

Intersecting Color Manifolds

Brian Funt and Hamidreza Mirzaei; Simon Fraser University; Burnaby, B.C., Canada

Abstract

Logvinenko's color atlas theory provides a structure in which a complete set of color-equivalent material and illumination pairs can be generated to match any given input RGB color. In chromaticity space, the set of such pairs forms a 2-dimensional manifold embedded in a 4-dimensional space. For single-illuminant scenes, the illumination for different input RGB values must be contained in all the corresponding manifolds. The proposed illumination-estimation method estimates the scene illumination based on calculating the intersection of the illuminant components of the respective manifolds through a Hough-like voting process. Overall, the performance on the two datasets for which camera sensitivity functions are available is comparable to existing methods. The advantage of the formulating the illumination-estimation in terms of manifold intersection is that it expresses the constraints provided by each available RGB measurement within a sound theoretical foundation.

Introduction

Logvinenko's color atlas [1] provides a way to enumerate a complete and unique set of color-equivalent stimuli (material reflectance and illuminant spectral pairs that are indistinguishable) for all possible RGB tristimulus values [2]. Specifically, it provides a mechanism by which we can generate a unique set of illuminant and material spectra that completely covers the entire color space without any redundancy. In his theory, both the material reflectance and illuminant spectra are each specified by 3 parameters, so a color-equivalent stimulus is specified by 6 parameters. For a given tristimulus value, the set of color-equivalent stimuli defines a 3-dimensional manifold, which he terms the material-lighting-invariance manifold.

Using this theoretical structure, we propose an illumination-estimation method. For an image of a scene under a single illuminant, two different RGBs from two different pixels define two different material-lighting-invariance manifolds. By assumption they have at least the single scene illuminant in common. To find the common illuminants, the material-lighting-invariance manifold is projected onto the 2D illuminant space. The common illuminants must lie within the intersection of the projected manifolds. However, the intersection is not, in general, a single value, but rather a set of values. Intersecting the sets of illuminants defined by the RGBs from other pixels further constrains the range of possible scene illuminants. Tests with real images show that the method's performance is comparable to that of other well-known methods. An advantage of the proposed method is that it is founded on the theoretical principles of the color atlas and exploits precisely the theoretical constraints the atlas provides.

Many strategies have been proposed for estimating the chromaticity of the scene illumination. The present approach has similarities to Forsyth's gamut mapping method [4] and the voting

methods (Color by Correlation [5] and Sapiro's Illuminant Voting [6]). It relates to gamut mapping in that it represents a set of constraints and derives information from their intersection. It relates to Color by Correlation and Illuminant Voting in that the intersection is implemented via voting for candidate illuminants.

Background

The camera or eye's response to light with the spectral power distribution $P(\lambda)$ light reflected from a Lambertian non-specular surface material with reflectance $S(\lambda)$ is modeled in the standard way as the triplet φ_i ($i=1, 2, 3$):

$$\varphi_i = \int_{\lambda_{\min}}^{\lambda_{\max}} P(\lambda)S(\lambda)R_i(\lambda)d\lambda \quad i=1,2,3 \quad (1)$$

where $R_i(\lambda)$ is the spectral sensitivity function of a sensor class.

Logvinenko [2] defines a light-color atlas Ap as a subset of light spectral functions (strictly positive ones) such that given any color stimulus (object material reflectance x illuminated by a given light) there is a unique element p in the light-color atlas such that, illuminating x by p , this element results in a metameric match to the given color stimulus. There is a similar definition for the object-color atlas Ax . An object-color atlas is defined as a set of object material functions Ax such that for each color stimulus there is a unique element in Ax such that, if illuminated with the same light, this element results in a metameric match to the given color stimulus.

Logvinenko further defines a general color atlas in terms of an object-color atlas Ax and a light-color atlas Ap . The color atlas A is such that for any object illuminated by any light, there is a unique element p from Ap and a unique element x from Ax that is color-equivalent to the input pair. Color equivalence means being indistinguishable in multiple-illuminant scenes; however, in the special case of single-illuminant scenes it corresponds to metamerism [2]. For an arbitrary color stimulus (x', p') therefore there is a unique pair (x, p) in the color atlas that is color equivalent to it. By virtue of the definition of a color atlas, there will be a one-to-one map between any two atlases. The set of all color-equivalent object/light pairs is called the object-color set. Logvinenko refers to the object-color set along with the set of coordinates systems defined by the family of color atlases as the object-color manifold.

For any given color stimulus (x, p) its coordinates in the general color atlas can be determined by 2-step color matching [2]. Step 1 is to find the unique element a_m from the object-color atlas such that

$$\Phi_i(x, p) = \Phi_i(a_m, p) \quad i=1,2,3 \quad (2)$$

Step 2 is then to find the unique element a_i from the light-color atlas such that

$$\Phi_i(a_m, p) = \Phi_i(a_m, a_i) \quad i=1,2,3 \quad (3)$$

Since the object-color atlas and the light-color atlas are each 3-dimensional, the resulting coordinates of the general color atlas are 6-dimensional.

The object reflectances and light spectra of the color atlases can be represented in terms of rectangular functions [2] that are a mixture of uniform gray and a rectangular component that takes only values 0 and 1, with at most 2 transitions between 0 and 1. An algorithm for computing these functions from CIE XYZ is described by Godau et al. [3]. These rectangular functions are very unlike typical reflectance and illuminant spectral functions and hence are not very suitable for our purposes. However, Logvinenko also proposes other parameterizations of the color atlas, one of which is based on a Gaussian representation of spectra. The Gaussian representation is given in terms of a 3-parameter set of spectral reflectance functions $g_m(\lambda; k_m, \sigma_m, \mu_m)$ and a similar 3-parameter set of spectral power distribution functions $g_l(\lambda; k_l, \sigma_l, \mu_l)$, both of which are Gaussian-like (see equations 19-26 of [2]) functions, where k , σ , and μ indicate the height, standard deviation and center (peak) of the Gaussian. The functions are not strictly Gaussians in that they wraparound from one end of the visual spectrum to the other. They also differ from the inverse Gaussians used by Macleod et al. [7].

Any given sensor response triplet φ_i ($i=1,2,3$) can be decomposed into a sextuplet $(k_m, \sigma_m, \mu_m, k_l, \sigma_l, \mu_l)$ representing a Gaussian-like material reflectance lit by a Gaussian-like illuminant such that

$$\int_{\lambda_{min}}^{\lambda_{max}} g_l(\lambda; k_l, \sigma_l, \mu_l) g_m(\lambda; k_m, \sigma_m, \mu_m) R_i(\lambda) d\lambda = \varphi_i \quad i=1,2,3 \quad (4)$$

An example of this kind of 3-parameter Gaussian metamer spectra and a rectangular metamer for a sample Munsell chip's spectral reflectance is shown in Figure 1.

This decomposition is not unique; however, each choice of illuminant uniquely defines a corresponding material and vice-versa. For a given sensor response triplet, the set of all such illuminant-material pairs defines a 3-dimensional manifold embedded in a 6-dimensional space.

Illumination-estimation generally means estimating only the chromaticity of the illuminant, since for color balancing the intensity of the illuminant does not matter. In the present case, disregarding the intensity of the illumination and any uniform scaling of the percent surface spectral reflectance function reduces the number of parameters from six to four $(\sigma_m, \mu_m, \sigma_l, \mu_l)$. We defining chromaticity in the standard way as $\tau_i = \square_i / (\square_1 + \square_2 + \square_3)$ ($i=1,2$) then similar to (4), any sensor chromaticity can be decomposed as

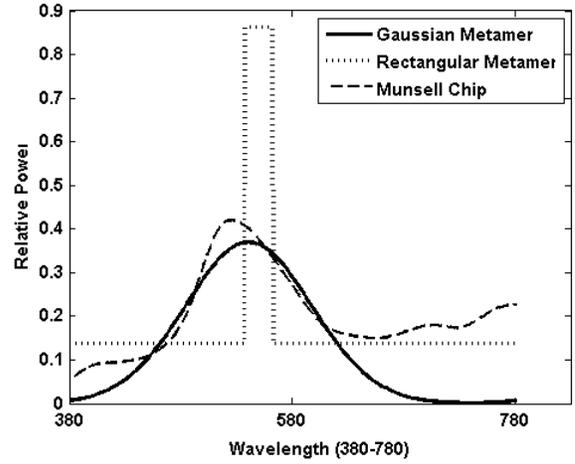


Figure 1. A sample Munsell chip's spectral reflectance illuminated by D65 and its rectangular and Gaussian metamers shown by dashed, dotted and solid lines respectively.

$$\int_{\lambda_{min}}^{\lambda_{max}} g_l(\lambda; 1, \sigma_l, \mu_l) g_m(\lambda; 1, \sigma_m, \mu_m) R_i(\lambda) d\lambda = n \tau_i \quad i=1,2 \quad (5)$$

where n is an arbitrary multiplier. The resulting parameters lie on a 2-dimensional manifold embedded in a 4-dimensional space. Figure 2 shows two examples of the manifolds defined by different input sensor chromaticity values.

Proposed Intersection Method

Given an RGB image of a single-illuminant scene, each distinct RGB value implies a manifold of illuminant-material pairs. Two different RGBs taken from the same single-illuminant image will define two different manifolds. For each RGB the set of possible illuminants consistent with it is represented by the illuminant components of the 4D manifold, so projection of the manifold onto the illumination axes results in the possible illuminants consistent with the given RGB. The illuminants consistent with two RGBs lie in the intersection of their projections. Each additional RGB will further limit the common intersection. To determine the overall scene illuminant, every distinct RGB from the image is used to limit the set of potential illuminant candidates as much as possible.

To compute the intersections, we use a Hough-transform-like approach. At first it might seem that the $(\sigma_m, \mu_m, \sigma_l, \mu_l)$ space would require a 4-dimensional Hough accumulator array, however, since we are only interested in the illuminant, a 2-dimensional array will suffice. The illuminant space is discretized and a 2-dimensional accumulator array is used. Given an image RGB, each $(\sigma_m, \mu_m, \sigma_l, \mu_l)$ point on the manifold defined by it increments the corresponding accumulator array cell (σ_l, μ_l) by a single vote. Once all the image's distinct RGBs have been processed, the accumulator array cell with the most votes determines the scene illuminant. If there is not a unique maximum, the illuminant estimates are averaged.

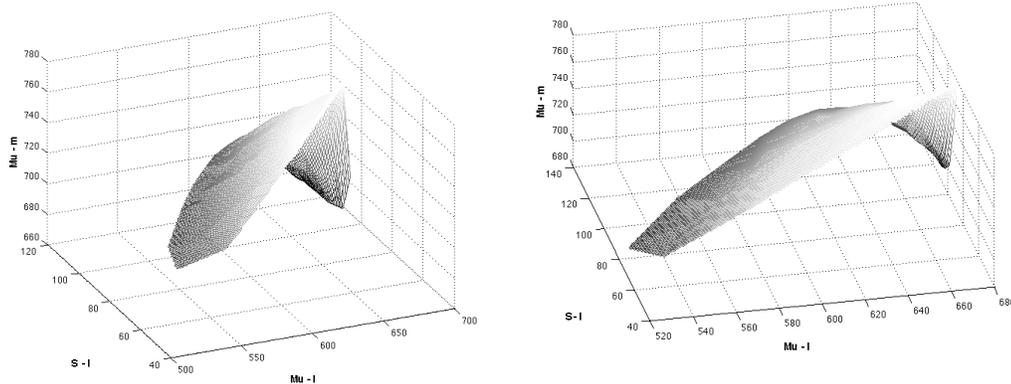


Figure 2. Plots of the 2-dimensional material-illuminant manifold embedded in the 4-dimensional space of $\mu_l, \sigma_l, \mu_m, \sigma_m$ for two different input chromaticity values. Axes are μ_l , S_l , μ_m , S_m for $\mu_l, \sigma_l, \mu_m, \sigma_m$, respectively. Three of the dimensions (i.e., μ_l, σ_l, μ_m) are represented by the axes. Note that the fourth dimension, σ_m , is not represented. Left and right panels illustrate the manifold for chromaticity values of (0.53, 0.13) and (0.66, 0.13) based on the SONY DXC-930 sensor sensitivity functions, respectively.

Implementation Details

The advantage of the light-color and object-color atlases is that they provide a means by which we can enumerate a unique set of illuminant and material spectra that completely covers the entire color space without any redundancy. In particular, we can enumerate the complete set of Gaussian illuminants from the light-color atlas by stepping through values of μ_l and σ_l , and are guaranteed that each will be of a different chromaticity. In practice, a step size must be chosen. Stepping uniformly in each of μ_l and σ_l discretizes the 2D space into bins. Experimentally, the ranges of both μ_l and σ_l were divided into 20 intervals each, for a total of 400 accumulator array bins.

To facilitate the voting procedure, we pre-compute a table containing all the possible material-lighting pairs (σ_l, μ_l) and (σ_m, μ_m) along with the corresponding (τ_1, τ_2) chromaticity values they would produce given sensor sensitivity functions $R_k(\lambda)$ according to Equation 4. The Gaussian illuminant spectrum $g_l(\lambda)$ and the Gaussian reflectance spectrum $g_m(\lambda)$ are both easily computed from the wrap-around Gaussian definition of equations 19-26 of [2]. Since only chromaticities are required, we can set both $k_m=1$ and $k_l=1$ and compute the chromaticity of the sensor response using Equation 5.

For the look-up table to be complete, all possible (σ_m, μ_m) , (σ_l, μ_l) pairs should be considered, which leads once again to a choice of step-size. Larger steps mean a coarser quantization; smaller steps mean a larger table. In the tests reported below μ and σ are discretized as:

$$380 \leq \mu \leq 780 \text{ in steps of } 5 \text{ nm}$$

$$0 \leq \sigma \leq 400 \text{ in steps of } 5 \text{ nm}$$

This choice of step-size leads to a table of about 10^7 entries,

each entry consisting of the 6-tuple $(\sigma_l, \mu_l, \sigma_2, \mu_2, \tau_1, \tau_2)$.

The voting process for each input chromaticity (τ_1^0, τ_2^0) then involves finding all the table entries of matching chromaticity. Two chromaticities are considered to match when $norm(\tau^0 - \tau_{table})$ is small.

Algorithm Summary

1. Pre-compute a table consisting of all the pairs of (σ_l, μ_l) and (σ_m, μ_m) along with their corresponding chromaticities at a given quantization step-size.
2. Given an image, compute its chromaticity histogram. The tests below are based on a 30x30 histogram. Binarize the histogram based on a threshold on the minimum number of occurrences. This eliminates spurious chromaticity values from noise.
3. For every distinct image chromaticity, find all the matching chromaticities in the table and return the corresponding σ_l and μ_l pairs.
4. For each (σ_l, μ_l) pair found in step 3, increment the accumulator array with a vote for that illuminant.
5. Once all distinct image chromaticities have been processed, the scene illuminant is determined by the (σ_l, μ_l) of the accumulator array bin with the maximum number of votes. Figure 3 shows an example of an accumulator array filled with votes. Given (σ_l, μ_l) , its chromaticity is easily computed via Equation 5.

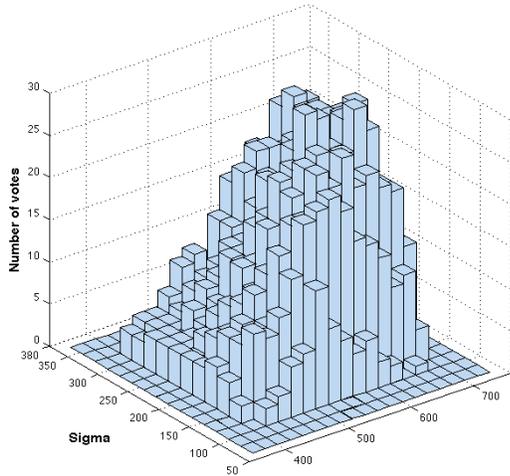


Figure 3. An example of the accumulator array of the number of votes generated by 30 image RGBs. The ranges of peak wavelength λ (right-hand axis) and spectral bandwidth σ (sigma axis), which are $[380\text{nm}, 780\text{nm}]$ and $[0, 400]$ respectively, have been divided into 20 intervals, giving 400 bins in total.

Results

The Projected Manifold Intersection method has been tested on both the SFU 321-image set [8][9] and the SFU HDR image set [10][11]. These are the only two color constancy datasets we are aware of that include the spectral sensitivity functions for the cameras used. Its performance is compared to that of MaxRGB algorithm, MaxRGB (bicubic preprocessing version of [11]), Greyworld, Do-nothing, Greyedge [12][1], and Color by Correlation Bright (Color by Correlation using bright pixels only [5] and tested on only 310 of the 321 images). The images in the 321-set are linear ($\gamma=1$) images of indoor scenes. The results are tabulated in Table I. The Do-Nothing error is the error obtained by simply assuming the scene illumination is always white (i.e., estimating its chromaticity as $r=g=b=1/3$). The error is the angular error in degrees between the estimated and ground-truth measures of the illumination's r, g, b .

For the HDR set, we computed a second look-up table based on the spectral sensitivity functions for the Nikon D700 camera used for that dataset. Table II illustrates the performance of the Manifold Intersection method in comparison to the other methods on HDR image database.

Table I: Performance of Projected Manifold Intersection in comparison to other well-known illumination-estimation methods on the 321 linear images of the SFU dataset. The error measure is angular error in degrees.

Method	Median	Average	Max
Do-Nothing	16	17	37
Grey World	7.1	9.8	7
Max RGB code of [1]	6.5	9.1	36
Grey Edge	3.7	6.1	28
SoG	4.0	6.0	25
MaxRGB (bicubic)	3.1	5.6	27
CbyC bright from Table V of [5]	3.2	6.6	
Projected Manifold Intersection	4.9	6.7	29

Table II: Performance of Projected Manifold Intersection in comparison to other well-known illumination-estimation methods on the 105 linear images of the SFU HDR dataset. The error measure is angular error in degrees.

Method	Median	Average	Max
Do-Nothing	15	15	30
Grey world	7.3	7.9	23
Grey Edge	3.9	6.0	25
SoG	4.0	6.0	25
MaxRGB (bicubic)	3.9	6.3	28
Projected Manifold Intersection	4.4	7.1	24

Conclusion

Logvinenko's color atlas theory provides a structure in which a complete set of color equivalent material and illumination pairs can be generated to match any given input RGB color. In chromaticity space, the set of such pairs forms a 2-dimensional manifold embedded in a 4-dimensional space. This manifold is the material-lighting-invariance manifold. For single-illuminant scenes, the illumination for different input RGB values must be contained in the intersection of the illuminant components of the corresponding manifolds. The proposed Projected Manifold Intersection method estimates the scene illumination based on calculating the intersection through a Hough-like voting process. Overall, the performance on the two datasets for which camera sensitivity functions are available is comparable to existing methods. The advantage of the formulating the illumination-estimation in terms of projected manifold intersection is that it expresses the constraints provided by each available RGB measurement within a sound theoretical foundation.

Acknowledgements

We wish to thank Alexander Logvinenko for his many helpful discussions. Funding was provided by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] A. Logvinenko, An object-color space, *J. of Vision*, vol. 9, no. 11, pp. 1-23, 2009.
- [2] A. Logvinenko, Object-colour space revisited, *International Conference on Computer Graphics, Imaging and Visualization. CGIV2010*, 2010.
- [3] C. Godau and B. Funt, The Logvinenko object color atlas in practice, *Color Research and Application*, published online: August 16, 2011, DOI: 10.1002/col.20680.
- [4] D.A. Forsyth. A novel algorithm for color constancy, *Int. Journal of Computer Vision*, 5(1): pp. 5–36, 1990.
- [5] S. Hordley and G. Finlayson, A Re-evaluation of Colour Constancy Algorithm Performance, *J. Opt. Soc. Am. A*, 23, 5 (2006).
- [6] G. Sapiro, Color and Illuminant Voting, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, v.21 n.11, pp.1210-1215, November 1999.
- [7] D. MacLeod and J. Golz, *A Computational Analysis of Colour Constancy*, in *Colour Perception: Mind and the Physical World*, Oxford University Press, p. 205, 2003.
- [8] K. Barnard, L. Martin, A. Coath, and B. Funt, A Comparison of Computational Color Constancy Algorithms, Part 2; Experiments with Images, *IEEE Transactions on Image Processing*, 11(9), pp. 985-996 (2002).
- [9] K. Barnard, L. Martin, A. Coath, and B. Funt, Simon Fraser University Computational Vision Lab, http://www.cs.sfu.ca/~colour/data/colour_constancy_test_images/index.html. Accessed: November 20, 2010.
- [10] B. Funt and L. Shi, HDR Dataset, Computational Vision Lab, SFU, Canada. http://www.cs.sfu.ca/~colour/data/funt_hdr/ Accessed: November 20, 2010.
- [11] B. Funt and L. Shi, The Rehabilitation of MaxRGB, *Proc. Eighteenth IS&T Color Imaging Conference*, San Antonio, Nov. 2010.
- [12] J. van de Weijer, T. Gevers and A. Gijsenij, Edge-Based Color Constancy, *IEEE Transactions on Image Processing*, 16, 9 (2007)
- [13] J. van de Weijer, T. Gevers and A. Gijsenij, <http://lear.inrialpes.fr/people/vandeweijer/software>, Accessed April 2, 2010.