

Principal-Components-Based Display Strategy for Spectral Imagery

J. Scott Tyo, *Member, IEEE*, Athanasios Konsolakis, David I. Diersen, and Richard Christopher Olsen, *Member, IEEE*

Abstract—A new pseudocolor mapping strategy for use with spectral imagery is presented. This strategy is based on a principal components analysis of spectral data, and capitalizes on the similarities between three-color human vision and high-dimensional hyperspectral datasets. The mapping is closely related to three-dimensional versions of scatter plots that are commonly used in remote sensing to visualize the data cloud. The transformation results in final images where the color assigned to each pixel is solely determined by the position within the data cloud. Materials with similar spectral characteristics are presented in similar hues, and basic classification and clustering decisions can be made by the observer. Final images tend to have large regions of desaturated pixels that make the image more readily interpretable. The data cloud is shown here to be conical in nature, and materials with common spectral signatures radiate from the origin of the cone, which is not (in general) at the origin of the spectral data. A supervised method for locating the origin of the cone based on identification of clusters in the data is presented, and the effects of proper origin orientation are illustrated.

Index Terms—Hyperspectral imagery, multidimensional image display, spectral imagery.

I. INTRODUCTION

RECENT YEARS have seen the development of spectral imagers that measure spectra at more than 200 wavelengths at every pixel in a scene [1], [2]. Many powerful processes have been developed to extract scene information from spectral images, and once classification algorithms or image processing tools have been applied, the results are almost invariably presented through some pseudocolor image. While many display methodologies are quite powerful, there is no standard method used to render spectral imagery in false-color images for presentation.

Currently, the use of false color displays is generally reserved as a tool for presenting data after processing. Once a scene has been classified by a particular algorithm, a specifically tailored colormap is created to emphasize the performance of the classification system or highlight a specific image feature. Commonly, the data are displayed as an initial processing step in an attempt to distinguish scene elements. Rarely is visualization in

and of itself used as a tool that allows the spectral analyst to perform identification before cueing more powerful processing strategies.

A significant amount of research has been performed on perceptually based color displays for use with scientific data [3], [4]. Despite the known advantages of such strategies, most colormaps in use with hyperspectral imagery today have been developed based on the mathematics of the images without considering the workings of the human vision system and how the information will be perceived by the viewer. The resulting color images are often difficult to interpret, and in some cases the display strategy obscures some image features in the process of highlighting others.

In this paper, an attempt is made to incorporate knowledge of human color processing into display strategies for spectral imagery in order to create images that present the rich information available from spectral sensors in a readily interpretable manner. A color display strategy based on the opponent-colors model for human color perception [5] is used to present noncovariant spectral information channels derived from a principal component (PC) transformation [6], [7]. The long-range goal of the project is to develop an invariant display strategy that presents important scene constituents in a consistent and intuitive fashion as environmental conditions change.

The remainder of this paper is organized as follows. Section II provides background material on color vision and spectral imagery as well as the relation between the two. Section III discusses three-dimensional (3-D) scatterplots in PC space and introduces the concept of a conical data cloud. The new opponent color mapping strategy for spectral data is presented in Section IV. Discussion of the associated pseudocolor mapping strategies is presented in Section V, and the conclusions are presented in Section VI.

II. BACKGROUND

A. Foundations of Human Color Vision

It is well known that human color vision is a 3-D system that operates by sampling the visible spectrum with three photopigments that have maximum sensitivity at approximately 440-nm (short-wavelength or *S* cones), 540-nm (medium-wavelength or *M* cones), and 575-nm (long-wavelength or *L* cones) [8]. It has also been known for many years that the human vision system processes color by means of an achromatic (sum) channel and two opponent-color channels, one which is red-green opponent, one which is blue-yellow opponent [9]. These opponent color

Manuscript received February 6, 2002; revised October 17, 2002.

J. S. Tyo is with the Electrical and Computer Engineering Department, the University of New Mexico, Albuquerque, NM 87131 USA (e-mail: tyo@ieee.org).

A. Konsolakis and D. I. Diersen are with the Electrical and Computer Engineering Department, U.S. Naval Postgraduate School, Monterey, CA 93943 USA.

R. C. Olsen is with the Physics Department, U.S. Naval Postgraduate School, Monterey, CA 93943 USA (e-mail: olsen@nps.navy.mil).

Digital Object Identifier 10.1109/TGRS.2003.808879




Fig. 1. Two common mapping strategies for spectral data. (A) Three widely spaced bands are chosen in hopes that they are minimally correlated $\lambda = (2.101 \mu\text{m}, 2.2008 \mu\text{m}, 2.13402 \mu\text{m})$. (B) Three specific bands are chosen to highlight a particular spectral feature $\lambda = (2.2307 \mu\text{m}, 2.2008 \mu\text{m}, 2.1309 \mu\text{m})$. It is expected that the mineral kaolinite will appear as magenta in panel (B), as the green intensity in the RGB image is chosen to lie at the bottom of a kaolinite absorption feature.

channels have been demonstrated to be “cardinal directions in color space” [5] and to be perceptually orthogonal.

Buchsbaum and Gottschalk [10] demonstrated that the achromatic, red-green opponent, and blue-yellow opponent color processing channels represent statistically noncovariant information pathways derivable from a PC analysis of the spectral sensitivities of the three classes of photoreceptors. The eigenvectors of the photoreceptor response covariance matrix have spectral characteristics which are approximately achromatic (A), red-green opponent ($R - G$), and blue-yellow opponent ($B - Y$). An analysis of the corresponding eigenvalues indicates that the A channel contains 95.7% of the information, the ($R - G$) channel 2.8%, and the ($B - Y$) channel 1.5%. Their result was the first mathematical explanation for why the human vision system processes color via opponent channels. The opponent colors formulation is used in following sections to develop a display strategy for high-dimensional spectral imagery based on a similar PC analysis.

Radiance and the photoreceptor responses are nonnegative, so the covariance matrix is also nonnegative. As the data lies in the first octant in (R, G, B) space, the data cloud can be circumscribed by a convex cone. This insight has been used to develop the hue, saturation, value (H, S, V) color representation system (the term value often replaced by lightness in reflective applications and by brightness in emissive applications [4]). The (H, S, V) representation assumes that color space is conical, with the A channel as the axis of the cone and the ($R - G$) and ($B - Y$) channels orthogonal to it and to each other. The color parameter hue is equivalent to angle within the cone, and varies from the perceived color red through yellow to green, then through blue back to red. The parameter saturation indicates the purity of a given hue and measures the distance of a given perceived color from the gray axis. A saturation value of 1 lies on the surface of the cone and represents a “pure” hue. A saturation of 0 lies on the axis of the cone and is perceived as gray. The parameter value measures the projection of a given perceived color onto the A axis [11]. The parameters hue, saturation, and value are the color features that are most readily identified by human observers, and can provide a perceptually uniform color space if properly mapped to the display device [4], [3]. A perceptually uniform space is one where a unit change in one of the color parameters (hue, saturation, or value) provides a unit change in

discriminability for and observer. (R, G, B) color spaces are not perceptually uniform [3].

B. Similarities Between Spectral Imagery and Human Color Vision

Human color vision is essentially a special case of spectral imagery (namely three-band multispectral in the visible with overlapping, broad spectral samples). The high-dimensionality associated with spectral imagery at each pixel in a scene is often reduced by means of a principal components [12], [7] or minimum (or maximum) noise fraction (MNF) transformation [13], the mechanisms of which are well known and not reviewed here.

The PC transformation is obtained by diagonalizing the covariance matrix. A projection of the data from band space onto the complete set of orthogonal eigenvectors of the covariance matrix rotates the data cloud so it is aligned with a set of basis vectors that are statistically noncovariant. These eigenvectors represent a convenient set of basis axes for a particular image, just as the A , ($R - G$), and ($B - Y$) channels do in human color vision. The features of the PC transform of spectral imagery data are almost identical to those of the A , ($R - G$), and ($B - Y$) opponent color channels in human color vision [10]. This similarity is the basis for the mappings discussed in Section IV below.

C. Choice of Radiance versus Reflectance Space

In the field of spectral imagery, there is always the question of whether to process data in radiance space or in reflectance space. A number of physics-based atmospheric compensation techniques exist that attempt to minimize the effect of the atmosphere on the final imagery data [14]. The primary argument for working in reflectance space is to be able to match the observed spectra with the underlying illumination bias removed to a spectral library composed of material reflectance spectra measured in the laboratory. When no such comparison is made, there is no real motivation to move out of radiance space. In fact, the computational complexity involved in making the transformation is alone enough to demotivate such actions. For the particular application considered here, i.e., display of spectral data, we argue that it is more intuitive to remain in radiance space. The reason for this is the similarity between color vision and spectral imagery discussed above, and the fact that human vision fundamentally works on radiance data.

D. Conventional Color Mappings

The data chosen for initial analysis are from a 1995 Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) [15] overflight of Cuprite Mining District, NV, covering a region that has been used extensively for remote sensing experiments. AVIRIS is a whisk broom sensor that builds up a spectral image by sampling the spectrum one pixel at a time. Cross track scanning is accomplished in the optics of the system; along track scanning is accomplished through the motion of the platform. AVIRIS has 224 spectral channels with spectral bands ranging from 400–2500 nm in wavelength. Different detectors are used in different portions of the spectrum, so the spectral bandwidth varies through the range. The data used in this study




Fig. 2. (A) Representative radiance spectrum of kaolinite taken from a region in the image known to contain that mineral on the surface. (B) Reflectance of Kaolinite from the U.S. Geological Survey spectral library (provided with ENVI). The vertical lines give the spectral locations corresponding to the red, green, and blue intensities in panel (B) of Fig. 1.

are distributed with the ENVI software package and are used in the ENVI tutorials [16]. The data cover 50 equally spaced, contiguous spectral bands from 1.99–2.48 μm (AVIRIS bands 173 to 222). All analysis performed here was accomplished using Matlab (version 6.0.0, R12) with the Image Processing Toolbox (version 2.2.2, R12). ENVI (version 3.2) running in IDL (version 5.2.1) was used to preprocess the data.

Data like those in Fig. 1 present an interesting problem: how can the hyperspectral imagery information from a portion of the spectrum not perceived by humans be presented in a manner that is understood with minimal training? Fig. 1 presents examples of two popular methods for mapping hyperspectral data into pseudocolor images. In Fig. 1(A), three widely spaced bands at wavelengths of 2.101 μm , 2.208 μm , and 2.3402 μm (corresponding to AVIRIS bands 183, 193, and 207) are used to specify the (R, G, B) triples at each pixel in the image. These bands were chosen because they cover much of the available spectrum, and are likely to be minimally correlated due to their spectral separation. No statistical analysis went into choosing these three wavelengths. Fig. 1(B) uses the spectral images at $\lambda = (2.2307 \mu\text{m}, 2.2008 \mu\text{m}, 2.1309 \mu\text{m})$ (corresponding to AVIRIS bands 196, 193, and 186) These wavelengths were chosen to highlight a particular absorption feature corresponding to the mineral kaolinite, whose representative radiance spectrum from this image is presented in Fig. 2. The red, green, and blue digital numbers were automatically stretched using a linear stretch that saturates the top and bottom

2% of the pixels in each color band and interpolates linearly between [16].

The mapping presented in Fig. 1(A) highlights features of the image and segments the image into regions corresponding to different materials. However, there are many possibilities for choosing the bands corresponding to the (R, G, B) intensities. The color scheme will change for each of these, and there is no *a priori* way to predict the perceived hue for a particular spectral signature. Furthermore, if an individual spectral feature does not overlap with one of the bands chosen for display, that information will not be represented in the final image. For highly specialized mappings like that presented in Fig. 1(B), it is possible to predict the perceived color for the spectrum that is highlighted, but it is not easy to predict the perceived color of materials with other spectral signatures.

PC analysis offers a potential solution to the problem of seemingly arbitrary choice of bands to map into an RGB image. The eigenvectors corresponding to the principal component images sample the entire spectral range. It is therefore more likely that prominent spectral features will be included in composing a pseudocolor image using principal components. Fig. 3 presents the first three eigenvectors of the covariance matrix for this image, and Fig. 4 presents the corresponding PC images. The first eigenvector is closely related to the average scene radiance and resembles the solar radiance distribution convolved with the atmospheric absorption and average scene reflectance. This is a common occurrence in PC analysis of spectral imagery, and it is closely related to the result obtained by Buchsbaum and




Fig. 3. First three eigenvectors of the covariance matrix for the image depicted in Fig. 1. Note that the first eigenvector is all positive, and the second and third have similar spectral frequency variation, though they are orthogonal. The 30th eigenvector has much more spectral variability, and the variance in the corresponding PC image is 1000 times lower than in the second and third PC images. All eigenvectors are normalized so that $\|\mathbf{v}\|_2 = 1$. The second and third eigenvectors have properties of “out of phase” functions like sine and cosine, as the extrema of the second PC tend to line up with zero-crossings of the third PC and vice-versa.

Gottschalk for human color vision in the visible spectrum [10]. The second and third eigenvectors each have zero-crossings, but vary reasonably slowly as a function of wavelength. The 30th eigenvector is also plotted in Fig. 3 for comparison. This eigenvector has little spectral structure, and the resulting image in Fig. 4(D) has little spatial information.

The normalized eigenvalues obtained by computing

$$\bar{\lambda}_i = \frac{\lambda_i}{\sum_j \lambda_j} \quad (1)$$

are $\lambda_1 = 0.922$, $\lambda_2 = 0.0476$, and $\lambda_3 = 0.0233$ (where λ_i represents the i th eigenvalue, not the wavelength). Summing these eigenvalues indicates that over 99.3% of the variance in the image is contained in the first three PCs. The 30th normalized eigenvalue is 2.93×10^{-5} . Unlike the intensity distribution in any particular band, the PCs provide spectral functions that sample the entire spectral range of interest, preventing any one feature from being completely missed. A shortcoming of PC-based display techniques is that spectral features that only exist in only a few pixels may not contribute enough to the covariance matrix to be represented in the first few eigenvectors.

A common method for presenting the PC images in a pseudo-color display is to map the first three PCs (or another set of three PCs) into a (R, G, B) image. An example of such an image is presented in Fig. 5(A). A second strategy that has been used by us in the past is to map the second PC into two of the three RGB intensities and the third PC into the remaining channel. An example of such an image is presented in Fig. 5(B). These mappings are discussed in greater detail in Section V below.

III. THREE-DIMENSIONAL SCATTERPLOTS IN PC SPACE

To set the stage for the new mappings proposed in Section IV, we now explore the 3-D scatterplot of the data cloud in principal components space. It is often helpful in hyperspectral imaging applications to consider the spectral radiance distribution across the N bands of the spectral imager at each pixel as a vector in N -dimensional space. For a discussion on the geometric interpretation of spectral data, the reader is referred to [17] and [18]. Scatterplots are projections of the high-dimensional spectral data onto a particular two-dimensional (2-D) planes. Usually, the axes are orthogonal dimensions derived from a PC [7] or MNF [13] transformation. Pixels that are similar project to the same general location in the scatterplot, and if they can be recognized as a distinct group, classification can be performed. Variation of the plane of projection makes it possible for an observer to find views of the data that maximally separate distinct classes of data, and supervised or unsupervised clustering decisions can be made based on some distance metric between pixels in a scatterplot.

The spectral radiance in each band of a spectral imager is, by definition, nonnegative. Because of this property, the data cloud presented by the scatterplot is guaranteed to lie in the portion of N -dimensional hyperplane where all variables are positive (analogous with the first quadrant of the 2-D plane). The data can therefore be circumscribed by a convex cone that emanates from an origin corresponding to the point of shade [17]. The color display strategies proposed below map data from this cone to the (H, S, V) cone discussed in Section II-A above. In order




Fig. 4. PC image numbers 1, 2, 3, and 30. The high spatial frequency information corresponding to geography is represented in the first PC image [panel (A)], while lower spatial frequency information corresponding to classification information is represented in the second and third PC images [panels (B) and (C)]. The spatial frequency content of panels (B) and (C) are similar, as are the overall spectral shapes of the second and third PCs. Panel (D) contains the 30th PC image. Notice that it has little low spatial frequency structure, as its eigenvalue indicates. There is little grouping information available in this principal component, as there is little spectral clustering in the corresponding eigenvector in Fig. 3.




Fig. 5. Pseudocolor representations obtained by conventional mappings of PC image data into (R, G, B) space. (A) $(P_1, P_2, P_3) \rightarrow (R, G, B)$. (B) $(P_2, P_2, P_3) \rightarrow (R, G, B)$.

to properly implement these mappings, it is essential to properly locate the tip of this cone.

A. Two-Dimensional Projections of Scatterplots

While it is only possible to present the data projected onto a single 2-D plane, the discussion above suggests that it might be beneficial to consider the data as a cloud in 3-D space for the purposes of creating a display. The scatterplot corresponding to




Fig. 6. Scatterplot of the data from Fig. 5 projected onto the three principal planes of the 3-D data cone. (A) $P_2 - P_3$ plane. (B) $P_2 - P_1$ plane. (C) $P_3 - P_1$ plane. The pixels highlighted in red were from one of two regions: the bright pixels were from the playa at the lower right of the image, and the dark pixels were from the dark region just left of center in the image. The pixels highlighted in green were from the mountainous region in the upper right of the image. The two groups of highlighted pixels were assumed to lie near lines emanating from the unknown origin of the conical data space. The red and green lines represent least-squares fits to the highlighted data clouds, and their near intersection at $P_1 = 0$ provides an estimate for the location of the origin. The geometric mean of the data cloud is at $P_1 = 3716$, $P_2 = -0.604$, $P_3 = 1.25$. In panel (C), the red and green clusters are located *behind* the majority of the data cloud, as can be seen from the top view in panel (A).

the PC images presented in Fig. 4 is shown in Fig. 6 projected onto the three principal planes of the space spanned by the first three eigenvectors. (Scatterplots like those in Fig. 6(B) and (C) are usually presented without extending the range of the P_1 axis back to 0, disguising the relationships that are discussed below.) An important first step in implementing the color mappings in Section IV is the proper location of the origin of the conical scatterplot in Fig. 6. If the origin is incorrectly placed, then classes that radiate out from the actual origin will appear at different angles with respect to the incorrectly computed origin as intensity varies. Locating the origin of the cone is equivalent to identification of the point of shade. The point of shade is the spectral signature of a pixel that has no reflectance, and the observed radiance is due solely to the path radiance. Proper orientation of the point of shade is important for a number of spectral processing tools, including subspace projection and linear unmixing techniques [19], [20]. The point of shade is usually assumed to be near the darkest pixels in the image [20]; however, in this particular image, there are no pixels with zero radiance (in fact, the range of values for P_1 is from a digital number of 1732 to 8627). If the darkest pixel in the image is chosen for the point of shade, then the origin is likely to be located near the lower intensity group of pixels highlighted in red. It is shown below that this choice of origin is incorrect, as the true point of shade is not included in the dataset.

B. Location of the Origin

In this section, a supervised method for locating the origin of the conical data cloud is introduced that depends on the ability to locate common classes of materials with variable intensities. Common classes are identified based on their location in a common direction in the conical data space. In the dataset under investigation two such classes can be identified, and they are highlighted in the scatter plots in Fig. 6. The points highlighted in green are from the mountainous regions in the upper center and upper right portions of the scene. The points highlighted in red come from two regions, one from the bright playa at the

right edge of the image and the other from the dark region just left of center. Linear regressions were fit to these classes and extrapolated back to the $P_1 = 0$ plane in order to locate a vertex for the cone. The red line extrapolates to $P_2 = -30$, $P_3 = 34$, the green line to $P_2 = -30$, $P_3 = 38$. The vertex of the cone was therefore placed at $P_1 = 0$, $P_2 = -30$, $P_3 = 36$. In this dataset, the origin is very close to the $P_1 = 0$ plane. However, for datasets that include the visible portions of the spectrum, it is expected that the origin (point of shade) will correspond to a nonzero spectrum that is biased toward short wavelengths [19].

The effects of improper and proper orientation can be observed by examining Fig. 6(A), Fig. 7, and Fig. 8. In Fig. 6(A), the two clusters highlighted in red that have been hypothesized as lying near the same radial line actually appear at two different angles with respect to the (improper) origin of the figure. In Fig. 7, the same three views of the scatterplot are presented, but the second and third PC images have been divided by the value of the first PC at each pixel. This transformation should convert the conical data space to a cylindrical space. All points with the same azimuthal and polar angles in the cone should fall on a line parallel to the axis of this cylinder. Note that the highlighted clusters fall on tilted lines in Fig. 7.

In contrast, Fig. 8 presents almost the same transformation, but in the latter figure, the offset PC values

$$P'_2 = P_2 + 30 \quad (2)$$

$$P'_3 = P_3 - 36 \quad (3)$$

are presented normalized to P_1 . All points that make the same angle

$$\tan\theta = \frac{\sqrt{P'^2_2 + P'^2_3}}{P_1} \quad (4)$$

with respect to the axis of the cone now appear at the same cylindrical radius in Fig. 8. Note in Fig. 8 that the two distinct red clusters in Fig. 6 have merged in the top view [panel (A)], and are stacked vertically in the two side views.




Fig. 7. Conical data normalized by P_1 without first shifting the origin. (A) $P_2 - P_3$ plane. (B) $P_2 - P_1$ plane. (C) $P_3 - P_1$ plane. There is a slight tilt to the data that will affect the opponent color mapping strategy presented in Section IV.




Fig. 8. Conical scatterplot of Fig. 6 converted to a cylindrical scatterplot. The offset values of the second and third PCs as given in (2) and (3) are normalized by the first PC. (A) $P_2 - P_3$ plane. (B) $P_2 - P_1$ plane. (C) $P_3 - P_1$ plane. All points with the same spectral character but different intensities should appear at the same cylindrical radius, but at different vertical positions.

IV. OPPONENT COLORS MAPPING STRATEGY

The discussion in Section II-B presents the close relationship between the nature of the first three PCs in the spectral image being investigated and the corresponding spectral channels in human vision. The spectrum of the first eigenvector is roughly achromatic in that it samples the mean pixel radiance. The resulting image contains the approximate overall reflectance distribution throughout the image, and is known to contain the high spatial frequency information corresponding to the scene topography in typical natural radiance scenes. This channel contains almost 93% of the scene variance, and is analogous to the A channel in human vision. The second and third PC images have a total of about 7% of the scene variance and contain lower spatial frequency information than the first PC image, as is the case for the $(R - G)$ and $(B - Y)$ channels in human color vision [21]. The high spatial frequency information in Fig. 4(A) is approximately equivalent to an overall reflectance map, as the first PC computes a positively weighted sum across the entire spectrum of interest. This image contains much of the scene topography that would be evident in a panchromatic image. This is to be expected in radiance images of natural scenes. In reflectance images, the first eigenvector is no longer tied to the illumination function, and this concentration of the high spatial frequency information in the first PC image may not hold. The second and third PC images in Fig. 4(B) and (C) present some classifica-

tion information, as regions with similar spectral features tend to group together. In natural scenes like this one, there tends to be spatial clustering of materials. In Section III, the eigenvectors were shown to be the basis for a conical data space, just as the A , $(R - G)$, and $(B - Y)$ channels are in color space. These similarities suggest an ergonomic mapping from high-dimensional data space to a 3-D pseudocolor space that aligns data cone discussed in Section III with the cone of the (H, S, V) color representation scheme.

To implement this mapping, the first PC image is mapped into the A channel, the second PC image is mapped into the $(R - G)$ channel, and the third PC image is mapped into the $(B - Y)$ channel. This mapping is equivalent to the following strategy:

$$\begin{aligned} \phi &= a \tan \left(\frac{P'_3}{P'_2} \right) \rightarrow H \\ \frac{\sqrt{P_2'^2 + P_3'^2}}{P_1} &\rightarrow S \\ P_1 &\rightarrow V \end{aligned} \quad (5)$$

where H , S , and V are hue, saturation, and value, and P_i is the i th principal component (offset for P_2 and P_3). A hue wheel is presented in Fig. 9 that provides the geometrical mapping of angles to hues. The transformation in (5) is similar to a pseudocolor mapping strategy that has been suggested for multispectral




Fig. 9. Four mappings which translate the first three PC images into the (H, S, V) color space. (A) Full implementation of (5) with saturation amplified so that the maximum value of S in the image is unity. (B) Suppression of saturation. Each pixel is fully saturated, but the first and third equations in (5) are still employed. (C) Only the hue is mapped. (D) A specialized mapping designed to locate regions of an image that are similar or dissimilar to a particular spectral signature. In this case, yellow pixels have a positive second PC, blue pixels negative. The third PC information is suppressed.

[22] and polarization [23], two other examples of multidimensional imagery.

Examination of (5) demonstrates why it was essential to properly locate the origin of the data cone in Section III-B. Use of the angles computed by (5) without reorienting the origin would result in different hues for the two red clusters highlighted in Fig. 6, even though it is apparent that they have the same spectral character (at least in the first three PCs). The saturation and hue of the pixels in the green cluster would vary continuously as the intensity (P_1) increased. The result of the tilt apparent in Fig. 7 is that pixels of like materials will appear with slightly different saturation and hue as intensity increases, thus producing a lack of color constancy [24].

Once the mapping described by (5) is accomplished, there are still several choices available for how the information will be displayed. Fig. 9(A) presents the full mapping, with the value

of S stretched so that the maximum saturation in the image is unity. In this image, the value of S obtained using (5) was amplified by a factor of 6.6. Since the amount of variance in the second and third PC images together is an order of magnitude smaller than that in the first PC image, it is expected that the maximum distance from the axis of the cone in any particular plane should be about an order of magnitude smaller than the corresponding value of P_1 (intensity). This image represents all of the information in the three PC images of Fig. 4 in a manner that is ergonomic with the human vision system. The geographical information (high spatial frequency) is maintained in the A channel, while the classification information from P_2 and P_3 is mapped into color. There are many pixels that are highly saturated (after stretching), but large regions of the image have low saturations, as these pixels have small projections on the P_2-P_3 plane.




Fig. 10. Five different material mean spectra obtained by averaging over a ground-truthed region of the data. (A) Mean spectrum. (B) Residual remaining after subtracting the projection of the mean spectrum onto the first PC shown in Fig. 3. Materials are identified by name in Table I.

A second possible mapping involves the suppression of saturation information. This method, shown in Fig. 9(B), presents each pixel as fully saturated (pure hue), and results in a “closest classification” image, i.e., each pixel is assigned to the material (hue) that it most closely resembles. Similar strategies are often employed for display when using an endmember selection strategy [25]. However, by superimposing the intensity information from P_1 , the geographical features can still be seen.

A third possible mapping, presented in Fig. 9(C), involves suppression of both saturation and value, and results in a typical classification image where each pixel is assigned to the material (hue) that it matches most closely with $S = 1$ and approximately constant intensity [$V = 0.75$ in Fig. 9(C)]. Fig. 9(D) presents a specialized mapping analogous with Fig. 5(B). In this figure, the data are projected onto the plane spanned by the first and second eigenvectors. Spectra with positive values of P_2 are presented as yellow, those with negative projections as blue. The value of saturation is determined by the angle made with the first eigenvector, just as in Fig. 9. This strategy is loosely equivalent to the spectral angle mapper algorithm. [19] The high spatial frequency information is preserved by mapping the first PC into value. (It should be noted that the (H, S, V) to (R, G, B) transformation used is the one provided with the Matlab Image Processing Toolbox. No effort has been made to make the display perceptually uniform, as such efforts are beyond the scope of this work.)

TABLE I
PROJECTION OF THE MEAN MATERIAL SPECTRA ON THE FIRST THREE PCs, AS WELL AS ANGLE IN THE $P_2 - P_3$ PLANE AS INDICATED IN THE SCATTERPLOT OF FIG. 6. THE COLOR COLUMN PRESENTS THE FINAL (H, S, V) TRIPLE USED TO REPRESENT THE MEAN OF THE SPECTRUM IN FIG. 9

Material	P_1	P_2	P_3	Angle	Sat.	Color
Calcite	4010	-96	-76.6	-143°	0.030	
Alunite	3810	169	217	52.8°	0.073	
Kaolinite	4060	294	-34.1	-5.46°	0.073	
Buddingtonite	3980	-79.0	301	104°	0.080	
Silica	4990	-217	29.2	170°	0.032	
Playa	6020	-196	69.0	159°	0.035	
Varnished Tuff	3590	-176	92.2	151°	0.056	

V. DISCUSSION

A. Comparison of New Mapping to Conventional Mappings

In Section II-D, several conventional color mappings were presented. In this section, a brief discussion of the relative merits of each method is presented here, and the opponent colors mapping method outlined above is compared with the more conventional approaches.

The class of mappings presented in Fig. 1 is possibly the most powerful set of methods. These mappings present different band intensities as the R, G, and B intensities in the final image. When a particular spectral feature is to be highlighted, a particular material uncovered, etc., then a specifically tailored color

mapping strategy should be developed. Numerous examples of such strategies have appeared in the literature. The images in Fig. 1 nicely present the high spatial frequency information that depicts the geography. This is expected because areas that have uniformly low reflectance will be dark in all three bands, and therefore dark in the final image. The high spatial frequency has therefore been mapped into the A channel as is desirable. The problem with these strategies is that it is not easy to predict *a priori* the effect of the mapping on the elements of the scene that are not being emphasized. Simple mapping of spectral bands into (R, G, B) space can overlook important spectral features, and different implementations will result in different hues in the final image.

The typical class of PC-based pseudocolor representations depicted in Fig. 5 solves the problem of using isolated spectral information, but does not take advantage of the knowledge of the human vision presented in Section II-A. The PC transformation separates the information into orthogonal variables, but the mappings used in Fig. 5 take this orthogonal information and present it in nonorthogonal visual pathways. While some of the high spatial frequency information is apparent in Fig. 5(A), much of it is mapped into the $(R - G)$ and $(B - Y)$ color channels, and is therefore suppressed. The smearing apparent in Fig. 5(A) is an artifact of how the viewer *perceives* the data. A further problem with this image is that overall reflectance information (albedo) that is captured in the first PC image is mapped into a color channel, and pixels with similar intensities (first PC values) are presented with different brightness values because the second and third PC images contribute when they should not. The mapping in Fig. 5(B) is more ergonomic, but instead of mapping the information into orthogonal channels, the second and third PCs are mapped into the $A - (B - Y)$ plane.

The image in Fig. 9(A) has one very important property that makes it easier to view, as compared with images like those in Fig. 4(A) or by techniques that assign a hue to each pixel that can be classified as a particular material [25]. Pixels that do not have a significant projection onto either P_2 or P_3 appear largely desaturated in Fig. 9(A). This makes the image easier to view because naturally occurring scenes tend to be largely desaturated with low dimensionalities in the visible portion of the spectrum [10]. Human observers are not used to examining images with large regions of highly saturated hues that vary rapidly. The set of images in Fig. 9 provides a way to go back and forth between highly saturated and less saturated images to make material classification more obvious.

There is no additional information presented in Fig. 9 beyond that in Fig. 5, as the same data went into preparing both sets of images. However, the available data are arranged more ergonomically in Fig. 9, allowing several important features of the image to be readily apparent.

B. Material Classification

From a classification point of view, the mappings presented in Section IV have the desirable property that materials that differ greatly in their spectral signature appear with large differences in hue, and materials with similar spectral characteristics appear with similar hues. No classification decision is made in implementing the algorithm; rather the image is presented in a way




Fig. 11. Scatterplot representation in the $P_2 - P_3$ plane with the point corresponding to each pixel colored with the same (H, S, V) triple as in the image in Fig. 9(A). Only one out of every twenty pixels is plotted here to aid in viewing the result. Unsaturated (gray) pixels are those near the origin, and particular hues radiate outwards. The lines are drawn from the origin to the mean of the classes in Table I. Note that the means each have a different value of P_1 which does not appear in this figure, since P_1 is orthogonal to this projection.

that allows the observer to segment the image as they see fit. The procedures outlined in this study are not a classification method *per se*; however, the manner in which various materials are presented is of interest. A number of pixels have been identified from ground-truth observations as containing specific materials. A set of seven regions of interest (ROIs) is distributed with this dataset. The mean spectrum in each of these regions is presented in Fig. 10. To find the region of the scatterplot of Fig. 6 corresponding to each of these mean spectra, a projection was taken on the first eigenvector. The residual spectra is defined here as the portion of the radiance vector at each pixel orthogonal to the P_1 axis. It is computed by

$$\mathbf{x}_{\text{resid}} = \mathbf{x} - (\mathbf{x}^T \mathbf{v}_1) \cdot \mathbf{v}_1. \quad (6)$$

Classification information is held in the residual spectra. For each of the materials, an angle in the $P_2 - P_3$ plane was computed by taking the appropriate projections (referenced to the origin found above), and the results are summarized in Table I.

The scatterplot of the data is shown again in Fig. 11 projected onto the $P_2 - P_3$ plane. Each point in the scatterplot is colored with the same (H, S, V) values as the corresponding pixel in Fig. 9(A). Lines radiating from the origin to the location of the means of the sample ROIs in Table I are also superimposed on the data. There is a group of pixels apparent in Fig. 11 that is shaded green that is not included in the ROIs in Table I. These pixels appear to come from a major scene constituent that was overlooked in the ROI selection process.

Two obvious problem with this display strategy from a classification point of view are that 1) materials with little spectral separation that do not dominate the image are presented with sometimes indistinguishable hues, even when the materials can be reliably differentiated using other methods; and 2) mixed pixels end up getting presented as “half way between” the two (H , S , V) values for the individual constituents. This could result in a spurious hue, or possibly a desaturated pixel, depending on the orientation of the vectors corresponding to the two constituents. Additionally, a target that only occupies a small number of pixels might not significantly impact the covariance matrix. Such a target spectrum may therefore end up with little or no projection onto the $P_2 - P_3$ plane.

VI. CONCLUSION

The principal component-based mapping strategy discussed here provides a straightforward method to perform an initial unsupervised classification without making any formal decisions. The resulting image is segmented spatially based on the projection of the local radiance distribution in the $P_2 - P_3$ plane. By visually inspecting the resulting image, the analyst can direct attention to appropriate areas of the scene for further processing. No classification decision is made by the algorithm, as final segmentation is done by the human observer.

The PC transformation presented here was derived from in-scene statistics. This can be problematic as spectral features from the scene that appear in a large number of pixels can leak into the first eigenvector. Furthermore, the strategy presented above is not invariant, as the specific transformation will vary from scene to scene as the statistics change. The question still remains as to whether a general set of basis vectors can be derived that is broadly applicable to classes of spectral imagery. If a set of general eigenvectors can be identified, the color mapping can be arranged so that materials are presented in a straightforward manner, i.e., water can always be mapped to blue, etc. Also, as a wider range of wavelengths is considered, it should be expected that more than three PCs will be necessary to capture an equivalent amount of the scene variance (99% or more) [26], [27].

The presentation strategy discussed here is best suited to depicting large-scale image features, not for identifying small, isolated targets. Targets which occur only at a few pixels in an image will have little effect on the overall covariance matrix, and will not contribute significantly to the second and third PCs. For this reason, such small targets *may* have small projections in the $P_2 - P_3$ plane. This is a subject that merits significant investigation.

The pseudocolor mapping strategy presented here provides two main benefits. First, it is offered as a way to take an initial look at the data. By performing an unsupervised PC transformation and analyzing the most significant PCs, initial classification decisions can be made. Better yet, if an invariant strategy is developed, the scene can be mapped to a standard color representation that allows ready interpretation. It is important to compare the PC eigenvectors that result in classes of images in order to determine if a standard transformation can be applied for particular remote sensing tasks, similar to the Tasseled Cap

Transform [28] developed for multispectral LANDSAT sensing of agricultural scenes. Detailed investigation of the relationship between the PC eigenvectors and image content is needed to determine how robust the strategy is. Comparison with other classification strategies is merited in order to quantify the usefulness of this presentation strategy in general. Second, an image composed using the PC-based strategy can form the backdrop for display of the results of more highly sophisticated classification algorithms. The transformation was illustrated here with 50-band hyperspectral data, but similar results have been obtained with full HYDICE and AVIRIS datasets in excess of 200 bands [27].

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their exceptionally insightful comments that helped strengthen the final version of this manuscript.

REFERENCES

- [1] G. Vane and A. F. H. Goetz, “Terrestrial imaging spectroscopy,” *Remote Sensing Environ.*, vol. 24, pp. 1–29, 1988.
- [2] A. F. H. Goetz, “Imaging spectrometry for remote sensing: Vision to reality in 15 years,” in *Proceedings of SPIE*, M. R. Descour, M. J. Mooney, D. L. Perry, and L. Illing, Eds. Bellingham, WA: SPIE, 1995, vol. 2480, Imaging Spectrometry, pp. 2–13.
- [3] P. K. Robertson, “Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data displays,” *IEEE Comput. Graph. Appl.*, vol. 8, pp. 50–64, Sept. 1988.
- [4] P. K. Robertson and J. F. O’Callaghan, “The generation of color sequences for univariate and bivariate mapping,” *IEEE Comput. Graph. Appl.*, vol. 6, pp. 24–32, Feb. 1986.
- [5] J. Krauskopf, D. R. Williams, and D. W. Heeley, “Cardinal directions of color space,” *Vision Res.*, vol. 22, pp. 1123–1131, 1982.
- [6] I. T. Jolliffe, *Principal Components Analysis*. New York: Springer-Verlag, 1986.
- [7] P. J. Ready and P. A. Wintz, “Information extraction, SNR improvement, and data compression in multispectral imagery,” *IEEE Trans. Commun.*, vol. 21, pp. 1123–1131, 1973.
- [8] J. J. Vos and P. L. Walraven, “On the derivation of the foveal receptor primaries,” *Vision Res.*, vol. 11, pp. 799–818, 1971.
- [9] G. Wyszecki and W. S. Stiles, *Color Science*. New York: Wiley, 1967.
- [10] G. Buchsbaum and A. Gottschalk, “Trichromacy, opponent colors coding and optimum color information transmission in the retina,” *Proc. R. Soc. Lond. B*, vol. 220, pp. 89–113, 1983.
- [11] B. Fortner and T. E. Meyer, *Number by Colors*. Santa Clara, CA: TELOS, 1997.
- [12] H. L. van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: Wiley, 1968.
- [13] A. A. Green, M. Berman, P. Switzer, and M. D. Craig, “A transformation for ordering multispectral data in terms of image quality with implications for noise removal,” *IEEE Trans. Geosci. Remote Sensing*, vol. 26, pp. 65–74, Jan. 1988.
- [14] B.-C. Gao and A. F. H. Goetz, “Column atmospheric water vapor and vegetation liquid water retrievals from airborne imaging spectrometer data,” *J. Geophys. Res. D*, vol. 95, pp. 3549–3564, 1990.
- [15] G. Vane, R. Green, T. Chrien, H. Enmark, E. Hansen, and W. Porter, “The airborne visible infrared imaging spectrometer,” *Remote Sens. Environ.*, vol. 44, pp. 127–143, 1993.
- [16] Research Systems, Inc., “ENVI v. 3.2 User’s Guide,” Boulder, CO.
- [17] J. W. Boardman, “Analysis, understanding, and visualization of hyperspectral data as convex sets in n -space,” in *Proceedings of SPIE*, M. R. Descour, M. J. Mooney, D. L. Perry, and L. Illing, Eds. Bellingham, WA: SPIE, 1995, vol. 2480, Imaging Spectrometry, pp. 14–22.
- [18] G. Healey and D. Slater, “Models and methods for automated material identification in hyperspectral imagery acquired under unknown illumination and atmospheric conditions,” *IEEE Trans. Geosci. Remote Sensing*, vol. 37, pp. 2706–2717, Nov. 1999.

- [19] S. M. Adler-Golden, R. Y. Levine, M. W. Matthew, S. C. Richtsmeier, L. S. Bernstein, J. Gruninger, G. Felde, M. Hoke, G. P. Anderson, and A. Ratkowski, "Shadow—Insensitive material detection/classification with atmospherically corrected hyperspectral imagery," in *Proceedings of SPIE*, S. S. Shen and M. R. Descour, Eds. Bellingham, WA: SPIE, 2000, vol. 4381, Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII, pp. 460–469.
- [20] P. G. Lucey and M. E. Winter, "Efficient materials mapping for hyperspectral data," in *Proceedings of SPIE*, S. S. Shen and M. R. Descour, Eds. Bellingham, WA: SPIE, 2001, vol. 4381, Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII, pp. 164–172.
- [21] K. T. Mullen, "The contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings," *J. Physiol. (London)*, vol. 359, pp. 381–400, 1985.
- [22] P. K. Robertson and J. F. O'Callaghan, "The application of perceptual color spaces to the display of remotely sensed imagery," *IEEE Trans. Geosci. Remote Sensing*, vol. 26, pp. 29–39, Jan. 1988.
- [23] J. S. Tyo, E. N. Pugh, and N. Engheta, "Colorimetric representations for use with polarization-difference imaging of objects in scattering media," *J. Opt. Soc. Amer. A*, vol. 15, pp. 367–374, 1998.
- [24] D. H. Brainard, B. A. Wandell, and E.-J. Chichilnisky, "Color constancy: From physics to appearance," *Curr. Dir. Psychol. Sci.*, vol. 2, no. 5, pp. 165–170, 1993.
- [25] M. E. Winter, "Fast autonomous spectral endmember determination in hyperspectral data," in *13th Proc. Int. Conf. Applied Geologic Remote Sensing*, Vancouver, BC, Canada, 1999.
- [26] J. S. Tyo, "Relationship between system condition and systematic errors in Stokes vector polarimeters," in *Proceedings of SPIE*, D. H. Goldstein and D. B. Chenault, Eds. Bellingham, WA: SPIE, 2002, vol. 4481, Polarization Measurement, Analysis, and Remote Sensing IV, pp. 22–30.
- [27] A. Konsolakis, "Invariant display strategies for hyperspectral imagery," Master's thesis, Naval Postgraduate School, Monterey, CA, 2001.
- [28] E. P. Crist and R. J. Kauth, "The tasselled cap demystified," *Photogramm. Eng. Remote Sens.*, vol. 52, pp. 81–86, 1986.



J. Scott Tyo (S'96–M'97) was born at the Frankfurt American Hospital in 1972. He received the B.S.E., M.S.E., and Ph.D. degrees in electrical engineering from the University of Pennsylvania, University Park, in 1994, 1996, and 1997, respectively.

He is currently an Assistant Professor in the Electrical and Computer Engineering Department, University of New Mexico (UNM), Las Cruces. From 1994 to 2001, he was an officer in the U.S. Air Force, leaving service at the rank of Captain. He joined the faculty at UNM in 2001. From 1996 to 1999, he was

a Research Engineer in the Directed Energy Directorate of the USAF Research Laboratory, Kirtland AFB, NM. From 1999 to 2001, he was a member of the faculty of the ECE Department, U.S. Naval Postgraduate School, Monterey, CA. Since joining the faculty at UNM, his research has focused on ultra wideband microwave radiating systems as well as microwave and optical remote sensing.

Dr. Tyo is a member of the Optical Society of America, Commissions B and E of URSI, SPIE, Tau Beta Pi, and Eta Kappa Nu.

Athanasios Konsolakis, photograph and biography not available at the time of publication.

David I. Diersen, photograph and biography not available at the time of publication.



Richard Christopher Olsen (M'95) received the B.S. degree in physics from the University of Southern California, Los Angeles, in 1974, and the M.S. and Ph.D. degrees in physics from the University of California at San Diego, La Jolla, in 1976 and 1980, respectively.

He is currently a Professor at the U.S. Naval Postgraduate School (NPS), Monterey, CA, with a joint appointment in Physics and Space Systems. His research interests are in a broad range of areas associated with the exploitation of remote sensing imagery,

particularly spectral imagery. His other interests are in data fusion, particularly the combination of data from radar and electro-optical sensors. He has had over 80 thesis students working in space science and remote sensing at NPS.