Data Analytics for SAR

D. Murphy, dmurphy@lanl.gov
M. Calef. mcalef@lanl.gov

We assess the ability of variants of anomalous change detection (ACD) to identify human activity associated with large outdoor music festivals as they are seen from synthetic aperture radar (SAR) imagery collected by the Sentinel-1 satellite constellation. We found that, with appropriate feature vectors, ACD using random-forest machine learning was most effective at identifying changes associated with the human activity.

Background and Motivation

Anomalous Change Detection (ACD) provides a means to separate common, uninteresting changes from rare, significant changes in co-registered images collected at two points in time. The essential idea is to approximate the distribution of changes and look for particular changes that are rare with respect to this distribution [1]. This method has been applied with good effect to optical and hyper-spectral imagery where changes in ambient light level between collections lead to pervasive differences between the images.

Synthetic Aperture Radar (SAR) emits a signal that penetrates clouds, and forms images based on the reflected signal. As such, SAR does not suffer from the same pervasive changes as optical and hyper-spectral imagery do. However, there can still be widespread changes in a scene that are uninteresting, such as harvested crops in farmland. This leads to the question we begin to explore: can ACD be effective for SAR imagery? Our test case is detecting human activity, in particular outdoor music festivals, in large SAR images.

Description

The Burning Man Festival held in the Black Rock Desert of Nevada is an instructive test case. SAR imagery of the site before and during the festival are shown in Figure 1.

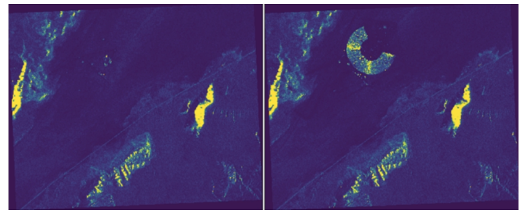

Figure 1 Before and during the Burning Man music festival.

If we consider just the per-pixel intensity of the reflected signal, we observe a distribution of changes and a compute an approximation of those changes as shown in Figure 2, where the x-axis is the intensity of a pixel in the first image, and the y-axis is the intensity of the same pixel in the second image.

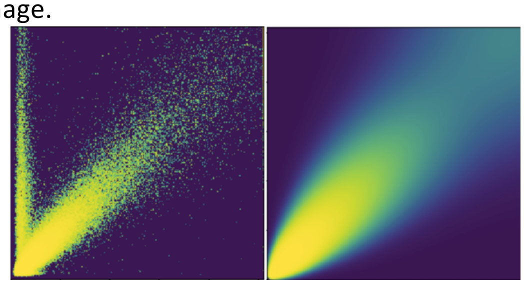

Figure 2: The observed changes in intensity on the left, and the approximated distribution of changes on the right.

The set of pixels transitioning from very dark to bright near the y-axis in the left image of Figure 2 are uncommon with respect to the approximated distribution of changes, and thus flagged as anomalous.

This simple approach was not sufficient for music festivals that occurred in settings with more complex ambient changes. Per-pixel intensity did not provide sufficient information. In response, we started building feature vectors based on information in a neighborhood of a particular pixel. In addition, we used a random forest to classify anomalous vs. not anomalous. For not-anomalous training data we used correlated data from the two images. For anomalous training data, we used randomly paired neighborhoods from the two images. The response operator characteristic (ROC) curves for Coachella are shown in Figure 3.

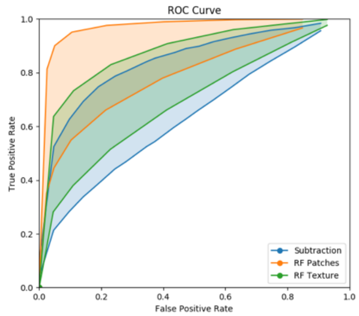

Figure 3: ROC curves for various change detection methods. The width of the ROC curve reflects uncertainty in the precise bounds of the music festival.

Anticipated Impact

The recent increase in open and commercial satellite data has made change detection in images a "big data" problem. Our work is helping to develop and characterize useful change detection methods that will be allow mining of the volumes of data, and ultimately deliver knowledge, not pixels, to the end user.

Acknowledgements

The authors gratefully acknowledge the guidance and support of Karen Miller.

Los Alamos Report LA-UR-17-28988. Funded by the Department of Energy at Los Alamos National Laboratory under contract DE-AC52-06NA25396.

References

[1] James Theiler and Simon Perkins, Proposed Framework for anomalous change detection, ICML Workshop on Machine Learning Algorithms for Surveillance and Event Detection 2006, pp 7-14