Design Goals and Solutions for Display of Hyperspectral Images

Nathaniel P. Jacobson, *Student Member, IEEE,* and Maya R. Gupta, *Member, IEEE*

***Abstract*—Design goals and solutions are proposed for the display of hyperspectral imagery on tristimulus displays. The requirements of a hyperspectral visualization depend on the task. We focus on creating consistent representations of hyperspectral data that can facilitate understanding and analysis of hyperspectral scenes, and may be used in conjunction with task-specific visualizations. Fixed linear spectral weighting envelopes are given, creating natural-looking imagery where hue, brightness, saturation, and whitepoint have meanings consistent with the human visual system interpretation of natural scenes. For Airborne Visible/Infrared Imaging Spectrometer images, hue interpretation of water and vegetation is also preserved. The proposed designs avoid the preattentive distractions of principal component analysis imagery, and appear to provide comparable or enhanced spectral and edge discriminability.**

***Index Terms*—Dimensionality reduction, hyperspectral imaging, image fusion, multidimensional signal processing, principal components analysis (PCA), remote sensing, visualization.**

I. Introduction

HYPERSPECTRAL and multispectral images contain many more image bands than can be displayed on a tristimulus display. Tristimulus displays include standard monitors and LCD displays, and any display has three image channels, commonly red, green, and blue. Human interaction with hyperspectral and multispectral images is vital: to direct computerized methods, to validate automated analysis, and to make appropriate decisions and interpretations [1]. The question arises: how to display hyperspectral and multispectral images in a way that best enables human interaction with the images and computed results?

All three-band representations of an originally N-band image necessarily represent a loss of information for $N > 3$. For any visualization scheme, a particular color value could represent many different N-dimensional vectors. Normal human vision suffers the same problem: many different visible spectra cause the same L, M, and S cone photoreceptor stimulations. These metameric spectra are often distinguished by context, shape, and other visual clues.

If all hyperspectral visualizations lose information, how does one judge whether a visualization is good? The best criterion is task-specific and requires a highly specified task and human tests to see how well a certain visualization enables the task. However, we propose that there are also task-independent goals that can be used to design and evaluate effective visualizations. In Section II, we list these goals and describe why they are useful to fast and accurate human interpretation of scenes. The primary goal we propose is *consistency*—that a spectra be displayed as the same color value whenever it appears. Consistent representations would allow humans to "learn" what different spectra look like, and be able to identify spectra and interpret scenes. The metamers that are inherent in all visualizations could be overcome by shape and contextual features, as in normal human vision.

We propose visualization solutions in Sections III and IV that satisfy the consistency goal and are designed to satisfy the other design goals proposed in Section II. Our visualization solutions are inspired by the successful human visual system, which converts broadband visible radiation into three signals roughly corresponding to red, green, and blue. In Section V, other image fusion solutions are reviewed and contrasted for hyperspectral visualizations. The design goals proposed in Section II are themselves objective criteria for comparing hyperspectral visualizations, and we propose additional objective criteria for measuring spectral discrimination in Section VI. Last, in Section VII we present the images rendered using the proposed visualization methods.

II. Goals for Hyperspectral Image Visualization

We propose the following design goals for displaying hyperspectral imagery. Not all goals will be important for all tasks, and it may not be possible to achieve all goals simultaneously. However, each of these goals would increase the effective transmission of information. In Sections III and IV, we present visualization solutions and show how they satisfy these goals in the following sections.

1) *Consistent Rendering:* Any given spectrum is always displayed as the same color value so that it can be easily recognized across images. This goal also facilitates comparison between different images. There is a further advantage if the rendered colors correspond to an existing color-association system. This constraint may be loosened to allow luminance scaling of the display image, which would result in a spectra being consistently displayed as the same hue and saturation.
2) *Edge Preservation:* Edges (at all spatial resolutions) of the original hyperspectral image are represented faithfully in the visualization. We discuss objective metrics for this goal in Section IV. A subgoal is the discriminability of different spectra in the visualization.

Manuscript received March 14, 2005; revised August 1, 2005. This work was supported in part by the National Science Foundation under Grant SBE-0123552.

The authors are with the Department of Electrical Engineering, University of Washington, Seattle, WA 98103 USA (e-mail: gupta@ee.washington.edu).

Digital Object Identifier 10.1109/TGRS.2005.857623

3) *Computational ease:* The visualization can be rendered quickly, enabling real-time interactivity.
4) *Equal-energy white point:* A pixel with the same reflectance value for each spectral band appears as a shade of gray. Thus, lack of color saturation is related to how closely an object resembles a gray body diffuse reflector.
5) *Smallest effective differences:* Visual distinctions are no larger than needed to effectively show relative differences. According to Tufte [2], who is a proponent of this design goal for general visualizations, minimal distinctions reduce visual clutter, and using smallest effective differences helps in designing secondary and structural elements such as arrows, pointer lines, highlights, legends, etc.
6) *Appropriate preattentive features:* The visualization minimizes preattentive features of the image that distract the viewer without reason. For example, a small bright saturated color region on a background of a different color will "pop-out" at the viewer and attract attention [3].
7) *Natural palette:* Visualizations use a palette and spatial distribution of colors that is consistent with natural imagery. This goal is partly based on the assumption that humans are well-designed to interpret natural scenes, and partly based on the misinterpretation caused by strongly saturated colors. In particular, large regions of strongly saturated colors rarely appear in nature, and have long been eschewed by design experts as confusing and distracting [4], [5]. In fact, strong background coloration can induce perceived differences in smaller color regions. These simultaneous contrast effects can make it difficult to accurately judge quantitative differences between colors [6]. For example, two small squares of the same color will actually look increasingly different if viewed against two backgrounds with strong differences in color. Conversely, two small squares of different colors can look the same when viewed against a background that strongly differs from the foreground. These simultaneous contrast effects are well-studied visual phenomena [7], [8], and their effect in maps can be explored on-line with Cynthia Brewer's ColorBrewer software (www.colorbrewer.org). Large regions of saturated color may also induce temporal chromatic adaption, where after staring at one part of the image other parts of the image then appear to be different colors due to the afterimage formed [7], [8].
8) *Wavelength Shift Invariance:* All wavelengths are accorded equal weight. This allows the visualization method to work equally well for any number of spectral bands, in any spectral range. This makes a visualization method easily adaptable to spectral zooming, spectral panning, or new instrumentation. A basic requirement is that a monochromatic spike $\delta(\lambda)$ at any wavelength λ is displayed with the same luminance. A further desired feature is that the perceived difference in a color property (such as hue) between two rendered monochromatic spectra $\delta(\lambda_1)$, $\delta(\lambda_2)$ depends only on the change in wavelength $\lambda_2 - \lambda_1$.

These goals are consistent with goals for scientific visualization proposed by other researchers. For example, Tyo *et al.* [9] propose "anthropometric optimality" which they describe as "information should be presented in a way that maximizes usefulness to the human observer." Our goals try to achieve this by taking into account human visual properties, such as our goals of a natural palette, appropriate preattentive features, equal-energy white-point, consistent rendering, and edge preservation. The goal of consistent rendering, or at least relative consistency with respect to luminance, satisfies the most common natural variation captured by Tyo *et al.*'s goal of "ecological invariance," that "representations remain qualitatively similar over natural variation." Robertson and Callaghan [10] argue that hyperspectral image displays should communicate data "as effectively and unamibiguously as possible." The above goals aim to support this larger goal.

III. Spectrally Weighted Envelopes

We show that a set of fixed spectral weighting envelopes create visualizations that perform well according to the goals stated in Section II, and will yield useful visualizations for a range of tasks.

The three displayed image bands R, G, and B, will be fixed linear integrations of the original hyperspectral image weighted by three different spectral envelopes. This is similar to how human photopic (daylight) vision works. In photopic human vision, three different cone types have different spectral sensitivity envelopes. Each cone absorbs incoming photons of a particular wavelength with a probability corresponding to the sensitivity of that cone to that wavelength. In this visualization method, the probabilistic sensitivity has been replaced by a deterministic weighting.

Consider an original N-band image, where $H_{ij}[\lambda_n]$ denotes the reflectance value of the nth hyperspectral band of the ith row and jth column pixel, where $n = 1, \ldots, N$. Let $\{R_{ij}, G_{ij}, B_{ij}\}$ be the RGB vector for the ith row and jth column pixel of the displayed image. Let $r[\lambda_n]$, $g[\lambda_n]$, and $b[\lambda_n]$ for $n = 1, \ldots, N$ be weights on the nth spectral band. Then the proposed visualizations are linear integrations of the form

$$R_{ij} = \sum_{n=1}^{N} r[\lambda_n] H_{ij}[\lambda_n] \tag{1a}$$

$$G_{ij} = \sum_{n=1}^{N} g[\lambda_n] H_{ij}[\lambda_n] \tag{1b}$$

$$B_{ij} = \sum_{n=1}^{N} b[\lambda_n] H_{ij}[\lambda_n]. \tag{1c}$$

The rest of this section focuses on designing the spectral envelopes r, g, b. The three-band visualization commonly formed from the first three principal components of PCA can be treated as a set of image-adaptive spectral envelopes r, g, b. However, we focus on designing fixed spectral envelopes in order to satisfy the *consistency* goal for representing hyperspectral images. Fixed spectral weightings also satisfy the design goal of *computational ease*.

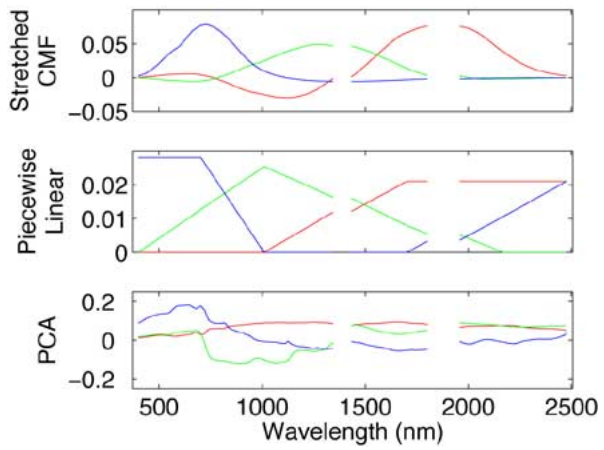

Fig. 1. Three spectral weighting envelopes are shown. Gaps are the result of AVIRIS bands that are noisy or contain no data. The top set of envelopes are the CIE 1964 color matching functions transformed to the sRGB color space with equal-energy white point and stretched across the available AVIRIS wavelengths. The second set of envelopes are piecewise linear, designed to map deep infrared spectra to magenta hues. The third set of envelopes are the PCA basis spectra extracted from the Moffett Field test image shown in Fig. 4(a).

If the total weight of each spectral weighting envelope is equal so that

$$\sum_{n=1}^{N} r[\lambda_n] = \sum_{n=1}^{N} g[\lambda_n] = \sum_{n=1}^{N} b[\lambda_n] \tag{2}$$

then the *equal-energy white point* design goal will be satisfied.

If the spectral weighting envelopes are wavelength shift invariant (and assumed to be zero outside of their defined range), they can be stretched, compressed, shifted, or cropped for spectral zooming, panning, or for use with different instruments. If the envelopes are derived from continuous functions, this can be achieved by resampling. Interpolation is needed for discrete envelopes.

This leaves an enormous set of possible choices for the spectral envelopes r, g, and b. We considered a number of reasonable designs, two of which are discussed below.

The first design is based on a stretched version of the CIE 1964 tristimulus color matching functions (CMFs) (http://cvision.ucsd.edu/), and this design is shown at the top of Fig. 1. The color matching functions for a particular wavelength were derived in experiments that determine how much of each of three primary colored lights (red, green, blue) had to be mixed together to give a viewer the same impression as a monochromatic light of that wavelength [11]. As noted by Polder and van der Hijden [12], a hyperspectral image that covers the visible range can be visualized as a colorimetrically correct image by applying the CIE color matching envelopes to the visible range hyperspectral image.

To use the color matching functions as a set of visualization envelopes, we took the CIE 1964 tristimulus color matching envelopes, transformed the envelopes to the sRGB colorspace [8] (sRGB is a standardized RGB space that can be used with most monitors), and multiplied by the D65 white point needed to illuminate the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) reflectance data for display in the sRGB color space. The wavelength scale for the envelopes was then stretched

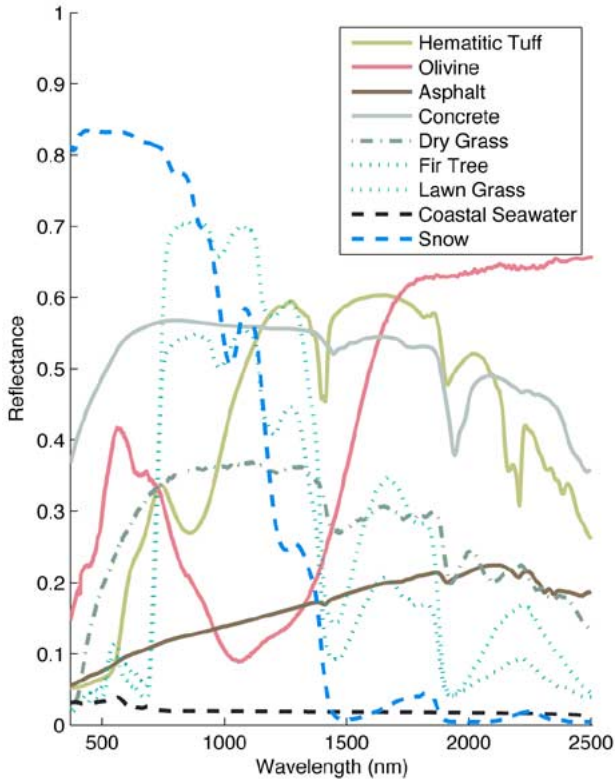

Fig. 2. Some example spectra from the USGS spectral library are shown. Line color is the rendered color for each spectra using the stretched CMF envelopes.

to cover the AVIRIS hyperspectral range. The stretch was performed by equating $\lambda = 400$ nm to the first valid AVIRIS band (band 4), equating $\lambda = 680$ nm to the last valid AVIRIS band (band 221), and linearly interpolating values between them. These envelopes reproduce what the eye might see if the entire AVIRIS range was compressed to the visible range. This process could be easily modified for a different instrument or spectral range.

It should be noted that the resulting color matching function envelopes (CMF envelopes) have some negative values. These are a consequence of the limited color gamut of tristimulus displays. Spectra that are strongly reflective in a negative envelope region will be rendered with a negative sRGB value (corresponding to a "real" but undisplayable color), which must be brought into the correct (positive) range in order to be displayed. For example, when using the Stretched CMF envelopes, spectra with high reflectance between 750–1250 nm (such as lawn grass; see Fig. 2) render to an sRGB pixel with negative red. We explain in the next section how negative output values, when they occur, are dealt with. Figs. 4(b) and 5(b) show examples of scenes rendered using this set of envelopes.

A second designed set of envelopes is piecewise linear, which provides more uniform hue variation over the spectrum, and does not produce out-of-gamut pixels. The design is shown in Fig. 1. Figs. 4(c) and 5(c) show examples of scenes rendered using this set of envelopes. We designed a number of other fixed envelopes, with similar but generally less optimal results.

For comparison, Fig. 1 also shows the image-adaptive basis functions corresponding to the first three principal components from PCA performed on the Moffett Field test image shown in

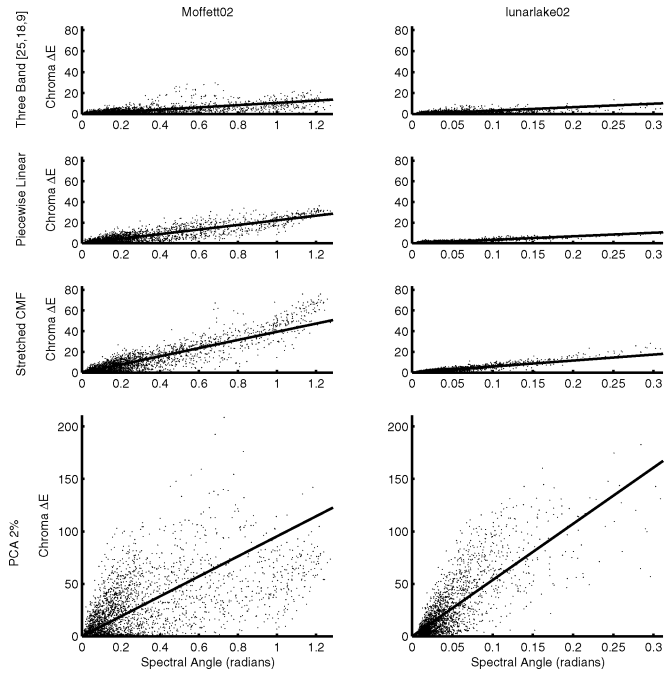

Fig. 3. Scatterplots of spectral angle versus chroma ΔE of the visualization (Euclidean distance in the (a, b) plane of the CIELab color space) are shown for two AVIRIS datasets (one in each column). Each row is a different visualization method. The gray points are different pairs of pixels, and each black line is the least-squares fit to each dataset. Note that the horizontal scale is different for the two scenes, so slopes can only be compared between different methods for visualizing the same scene.

TABLE I
SLOPE m

	Three Band [25,18,9]	Piecewise Linear	Stretched CMF	PCA 2%
Moffett01	13.7	28.9	46.0	195.6
Moffett02	10.7	22.3	39.5	95.3
Moffett03	9.3	21.5	37.7	92.5
Lunar01	27.5	36.3	54.7	653.1
Lunar02	32.8	34.1	58.0	536.3
Cuprite01	33.9	32.9	51.7	822.9
Cuprite02	28.9	34.3	56.2	608.2
Cuprite03	29.4	32.9	52.8	888.2
Cuprite04	29.4	36.3	52.5	647.8
Jasper01	20.4	28.5	42.1	202.0
Jasper02	20.0	28.7	43.6	162.8
Jasper03	22.4	27.6	42.0	183.5
Jasper04	23.0	27.3	42.9	218.7
Jasper05	25.2	28.0	43.9	197.7
Average	23.3	30.0	47.4	393.2

TABLE II
CORRELATION p

	Three Band [25,18,9]	Piecewise Linear	Stretched CMF	PCA 2%
Moffett01	0.69	0.85	0.85	0.78
Moffett02	0.67	0.85	0.89	0.52
Moffett03	0.55	0.86	0.87	0.67
Lunar01	0.26	0.83	0.86	0.82
Lunar02	0.53	0.84	0.92	0.79
Cuprite01	0.42	0.67	0.76	0.63
Cuprite02	0.51	0.81	0.85	0.74
Cuprite03	0.36	0.68	0.82	0.80
Cuprite04	0.46	0.82	0.86	0.85
Jasper01	0.27	0.86	0.86	0.59
Jasper02	0.32	0.88	0.88	0.78
Jasper03	0.25	0.78	0.81	0.70
Jasper04	0.33	0.77	0.81	0.52
Jasper05	0.35	0.85	0.87	0.71
Average	0.43	0.81	0.85	0.71

Fig. 4(a). Code to render AVIRIS images with the fixed spectral envelopes is available.[1]

IV. GAMUT FITTING WITH HUE PRESERVATION

After the fixed integration envelopes are applied, some images may have pixel values that are outside the sRGB gamut (colors that are nondisplayable). In other cases, the entire image may lie in a small subset of the sRGB gamut, for example, the entire image may be very dark. In this section we discuss how the images are better fit to the sRGB gamut so that hue is preserved and/or contrast enhanced. All methods used involve linear transformations of luminance values, and represent a loosening of the *consistent rendering* goal. No other goals are compromised.

First, in the case that some pixel values are out-of-gamut, we transform the colors in the image to the CIELab color space (a device-independent, approximately perceptually uniform color representation [8]), then shift all the pixels' luminance values (the L channel) so that no more than 2% of pixels have negative values or values too large to be rendered. This method was chosen to best satisfy the goal of *consistent rendering*, as it preserves the hue and saturation of each pixel, and maintains the original level of luminance contrast in the image. The value of 2% was selected as an acceptable compromise between bringing pixels into the gamut and preventing the image from becoming too bright. Any remaining out-of-gamut (nondisplayable) color values are clipped to the gamut in sRGB.

For the images shown in this paper, the sRGB values of pixels that remain out of gamut have simply been set to 0 or 255. The optimal solution would be a hue-preserving gamut clipping algorithm. This is nontrivial because the hue-preserving clipping would need to be performed in the CIELab color space, but the gamut limits are defined in the sRGB color space. We hypothesize that the difference gained by using a hue-preserving clipping algorithm, which would affect less than 2% of the pixels of any image, would be relatively small.

For images that did not exceed the displayable gamut, we applied the following luminance enhancement: a linear scaling factor and offset for the L channel in the CIELab space were determined iteratively such that at least one pixel had an sRGB value of 0 and at least one pixel had an sRGB value of 255, without any gamut clipping. This enhancement increases the luminance contrast of the image while retaining consistent rendering of hue and saturation. This should cause luminance edges to appear stronger, and texture to be more distinct.

[1]http://ee.washington.edu/research/guptalab/projects/project4/index.html

V. Other Visualization Methods

The most common way to display a hyperspectral image is to select three of the original N bands and map them to RGB, as done in the image browser for AVIRIS.[2] The three bands can be picked by an adaptive method based on the image content [13], to highlight expected features, or to approximate a conventional color image. If the band-selection is image adaptive, then the meaning of display colors can be drastically different for different images.

A standard method for reducing the dimensionality of a dataset is principal components analysis (PCA), first proposed for this application by Ready and Wintz in 1973 [14]. For an N-band hyperspectral image, the first three principal components are the three N-dimensional, orthonormal basis vectors that capture the most data variance; note that spatial information is not taken into account. The N-band image pixels are linearly projected onto these three N-dimensional basis vectors to create three image bands. The three bands can be mapped to RGB, HSV [15], etc. for display. Figs. 4(a) and 5(a) show examples of PCA visualizations.

PCA has a number of disadvantages for hyperspectral visualization.

1) The meaning of display colors changes for every image, and thus the visualizations can be difficult for users to interpret. Consistent displays allow humans to ascribe consistent meaning to colors, speeding tasks like spotting target objects and understanding scenes.
2) The standard implementation of PCA visualization, as used (by default) in ENVI, is to linearly stretch each of the three display bands so that 2% of pixels are at the minimum and 2% are at the maximum display value. This leads to visualizations with very saturated regions. Although this creates vivid contrasts between pixels in different regions, the contrast between pixels in the same region often suffers as they are crammed into a corner of the perceptual colorspace. This is especially noticeable for pixels near the blue and red primaries.
3) The PCA bands correspond to the maximum data variance, but mapping those three bands to R, G, and B does not transfer to maximum human visual discriminability. There has been some promising recent work [15] on mapping PCA image bands to an opponent color space which may reduce this disadvantage. The amount of variance in each of the principal components will differ between images, so mapping the principal components in a fixed way to the R, G, B channels still will not optimally transfer the information.
4) The highly saturated colors produced by PCA can be preattentive, that is, they "pop-out" at the viewer and draw attention. Since PCA maps colors without a fixed semantic meaning, bright saturated color regions can preattentively attract the viewer's attention without a good reason, and be distracting from the viewer's task [3].
5) PCA has high computational complexity, since it requires finding the eigenvectors of the covariance matrix (size $N \times N$) of the data matrix (size $IJ \times N$), where I, J

[2]aviris.jpl.nasa.gov/html/aviris.quicklooks.html

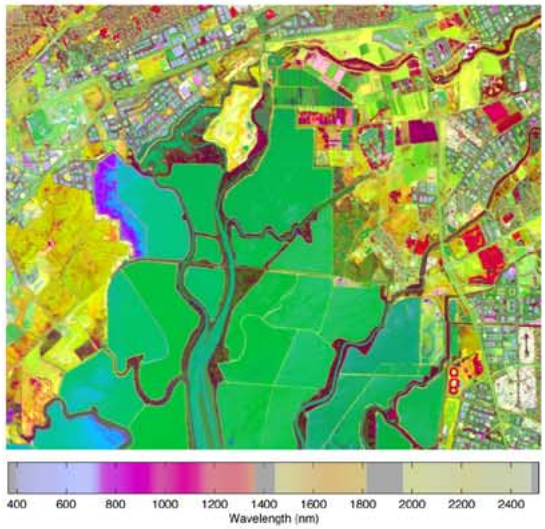

(a) PCA visualization, with display mapping $(P_1, P_2, P_3) \rightarrow (R, G, B)$

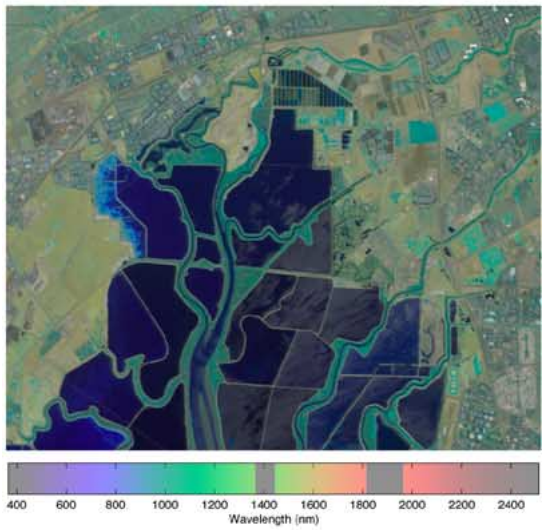

(b) Stretched CMF envelopes visualization

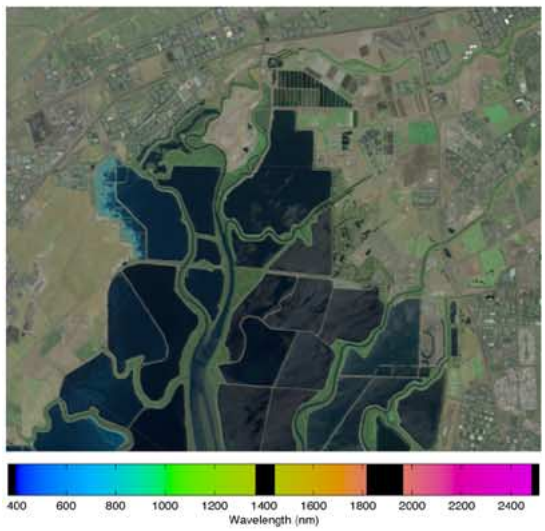

(c) Piecewise Linear envelopes visualization

Fig. 4. Visualizations of Moffett field.

are the image width and height, and N is the number of spectral bands.

Many techniques have been developed for the general problem of image fusion [16]–[18], but few of them are applicable to the display of hyperspectral images. Grayscale image fusion techniques are well-developed for fusing images with different focal lengths and lighting conditions. When applied across spectra, such grayscale fusions remove all spectral information and put an unnecessary limit on the information that can be delivered by a color display.

The fusion of images with high spectral resolution and images with high spatial resolution [19] is a developed field with well-defined quality metrics. However, this class of image fusion techniques do not reduce dimensionality, so they are not directly applicable to hyperspectral visualization.

Our method of reducing the dimensionality of hyperspectral images by spectral weighting envelopes is algorithmically similar to "artificial color" as proposed by Caulfield [20]. In that framework, artificial colors are spectral discriminants that are produced by applying overlapping spectral sensitivity curves to hyperspectral data. Caulfield's artificial colors are proposed as preprocessing for pattern recognition.

VI. Quantitative Metrics of Visualization Fidelity

Creating a useful metric for comparing visualization methods for hyperspectral imaging is difficult. Image fusion methods for some tasks, such as merging images at different focal depths, can compare the fused result with known correct results. However, when dimensionality reduction is involved, there is no "correct" image to compare results with.

We first considered a class of metrics which quantify discriminability in the fused color image. For example, a discriminability metric might count the number of unique discriminable colors in a visualization. Or, a metric might measure the total edge energy in a visualization at multiple resolutions by summing energy in wavelet bands. The disadvantage of this class of metrics is that it does not enforce the reliable presentation of the original hyperspectral data. Arbitrary image transformations, such as adding noise, or extreme (up to 20% saturation) contrast scaling, can improve the performance according to these metrics despite obviously degrading image utility.

A more useful class of metrics require some comparison with the original hyperspectral data. The metric should reflect how well a human can discriminate information in the visualization that was present in the original hyperspectral data.

To this end, a good metric will measure how well the perceptual distance between two data points in a visualization corresponds to a relevant distance between those two points in the original data space. A metric that does this well will be directly useful in gauging how well a visualization will help humans evaluate automated classification, identification, and detection algorithms, and aid humans in making decisions based on such automated algorithms.

We propose that the visual segmentation and classification tasks as performed by a human user are analogous to the Spectral Angle Mapper Algorithm [21], which is an established clas-

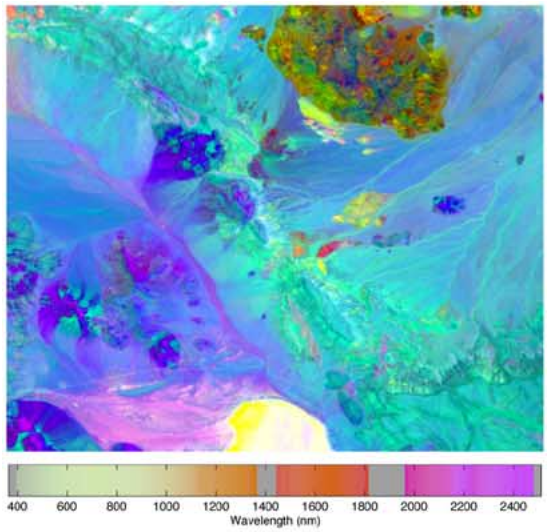

(a) PCA visualization, with display mapping $(P_1, P_2, P_3) \rightarrow (R, G, B)$

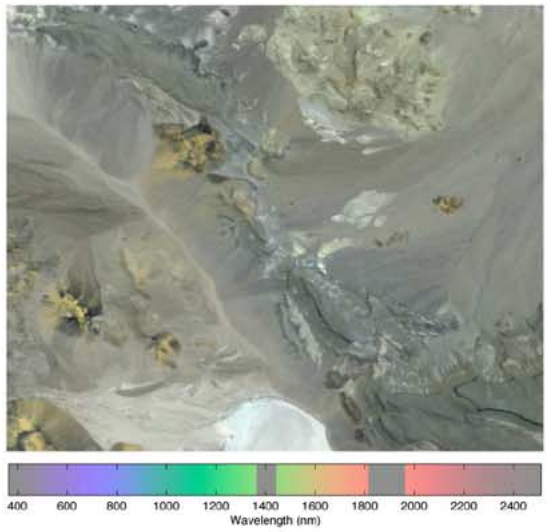

(b) Stretched CMF envelopes visualization

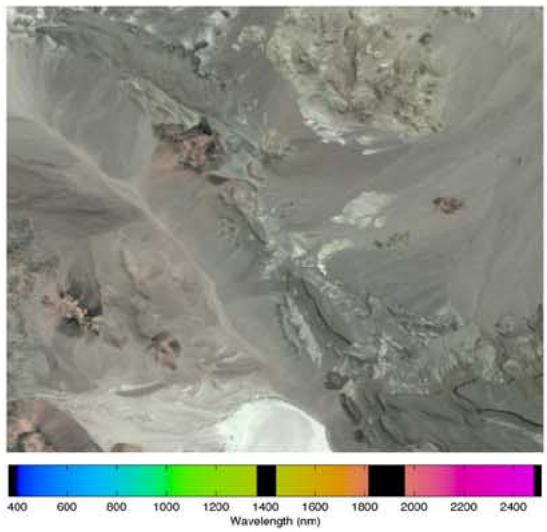

(c) Piecewise Linear envelopes visualization

Fig. 5. Visualizations of Lunar Lake.

sifier for hyperspectral imagery. The spectral angle measures the distance between two hyperspectral vectors x and y

$$\angle xy = \arccos\left(\frac{x \cdot y}{|x||y|}\right). \tag{3}$$

Since spectral angle is normalized, it is insensitive to changes in radiant intensity, which is subject to ground texture and shadows. In a visualization rendered using a spectral weighting envelope as proposed in this work, the spectral angle between hyperspectral pixels should correspond to a distance in hue and saturation between displayed pixels.

The perceptual difference between colors is often measured by Euclidean distance (ΔE) in the CIELab color space. Overall radiant intensity should be proportional to luminance. To obtain a property of the visualization that is analogous to spectral angle, we ignore luminance, measuring Euclidean distance in the chroma (a, b) plane $(\Delta E_{a,b})$. In the (a, b) plane of CIELab, saturation and hue correspond to polar coordinates (r, θ).

We have formulated two metrics based on the relationship between spectral angle and CIELab (a, b) distance for pairs of pixels, and tested these metrics with a set of AVIRIS images. Although luminance is very effective for discriminating shapes and edges, it is not well suited to pixel-level spectral classification due to its dependence on radiant intensity factors.

Our first metric is the correlation ρ of the spectral angle between two data points and their nonluminance color difference in the visualization (the difference between the hue and saturation of the rendered colors).

Computationally, we sample the set of k pixels pairs consisting of all pairs with purely horizontal or purely vertical displacement of $1, 2, 4, \ldots, 128, 512$ pixels. For a 512×614 image, $k = 5\,135\,462$. For each of these pairs of pixels, calculate the spectral angle in the original hyperspectral space, and form a vector X of all of the spectral angles. Then, for each pair of pixels, transform the rendered sRGB coordinates to CIELab, and calculate the Euclidean distance in the chroma (a, b) plane between the two pixels, forming a vector Y with all of these perceptual nonluminance color differences. Then ρ is the correlation between vectors X and Y

$$\rho = \frac{\frac{1}{k}(X^T Y) - \bar{X}\bar{Y}}{(std(X))\,(std(Y))} \tag{4}$$

where X and Y are vectors that contain all the pairs of pixels that fall along the same horizontal or vertical line, $\bar{X}$ is the mean of X, and $std(X)$ is the standard deviation of X.

The correlation ρ measures the effectiveness of using perceived hue and saturation distances to predict differences in spectral angle. In the absence of perceptual errors, a perfect correlation of $\rho = 1$ would give a human user the same classification-relevant information that is used by a computer running the SAM algorithm on the original hyperspectral data.

A good correlation between chroma distance and spectral angle is only useful if the chroma distance is large enough that it can be accurately perceived and compared with the chroma distance between another pair of pixels. Thus, our second metric is the average ratio m, between chroma (a, b) distance and spectral angle, which is the slope of the line that is anchored at the origin and is the least-squares fit to the data pairs (X, Y). Spectral angle and chroma distance could be highly correlated $(\rho = 1)$ but if the linear regression slope m is too small, then different chroma differences would not actually be perceptually distinguishable.

According to these metrics, the Stretched CMF envelopes generally performed best of all the fixed spectral envelopes designed. Results comparing the Stretched CMF envelopes to Three Bands and PCA are shown in Fig. 3 and given in Tables I and II. In general, it is seen that the slopes m are higher for PCA images, corresponding to larger color contrasts with change in spectral angle, but the correlation is as high or higher for the proposed Stretched CMF envelopes, corresponding to a stronger correlation between changes in spectral angle and perceived change in hue and saturation in the visualization. PCA is not designed to capture information about spectral angles directly, and thus the spread of chroma distance for a particular spectral angle difference is not surprising.

VII. Resulting Images

To illustrate our proposed visualization technique, we used hyperspectral images from the National Aeronautics and Space Administration Jet Propulsion Laboratory AVIRIS system, which captures 224 spectral bands, ranging from 400–2500 nm, with nominal bandwidth of 10 nm [22]. We used the AVIRIS reflectance data which is corrected to compensate for atmospheric absorption and the spectrum of the sun. The reflectance data can be used to match the spectrum of an image pixel to the spectrum of a known material. We compare results for several images from the 1997 flight lines of Moffett Field CA, Jasper Ridge CA, Cuprite NV, and Lunar Lake NV.

The image bands 107–114 (1353–1423 nm), 153–168 (1811–1948 nm), and 222–224 (2486–2506 nm) were not used, as they were determined (by inspection) to contain significant noise. This observation is supported by the signal-to-noise curve from the 1997 AVIRIS calibration given in [23, Fig. 9]. In addition, bands 33, 34, 97, and 98, which occupy regions of overlapping spectral sensitivity, were not used.

Figs. 4 and 5 show visualizations using PCA, the Stretched CMF envelope, and Piecewise Linear envelopes (both plotted in Fig. 1). Colorbars show the rendered color of monochromatic light at each wavelength. As is done in RSI's ENVI software, each PCA image band has been linearly stretched so that 2% of pixels are saturated.

For both proposed integration envelopes, the spectrally weighted colors not only carry consistent meaning, but also have specific and intuitive semantic interpretations for materials in an image. As shown in Fig. 2, vegetation spectra reflects strongest in the range of 650–1400 nm [24], and the Stretched CMF envelope maps vegetation to green (this is also true for the Piecewise Linear envelopes). Water has the strongest reflectance for wavelengths between 400–600 nm [24] and appears blue when visualized with both envelope sets. Because of the *consistent rendering* property, these colors provide classification cues that are independent of region shape.

In the Moffett field scene shown in Fig. 4, the main features of water, live vegetation, dry vegetation and buildings can be identified in each visualization.

In the PCA visualization [Fig. 4(a)] the areas rendered red "pop out" strongly. In order to classify the red regions a viewer must rely on the shape of the red regions. Since the red regions appear in shapes that one recognizes as fields, baseball diamonds, and golf courses, one hypothesizes that red regions are vegetation in this image. However, shape cues can be misleading, since the upper right of the image has a few red river-like features that might lead one to believe that red corresponds to water in this image, although the "green" material has more water-like shapes. Shape can also be used to identify buildings and other man-made features in the image. In short, since the color associations in PCA visualizations are not consistent between scenes, we have to rely on shape cues alone to mentally classify materials.

In the stretched CMF visualization [Fig. 4(b)], the colors are more natural looking, and no colors "pop out" strongly. Mentally classifying regions is aided by region color. The green regions have strong reflectance around 1100 nm, which is consistent with healthy vegetation such as lawn grass (see Fig. 2). The blue regions are fairly dark, and their hue suggests a dominant reflectance around 700 nm. This is consistent with the reflectance spectrum of water. The blue regions in the lower middle of the image show a brighter gray material that occurs along with the water, a feature that is nearly invisible in the PCA visualization. Our best classification cue here is shape and general knowledge, which would label it as mud. The green-yellow regions appear similar to the reflectance spectrum of dry grass. A guess that the small rectangular shapes are buildings is supported by their gray to white colors, consistent with concrete or shiny metal roofing.

In the Lunar Lake scene shown in Fig. 5, there is no open water and little vegetation, so the identifiable features are rock and soil types. As shown by the range of the spectral angle between hyperspectral pixels in the two scenes (see Fig. 3) the materials in this scene are less varied.

Generally, features are better segmented using the PCA images than the fixed envelope visualizations. Conversely, there are features such the texture of the evaporation ponds in the Moffett Field image that are clear in the fixed envelope visualization but not in the PCA image. Due to the *consistent rendering* property, the fixed envelope visualizations have the added advantage of using hue and saturation as classification cues, in addition to shape cues. Also, due to the *wavelength shift invariance* property, average reflectivity can be gauged from the brightness of the fixed envelope visualizations, which is not possible with the PCA visualizations.

As discussed in item 3) of Section V, PCA finds a subspace that maximize the variation of the data, where the image data is treated as a vector and spatial relationships between data points are not taken into account beyond the notion that all the data points are from one image. Hence, PCA should not be expected to preserve edges in the visualization.

VIII. Discussion

In this paper, we presented design goals and solutions for the display of fused hyperspectral imagery on tristimulus devices, without reference to a specific task or application. Our approach uses a set of fixed linear basis functions to create a consistent visualization. We have demonstrated two sets of such basis functions with arguably beneficial properties. Their success at meeting the design goals is given below.

Both spectral weighting envelope sets give *consistent rendering* of hue and saturation, with minor luminance scaling used to take full advantage of the dynamic range of the display. As seen in the Lunar Lake figure (Fig. 5), the consistent renderings cover a much smaller volume of the colorspace than the PCA rendering. One could argue that this is a suboptimal utilization of the full colorspace. However, as advocated by Robertson and O'Callaghan [10, pg. 56], scene understanding is often best enhanced by the "ability to associate perceptually meaningful color attributes." That is, one knows instantly what is and what is not in the scene by the colors rendered. Additionally, the sparsity of bright saturated colors in the proposed renderings allows bright saturated colors to be used for marking important items in the image. A disadvantage to requiring a visualization method to use a maximum number of colors is that in cases where there is not a large amount of differentiable information in the image the visualization will be very sensitive to noise. Lastly, targets that should stand out as substantially different from a mildly varying background can be "lost" in a visualization that maximizes the number of colors used.

The stretched CMF envelopes have better *edge preservation*, as measured by the relationship between chroma distance and spectral angle shown in Fig. 3 and Table I. It should be noted that PCA creates the strongest edges, although the edges are not as highly correlated with spectral angle. Spectral weighting envelope visualizations are *computed quickly*, requiring only one multiplication and addition per channel per pixel. Both spectral weighting envelope sets have an *equal-energy white point*, so that color saturation in a rendered pixel indicates how nonuniform its reflectance spectrum is. Both spectral weighting envelopes are designed to use *smallest effective differences* and *appropriate preattentive features* by using minimal contrast stretching, which is restricted to linear transformations of luminance. A *natural palette* is created by using spectral envelopes with few inflection points. An equal-energy white point also helps, since it promotes a balance of colors. *Wavelength shift invariance* is not quite achieved by the current designs; it remains an option for future work.

IX. Conclusion

The proposed visualization goals are useful for designing, evaluating, and using scientific visualizations. We have shown that fixed spectral envelopes can satisfy many goals, including consistency, while still displaying edge information which correlates with spectral angles.

The proposed visualizations can be used separately as a tool to interpret and contrast scenes, or can be used in conjunction with task-specific visualizations. For example, these visualizations could be used as a base image upon which to mark where particular endmembers are found. The desaturated nature of the proposed visualizations leaves room in the color gamut to use bright saturated colors for such markings, or to point out locations of interest.

REFERENCES

[1] M. Gladwell, "The picture problem," *The New Yorker*, pp. 74–81, Dec. 2004.

[2] E. R. Tufte, *Visual Explanations*. Cheshire, CT: Graphics, 1997.

[3] C. G. Healey, K. S. Booth, and J. T. Enns, "Visualizing real-time multivariate data using preattentive processing," *ACM Trans. Model. Comput. Simul.*, vol. 5, no. 3, pp. 190–221, Jul. 1995.

[4] E. R. Tufte, *Envisioning Information*. Cheshire, CT: Graphics, 1990.

[5] E. Imhof, *Kartographische Geländedarstellung*. Berlin, Germany: De Gruyter, 1965.

[6] G. M. Johnson and M. D. Fairchild, *Digital Color Imaging*. Boca Raton, FL: CRC, 2003, ch. 2, pp. 141–148.

[7] M. Fairchild, *Color Appearance Models*. Reading, MA: Addison-Wesley, 1998.

[8] M. Stone, *A Field Guide to Digital Color*. Wellesley, MA: AK Peters, 2003.

[9] J. S. Tyo, E. N. Pugh, Jr., and N. Engheta, "Colorimetric representations for use with polarization-difference imaging of objects in scattering media," *J. Opt. Soc. Amer. A*, vol. 15, no. 2, 1998.

[10] P. K. Robertson and J. F. O'Callahan, "The application of perceptual color spaces to the display of remotely sensed imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 26, no. 1, pp. 49–59, Jan. 1998.

[11] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. New York: Wiley, 2000.

[12] G. Polder and G. W. A. N. van der Heijden, "Visualization of spectral images," presented at the *SPIE Conf. Visualization and Optimization Techniques*, 2001.

[13] H. Du, H. Qi, X. Wang, R. Ramanath, and W. E. Snyder, "Band selection using independent component analysis for hyperspectral image processing," in *Proc. 32nd Applied Imagery Pattern Recognition Workshop*, 2003, pp. 93–98.

[14] P. J. Ready and P. A. Wintz, "Information extraction, SNR improvement, and data compression in multispectral imagery," *IEEE Trans. Commun.*, vol. 21, no. 10, pp. 1123–1131, 1973.

[15] J. S. Tyo, A. Konsolakis, D. I. Diersen, and R. C. Olsen, "Principal-components-based display strategy for spectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 3, pp. 708–718, Mar. 2003.

[16] S. G. Nikolov, P. R. Hill, D. R. Bull, and C. N. Canagarajah, *Wavelets in Signal and Image Analysis*. Dordrecht, The Netherlands: Kluwer, 2001, vol. 19, Computational Imaging and Vision, pp. 213–244.

[17] T. A. Wilson, S. K. Rogers, and M. Kabrisky, "Perceptual-based image fusion for hyperspectral data," *IEEE Trans. Geosci. Remote Sens.*, vol. 35, no. 4, pp. 1007–1017, Jul. 1997.

[18] C. Pohl and J. L. van Genderen, "Multisensor image fusion in remote sensing: Concepts, methods and applications," *Int. J Remote Sens.*, vol. 19, no. 5, pp. 823–854, 1998.

[19] T. Ranchin, B. Aiazzi, L. Alparone, S. Baronti, and L. Wald, "Image fusion: The ARSIS concept and some successful implementation schemes," *ISPRS J. Photogramm. Remote Sens.*, 2003.

[20] *Proc. SPIE Conf. Nature's Alternative to Hyperspectral Imaging and Why Nature Is Right*, vol. 4787, 2002, pp. 132–136.

[21] F. A. Kruse, A. B. Letkoff, J. W. Boardman, K. B. Heidebrecht, A. T. Shapiro, P. J. Barloon, and A. F. H. Goetz, "The Spectral Image Processing System (SIPS)—Interactive visualization and analysis of imaging spectrometer data," *Remote Sens. Environ.*, vol. 44, pp. 145–163, May–Jun. 1993.

[22] G. Vane, R. O. Green, T. G. Chrien, H. T. Enmark, E. G. Hansen, and W. M. Porter, "The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sens. Environ.*, vol. 44, pp. 127–143, Jun. 1993.

[23] R. O. Green *et al.*, "Imaging spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sens. Environ.*, vol. 65, pp. 227–248, 1998.

[24] R. N. Clark, G. A. Swayze, R. Wise, K. E. Livo, T. M. Hoefen, R. F. Kokaly, and S. J. Sutley, "USGS Digital Spectral Library splib05a,", USGS Open-File Rep. 03-95, Ver. 1.0, Sep. 2003. [Online]. Available: http://pubs.usgs.gov/of/2003/ofr-03-395.

Nathaniel P. Jacobson (S'05) received the B.S. degree in electrical engineering from the University of Washington, Seattle, in 2004. He is currently pursuing the M.S.E.E. degree at the University of Washington.

His research interests include the visualization of high-dimensional data and modulation filtering.

Maya R. Gupta (M'02) received the B.S. degree in electrical engineering and the B.A. degree in economics from Rice University, Houston, TX, in 1997, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1999 and 2003, respectively.

Since 2003, she has been a Professor in the Department of Electrical Engineering, University of Washington, Seattle. She worked from 1999 to 2003 for the Color Image Processing Group of Ricoh's California Research Center and has also worked for AT&T Labs, NATO, Hewlett Packard, and Microsoft. Her research interests are in color processing and nonparametric statistical learning.