J6.2   THE APPLICATION OF SUPPORT VECTOR MACHINES TO ANALYSIS OF GLOBAL SATELLITE
DATA SETS FROM MISR

Michael J. Garay*
University of California, Los Angeles, California

Dominic Mazzoni, Roger Davies, and David J. Diner
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California

1. INTRODUCTION

The Multi-angle Imaging SpectroRadiometer (MISR) is one of a suite of five instruments onboard NASA's Terra EOS satellite, launched in December 1999. Typical satellite imagers view the earth from a single direction, but MISR's cameras image the earth simultaneously from nine different directions in four spectral bands. In this way, MISR provides unique multiangle information about solar radiation scattered from clouds, aerosols and other terrestrial surfaces. One of the primary goals of the MISR mission is to improve our understanding of how clouds and aerosols affect the earth's global energy balance.

MISR has now been operational for over four years, acquiring more than 50 Terabytes of high quality data at a resolution of 1.1 km x 1.1 km, with a subset of data at an even higher 275 m x 275 m resolution. In order to make use of this amount of data on a global scale, we have applied Support Vector Machines (SVMs) to MISR imagery classification. SVMs are a type of supervised learning algorithm, in the same category as Artificial Neural Networks, Decision Trees, and Naïve Bayesian Classifiers. Using SVMs, we have developed a global cloud mask, a global cirrus cloud detector and a combination global aerosol detector and classifier.

2. BACKGROUND

This section describes the data available from MISR, other applications of machine learning to satellite image classification, with specific emphasis on cloud detection and classification, and provides a brief introduction to Support Vector Machines and the existing MISR cloud detection algorithms.

2.1 MISR Data

The Terra satellite, which carries MISR, is in a sun-synchronous, polar orbit that crosses the equator every 99 minutes during the descending portion of the orbit at approximately 10:30 a.m. local time. The MISR instrument consists of nine pushbroom cameras with an average swath width of about 360 km (see Fig. 1). The cameras sample data in four spectral bands – blue, green, red and near-infrared – from nine different directions: ±70.5°, ±60°, ±45.6°, ±26.1° and 0° (nadir).

* *Corresponding author address*: Michael J. Garay, Univ. of California, Los Angles, Dept. of Atmospheric & Oceanic Sciences, Los Angeles, CA; E-mail: garay@atmos.ucla.edu

MISR has a repeat cycle of 16 days, but, because of overlap, the instrument obtains global, multiangle coverage of the entire earth in nine days at the equator and two days at the poles. For more detailed information on MISR see Diner et al. (1998) or the MISR website http://www-misr.jpl.nasa.gov.
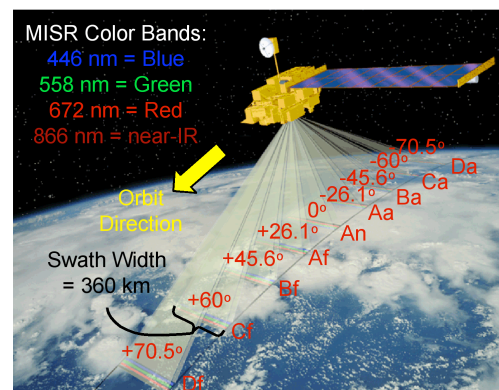

Figure 1. Computer-generated image of NASA's Terra EOS satellite with the MISR instrument onboard. Sampling is as high as 275 m x 275 m per pixel. To reduce the data volume, most images are obtained at a resolution of 1.1 km x 1.1 km. *Image courtesy of Shigeru Suzuki and Eric M. De Jong. Solar System Visualization Project. JPL Image P-49081.*

As MISR passes over the earth, the most forward looking camera (Df +70.5°) images a point 2800 km away from the point imaged by the most aftward camera (Da –70.5°). It takes seven minutes for a point on the surface to be imaged by all nine MISR cameras as the instrument passes over. When the data is processed, the cameras are registered so that each point on the surface is effectively seen from nine different viewing directions by a "virtual" MISR instrument (see Fig. 2). In actuality, the cameras are registered either to an idealized earth "ellipsoid" or to terrain provided by a digital elevation model (DEM). Because they are physically closer to the instrument, objects above the surface, such as clouds and some types of aerosol, appear to move relative to the surface from one camera to the next due to the effect of parallax.

The processed MISR data is stored in Hierarchical Data Format (HDF) and is available from the NASA Langley Research Center Atmospheric Sciences Data Center.

2.2 Other Applications of Machine Learning to Satellite Cloud Imagery Classification

Beginning as early as Shenk et al. (1976), artificial intelligence methods have been applied to the problem of cloud detection and classification in satellite imagery. Computer pattern recognition was used by Ebert (1987), Garand (1988) and Ebert (1989). Neural network approaches, by far the most popular, were adopted by Key et al. (1989), Lee et al. (1990), Peak and Tag (1992; 1994), Bankert (1994), Bankert and Aha (1996), Miller and Emery (1997) and Tian et al. (2000). This list is by no means exhaustive. Usually these artificial intelligence approaches have been applied only to a small range of scenes or within a geographically limited region. Few artificial intelligence methods have been applied to global satellite data.
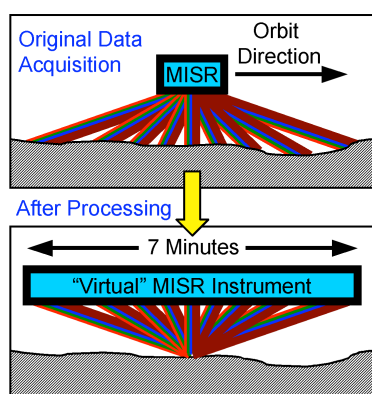
Figure 2. When data is initially acquired, each MISR camera views a different portion of the earth. After processing, all nine MISR cameras are registered to the same point on the surface.

2.3 Support Vector Machines

Support Vector Machines are similar to other supervised learning algorithms (Artificial Neural Networks, Naïve Bayesian Classifiers, etc.) in that they learn to classify new data based on fully labeled training data. In the applications discussed here, SVMs have been used to classify individual pixels in MISR images. Each pixel is classified using "feature vectors" that incorporate information from multiple MISR cameras and spectral bands.

This paper is not an introduction to SVMs; for the purposes of this discussion SVMs can be considered to be "black boxes" that take labeled feature vectors as input and produce deterministic classification algorithms that can be applied to classify feature vectors from novel data. Readers who are interested in the details of SVMs should refer to Cortes and Vapnik (1995) or the excellent website http://www.kernel-machines.org for more information.

Most of the techniques used to develop the MISR classifiers described here could be applied equally well to any other supervised classification technique.

However, SVMs have some important properties that make them an appropriate choice for this type of work. First, the SVM training algorithm is deterministic and avoids becoming trapped in local minima during optimization. Second, the training algorithm inherently incorporates a technique to balance maximal accuracy with maximal generalization performance. While not a panacea, this does help to mitigate problems of overfitting often seen with other learning classifiers (Schölkopf and Smola 2002). Finally, because SVM training depends only on the distances between feature vectors, which can be computed in a single preprocessing stage that is inherently parallelizable, it is possible to explore very large feature vectors (hundreds of features), making the difficult feature selection stage less important.

2.4 Existing MISR Cloud Detectors

Cloud detection and screening are important for the MISR mission's scientific objectives, some of which require accurate retrievals of aerosol and surface properties. Three cloud detectors – called cloud "masks" – have been developed by the MISR science team specifically for use with MISR. These cloud masks are designed to classify each 1.1 km x 1.1 km MISR pixel individually as either cloudy or clear, with high confidence or low confidence.

The Radiometric Camera-by-camera Cloud Mask (RCCM) uses the mean and variance of the red and near-IR radiances from each of MISR's nine cameras individually to identify cloudy and clear regions. A statistical method provides dynamic thresholds that separate cloudy pixels from clear pixels based on the input data (Diner et al. 1999b). As shown in Fig. 3 the RCCM is extremely accurate over water, where good thresholds have been implemented. However, the RCCM is less reliable over land, where the thresholds are still in the development stage. Additionally, the RCCM has difficulty correctly classifying specular reflection from bodies of water, commonly called "sun glint" or "sun glitter." For this reason, a conservative, geometrically derived "glitter mask," specific for the geometry of each camera, is used to designate areas in an image where sun glint might be expected to cause problems for the RCCM retrieval.

The Stereoscopically Derived Cloud Mask (SDCM) takes advantage of MISR's multiangle capability. A feature matching algorithm is used with combinations of MISR's cameras to determine the height of observed reflecting surfaces (Moroney et al. 2002). If this height is greater than the height of the terrain at that location, given by the digital elevation model, then the associated pixel is assumed to be cloudy (Diner et al. 1999a). Figure 3 shows that the SDCM is only slightly less accurate than the RCCM, but the SDCM performs equally well over land and water. However, there are times when the feature matching algorithm is unable to find a satisfactory match between features in multiple cameras. This can occur, for example, when clouds in

the scene are extremely homogeneous horizontally. As a result, roughly one third of the MISR image pixels are left unclassified by the SDCM.

Finally, the Angular Signature Cloud Mask (ASCM) was developed based on the Band-Differenced Angular Signature (BDAS) approach, described by Di Girolamo and Davies (1994). BDAS utilizes both spectral and angular changes in reflectivity to distinguish clouds from the background. The ASCM has only recently become a MISR operational product and was not used for comparison in the work described here.

Each of these cloud masks works extremely well under certain conditions, but no single one appears to be robust enough for highly accurate *global* cloud detection. For this reason, an artificial intelligence approach was applied to the MISR global cloud detection problem. This was the beginning of work that led to the further development of a cirrus cloud detector and a combined global aerosol detector and classifier.

3. APPROACH

Our initial goal was specifically to explore the ability of SVMs to produce a global cloud detection algorithm more accurate and robust than the two MISR cloud masks (RCCM and SDCM) operational at the time. This led, in a logical way, to the development of new SVM applications for other MISR imagery classification tasks.

3.1 Global Cloud Mask

Rather than attempting to build a cloud mask from first principles, we decided instead to leverage the strengths of the existing cloud masks to develop the new SVM cloud mask. We began by making the unconventional choice of constructing feature vectors out of a large number of "raw" features, instead of using a smaller number of higher-level features. This was motivated, in part, by the observation that SVMs were successful at classifying images of handwritten digits simply using grayscale pixel values as inputs, with no higher-level features (DeCoste and Schölkopf 2002).

For each 1.1 km x 1.1 km pixel of MISR data, we constructed a feature vector consisting of the red-band radiances of each of the 5x5 square of pixels centered at that point, using the three most nadir-pointing MISR cameras (Af +26.1°; An 0°; Aa –26.1°). Then, the radiances from the green, blue and near-infrared bands were incorporated for the smaller 3x3 region of adjacent pixels. Next, 25 features containing the standard deviation of the higher-resolution data within each pixel in the 5x5 region from the nadir camera (An) red band were added. Finally, one additional feature, the normalized solar zenith angle, brought the total set of inputs to the feature vector to 182 per pixel. Because there are well-known difficulties in cloud detection due to persistent snow and ice at high latitudes, we confined our "global" cloud detector to latitudes less than 60° in both hemispheres.

Since our goal was to train a single classifier that would be useful globally, it was necessary to obtain a large number of training labels for a variety of scene types. Faced with the daunting possibility of needing to individually label potentially tens of thousands of pixels as cloudy or clear around the globe, we instead settled on an approach that piggybacked on the existing cloud masks. We constructed several million feature vectors from randomly distributed MISR pixels. Then, we used those pixels where the RCCM and SDCM were most confident to label a subset of these globally distributed pixels. Specifically, we used the highest-confidence RCCM labels over glint-free ocean and the highest confidence SDCM labels over land and shallow water. Because each MISR pixel is already categorized as being located over ocean, land, or shallow water, based on a geographic land/water mask, this was a natural distinction to make.

Due to inaccuracies in both the RCCM and SDCM, it was likely that many of our training labels were incorrect. Even so, we expected that the SVM would be tolerant of these incorrect labels and still be capable of appropriate generalization. To calibrate the training process and validate the results, approximately 3,500 pixels were hand-labeled by an expert as cloudy or clear based on visual inspection of the images. Finally, the SVM training parameters were chosen that maximized the resulting classification accuracy on this set of test vectors.

3.2 Global Cirrus Cloud Detector

MISR is an ideal instrument for studying cirrus clouds, which play an important role in the earth's global energy balance but often go undetected by nadir-pointing instruments because they are optically thin. Scenes that appear to be clear in a nadir view can sometimes appear completely cloud covered in MISR's 70° view because the pathlength of radiation through the cloud is increased by a factor of three at this angle.

Based on our initial success with the global cloud classifier, the feature vectors used for the SVM cirrus cloud detector were exactly the same as those described for that case. The primary difference was that radiances from the 45.6°, 60° and 70.5° cameras were chosen to exploit the increased pathlength of radiation through clouds at these angles relative to nadir-pointing instruments. The forward-pointing cameras (Bf +45.6; Cf +60°; Df +70.5°) were used in the northern hemisphere and the aftward-pointing cameras (Ba –45.6°; Ca –60°; Da –70.5°) were used in the southern hemisphere to make the most advantage of forward scattering in either hemisphere. Labels for each feature vector were designated relative to the center pixel of the 60° camera; i.e., a vector was only labeled as cloudy if the central pixel in the 60° camera contained part of a cirrus cloud.

Because no MISR cloud mask currently detects only cirrus clouds, it was not possible to gather training data

from existing classifiers as was done in developing the original global cloud classifier. Instead, all the training labels for this case were provided by an expert using an interactive, graphical tool in development at the Jet Propulsion Laboratory called PixelLearn. This tool, originally devised for the MISR work described here, but later adapted to work with MODIS and Hyperion data, allows the user to explore multispectral, multiangular images, click on pixels and regions to label them as one of several classes, then interactively train an SVM to classify this data and immediately see the results of the SVM classification based on the labels provided. This tool made it possible to build an initial cirrus detector very rapidly and then focus the majority of the subsequent effort on reducing errors in those few places where it was clear that the classifier was making blunders, such as in regions with thin stratus clouds, by providing additional training labels.

3.3 Global Aerosol Classifier

To increase the level of complexity of the classification task, we next focused our attention on constructing a pixel-level classifier that would not only identify aerosols – in much the same way a cloud mask identifies clouds – but that would also classify aerosols into a variety of types. This task is slightly more complicated than the "binary" classification task of cloudy vs. clear pixels.

We began by adapting the original global MISR SVM cloud detector, which used only the three near-nadir cameras. Our initial expectation was that it would be difficult for the SVM to discriminate between aerosols and clouds. A number of training labels were provided by an expert for a variety of aerosol types, including smoke, dust and a generic "pollution" category. To our surprise, we found that the SVM was able to separate most clouds from aerosols with ease, apparently due to the spectral signatures present in the aerosols, which are not present in most clouds. Additionally, the combination of spectral and multiangle data allowed the aerosols themselves to be discriminated from one another. To increase the likelihood of aerosol detection by increasing the pathlength of radiation through the aerosol, the near-nadir cameras used initially were replaced by the most oblique (Bf, Cf and Df) cameras. We restricted this study to the northern hemisphere at latitudes less than 60°, to allow for preferential forward scattering into the forward bank of MISR cameras.

Aerosol classification was performed in a stepwise fashion, with pixels being assigned a single class by an SVM developed specifically to discriminate that class from all others. The complete algorithm was capable of distinguishing six different classes: clouds, clear water, clear land, smoke, dust and general pollution. Once properly trained based on expert labels, the algorithm was allowed to run independently on unclassified MISR data to identify new scenes containing a significant amount of aerosol.

4. RESULTS

The following section describes the results of the three classification tasks in a variety of ways. Examples have been selected to be illustrative of the capabilities of the SVM classifiers, rather than comprehensive assessments of their performance.

4.1 Global Cloud Mask
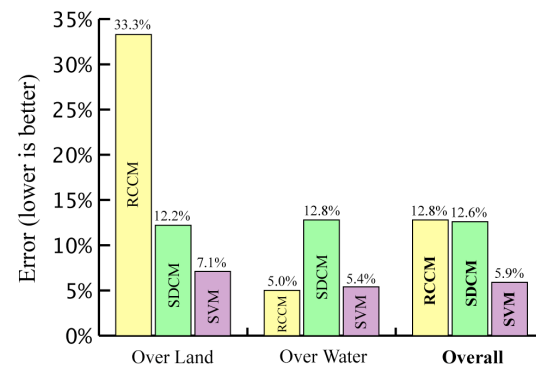
Figure 3: The graph depicts the performance of MISR's RCCM (yellow) and SDCM (green) cloud masking algorithms compared with the SVM global cloud mask (violet). The y-axis shows the error relative to expert labels. Smaller numbers indicate better performance. Performance is compared for three cases: just land scenes, just water scenes and all scenes.

Once training was completed for the SVM global cloud mask, using the approach described in Section 3.1, the performance of all three cloud masks was assessed using 3,500 labels provided by an expert. Figure 3 shows the results of this comparison. Results were obtained for classification over only land and only water, as designated by the MISR geographic land/water mask, and for all pixels independent of their land/water designation. Classifier performance is shown in Fig. 3 in terms of overall percentage error. A classifier that matched the expert labels perfectly would have an error of 0%. A penalty is assessed equivalently whether a cloudy pixel is misclassified as clear or vice versa. However, failure to classify a pixel as either cloudy or clear did not result in a penalty in this case.

Over land, where the RCCM thresholds were still being actively developed, the RCCM performed the worst of the classifiers, misclassifying one third of the pixels. The SDCM performed significantly better, misclassifying only 12.2% of the pixels relative to the expert labels. The SVM classifier, which was trained based on inputs from these other two classifiers over all scenes, was able to generalize appropriately and cut the classification error nearly in half relative to the SDCM, only misclassifying 7.1% of the land pixels.

The situation over water reveals a completely different story. The thresholds used by the RCCM are highly accurate over water, as evidenced by the extremely low 5% classification error for the RCCM. Because the

SDCM classifies pixels based on height above the surface, the performance of this detector was relatively unchanged. The SVM detector was only slightly less accurate than the RCCM, with an error rate of 5.4%. Comparison the SDCM results over land and water shows a difference of 0.6%, which can be taken as a rough estimate of the confidence of this assessment, since the performance of the SDCM is expected to be independent of the type of surface, purely on physical grounds. Therefore, it can be fairly said that the RCCM and the SVM classifier essentially performed equivalently for water pixels.

Globally, the RCCM and SDCM performed about equally well relative to the expert labels. In either case, approximately 13% of the pixels were misclassified by the operational MISR cloud masks: 12.8% for the RCCM and 12.6% for the SDCM. However, the performance of the RCCM is heavily biased due to its poor performance over land. Improved land thresholds have been incorporated into the MISR operational product since this study was conducted. The SVM classifier shows significant improvement over the other two methods, reducing the global pixel level error by more than a factor of two, to 5.9%. This result shows that an SVM trained on the existing cloud masks, both of which have a larger global error, is capable of generalizing in such a way as to improve the overall effectiveness of the cloud classifier. Additionally, if those pixels were eliminated where the SVM output was less certain (where the output value fell between -1.0 and 1.0), the remaining 70% of the pixels were classified with only 3% error.

4.2 Global Cirrus Cloud Detector

Due to the lack of ground truth data, the assessment of the performance of the cirrus cloud detector was primarily a visual one. Fig. 4 (found at the end of the paper) is an example scene containing some thin cirrus. The left-hand panel shows the typical nadir-looking view from MISR's An camera. The red, green and blue channels have been combined to provide a realistic RGB image. The center panel shows the same scene from MISR's Cf (+60°) camera. Radiation passing through the clouds in this image has twice the pathlength compared to radiation passing through the clouds in the left-hand panel. Note the cloud in the center left of the Cf image is slightly brighter than in the left-hand image. Very careful inspection of the image will also show that this same cloud has moved relative to the ground between these two images. This is due to the effect of parallax because the cloud is above the ground and closer to the instrument. The image on the right shows the result of applying the global cirrus cloud detector. The global SVM cloud detector was used to detect all the clouds in the scene; colored blue here. Cirrus clouds detected using the new SVM cirrus cloud detector appear in light yellow.

4.3 Global Aerosol Classifier

Here we are interested in showing the advantage derived from using MISR's multiple views, as opposed to the traditional single view for the global aerosol classifier. One of the advantages of the machine learning approach is that we did not need to add any information about how these additional features from the additional cameras were related. Simply augmenting each feature vector with additional information tended to improve the overall accuracy.

As specific evidence of this, the two confusion matrices shown in Fig. 5 compare the ability of a multi-class SVM to distinguish between six classes when given feature vectors containing only one MISR camera vs. three MISR cameras. In a confusion matrix, each row represents a class of true pixels according to expert labels and each column represents the classification decision made by the machine learning classifier. A perfect classifier, then, would show 100% correct classification down the main diagonal.

Using single Nadir view: 29.9% error

	Pollution	Dust	Smoke	Land	Water	Cloud
Pollution	44.2%	0%	20.7%	35.1%	0%	0%
Dust	0%	53.3%	3.1%	43.6%	0%	0%
Smoke	25.5%	39.6%	23.4%	0%	0%	11.5%
Land	0%	0%	0%	100.0%	0%	0%
Water	0%	0%	0%	0%	100.0%	0%
Cloud	0%	0.1%	0%	0%	0%	99.9%

Using three oblique angles from MISR: 15.6% error

	Pollution	Dust	Smoke	Land	Water	Cloud
Pollution	33.8%	0%	52.4%	13.8%	0%	0%
Dust	0%	96.8%	0%	3.2%	0%	0%
Smoke	14.9%	1.9%	75.6%	0%	0%	7.6%
Land	0%	0%	0%	100.0%	0%	0%
Water	0%	0%	0%	0%	100.0%	0%
Cloud	0%	0%	0%	0%	0%	100.0%

Figure 5: The two confusion matrices shown above describe the overall performance of the SVM aerosol detector and classifier. Correct classifications fall along the diagonal, while incorrect classifications lie off the diagonal.

The captions above each of the confusion matrices in Fig. 5 show that increasing the number of input cameras to the SVM aerosol classifier results in an improvement of nearly a factor of two in the error rate. Inspection of the diagonal, highlighted in green, shows that in all but one case the addition of angular information results in improved classification accuracy. The category labeled "pollution" is a generic category intended to represent all aerosol pixels not clearly distinguishable into one of the other two classes, smoke and dust. The SVM with the larger number of inputs attempts to classify pixels labeled as "pollution" into the category "smoke" 52.4% of the time, which accounts for the decreased accuracy in this one category. Improvements in other areas of the matrix more than compensate for this shortcoming, however. The confusion matrices highlight the usefulness of the multiangle approach to aerosol classification.

5. DISCUSSION

This work has shown that our approach to machine learning classification can be very effective for many global image pixel classification problems. The first key to our approach is the use of large feature vectors containing many raw features from different camera angles and spectral bands throughout a large spatial neighborhood, as opposed to a smaller number of higher-level features. It is also important to gather a large number of training examples, either by leveraging existing classifiers, or using an interactive tool to quickly label many images. Finally, we chose to use a robust machine learning technique; specifically support vector machines, to train a classifier that generalizes well. This technique is effective for both binary classifiers – such as the global cloud and cirrus detectors – and multi-class classifiers – such as the aerosol-type classifier.

From a scientific standpoint, the work with the SVM classifiers has been illuminating. Realizing that none of the existing MISR cloud classifiers was truly able to simultaneously take advantage of MISR's multiangle capabilities and the time-tested threshold-based approach to cloud detection in satellite imagery, we discovered that a machine learning algorithm was able to extract the best aspects of both approaches, yielding promising global cloud classifier. Although the effects of parallax have not yet been accurately accounted for, the work on the cirrus cloud detector provided useful insights into the problem of image registration for MISR's large view angle cameras and showed the utility of these cameras due to the increase in photon pathlength through thin clouds. Additionally, the aerosol detector provided some surprising information about the ability of a multiangle imager, making use of a limited number of shortwave spectral bands, to discriminate aerosols from clouds and different aerosol types from one another. The interactive SVM trainer also showed the importance of providing accurate positive training examples and the necessity of negative training examples to eliminate false positive identifications.

6. FUTURE WORK

Work currently underway at JPL involves developing a cloud classifier that is able to separate various common cloud types in MISR imagery. This task is particularly difficult because most cloud classes are based on large regional features, while our classifiers so far have been based on relatively local features.

ACKNOWLEDGEMENTS

REFERENCES

Bankert, R. L., 1994: Cloud classification of AVHRR imagery in maritime regions using a probabilistic neural network. *J. Appl. Meteor*., **33**, 909-918.

Bankert, R. L., and D. W. Aha, 1996: Improvement to a neural network cloud classifier. *J. Appl. Meteor*., **35**, 2036-2039.

Cortez, C., and V. Vapnik, 1995: Support vector networks. *Machine Learning*, **20**, 273-279.

DeCoste, D., and B. Schölkopf, 2002: Training invariant support vector machines. *Machine Learning*, **46**, 161-190.

Di Girolamo, L., and R. Davies, 1994: A band-differenced angular signature technique for cirrus cloud detection. *IEEE Trans. Geosci. Remote Sens*., **32**, 890-896.

Diner, D. J., and coauthors, 1998: Multiangle Imaging Spectroradiometer (MISR) instrument description and experiment overview. *IEEE Trans. Geosci. Remote Sens*., **36**, 1072-1087.

Diner, D. J., R. Davies, L. Di Girolamo, A. Horvath, C. Moroney, J.-P. Muller, S. R. Paradise, D. Wenkert, J. Zong, 1999a: MISR Level 2 cloud detection and classification Algorithm Theoretical Basis. JPL Internal D-11399 Rev. D., Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 102 pp. [Available online at: http://eospso.gsfc.nasa.gov/eos_homepage/for_scientists/atbd].

Diner, D. J., L. Di Girolamo, and E. E. Clothiaux, 1999b: MISR Level 1 cloud detection and classification Algorithm Theoretical Basis. JPL Internal D-13397 Rev. B, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 38 pp. [Available online at: http://eospso.gsfc.nasa.gov/eos_homepage/for_scientists/atbd].

Ebert, E., 1987: A pattern recognition technique for distinguishing surface and cloud types in the polar regions. *J. Climate Appl. Meteor*., **26**, 1412-1427.

Ebert, E., 1989: Analysis of polar clouds from satellite imagery using pattern recognition and a statistical cloud analysis scheme. *J. Appl. Meteor*., **28**, 382-399.

Garand, L., 1988: Automated recognition of oceanic cloud patterns. Part I: Methodology and application to cloud climatology. *J. Climate*, **1**, 20-39.

Key, J., J. A. Maslanik, and A. J. Schweiger, 1989: Classification of merged AVHRR and SMMR arctic data with neural networks. *Photogramm. Eng. Remote Sens*., **55**, 1331-1338.

Lee, J. R. C. Weger, S. K. Sengupta, and R. M. Welch, 1990: A neural network approach to cloud classification. *IEEE Trans. Geosci. Remote Sens*., **28**, 846-855.

Miller, S. W., and W. J. Emery, 1997: An automated neural network cloud classifier for use over land and ocean surfaces. *J. Appl. Meteor*., **36**, 1346-1362.

Moroney, C., R. Davies, and J.-P. Muller, 2002: Operational retrieval of cloud-top heights using MISR data. *IEEE Trans. Geosci. Remote Sens*., **40**, 1532-1540.

Peak, J. E., and P. M. Tag, 1992: Toward automated interpretation of satellite imagery for Navy shipboard applications. *Bull. Amer. Meteor. Soc*., **73**, 995-1008.

Peak, J. E., and P. M. Tag, 1994: Segmentation of satellite imagery using hierarchical thresholding and neural networks. *J. Appl. Meteor*., **33**, 605-616.

Schölkopf, B., and A. Smola, 2002: *Learning with Kernels*. MIT Press, Cambridge, MA, 610 pp.

Shenk, W. E., R. J. Holub, and R. A. Neff, 1976: A multispectral cloud type identification method developed for tropical ocean areas with Nimbus-3 MRIR measurements. *Mon. Wea. Rev*., **104**, 284-291.

Tian, B., M. Shaikh, M. R. Azimi-Sadjadi, T. H. Vonder Haar, and D. Reinke, 1999: A study of cloud classification with neural networks using spectral and textural features. *IEEE Trans. Neural Networks*, **10**, 138-151.
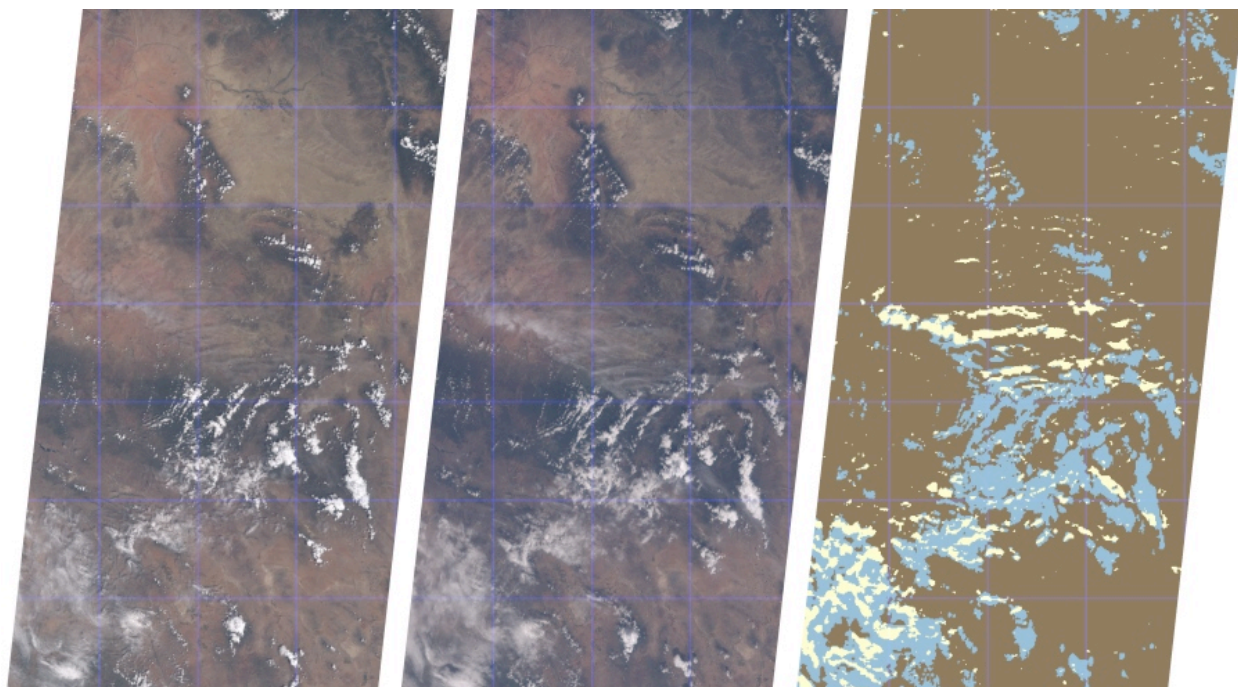
Figure 4: This figure shows two views of a MISR scene containing high, thin cirrus and the result of applying the SVM cirrus cloud detector. The left image shows the nadir, An (0°) camera RGB view. The center image is the forward-looking Cf (+60°) RGB view. Note in this image the increased brightness of the cloud slightly to the right of center due to the increased photon pathlength through the cloud at this larger angle relative to the An camera on the left. The right image shows the results of the SVM cirrus cloud detector. The key is shown below.

KEY	
Cirrus Clouds	
Other Clouds	
Land	