Environmental Challenges to Color Constancy

Larry Arend
NASA Ames Research Center, M/S 262-4, Moffett Federal Airfield, CA 94035
larend@mail.arc.nasa.gov

ABSTRACT

Theories of human color constancy have been based on experiments with relatively simple laboratory stimuli. Even recent “nearly natural” stimuli are optically much simpler than natural visual environments. I review here some of the complexity of natural visual environments. I argue that several kinds of optical structure exploited by theories of human color constancy may not occur in most natural scenes. Continued progress in color constancy research will require better descriptions of natural visual environments and of human color constancy performance within them. Both pose large challenges.

Keywords: color constancy, environmental optics, perception, surface color, illuminant color.

INTRODUCTION

The research problem of human color constancy is fundamentally about perception in the natural world. The whole problem consists of recognition that the retinal signal does not directly specify the intrinsic colors of surfaces due to the optical structure of the natural visual environment. The retinal signal is jointly determined by some combination of distal optical variables, some of which (the ones we wish to recover) are intrinsic color properties of the surface and some of which are contingent, i.e., accidents of the transient optical environment in which the surface is being viewed.

As in other scientific enterprises we have created simplified laboratory models of the problem in order to achieve experimental control and facilitate theory. Several recent efforts to introduce more complicated laboratory models have slightly narrowed the gap between the laboratory and natural visual environment (Gilchrist & Jacobsen, 1983, 1984; Kraft & Brainard, 1999; Hurlbert, 1999). In spite of these advances the laboratory stimuli still do not capture a number of environmental factors that complicate perception of surface colors. As a result we still know very little about the degree of human color constancy in natural viewing environments, let alone the visual computations involved.

This paper reviews some of the environmental challenges to perception of stable surface colors. I will argue that several kinds of optical structure exploited by theories of human color constancy may be missing from most natural scenes. This is not intended to be a criticism of the past work, but rather a discussion of how difficult this problem is.

PROPOSED ENVIRONMENTAL BASES FOR COLOR CONSTANCY ALGORITHMS

Proposed algorithms for computing invariant descriptors of the surface color from the retinal signal typically exploit some particular supposed quantitative structure of the optical environment (see Hurlbert, 1998, for a more detailed discussion of computational models and related optical assumptions). The laboratory models one uses to study surface color perception have, by construction, the optical structure required by one's theory, and the algorithms succeed in their laboratory settings. The question to be addressed here is, “Do natural environments typically have the kinds of structure that various theories require?”, i.e., how good are the laboratory models as representations of the natural visual environment?

As examples I will discuss three kinds of structure that have been proposed in prior color constancy research. These were chosen partly for their importance in the history of color constancy theory and partly for the arguments that they may not occur in fairly common natural scenes. At this stage of color constancy research the arguments are necessarily based more on examples and optical principles than on direct statistical description of natural optical environments, for reasons discussed below.

EDGE RATIOS

Probably the most widely held assumption in constancy modeling has been that surface variables and illumination variables interact multiplicatively.

The information available to the observer is generally considered to be the spatiotemporal pattern of excitations of the four photoreceptor classes, ρ_k, at every point in the images on the retinas. In color constancy theory, most efforts have been based on highly simplified equations, capturing only the variation of spectral reflectance, $S(\lambda)$, the spectral power distribution of the illuminant, $E(\lambda)$, and the spectral sensitivities of the receptors, $R_k(\lambda)$:

$$\rho_k = \int E(\lambda)\, S(\lambda)\, R_k(\lambda)\, d\lambda \qquad k = 1,...,4 \qquad (1)$$

There are at least three issues regarding the multiplicative interaction of illumination and reflectance. Two are fundamental, concerning the basic mathematical structure of the color constancy problem. The third concerns accuracy of proposed algorithms.

Adequacy of cone ratios as approximations. On the basis of Eq. 1 color constancy theorists have argued that various forms of the von Kries adaptation algorithm could approximately compensate for changes in the illuminant color. For example, within each receptor class the ratio of excitation by the light from one reflectance to the excitation from another would remain approximately constant over illumination changes, for a limited class of surfaces, illuminants, and receptor sensitivities.

The formal mathematical limitations of the von Kries formulation as a representation of the Eq. 1 have been extensively discussed elsewhere (see, e.g., Brill & West, 1981; Worthey, 1985; West & Brill, 1982; Ives, 1912) and will not be addressed here. The invariance is only approximate and the departures from invariance are determined jointly by the surface spectral reflectance distributions, illuminant spectral power distributions, and receptor spectral sensitivities.

Additive processes, spatially uniform and nonuniform. A second fundamental issue has received less attention. This is the widespread occurrence in natural scenes of optical effects in which illumination interacts in an additive way rather than multiplicative. Additive processes alter the mapping from reflectance ratios to cone excitation ratios, undermining the utility of the latter in calculation of invariant descriptors of surface color.

The obvious examples of additive processes are fog, dust, and haze in the open air. These add light to the retinal image which is unmodified by the reflectances of the surfaces beyond the haze (Fig. 1). The spectrum of this additive light is a function of both the ambient illumination and the scattering agents.

Haze is common in many climates, but additive phenomena are even more general. Most natural surfaces are neither Lambertian reflectors nor mirrors, but rather exhibit some intermediate degree of specularity. The light from many view angles contains a portion that has been multiplicatively modified by the body reflectance of the surface, but also another that is not affected by the body reflectance, with the spectral distribution of the local integrated illuminant (Klinker, et al, 1987; Klinker, 1988; Maxwell, 1996). The degree to which this component depends on the viewing angle varies with the material, ranging from independence to strong specular behavior.


Figure 1. Fog alters the cone excitation ratios corresponding to reflectance changes.

With respect to surface color perception, eq. 1 applies only to the body reflectance component of the light. The specular component plays the same additive role as light from atmospheric scatter , but generally with less spatial uniformity. The size and geometrical properties of the specular component vary from surface to surface, forming a sort of optically superposed image of a "gray scene" in which the reflected light has approximately the spectral distribution of the incident light. At a spatial change of reflectance, the ratio of excitations produced by the specular light within each receptor system is independent of the ratio of excitations produced by the body reflected light, depending on the physics of the surface s rather than the pigments. Consequently the mapping from the two body reflectances to cone excitation ratios is altered differently from location to location in the image. Unless some way can be found to calculate the specular component out of the cone excitations they are not reliably related to body reflectances of 3D, non-Lambertian surfaces.

One often-overlooked additive source is light scatter in the internal media of the eye. There is some scatter even in young healthy eyes, and the scatter grows in the aging process. To my knowledge there has been no direct study of color constancy performance as a function of aging, but the anecdotal evidence is intriguing. There are reports (Peli, 2000, personal communication) that many patients with developing cataracts do not notice visual changes until they interfere with acuity in daily life. They don't complain of degraded perception of surface colors, suggesting that their severe loss of retinal contrast produces no dramatic changes of surface-color appearance[1].

In all of these additive light situations the potential value of ratio responses to color constancy computations is reduced.

Reflectance and illuminant populations. The third issue concerns theoretical estimates of accuracy of the von Kries algorithm in natural scenes. The mathematical limitations noted above depend upon the populations of surfaces and lights in question. Most color constancy theorists have assumed as their illumination population the daylights of Judd, et al (1964) and the Munsell set of reflectances. For example, Foster and Nascimento (1994) showed that ratios of cone excitations from two surfaces are approximately invariant over illumination changes for Munsell surfaces and Judd, et al illuminations. When they used more spectrally pure illuminant and/or reflectance spectra the departures from invariance increased substantially. It's important then to ask how valid the Munsell and Judd et al populations are as models of natural scene surfaces and lights.

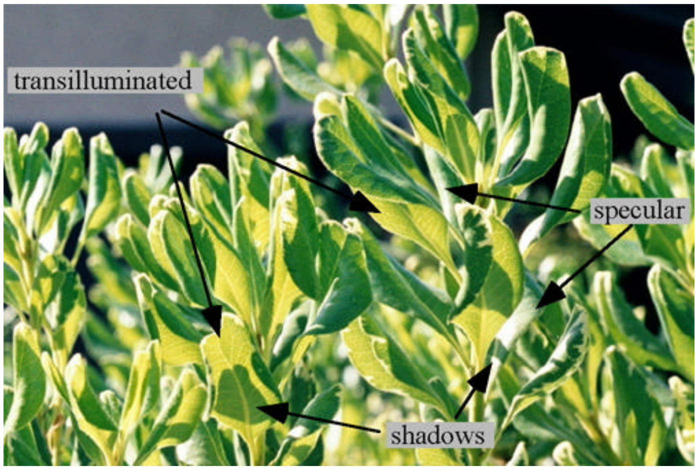

Figure 2. A bush with shading, shadows, specularity and transillumination. The shadows of the leaves have sharp edges (steep gradients).

The illumination that must be dealt with in computing the color of a surface in a natural scene is the local light that falls on the particular patch of surface in question, which is generally not direct skylight. Consider the surface of a leaf on the bottom side of the forest canopy. There is no direct line-of-sight from the surface to the sky, yet there is sufficient illumination to see the leaf. The illumination is the integrated light reflected to the surface by all the surrounding surfaces (in this case, primarily leaves and bark). If we move the leaf to other nearby locations it will receive light from smaller or larger regions of sky in addition to reflected light, some including direct sunlight, some not. None of these energy spectra is within the gamut of the Judd, et al daylights. Those with higher deviations from the equal energy spectrum can produce larger departures from surface color invariance in von Kries algorithms[2].

SHALLOW ILLUMINATION EDGES

Land and his associates (Land & McCann, 1971; Land, 1977) proposed that spatial illumination changes have shallower gradients than reflectance changes, allowing fairly simple algorithms to discriminate the two types of image gradient. While this is occasionally true in natural scenes, sharp illumination edges are very common in direct sunlight (Fig. 2) and shallow reflectance gradients are not unusual (Fig. 3). The spatial distributions of reflectance and illumination are discussed further below.

MULTIPLE $S(\lambda)$ SAMPLES WITHIN $E(\lambda)$ REGION

Many constancy models have assumed that typical natural scenes include multiple surface reflectances within individual regions of uniform illumination. This may occasionally occur, but there are optical reasons to doubt that it is common enough to be relied upon.

Illumination and reflectance tend to vary together over

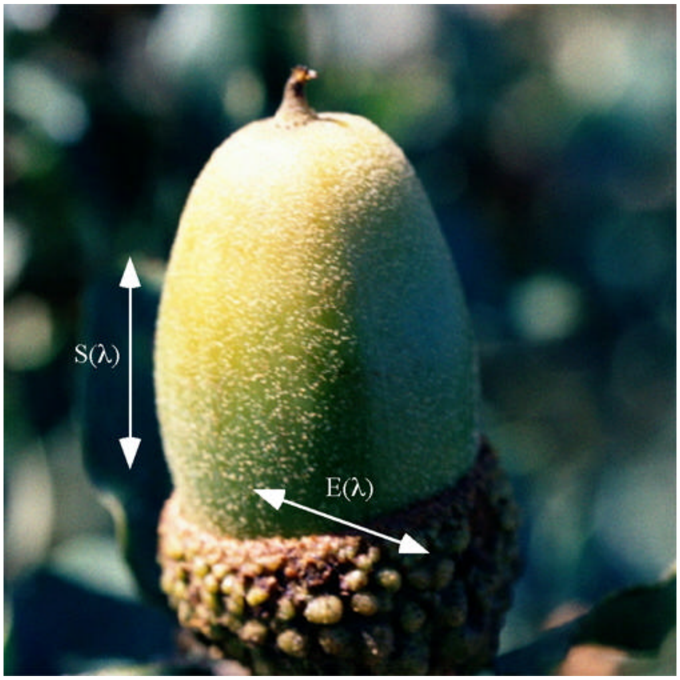

Figure 3. An acorn with shallow gradients of reflectance.

---

[1] Following cataract surgery patients do notice the blueness of colors, indicating that the gradual nature of the changes during cataract development is important.

[2] Similarly, purer (higher chroma) surfaces produce larger departures from invariance.

space, at roughly the same scales, in natural scenes due to the correlation of both variables with object geometry. While the reflectances of natural object surfaces are seldom exactly spatially uniform, large spatial color changes usually coincide with either sharp curvature of the surface of the object or occlusion at the terminal boundaries of objects. Those geometrical events also frequently produce changes of local illumination color due to several optical processes: shading, shadows, occlusion, and covariance of physical processes.

Visual scales. Before discussing the optics we need to briefly consider spatial scale. The book *Powers of Ten* (Morrison, et al, 1982) is a set of visualizations of the spatial scales of the known universe. Human visual ecology spans roughly the middle eight orders of magnitude, from about 10^{-4} to 10^{4} m. We perceive objects as big as mountains and as small as nearby grains of sand, with the texture at one scale being object shapes at larger scale. At these scales the visible structures are large relative to the wavelength of visible light. Visual scenes consist of nested geometric structures at different spatial scales but with the same photometric and colorimetric laws, so the following arguments apply at all visual scales.

Shading. Object shading is variation of the local effective illumination as the orientation of the surface changes relative to the incident light. Generally both the power spectrum and amount of illumination vary. Parts of objects with different reflectances often have different orientations (e.g., the fruit surfaces in Figs. 3, 7)[3], producing different effective illuminations.

Shadows. Throughout most of a clear day an object in sunlight tends to cast a shadow on nearby surfaces that is roughly the size of the object, due to the small angular subtense of the sun (e.g., the cactus leaves and pebbles in Fig. 5). Consequently the local spatial distributions of reflectance and shadows are nearly identical (though not coincident).

Figure 4. Patterns of shading, shadows, occlusions and specularity in a Eucalyptus tree.

Occlusion. Objects with relatively uniform surface color can have large spatial color changes at their terminal boundaries where they occlude more distant objects. The occluding and occluded objects are also generally in different illuminants, so that regions of relatively uniform reflectance coincide with regions of relatively uniform illumination (e.g., the Eucalyptus leaves in Fig. 4, the cactus leaves in Fig. 5).

Covariance of physical processes. This complicating factor is less common but not rare. Many physical and biological processes are dependent on illumination amount and spectral distribution. Lichen growth, fading of pigments, fruit ripening, etc. are all influenced by mean illumination, and all of them affect the local surface reflectance. When these processes are at work, the shading phenomena described above produce reflectance patterns that mimic the shading patterns (Figs. 3, 6, 7).

Another type of covariance is relatively rare, but is interesting for the striking illusions produced. Processes affected by gravity can also produce reflectance patterns that are correlated with local shading . Dust, pollen, pollution particles, etc. accumulate on object surfaces in much the same places that are illuminated from overhead. The deposits are not generally the same reflectance as the underlying object, so the reflectance of the object-plus-coating varies with the depth of the coating. The illusory perceived shading can be quite compelling, even in shadows and on cloudy days when real shading is weak.

The combined effect of these optical phenomena is to produce frequent cases in natural scenes in which each illumination falls on only a single surface reflectance.

---

[3] Chromatic reflectance gradients in Figs. 3, 6, and 7 are unfortunately not visible in the grayscale print. They are labeled.

Figure 5. A sunny Arizona scene with shading, shadows, and occlusions at several spatial scales.

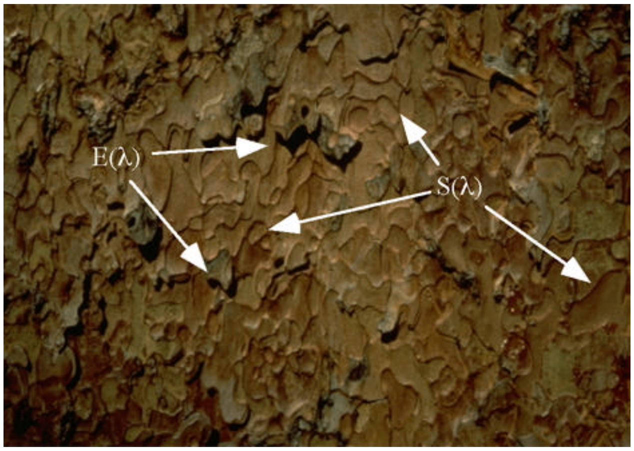

Figure 6. The red “shadow” pattern on the bark scales is a concentration of red growth (reflectance gradient).

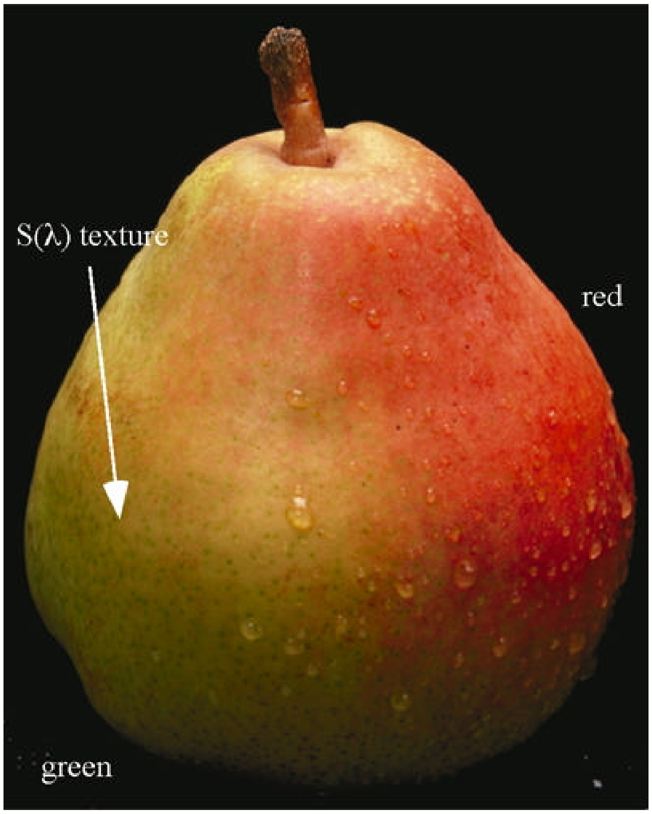

Figure 7. A texture made up of several intermixed reflectances on a relatively smooth surface.

Figure 8. A tree-trunk with shading, shadows, and occlusions at several spatial scales.

SUMMARY

To sum up so far, the optics of the natural visual environment are more complicated than those represented in the laboratory models used in color constancy research. The significance for color constancy theory is that optical structure needed by various proposed color constancy algorithms may not be reliably present in real natural scenes.

RESEARCH GAPS

The preceding arguments highlight two gaps in our research efforts:

1) We need better, more complete descriptions of natural visual environments. The examples and discussions of optical principles above indicate that the structure needed by constancy theories is sometimes lacking in natural settings. Systematic measurements are needed to document how widespread the problems are.

2) We need measurements of human color constancy performance within scenes that lack the optical structure required by the theories and within scenes that have the structure.

Both gaps pose large challenges, but there are encouraging precedents for both.

MEASUREMENT OF HUMAN COLOR CONSTANCY IN NATURAL ENVIRONMENTS

There have been some successful preliminary efforts to measure the apparent surface colors of objects in natural scenes. Gilchrist & Jacobsen (1983) measured the lightnesses of surfaces in natural scenes viewed through veiling luminances and found that lightness constancy could be quite good in spite of the reduced luminance contrast of the scene. While most of the visible surfaces were manmade, the scenes involved natural shading on 3D objects. Several of the early studies have indicated that complex natural optics may improve the approximation to color constant perception rather than interfere.

Kraft and Brainard (1999) extended the complexity of controlled laboratory stimuli by adding specular and 3D objects to more traditional flat, matte, uniform color patches. While they did not directly measure the perceived surface colors of any of the 3D surfaces, their apparatus and methods point the way to future experiments with still more complex scenes and observer tasks.

Advances in size and weight of instruments and the advent of high quality digital cameras should make it possible in the near future to measure human color constancy in well-described natural settings.

SYSTEMATIC DESCRIPTION OF ENVIRONMENTAL OPTICS IN RELATION TO HUMAN SURFACE COLOR PERCEPTION.

Appropriate quantitative description of natural visual environments is technically difficult, and relatively few data exist to date. At first glance the problem seems hopeless; how can we possibly reduce the population of all natural visual scenes to a manageable set? Deserts are different from meadows, forests, swamps, and mountains, and all vary with season, time-of-day, and weather. Even if we could make a principled choice of a small set of scenes, how could we produce a manageably compact description of their surfaces, lights, and geometry?

The task is obviously difficult, but developments in several related disciplines suggest the problems may not be completely intractable.

Plant and animal ecologists have developed methods for efficiently describing natural optical environments. For example, Endler (1993) was able to describe illumination in a wide variety of forest habitats from a very compact set of data. He obtained satisfactory approximations of illuminance spectra at various locations in nine kinds of forest around the world. He used only the irradiance spectra of four basis light sources (“sun”, “clouds”, “blue sky”, and “canopy: leaves, bark”), weighted by the respective proportions each contributed within the geometry of location and forest in question.

Computer graphics researchers have long sought descriptions of surfaces that generate visually convincing synthetic images. One aspect of this research has been measurement of actual surfaces (especially those that earlier graphic algorithms had trouble representing). The appearance of surfaces usually depends on the angles of view and illumination, in some cases dramatically so. Bidirectional Reflectance Distribution Functions (BRDF) are measurements of the reflectances of a surface at a suitable sample of all combinations of spectral band, illumination angle and viewing angle. Several laboratories are developing extensive libraries of BRDFs (see Dana, et al, 1999, for a review; URL1; URL2; URL3; URL4). Researchers have had some success at compact representation of both measured and modeled BRDFs (Rusinkiewicz, 1997; Koenderink, et al, 1996).

While these techniques from ecology and computer graphics may not lend themselves to direct application in color constancy research, the subject matter (appearance of natural scenes) is close enough to suggest that the techniques could be adapted to the human perceptual problem.

WHAT DOES "PERCEIVED SURFACE COLOR" MEAN ?

What do we mean by the perceived color of a surface? Even without getting into philosophical and deeper psychological aspects of this question there are problems at the level of simple optics. Once we look beyond flat, matte materials complicated and confusing questions arise.

One of the simpler issues concerns spatial inhomogeneity of spectral reflectance. Returning to problems of scale, most natural materials are spatially inhomogeneous when examined by the unaided human eye at close range. When the texture consists of variation of one pigment to another there is nothing very odd about surface color; the descriptions of the surface's spectral reflectance and apparent surface color are merely compound and awkward (Fig. 7). The light on the retina at a longer view distance is a spatial average of light from several different spectral reflectances, but the optically fused color has the same optical behavior as the unresolved component reflectances and it's fairly clear what we mean by the color of the surface.

There is another common kind of texture that is not so straightforward. Often the texture is produced by uneven surface geometry. Ridges and bumps on the surface are shaded and cast local shadows. A forest viewed from a distance may appear to be a very dark green that on closer inspection turns out to be lighter green trees interspersed with shadows[4]. What is the intrinsic color of the "surface" of the distant forest? A tree-trunk with light-colored ridges appears darker gray from a distance (Fig. 8). A bush may appear to have very light foliage from a distance where the individual leaves can't be resolved, but on closer view be a mix of body color, specular reflection, transillumination, and shadow (Fig. 2).

The problem with these colors is that, unlike the preceding speckled pigment case, the optically unresolved surface color is dependent on lighting geometry. The signal from the surface varies with time of day and weather. The kinds of image structure that computational algorithms rely on to deal with shading and highlights is below the spatial resolution of the visual system.

The same problems occur at smaller scales. We have already considered specularity as an illumination problem. It can also be considered to be a spectral reflectance problem. BRDFs of many materials include complicated specular properties attributable to the microstructure of the surface. Some materials reflect part of the illuminant at the surface without modification by the pigment in the body of the material. What is the intrinsic surface color of the material? The material is never seen without the superposed additive light component. What is the "intrinsic color" of such a surface? What perceived surface color would represent color constancy?

Some natural materials have reflective anisotropies of other kinds. Fur and hair look different under various angles of view and illumination. The fur on a bear's shoulder may appear brown or black, depending on the animal's posture and the light. Many plant structures also have color appearance that depends on complicated optics. Seed structures (Fig. 9) often have many small, closely packed filaments. The black appearance when viewing from the surface normal is attributable to light trapping in the tightly packed seed structure. Still other materials (e.g., weathered wood) have different spectral reflectances at the peaks and troughs of their corrugated surfaces. Like corduroy cloth, they have different apparent surface color depending on illumination and view directions. In cases of corrugations finer than the resolution of the eye at the viewing distance, the apparent color of the surface may darken as the illumination angle changes from illuminated troughs to shadowed troughs. What do we mean by "color constancy" for these surfaces?

Figure 9. A seed structure which is darker viewed from the surface normal.

At present there is no firm evidence that these complications are the rule rather than the exception in natural scenes. Nevertheless, one of the most striking aspects of preparing this paper was noting how little of the natural visual environment lends itself to analysis in terms of concepts familiar from laboratory color constancy research.

---

[4] The Black Hills of South Dakota are named for the appearance of their conifer forests from a distance (Smith & Young, 1904).

ACKNOWLEDGEMENTS

Supported by NASA Aerospace Operations Systems Program. Preliminary work on these concepts was conducted in 1995, while the author was a fellow at the Zentrum für interdisziplinäre Forschung, Universität Bielefeld, Germany.

REFERENCES

Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. Achromatic loci. *J. Opt. Soc. Am. A*, 15, 307-325.

Brainard, D. H., Brunt, W. A., & Speigle, J. M. (1997). Color constancy in the nearly natural image. 1. Asymmetric matches. *J. Opt. Soc. Am. A,* 14, 2091-2110.

Brill, M. and G. West (1981). Contributions to the theory of invariance of color under the condition of varying illumination. *J. Math. Biology* , 11, 337-350.

Dana, K. J., van Ginneken, B., Nayar, S. K., & Koenderink, J. J. (1999). Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics, 18,* 1-34.

Endler, J. A. (1993). The color of light in forests and its implications. *Ecol. Monographs,* 63, 1-27.

Foster, D. H. & S. M. C. Nascimento (1994). Relational colour constancy from invariant cone excitation ratios. *Proc. Roy. Soc. B,* 257, 115-121.

Gilchrist, A. L., & Jacobsen, A. (1983). Lightness constancy through a veiling luminance. *J. Exp. Psychol.: Hum. Percept. Perform.,* 9, 936-944.

Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception*, 13, 5-19.

Hurlbert, A. (1998). Computational models of color constancy, in *Perceptual Constancy* Eds., V.. Walsh & J. Kulikowski (Cambridge, UK: Cambridge, UK), pp. 283-322.

Hurlbert, A. (1999). Is colour constancy real? *Current Biology* 9(15), R558-R561.

Ives, H. E. (1912). The relation between the color of the illuminant and the color of the illuminated object. *Trans. Illum. Eng. Soc.* 7, 62-72.

Judd, D. B., D. L. MacAdam, & Wyszecki, G. (1964). Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, 54, 1031-1040.

Klinker, G. J. (1988). *A physical approach to color image understanding.* Ph.D. Thesis, CMU-CS-88-161.

Klinker, G. J., Shafer, S. A., & Kanade, T. (1987). Using a color reflection model to separate highlights from object color. *Proc. of the First International Conf. on Comp.Vis., London, IEEE,* 145-150.

Koenderink, J. J., van Doorn, A. J., & Stavridi, M. (1996). Bidirectional reflection distribution function expressed in terms of surface scattering modes. *Proceedings of the European Conference on Computer Vision, 2,* 28–39.

Kraft, J. M. & D. H. Brainard (1999). Mechanisms of color constancy under nearly natural viewing. *Proc. National Acad. Sci. USA*, 96, 307-312.

Land, E. H. (1977). The Retinex theory of color vision. *Scient. Amer.,* 237, 108-128.

Land, E. H. & J. J. McCann (1971). Lightness and Retinex theory. *J. Opt. Soc. Am.*, 61, 1-11.

Maxwell, B. A. (1996). *Segmentation and interpretation using multiple physical hypotheses of image formation* (Technical Report CMU-RI-TR-96-28): Carnegie Mellon University Robotics Institute.

Morrison, P., P. Morrison, Office of Charles and Ray Eames. (1982). *Powers of Ten.* New York, Scientific American Library.

Peli, E. (2000) Personal communication.

Rusinkiewicz, S. (1997). *A survey of BRDF representation for computer graphics*. http://www-graphics.stanford.edu/~smr/cs348c/surveypaper.html

Smith, G. M. and C. M. Young (1904). *The History and Government of South Dakota*. http://blackhills-info.com/history/index.htm.

West, G. & M. H. Brill (1982). Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy. *J. Theor. Biolog.,* 15, 249.

Worthey, J. A. (1985). Limitations of color constancy. *J. Opt. Soc. Am. A,* 2, 1014-1026.

Worthey, J. A. & M. H. Brill (1986). Heuristic analysis of von Kries color constancy. *J. Opt. Soc. Am. A,* 3, 1708-1712.

URL1. http://www.cs.columbia.edu/CAVE/curet/

URL2. http://www.ciks.nist.gov/links.htm

URL3. http://www.ciks.nist.gov/appmain.htm

URL4. http://www.graphics.cornell.edu/online/measurements/