How Does the Brain Arrive at a Color Constant Descriptor?

Marc Ebner

Universität Würzburg, Lehrstuhl für Informatik II
Am Hubland, 97074 Würzburg, Germany
ebner@informatik.uni-wuerzburg.de
http://wwwi2.informatik.uni-wuerzburg.de/staff/ebner/welcome.html

Abstract. Color is not a physical quantity which can be measured. Yet we attach it to the objects around us. Colors appear to be approximately constant to a human observer. They are an important cue in everyday life. Today, it is known that the corpus callosum plays an important role in color perception. Area V4 contains cells which seem to respond to the reflectance of an object irrespective of the wavelength composition of the light reflected by the object. What is not known is how the brain arrives at a color constant or approximately color constant descriptor. A number of theories about color perception have been put forward. Most theories are phenomenological descriptions of color vision. However, what is needed in order to understand how the visual system works is a computational theory. With this contribution we describe a computational theory for color perception which is much simpler in comparison to previously published theories yet effective at computing a color constant descriptor.

1 Motivation

The measured color varies with the type of illuminant used. The energy Q measured by a sensor is proportional to the reflectance R at the corresponding object point and is also proportional to the irradiance E at the corresponding object point, i.e. we have

$$Q(\lambda) \propto R(\lambda)E(\lambda)$$

for wavelength λ. The fact that the measured color varies with the type of illuminant is observed by many amateur photographers all around the world. One simply has to compare a photograph of the same scene once using incandescent light and once using sunlight. Professional photographers are well aware of this and can use filters to change the color balance [1,2]. Digital cameras apply post-processing algorithms which can change the color balance such that the resulting photograph looks more natural. In contrast, the color observed by a human observer stays remarkably constant [3]. This phenomenon has been investigated in detail by Land [4,5]. Obviously, it is of great interest to learn what algorithm is actually used by the human visual system in order to arrive at a color constant descriptor which remains constant (or at least approximately constant) irrespective of the illuminant used. Several psychophysical models of color perception have been put forward. However, such models do not explain how or why the color perceived by an observer could depend on either average apparent reflectance or the average luminance. Such models are phenomenological descriptions of color vision. What is needed is a computational theory of color vision [6].

Quite a large number of color constancy algorithms have been developed, from Land and McCann's Retinex theory [7] and its many variants [8,9,10,11,12], Buchsbaum's gray world hypothesis [13], gamut constraint methods [14,15,16], color cluster rotation [17] to comprehensive normalization [18] and computation of intrinsic images [18]. Most color constancy algorithms assume that the illuminant is constant within the image. A notable exception is Land and McCanns Retinex algorithm together with the variants of Horn [8], Blake [11] and Rahman et al. [12].

2 Iterative Computation of Local Space Average Color

Land's alternative formulation of the Retinex algorithm [9] as well as the algorithm of Rahman et al. [12] require that some form of averaging of image pixel be carried out. Land [9] assumes that input from several receptors is averaged. The algorithm of Rahman et al. [12] computes the blurred image using a convolution. Local space average color may also be computed iteratively as Ebner has shown [19,20,21].

The algorithm of Ebner assumes that a grid of processing elements exists with one processing element per image pixel. Each processing element is connected only to its nearest neighbors. Let $N(x, y)$ be the neighboring processing elements of the element located at position (x, y) of the image, i.e.

$$N(x, y) = \{(x', y')|(x', y') \text{ is neighbor of element } (x, y)\}.$$

Each processing element computes local space average color $\mathbf{a}(x, y)$

$$\mathbf{a}(x, y) = [a_r(x, y), a_g(x, y), a_b(x, y)].$$

Let us assume that we already have some average stored at each processing element. The following update equations are then iterated

$$\mathbf{a}'(x, y) = \frac{1}{|N(x, y)|} \sum_{(x',y') \in N(x,y)} \mathbf{a}(x', y')$$
$$\mathbf{a}(x, y) = \mathbf{c}(x, y) \cdot p + \mathbf{a}'(x, y) \cdot (1 - p)$$

where p is a small percentage. The first operation simply takes the local space average color which is available from neighboring elements and averages this data. In other words, we get a new average based on the data stored at neighboring elements. The color which is available at the current element is then slowly faded into the average using the second operation. If these two operations are iterated indefinitely, the data simply diffuses between neighboring elements. This process converges to local space average color irrespective of the data stored initially inside the processing elements.

The extent over which local space average color is computed, is determined by the parameter p. For a small value of p, local space average color is computed over an extensive area whereas for a large value of p, local space average color is computed over a small area. The iterative computation of local space average color is very similar though not identical to the convolution of the input image with an exponential kernel.

$$\mathbf{a}(x,y) = \int\int \mathbf{c}(x',y') e^{-\frac{|x-x'|+|y-y'|}{\sigma}} dx' dy'$$

The correspondence between the parameter σ and the parameter p is given by

$$\sigma = \sqrt{\frac{1-p}{4p}}.$$

Instead of using a grid of processing elements in order to compute local space average color, a resistive grid may also be used. With a resistive grid, adjacent points are simply connected through a resistor.

3 The Gray World Assumption

Local space average color may then be used to compute a color constant descriptor using the gray world assumption. The gray world assumption was originally proposed by Buchsbaum [13]. It is based on the assumption that on average, the world is gray. Buchsbaum assumed overlapping response characteristics of the sensors. We will derive the gray world assumption using non-overlapping response characteristics. Let $c_i(x,y)$ be the measured intensity of color channel i at position (x,y) of the image. The measured intensity is proportional to the reflectance and the irradiance.

$$c_i(x,y) = R_i(x,y) L_i(x,y).$$

Buchsbaum assumed that the illuminant is constant over the entire image, i.e. we have $L_i(x,y) = L_i$. This gives us

$$c_i(x,y) = R_i(x,y) L_i.$$

Thus, the illuminant scales the reflectances. A color constant descriptor can be computed once an estimate of the illuminant is available.

Let us now compute space average color over all image pixels. Space average color $\mathbf{a} = [a_r(x,y), a_g(x,y), a_b(x,y)]$ of an image with n pixels is given by

$$a_i = \frac{1}{n} \sum_{x,y} c_i(x,y) = \frac{1}{n} \sum_{x,y} R_i(x,y) L_i = L_i \frac{1}{n} \sum_{x,y} R_i(x,y).$$

We now assume that several differently colored objects are located inside the image. Since we don't know anything about the colors of the objects we simply assume that all colors are equally likely, i.e. we assume that the reflectances are uniformly distributed over the range $[0, 1]$. If we have a sufficiently large number of different colors inside the image, then we obtain for the expected value of the sum

$$\frac{1}{n}\sum_{x,y} R_i(x,y) = \frac{1}{2}.$$

We now see that the color of the illuminant can be estimated by computing global space average color.

$$L_i \approx 2a_i$$

We use the gray world assumption locally in order to estimate the color of the illuminant at each image pixel (x, y)

$$L_i(x,y) \approx 2a_i(x,y)$$

where local space average color $\mathbf{a}$ is computed iteratively as described above. Note that because the illuminant is estimated for each image pixel, the algorithm also works for a spatially varying illuminant, i.e. multiple light sources, provided that the environment is sufficiently diverse.

We then compute a color constant descriptor o_i by dividing each image pixel by twice local space average color.

$$o_i(x,y) = \frac{c_i(x,y)}{2a_i(x,y)} \approx \frac{c_i(x,y)}{L_i(x,y)} \approx \frac{R_i(x,y)L_i(x,y)}{L_i(x,y)} = R_i(x,y)$$

Figure 1 shows the results of this algorithm for a sample image. The image was taken with a Canon 10D. The white balance was set to 6500K and a yellowish illuminant was used. The image shown on the left is the input image and the image on the right is the output image. The color cast is removed nicely as can be seen in the output image.

Fig. 1. The input image is shown on the left. Local space average color is computed using an exponential kernel. The output image is shown on the right.

4 Usage of Color Shifts

Experiments done by Helson [22] indicate that human subjects appear to use color shifts in order to estimate the color of achromatic samples illuminated by colored light. The perceived color of the sample depends on the color of the illuminant as well as on the color of the background. He found that a bright patch located on a gray background will have the color of the illuminant. A dark patch will have the complementary color of the illuminant. Patches which have an intermediate reflectance will appear achromatic. The algorithms of Land [9], Horn [8,23], Moore et al. [24] and Rahman et al. [12] do not reproduce this behavior. This is because the ratio between the color of the pixel and local space average color is computed. If one computes this ratio, the color of the illuminant falls out of the equation. The stimuli will always appear to be achromatic for all settings that Helson investigated. A more extensive discussion is given in [25].

Ebner [19] has developed a computational algorithm for color constancy based on the use of color shifts. As we have seen above, we need to divide each image pixel by twice the space average color in order to obtain a color constant descriptor. However, we may also obtain a color constant descriptor if we shift local space average color onto the gray vector. The gray vector runs through the color space from black through gray and onto white. According to the gray world hypothesis, the average color of image pixels should be located on the gray vector. If the average color is not located on the gray vector it has to be corrected such that the gray world assumption is fulfilled. Let $\mathbf{w} = \frac{1}{\sqrt{3}}[1, 1, 1]^T$ be the normalized gray vector and let $\mathbf{c} = [c_r, c_g, c_b]^T$ be the color of the current pixel. Let $\mathbf{a} = [a_r, a_g, a_b]^T$ be local space average color computed for the same pixel. We first compute the component $\mathbf{a}_\perp$ of local space average color which is perpendicular to the gray vector. The vector $\mathbf{a}$ is projected onto the white vector $\mathbf{w}$ and the projection is then subtracted from $\mathbf{a}$. This gives us $\mathbf{a}_\perp$.

$$\mathbf{a}_\perp = \mathbf{a} - (\mathbf{a} \cdot \mathbf{w})\mathbf{w}$$

This vector points from the gray vector to the local space average color. We then subtract this vector from the color of the current pixel, i.e. we compute

$$\mathbf{o} = \mathbf{c} - \mathbf{a}_\perp.$$

Figure 2 illustrates how the shift is applied for two vectors $\mathbf{c}$ and $\mathbf{a}$. If we look at the individual components, i.e. color channels, we obtain

$$o_i = c_i - a_i + \frac{1}{3}(a_r + a_g + a_b).$$

Let $\bar{a} = \frac{1}{3}(a_r + a_g + a_b)$, then we have

$$o_i = c_i - a_i + \bar{a}.$$

The result of this operation is that the local space average color is shifted onto the gray vector and a color cast is removed. Since the shift is performed perpendicular to the gray axis the average intensity of the image pixels is not changed. This algorithm shows the same behavior as described by Helson.

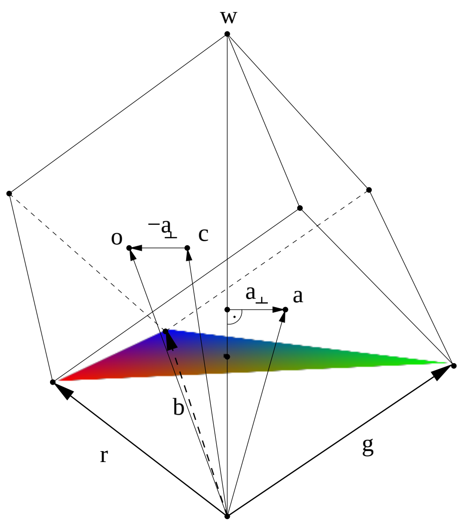

Fig. 2. The component $\mathbf{a}_\perp$ of local space average color $\mathbf{a}$ which is perpendicular to the gray vector $\mathbf{w}$ is subtracted from $\mathbf{c}$. The result is a color corrected image.

5 A Computational Theory of Color Perception

Theoretical models for color perception have been developed by Judd [26] and by Richards and Parks [27] among others. These are psychophysical models of color perception. They do not explain how or why the color could depend on either average apparent reflectance or the average luminance. They are phenomenological descriptions of color vision. What is needed is a computational theory of color vision [6]. The algorithms which would lend themselves to a biological realization are the parallel algorithms of Land and McCann [7], Land [9], Horn [8], Blake [11] and Ebner [19,20,21].

Of course, as of today, it is not yet known how the human visual system actually computes color constant descriptors. We do know that color constant cells have been found inside visual area V4 [3,28]. Area V4 may be subdivided into two subareas V4 and V4α [29]. V4 has a retinotopic organization whereas area V4α does not have a retinotopic organization. Cells found inside visual area V4 have very large receptive fields. These may be the cells which respond to either local or global space average color. They respond to the color of objects irrespective of the wavelength composition of the light reflected by the object. Area V4 also has callosal connections. The corpus callosum connects both hemispheres of the brain. Land et al. [30] have shown that an intact corpus callosum is required for accurate color perception.

Currently, we do not know how the processing of color information is actually done in V4. The computation of local space average color could either be done in space or in time [10]. Algorithms which perform an integration over time include the algorithms of Horn [8], Blake [11] and Ebner [19,20,21]. If the human visual system uses integration over time, then recurrent neurons are required which only have to be connected to their nearest neighbors. Instead of computing local space average color iteratively, local space average color could also be computed by consecutively applying a Gaussian blur. In this case, the neurons would have to form a hierarchy consisting of neurons with receptive fields of increasing sizes. The first neuron of the hierarchy would have a very small receptive field. The second neuron would have a sightly larger receptive field and so on. The neuron located at the top of the hierarchy would receive a completely blurred image as input, in other words, global space average color. This method would resemble the algorithm of Rahman et al. [12].

Hurlbert [10] suggested that it is also possible that the rods in the periphery of the retina are used to compute a spatial average over the image boundary. D'Zmura and Lennie [31] suggested that color constancy might be due to an adaptation mechanism. The retina is exposed to different parts of the scene as the eye, head and body moves. Space average color could be computed in the course of time, i.e. as the retina is exposed to different parts of the scene, by averaging the data per receptor and the adapted state at any point of the retina would be a function of this space average color. However, the experiments of Land and McCann [7] who have also experimented using short exposure times suggest that color constancy is an ability which exists even if the image is only perceived for a fraction of a second. The ability to perceive colors as constant is not dependent on long exposure times.

The first visual area where color constant cells have been found is V4. Assuming only local connections, i.e. that a highly parallel algorithm is employed, either the algorithm of Land [9], Horn [8] or Ebner [19,20] could be used by the visual system. A hierarchy of neurons, which is just used to compute a blurred image, would require an unnecessarily large neural architecture. Why should evolution favor this type of architecture if the same can be achieved using much simpler means? If the algorithm of Horn [8] is realized by the visual system, we would first need to construct a Laplacian operator. Local differencing could be used to implement a Laplacian operator. The output of the Laplacian operator would already be a color constant descriptor because the response of the photoreceptors is logarithmic (or nearly logarithmic) [32]. In order to implement the algorithm of Horn, we would now need a thresholding operation and an integration step. The integration would most likely be done in V4. Livingstone and Hubel [33] assume that such an algorithm is used by the visual system. Instead of operating on the cone channels red, green, and blue the Retinex algorithm is assumed to operate either in a longitude-latitude spherical polar coordinate system or inside a rotated coordinate system. In the spherical coordinate system of Livingstone and Hubel, radius denotes the dark-light scale, longitude the red-green axis and latitude the blue-yellow axis. That the Retinex algorithm can also be applied inside a rotated coordinate system was also noted by Land [9].

It may also be that the actual color signals inside the rotated coordinate system are averaged instead of averaging the thresholded output of the Laplacian. This would be essentially be the algorithm described by Ebner [19,20]. The advantage of this algorithm is that no threshold has to be set. In practice it is usually very difficult to properly choose a threshold. For this algorithm, local space average color would be subtracted from the color of the given pixel. Figure 3 shows the resulting architecture.

Livingstone and Hubel suggest that cells found inside the blobs of V1 act as building blocks which contribute to long-range interactions occurring in V4. It should be noted that in both models, the model of Horn [8] and the model of Ebner [19,20] no long range interactions are necessary. Only local connections between cells are required. The reason why distant areas may influence the color of a given point is most likely due to iterative propagation of data from one cell to the next.

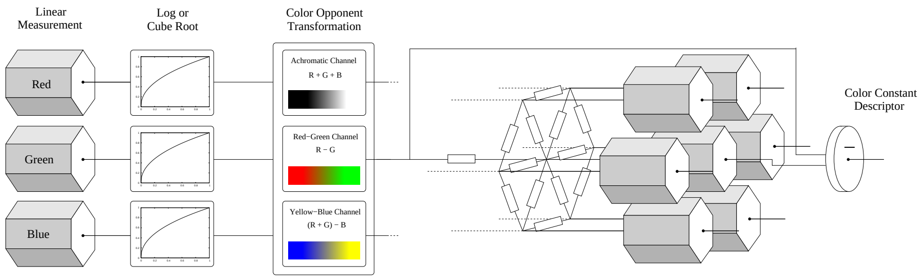

Fig. 3. Proposal of how the human visual system may arrive at a color constant descriptor.

We now dicuss our model in full. First, the cones of the retina measure the incident light for the three different color bands red, green, and blue. There is some dispute about whether the relationship between lightness and reflectance is logarithmic or follows a cube root or square root relationship. A logarithmic relationship was proposed by Faugeras [34]. See Hunt [35] for a discussion on why the relationship may either be a cube root or square root relationship. The reader should take note of the fact that the perceptually uniform CIE $L^*a^*b^*$color space (see [36,37]) also uses a cube root transformation. With a suitable scaling factor and offset, all three functions are a possible approximation. Let us assume for simplicity that the first step is the application of a logarithmic or other closely related function. Then a coordinate transform follows. This coordinate transform is most likely carried out by the color opponent cells. The color space is now described by the three axes red-green, blue-yellow and black-white. Local space average color is then computed using interconnected neurons. We only require that the neurons be connected to their nearest neighbors. The smoothing step could be carried out through resistive coupling of neurons. It is known that gap junctions behave mainly as pure resistors [32]. Such gap junctions could be used to diffuse color information to adjacent neurons. Once local space average color has been computed, it is subtracted from the color of the current pixel. The result is a color constant descriptor.

6 Conclusions

Most theories of color vision are phenomenological descriptions, i.e. they try to explain why we perceive colors the way we do. However, what is needed is a computational theory of color perception that can be mapped to what is known about the visual system. The computational theory presented above is very simple yet effective at computing color constant descriptors. It estimates the illuminant locally for each point of the scene and hence also handles scenes with multiple illuminant. The theory explains that only locally interconnected neurons are required in order to compute local space average color and thereby a color constant descriptor. Local space average color can be computed by a set of interconnected neurons each receiving input from a particular point of the retina. Only resistive coupling between such neurons is required. According to this theory, the long range connections through the corpus callosum simply connect adjacent neurons of the left and right hemispheres of the brain. Due to the logarithmic response of the receptors one needs to subtract the computed local space average color from the input signal, i.e. a negative coupling between the output from local space average color and the input signal is all that is required. The result is a color constant descriptor.

References

1. Hedgecoe, J.: Fotografieren: die neue große Fotoschule. Dorling Kindersley Verlag GmbH, Starnberg (2004)
2. Jacobsen, R.E., Ray, S.F., Attridge, G.G., Axford, N.R.: The Manual of Photography. Photographic and Digital Imaging. Focal Press, Oxford (2000)
3. Zeki, S.: A Vision of the Brain. Blackwell Science, Oxford (1993)
4. Land, E.H.: The retinex. American Scientist **52** (1964) 247–264
5. Land, E.H.: The retinex theory of colour vision. Proc. Royal Inst. Great Britain **47** (1974) 23–58
6. Marr, D.: Vision. W. H. Freeman and Company, New York (1982)
7. Land, E.H., McCann, J.J.: Lightness and retinex theory. Journal of the Optical Society of America **61**(1) (1974) 1–11
8. Horn, B.K.P.: Determining lightness from an image. Computer Graphics and Image Processing **3** (1974) 277–299
9. Land, E.H.: An alternative technique for the computation of the designator in the retinex theory of color vision. Proc. Natl. Acad. Sci. USA **83** (1986) 3078–3080
10. Hurlbert, A.: Formal connections between lightness algorithms. J. Opt. Soc. Am. A **3**(10) (1986) 1684–1693
11. Blake, A.: Boundary conditions for lightness computation in mondrian world. Computer Vision, Graphics, and Image Processing **32** (1985) 314–327
12. Rahman, Z., Jobson, D.J., Woodell, G.A.: Method of improving a digital image. United States Patent No. 5,991,456 (1999)
13. Buchsbaum, G.: A spatial processor model for object colour perception. Journal of the Franklin Institute **310**(1) (1980) 337–350
14. Forsyth, D.A.: A novel approach to colour constancy. In: Second International Conference on Computer Vision (Tampa, FL, Dec. 5-8), IEEE Press (1988) 9–18

15. Finlayson, G.D.: Color in perspective. IEEE Transactions on Pattern Analysis and Machine Intelligence **18**(10) (1996) 1034–1038
16. Barnard, K., Finlayson, G., Funt, B.: Color constancy for scenes with varying illumination. Computer Vision and Image Understanding **65**(2) (1997) 311–321
17. Paulus, D., Csink, L., Niemann, H.: Color cluster rotation. In: Proc. of the Int. Conf. on Image Processing (ICIP), IEEE Computer Society Press (1998) 161–165
18. Finlayson, G.D., Schiele, B., Crowley, J.L.: Comprehensive colour image normalization. In Burkhardt, H., Neumann, B., eds.: 5th European Conference on Computer Vision (ECCV '98), Freiburg, Germany, Berlin, Springer-Verlag (1998) 475–490
19. Ebner, M.: Color constancy using local color shifts. In Pajdla, T., Matas, J., eds.: Proceedings of the 8th European Conference on Computer Vision, Part III,Prague, Czech Republic, May, 2004, Berlin, Springer-Verlag (2004) 276–287
20. Ebner, M.: A parallel algorithm for color constancy. Journal of Parallel and Distributed Computing **64**(1) (2004) 79–88
21. Ebner, M.: Evolving color constancy. Special Issue on Evolutionary Computer Vision and Image Understanding of Pattern Recognition Letters **27**(11) (2006) 1220–1229
22. Helson, H.: Fundamental problems in color vision. I. the principle governing changes in hue, saturation, and lightness of non-selective samples in chromatic illumination. Journal of Experimental Psychology **23**(5) (1938) 439–476
23. Horn, B.K.P.: Robot Vision. The MIT Press, Cambridge, Massachusetts (1986)
24. Moore, A., Allman, J., Goodman, R.M.: A real-time neural system for color constancy. IEEE Transactions on Neural Networks **2**(2) (1991) 237–247
25. Ebner, M.: Color Constancy. John Wiley & Sons, England (2007)
26. Judd, D.B.: Hue saturation and lightness of surface colors with chromatic illumination. Journal of the Optical Society of America **30** (1940) 2–32
27. Richards, W., Parks, E.A.: Model for color conversion. Journal of the Optical Society of America **61**(7) (1971) 971–976
28. Zeki, S., Marini, L.: Three cortical stages of colour processing in the human brain. Brain **121** (1998) 1669–1685
29. Zeki, S., Bartels, A.: The clinical and functional measurement of cortical (in)activity in the visual brain, with special reference to the two subdivisions (V4 and V4α) of the human colour centre. Proc. R. Soc. Lond. B **354** (1999) 1371–1382
30. Land, E.H., Hubel, D.H., Livingstone, M.S., Perry, S.H., Burns, M.M.: Colour-generating interactions across the corpus callosum. Nature **303** (1983) 616–618
31. D'Zmura, M., Lennie, P.: Mechanisms of color constancy. In Healey, G.E., Shafer, S.A., Wolff, L.B., eds.: Color, Boston, Jones and Bartlett Publishers (1992) 224–234
32. Herault, J.: A model of colour processing in the retina of vertebrates: From photoreceptors to colour opposition and colour constancy phenomena. Neurocomputing **12** (1996) 113–129
33. Livingstone, M.S., Hubel, D.H.: Anatomy and physiology of a color system in the primate visual cortex. The Journal of Neuroscience **4**(1) (1984) 309–356
34. Faugeras, O.D.: Digital color image processing within the framework of a human visual model. IEEE Transactions on Acoustics, Speech, and Signal Processing **ASSP-27**(4) (1979) 380–393
35. Hunt, R.W.G.: Light energy and brightness sensation. Nature **179** (1957) 1026–1027
36. Glasser, L.G., McKinney, A.H., Reilly, C.D., Schnelle, P.D.: Cube-root color coordinate system. Journal of the Optical Society of America **48**(10) (1958) 736–740
37. International Commission on Illumination: Colorimetry, 2nd edition, corrected reprint. Technical Report 15.2, International Commission on Illumination (1996)