_{30} - 3n_{12})^2 + (3n_{21} - n_{03})^2 \quad (6c)$$

$$Hu_4 = (n_{30} + n_{12})^2 + (n_{21} + n_{03})^2 \quad (6d)$$

$$Hu_5 = (n_{30} - 3n_{12})(n_{30} + n_{12})[(n_{30} + n_{12})^2 - 3(n_{21} + n_{03})^2] + (3n_{21} - n_{03})(n_{21} + n_{03})[3(n_{30} + n_{12})^2 - (n_{21} + n_{03})^2] \quad (6e)$$

$$Hu_6 = (n_{20} - n_{02})[(n_{30} + n_{12})^2 - (n_{21} + n_{03})^2] + 4n_{11}(n_{30} + n_{12})(n_{21} + n_{03}) \quad (6f)$$

$$Hu_7 = (3n_{21} - n_{03})(n_{30} + n_{12})[(n_{30} + n_{12})^2 - 3(n_{21} + n_{03})^2] - (n_{30} - 3n_{12})(n_{21} + n_{03})[3(n_{30} + n_{12})^2 - (n_{21} + n_{03})^2] \quad (6g)$$

$u_{20} - u_{02}$	u_{11}	ϕ	Note
Zero	Zero	0	
Zero	Positive	+45°	
Zero	Negative	-45°	$\varepsilon = \frac{2u_{11}}{u_{20} - u_{02}}$
Positive	Zero	0	
Negative	Zero	-90°	
Positive	Positive	$(1/2) \tan^{-1} \varepsilon$	$(0 < \phi < 45^\circ)$
Positive	Negative	$(1/2) \tan^{-1} \varepsilon$	$(-45 < \phi < 0)$
Negative	Positive	$(1/2) \tan^{-1} \varepsilon + 90^\circ$	$(45^\circ < \phi < 90^\circ)$
Negative	Negative	$(1/2) \tan^{-1} \varepsilon - 90^\circ$	$(-90^\circ < \phi < -45^\circ)$

Table 1 – Ellipse tilt angle ϕ for various cases of the signs of the second moments [6]

4. PATTERN VECTOR DERIVATION

A pattern vector comprised of 15 similarity measures is used to establish a measure of consistency between the DEM-derived images and the corresponding WxR-derived images. These similarity measures were selected for their sensitivity to certain translation, rotation, and scale differences that may result from errors such as imprecise sensor calibration, reversed / mis-ordered navigation data, unit conversion errors, random errors, voids, and others. Table 2 shows the 15 measures included in the pattern vector.

Vector element	Symbol	Description
Normalized correlation	r	See Equation (1)
Mean-square error	mse	See Equation (2)
Area	m_{00}	See Equation (3)
Centroid	\bar{x}	See Equation (4)
Centroid	\bar{y}	See Equation (4)
Ellipse Tilt Angle	ϕ	See Table 1
Shape	m_{20}	See Equation (3)
Shape	m_{02}	See Equation (3)
Circular deviation	n_{11}	See Equation (5)
Shape	n_{02}	See Equation (5)
Shape	n_{20}	See Equation (5)
Shape	n_{22}	See Equation (5)
Shape	Hu_1	See Equation (6a)
Shape	Hu_2	See Equation (6b)
Shape	Hu_3	See Equation (6c)

Table 2 – Pattern vector elements

5. CLASSIFICATION

As shown in Table 2, the first two elements in the pattern vector are the normalized correlation and mean-square error. The next 13 elements are determined by computing the aforementioned moments for the WxR- and DEM-derived images and populating the remaining pattern vector elements with the absolute difference of the moment values. A two-category discriminant function [7] classifies each WxR- and DEM-derived image pair as being either consistent or anomalous (an error is present) by calculating the scalar product of the pattern vector x and a pre-defined weighting vector w as shown in equation 7.

$$d = w_1x_1 + w_2x_2 + \dots + w_{15}x_{15} \quad (7)$$

The weighting vector w is derived through a least squares error algorithm supplied with training sets (pattern vectors) that contain both real and simulated data. The classification training labels are either “1”, meaning the vector belongs to the class of non-matching cases, or “-1”, meaning the vector does not belong in the class of non-matching cases. This provides for a simple decision rule – if d is less than zero, the WxR- and DEM-derived images are considered consistent, otherwise an off-nominal error is present.

6. TEST RESULTS

Preliminary results were obtained by applying the methodology to the SHADE-processed flight data collected in JNU and RNO. The JNU sample is comprised of 90 WxR scans taken while flying down the Gastineau channel with terrain above the aircraft altitude on both sides of the flight path. The aircraft flies straight-and-level at about 800 ft above the channel floor. During the end of the segment, a slight left-hand turn occurs. The RNO sample is comprised of 63 WxR scans taken on final approach to runway 16R at the Reno-Tahoe International Airport.

Additionally, the RNO sample was augmented with cases of intentionally-biased DEM data. Specifically, corrupted samples were created by injecting along-track, cross-track, and up/down biases into the reference DEM and re-generating the DEM-derived SHADE images. Along-track and cross-track biases ranged from -10 range bins (-1448 meters) to +10 range bins (+1448 meters). Up/down biases ranged from -1.0 range bins (-144.8 meters) to +1.0 range bins (+144.8 meters).

For along-track and cross-track discriminant function training purposes, the first 34 scans from the reference RNO data along with 6 scans (out of a total of 90) from the reference JNU data were labeled as “non-anomalous” (i.e. fault-free). The anomalously labeled data included the extreme bias cases for along-track and cross-track errors and six artificially rotated JNU sets. Also included was scan 59 from the RNO set, which shows an erroneous all-white DEM-derived image. For reasons that will be discussed later, the data sets that were biased up/down were not included in the training data. Figure 4 shows the results of applying the machine vision method to scans 1-5 of the RNO data after inserting the along-track biases.

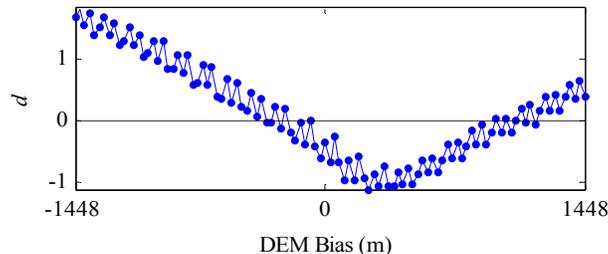


Figure 4. Classification of RNO data in the presence of along-track DEM biases, scans 1 – 5

Considering that figure 4 contains the first 5 scans from each test case, it is apparent that the maximum “non-anomaly” case was not at zero bias as expected, but closer to +180 meters. This suggests that there may be a synchronization error between the recorded WxR data and the position and attitude data recorded and used by the DEM-derived thread. The plot does, however, show a positive correlation between the amount of bias (error) and the results from the classification function. Recall the classification rule used for these initial tests is that when the classification function is found to be

greater than zero, the comparison is deemed “anomalous”, whereas when the classification function is found to be less than zero, the WxR- and DEM-derived images are deemed to be in agreement.

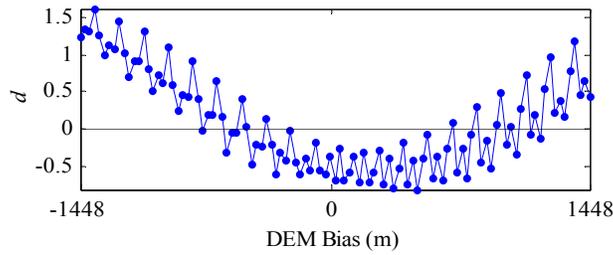


Figure 5. Classification of RNO data in the presence of cross-track DEM biases, scans 1 – 5

Figure 5 shows the results of applying the machine vision method to scans 1-5 of the RNO data after inserting the cross-track biases. Similar to the along-track bias cases, the performance in the presence of cross-track biases also appears to show the maximum “non-anomaly” case closer to +180 meters. Also as with the along-track bias cases, the cross-track results shown in Figure 5 exhibit a positive correlation between the degree of bias (error) and the results from the classification function. The technique was then applied to the JNU data – the first scan is shown in Figure 6.

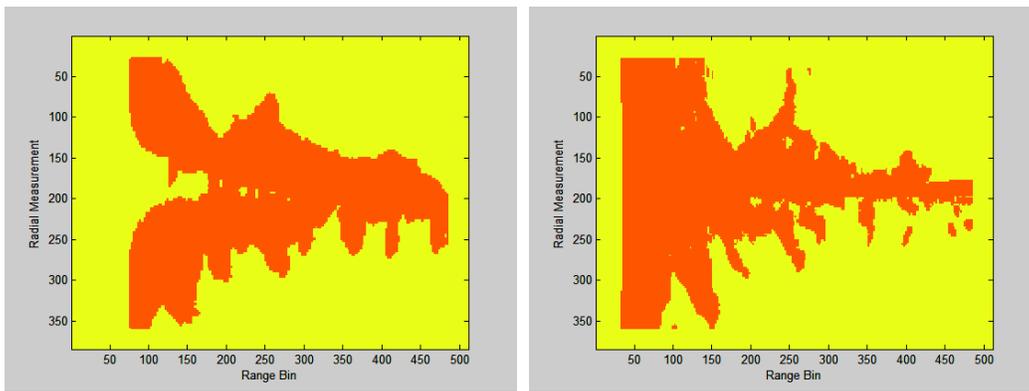


Figure 6. JNU images (a) WxR-derived Scan, (b) DEM-derived Scan

The JNU WxR- and DEM-derived images are presumed to be reasonably correct with regards to along-track, cross-track, and up-down bias. Using the same training set as was used for the RNO analysis, Figure 7 shows eight false positives – six relatively close to zero (scans 18, 19, 21, 23, 24, and 25) and two that are greater than one (scans 89 and 90).

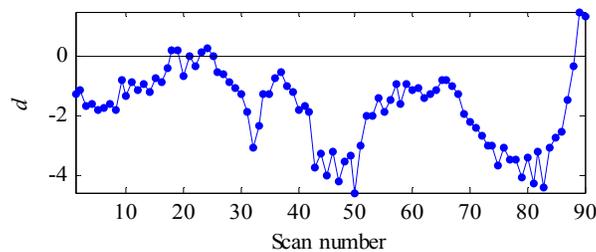


Figure 7. Classification of JNU scans 1 – 90

The detected anomalies in the last two scans may be due to excessive ground clutter in the immediate area or measurement artifacts induced by the aircraft heading change at the end of the flight segment. For example, Figure 8 shows the image pair for the final scan (scan 89 is very similar in appearance).

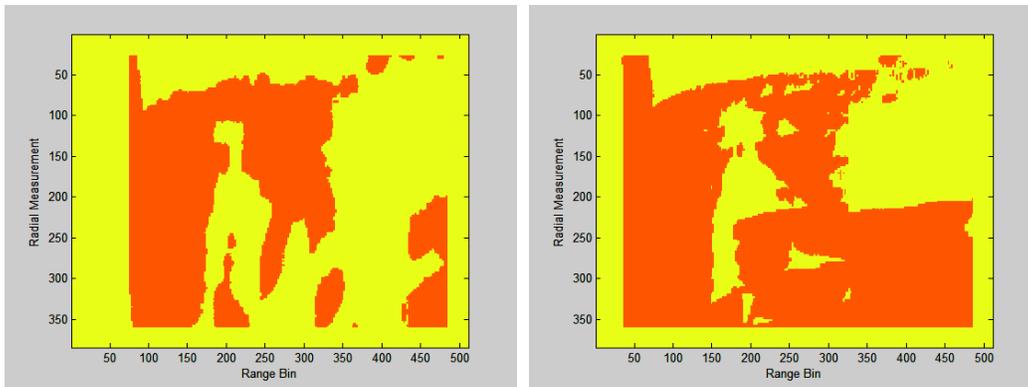


Figure 8. JNU Image for scan 90 (a) WxR-derived scan, (b) DEM-derived scan

Although some similarities can be seen in the Figure 8 images, there are large sections that are obviously different. The other six false-positives, grouped around scan 20, do not show such obvious differences. Although the classification values for these six scans are relatively close to zero, further investigation is necessary to determine the source of the discrepancy.

Figure 9 shows the classification results using the training data described above on the RNO up/down bias data. Notably all bias cases are classified as non-anomalous.

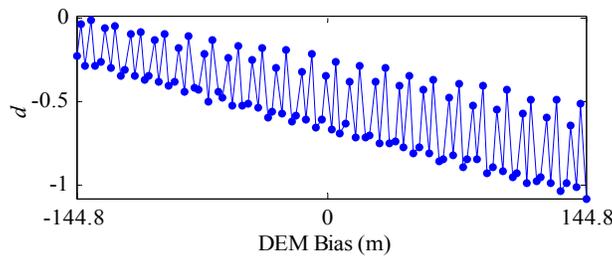


Figure 9. Classification of RNO data in the presence of up/down DEM biases, scans 1 – 5

The up/down biases essentially result in scaling differences in the images that, given the relative magnitude of the bias (<144.8 m), results in a relatively small change in appearance. An early attempt at including these up/down bias cases in the training set resulted in poor discriminant function performance on the cross-track and along-track bias cases. Figure 10 shows the impact of a 144.8 m up bias on the DEM-derived image.



Figure 10. RNO images (a) no vertical bias, (b) 144.8 m up bias, and (c) difference

Figure 10c shows the difference image resulting from a subtraction of the non-bias image from the biased image. The majority of the difference image consists of speckles comprised of a few pixels each. The left side of Figure 10c shows a few larger difference regions, but as a percentage of the full DEM-derived image these regions are relatively small.

Finally the technique was augmented to include four additional pattern discriminators – negative along-track bias, positive along-track bias, negative cross-track bias, and positive cross-track bias. The same training data was used for these four new discriminators as was used earlier for the single discriminator except that the anomalous label “1” was set for only the training series applicable to the particular case and the artificially rotated JNU cases were eliminated. As an example, only the data for the extreme negative along track bias was labeled anomalous for the negative along-track discriminator.

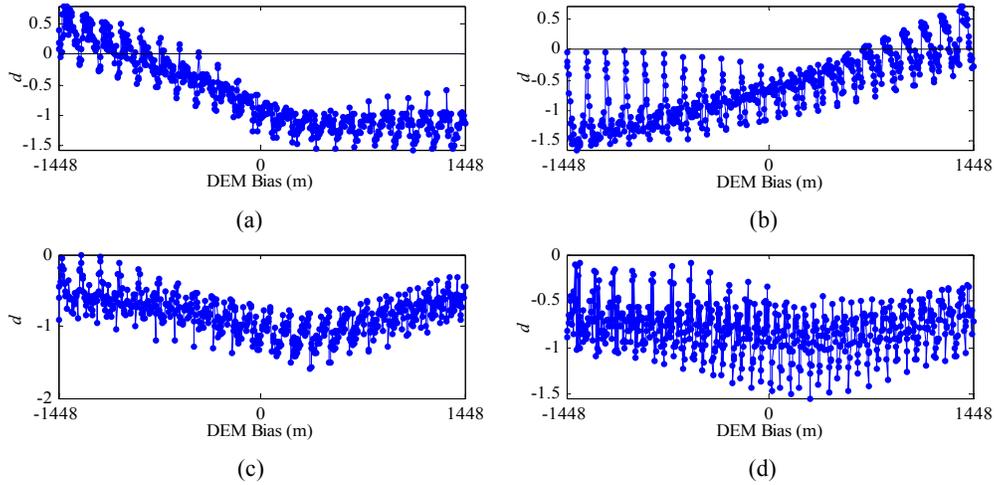


Figure 11. Discriminator results from RNO in the presence of DEM biases, scans 1 – 30 (a) negative along-track, (b) positive along-track, (c) negative cross-track, (d) positive cross-track

Figure 11 shows the results from each of the four added discriminators for RNO along-track biased data, scans 1 – 30. Figures 11a and 11b show trends consistent with the along-track test cases that progress from negative bias to positive bias. The results from the cross-track discriminators, shown in figures 11c and 11d, are generally below the zero threshold value.

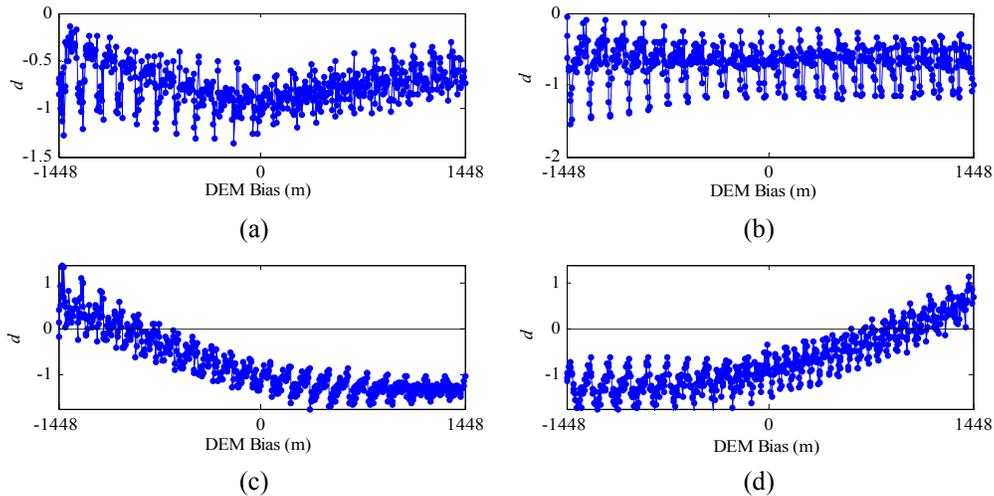


Figure 12. Discriminator results from RNO in the presence of DEM biases, scans 1 – 30 (a) negative along-track, (b) positive along-track, (c) negative cross-track, (d) positive cross-track

Figure 12 shows the results for RNO cross-track biased data, scans 1 – 30. In similar fashion to the along-track results shown in Figure 11, Figure 12 shows that the cross-track detectors provide trending consistent with the cross-track bias data and the along-track detectors computing values that fall below the zero threshold.

7. CONCLUSION AND FUTURE WORK

This paper presented a method that employs the SHADE algorithm along with machine vision and pattern analysis techniques for detecting anomalies (i.e. off-nominal behavior) between the onboard radar-derived images and synthetically-derived images. The method was applied to sequences of data recorded during flight tests conducted in Reno, Nevada, and Juneau, Alaska. Preliminary results suggest that this approach may prove useful for monitoring the integrity of synthetic vision systems. Moreover, results show that it may be possible to not only provide a measure of consistency between WxR- and DEM-derived images, but to also provide insight into the nature of the discrepancy (e.g. along-track and/or cross-track bias detectors).

Two other algorithms are planned for investigation - the first will extract individual features from WxR-derived images that are then compared, on a region-by-region basis, with the DEM-derived images using the same similarity measures and pattern classifier operations discussed previously in this paper. The other algorithm to be considered does not compare the images, but rather tracks features seen within the images over repeated scans. This approach exploits the fact that, for reasonably fast update rates, extracted features should be traceable across image sequences and consistent with the aircraft's position as derived from onboard navigation systems. Using this approach, "open loop" integrity monitoring can occur at the SHADE frame rate, while full registration with onboard WxR, which is more computationally demanding, can be performed in a parallel task executing at a relatively slower frame rate. By providing continuous open loop integrity monitoring, the parallel task can be scheduled according to the slower scan rate of the WxR radar, or at a rate that provides acceptable "closed loop" comparison.

REFERENCES

1. Young, S., Uijt de Haag, M., and Sayre, J., "Using X-band Weather Radar Measurements to Monitor the Integrity of Digital Elevation Models for Synthetic Vision Systems," in Proceedings of SPIE Vol. 5081 *Enhanced and Synthetic Vision 2003*, edited by Jacques G. Verly, SPIE, Bellingham, WA, 2003, pp. 66-76.
2. Young, S., Kakarlapudi, S., and Uijt de Haag, "A Shadow Detection and Extraction Algorithm Using Digital Elevation Models and X-Band Weather Radar Measurements," *International Journal of Remote Sensing*, in press 2005.
3. Haralick, R. M., Sternberg, S. R., Zhuang, X., "Image Analysis Using Mathematical Morphology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-9, No. 4, July 1987
4. Lewis, J. P., "Fast Normalized Cross-Correlation," *Vision Interface*, 1995
5. Hu, Ming-Kuei, "Visual Pattern Recognition by Invariant Moments," *IRE Transactions on Information Theory*, Volume: 8, Issue: 2, Feb 1962
6. Teague, M. R., "Image Analysis via the General Theory of Moments," *Journal Optical Society of America*, Vol 70, No. 8, 1980
7. Duda, R. O., Hart, P. E., Stork, D. G., *Pattern Classification*, 2nd edition, John Wiley & Sons, Inc. 2001