AN ALGORITHM FOR DE-SHADOWING SPECTRAL IMAGERY

Steven M. Adler-Golden[1], Michael W. Matthew[1], Gail P. Anderson[2], Gerald W. Felde[2], and James A. Gardner[3]

1. INTRODUCTION

The interpretation of visible-near infrared spectral imagery of the earth's surface can be complicated by illumination variations due to shadowing, sloped terrain, or other causes. In atmospherically "corrected" (spectral reflectance) data, these variations mainly affect the spectrum amplitude. However, in deep shadows where the illumination is mainly skylight, the spectrum is not only dimmer but also skewed to short (blue-violet) wavelengths. Either or both of these effects can impede the classification of surface materials and the detection of targets with standard methods. Recently, some classification and detection algorithms have been developed that are insensitive to illumination (Healy and Slater 1999; Adler-Golden et al., 2001). However, for many applications, and for visual inspection of the image, it is very beneficial to normalize all the pixels to a common illumination, such as full sun and skylight. This requires the ability to characterize the illumination for each pixel, and in particular to identify shadows and quantify their depths. Knowledge of shadow depths can also be valuable as a source of elevation information for terrain, surface objects, and clouds.

Several methods have appeared in the literature for identifying and correcting for shadowing. The standard approach is based on linear unmixing of atmospherically corrected data (Boardman, 1993) using spectral endmembers, which represent the distinguishable pure materials in the scene. The shadow is defined as a "black" (zero reflectance) endmember, and a sum-to-unity constraint is imposed on the endmember weights. A positivity constraint may also be imposed. Upon unmixing, the reflectance spectrum is approximately "de-shadowed" by dividing by one minus the shadow weight (or, equivalently, the sum of the non-shadow endmember weights), provided that this quantity is not too small. There are several limitations to this approach. One is that by defining the shadow as black, the blue-skewed skylight illumination of the shadow is ignored, leading to under-correction for shadow at long wavelengths and overcorrection at short wavelengths. Another is the dependence of the results on the specific endmembers used, and hence on the method of endmember selection (e.g., by clustering, convex hull, convex cone, or other methods).

Recently, more sophisticated shadow removal methods have been developed as part of atmospheric correction processing. Davenport and Ressl (1999) obtained excellent results with visible channels of the HYDICE sensor using an algorithm that models the illumination spectrum as a function of shadow depth, which is extracted using an endmember-finding and fitting method. Another method, developed by Portigal (2002), avoids endmember complications by determining the shadow fraction with a matched filter. The filter is constructed using a shadowed pixel in the radiance image; the filter output to adjust the amount of direct solar illumination in the atmospheric correction while maintaining full sky illumination. Potential pitfalls include sensitivity to the choice of shadow spectrum and limitations of the full-sky assumption.

In this paper we describe a rather simple matched filter-based de-shadowing algorithm for spectral reflectance images that builds upon several ideas from the previous work. At present the algorithm is partially supervised, in that the best results are achieved with manual adjustment of a parameter that controls the sky brightness. The de-shadowing method can also be applied to the general illumination-compensation problem, including terrain "flattening," a topic that will be reserved for a future paper. The algorithm uses iterative application of the matched filter, an approach that may be suitable for a broader class of problems that involves a non-linear effect (in this case, the skylight blue-skewing).

All of the de-shadowing approaches mentioned presume that differences caused by illumination variation

---

1 Spectral Sciences, Inc., 99 South Bedford Street, Suite 7, Burlington, MA 01803 (sag@spectral.com)
2 AFRL/Space Vehicles Directorate, 29 Randolph Street, Hanscom AFB, MA 3010
3 University of Arizona, Tuscon, AZ 85721

are spectrally distinguishable from differences in material composition. In the black shadow approximation, this amounts to an assumption that all pure spectra in the image are linearly independent. As this is often untrue, particularly in complex scenes, many instances of mis-correction (i.e., when a fully illuminated pixel is assigned as a shadowed pixel of another material or combination of materials) are to be expected. Nonetheless, we have found that useful and, for the most part, visually realistic de-shadowed images can be obtained with this method under a wide variety of conditions.

2. DESCRIPTION OF METHOD

The method consists of (1) atmospherically correcting the original radiance data to reflectance with the usual assumption of full sun and sky illumination, (2) applying a matched filter for a black surface to the reflectance data to infer an initial shadow fraction for each pixel, (3) rebalancing the data to remove the skylight skewing effect, (4) refining the shadow fraction by reapplying the matched filter, and (5) using the shadow fraction, de-shadowing the data by dividing by the ratio of the inferred to full illumination. The procedure was implemented in IDL code using the Research Systems, Inc. ENVI software.

Atmospheric Correction. A number of atmospheric correction algorithms have been developed. For de-shadowing, especially accurate results are needed for dark pixels. We have employed the FLAASH algorithm (Matthew *et al*., 2000, Adler-Golden *et al*., 1999), which uses MODTRAN4 (Berk *et al*., 1998) for radiation transport modeling and includes an adjacency effect correction and an automated aerosol retrieval method based on the work of Kaufman *et al*. (1997).

Matched Filtering.   A matched filter for a zero reflectance target is generated and applied to the reflectance image. In the following discussion we provide some mathematical motivation for the matched filtering step and qualitatively compare the matched filter and endmember unmixing approaches.

As described in the Introduction, each pixel reflectance spectrum y is assumed to be a mixture of linearly independent, fully-illuminated endmember spectra, e_i:

$$y = \sum_i w_{iy} e_i \tag{1}$$

where w_{iy} is the ith endmember weight for the pixel (here bold face denotes a column vector). Let a be a fully illuminated pixel, whose weights w_{ia} sum to 1. If shadow is considered as a zero reflectance endmember, then the shadow fraction f can be written as the difference between the total weights for pixel a and pixel y:

$$f_y = \sum_i w_{ia} - \sum_i w_{iy} = -\sum_i (w_{iy} - w_{ia}) \tag{2}$$

In unconstrained least-squares unmixing, the ith endmember weight may be extracted from each spectrum by taking its dot product with a filter vector v_i that is orthogonal to all endmembers except e_i; that is, $w_{iy} = v_i^T y$ and $w_{ia} = v_i^T a$. Eq. (2) can then be written as

$$f_y = -\sum_i v_i^T (y - a) = g^T (y - a) \tag{3}$$

where g is a filter vector. If we furthermore assume that shadow is rare, the scene mean spectrum can be used for a. Thus, according to Eq. (3) the shadow fraction can be unmixed by applying a linear filter to mean-subtracted spectra. One such linear filter is the matched filter. Instead of being constructed from explicit endmembers, the matched filter q is defined in terms of the covariance matrix C:

$$q = C^{-1}(t - a)/[(t - a)^T C^{-1}(t - a)] \tag{4}$$

where t is the target spectrum to be detected and a is the scene mean spectrum.   Taking $t = 0$ for shadow, the shadow matched filter is given by

$$q_{shadow} = - C^{-1} a / (a^T C^{-1} a) \tag{5}$$

The matched filter, like the CEM filter (Farrand and Harsanyi, 1997), yields a minimum RMS target abundance for

the scene as a whole.   However, instead of using the covariance and mean-subtracted reflectance signal, the CEM filter uses the correlation matrix and the reflectance signal. This results in a singularity when $t = 0$ and makes the CEM filter unusable for shadow detection

In the limit of zero shadow abundance and linear independence of all materials in the scene, the outputs of both the matched filter and the endmember-based filter g should be zero, and the two filters should be equivalent. In real scenes the outputs differ. In particular, the matched filter generates negative as well as positive shadow outputs, whereas the endmember filter, if constructed from outlying endmembers (such as from a convex-hull or convex-cone algorithm), should in principle generate only non-negative outputs. A negative output may be interpreted as indicating either an anomalously bright material or excess illumination in the pixel compared with the nominal full sun plus sky illumination of a horizontal surface. With flat terrain, it would generally not be appropriate to "correct" the illumination for a negative shadow output. However, with sloped terrain it may be desirable to do so in order to remove the excess illumination of surfaces tilted towards the sun.

As the amount of shadow in the scene increases, the behavior of the matched filter increasingly deviates from what would be obtained with linear unmixing. As the mean spectrum becomes contaminated with shadow, the filter's response to partial shadow is reduced and the negative outputs become larger and more numerous. Very dark materials that can be confused with shadow, such as bodies of water, produce the same contamination effect. A remedy is to remove as much water and deep shadow from the scene as is practical before constructing the matched filter. A simple method is to set a low threshold for the average spectral reflectance, such as 0.03, and to eliminate all pixels falling below this threshold when calculating the mean and covariance.

Spectral Rebalancing. Application of the matched filter to the image yields a first estimate of the shadow fraction f for each pixel. The spectra are then rebalanced to simulate illumination by a spectrally uniform source, namely, f times the direct sun plus sky spectrum, as follows. Let $d = d(\lambda)$ be the spectrum of the direct sun illumination and $s = s(\lambda)$ the spectrum of the sky illumination, and define f as the fractional illumination of each pixel by the direct sun. Then the rebalancing factor applied to the spectrum is

$$F(\lambda) = f(d(\lambda) + s(\lambda))/(fd(\lambda) + s(\lambda)) = f(1 + s(\lambda)/d(\lambda))/(f + s(\lambda)/d(\lambda)) \qquad (6)$$

The spectral ratio $s(\lambda)/d(\lambda)$ is, strictly speaking, a quantity that depends on the detailed three-dimensional geometry of the terrain and surface objects. However, we assume that it can be adequately modeled with the full sky to sun illumination ratio as calculated from MODTRAN (Adler-Golden et al., 2001) or some other radiation transport code, multiplied by an empirical scaling factor that accounts for the average sky fractional subtense in the shadows. The ratio can often be reasonably approximated by the power law expression $c\lambda^{-N}$, where c is an empirical constant and N is an Angstrom-law exponent for the atmospheric scattering.

Re-application of the Matched Filter. The matched filter is applied to the rebalanced spectra, resulting in refined shadow fraction estimates. The rebalancing and filter application steps may be iterated a few times to insure convergence.

De-shadowing.   The de-shadowed spectrum for each pixel is obtained by multiplying the rebalanced spectrum by the factor $1/(1 - f)$.

3. SAMPLE RESULTS

We have applied the de-shadowing method to a total of around a half-dozen images to date, including AVIRIS, HYDICE and Landsat TM data. In this discussion we focus on two atmospherically corrected scenes: a forested scene taken by AVIRIS that contains a broken cloud field and an urban scene taken by HyMap that contains building shadows.

3.1 AVIRIS Scene

The AVIRIS scene, f960628t01p02_r04_sc07, is of the White Mountains region near North Conway, NH. Good removal of the cloud shadows was obtained using either the original spectral data or the data binned to the six

Landsat TM spectral channels. For simplicity we focus here on the Landsat TM representation, which allows us to show the complete data cube in two color images: a true-color image of the blue, green and red channels 1 to 3 and a false-color image of the infrared channels 4 to 6. The scene was atmospherically corrected using FLAASH. Little difference was found in performing the atmospheric correction or de-shadowing steps before or after the spectral binning.

Figure 1 shows the reflectance image prior to de-shadowing. To emphasize the darker pixels, a non-linear scaling has been applied to these and other color images in this paper, which sets the mean brightness level at 50%. Figure 2 shows the retrieved shadow fraction as a grayscale image; black denotes 100% shadow and white denotes zero shadow. Figure 3 shows the final de-shadowed image using $c=0.07$ and $N=2$ in the power-law expression for the full sky to sun illumination ratio; $N=2$ represents a compromise between the wavelength dependences of Rayleigh and aerosol scattering. Despite some residuals, appearing as a slight darkening in the infrared and a slight grayish cast in the visible, the cloud shadow removal is quite effective. Contrast is slightly enhanced compared to Figure 1 due to the non-linear display. The most serious artifacts are an artificial brightening of water bodies (streams and a small pond) and paved roads. Details of the scene shown in Figure 4 highlight the dramatically improved visualization of surface features upon de-shadowing.


Figure 1. AVIRIS White Mountains scene displayed in Landsat TM channels; (blue, green, red) = visible channels 1,2,3 at left, infrared channels 4,5,6 at right.


Figure 2. Retrieved shadow image for the White Mountains scene.
The effect of ignoring the sky illumination of shadows (by setting $c=0$) is seen in Figure 5. In the visible the

illumination is underestimated and the image is noticeably over-corrected. In the infrared the illumination is slightly overestimated and the channels are slightly under-corrected.

Convergence of the algorithm was investigated by comparing the results for one and two iterations of the spectral rebalancing and matched filtering steps. Differences of a few percent in the de-shadowed spectra were found (Figure 6); this effect is too small to be noticed in the images, however.


Figure 3. White Mountains scene after de-shadowing, displayed as in Fig. 1.


Figure 4. Details from Figures 1 and 3.

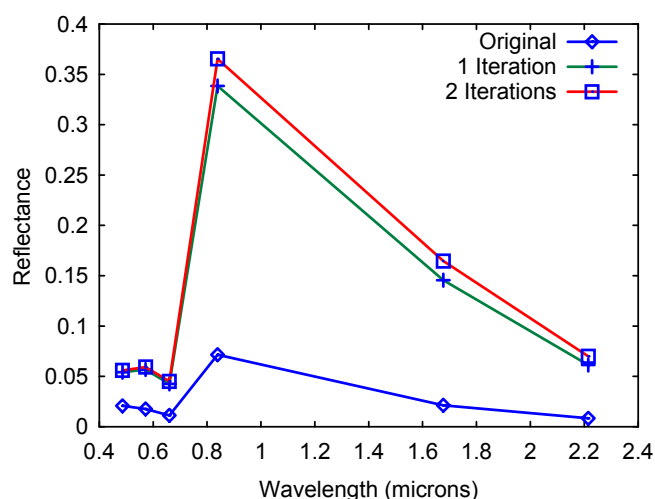Figure 5. De-shadowed image ignoring the sky illumination.

Figure 6. Shadowed vegetation pixel in the White Mountains scene before (original) and after de-shadowing (1 and 2 iterations).

3.2. HyMap Scene

An atmospherically corrected HyMap image of downtown San Francisco from the archives of Analytical Imaging and Geophysics LLC was kindly provided to us by the Boeing Company. Special handling was needed to process this data cube because of what appears to be a sizable baseline of residual atmospheric scattering. To minimize its effect on the shadow detection, the matched filter was restricted to wavelengths from 1.4-2.5 μm. While we did not attempt to create a full de-shadowed data cube, we were able to generate a good de-shadowed true-color image by baseline-subtracting the displayed red, green and blue channels, using their minimum values as the baseline, and by increasing the green and blue sky-to-sun spectral ratios relative to the N=2 predictions (this might be compensating for an especially clear sky, dominated by Rayleigh scattering).

The results from a portion of this scene are shown in Figure 7. The removal of the building shadows is quite striking and gives the scene a strangely flat appearance. The water is brightened, as usual, but otherwise there appear to be no major problems. There is a tendency for under-correction in the deepest shadows, next to and between tall buildings, where there would be less than the usual amount of skylight. Even so, the shadow removal is sufficient to discern a wealth of new detail.

Figure 7. HyMap scene of San Francisco before (left) and after (right) de-shadowing.

4. CONCLUSIONS

A simple de-shadowing method based on iterative matched filtering has been developed for atmospherically corrected spectral imagery. In initial tests with a handful of data cubes, generally realistic and quantitative de-shadowed images and spectra have been obtained for non-water surfaces. The method has the basic shortcomings inherent in a purely spectrally-based approach, which does not take advantage of spatial context or ancillary information that might distinguish between shadowed and fully lit surfaces that are spectrally similar. In addition, the assumption that all of the shadows are illuminated by a uniform skylight spectrum is not very realistic, and can cause problems in very deep shadows as well as in shadows that are lit by sources other than the sky, such as a transmissive plant canopy. Nevertheless, this simple model for shadow illumination provides a dramatic improvement over ignoring the illumination completely, as is done with standard linear unmixing. Although the initial applications have involved removal of cast shadows, the de-shadowing method also has potential for terrain "flattening," which we hope to explore in the future. Further work is needed to test the algorithm on additional data and against other algorithms such as those of Davenport and Ressl (1999) and Portigal (2002).

5. ACKNOWLEDGEMENTS

6. REFERENCES

Adler-Golden, S.M., R.Y. Levine, M.W. Matthew, S.C. Richtsmeier, L.S. Bernstein, J. Gruninger, G. Felde, M. Hoke, G.P. Anderson, and A. Ratkowski, 2001, "Shadow-Insensitive Material Detection/Classification with Atmospherically Corrected Hyperspectral Imagery," Proc. of SPIE, Orlando, FL, Vol. 4381.

Adler-Golden, S.M., M.W. Matthew, L.S. Bernstein, R.Y. Levine, A. Berk, S.C. Richtsmeier, P.K. Acharya, G.P. Anderson, G. Felde, J. Gardner, M. Hoke, L.S. Jeong, B. Pukall, A. Ratkowski and H.K. Burke, 1999, "Atmospheric Correction for Short-wave Spectral Imagery Based on MODTRAN4," in Summaries of the Eighth Annual JPL Earth Science Workshop, Vol. I, available at http://makalu.jpl.nasa.gov.

Berk, A., L.S. Bernstein, G.P. Anderson, P.K. Acharya, D.C. Robertson, J.H. Chetwynd, and S.M. Adler-Golden, 1998, "MODTRAN Cloud and Multiple Scattering Upgrades with Application to AVIRIS,' Remote Sens. Environ., **65**, 367-375.

Boardman, J.W., 1993, "Automated Spectral Unmixing of AVIRIS Data Using Convex Geometry Concepts," in Summaries, Fourth JPL Airborne Geoscience Workshop, JPL Publication 93-26, v. 1, pp. 11-14.

Davenport, M. R. and W. Ressl, 1999, "Shadow Removal in Hyperspectral Imagery," Proceedings, International Symposium on Spectral Sensing Research, Las Vegas, October 1999.

Farrand, W.H. and J.C. Harsanyi, 1997, "Mapping the Distribution of Mine Tailings in the Coeur d'Alene River Valley, Idaho, Through the Use of a Constrained Energy Minimization Technique," Remote Sens. Environ., vol. 59, pp. 64-76.

Healy, G. and D. Slater, 1999, "Models and Methods for Automated Material Identification in Hyperspectral Imagery Acquired Under Unknown Illumination and Atmospheric Conditions," IEE Trans. Geosci. Remote Sensing, Vol. 37, pp. 2706-2717.

Kaufman, Y.J., D. Tanre, L.A. Remer, E.F. Vermote, A. Chu, and B.N. Holben, 1997, "Operational Remote Sensing of Tropospheric Aerosol over Land from EOS Moderate Imaging Spectroradiometer," J. Geophys. Res. **102**, 17051.

Kaufman, Y.J., A.E. Wald, L.A. Remer, B.-C. Gao, R.-R. Li, and L. Flynn, 1997, "The MODIS 2.1-μm Channel—Correlation with Visible Reflectance for Use in Remote Sensing of Aerosol," IEEE Transactions on Geoscience and Remote Sensing, **35**, pp. 1286-1298.

Matthew, M.W., S.M. Adler-Golden, A. Berk, S.C. Richtsmeier, R.Y. Levine, L.S. Bernstein, P.K. Acharya, G.P. Anderson, G.W. Felde, M.P. Hoke, A. Ratkowski, H.-H Burke, R.D. Kaiser, and D.P. Miller, 2000, "Status of Atmospheric Correction Using a MODTRAN4-based Algorithm," SPIE Proceeding, Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VI, Vol. 4049, pp. 199-207.

Portigal, F., 2002, "Adaptive Reflectance Calibration" method; see http://www.ultraspectral.com.