Spectral Based Color Image Editing (SBCIE)

Joyce Farrell
Hewlett-Packard Laboratories

Abstract

Spectral based color image editing systems include methods for *generating* spectral representations of surfaces and illuminants from input signals of image capture devices (such as digital cameras and scanners), *editing* spectral representations of surfaces and illuminants in a scene, and *transforming* spectral representations of surfaces and illuminants to output signals of image devices (such as CRT and LCD displays and color printers).

Introduction

Computer-assisted color image editing systems include methods for image capture, manipulation and reproduction. Until recently, color image editing systems were commercially available only as part of large graphics systems. Today, such systems are widely available in desktop publishing software or embedded in drivers for digital cameras, scanners and printers, providing users with the means to 1) generate digital representations of color images, 2) edit digital representations to make the image look visually appealing, and 3) reproduce digital representations on different media. Whether they are part of an elaborate publishing system or a software driver for imaging peripherals, all color image editing systems are dependent on and limited by devices that either capture or generate color images.

Most image capture devices have only three color sensors and therefore compress a multidimensional spectral signal into the output of three color channels -- usually, R, G and B. Most color image editing systems convert the device RGB signals into a human-based representation, such as CIE XYZ, and operate on this representation. For example, Neugebauer [1] and Schreiber [2] described color image editing systems based on converting device-dependent image representations, such as scanner or display RGB values, into human-based tristimulus signals, such as XYZ. These color image editing systems provide a means of changing and displaying the human-based tristimulus signals to enable interactive editorial color corrections. The objective of these systems is to visualize or simulate enhancements to the printed output before the actual image is printed.

This paper presents a different approach to color image editing. Rather than convert device-dependent signals or vectors into human-based vectors based on the CIE XYZ color matching functions, as is traditionally performed in today's color image editing systems, it is possible to use device-dependent sensor responses to estimate device-independent spectral reflectances of surfaces and illuminants in the captured and/or rendered scene. To illustrate this approach, I describe several methods for generating spectral representations of surfaces and illuminants in a scene. I introduce the concept of spectral-based color image editing (SBCIE) systems that embody methods to generate, edit and reproduce spectral representations of surfaces and illuminants based on the analysis of digital images generated by scanners, digital cameras and displays.

Background

In this section, I introduce the linear algebraic notation for spectral representations of illuminants and surfaces used throughout this paper. Then I describe the more practical low-dimensional linear spectral representations that can be used to approximate spectral representations. These low-dimensional linear models are necessary because most often we do not have complete spectral information about the illuminants and surfaces in a scene. I describe how the accuracy of spectral representations are limited by the dimensionality of our image capture data and by our apriori knowledge of scene illumination.

The spectral power distribution of illuminants and the spectral reflectances of surfaces are represented as functions of wavelengths. For example, the spectral power distribution (SPD) of an illuminant can be described by a single vector, **e**, with *n* entries representing the amount of energy emitted over a range of wavelengths (e.g. n = 81 when the wavelengths range from 380 nm to 780 nm in 5 nm steps).

The spectral radiance factor of a surface [4] is the wavelength composition of the light reflected and/or emitted from the surface. The spectral radiance factor of a surface can be represented by a *nxn* matrix, **S**. If the surface is diffuse and does not fluoresce, **S** has values between 0 and 1 along the diagonal and no values in off-diagonal positions in the matrix. (A glossy or specular surface may have values that exceed 1.0.) When surfaces do not fluoresce, their spectral radiance factor can be represented by a vector, **s**, corresponding to the diagonal component of **S**. A fluorescent surface, however, will absorb light in one wavelength and emit light in a longer wavelength. Thus the spectral radiance factor of fluorescent surfaces cannot be described by a single vector because it will have entries in the off-diagonal positions in the matrix **S**. The complete characterization of the spectral radiance factor of a surface that has both diffuse and fluorescent component requires the full *nxn* matrix **S**.

Low-dimensional linear spectral representations of illuminants and surfaces

The spectral representations of many illuminants, **e**, can be approximated by a linear combination of a smaller set of spectral basis functions, $\mathbf{B}_i$[5]:

$$e \approx \sum_{i=1}^{N} wiBi$$

where w_i are the weights chosen to minimize the error between the illuminant SPD and its linear model approximation, and N is the dimensionality of the spectral representation.

The spectral representations of diffuse surfaces, **s**, can also be described by a linear combination of a smaller set of spectral basis functions [6-9].

$$s \approx \sum_{i=1}^{N} wiBi$$

These low-dimensional linear models of illuminant spectra and surface reflectances are obviously more efficient and serve to reduce the amount of data that must be stored to reconstruct the spectral representations. Moreover, we often generate the spectral representations of illuminants and surfaces from low-dimensional spectral data, such as camera and scanner RGB values. In this case, the dimensionality of the image capture data will limit the dimensionality of our spectral reflectances for reconstruction purposes.

The accuracy of our spectral representations of surfaces and illuminants will be limited by 1) the inherent dimensionality of the illuminant spectra and surface reflectances, 2) the dimensionality of the image capture data, and 3) our apriori knowledge about surfaces and illuminant spectra. If we have enough spectral channels in our image capture device and knowledge of the illuminant, we can build an accurate representation of surfaces and illuminants. In practice, most applications do not have enough sensor data and apriori knowledge to build complete representations of surfaces and illuminants. In these cases, we are forced to make low-dimensional estimates of the surface reflectance matrices and illuminant vectors [3, 10-12] which approximate the complete representations.

Spectral-Based Color Image Editing Systems

Spectral-based color image editing (SBCIE) systems are based on methods for generating, editing and reproducing spectral representations of surfaces and illuminants. To illustrate these methods, I give several examples of how to generate spectral representations of surfaces and illuminants based on the analysis of digital representations of scenes generated by scanners, digital cameras and displays.

Scanners

It is possible to use scanner RGB values to estimate the first three principal components or basis functions of surface spectral reflectances. Rather than map the scanner RGB values into tristimulus values (as traditional colorimetric approaches have done [13,14]), the scanner RGB values are used to estimate the weight factors of three reflectance basis functions. The estimated weights and the corresponding spectral basis functions are then used to build three-dimensional linear representations of surface spectral reflectances.

To illustrate how to generate spectral representations of scanned surfaces from scanner data, I introduce the following notation. Let **R** be a *3xM* matrix of scanner responses to *M* spectral surfaces, **T** be a *3xn* matrix describing the spectral responsivities of a three-channel scanner where *n* defines the range of wavelength samples, **B** be a *nx3* matrix defining three spectral basis functions for surface reflectances and **W** be a *3xM* matrix of basis weights.

Then,

$$\mathbf{R} = \mathbf{TBW}$$

Note that **T** and **B** can be combined to form a *3x3* matrix, **C**, and **W** can then be solved by the regression equation:

$$\mathbf{W} = \mathbf{C}^{-1}\,\mathbf{R}$$

Having solved for **W**, our estimates of the spectral reflectances of the scanned surfaces are calculated by:

$$\mathbf{S} = \mathbf{BW}$$

This example is based on 3-channel output - the more the channels , the better one will do. For example, we get significant improvement with increasing to 4 spectral channels[15,16]. We can also, improve our estimates of the surface reflectances by judicious selection of spectral basis functions. For a review of methods for selecting the appropriate basis functions, see Sherman and Farrell [17]. For a description of a method for increasing the number of channels in a scanner see Farrell, Sherman and Wandell [18].

Digital Cameras

If we know both the spectral sensitivities of the color sensors in a digital camera and the spectral power distribution (SPD) of the illuminant, we can use the method described above to estimate the spectral reflectances of the surfaces in the scene captured by the digital camera:

Let **R** be a *3xM* matrix of camera responses to *M* spectral surfaces, **T** be a *3xn* matrix describing the spectral responsivities of a three-channel camera (illuminant included), and **B** be a *nx3* matrix defining three spectral basis functions for surface reflectances. We then solve for the weights on the spectral basis functions, **W**, and estimate the spectral reflectances, **S**, by **BW**.

There are many instances in which the illuminant SPD is not known, however. In fact, it is most often the case that we do not know what part of the color signal recorded by a digital camera reflects the scene illuminants and what part reflects the surface reflectances. If we knew the illuminant, we could solve for the surfaces [10]. If we knew the surfaces, we could solve for the illuminant [3]. When we cannot measure the illuminant SPD, we are forced to estimate it from the distribution of color pixel values in the captured image.

Illuminant estimation is an important problem that is beyond the scope of this paper. I refer the reader to the "subspace algorithm" by Maloney and Wandell [11] and the "gray world algorithm" by Buchsbaum [10]. Other algorithms which are based on formal statistical theories include the "maximum likelihood estimation algorithm" and the "covariance matching algorithm" by Trussell and Vrhel [22], and the "Bayesian color-constancy algorithm" by Brainard and Freeman [24]. All these methods rely in one way or another on restricting the dimensionality of linear vector space for representing the illuminant spd. The performance of these algorithms for estimating the illuminant SPD depends greatly on the number of classes of color sensors. The accuracy of the illuminant SPD estimations increases with the number of color sensor classes.

Displays

Color images of scenes are often rendered on emissive (e.g. CRT) or reflective (e.g. LCD) displays. Even though we may not know how the images were generated, we nonetheless have the perception that the displayed images are realistic depictions of actual scenes. In fact, one of the key areas in computer graphics is devoted to developing realistic spectral representations of surfaces and iluminants to be rendered on the display. If we begin with a known spectral representation, the rendering is straightforward and our work is done. The real problem is in generating the spectral representations of surfaces and illuminants, either from sensor data (as illustrated above) or from rendered images.

Let's assume that we know nothing about how a displayed image was generated but nonetheless wish to generate a spectral representation of the surfaces and illuminants that is consistent with our perception of the scene. One method for doing this is to have the operator select a region of the scene that corresponds to a white surface. Since a white surface has known spectral reflectance, we can estimate an illuminant spectral power distribution that would be consistent with the displayed tristimulus values for the white surface. Let **Y** be a *3x1* vector containing the displayed tristimulus values for the rendered white surface. Let **R** be a *3xN* matrix containing the linear RGB values for the white surface. And let **T** be a *3x3* matrix that maps **R** into **Y.**

$$\mathbf{Y = TR}$$

T is determined by the multiplication of two matrices: a *3xn* matrix containing the CIE XYZ color matching functions, **H**, and a *nx3* matrix containing the spectral power distribution of the three display phosphors, **P**.

$$\mathbf{T = HP}$$

Our first task is to find an illuminant that when reflected from a white (Lambertian) surface would generate **X**. In other words, we wish to solve for the illuminant spd, **E**, where the surface reflectance (**S**), human color matching functions (**H**), and tristimulus values (**Y**) are known:

$$\mathbf{Y = HSE}$$

The method of solving for **E**, given **S**, **H** and **Y** was first introduced by Buchsbaum [10]. In his algorithm, however, he used what is commonly referred to as the "grayworld assumption". This assumption states that the sensor response to a gray surface can be approximated by the sensor responses averaged over an entire image or scene. In this example, the average sensor response corresponds to the tristimulus values, **X**, averaged across the entire image. It is not difficult to demonstrate that this assumption does not hold for many images. And when this assumption is not valid, this method for color correction will fail. When the assumption is valid, however, the method works quite well. Thus, by asking observers to identify a known surface (be it white, gray, or purple) we can circumvent the grayworld assumption and use the tristimulus values for this known surface, **Y**, to solve for **E**.

When **E** is generated by daylight illumination, we can solve for **E** given **Y**, **H**, **S** and **D** where **Y** is a *3x1* vector containing the tristimulus values for a known surface, **S**, and **D** is a *nx3* matrix containing three spectral basis functions for daylight [5]. Equation 8 can be rewritten as

$$\mathbf{Y = HSDF}$$

where **F** is a *3x1* vector containing the eigenvalues or weights for **D**. Since **E** is estimated by **DF**, and **D** is known, our computational task is to estimate **F**. To simplify the calculation, combine **H**, **S** and **D** into a *3x3* matrix, **G**. Then, we estimate **F** by

$$\mathbf{F} = \mathbf{G}^{-1}\,\mathbf{Y}$$

Scene illuminants are often generated by some combination of fluorescent, tungsten and daylight illuminations, however, and we need more than three spectral basis functions to describe these more complex illuminations. One way to solve for **E** in this case is to provide users with a database of illuminant SPDs and to iteratively search through the database to find **E** that minimizes the difference between the predicted tristimulus values**, Y'**, and the actual displayed tristimulus values, **Y.**

Now, having estimated **E**, our task is to estimate the spectral reflectances of surfaces, **S**, that are consistent with **E** and **X**, where **X** is a *3xM* matrix containing the tristimulus values for all surfaces depicted in the scene. We select three spectral basis functions, **B**, with which to represent **S**, and use the three tristimulus values for each surface to estimate the weights, **W** for **B**. Again, we can simplify the calculation by combining **H**, **E** and **B** into a *3x3* matrix, **A**. Then, we estimate **W** by:

$$\mathbf{W} = \mathbf{A}^{-1}\,\mathbf{X}$$

Again, as in the previous examples, the spectral reflectances of the *M* surfaces are estimated by:

$$\mathbf{S} = \mathbf{BW}$$

Transforming and rendering spectral representations

Once we have generated the spectral representations, we can transform (edit) and render (print or display) them. The subsequent transformation and rendering of spectral representations of surfaces and illuminants in a scene completes the color reproduction system.

All transformations in SBCIE are essentially either changes in the entries of surface matrices, **S**, and/or illuminant vectors, **e**. This can be achieved most simply by providing a database of surface matrices and illuminant vectors from which the user can select. Alternatively, based on low-dimensional linear models, we can change the weights and basis vectors that are used to generate the surface matrices and illuminant vectors (see the discussion of *Low-dimensional linear models* above)

Surface transformations are the easiest to comprehend and visualize. We often want to change the skin tone of a person, the color of a dress, the saturation of grass, the color of drapes or paint, and so on. Illuminant transformations are difficult to visualize but are nonetheless powerful tools in SBCIE systems. For example, the perceived color balance in an image can be altered by rendering the surfaces under different illuminants [23]. Display operators can perform such illuminant transformations until the image appears to look visually appealing.

Other applications include image compositing and splining. Here the goal is to render different surfaces under the same illuminant. This is a challenge for any color image editing system that attempts to merge images based on surfaces from the same physical object but captured or rendered under different illuminants.

Once we have spectral reflectances of surfaces and illuminants, it is a straight forward process to render these representations on calibrated display devices.[19,20]. To render spectral representations on printers, one typically uses a device calibration look-up table (LUT) [21].

Conclusions

In this paper, I introduce the concept of spectral-based color image editing (SBCIE) systems as a general framework for image correction, composition, and enhancement. These systems are feasible because 1) surface reflectances are low-dimensional, 2) illuminants are often known or measured, and 3) image capture devices are linear. SBCIE systems enable operators to correct images that were captured under bad illumination, combine images captured under different illumination, select illuminants that make the surfaces in an image look more visually appealing, and create realistic effects simulating how a scene will look like under a desired lighting. Because the image manipulations and adjustments correspond to changes that we are all familiar with, such as changes in lighting and surface properties, the image manipulations are intuitive to the user.

Acknowlegements

I thank Brian Wandell and Doron Sherman for many helpful discussions and insights about linear spectral models and for their advice and support for SBCIE systems.

References

[1] H. E. J. Neugebauer, Color correction selector, *US Patent 2790844*, 1957.

[2] W. F. Schreiber, Color reproduction system, *US Patent 4500919*,1985.

[3] J. E. Farrell and B. A. Wandell, Rethinking the White Point, *IS&T and SID's Color Imaging Conference*, pp. 65-67, 1993.

[4] G. Wyszecki and W. S. Stiles, *Color Science*, New York: John Wiley and Sons, second edition , 1982.

[5] D. B. Judd, D. L. MacAdam and G. W. Wyszecki, Spectral distribution of typical daylight as a function of correlated color temperature, *Journal of the Optical Society of America,* Vol. 54, p. 1031, 1964.

[6] J. Cohen, Dependency of the spectral reflectance curves of the munsell color chips, *Psychonomic Science*, Vol 1, pp. 369-370, 1964.

[7] L. T. Maloney, Evaluation of linear models of surface spectral reflectance with small numbers of parameters, *Journal of the Optical Society of America A*, Vol. 3, pp. 1673-1683, 1986.

[8] B. A. Wandell and D. H. Brainard, Towards cross-media color reproduction, T*opical Meeting on Applied Vision, 1989 Technical Digest Series*, Vol. 16, pp. 132-137 (Optical Society of America, Washington, D. C. 1989).

[9] J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, Characteristic spectra of munsell colors, *Journal of the Optical Society of America,* Vol. 6, pp. 318-322, 1989 .

[10] G. Buchsbaum,, A spatial processor model for object color perception, *Journal of the Franklin Institute,* Vol. 310, pp. 1-26, 1980.

[11] L. T. Maloney and B. A. Wandell, Color constancy: a method for recovering surface spectral reflectance, *Journal of the Optical Society of America: Part A*, Vol. 1, pp. 29-33, 1986.

[12] D. Marimont and B. Wandell, Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A*, Vol. 11, pp. 1905-1913, 1992.

[13] P. C. Hung. Colorimetric calibration for scanners and media, *Proceedings of SPIE*, Vol. 1448, pp. 164-174, 1991.

[14] P. G. Roetling, J. E. Stinehour, and M. S. Maltz, Color characterization of a scanner, in *Proceedings of IS&T's Seventh International Congress on Advances in Non-Impact Printing Technologies*, Vol. 1, pp. 433-440, 1991.

[15] B. A. Wandell and J. E. Farrell, Water into Wine: Converting Scanner RGB to Tristimulus XYZ, *Proceedings of SPIE* Vol. 1909, pp. 92-100.

[16] M. J. Vrhel and H. J. Trussell, Filter considerations in color correction, in *SPIE Symposium on Electronic Imaging*, San Jose, CA, February 1992.

[17] D. Sherman and J. E. Farrell, When to use linear models for color calibration, *Proceedings of IS&T and SID's Second Color Imaging Conference*, pp. 33-36. 1994.

[18] J. E. Farrell, D. Sherman, B. Wandell, How to turn your scanner into a colorimeter, *Proceedings of IS&T's Tenth International Congress on Advances in Non-Impact Printing Technolgies*, pp. 579-581, 1994.

[19] B. A. Wandell, The synthesis and analysis of color images, *IEEE Transactions on Pattern Analysis and Image Processing*, Vol. PAMI-9, pp. 2-13, 1987.

[20] D. H. Brainard, Calibration of a computer controlled color monitor, *Color Research and Application*, Vol. 14, No. 1, pp. 23-34, 1989.

[21] P. C. Hung, Color rendition three-dimensional interpolation. In *Proceedings of the SPIE*, Vol 900, pp. 111-115, 1988.

[22] H. J. Trussell and M. J. Vrhel, Estimation of illumination for color correction. In Proceedings of ICASSP, pp. 2513-2516, 1991.

[23] M. J. Vrhel and H. J. Trussell, Color correction using principal components, Color Research and Application, Vol. 17, pp. 329-337, 1992.

[24] D. H. Brainard and W. T. Freeman, Bayesian method for recovering surface and illuminant properties from photosensor responses. In Proceedings of the IS&T/SPIE Symposium of Electronic Imaging Science & Technology, San Jose, CA, 1994.