LEARNING CLASSIFIERS FOR SCIENCE EVENT DETECTION IN REMOTE SENSING IMAGERY

Rebecca Castano[(1)], Dominic Mazzoni[(1)], Nghia Tang [(1)], Thomas Doggett[(2)], Steve Chien[(1)], Ronald Greeley[(2)], Ben Cichy[(1)], Ashley Davies[(1)]

[(1)]Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA, 91109, United States,
Email: <firstname.lastname>@jpl.nasa.gov

[(2)] *Department of Geological Sciences, Arizona State University, Box 871404,Tempe, AZ 85287-1404*
Email: <firstname.lastname>@asu.edu

ABSTRACT

We describe four pixel-based classifiers that were developed to identify events in hyperspectral data onboard a spacecraft. One of the classifiers was developed manually by a domain expert, while the other three were developed using machine learning methods. The top two performing classifiers were uploaded to the Earth Observing-1 (EO-1) spacecraft and are now running on the satellite. Classification results are used by the Autonomous Sciencecraft Experiment of NASA's New Millennium Program on EO-1 to automatically target the spacecraft to collect follow-on imagery. This software demonstrates the potential for future deep space missions to identify short-lived science events and make decisions onboard.

1. INTRODUCTION

Data analysis onboard a spacecraft can enable the detection of and reaction to dynamic events. For example, timely reaction to an eruption event on Io would be possible if the event was detected onboard and the spacecraft equipped to react. With ground-based analysis such reactions are not possible

The Autonomous Sciencecraft Experiment (ASE) is a JPL-led, New Millennium Program mission containing new technology in the form of software which has been flying on the Earth Observer-1 (EO- 1) satellite since the fall of 2003 [1, 2]. This new technology facilitates autonomous science-driven capabilities. Among the ASE flight software is a set of onboard science algorithms designed for autonomous data processing, to identify observed science events [1, 2, 3]. Using the output from these algorithms, ASE has the ability to autonomously modify the EO-1 observation plan, retargeting itself for a more in-depth observation of a scientific event in progress with current response times on the order of hours. Several onboard science algorithms are associated with ASE for detecting dynamic events. Events detected include volcanic eruptions[4], floods[5] and cryosphere events[6].

In using learning algorithms to develop classifiers, we focused on cryosphere event detection. The cryosphere is the component on the surface of a planetary body composed of ice. The ice may exist in a variety of forms including snow, permafrost, floating ice, and glaciers. The cryosphere dynamically interacts with the atmosphere and can significantly affect the climate on a planetary body. Ices, including water ice and CO_2, are found throughout the solar system. In addition to the water ice caps on Earth and the water and CO_2 at the Martian poles [7], ice is thought to exist in the permanently shadowed polar craters of the moon [8] and Mercury [9], and is a major component of the moons in the outer solar system and as well as comets. Cryosphere events studied by the science team include the formation and break-up of sea ice as well as the freezing and thawing of lakes.

Four pixel-based classifiers to identify cryosphere events using hyperspectral images were developed. These classifiers include an expert-derived manually constructed classifier, a classifier derived by searching exhaustively over combinations of thresholded band ratios, a Decision Tree, and a Support Vector Machine (SVM). A set of scenes was labeled by hand by a domain expert to provide training and testing data. Performance results on the test data indicate that the SVM and manual classifiers outperformed the Decision Tree and band-ratio classifiers with the SVM yielding slightly better classifications on selected classes than the manual classifier.

In this paper, we first describe the background of the Autonomous Sciencecraft Experiment and the Earth Observing -1 spacecraft on which it is flying. We then explain the four classifiers that were developed for cryosphere classification and present performance results on a set of test images.

Proc. 'ISAIRAS 2005 Conference', Munich, Germany,
5-8 September 2005 (ESA SP-603, September 2005)

1.1 Spacecraft and Instrument

EO-1 is part of NASA's New Millennium Program, designed to validate new technologies for remote sensing. It was launched from Vandenberg Air Force Base on 21 November 2000 and placed in a sun-synchronous orbit with an altitude of 705 km and a 10:01 AM descending node, giving it an equatorial crossing time that is one minute behind Landsat-7 and a 16-day repeat path orbital cycle. With observations up to two paths off nadir, EO-1 is able to image the same location as many as 5 times every 16 days in daylight. The EO-1 payload is comprised of three instruments: Hyperion, Advanced Land Imager (ALI) and the Linear Etalon Imaging Spectral Array (LEISA) Atmospheric Corrector. We analyze data from the Hyperion instrument onboard the spacecraft.

The Hyperion instrument [10] consists of two push-broom imaging spectrometers, covering the visible/near infrared (VNIR) and short-wave infrared (SWIR), respectively, which share a common telescope, producing hyperspectral images with a 30 m/pixel spatial resolution and 10 nm/band spectral resolution. Hyperion images are 7.5 km in width, with an along track length that depends on the duration of the data collect, but typically 60 km (8 seconds) or 90 km (12 seconds). ASE analyzes a 7.5 km by 15 km subset of the captured image when detecting cryospheric events. The VNIR spectrometer has 50 calibrated bands, ranging from 0.43 to 0.93 μm, and the SWIR spectrometer has 148 calibrated bands, ranging from 0.91 to 2.4 μm.

There are two identical processors onboard the EO-1 spacecraft, one for the primary spacecraft operations and the other for the payload. ASE uses the payload processor. It is a Mongoose V CPU with a processor speed of 8 MIPS and 256 MB of RAM. With this hardware constraint, the Hyperion data cannot be fully processed from Level 0 (raw) data to Level 1 (radiometrically calibrated) data [11]. Instead the data are partially processed to an onboard product designated Level 0.5, using data from a dark calibration image collected within a few minutes of the actual image.

Features of Level 1 data processing [11] not performed in the onboard processing include smear and echo correction to the SWIR bands, as well as interpolation between pre- and post- dark calibration images before dark image subtraction. While both Level 0.5 and Level 1 data are identical in VNIR, they diverge in SWIR, where the lack of smear and echo correction in Level 0.5 gives higher values than in the fully processed data. Because Level 0.5 data are not fully calibrated, the radiance and reflectance values for SWIR bands calculated onboard the spacecraft can be considered as pseudo-radiance and pseudo-reflectance.

2. APPROACH

The approach to onboard identification of cryosphere events is to classify pixels in the image independently based on the available spectral information. Pixels are classified as belonging to one of five classes: water, ice, land, snow, or cloud. A sixth class of unknown or unclassified is also allowed for some of the classifiers. The challenge is to achieve a sufficiently high accuracy under conditions of limited CPU, few available spectral bands, and incomplete calibration.

As mentioned earlier, we developed four pixel-based classifiers. Each of the classifiers was designed to run in the computationally constrained operating environment of the spacecraft. Onboard constraints permit access to only 12 of the bands of the Hyperion instrument, although these 12 are selectable from the full complement.

2.1 Manual Classifier

The manually constructed cryosphere was developed by a domain expert [6]. This classifier was designed to be run in sequence after a first classifier identified all the cloud pixels. The preliminary pass cloud classifier is described in [12]. The manual cryosphere classifier utilizes a total of seven bands. The classifier was derived empirically using 175 spectra. The spectra were taken over several images from snow, water, ice, land and cloud pixels, focusing on clouds not detected by [12]. Image regions from the training data set for each of the five classes (land, water, snow, ice, cloud) were identified visually. All of the scenes were processed from Level 0 to Level 0.5 on the ground, and the spectra converted to reflectance before plotting in a spreadsheet to determine a sequence of band ratios that could separate the spectra into the classes defined. Band ratios and thresholds were selected to best distinguish the pixels selected. The set of training images were then classified and regions incorrectly classified were identified visually. The spectra were studied and a new decision layer was added to the classifier to correct significant misclassifications. This procedure was iterated until a sufficiently high accuracy was achieved. The final version of the classifier employs a series of twenty steps to classify snow, water, ice and land. Full details of the manual classifier are provided in [6].

2.2 Best Automated Ratio

The first automated algorithm developed was to identify the best thresholds of ratios of bands using an exhaustive search of all possible options. This was to be used as a point of comparison for the other algorithms. It is effectively a very simple decision tree.

Bands that exhibited a high percentage (>1 %) of noisy pixels in the training set were not used. A noisy pixel is defined as a pixel whose value was invalid. This happens most commonly when the signal is very low and is overwhelmed by the noise in the dark image. There is a clear correlation between noisy bands and H_20 (and to a lesser degree O_2 and CO_2) absorption bands. Also, surfaces like water, shadow and lava that are less bright suffer from a low signal-to-noise ratio.

After removing the noisy bands 150 bands remained for use in the experiment. With 150 bands, there were 11175 possible band ratios to consider. For each ratio, the optimal threshold was determined where the decision was a one-vs-all classification. For example, the optimal threshold for classifying pixels as water or not water was determined for every possible ratio. This was done for each of the five classes. For each class, the ratio with the highest classification accuracy was selected for the classifier. The classifier was constructed by sequentially applying the individual ratio/threshold classifiers. For example, the water vs. other classifier was applied. The classifier for snow vs. other was then applied to all pixels not classified as water, etc. The order was determined by applying the individual classifiers in order of accuracy. The evaluation was done with both radiance and reflectance data. The final classifier is shown in Fig. 1.

In addition to the band ratio classifier, thresholding of individual bands and thresholding of normalized differences between bands were both considered. Thresholding of individual bands is undesirable because it is more sensitive to the uncalibrated data available onboard than band ratios. The normalized band differences were not used because they yield exactly the same result as the band ratios. This can be seen as follows. If $\vec{x}$ and $\vec{y}$ represent the spectral data at two different pixels and B_1^x and B_1^y represent the values of Band 1 for pixel $\vec{x}$ and pixel $\vec{y}$ respectively then

$$\frac{B_1^x}{B_2^x} > \frac{B_1^y}{B_2^y} \Leftrightarrow \left(\frac{B_1^x - B_2^x}{B_1^x + B_2^x}\right) > \left(\frac{B_1^y - B_2^y}{B_1^y + B_2^y}\right). \quad (1)$$

```
If ( B16^refl / B36^refl  > 1.42)
      Pixel = Water
Elseif ( B18^raw / B44^raw < 1.37)
      Pixel = Land
Elseif ( B37^rad / B43^rad > 1.14)
      Pixel = Ice
Elseif ( B188^raw / B217^raw < 1.33)
      Pixel = Cloud
Elseif ( B120^refl / B140^refl  >  0.93
      Pixel = Snow
Else
      Pixel = Unclassified
```

Fig. 1. Automated Ratio Classifier. This is the best classifier that was identified using a single ratio per class. Band values are either raw, reflectance, or radiance.

It can be seen from the classifier in Fig. 1, that at most four divisions for the band ratios and four conditional statements must be evaluated for this classifier.

2.3 Decision Tree

The second automated classifier developed was a decision tree classifier. A decision tree is a logical model represented as a binary (two-way split) tree that shows how the value of a target variable (pixel class) can be predicted by using the values of a set of predictor variables (pixel data). A decision tree is built through the process of binary recursive partitioning. This is an iterative process of splitting the data into partitions, and then splitting it up further on each of the branches, until each branch has data of only one class. A tree that is fully grown from the training set usually suffers from over fitting of the data, resulting in poor performance on a test data set. The solution is to prune the tree before full depth is reached. Decision tree methods are effective for a number of classification tasks and a benefit is that they generation rules that can be easily understood and explained. In this work, the standard C4.5 algorithm [13] was used to develop the decision tree classifier.

A number of experiments with the decision trees were conducted. Three different sets of features were used in these experiments:
- 11 selected bands
- 55 band ratios (all ratios of the 11 bands) and
- Principle Component Analysis (PCA) Eigenvectors derived from the training data

Refinement with respect to feature selection was done incrementally. Upon completion of a test, additional features were added to the training set and tested. If

the additional training features produced a better classifier then the original, they were kept, other wise they were not used. The selection of which features were added was done randomly. While it would probably not be computationally feasible to compute the PCA onboard, these features were used in the experiments to determine if increased accuracy would be possible using this method if more computational power were available.

A number of decision trees were trained and tested using combinations of the above feature sets. In several cases, the decision trees had particular difficulty distinguishing between two of the classes. Which two classes were confused varied across experiments and combinations of features, however common problems were distinguishing between land and cloud and distinguishing between cloud and snow. To address the issue, a secondary decision tree was trained on only the two classes in question. This second classifier was then chained to the first and applied as a discriminator only when the first classifier identified a pixel as belonging to one of the two poorly distinguished classes.

We found that the best performance was achieved using the 11 selected bands in combination with a secondary cloud vs. land classifier. Additional information did not increase accuracy, and in many cases, led to reduced accuracy. For example, one of the poorest classifiers employed both the 11 individual bands and the 11 PCA bands.

The maximum number of operations required is a function of the depth of the final tree. The depth of both the primary and secondary trees is 12. Therefore a worst case of 24 branches could be taken to do a single pixel classification. On the average, the decision tree is able to come to a decision at a much earlier branch. Since the tree did not use band ratios, only a comparison is required at every level of the tree. Thus, the worst case complexity is 24 floating point comparison operations.

2.4 Support Vector Machine (SVM)

The final cryosphere classifier used a Support Vector Machine (SVM). SVMs [14] are a family of classifiers that identify the optimal linear separator between classes in a (possibly) high dimensional space. There are two aspects of SVM's that make them appealing. First, SVMs belong to a special category of machine learning techniques that produce empirically derived classification algorithms by explicitly attempting to maximize the *margin*, the boundary which separates one type of classified of data from another. Maximizing the margin prevents the algorithm from over fitting the training data, which can lead to poor algorithm performance on new input data sets. Second, nonlinear decision boundaries are supported by mapping the input feature space to a higher dimension (possibly infinite) where the features are linearly separable. The key is that the mapping is carried out implicitly so that all that is necessary is to compute dot products in the original input feature space.

To construct the SVM classifier, the same 11 bands that were used with the manual classifier were used. In this work, emphasis was not placed on band selection.

A primary disadvantage of support vector machines is that the time to classify points is proportional to the number of support vectors, which for complicated problems can be as large as the number of training examples. Since we wanted to train on thousands of example pixels, this would result in unreasonably slow classification. Instead, we decided not to use a kernel function, leaving us with a linear support vector machine that can be evaluated quickly. Exactly 11 multiplies and 11 additions are required for each of the five classes, plus a comparison operation to determine the correct classification after each of the five evaluations. Evaluating the SVM was thus not substantially more computationally challenging than the other methods.

The SVM approach is somewhat more sensitive to bad or missing data than the other methods. While a decision tree can potentially ignore certain branches or take default branches when particular spectral bands are missing, the SVM approach considers all spectral bands simultaneously and gives meaningless results if any of them are missing or bad. Thus, any time a band was missing, we marked that pixel as unclassified. This typically only represented 1% of the pixels in an image.

3. EXPERIMENTAL RESULTS

The training set consisted of 50 images labelled by a domain expert. The test set was also 50 images with just over 91 million pixels. The results on the test data are shown for the manual classifier, the decision tree and the SVM. The results shown for the best ratio classifier are for a slightly different test set that is comparable. The results are shown in terms of precision and recall which are defined as follows. For a given classifier, we can define the true positives (TP) for each class as those pixels that were correctly classified, e.g. all pixels that are labelled as water and classified as water. The false positives (FP) are those pixels that are classified as belonging to a class, but labelled as belonging to another class, e.g. classified as water, but labelled as land. Finally, the false negatives (FN) or misses are defined as those pixels that are

labelled as belonging to the class, but are otherwise classified, e.g. labelled as water, but classified as ice. With these definitions, the precision of a classifier for a given class is defined as

$$P = \frac{TP}{TP + FP} * (100), \quad (2)$$

while the recall of a classifier for a given class is defined as

$$R = \frac{TP}{TP + FN} * (100) \quad (3)$$

The precision indicates the percent of pixels classified as a belonging to a particular class that truly do belong to the class according to the labels. In contrast, recall indicates the percent of pixels labelled as a class that were actually classified as belonging to that class. Precision and recall results for the four classifiers on the test data are shown in Tables 1 and 2.

Table 1. Precision on test set.

	Expert Derived	Best Ratio*	Decision Tree	SVM
Cloud	77.2	77.4	80.2	83.3
Ice	53.3	34.2	31.0	39.9
Land	93.7	93.0	92.5	95.3
Snow	71.4	76.5	72.5	78.5
Water	81.0	59.1	73.5	63.6

Table 2. Recall on test set.

	Expert Derived	Best Ratio*	Decision Tree	SVM
Cloud	55.1	77.4	57.2	65.7
Ice	49.7	34.2	62.4	47.8
Land	95.7	93.0	93.0	95.2
Snow	84.3	76.5	76.7	77.4
Water	91.0	59.1	91.2	66.1

Within the test set, the number of pixels belonging to each of the five classes was not evenly distributed. The distribution of pixels over the five classes is shown in Table 3. The overall accuracy of each classifier is given in Table 4. The expert-derived, best ratio, and SVM classifiers all have an unclassified label that can be assigned to a pixel. In these tests, water pixels were the most common class in the set of unlabeled pixels. In particular, this was a factor with the SVM as it requires valid values for all elements of the feature vector and, as explained earlier, the signal to noise ratio for water is low leading to a higher number of bad pixels for this class.

It should be noted, that these results do not necessarily indicate the full power of the learning methods as they were implemented for the onboard computational environment. In particular, a non-linear kernel would be used on a desktop machine to achieve a higher level of accuracy. The overall accuracy is very similar for the four classifiers, but it can be seen that the simplest approach cannot reach the level of performance of an expert derived or more advanced learning approaches.

Table 3. Distribution of pixels in test set.

Class	Abundance in test set (%)
Cloud	17
Ice	2
Land	59
Snow	14
Water	8

Table 4. Overall performance accuracy of the four classifiers.

	Expert Derived	Best Ratio*	Decision Tree	SVM
With Unclassified Class	82.9	76.6	80.9	81.3
Unclassified assigned to Water	83.0	76.7	80.9	83.8

An example image is shown in Fig. 2. A number of observations can be made concerning these results. This accuracy assumes that all classes are equally important, that is each pixel is equally weighted independent of its label.

3.1 Onboard Application of the Classifiers

The expert derived classifier and the SVM were both uploaded to the EO-1 spacecraft and are now running onboard. These classifiers are being used to identify cryosphere events such as lake freezing or thawing or sea ice breaking up. As described in Section 3, due in part to the use of uncalibrated data and limited onboard processing power, the expected accuracy at the pixel level for these classifiers is just over 80%. In addition, there is uncertainty in the along track error of the spacecraft. The one second timing uncertainty results in approximately a 7 km or 230 pixel uncertainty. Combined, these two factors make onboard pixel differencing at the spectral or label level infeasible.

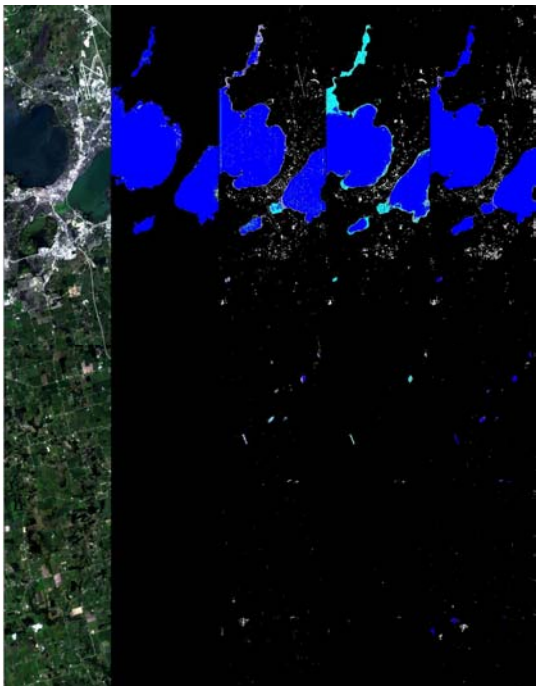

Fig. 2. Example Image. The left pane shows a false color image of a scene near Madison, Wisconsin taken in the summer. The next pane shows the expert labeling, where blue is water and black is land. The third pane contains the results of the expert-derived classifier. The fourth and fifth panes are the automated ratio and SVM results respectively. The cyan pixels that are seen in both the expert-derived and automated ratio results are pixels that have been misclassified as ice. Note the blue specs in all three of the classifier results are actual small bodies of water that the expert did not label.

A robust method for identifying the change events was implemented that uses the ratio of pixels belonging to various classes to assess the presence of change. The presence of sufficient valid data is first assessed (e.g. the scene is not completely cloud covered). If these criteria are met, a second criterion is analyzed to determine the presence of the cryosphere event. Fig. 3 shows an example used to determine sea ice break-up. The criteria, including thresholds, were determined by the science domain expert.

On September 22, 2004, the SVM classifier was run onboard EO-1 and correctly classified the image of Lake Winnibigoshish, Wisconsin, shown in Fig. 4 as cloudy. On December 1, 2004, the SVM classifier analyzed a scene of South Georgia Island near Antarctica and successfully identified open water

$$\text{If} \left(\frac{|Cloud| + |Unclassified|}{Total\_pixels} < 60\% \right)$$

$$\text{And} \left(\frac{|Snow| + |Ice|}{|Snow| + |Water| + |Ice|} < 86\% \right)$$

```
Then
      Event = True
Else
      Event = False
```

Fig. 3. Trigger criteria for sea ice break-up.

indicating sea ice break-up (See Fig. 5). As a result of the onboard analysis a reimage of the scene was triggered and collected on December 3, 2004. This was the first use of the SVM classifier being employed as an onboard science data analyzer to automatically trigger a spacecraft reaction. The SVM classifier is now in regular use as part of the ASE software onboard EO-1.

Fig. 4. Image of Lake Winnibigoshish, Wisconsin taken September 22, 2004. The scene was correctly classified as cloudy by the onboard SVM classifier.

4. CONCLUSIONS

Four classifiers for hyperspectral data were developed and tested. The machine learning classifiers had comparable results to the manually developed classifier, although they took an order of magnitude less of the experts' time to derive. The classifiers are presently being used on the EO-1 spacecraft to automatically trigger the collection of follow-up images of scenes based on the identification of cryosphere events. One key to the success of the onboard classifiers is that they do not need to operate at the fidelity of ground-based science data analysis methods. The objective of the data analysis on a

spacecraft is to detect events. Once the data has been downloaded, the full data is available for thorough scientific analysis on the ground.

In situ analysis can be used to trigger the collection of additional data, as with ASE onboard EO-1, but can also be used to change the data collection rate of an instrument when a dynamic event is detected or to identify high priority data for downlink. The use of machine learning classifiers for onboard data analysis shows great potential for use in future deep space missions, where the round trip messaging times make the reaction to dynamic events difficult to impossible with the traditional ground-in-the-loop approach.

ACKNOWLEDGEMENTS

This work was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Funding was provided by the NASA's New Millennium Program and the Interplanetary Network Directorate Technology Development program.

REFERENCES

[1] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication* . April 2005.

[2] Davies, A. G., V. Baker, R. Castano, S. Chien, B. Cichy, T. Doggett, J. M. Dohm, R. Greeley, R. Lee, R. Sherwood (2004) Autonomous Sciencecraft Experiment (ASE) Operations On EO-1 In 2004, *Lunar. Plan. Sci. Conf. XXXV*, abstract #1700.

[3] Sherwood, R., S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, and G. Rabideau, "Preliminary Results of the Autonomous Sciencecraft Experiment," *IEEE Aerospace Conference*. Big Sky, MT. March 2004.

[4] Davies, A. G., S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castaňo, B. Cichy, G. Rabideau, D. Tran and R. Sherwood (2005) Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE), *Rem. Sens. Environ*., in review.

[5] Ip, F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, B. Castaňo, S. Chien, B. Cichy, R. Greeley, R. Sherwood (2005) Development and Testing of the Autonomous Spacecraft Experiment (ASE) floodwater classifiers: Real-time Smart Reconnaissance of Transient Flooding, *Rem. Sens. Environ*., in review.

[6] Doggett, T., R. Greeley, S. Chien, R. Castano, B. Cichy, A. G. Davies, G. Rabideau, R. Sherwood, D. Tran, V. Baker, J. Dohm and F. Ip, "Autonomous Detection of Cryospheric Change with Hyperion on-board Earth Observing-1", *Rem. Sens. Environ*., in review.

[7] Thomas, P., S. Squyres, K. Herkenhoff, A. Howard and B. Murray, "Polar Deposits of Mars," in *Mar*s, edited by H. Kieffer, B. Jakosky, C. Snyder and M. Matthews, pp. 767-798, University of Arizona Press, Tucson, Arizona, 1992.

[8] Feldman, W.C., S. Maurice, A.B. Binder, B.L. Barraclough, R.C. Elphic and D.J.Lawrence, "Fluxes of fast and epithermal neutrons from Lunar Prospector: Evidence for water ice at the lunar poles," *Science*, 281, 1496-1500, 1998.

[9] Slade, M.A., B.J. Butler and D.O. Muhleman, "Mercury radar imaging: Evidence for polar ice," *Science*, 258, 635-640, 1992.

[10] Pearlman, J.S., P.S. Barry, C.C. Segal, J. Shepanski, D. Beiso and S.L. Carman, "Hyperion,

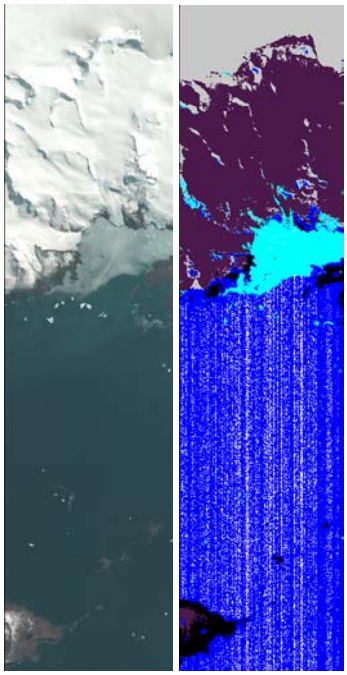

Fig. 5. Image of South Georgia Island near Antarctica taken December 1, 2004. The left is the false color image while the right shows the resulting SVM classification, where blue is water, black is land, cyan is ice, purple is snow, gray is cloud, and white is unclassified. Open water was correctly identified indicating sea ice break-up and triggering another image of the scene to be taken on December 3, 2004.

a Space-Based Imaging Spectrometer," *IEEE Trans. Geosci. Rem. Sens.*, 41, 1160-1172, 2003.

[11]Barry, P. *EO-1 / Hyperion Science Data User's Guide, Level 1_B*, Redondo Beach, CA: TRW Space, Defense & Information Systems, 2001.

[12]Griffin, M.K., H.K. Burke, D. Mandl and J. Miller, "Cloud Cover Detection Algorithm for EO-1 Hyperion Imagery," *17th SPIE AeroSense 2003, Orlando FL, Conference on Algorithms and Technologies for Multispectral, Hyperspectral and Ultraspectral Imagery IX*, 2003.

[13] Quinlan, J. R. *C4.5: Programs for Machine Learning.* San Mateo, CA. Morgan Kauffmann, 1993.

[14] Cortez, C., and V. Vapnik, Support vector networks, *Machine Learning*, **20**, 273 – 279, 1995.