

Removing shadows from hyperspectral images leaves nowhere to hide

Edward Ashton, Brian Wemett, Robert Leathers, and Trijntje Downes

A novel method for illumination suppression using a two-stage correction process allows for robust image reconstruction.

Varying levels of illumination due to shadows and cloud cover, among others, are a known problem for many hyperspectral (i.e., not limited to visible light) segmentation and targeting algorithms. The primary advantage to hyperspectral imaging is that, because an entire spectrum is acquired at each point, the operator needs no a priori knowledge of the sample and post-processing allows all available information from the data set to be mined. Shadows do not simply lower the overall intensity of the radiance spectrum, they also cause a change in the spectral shape because the color of the light in a shadowed region is different from that in a sunny area. Our objective is to develop a practical method to identify shadows in a scene and hence appropriately adjust the illumination level and color to better match similar materials in non-shadowed regions.¹

The simplest approach to address this problem is by normalizing the magnitudes of all spectral vectors in a data cube. This does not, however, account for the nonuniform effects of shading on different regions of the spectrum. It also reduces by one the effective dimensionality of the data cloud without reducing the actual number of dimensions, thus causing severe problems for any algorithm reliant on covariance-matrix inversion. Alternatively, one can convert the data from Cartesian space to a hyperspherical coordinate system. In this approach, each N -dimensional spectral vector is converted to an $(N-1)$ -dimensional spectral angle, again reducing the dimensionality. This technique requires the conversion of target spectra to hyperspherical coordinates and also does not take into account any differential interference caused by shading.

N -dimensional hyperspherical transformation creates $N-1$ angles and one scalar magnitude, or radius, in the N th band.




Figure 1. Definition of polar coordinates: radius (r) and angle (θ).




Figure 2. Spectral signature (in arbitrary intensity units) of grass in shadow and in the open.

This magnitude is the equivalent of the spectral illumination value. Figure 1 shows the 2D equivalent, where we have one angle, θ , and one magnitude, r . Starting from this simple approach, we subsequently segment the hypersphere by application of the k -means algorithm. The resulting segmentation is largely dependent on the illumination band. A shadow map can therefore be created by selecting the class with the lowest magnitude values.

Continued on next page

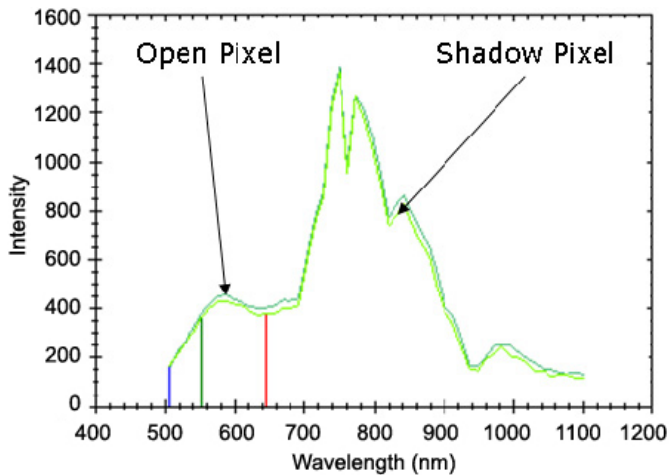


Figure 3. Spectra resulting from the conversion back to Cartesian coordinates after correction, in arbitrary intensity units.

The ratio of the resulting non-shadow-to-shadow mean vectors is used as a coarse correction to the shadow-area radiance. A weight mask is applied to correct the image along the shadow edges because a full correction is not needed. Finally, a fine correction is computed that depends on the type of material in the shadow area. Each shadow pixel is assigned to one of the non-shadow classes through a simple squared-error-distortion measure, which is possible because of the coarse correction step.

Mean vectors are calculated for each material type in shadow and non-shadow areas. Each shadow pixel is then corrected on the basis of its assigned class. This step is critical to ensure proper coloring and to correct for the differential interference inherent to shading. The data is now transformed back to Cartesian coordinates using the corrected pixels. Figure 2 shows the spectra of two very similar materials, one in shadow and a very different one in full sunlight. This highlights the problem of using spectral-based signature detection.

The signatures of the corrected spectra after conversion back to Cartesian coordinates are closer to each other, as demonstrated in Figure 3. In addition, we do not have the negative side effect of losing a dimension. Figures 4 and 5 display enlargements of the original and the illumination-suppressed image. In the latter, the tree shadows have been removed, with natural grass visible underneath. The entire parking lot is now visible where before it was partially in shadow. There are still some residual effects around the edges that need to be softened because the edge of the shadow region is smeared out.

In preliminary tests we found the new method to robustly correct spectra in shadow areas. The results are more convincing than for the two simpler methods used for comparison. The



Figure 4. Original image.



Figure 5. Illumination-suppressed image.

two-stage correction ensures that features in the shadows are corrected at least coarsely, even if a similar material does not exist in full illumination. Such features will not be corrected by the ‘fine-tuning’ second stage, and what the detriment in detection efficiency will be has yet to be determined. No method can completely remove the effects of shadows since the effective signal-to-noise ratio is degraded. Although in deep shade the data could be unusable, we believe that we can properly correct for the apparent change in color and reflectivity resulting from light shade. A comprehensive detection study is needed to verify

Continued on next page

the proper spectral correction. We also need to determine how to distinguish between shadows and dark objects. We are exploring improvements to the algorithm by including spatial processing.

Author Information

Edward Ashton and Brian Wemett

VirtualScopics, Inc.
Rochester, NY
<http://www.virtualscopics.com>

Robert Leathers and Trijntje Downes

Optical Sciences Division
US Naval Research Laboratory
Washington, DC

References

1. E. A. Ashton, B. D. Wemett, R. A. Leathers, and T. V. Downes, *A novel method for illumination suppression in hyperspectral images*, **Proc. SPIE 6966**, p. 69660C, 2008.