Water Across Synthetic Aperture Radar Data (WASARD): SAR Water Body Classification for the Open Data Cube

[1]Zachary Kreiser, [2]Brian Killough, [1]Syed R Rizvi

[1]Analytical Mechanics Associates, Hampton, VA
[2]NASA Langley Research Center, Hampton, VA

ABSTRACT

The detection of inland water bodies from Synthetic Aperture Radar (SAR) data provides a great advantage over water detection with optical data, since SAR imaging is not impeded by cloud cover. Traditional methods of detecting water from SAR data involves using thresholding methods that can be labor intensive and imprecise. This paper describes Water Across Synthetic Aperture Radar Data (WASARD): a method of water detection from SAR data which automates and simplifies the thresholding process using machine learning on training data created from Geoscience Australia's WOFS algorithm. Of the machine learning models tested, the Linear Support Vector Machine was determined to be optimal, with the option of training using solely the VH polarization or a combination of the VH and VV polarizations. WASARD was able to identify water in the target area with a correlation of 97% with WOFS.

Index Terms— Sentinel-1, Open Data Cube, Earth Observations, Machine Learning, Water Detection

1. INTRODUCTION

Water classification is an important function of Earth imaging satellites, as accurate remote classification of land and water can assist in land use analysis, flood prediction, climate change research, as well as a variety of agricultural applications [2]. The ability to identify bodies of water remotely via satellite is immensely cheaper than contracting surveys of the areas in question, meaning that an application that can accurately use satellite data towards this function can make valuable information available to nations which would not be able to afford it otherwise.

Highly reliable applications for the remote detection of water currently exist for use with optical satellite data such as that provided by LANDSAT. One such application, Geoscience Australia's Water Observations from Space (WOFS) has already been ported for use with the Open Data Cube [6]. However, water detection using optical data from Landsat is constrained by its relatively long revisit cycle of 16 days [5], and water detection using any optical data is constrained in that it lacks the ability to make accurate classifications through cloud cover [2]. The alternative solution which solves these problems is water detection using SAR data, which images the Earth using cloud-penetrating microwaves.

Because of its advantages over optical data, much research has been done into water detection using SAR data. Traditionally, this has been done using the thresholding method, which involves picking a polarization band and labeling all pixels for which this band's value is below a certain threshold as containing water. The thresholding method works since water tends to return a much lower backscatter value to the satellite than land [1]. However, this method can be flawed since estimating the proper threshold is often imprecise, complicated, and labor intensive for the end user. Thresholding also tends to use data from only one SAR polarization, when a combination of polarizations can provide insight into whether water is present. [2]

In order to alleviate these problems, this paper presents an application for the Open Data Cube to detect water from SAR data using support vector machine (SVM) classification.

2. PLATFORM

WASARD is an application for the Open Data Cube, a mechanism which provides a simple yet efficient means of ingesting, storing, and retrieving remote sensing data. Data can be ingested and made analysis ready according to whatever specifications the researcher chooses, and easily resampled to artificially alter a scene's resolution. Currently WASARD supports water detection on scenes from ESA's Sentinel-1 and JAXA's ALOS. When testing WASARD, Sentinel-1 was most commonly used due to its relatively high spatial resolution and its rapid 6 day revisit cycle [5]. With minor alterations to the application's code, however, it could support data from other satellites.

3. METHODOLOGY

Using supervised classification, WASARD compares SAR data to a dataset pre-classified by WOFS in order to train an SVM classifier. This classifier is then used to detect water in other SAR scenes outside the training set. Accuracy was measured according to the following metrics:

- *Precision*: a measure of what percentage of the points WASARD labels as water are truly water
- *Recall*: a measure of what percentage of the total water cover WASARD was able to identify.
- *F1 Score*: a harmonic average of the precision and recall scores

Both precision and recall are calculated at the end of the training phase, when the trained classifier is compared to a testing dataset. Because the WOFS algorithm's classifications are used as the truth values when training a WASARD classifier, when precision and recall are mentioned in this paper, they are always with respect to the values produced by WOFS on a similar scene of Landsat data, which themselves have a classification accuracy of 97% [6].

Visual representations of water identified by WASARD in this paper were produced using the function wasard_plot(), which is included in WASARD.

3.1 Algorithm Selection

The machine learning model used by WASARD is the Linear Support Vector Machine (SVM). This model uses a supervised learning algorithm to develop a classifier, meaning it creates a vector which can be multiplied by the vector formed by the relevant data bands to determine whether a pixel in a SAR scene contains water. This classifier is trained by comparing data points from selected bands in a SAR scene to their respective labels, which in this case are "water" or "not water" as given by the WOFS algorithm. The SVM was selected over the Random Forest model, which outperformed the SVM in training speed, but had a greater classification time and lower accuracy, and the Multilayer Perceptron Artificial Neural Network, which had a slightly higher average accuracy than the SVM, but much greater training and classification times.

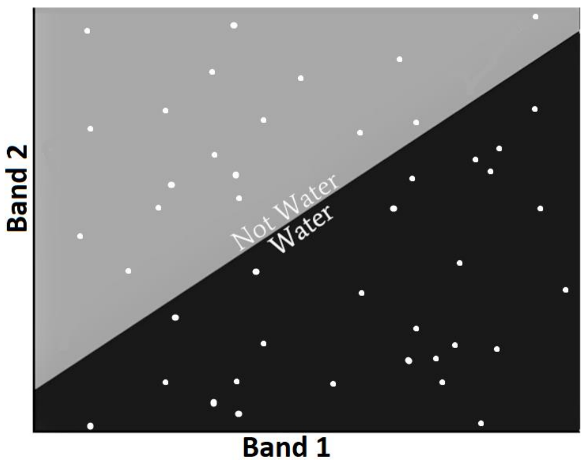

Figure 1: Visual representation of the SVM Classifier. Each white point represents a pixel in a SAR scene.

In Figure 1, the diagonal line separating pixels determined to be water from those determined not to be water represents the actual classification vector produced by the SVM. It is worth noting that once the model has been trained, classification of pixels is done in a similar manner as in the thresholding method. This is especially true if only one band was used to train the model.

3.1 Feature Selection

Sentinel-1 collects data from two bands: the Vertical/Vertical polarization (VV) and the Vertical/Horizontal polarization (VH). When 100 SVM classifiers were created for each polarization individually, and for the combination of the two, the following results were achieved:

Bands Used	Average Precision	Average Recall
VV	0.921	0.977
VH	0.948	0.982
VV,VH	0.960	0.980

Figure 2: Accuracy of classifiers trained using different polarization bands. Precision and Recall were measured with respect to the values produced by WOFS.

Figure 2 demonstrates that using both the VV and VH bands trades slightly lower recall for significantly greater precision when compared with the VH band alone, and that using the VV band alone is inferior in both metrics. WASARD therefore defaults to using both the VV and VH bands, and includes the option to use solely the VH band. The VV polarization's lower precision compared to the VH polarization is in contrast to results from previous research and may merit further analysis [4].

3.2 Training a Classifier

The steps in training a classifier with WASARD are

1. Selecting two scenes (one SAR, one optical) with the same spatial extents, and acquired close to each other in time, with a preference that the scenes are taken on the same day.
2. Using the WOFS algorithm to produce an array of the detected water in the scene of optical data, to be used as the labels during supervised learning
3. Data points from the selected bands from the SAR acquisition are bundled together into an array with the corresponding labels gathered from WOFS. A random sample with an equal number of points labeled "Water" and "Not Water" is selected to be partitioned into a training and a testing dataset
4. Using Scikit-Learn's LinearSVC object, the training dataset is used to produce a classifier, which is then tested against the testing dataset to determine its precision and recall

The result is a wasard_classifier object, which has the following attributes:

1. f1, recall, and precision: 3 metrics used to determine the classifier's accuracy
2. Coefficient: Vector which the SVM uses to make its predictions. The classifier detects water when the dot product of the coefficient and the vector formed by the SAR bands is positive
3. Save(): allows a user to save a classifier to the disk in order to use it without retraining
4. wasard_classify(): Classifies an entire xarray of SAR data using the SVM classifier

All of the above steps are performed automatically when the user creates a wasard_classifier object.

3.3 Classifying a Dataset

Once the classifier has been created, it can be used to detect water in an xarray of SAR data using wasard_classify(). By taking the dot product of the classifier's coefficients and the vector formed by the selected bands of SAR data, an array of predictions is constructed. A classifier can effectively be used on the same spatial extents as the ones where it was trained, or on any area with a similar landscape. While testing WASARD, it was observed that a classifier trained on one lake in Vietnam detected water accurately across the entire nation.

3.4 Noise Reduction

One major drawback of SAR data is that the intensity of the returned signals from a SAR satellite's radar pulses will vary in frequency from pixel to pixel due to the waves falling out of phase after hitting the Earth's surface. This results in an image which has pixels that vary in intensity across a homogenous area, referred to as *speckle* [1]. This speckle noise can potentially cause the classifier to mislabel points, and so reducing speckle is necessary to ensuring WASARD makes the most accurate classifications possible. Since speckle shows up as points normally too small to be separate bodies of water, WASARD reduces noise by scanning over the classified dataset with a moving window and removing isolated pixels labeled as water.

5. RESULTS AND DISCUSSION

WASARD was chiefly tested on bodies of water in Southern Vietnam, which is appropriate in large part due to the region's tropical climate and jungle covered landscape, which contribute to the area being frequently covered by heavy clouds. Since this weather makes water detection with optical data less effective, it is an ideal region on which to use WASARD.

5.1 Accuracy of Classifications

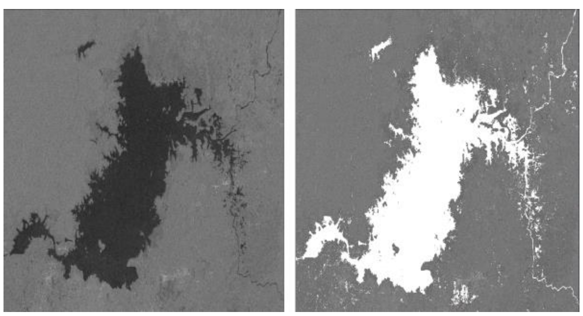

Figure 3: Water Detection by WASARD (Raw SAR image on left, WASARD's predictions denoted in white on the right hand side).

The comparison in Figure 3 demonstrates WASARD's water detection on a reservoir in Southern Vietnam. This classifier has a precision score of .963, and recall of .983. The Support Vector Classifier is represented by the following equation:

SVC = (VH Coefficient) (VV Band Value) + (VH Coefficient) (VH Band Value) + Bias Constant, where SVC < 0 is water and SVC > 0 is non-water.
*For Figure 3: SVC = -45.899 * VH - 1.271 * VV + 1.007*

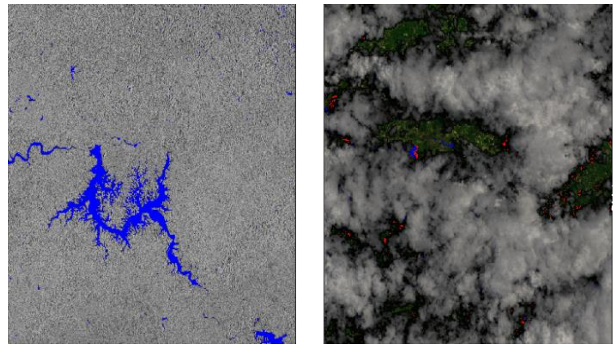

Figure 4: WASARD classification (left) vs WOFS classification (right). Pixels identified as water by each algorithm are denoted in blue.

Figure 4 demonstrates WASARD's utility in identifying water no matter the conditions. The Sentinel scene on the left and the Landsat scene on the right have the same latitudinal and longitudinal extents and were acquired on the same day, but the region was covered by clouds at the time of the Landsat acquisition. The WOFS data here is unusable since the clouds prevent accurate classification, but WASARD is still able to perform. Only one clear Landsat scene is needed to train a WASARD classifier, making WASARD invaluable in areas frequently covered by clouds.

On the areas studied, WASARD was able to correctly classify points with an accuracy which is comparable or superior to existing methods for SAR data [1]. WASARD's classifications had a precision of .960 and recall of .980 on average, with an overall correlation with WOFS of 97%. These numbers are in comparison to the labels produced by WOFS on a Landsat scene with the same spatial extents and as close temporally as possible. It is possible, therefore, that these results are slightly lower than WASARD's true accuracy due to changes in water cover between when the Landsat scenes and SAR scenes were acquired. Overall, it would seem that WASARD offers a viable alternative to classification with WOFS for when weather conditions preclude WOFS's use.

5.2 Applications

Figure 5 shows a composite of WASARD's classifications in Buon Tua Sarh in Dak Nong, Vietnam. This composite was built from 17 scenes spanning 22 months. It is clear that there is a large amount of variation in the presence of water in this region, as there are large portions colored blue where water was found 80-100% of the time, as well as significant portions colored yellow or orange where water was found 20-60% of the time. Closer research reveals that this variation is due to the body of water being a reservoir whose water level is controlled by a dam at its northernmost point. While the example of identifying a body of water with a dam is somewhat trivial, it is a valid demonstration of how time series data might be used to analyze flooding patterns in a region. WASARD makes running time series analysis on SAR scenes easy with the included function wasard_time_plot().

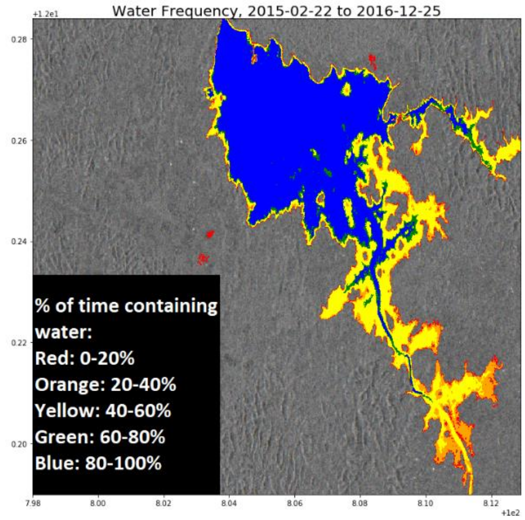

Figure 5: Time series representation of water identified by WASARD from February 2015 to December 2016. 17 scenes were analyzed to create this image.

5.3 User Friendliness

In addition to its accuracy and its functionality despite cloud cover, WASARD's value also stems from the ease with which it can be used. Every step of generating an effective classifier is automated. WASARD automatically handles all of the following:

- Selecting a clear Landsat scene and a corresponding SAR scene on which to train the SVM
- Adjusting the resolution of the SAR data to fit that of the Landsat data
- Creating a training dataset, and fitting the classifier to the data.

All the user need do is feed in two datasets with identical spatial extents, one of SAR data and one from Landsat. Additionally, since WASARD operates in conjunction with the Open Data Cube, it allows the user to analyze patterns of water cover without the use of a GIS. This user-friendliness contrasts most current methods of detecting water from SAR, which involve the user creating histograms to determine thresholds [2] and using a GIS to run analysis on water cover patterns [3].

5.4 Tradeoffs

Despite its benefits, WASARD is constrained in a few areas. WASARD struggles to return precise measurements in flat areas without vegetation such as deserts due to the way smooth, flat terrain reflects microwaves in a similar fashion to water. This same constraint applies to any smooth, flat surface such as roofs, as well as terrain such as mountains that may cause the wave to be reflected away from the satellite's sensor. Incorporating a digital elevation model may be able to alleviate this problem [1].

Additionally, due to the random selection of points to be used for training data, WASARD will create a slightly different classifier each time it is trained on a given area. This was tested by calculating the precision and recall scores from 100 classifiers trained back to back on the same area. Standard deviations of .008 for precision and .004 for recall were found. Therefore, trial and observation of multiple classifiers is recommended to find the optimal one, which can then be saved and reused. Included in WASARD is a function get_best_classifier() which automates this process, training a given number of classifiers and returning them to the user in a list sortable by precision, accuracy, or f1 score.

6. CONCLUSION

WASARD is an effective application for use with the Open Data Cube that addresses the problem of threshold estimation, one of the greatest obstacles to detecting water using SAR data. WASARD solves this problem by using the WOFS algorithm to produce a dataset which can then be used to train an effective classifier for SAR data, with no intermediate steps required by the user. Identifying and removing false positives caused by speckle further improves WASARD's predictions to the point that it can match or defeat existing methods in both its simplicity and accuracy. Demonstrations shown in this paper were all performed on data from ESA's Sentinel. Preliminary tests done on JAXA's ALOS were promising but constrained by lack of available data. Despite the fact that it is limited by its decreased performance in desert landscapes, WASARD offers an effective and easy to use SAR water detection application for the Open Data Cube in tropical environments such as South Vietnam.

7. REFERENCES

1. S. Hong *et al.*, "Water Area Extraction Using RADARSAT SAT Imagery Combined with Landsat Imagery and Terrain Information," *Sensors*, vol. 15(3), pp. 6652-6657, 2015.
2. A. Behnamian *et al.*, "Semi-Automated Surface Water Detection with Synthetic Aperture Radar Data: A Wetland Case Study," *Remote Sensing*. 23 November 2017.
3. R. Rahman, P. Thakur, "Detecting, mapping, and analyzing of flood water propagation using synthetic aperture radar (SAR) satellite data and GIS: A case study from the Kendrapara District of Orissa State of India," *The Egyptian Journal of Remote Sensing and Space Science*. 24 October 2017.
4. M.A. Clement *et al.*, "Multi-temporal synthetic aperture radar flood mapping using change detection," *Journal of Flood Risk Management*. 27 APR 2017.
5. L. Zeng *et al.*, "Analyzing changes of the Poyang Lake water area using Sentinel-1 synthetic aperture radar imagery," *International Journal of Remote Sensing*, vol. 38(23), 13 August 2017.
6. N. Mueller *et al.*, "Water observations from space: Mapping surface water from 25 years of Landsat imagery across Australia," *Remote Sensing of Environment*, vol. 174, pp. 341-352, 1 March 2016.