Linear Models and Computational Color Constancy

Brian Funt
School of Computing Science, Simon Fraser University, Vancouver, Canada

Abstract

Linear models appear in many recent color constancy theories; however, they can play two quite different roles. They may either occur as a direct component of the computational strategy, as in the case of the Maloney-Wandell algorithm; or they may appear as necessary part of the theoretical development, but not as part of the computation perse, as in the case of algorithms based on spectral sharpening (Finlayson, Drew, Funt). This paper surveys recent work in color constancy and concludes that in many circumstances the second role may be more appropriate than the first.

Introduction

For us to build machines that reproduce colors accurately or to make effective use of color in robotics requires that we understand human color perception; and the last decade has produced many interesting new computational theories of color coming from both computer science and psychology. A central concern of these theories is to describe how color depends or does not depend on the incident illumination. A colored surface cannot be seen unless we shine some light on it, but then the spectrum of the reflected light depends on the product of the spectrum of the incident light's spectrum and the surface's reflectance. Since the spectrum of the light energy reaching the eye has the two factors of illumination and reflectance confounded into one the effect of the illumination must be taken into account in order to determine the true surface properties.

When light of *spectral power distribution* (*spectrum* for short) $E(\lambda)$ reflects off a matte surface of *percent surface spectral reflectance* (*reflectance* for short) $S(\lambda)$ it produces a *color signal* $C(\lambda)$,

$$C(\lambda) = E(\lambda) \cdot S(\lambda) \tag{1}$$

where the product is formed by multiplication on a wavelength-by-wavelength basis over the visible wavelength range of 400-700 nanometers. Each channel of spectral sensitivity $R_k(\lambda)$ $(k=1...3)$

$$\rho k = \int_{visible} C(\lambda) \cdot R_k(\lambda) d\lambda \tag{2}$$

Since the visual system at any given location in the visual field makes only 3 measurements of the full color signal, there is a substantial amount of data reduction and possible data loss involved in Equation 2. In contrast, even a coarse 10nm sampling of the spectrum over the 400-700nm range results in 31 measurements. If the illumination $E(\lambda)$ were known and the incoming color signal $C(\lambda)$ were measured in full, then equation (1) could be directly solved for the surface reflectance properties $S(\lambda)$. However, the visual system clearly faces a difficult problem in recovering any reliable illumination-independent information about $S(\lambda)$ since $E(\lambda)$ is unknown and $C(\lambda)$ is so crudely reduced to only 3 values.

Computational Color Constancy

The various computational color constancy algorithms can be divided along the dimensions: statistical assumptions about the distribution of surface colors, assumptions about reflectances and illuminants, image gamut assumptions, peeking methods, requirements about multiple illuminants, and inclusion or discovery of known surfaces. Statistical assumptions about scenes include a variety of grey-world assumptions, for example that averaged over the entire scene the surface reflectance is grey or that it matches some other known average reflectance[3,17] or that somewhere in every scene will be found a surface patch that maximally stimulates each of the long, medium and short wave sensors (retinex with reset).[22] Brainard and Freeman[1] assume that the scene's illumination and reflectance spectra are drawn from known probability distributions. All of these methods work well when the statistical assumptions are met but can easily be made to fail by designing reasonable scenes in which the assumptions are not met.

The 'peeking' methods involve some method of obtaining an indirect glimpse of the illumination. For example, specularities usually reflect the incident illumination unchanged and since the specular component will be common to surfaces of different color it is possible[27,23] to extract it. Similarly, light interreflected between two surfaces shares a common component due to the incident illumination and so the incident illumination can be calculated[14]. Both the specularity and interreflection methods are vulnerable to scenes in which there are multiple sources of illumination. While a problem for these methods, multiple sources can also lead to further information if it can be determined that the same surface is being viewed under 2 or more illumination conditions. This information has been exploited to obtain surface shape[30,26] and for color constancy[10,5] Collections of surfaces of known reflectance have been used in volumetric color constancy[2] or supervised color constancy[25] to solve for the illumination. Recent methods[16,18] of color-based object recognition, that do not require color constancy à priori can be used to find such collections of surfaces from which the illumination properties can then be estimated.

Forsyth[12] introduced the idea of a canonical gamut of colors as a constraint for color constancy. A large set of

surfaces is examined under some standard illumination condition called the canonical illuminant and the set of all ρ_k sensor responses (i.e. 'colors') that arise is tallied. The 3D convex hull of this set then defines the canonical gamut, which is intended to represent the complete set of sensor responses that could arise under the canonical illuminant. When a different collection of surfaces is viewed under some unknown illuminant, the convex hull of that collection's sensor responses is obtained and compared to the canonical gamut. To the extent that this collection of surfaces is a representative one, the differences between the observed gamut and the canonical gamut will be due to the differences in the illuminations. Comparing the gamuts is not as simple as finding the mapping that transforms one gamut into another, but intuitively it is somewhat like that. The actual process involves taking each hull point of the canonical gamut and calculating the set of mappings projecting it into the observed gamut. The intersection of all the possible mappings for all the canonical hull points defines the possibilities as to what the unknown illumination might be.

Finite-Dimensional Linear Models

Assumptions about reflectances and illuminants can be expressed in terms of the dimension of the linear models required to approximate them accurately. In general, solving for color constancy is a very underconstrained problem and even approximate solutions might be impossible were it not for the fact that the spectra of lights and reflectances appear to be quite constrained. These constraints are captured by finite-dimensional linear models which are used in many recent color constancy theories[17,24,19].

Statistical analyses done using principal component analysis of databases of reflectances and illuminations[4,20] have shown that there is a great deal of correlation between the power at different wavelengths in these spectra. This means that they can be described well by a smaller number of parameters than used to describe the data originally (e.g. 31). Surprisingly, as few as 3 parameters do quite well for daylight illuminations and 4-6 do well for reflectances. Illumination spectra are approximated as a linearly weighted sum of the first n basis spectra obtained via the principal component analysis:

$$E(\lambda) \approx \sum_{i}^{n} \varepsilon_i E_i(\lambda) \qquad (3)$$

A similar equation holds for reflectances. Reducing the number of parameters in this way of course helps in terms of data reduction, but more importantly it means that the number of unknowns to solve for drops to the point where it may match the number of knowns which at most is 3 per image location.

One very interesting theory that exploits the reduction in unknowns is that of Maloney and Wandell[24] They assume, albeit somewhat unrealistically, that reflectances can be reasonably represented using only 2 parameters. Under this assumption, and because of the linearity found in Equation 2, the set of sensor responses obtained from a collection of differently colored surfaces under a single illuminant must fall on a plane in color space. Color space is the 3-dimensional space defined by the set of sensor responses ρ_k *(k=1...3).* In the Maloney-Wandell theory, the illumination defines the orientation of the plane spanned by the sensor responses found in the image so that the illumination can be determined by fitting a plane to the data and solving for its orientation.

3-Parameter Models

A particularly appealing feature of low-dimensional models for spectra and reflectances is that information about the entire function is encoded in a few parameters. There may well be situations in which we need to deal with the full spectrum, but it is far from clear that color requires the full spectrum. A trichromatic system condenses the spectrum to 3 values and we know from color matching experiments that 3 primaries suffice for color mixing.

Both the Maloney-Wandell and Forsyth theories can be simplified to a certain extent by a technique we call *spectral sharpening*.[6] Spectral sharpening creates a new set of sensor sensitivity functions through a fixed, linear combination of the original sensor sensitivity functions $R_k(\lambda)$. The sharpened sensors are generally more narrowband than the original sensors and this means that adjustments for changes in the illumination can be modeled accurately by simply scaling the responses of each sensor independently.

Scaling the sensor responses independently corresponds to a von Kries type method of adaptation. It also equates to transforming the triple of sensor reposes by a diagonal matrix in order to model a shift in the spectrum of the illumination. In the past this scaling or diagonal transformation has been applied directly to the cone responses. Even worse it has been applied to the CIE XYZ coordinates as in the case of the CIELAB. The accuracy with which a diagonal transform models illumination change depends very much upon the coordinate system in which it is performed. Spectral sharpening finds the optimal coordinate system. Several definitions of optimality have been tried which all lead to quite similar sets of sharpened sensors.

Extensive testing using typical illuminants and reflectances has shown that the simple 3-parameter, diagonal scaling model performs very nearly as well as a full 9-parameter linear transformation in mapping between the two sets of sensor responses obtained from a single surface under two different illuminants. In addition, it works approximately as well (occasionally slightly better, occasionally slightly worse) as when 3-dimensional linear models are used to model the full spectra of the illuminants. In other words, even when the spectra of the two illuminants and the sensor response created under the first illuminant are given, a 3-dimensional model of illumination is in general no more effective in predicting what the sensor response will be under the second illuminant than a 3-parameter di- agonal scaling model applied to the sharpened sensor responses.

For the case in which illumination is modelled perfectly by a 3-dimensional linear model and reflectance perfectly by a 2-dimensional model, we prove that a diagonal model completely accounts for any illumination change[7]. The requirements of 3D illumination and 2D reflectance are exactly those imposed by the Maloney-Wandell algorithm in order for it to work perfectly. As a result the Maloney-Wandell theory which was initially formulated in

terms of finite-dimensional linear models, can be re-stated in terms of a simple diagonal, von-Kries-like, direct scaling of sharpened sensors.

This result is both surprising and, in a way, to be expected: surprising in the sense that finite-dimensional models appear at the core of the algorithm; expected in that, as Brian Wandell[29] argues, the algorithm always did rely on exactly 3-parameters in describing illumination—namely, the coefficients of the 3-dimensional linear model—so the diagonal model offers simply a different set of 3 parameters.

Using finite-dimensional models Healey et. al[18] develop an interesting algorithm for illumination-invariant object recognition. They show that for illuminations that are 3-dimensional, a change in illumination causes a simple linear transformation in their texture measure which can then be factored out. As in the case of the Maloney-Wandell algorithm, we show[11] that their texture algorithm can be simplified by reformulating it in terms of 3-parameter diagonal transformations applied to spectrally sharpened sensor responses.

Sharpened sensors can be used also in the Forsyth theory so that it is both simplified[8] and so that constraints on the gamut of possible illuminations[10] can be incorporated in addition to the constraints on the gamut of surface reflectances. Although retinex[22] is not formulated in terms of finite-dimensional models, its calculations will also discount the illuminant more effectively when performed in sharpened sensor space.

I believe that other theories that depend on 3-dimensional linear models such as Gershon's[17] and Brainard and Freeman's[1] can be similarly recast in terms of diagonal transformations, but have not as yet done so. That some theories can be recast in terms of diagonal transformations is not in any way intented as a criticism of those theories (and if it were it would apply equally to the several situations[14,19,13] in which I have used 3-dimensional linear models) but rather as evidence for the effectiveness of spectral sharpening.

Conclusion

Finite-dimensional models have played an important role in the development of recent color theories, but in hindsight they are perhaps not really needed. Color is adequately described by a 3-dimensional coordinate system. The perceived need for finite-dimensional models stemmed more from the lack of a good basis for the color coordinate system than a need actually to represent full spectra. With the exception of circumstances in which linear models of dimension 4 or greater are used (e.g. as in computer graphics[9]) in general for the common 3-dimensional case, once the cone responses are converted to the new sharpened sensor basis many color manipulations can be performed in the new basis without recourse to finite-dimensional descriptions of the underlying spectra.

References

1. Brainard, D. H., and Freeman, W. T., "Bayesian method for recovering surface and illuminant properties from photosensor responses", *SPIE Human Vision, Visual Processing and Digital Display V*, 364-376, Feb. 1994; (see page 344, this publication).
2. Brill, M. H., "A Device Performing Illuminant—Invariant Assessment of Chromatic Relations*" J. Theor. Biol.*, **71**:473-478, 1978.
3. Buchsbaum, G., "A Spatial Processor Model for Object Colour Perception," *J. Franklin Inst*. **310**: 1-26, 1980.
4. Cohen, J., "Dependency of The Spectral Reflectance Curves of The Munsell Color Chips," *Psychon. Sci.* **1**: 369-370, 1964.
5. D'Zmura, M., "Color Constancy: surface color from changing illumination," *J. Opt. Soc. Am. A*, **9**: 490-493, 1992.
6. Finlayson, G., Drew, M., and Funt, B., "Spectral Sharpening: Sensor Transformations for Improved Color Constancy," *J. Opt. Soc. Am. A,* **11**: 1553-1563, 1994.
7. Finlayson, G., Drew, M., and Funt., B., "Color Constancy: Generalized Diagonal Transforms Suffice," *J. Opt. Soc. Am. A,* **11**(11): 3011-3020, 1994.
8. Finlayson, G., "Color Constancy in Diagonal Chromatic ity Space," *IEEE Proc. Fifth Intl. Conf. on Comp. Vision, (in press)*, June 20-23,1995.
9. Peercy, M., "Linear Color Representations for Full Spectral Rendering", *Proc. SIGGRAPH'93*, 191-198, Aug. 1993.
10. Finlayson, G., Funt, B. and Barnard, J., "Color Constancy Under a Varying Illumination," *Proc. Fifth Intl. Conf. on Comp. Vis.,* Jun 1995.
11. Finlayson, G., Chatterjee, S. and Funt, B. "Color Angle Invariants for Object Recognition," *Proc. Third Color Imaging Conference,* Nov. 1995.
12. Forsyth, D., "A novel algorithm for color constancy," *Intl. Journal of Computer Vision* **5**:5-36, 1990.
13. Funt, B. V. and Drew, M. S, "Color constancy computation in Near-Mondrian scenes using a finite dimensional linear model", *IEEE Proc. Computer Vision and Pattern Recognition Conference*, 544-549, June 1988.
14. Funt, B. V., Drew, M. S. and Ho, J., "Color constancy from mutual reflection," *Int. J. Computer Vision* 6: 5-24, 1991.
15. Funt, B. V. and Drew, M. S., "Color space analysis of mutual illumination," *IEEE Trans. Patt. Anal. and Mach. Intell.*, **15**(12):1319-1325, 1993.
16. Funt, B. V., and Finlayson G. D., "Color Constant Color Indexing," *IEEE Trans. Patt. Anal. and Mach. Intell.,* **17**(5), May 1995.
17. Gershon, R., Jepson, A., and Tsotsos, J., "From RGB to Surface Reflectance: Computing Color Constant Descriptors in Images," *Perception* 755-758, 1988.
18. Healey, G. and Slater, D. "Global Color Constancy: recognition of objects by use of illumination invariant properties of color distributions," *J. Opt. Soc. Am. A,* **11** (11): 3003-3010, Nov 1994.
19. Ho, J., Funt, B. V., Drew, M. S., "Separating a Color Signal into Illumination and Surface Reflectance Components: Theory and Applications," *IEEE Trans. Patt. Anal. and Mach. Intell.* **12**:966-977, 1990.
20. Judd, D, MacAdam, D, and Wyszecki, G., "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature," *J. Opt. Soc. Am.* **54**:1031-1040, 1964.
21. Kodepudy, R. and Healey, G., "Use of Invariants for Recognition of Three-Dimensional Color Textures," *J. Opt. Soc. Am. A,* **11**(11):3037-3049, Nov 1994.
22. Land, E. H., and McCann, J. J., "Lightness and retinex theory," *J. Opt. Soc. Am,* **61**:1-11, 1971.
23. Lee, H, "Method for Computing the Scene-illuminant Chromaticity from Specular Highlights," *J. Opt. Soc. Am. A*, **3**(10):1694-1699, Oct. 1986.

24. Maloney, L. and Wandell, B., "Color constancy: a method for recovering surface spectral reflectance," *J. Opt. Soc. Am. A,* **3**:29-33, 1986.
25. Novak, C. L., Shafer, S. A., "Supervised Color Constancy for Machine Vision*", SPIE* Vol. **1453** *Human Vision, Vi-Visual Processing, and Digital Display II*, 353-368, 1991.
26. Petrov, A. P., "On Obtaining Shape from Color Shading," *Color Res. Appl.* **6** (18): 375-379.
27. Shafer, S., "Using color to separate reflection components," *Color Res. Appl.*, **10**:210-218, 1985.
28. Swain, M., and Ballard, D., "Color Indexing," *Int. J. Comp Vision*, **7**:11-32, 1991.
29. Wandell, B. Personal Communication, 1993.
30. Woodham, R., Iwahori, Y. and Barman, R., "Photometric Stereo: Lambertian Reflectance and Light Sources with Unknown Direction and Strength," *TR91-18 Univ. B.C.*, Aug. 1991.

published previously in the IS&T 1995 Color Imaging Conference Proceedings, page 26