SAND2016-2159C

Index for Surface Coherence (ISC): A Method for Calculating Change Susceptibility

Jonathan Tran[1]

[2]Sandia National Laboratories, P.O. Box 5800, MS0512, Albuquerque, NM 87185-116, USA

ABSTRACT

Coherent change detection (CCD) provides a way for analysts and detectors to find ephemeral features that would otherwise be invisible in traditional synthetic aperture radar (SAR) imagery. However, CCD can produce false alarms in regions of the image that have low SNR and high vegetation areas. The method proposed looks to eliminate these false alarm regions by creating a mask which can then be applied to change products. This is done by utilizing both the magnitude and coherence feature statistics of a scene. For each feature, the image is segmented into groups of similar pixels called superpixels. The method then utilizes a training phase to model each terrain that the user deems as capable of supporting change and statistically comparing superpixels in the image to the modeled terrain types. Finally, the method combines the features using probabilistic fusion to create a mask that a user can threshold and apply to a change product for human analysis or automatic feature detectors.

Keywords: false alarm reduction, SAR, coherent change detection, probabilistic fusion, superpixel neighborhoods

1. INTRODUCTION

Synthetic aperture radar (SAR) coherence change detection (CCD) is a very powerful mechanism for detecting very small changes in a scene that would otherwise be naked to the human eye when looking at a traditional SAR image. By utilizing both the magnitude and phase difference,[1] sub-wavelength changes can be detected. These small changes are very useful in identifying temporal change caused by human activity, disturbances such as tire tracks from driving,[1] or even bipedal travel.[2] This sensitivity does not come without a cost, however. Regions containing vegetation and low or no return areas manifest as disturbances.[3] These disturbances can be very distracting when passed to an analyst or an algorithm for ephemeral feature detection. To mitigate these disturbances, many false alarm mitigation techniques have been developed. These techniques utilize clutter to noise ratio (CNR),[2] log-likelihood estimation of the noise[4,5] or a maximum-likelihood derivation of the CNR.[6] Other techniques take advantage of three or more passes[7,8] While these do offer good suppression of redundant disturbance information and areas with inherent low coherence, they tend to have problems in completely eliminating low coherence regions (causing additional speckle behavior that could be interpreted by a detector as human activity), require a certain range of CNR, or in the worst cases, they erode the disturbance of interest. The approach used in this method tries to eliminate the regions of false alarms by modeling and combining information from both the coherence and SAR magnitude domains. This allows a user to choose the terrain types to suppress and thus relies less on a system's performance parameters.

The method proposed utilizes median CCD and median SAR radar cross section (RCS) products to provide a stable observation of the scene. The median products are then segmented using the SLIC superpixel segmentation (SPS)[9,10] to produce pixels with similar statistics. With the superpixels and median products, the method looks to model terrain that a user defines as capable of supporting change. After a training process to characterize the models of each of the terrains, the algorithm uses these models to perform statistical comparison of each superpixel in the image. The results are then combined using probabilistic fusion.

[1] Contact jtran@sandia.gov for any questions or additional information

[2] Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000

(a) Test image - median RCS product

(b) Test image - median CCD product

Figure 1. Test image feature products

Despite the grouping from superpixel segmentation, the result can still exhibit a speckle appearance from a superpixel to superpixel basis. However, this can be alleviated by incorporating a superpixel's neighbors in evaluating its distribution. By including a superpixel's neighbors, spatial stability and greater context is achieved.

2. SAR PRODUCTS

The SAR CCD and RCS products used in this study were collected using a Ku-Band radar with the parameters shown in Table 1. The subsequent products were all derived from these products. The test images are subsections of the whole images with the chips being about 1200 X 1200 pixels.

Table 1. Radar Parameters

Radar Parameter	**Value**	**Radar Parameter**	**Value**
Frequency	16.8 GHz	Wavelength	0.018 m
Grazing Angle	40.3°	Squint Angle	-89.7°
Range Resolution (ground)	5.5 in.	Azimuth Resolution (ground)	5.1 in.
Image Size (pixels, row X column)	3459 X 4864	Range	2.9 km

2.1 Median Products

When using the CCD and RCS products, each individual image will show all activity that has changed from the previous scene. While this is normally the desired effect, suppressing false alarms requires a more stable observation of the scene and identifying changes that are always apparent in the scene. This has an effect of mitigating some of the speckle and bolstering the static features. The temporal effect of using median products also eliminates the problem of suppressing activity that could be of interest. To create these products, a stack of RCS or CCD images is collected, a pairwise registering is performed,[11] and then the median for each pixel in the stack is found. An example of a median RCS and CCD product are shown in Figure 1a and Figure 1b, respectively. These are also the images used in this study.

2.2 Superpixel Segmentation

To classify different terrain types in a scene, a segmentation must be done to group pixels with similar statistics. In this method, superpixel oversegmentation is used. Superpixels captures image redundancy

thus reducing the complexity necessary for operating on the pixel grid[12].[9] The method used for segmenting the image is the SLIC superpixel segmentation, which allows a user to define how compact the superpixel appears and the number of superpixels in the image therefore creating an almost uniform grid of pixel groups[9].[10] This approach provides pixel groups for statistics evaluation and ensures the statistics of those pixels groups will be similar. A truly uniform segmentation would provide pixel groups and reduce the computing complexity, but the pixels of those groups would be visually and statistically very dissimilar.

3. TERRAIN FEATURE IDENTIFICATION

Identifying regions of change support entails identifying a wide enough variety of terrain features that support change. To do this, several images across an area should be used to ensure diversity. In the route chosen for this study, 16 different terrain types were identified, but only eight were determined to support change. As seen in Figure 2a - 2d, the terrain chosen for training includes soil, gravel, and desert (indicated by cooler colors: blue, green, teal, etc.; warmer colors indicate superpixels that do not support change). The three terrain types are further broken into subcategories, however. This is due to variation in desert, soil, and gravel across a route. For example, the desert in one patch may look very different in signature to the desert in another patch, but they are both regions that ultimately support change. As a result, they should both be modeled and grouped separately. While regions that do not support change were selected and grouped, they do not undergo the same modeling and training process. They will be used to see how well the classes that were modeled discriminate data that is not considered a target. The justification for this will be explained further in section 5.

After selecting and grouping the image appropriately, 10% of each of the superpixels groups are taken to represent the terrain class. Using the formula for minimum sampling size

$$n = (\frac{Z_a \sigma}{E})^2 \tag{1}$$

where Z_a is the Z-score for a given confidence interval, σ is the population standard deviation, E is the margin of error, the size of a superpixel is about 600 pixels, and at least 200 superpixels were used to train a feature, 10% is more than enough to statistically represent the feature.

The pixels in each of the 10% superpixels are then binned into histograms and fit with distribution curves. To find the best representation of the data, the data is scaled and fit with a variety of different distribution types. The best scaling and distribution type, in a minimum mean square error (MMSE) sense, are saved for each feature class. As of now, the distribution types used are Gaussian, Beta, Gamma, Log-Normal, Exponential, and Rayleigh. While some of data may not fall into one dominant distribution type, the complexity avoided by not assuming a mixed distribution is significant. Additionally, the data chosen should be grouped in such a way that multiple modes are not apparent in the distribution. If the data does look like a mixture of different distribution types, the data may need to be broken into separate classes. The results for fitting the classes used in this study are shown in Figure 3 - 10.

4. STATISTICAL FEATURE COMPARISON

Once representations of the terrain classes have been found, statistical comparisons for the rest of the image can be made. To do this, every superpixel in the image is scaled and fit with the distribution associated with each terrain class. These distributions can then be compared using Kullback-Leibler Divergence.

4.1 Kullback-Leibler Divergence

Kullback-Leibler (KL) Divergence is a simple method for measuring the distance between two curves. The equation for evaluating the KL divergence is given by

$$D_{KL}(P||Q) = \sum_k P(k) \ln \frac{P(k)}{Q(k)} \tag{2}$$

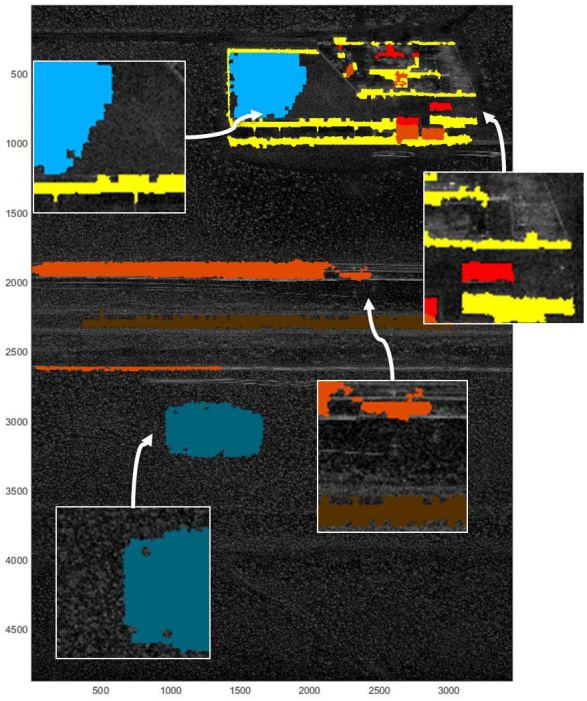

(a) Training patch 1 - Training superpixels w/zoomed overlays

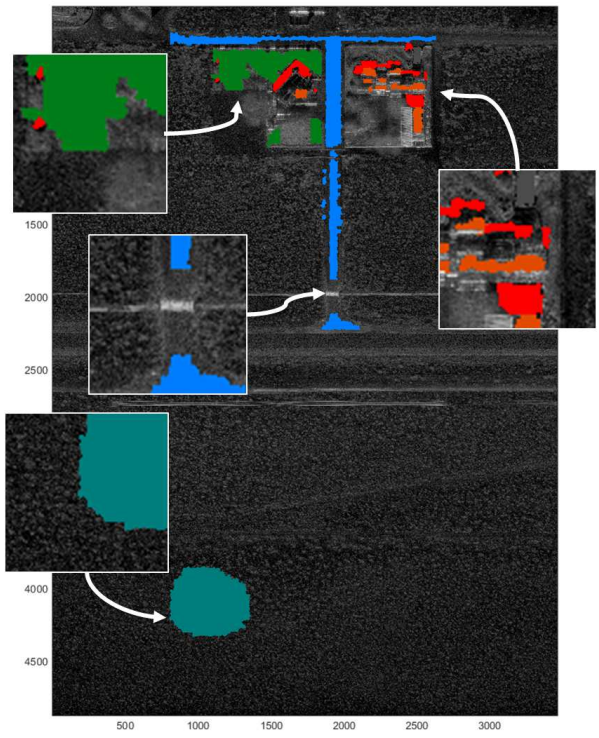

(b) Training patch 2 - Training superpixels w/zoomed overlays

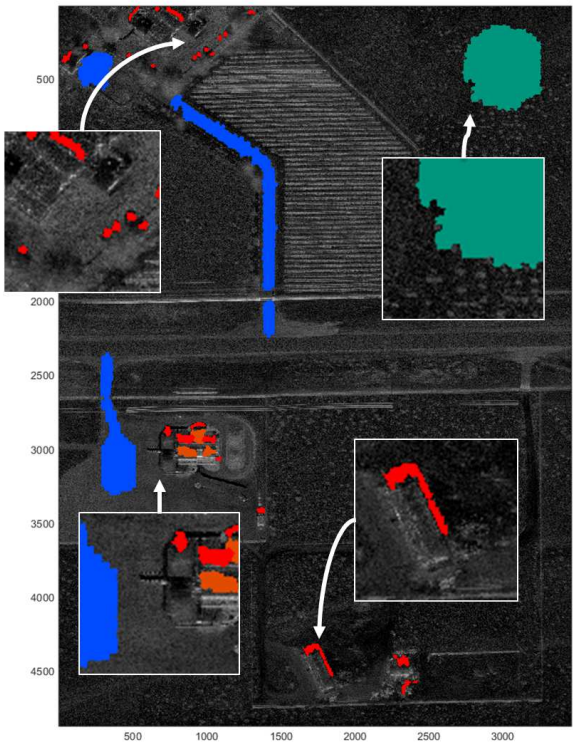

(c) Training patch 3 - Training superpixels w/zoomed overlays

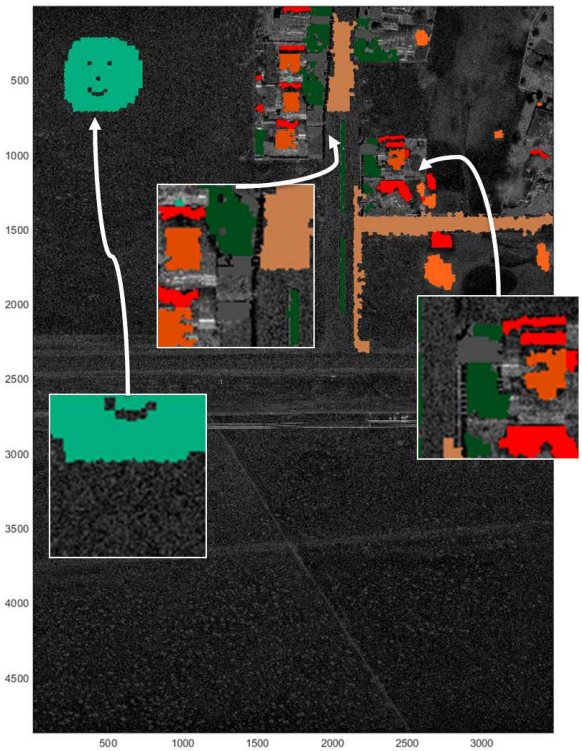

(d) Training patch 4 - Training superpixels w/zoomed overlays

Figure 2. Training superpixels selected across a variety of patches

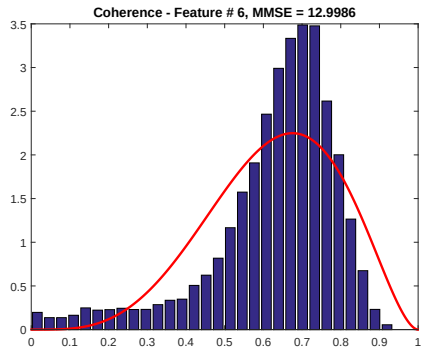

(a) Coherence Feature #6

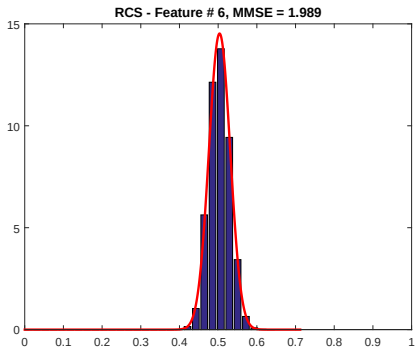

(b) Magnitude Feature #6

Figure 3. Feature Curve Fit for Class 6

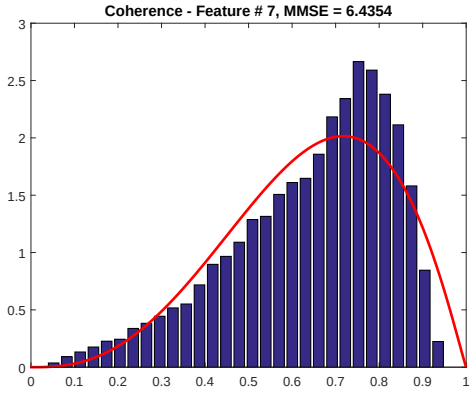

(a) Coherence Feature #7

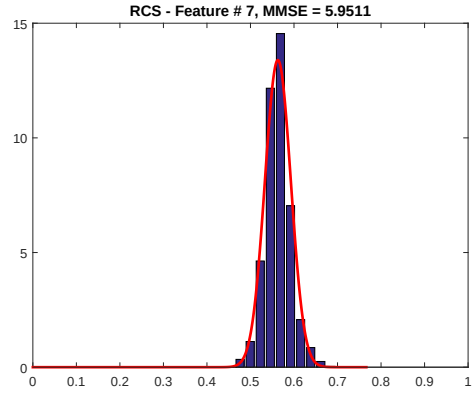

(b) Magnitude Feature #7

Figure 4. Feature Curve Fit for Class 7

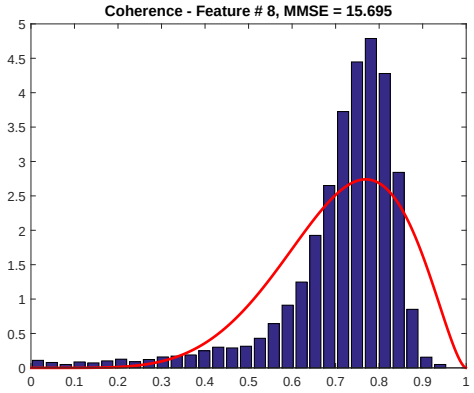

(a) Coherence Feature #8

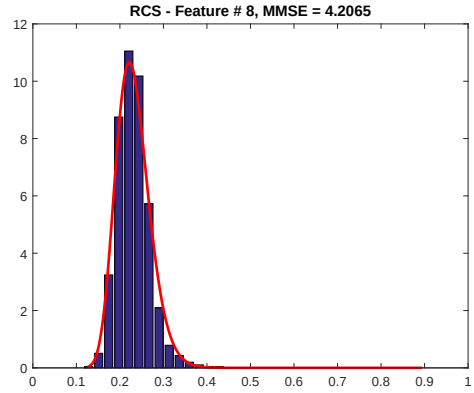

(b) Magnitude feature #8

Figure 5. Feature Curve Fit for Class 8

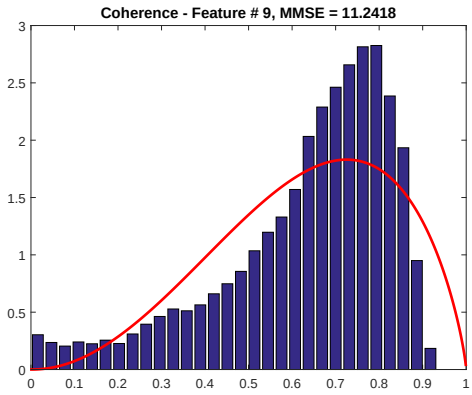

(a) Coherence feature #9

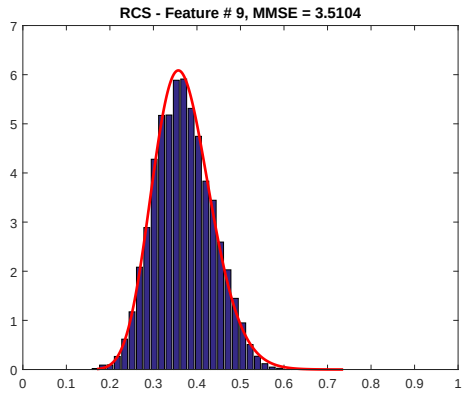

(b) Magnitude feature #9

Figure 6. Feature Curve Fit for Class 9

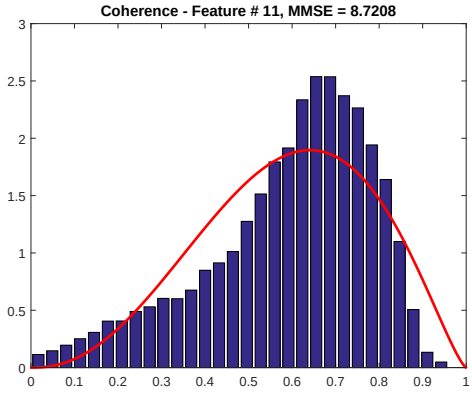

(a) Coherence feature #11

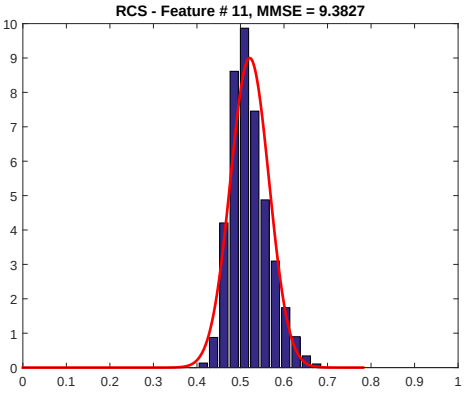

(b) Magnitude feature #11

Figure 7. Feature Curve Fit for Class 11

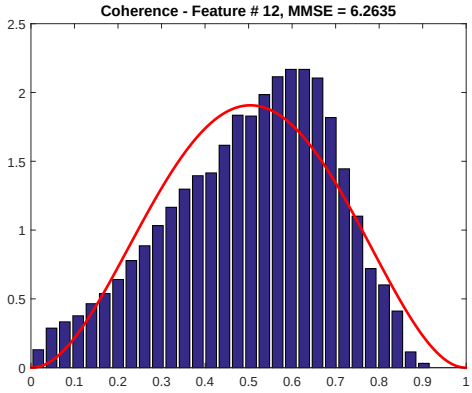

(a) Coherence feature #12

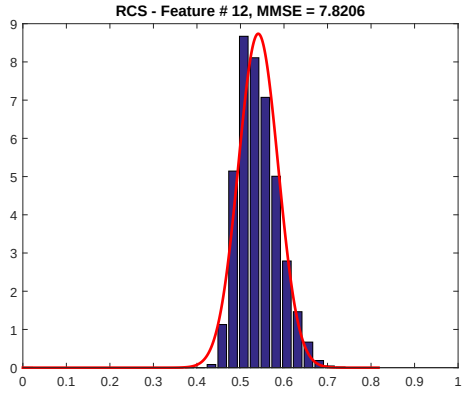

(b) Magnitude feature #12

Figure 8. Feature Curve Fit for Class 12

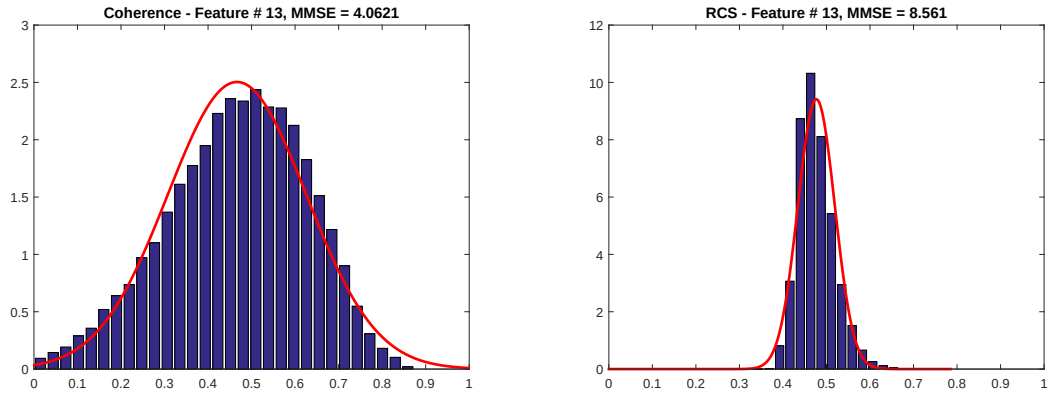

(a) Coherence feature #13 (b) Magnitude feature #13

Figure 9. Feature Curve Fit for Class 13

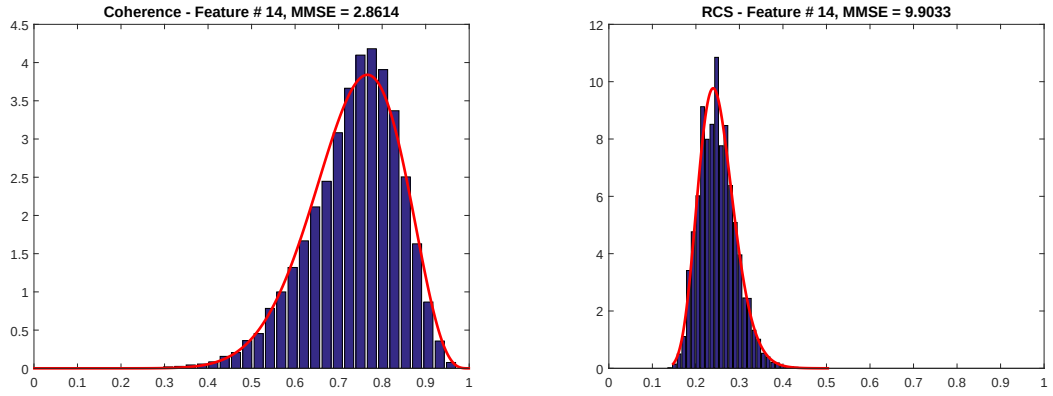

(a) Coherence feature #14 (b) Magnitude feature #14

Figure 10. Feature Curve Fit for Class 14

where $P(k)$ is the probability density function (PDF) of target curve and $Q(k)$ is the PDF of the curve to evaluate. Another popular way of interpreting KL divergence is energy or the number of extra bits necessary to encode the PDF, Q, given the coding scheme is optimized for P. Yet another way of interpretation is the amount of information lost when using Q to approximate P. All of these interpretations are relevant and give better confidence in using this method for comparing the target curve to each superpixel curve. However, it is must be noted that KL divergence is not symmetric. This means that $D(P||Q)$ is not equal to $D(Q||P)$ and the designation of the target curve is important. To relieve this, Equation 3 is used:

$$D_{KL} = \frac{1}{2}D_{KL}(P||Q) + \frac{1}{2}D_{KL}(Q||P) \tag{3}$$

Since no information about the data is known a priori, each superpixel is assumed to belong to every terrain class. As a result, each superpixel is scaled and fit with the distribution associated for each terrain class. These distributions are then compared to the ideal distribution (distributions found with the training data) of each class using Equations 2 and 3. The result is higher KL scores for superpixels that do not match the terrain class and lower scores for those that do.

5. PROBABILISTIC FUSION

After finding KL scores for each superpixel of each feature (magnitude and coherence), a method for combining these features is necessary. Probabilistic fusion is such a method. Probabilistic fusion uses a statistical approach to combine features from different sources to bolster the identification of a target[13,14] This method

requires a distribution curve that represents the KL scores of target superpixels. From the curve, a distribution of probabilities can be achieved using Equation 4

$$F_i(z_i) = 1 - Prob(Z_i \leq z_i) \tag{4}$$

where Z_i is a random variable (distribution curve found previously) representing the i^{th} feature scores and $F_i(z_i)$ is the probability that the i^{th} score will take on a value of z_i or less. From this equation, it is easy to see that the probability values will fall between $[0, 1]$ where a value of 1 indicates a superpixel matches the target characteristic perfectly and a value of 0 represents a non-target. For a distribution curve that follows the data very well, the probabilities will have a uniform distribution. When the data does not track the distribution curve, the histogram distribution of the data will show a peak near 0 indicating none of those scores indicate a target like score.

Further, the probabilities can be mapped using the following equation:

$$Y_i = Y_i(Z_i) = -\log(F_i(Z_i)) \tag{5}$$

By mapping, the thresholding for non-targets becomes much more intuitive with non-target pixels skewing towards infinity and target superpixels skewing towards zero. The next step of this method is the merging. This is accomplished by simply summing all of the Y_i variables over i. The fused measure is then defined as such:

$$S_f = \sum_{i=1}^{N} Y_i \tag{6}$$

where N is the number of features evaluated (in this case, two: coherence and magnitude). This extra mapping also makes the target probability distribution have an exponential distribution. Because the sum of two exponential distributions results in gamma distribution, the combined training data scores can then be fit with a gamma distribution and used to acquire a statistically significant threshold. However, this last step is not necessary as a simple Otsu thresholding will suffice.

For this application, the KL-score distribution is achieved by obtaining the scores of the superpixels selected in the training process that were not used to characterize the terrain classes, as mentioned in Section 3. The scores are then binned into histograms and autonomously fit using the same method mentioned in Section 3. The results of the fit and the corresponding probability distribution is seen in Figures 11 - 18. In addition, these figures also show how non-target superpixels will appear and their probability distributions when compared to each classes KL distribution. The non-target superpixels were chosen by selecting superpixels that are mostly likely to cause false alarms or not support change (vegetation, buildings, etc.; indicated by the warmer colors in Figure 2a - 2d). Inspection of the figures show that the distribution of targets match the data fairly well, and as a result, the probability distributions are fairly uniform. The non-target distributions do not match the data very well and subsequently the probabilities have distributions with peaks at 0.

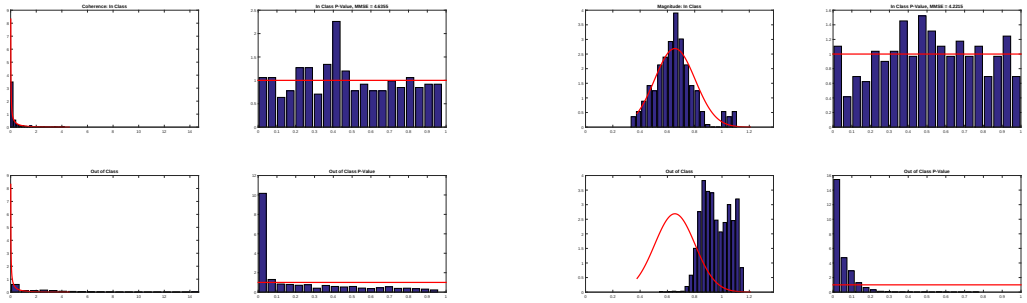

(a) Coherence (b) Magnitude

Figure 11. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 6

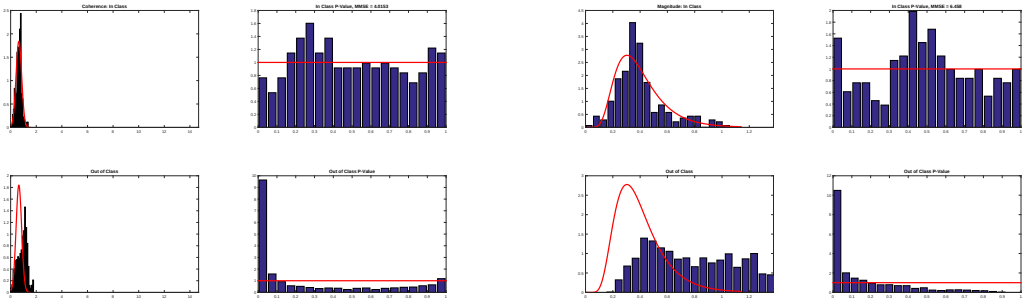

(a) Coherence (b) Magnitude

Figure 12. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 7

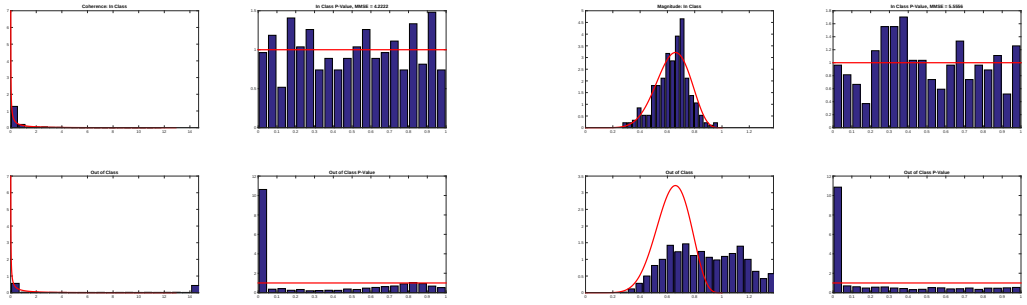

(a) Coherence (b) Magnitude

Figure 13. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 8

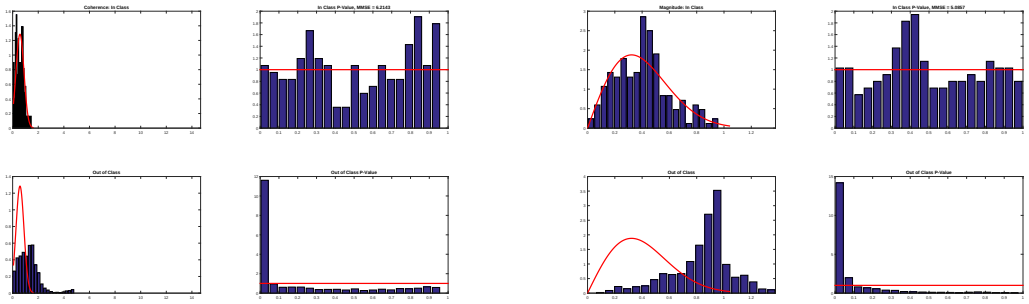

(a) Coherence (b) Magnitude

Figure 14. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 9

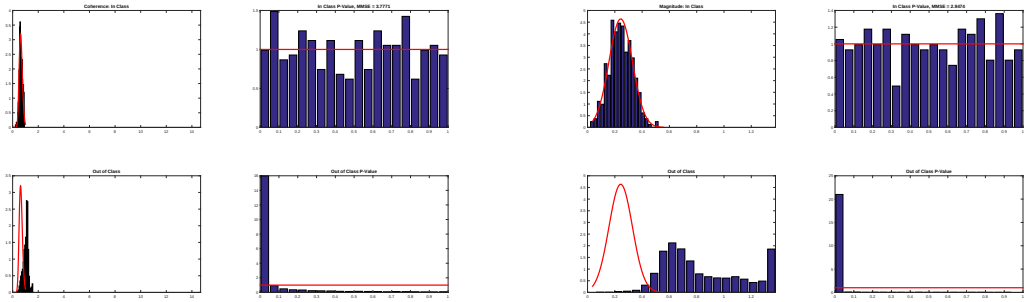

(a) Coherence (b) Magnitude

Figure 15. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 11

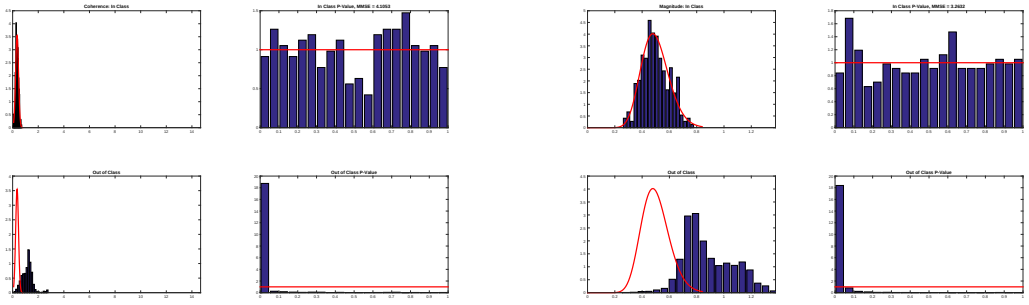

(a) Coherence (b) Magnitude

Figure 16. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 12

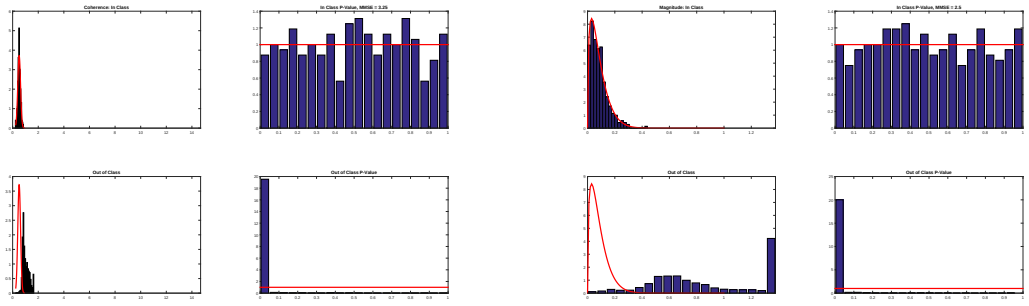

(a) Coherence (b) Magnitude

Figure 17. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 13

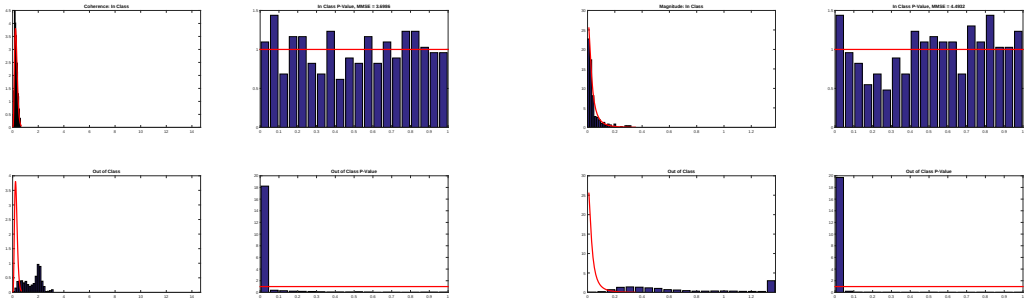

(a) Coherence (b) Magnitude

Figure 18. Histogram and distributions of target & non-target KL-scores and target & non-target probabilities - Terrain Class 14

Since this method assumes that all superpixels belong to every terrain class, an extra step to find the minimum is necessary to achieve the final fused map. The fused measure for this application is a data block with dimensions of number of superpixel by number of classes data block. By taking the minimum of each superpixel, it is ensured that the superpixel is assigned the best possible target class. The final result after taking the minimum is seen in Figure 19. The features that do not support change or non-targets generally have a signature that is dissimilar from the target features, the non-targets exhibit very high fused scores. In fact, the scores are often so high that displaying the results directly will skew the dynamic range such that nothing is discernible. To fix this, the upper max is clipped to 20. Figure 19 shows the clipped result.

6. CLASS REFINEMENT WITH SUPERPIXEL NEIGHBORHOODS

Inspection of Figure 19 shows that the result of the algorithm produces fairly good results. However, the result seems to exhibit a speckly behavior. While regions that support change will generally fall under a threshold, a few superpixels may appear to have non-target behavior. The adverse is also true. In non-target regions, there appear to be superpixels that do not follow the trend of the non-target. This can result in a mask that does not properly hide all regions of non-target. A solution to a similar problem was proposed by Fulkerson.[15] The simple yet very effective method proposed takes advantage of superpixel neighbors to stabilize the anomalous superpixels. Take the histogram distribution of superpixel s_i, where $s_i \in S$ and S represents all the superpixels in the image, to be H_i^0. H_i^N is now defined to be the histogram attained by merging H_i^0 will the histograms of all the superpixels that are less than or equal to N superpixels away. Another way of expressing this is:

$$H_i^N = \sum_{s_j | D(s_i, s_j) \leq N} H_j^0 \tag{7}$$

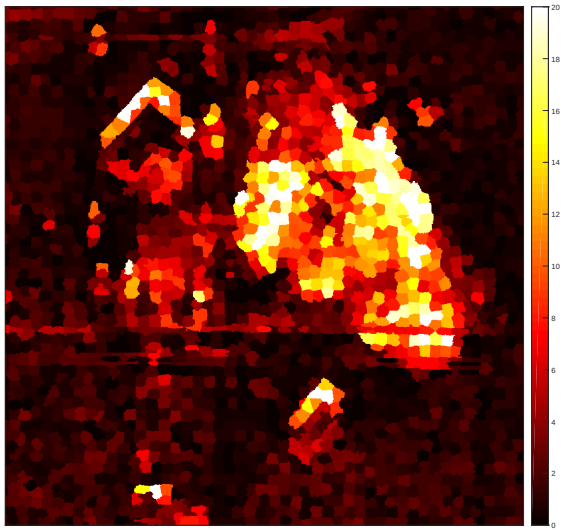

Figure 19. Final result of minimum fused measures

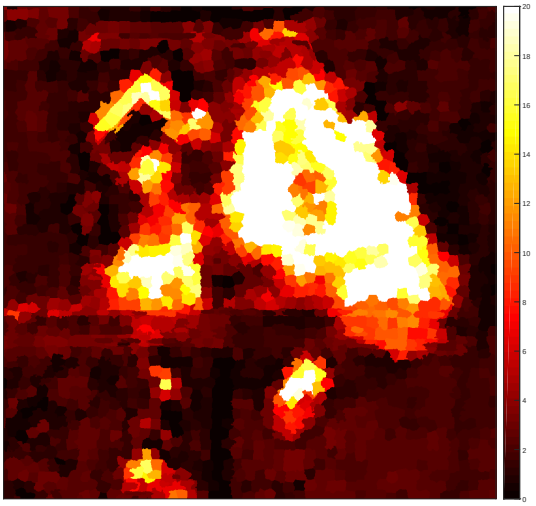

(a) Minimum fused measures w/superpixel neighborhood of N = 1

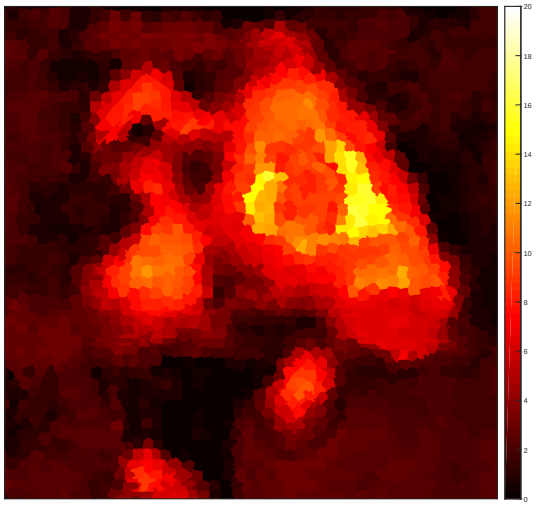

(b) Minimum fused measures w/superpixel neighborhood of N = 2

Figure 20. Comparison of incorporating superpixel neighbors

where $D(s_i, s_j)$ is the length of the shortest path between two superpixels. By incorporating a superpixel's neighbors, the distribution of a superpixel gains spatial consistency and as a result, becomes a better description of the superpixel.[15] The result of incorporating neighbors within one superpixel can be seen in Figure 20a. The result shows that the target and non-target regions become more homogeneous making it easier to discern regions that need to be ignored. This addition of neighboring superpixels does not come without cost, however. As N increases, this smoothing effect begins to blur the boundaries of the image's features. This is apparent by comparing the results seen in Figure 20b with Figures 19 and 20a. While the superpixel speckle has been eliminated from the homogeneous regions, the edges of the buildings are no longer apparent and shapes inside of the building are no longer distinguishable. To combat this effect, Fulkerson also incorporates conditional random fields (CRF). However, in this algorithm the CRF was not implemented. This is an area for potential future work.

7. CONCLUSIONS AND FUTURE DIRECTIONS

The proposed method offers a unique approach that factors in radar phenomenology while allowing a user to have freedom in defining the target classes. This approach also has the advantage of only requiring targets to be trained. By using a null hypothesis approach, probabilistic fusion assumes all superpixels are targets resulting in superpixels that do not follow the characteristics of a trained class to have high fused scores. While this approach does offer good results, it necessitates a system that provides a high SNR and good dynamic range. If non-target terrain is indistinguishable from target class terrain in the magnitude domain, it may have a hard time discriminating depending on what the terrain's signature is in the coherence domain. If additional features are introduced, this could provide additional distinguishing characteristics. In this effort, additional features were not explored, but they can easily be incorporated by using Equation 6. In addition, these features can be weighted. If a particular feature is known to be a better distinguisher, this can have a greater weighting. This approach uses median magnitude images and coherence images. However, it may be possible to use a regular magnitude or coherence product, but it is very important to find a training image that does not contain anomalous activity. While this could provide more accurate characterization of certain terrain, it may become a more difficult task due to the speckly behavior of SAR that is alleviated by using the median products. Another area of improvement to this algorithm is the incorporation of conditional random fields (CRF). As mentioned in the end of Section 6, incorporating superpixels neighbors to classify a superpixel's distribution can have a desired smoothing effect. However, this smoothing effect blurs the well defined edges of different terrain classes. To mitigate this blurring, CRF can be incorporated. Lastly, this method requires human supervised training. This training while minimal requires someone to find patches that are diverse enough in terrain classes and pick a subset of those superpixels. This is highly effective, but classes that should be deemed targets but are not originally modeled may produce bad scores. In addition, the training parameters are only valid at the geometry at which the training images were taken. As the radar geometry differs, the scattering phenomenology of the terrain changes. This requires the training parameters to be updated. While the training process cannot be avoided, it along with the iterative process of making sure all target classes are modeled can be eased by using an unsupervised training process. This was not explored in this effort, but offers an avenue for potential future work.

Although this method is not fully developed, it provides a product that is very good alternative to the current false alarm mitigation techniques that only take advantage of coherence. By incorporating the magnitude statistics in classifying non-target or false alarm areas, the classification process is able to properly distinguish terrain that looks similar in either SAR or coherence domain, but not both domains. Additional exploration of this technique could lead to even better false alarm suppression.

Acknowledgement

This work was supported by Pattern ANalytics To support High performance Exploitation and Reasoning (PANTHER), a Laboratory Directed Research and Development (LDRD) Project at Sandia National Laboratories. For additional information about PANTHER, please contact Kristina Czuchlewski, Ph.D., krczuch@sandia.gov.

REFERENCES

[1] Jakowatz, C. V., Wahl, D. E., Eichel, P. H., Ghiglia, D. C., and Thompson, P. A., [*Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*], Kluwer Academic Publishers, Boston, MA (1996).

[2] Yocky, D. A. and Johnson, B. F., "Repeat-pass dual-antenna synthetic aperture radar interferometric change-detection post-processing," *Photogrammetric engineering and remote sensing* **64**(5), 425–429 (1998).

[3] Zebker, H., Villasenor, J., et al., "Decorrelation in interferometric radar echoes," *Geoscience and Remote Sensing, IEEE Transactions on* **30**(5), 950–959 (1992).

[4] Preiss, M., Gray, D., Stacy, N. J., et al., "Detecting scene changes using synthetic aperture radar interferometry," *Geoscience and Remote Sensing, IEEE Transactions on* **44**(8), 2041–2054 (2006).

[5] Newey, M., Barber, J., Benitz, G., and Kogon, S., "False alarm mitigation techniques for sar ccd," in [*Radar Conference (RADAR), 2013 IEEE*], 1–6, IEEE (2013).

[6] Wahl, D. E., Yocky, D. A., and Jakowatz Jr, C. V., "A new maximum-likelihood change estimator for two-pass sar coherent change detection.," tech. rep., Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States) (2014).

[7] Barber, J. and Kogon, S., "Probabilistic three-pass sar coherent change detection," in [*Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on*], 1723–1726, IEEE (2012).

[8] Simonson, K. M., Perkins, D. N., and Lizarraga, I., "The use of multi-pass SAR change detection to highlight anomalous activity patterns," in [61^{st} *Annual Meeting of the MSS Tri-Service Radar Symposium*], (July 2015).

[9] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S., "Slic superpixels compared to state-of-the-art superpixel methods," *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **34**(11), 2274–2282 (2012).

[10] Moya, M. M., Koch, M. W., Perkins, D. N., and West, R. D. D., "Superpixel segmentation using multiple sar image products," in [*SPIE Defense+ Security*], 90770R–90770R, International Society for Optics and Photonics (2014).

[11] Perkins, D. and Gonzales, A., "Registering coherent change detection products associated with large image sets and long capture intervals," (Apr. 2014).

[12] Ren, X. and Malik, J., "Learning a classification model for segmentation," in [*Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*], 10–17, IEEE (2003).

[13] Simonson, K. M., "Probabilistic fustion of ATR results," SAND Report SAND98-1699, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550 (August 1998).

[14] Cox, D. R. and Hinley, D. V., [*Theoretical Statistics*], Chapman and Hall (1974).

[15] Fulkerson, B., Vedaldi, A., and Soatto, S., "Class segmentation and object localization with superpixel neighborhoods," in [*Computer Vision, 2009 IEEE 12th International Conference on*], 670–677, IEEE (2009).