An operational pixel classifier for the Multi-angle Imaging SpectroRadiometer (MISR) using Support Vector Machines

Dominic Mazzoni, Michael Garay, Roger Davies
Dominic.Mazzoni@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109

Abstract

The Multi-angle Imaging SpectroRadiometer (MISR) data products now include a scene classification for each 1.1-km pixel that was developed using Support Vector Machines (SVMs), a cutting-edge machine learning technique for supervised classification. Using a combination of spectral, angular, and texture features, each pixel is classified as land, water, cloud, aerosol, or snow/ice. The classifiers were trained by MISR scientists who labeled hundreds of scenes using a custom interactive tool that showed them the results of the training in real time, making the process go significantly faster. Preliminary validation shows that the accuracy of the classifier is approximately 81% globally at the 1.1-km pixel level. Applications of this classifier include global studies of cloud and aerosol distribution, as well as data mining applications such as searching for smoke plumes. We believe this is the largest and most ambitious operational use of machine learning techniques for a remote-sensing instrument to date, and the success of this system will hopefully lead to further use of this approach.

Overview: The Multi-angle Imaging Spectoradiometer (MISR) (Diner et al., 1998) provides a unique view of the Earth by capturing images from nine cameras at fixed angles ranging from nadir to $70.5°$ in both the forward and aft directions. MISR exploits the multiple angular views of each scene to determine the height of clouds using stereoscopic pattern matching, and to compute aerosol microphysical properties, among other things. While MISR data products include masks that distinguish between cloudy and clear pixels, no data product had yet attempted to make distinctions between clouds and aerosols, nor between clouds and highly reflective snow or ice. Several years ago we began a project to investigate the use of machine learning techniques to improve the accuracy of existing MISR classifiers and add new classifiers, and this has resulted in the operational system we describe in this paper.

In the paradigm of supervised classification, items (in this case, MISR pixels) are classified into one of several categories using a model that is trained on many labeled examples provided by an expert. Methods for supervised classification include maximum likelihood, k-nearest-neighbor, decision trees, artificial neural networks, and support vector machines (SVMs). We focused on SVMs, one of the newest and most promising learning techniques[1]. SVMs are specifically constructed to minimize a statistical bound on the generalization error, resulting in models that extrapolate to new examples quite well and avoid overfitting. SVM training is also deterministic, not randomized like training a neural network. One downside of SVMs is that they can be slow to evaluate new examples, and this is a challenge we needed to overcome in this task.

[1]Modern SVMs were introduced by Cortes & Vapnik (1995); interested readers are referred to tutorials such as Burges (1998) or books such as Schölkopf & Smola (2002) for a thorough treatment of the subject.

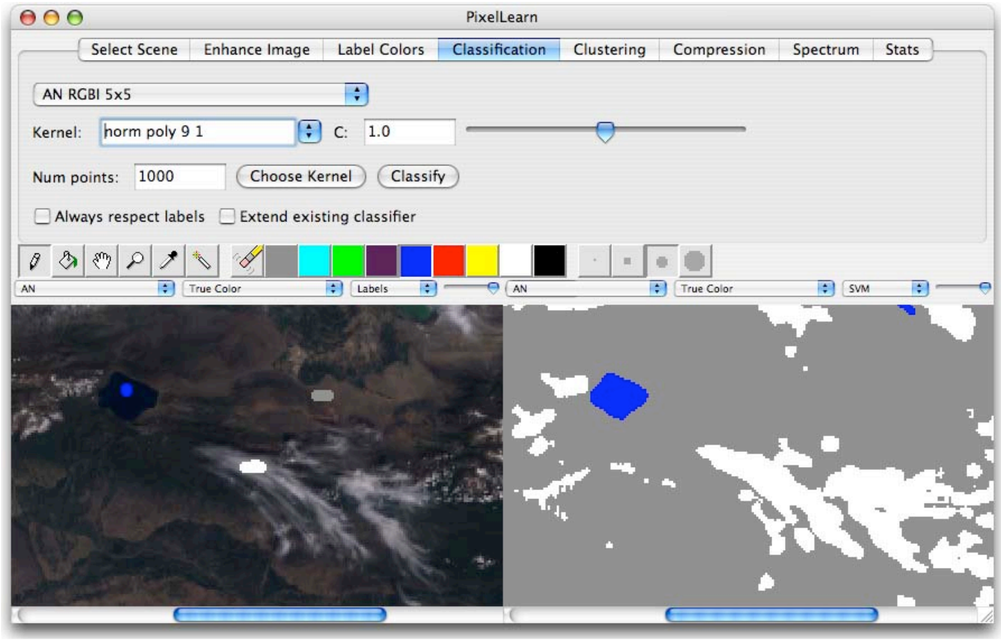

Figure 1: The PixelLearn tool we developed allowed scientists to explore MISR data and interactively label the pixels. In the example above, the user has labeled three regions containing water, clouds, and land in the lower-left panel, and a support vector machine classification of the image appears in the lower-right image. For advanced users, the controls in the top panel control all aspects of SVM training.

Related Work: While many others have applied various machine learning methods, including support vector machines, to classify pixels from remote sensing instruments (e.g., Welch et al. (1992), Bankert (1994b), Bankert (1994a), Bankert & Aha (1995), Tian et al. (2000), Saitwal et al. (2003), Baum et al. (1997), Azimi-Sadjadi & Zekavat (2000), Li et al. (2003), Lee et al. (2004)), very few have resulted in an application in an operational environment. One successful example is Bankert's real-time classifier for GOES images available on the web[2], which uses a nearest-neighbor classification algorithm and a large collection of training data. More details about his approach are described in their paper on Advanced Very High Resolution Radiometer (AVHRR) cloud-type classification (Tag et al., 2000). An SVM-based classifier is currently running onboard NASA's EO-1 spacecraft, as part of the Autonomous Sciencecraft Experiment (Mazzoni et al., 2005). When images of certain targets are acquired by the Hyperion instrument, the classifier applies a trained SVM to 12 of its 242 spectral bands to classify each pixel as land, water, cloud, ice, or snow. Simple heuristics are then used to determine whether or not the image should be saved for downlink or discarded to make room for more interesting images.

Methodology: An SVM, like any other supervised classification algorithm, learns to classify each example given a feature vector of numbers that describe the example. In order to exploit texture

[2] `http://www.nrlmry.navy.mil/sat-bin/clouds.cgi`

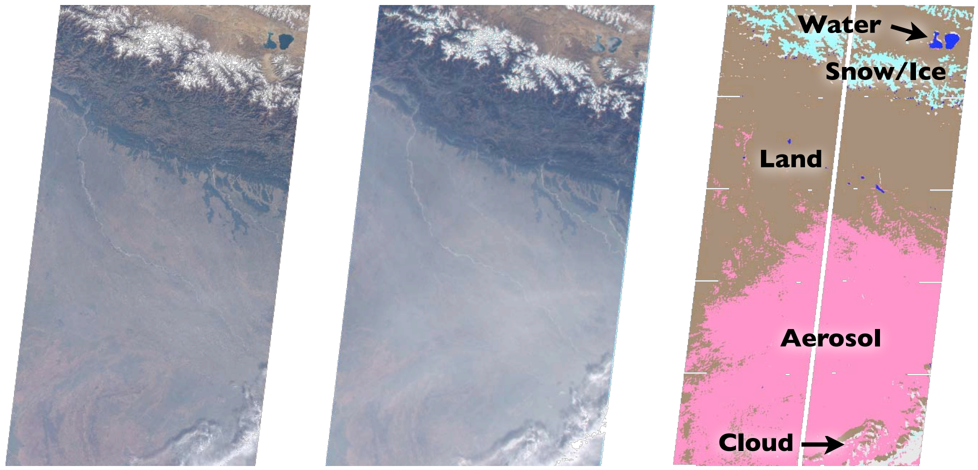

Figure 2: An example of the operational MISR pixel classifier in action. The left image is from MISR's nadir-pointing camera, and the middle image is from MISR's 45-degree forward-pointing camera. The image on the right shows a representation of the SVM classifier's output. The SVM classifier has correctly identified the snow on the mountains, a lake, pollution aerosols in a valley, and clouds in the lower part of the scene.

information, we build our feature vectors using radiances from every pixel in a 5×5 neighborhood surrounding the pixel we wished to classify. (Surprisingly, we found it was more computationally efficient and just as effective to simply include all 25 values in the feature vector, rather than computing texture coefficients and constructing a smaller feature vector.) We used radiance values from four of MISR's spectral bands and five of its nine cameras for each pixel. (Using all nine cameras did not help; empirically, five was optimal for this classification.) We also included the solar zenith angle in the feature vector – our final total was 154 features per pixel.

It is very common for the process of collecting training data to be totally complete before any machine learning algorithms are run. However, we were very concerned that at best, this is an inefficient use of time, and at worst, it can lead to poor accuracy. Instead we developed a custom tool, called *PixelLearn* (Figure 1), for interactively labeling multispectral, multiangle datasets. PixelLearn includes a built-in fast SVM training algorithm, and it constantly updates the classification on one half of the screen while the user adds more labels on the other half.

At the end of the interactive phase, we had obtained hundreds of thousands of labeled pixels spanning hundreds of scenes. We then trained a single best SVM classifier based on all of these examples (which took days of computation time). This SVM was rather complex, and it was about an order of magnitude too slow to run as part of the MISR data products system. To address this, we developed an algorithm for constructing a smaller number of so-called "reduced-set vectors" to approximate the support vectors from the original SVMs, based on an approach suggested by Burges (1996). The end result of this process was an SVM that was 20 times faster to evaluate, with a loss of accuracy of about 0.5%.

Operational Implementation: We integrated the SVM classifier into the operational data processing system at the Langley Atmospheric Science Data Center, as part of release 4.0 of the MISR processing software, which started processing new data on December 1, 2005. The code runs as part of the level 2 processing and generates several new fields as part of the top-of-atmosphere/cloud classifiers data file. See Figure 2 for an example of the classifications it generates, and for more information, see the MISR Data Products Specifications (Bull et al.).

We evaluated the final classifier by having two human experts independently label millions of pixels, and then comparing the results to those given by the SVM classifier. In addition, two human experts each labeled some of the same scenes independently, making it possible to compare the error rates of the SVM with the degree of disagreement between humans. We validated on four complete MISR orbits, chosen so that they were spread out reasonably in both space and time. Overall, the SVM's five-class classification decision agreed with the human expert 80.9% of the time, at the 1.1-km pixel level. When each block of 16×16 pixels was replaced by the majority class, the agreement improved to 84.9%, indicating that a small but significant fraction of the errors were isolated blunders and not gross misclassifications. The level of agreement between two humans, at 93.0% and 96.3% for the 1.1-km and 17.6-km resolutions, respectively, quantify the degree of subjectivity present in these experiments. These serve as an upper bound for the maximum performance once could expect from any classifier.

Conclusions and Future Work: The project described here is the most ambitious and large-scale application of machine learning technology to an operational remote-sensing application to date. We found that using an interactive application is an effective way to make efficient use of expert scientists' time in collecting training labels. We showed that using reduced-set techinques, it is possible to create an SVM with limited computational requirements.

We have begun two new studies that are directly benefitting from this work. In one task, we are searching through years of archived data to collect statistics on the distribution and heights of aerosol plumes. These statistics will be used to validate and initialize transport models to better understand the effect of certain types of aerosols (e.g. smoke from boreal forest fires) on air quality. The SVM classifier is an important part of our technique for searching for these plumes in terabytes of data and outlining them for automatic feature extraction, a task that is too time-consuming for humans to do manually. In the other task, we are investigating the use of the SVM classifier to help discriminate between clouds and aerosols in order to determine when to run the MISR algorithm to retrieve the aerosol optical depth and microphysical properties.

Acknowledgements: The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Thanks to MISR principal investigator Dave Diner for his support of this work. Other investigators on this project included Dennis DeCoste, Kiri Wagstaff, Amy Braverman, and Benyang Tang, and the operational code was developed and integrated by David Nelson.

References

Azimi-Sadjadi, M. R. & Zekavat, S. A. (2000). Cloud classification using support vector machines. In *Proceedings of the 2000 IEEE Geoscience and Remote Sensing Symposium (IGRASS 2000)*, volume 2, (pp. 669–671)., Honolulu, HI.

Bankert, R. (2006). Naval research laboratory monterey GOES cloud classification (website). `http://www.nrlmry.navy.mil/sat-bin/clouds.cgi`.

Bankert, R. L. (1994a). Cloud classification of AVHRR imagery in maritime regions using a probabilistic neural network. *Journal of Applied Meteorology*, *33*, 909–918.

Bankert, R. L. (1994b). Cloud pattern identification as part of an automated image analysis. In *Preprints, 7th American Meteorological Society Conference on Satellite Meteorology and Oceanography*, (pp. 441–443)., Boston, MA.

Bankert, R. L. & Aha, D. W. (1995). Automated identification of cloud patterns in satellite imagery. In *Preprints, 14th Conference on Weather Analysis and Forecasting, American Meteorological Society*, (pp. 313–316)., Dallas, TX.

Baum, B. A., Tovinkere, V., Titlow, J., & Welch, R. (1997). Automated cloud classification of global AVHRR data using a fuzzy logic approach. *Journal of Applied Meteorology*, *36*, 1519–1540.

Bull, M., Matthews, J., Moroney, C., & Smyth, M. MISR data products specifications. `http://eosweb.larc.nasa.gov/PRODOCS/misr/DPS/`.

Burges, C. J. C. (1996). Simplified support vector decision rules. In *Proceedings of the Thirteenth International Conference on Machine Learning*, (pp. 71–77).

Burges, C. J. C. (1998). A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, *2*(2), 121–167.

Cortes, C. & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*, 273–297.

Diner, D., Beckert, J., Reilly, T., Bruegge, C., Conel, J., Kahn, R., Martonchik, J., Ackerman, T., Gordon, H., Muller, J.-P., Myneni, R., Sellers, R., Pinty, B., & Verstraete, M. (1998). Multiangle imaging spectroradiometer (MISR) instrument description and experiment overview. *IEEE Trans. Geosci. Remote Sens.*, *36*, 1072–1087.

Lee, Y., Wahba, G., & Ackerman, S. (2004). Cloud classification of satellite radiance data by multicategory support vector machines. *Journal of Atmospheric and Oceanic Technology*, *21*(2), 159–169.

Li, J., Menzel, W. P., Yang, Z., Frey, R. A., & Ackerman, S. A. (2003). High-spatial-resolution surface and cloud-type classification from MODIS multispectral band measurements. *Journal of Applied Meteorology*, *42*, 204–226.

Mazzoni, D., Tang, N., Doggett, T., Chien, S., Greeley, R., & Cichy, B. (2005). Learning classifiers for science event detection in remote sensing imagery. In *Proceedings of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 2005)*.

Saitwal, K., Azimi-Sadjadi, M. R., & Reinke, D. (2003). A multichannel temporally adaptive system for continuous cloud classification from satellite imagery. *IEEE Transactions on Geoscience and Remote Sensing*, *41*(5), 1098 – 1104.

Schölkopf, B. & Smola, A. J. (2002). *Learning with Kernels*. Cambridge, MA: The MIT Press.

Tag, P. M., Bankert, R. L., & Brody, L. R. (2000). An AVHRR multiple cloud-type classification package. *Journal of Applied Meteorology*, *39*, 125 – 134.

Tian, B., Azimi-Sadjadi, M. R., Vonder Haar, T. H., & Reinke, D. (2000). Temporal updating scheme for probablistic neural network with application to satellite cloud classification. *IEEE Transactions on Neural Networks*, *11*(4), 903–920.

Welch, R., Sengupta, S., Goroch, A., Rabindra, P., Rangaraj, N., & Navar, M. (1992). Polar cloud and surface classification using AVHRR imagery: An intercomparison of methods. *Journal of Applied Meteorology*, *31*(5), 405–420.