Appears in Proceedings of the Workshop on AI and Space, International Joint Conference on Artificial Intelligence, Barcelona, Spain, July 2011.

USING SUPPORT VECTOR MACHINE LEARNING TO AUTOMATICALLY INTERPRET MODIS, ALI, AND L-BAND SAR REMOTELY SENSED IMAGERY FOR HYDROLOGY, LAND COVER, AND CRYOSPHERE APPLICATIONS

J. Doubleday, D. McLaren, S. Chien, Y. Lou

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91101
{firstname.lastname}@jpl.nasa.gov

Abstract –

We report on efforts to automatically interpret remotely sensed radar and multispectral imagery using machine-learned classifiers. Specifically we utilize Support Vector Machine (SVM) learning techniques for L-band SAR, EO-1/ALI, and MODIS Imagery. We share our qualitative and quantitative results thus far and discuss challenges experienced.

Keywords: SVM, classification, remote sensing, polarimetric SAR, MODIS, ALI, disaster management

1 INTRODUCTION

Remotely sensed imagery has been found useful in a wide range of applications. In particular, remotely sensed imagery is frequently used in monitoring, responding to, and mitigating natural disasters [Tralli et al. 2005, Chien et al. 2011a].

However, timely acquisition, processing, delivery and analysis of the remote sensed data is critical in these applications. This paper discusses efforts to automate the interpretation of raw imagery into directly usable products. Automation of this analysis step will allow for use of more data sources, enable more timely analysis, and enable automated monitoring to detect disasters as they evolve. The end goal of such automated classification methods is to enable 24/7 automated imager interpretation, feeding directly into phenomenological models such as hydrological/flooding models or forest fire progression models to enable more precise tracking, prediction, and mitigation for natural disasters using an adaptive sensing sensorweb paradigm [Chien et al. 2011b].

Specifically we utilize support vector machine learning (SVM) techniques in an attempt to automate analysis of the remotely sensed imagery. In this paper we describe efforts to apply SVM for surface water extent mapping, vegetation mapping, and snow/ice extent mapping. We describe results, related work, and future areas for work.

2 SURFACE WATER EXTENT CLASSIFICATION OF MULTI-SPECTRAL MODIS DATA

We have used SVM techniques and band ratio techniques to automatically classify MODIS flood extent maps of Thailand. We have developed algorithms, using data from MODIS bands 7, 2, and 1 (red, green, blue) available from the MODIS Rapid Response System [GSFC]. Both methods classify pixels into three categories: cloud, water, and land. A permanent water mask was created by running a trained SVM (see below) over 4 relatively clear images from Thailand's dry season, and keeping pixels that were classified as water in at least 2 of the scenes. This mask was used to screen out permanent water pixels.

The first classification algorithm used was a simple thresholding technique to classify flooding, primarily used to request observations of Thailand flood targets with the EO-1 ALI instrument. This algorithm created a composite image by iterating through images from 8 MODIS overflights, filling in pixels that may be obscured by clouds in some images. Clouds were identified as sufficiently bright pixels with the ratio *g/b* between low and high values. Three color thresholds were chosen to mark the highest expected color values for a water pixel, W_r, W_g, and W_b. The RGB score components were then computed as in eq. 1,

$$s_i = 255 - (\frac{255}{W_i} * v_i) \qquad (1)$$

where i is r, g, or b, and v_i is the value of a pixel in one band. The final score for each pixel was computed as S in eq 2.

$$S = \max(0, (s_r + s_g + s_b) * k^d) \quad (2)$$

\where d is the age of the pixel data in days and k is a decay factor. Finally, all pixels with score exceeding a user-specified threshold were classified as flooded. The exact values of these parameters were tuned based on manual experimentation.

The second classification algorithm was a support vector machine (SVM) with a 2nd-degree polynomial kernel,

trained with the Pixellearn [MLIA] program on manually labeled MODIS data. A confusion matrix against manually labeled data is shown in Tab. 1. The SVM identified 96.8% of all labeled pixels correctly.

Table 1: Confusion Matrix for MODIS SVM trained on scene from 2010-078

	Background	Land	Water	Cloud
Background	0	4372895	255495	510805
Land	0	3260	97	10
Water	0	126	2903	0
Cloud	0	64	0	2706

Both algorithms were run for the composite image shown in Fig. 1, which is the image generated by the thresholding algorithm from 8 overflights prior to the day 2010-308. Fig. 2 shows the results for both the SVM and thresholding algorithms.

The thresholding algorithm generated a composite image using data captured by both Terra and Aqua, from 2010-305 to 2010-308, and attempts to screen out cloudy pixels (Figure 2 bottom). The SVM was run on this composite image to provide a more equal comparison (Fig. 2 top). It can be seen that the thresholding algorithm left in pixels that the SVM correctly classified as cloudy, and classified them as land instead. It also generated more false positive flood pixels (especially in western Thailand) and identified more pixels as flooded overall. The SVM classified 0.7% of the pixels as flooded, while the thresholding counted 1.0% as flooded.

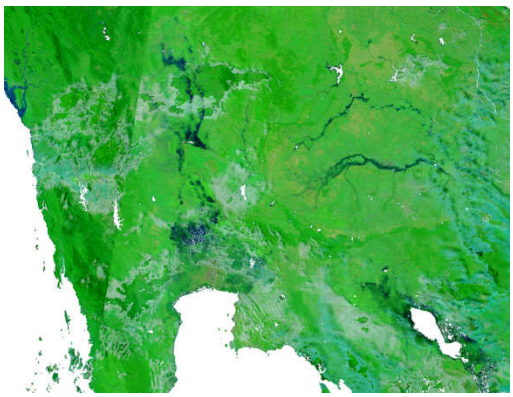

Figure 1: Composite of 8 MODIS 721 images

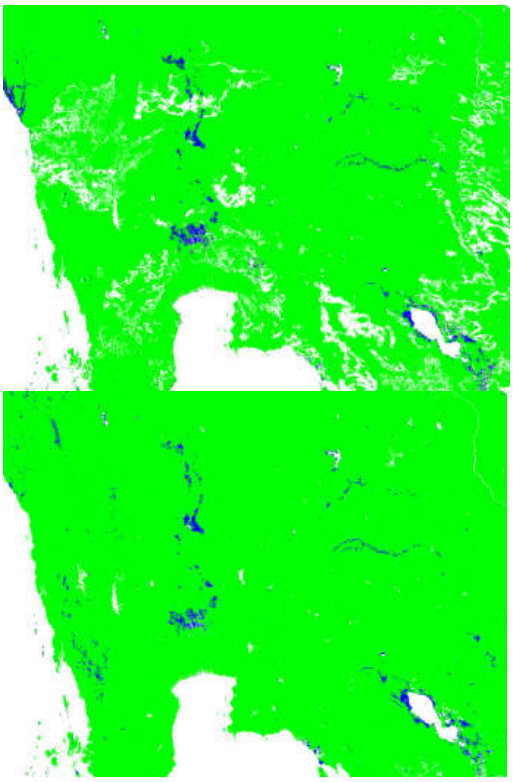

Figure 2: SVM classification results (top) and thresholding results (bottom); blue: flooded, green: land, white: cloud & masked out areas

3 SURFACE WATER EXTENT CLASSIFICATION OF MULTI-SPECTRAL ALI DATA

We have used SVM techniques to learn classifiers to automatically detect flooded areas in Earth Observer One (EO-1) Advanced Land Imager (ALI) data. Several scenes of ALI data from regions of Thailand and Siem Reap/Tonle Sap area were collected and hand labeled for water (large lakes or catchments), developed areas, undeveloped ground, cloud, and finally cloud shadowed regions. In the interest of producing products that may be useful in flood mitigation, labels for ground and water were chosen aggressively through partially clouded observations. Water labeling was assisted by first generating a binary image of a threshold of ratio of ALI bands nearest 550nm and 860nm, and using this map to label easily identifiable water structures (river, ponds, lakes).

Only a small portion of ALI scenes were labeled as such – the remainder of the scenes contained large swaths of easily identifiable features (clouds, ground, water bodies) that remained unlabelled and were treated as a hold-out set in qualitative evaluation.

For all flood classification efforts for ALI, a feature set of all possible band ratios (excluding reciprocals) of the 9 multispectral bands was utilized (resulting in 36 ratios).

Labeling, training, validation (quantitative and qualitative) and kernel-parameter selection, was done through the Pixellearn tool [MLIA]. Cross-validation accuracy was used as the basis for a classifier score in a typical grid search for parameters and kernels. Kernels included linear, polynomial (inhomogeneous, across several small, prime degrees) and RBF with varying gamma. Search was run on each labeled scene separately – the optimal selection of parameters and kernel varied. In some cases where labels were limited, the linear kernel proved optimal or as optimal as more complex kernels, while in others RBF would achieve the best cross-validation accuracy. However, often for the linear kernel the generalization to the remainder of unlabelled scene was poor, while selecting a particular RBF parameterization, training and validating across several scenes, would show poor generalization across scenes. Through manual and qualitative selection, a polynomial kernel of degree 5 and a cost factor (C) of 300 was chosen for best generalization to unlabelled data across several scenes.

Testing against a hand labeled hold-out set of obvious features (seen in Fig. 3) yields the fairly good accuracies as seen in Tab. 2. Unfortunately, we have not been able to perform direct comparison of these results against ground truth measurements.

Table 2. Confusion matrix of single hold-out scene (Siem Reap, 2010-293)

	Background	cloud	water	land	shadow	urban	total	cloud	water	land	shadow	urban
Background	0	756566	1076468	1066693	317514	0	3217241	24%	33%	33%	10%	0%
cloud	0	546	0	0	0	0	546	100%	0%	0%	0%	0%
water	0	25	3754	48	111	0	3938	1%	95%	1%	3%	0%
land	0	0	0	990	65	0	1055	0%	0%	94%	6%	0%
shadow	0	0	0	12	385	0	397	0%	0%	3%	97%	0%
urban	0	0	0	97	1	0	98	0%	0%	99%	1%	0%

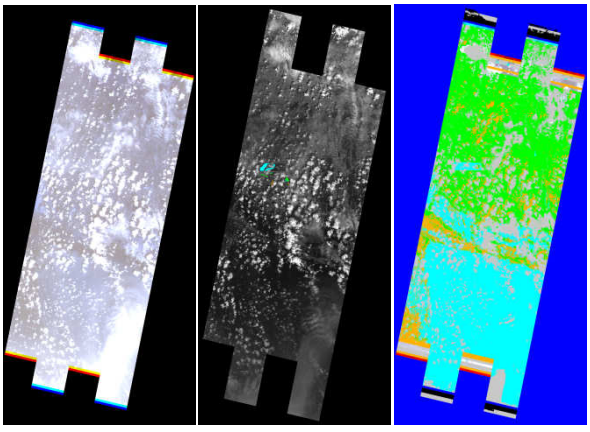

Figure 3: Three views of ALI scene over Siem Reap (EO1A1270512010293110PF). Left: an RGB composite. Center: Band-9 with hold-out cross-validation labels. Right: SVM classification results. Color map: blue-water, green-land, grey-cloud, orange-shadow, (darkblue-invalid).

4 AUTOMATED CLASSIFICATION OF L-BAND SAR DATA INTO VEGETATION CLASSES

The feasibility of generating a vegetation classification scheme using our real-time L-band SAR data has been studied using a multi-class support vector machine (SVM). Using UAVSAR data collected over forested and wetland areas of Maine and Vermont, spread out over approximately 12 days of varying weather, we trained a support vector machine using land cover maps from NLCD2001 with classes collapsed down to 7 coarser vegetation classes: dense (forest), herbaceous wetland, forested wetland, medium (shrub), low (grass), bare land and water. Approximately 5000 training samples (pixels) were chosen at random across the Maine dataset and included features of backscatter for HH, VV, HV, HHVV phase, 7x7 pixel averages

Table 3. Confusion matrix for vegetation classification vs NLCD2001 condensed classes, evaluated over ~15 30kmx30km images from Maine and Vermont

		Classifier results:						
NLCD		1000	2000	2100	2200	3000	4000	count
water	1000	54%	4%	2%	19%	11%	10%	682,069
dense veg (forest)	2000	0%	46%	21%	12%	11%	10%	16,382,691
wetland (herbaceous)	2100	1%	31%	30%	17%	12%	10%	1,936,352
wetland (woody)	2200	6%	18%	17%	24%	17%	19%	478,129
moderate veg (shrub)	3000	3%	26%	15%	13%	20%	23%	450,214
light veg (grass)	4000	8%	19%	10%	18%	19%	26%	51,743
urban	5000	5%	16%	8%	27%	23%	21%	323,258
urban (low intensity)	5100	6%	19%	9%	21%	22%	22%	519,914
developed (open space)	5200	3%	28%	13%	16%	20%	20%	775,630
Total Accuracy: 8668968 / 19981198 (0.433856)								

and variance of some of these quantities for 19 total features. Training yielded ~900 support vectors using a Guassian radial basis kernel.

Tab. 3 shows fairly poor classification performance when compared against NLCD2001 data, with the best accuracy at 54% for the water class of which 19% were misclassified as woody wetland, while dense vegetation had the next best accuracy of 46%, with 21% misclassified as herbaceous wetland. These figures are fairly underwhelming, however qualitative inspection of results look more promising and draw doubt upon our validation set. Fig. 4a and 4b show the contrast in NLCD2001 ground truth to classifier output, while optical imagery in Fig. 4c shows a better visual correlation with the classifier output – of note is the abundance of light vegetation (open spaces or grassy areas, brown) consistent

with the output, opposed to the majority of dense forest indicated in the ground truth in Fig. 4a. We suspect the NLCD2001 data is ultimately too noisy for training and validation purposes either because of age, resolution/aliasing artifacts (NLCD2001 is 30m resolution while UAVSAR is better than 10m allowing for individual tree crown identification) or other general fidelity problems. Future work will identify better land cover data and investigate improvements in cross-validation, while at the moment we find it interesting that with this potentially noisy dataset we obtain results that compare well, at least qualitatively with more recent optical data.

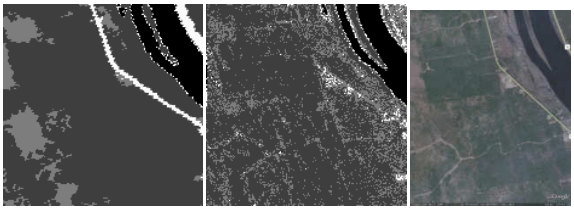

Figure 4. Vegetation classification data for scene in Maine, collected Sept 2009. (a) NLCD2001 land cover map, (b) classification results, (c) optical imagery courtesy Google Earth. Color scheme for (a) and (b), black: water, dark-gray: dense veg (including wetlands), gray: moderate veg, light-gray: light-veg, white: developed areas

Finally we would note that the resulting SVM, with ~900 support vectors would be too complex to be run on an in-flight platform. We anticipate replacing noisy training data with more accurate data would result in a reduced set and better cross-validation results. Other approaches to reach a feasible computation time would include a hardware based SVM evaluator or reduced set SVM [Tang and Mazzoni 2006] or progressive SVMs [Wagstaff et al. 2010].

5 SNOW/ICE CLASSIFICATION OF L-BAND SAR DATA

We again employ a support vector machine to classify pixels of ground projected data using the same feature set as presented in section 5 for vegetation classification, but with target classes of snow or ice versus land versus water. Data collected over Iceland's Hofsjokull on Jun 12th 2009 together with various optical remote sensing data including Landsat7 taken May 16 2009 (Fig. 5a) to serve as our training and validation data, and ground truth. Given the time of year, and large temporal separation between the optical and UAVSAR data collection, we generated a conservative hand labeled training and validation set (Fig. 5b), taking into consideration late-summer data to find a year-round snow pack. Training on ~300 randomly sampled pixels from the hand-labeled set resulted in the classification image of Fig. 6, and the confusion matrix in Tab. 4.

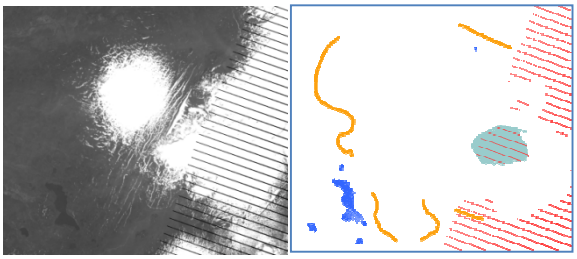

Figure 5. On left (a) Landsat7 band 1 image collected 2009-05-16, snow and ice from Hofsjokul on right, and seasonal snow on volcanic peak at center top. On right (b) is a conservative hand-labeling, water: blue, land: orange, snow/ice: gray, red: no-data.

Figure 6. Result of Snow/Ice SVM classifier: pink: land, dark-gray: water, gray-blue: snow/ice

Table 4. Confusion matrix of snow/ice classifier against hand labels, data collected Jun 12th 2009.

	water	ice	land	count
no label	7%	10%	83%	16.7M
water	91%	3%	6%	145884
ice	3%	90%	7%	440706
land	1%	2%	96%	311201
Total Accuracy: 829007 / 897791 (0.923385)				

Overall classification accuracies (>90%) are much improved as compared to the vegetation classifier and the number of support vectors dramatically reduced (~50) potentially allowing for in-flight evaluation on traditional processor architectures.

6 RELATED AND FUTURE WORK, CONCLUSIONS

SVM has been used in a number of previous projects to classify remotely sensed imagery. SVM classifers were learned on the ground and uploaded to classify cryosphere events using Hyperion data for the Earth Observer One mission [Castano et al. 2006]. Mazzoni and others [Mazzoni et al. 2007a, 2007b] used SVM to automatically classify atmospheric features in MODIS/MISR data.

Detection of sulfur springs against an ice backdrop [Mandrake et al. 2009] has been demonstrated with some success with Hyperion data. Similar techniques were applied to Mars Odyssey THEMIS data [Castano et al. 2007].

In further work, we would like to apply SVM to learn smoke products from ALI and MODIS imagery for forest fires. We would also like to apply SVM techniques to learn ash plume detection methods. It is expected that both of these problems will require the inclusion of texture features to achieve reasonable accuracy.

We have presented a number of applications of SVM learning to learn automatic classifiers for remotely sensed data. Data sets used include MODIS, ALI, and UAVSAR L-band SAR instrument data. Applications include surface water extent mapping, vegetation cover classification, and cryosphere feature classification.

ACKNOWLEDGEMENT

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration.

REFERENCES

R. Castano, D. Mazzoni, N. Tang, T. Doggett, S. Chien, R. Greeley, B. Cichy, A. Davies, "Onboard Classifiers for Science Event Detection on a Remote Sensing Spacecraft," Proceedings of the International Conference on Knowledge Discovery and Data Mining (KDD), p. 922-930, August 2006.

R. Castano, K. L. Wagstaff, S. Chien, T. M. Stough, and B. Tang. On-board Analysis of Uncalibrated Data for a Spacecraft at Mars. Proceedings of the Thirteenth International Conference on Knowledge Discovery and Data Mining (KDD), p. 922-930, August 2007.

S. Chien and V. Tanpipat, "Remote Sensing of Natural Disasters," in Encyclopedia of Sustainability Science and Technology, Springer, 2011.

S. Chien, J. Doubleday, D. Mclaren, D. Tran, C. Khunboa, W. Leelapatra, V. Plergamon, V. Tanpipat, R. Chitradon, S. Boonya-aroonnet, P. Thanapakpawin, A. Meethome, C. Raghavendra, D. Mandl, "Combining Space-based and In-situ measurements to track flooding in Thailand," International Geoscience and Remote Sensing Symposium, Vancouver, BC, July 2011.

U.S. EPA, Multi-Resolution Land Characteristics Consortium 2011 (MLRC) "2001 National Land Cover Data (NLCD 2001)" http://www.epa.gov/mrlc/nlcd-2001.html

NASA GSFC, "MODIS Rapid Response System Subset – FAS_Indochina", http://rapidfire.sci.gsfc.nasa.gov/subsets/index.php?subset=FAS_Indochina

Machine Learning and Instrument Autonomy (MLIA) Group, Jet Propulsion Laboratory, PixelLearn Users Quide, 2008.

L. Mandrake, K. L. Wagstaff, D. Gleeson, U. Rebbapragada, D. Tran, R. Castano, S. Chien, R. T. Pappalardo, "Onboard SVM Analysis of Hyperion Data to Detect Sulfur Deposits in Arctic Regions," Proceedings of the IEEE WHISPERS Workshop, August 2009.

D. Mazzoni, J. A. Logan, D. Diner, R. Kahn, L. Tong, and Q. Li. A data-mining approach to associating MISR smoke plume heights with MODIS fire measurements. Remote Sensing of the Environment 107(1-2), 2007, p. 138-148.

D. Mazzoni, M. J. Garay, R. Davies, and D. Nelson. An operational MISR pixel classifier using support vector machines. Remote Sensing of the Environment 107(1-2), 2007, p. 149-158.

B. Tang and D. Mazzoni, "Multiclass reduced-set support vector machines," In Proceedings of the 23th International Conference on Machine learning, pages 921–928, New York, NY, USA, 2006. ACM Press.

D. M. Tralli, Blom R. G., Zlotnicki V., Donnellan A., Evans D. L., Satellite remote sensing of earthquake, volcano, flood, landslide and coastal inundation hazards (2005) ISPRS Journal of Photogrammetry and Remote Sensing, 59 (4), pp. 185-198.

K. L. Wagstaff, M. Kocurek, D. Mazzoni, B. Tang, "Progressive Refinement for Support Vector Machines," Data Mining and Knowledge Discovery, 20(1):53-69, doi:10.1007/s10618-009-0149-y, January 2010.