

# Reflectance, illumination and edges

John McCann, Carinna Parraman, Alessandro Rizzi

McCann Imaging, Belmont, MA, USA, Centre for Fine Print Research, University of the West of England, Bristol, UK,  
Dip. Tecnologie dell'Informazione - Università degli Studi di Milano, Milano, Italy

## Abstract

We studied color constancy using a pair of 3-D Color Mondrian displays made of two identical sets of painted wooden shapes. There are only 6-chromatic, and 5-achromatic paints applied to nearly 100 block facets. The three-dimensional nature of these test targets adds shadows and multiple reflections not found in flat Mondrians. Observers viewed one set in uniform illumination--Low-Dynamic-Range(LDR); the other in highly directional non-uniform illumination--High-Dynamic Range(HDR). Both 3-D Mondrians, were side-by-side, in the same room, at the same time. We used two measurement techniques to evaluate how well the appearances correlated with the object's reflectances. First, we asked observers to compare the appearances of individual three-dimensional surfaces having identical reflectances, and recorded these changes in appearance using magnitude estimation. Second, an author painted a reproduction of the pair of Mondrians using watercolors. We measured the watercolor reflectances of the corresponding areas to quantify the change in appearances. Both measurements give us important data on how reflectance, illumination and image structure affect color constancy. A constant paint does not exhibit perfect color constancy, but rather shows significant shifts in lightness, hue and chroma in response to non-uniform illumination.

## Introduction

Human color constancy involves the content of the scene. It depends on the reflectances of objects, the spectral content and the spatial distribution of the illumination, and the arrangement of the scene. There are a number of models of color constancy used to predict colors from the array of radiances coming to the eye, or camera. They not only use a variety of image processing assumptions, they have different sets of required information, and different goals for the model to calculate. Table 1 lists the names, goals (result of the calculation), required information (inputs to calculation), and references.

Land's 1971 Color Mondrian [1] used a flat array of matte colored papers. He varied the amounts of uniform R,G, and B illumination over the entire array of more than 100 papers. He measured the light coming from a paper, then moved to a second paper and changed the illumination so that the second paper sent the same stimulus at a pixel to the eye. This

experiment demonstrated that identical retinal stimuli give rise to all colors. A red paper still looked red when its illumination was altered so that it was the same light stimulus as a green paper. The quanta catch of the retina at a pixel does not correlate with appearance. Color constancy measurements showed that color appearance correlates with the reflectance of the paper in Land's Color Mondrian. [2]

Land's Retinex model simply requires, as input, the spectral radiances at each pixel in the field of view. Its goal is to calculate the appearance of all colors in the scene. It builds color appearances out of spatial comparisons. Land said "... the function of retinex theory is to tell how the eye can ascertain reflectance in a field in which the illumination is unknowable and the reflectance is unknown." [1] Later retinex papers restated the language using *edges* and *gradients*, instead of *illumination* and *reflectance*. This was a result of studies of real life scenes in which: gradients in reflectance are difficult to see, and shadows with abrupt edges in illumination are highly visible.[3] (Table 1-row 1).

Model	Calculation Goal [Output]	Given Information [Input]	Reference
Retinex	appearance	radiance array of entire scene	Land,1971 <sup>1</sup>
Discount Illumination CIELAB	appearance	pixel's radiance + pixel's irradiance	CIELAB 1976 <sup>5</sup>
Discount Illumination CIECAM	appearance	pixel's radiance + pixel's irradiance + 4 scene coefficients	CIECAM 1997, 2002 <sup>5</sup>
GrayWorld Computer Vision	reflectance	radiance array of entire scene	Ebners <sup>6</sup>

Table 1 lists four classes of color constancy models. The first two columns list the names and the goals of the model's calculation. The third column identifies the information from the scene required to do the calculation. The fourth column lists references that describe the details of the calculation.

The late 19th century discussions of constancy began with the appearances of objects in different spectral illuminations. Helmholtz proposed the idea that humans discount the illumination,[4] so that

appearances correlated with recognizing the object, namely its reflectance. This principle is incorporated in pixel based color appearance models such as CIELAB and CIECAM. [5] These models use physical measurements of the illumination to normalize radiances from objects and remove the spectral information contained in the illumination. These models cannot predict color appearance without measurements of illumination at the pixel of interest as input. CIECAM requires 4 additional scene-dependent coefficients  $c$ ,  $N_c$ ,  $F_{LL}$ , and  $F$ . [5] (Table 1-rows 2, 3).

Computer vision has worked to remove this illumination measurement limitation by calculating illumination from scene data. The image processing community has adopted this approach to derive the illumination from the array of all radiances coming to the camera. [6] One of the important assumptions used in many of these, so-called Gray-World algorithms, is that scenes have a constant average reflectance. If true, then Gray-World algorithms can use the average radiance of all pixels to measure the spectral distribution of the illuminant. As long as the illumination is constant for all pixels in the scene, then each pixel's radiance divided by the calculated illumination will equal that pixel's reflectance. Computer-vision Gray-World models measure success by how well they can calculate an object's reflectance in different spectral illuminants. In order to use these GrayWorld models in a discussion about human vision, we need a separate psychophysical experiment to test whether appearances correlate with reflectance for the image in question. One cannot use such models for vision in situations where appearance deviates from reflectance. (Table 1-row 4) [6].

All four models listed in Table 1 do well calculating appearances in the flat uniformly illuminated Color Mondrian. The experiments in this paper present a different set of requirements for color appearance models. Here, with a restricted set of reflectances and highly variable illumination, we have more information to help sort out the importance of reflectance and illumination, as well as edges and gradients in modeling human vision.

## Experimental Method

The experiments in this paper are designed to study the interplay of reflectance, illumination and spatial content in human color appearance. We replaced the flat array of color papers used in Land's Mondrian with a collection of three-dimensional painted blocks. We replaced the 100 plus color papers used in Land's Mondrian with eleven reflectances: 6 chromatic and 5 achromatic. We replaced the spatially uniform illumination with a pair of different illuminants: one as uniform as possible, and the other highly directional. The uniform illumination has a low dynamic range (LDR), while the directional one has a high dynamic range (HDR).[7] While Land used many

reflectances to make a complex array of radiances, we used the shadows and gradients created by the 3-D objects to generate complexity. The three-dimensional nature of these test targets adds shadows and multiple reflections. These properties enrich the targets and make them more like real scenes. Here we measure the effects of illumination on constant reflectances. Human vision models need this data to assess how well their predictions match appearances.

## Two identical 3-D Mondrians

The experiment used two identical sets of objects in uniform and in non-uniform illumination in the same room at the same time. We painted each of the flat surfaces with one of eleven different paints (R, Y, G, C, B, M, W, G1, G2, G3, K). Figures 1 & 2 show photographs of the two scenes.



Figure 1 shows the low-dynamic-range (LDR) scene.



Figure 2 shows the high-dynamic-range (HDR) scene.

Figure 3 (left) shows a circular test target with 11 painted sections. Figure 3 (right) lists the Munsell chip closest to each paint, evaluated in daylight, and the measured Yxy in LDR light.




Figure 3 shows the flat painted test target; the paint designations, the Munsell designation and Y,x,y values

### Characterization of LDR & HDR Illuminations

Above, in Figure 1 we see a photograph of the LDR Mondrian in illumination that was as uniform as possible. The blocks were placed in an illumination cube. It had a white floor, translucent top and sides, and a black background. We directed eight halide spotlights on the sides and top of the illumination cube. The combination of multiple lamps, light-scattering cloth and highly reflective walls made the illumination nearly uniform. Departures from perfect uniformity came from shadows cast by the 3-D objects, and the open front of the cube for viewing.

Figure 2 is a photograph of the HDR Color Mondrian illuminated by two different lights. One was a 150W tungsten spot light placed to the side of the 3-D Mondrian at the same elevation. It was placed 2 meters from the center of the target. The second light was an array of WLEDs assembled in a flashlight (orange stick). It stood vertically and was placed quite close (20 cm) on the left. Although both are considered variants of white light they have different color appearance. The placement of these lamps produced highly non-uniform illumination and increased the dynamic range of the scene. In the HDR 3-D Mondrian, the black back wall had a 10 cm circular hole cut in it. Behind the hole was a small chamber with a second black wall 10 cm behind the other. We placed the flat circular test target on the back wall of the chamber. The angle of the spotlight was selected so that no direct light fell on the circular target. That target was illuminated by light reflected from the walls of the chamber. The target in the chamber had significantly less illumination than the same paints on the wooden blocks. The target in the chamber significantly increased the range of the non-uniform display. However, human observers had no difficulty seeing the darker circular target.

One way of assessing the uniformity of illumination is to make a third set of blocks, all painted middle grey. We photographed this actual Gray-World (Figure 4 left). It shows that with 3-D objects uniform

illumination is extremely difficult to achieve. Despite the use of 8 light sources and light diffusers, the three-dimensional objects cast faint shadows on other block faces. Perfectly uniform illumination requires that the object be in the center of a perfect integrating sphere. The HDR GrayWorld Figure 4 (right) shows a much wider range of luminances and apparent lightnesses.




Figure 4 (left) LDR Gray-World 3-D Mondrian. Figure 4 (right) HDR Mondrian shows that directional illumination can change a pixel's radiance as much as reflectance. Although all block faces have the same middle-grey reflectance, the camera digits in bright illumination are as high as 210, while in the shadows they are as low as 3.

### Magnitude Estimation Measurements

Observers compared the HDR and LDR Mondrians. [8,9] They were given a four-page form that identified a selection of 72 areas in the displays. The observers were shown the painted circular test target (Figure 3) placed on the floor of the display, in uniform light. This standard was explained to be the appearance of “ground truth”. They were told that all the flat surfaces had the same paints as the standard.




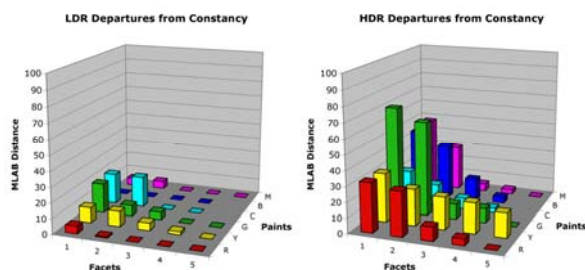
Figure 5 (left) shows the “ground truth” reflectance samples and illustrates the strategy for magnitude estimation of hue shifts; (center) lightnesses; (right) chroma.

Observers were asked if the selected areas had the same appearance as “ground truth”. If not, they were asked to identify the direction and magnitude of the change in appearance. The observers recorded the estimates on the forms. Observers were asked to estimate hue changes starting from each of the six patches of colors [R, Y, G, C, B, M]. Participants were asked to consider the change in the hue as a percentage difference between the original hue R and the hue direction Y. For example, 50%Y indicates a hue shift to a color halfway between R [Munsell 2.5R] and Y [Munsell 2.5Y] (Figure 5 left). 50% Y is Munsell 2.5YR. 100%Y meant a complete shift of hue to Y. Observers estimated lightness differences on a Munsell-like scale indicating either increments and

decrements, for the apparent lightness value (Figure 5 center). Observers estimated chroma by assigning paint sample estimates relative to 100% (Figure 5 right).


We know the Munsell Notation of chips of the 11 painted “ground truth” samples. We know the direction and magnitude of changes in appearance. We performed linear estimations to calculate the Munsell designation of a matching Munsell chip to each magnitude estimate. We used distance in the Munsell Book as described in the MLAB Color space, [10,11] as the measure of change in appearance.

Figure 6 shows one observer’s results of selected areas in the pair of 3-D Mondrians. Figure 6 plots the color change estimates for one observer, 6 colored paints and 28 color facets. We converted the observer estimates to Munsell chip designation, and then to MLAB Color space. We assumed that the Munsell Book of Color is, as intended, equally spaced in Color. MLAB converts the Munsell designations to a format similar to CIELAB, but avoids its large departures uniform spacing. [11] When the observer reports no change in appearance from illumination MLAB distance is zero. A change as large as white to black (Munsell 10/ to Munsell 1) is MLAB distance of 90. The plots in Figure 6 show each magnitude estimate for each color paint. For example, we asked an observer to evaluate 5 facets with red paint in HDR. We used the circular target in front of the display as reference;  $R$  equals  $2.5R4/14$ . The most different evaluation, for the red paint in the shadowed room (Figure 6, right), was  $R=2.5R2.4/8.4$ . The distance between  $2.5R4/14$  and  $2.5R2.4/8.4$  was 32.25 in MLAB units. The least different evaluation was no change, or 0.0 MLAB distance. The other three evaluations were 4.34, 9.25, and 29.13 MLAB units. The average value is of little statistical importance because the change depends on the illumination falling on each particular facet. These faces are not samples of the same population. The values for Y, G, C, B, M areas are plotted behind R. The corresponding MLAB values for the LDR scene are plotted on the left side of Figure 6.



**Figure 6** plots the MLAB distance between the ground truth and one observer’s color magnitude estimates for 28 surfaces in each illumination.

As seen above in Figure 6(left), there are small changes in appearance due to the lack of perfectly uniform illumination in the LDR scene. By comparison the HDR scene (right) shows much larger changes in appearance for the same reflectance. Here, the largest change is a green with a 70.3 MLAB change. We observed similar changes in appearance for gray samples. We observed similar results in all other observers’ data. These 3-D Mondrians provide an excellent test bed for evaluating approaches to color appearance. Figure 7 left shows a region of the 3-D Mondrian that has white paint. In the HDR illumination observers reported white, gray, pink and yellow appearances from the same white reflectance. The discrepancies from white were caused by shadows and light reflected from colored surfaces.



**Figure 7** (left) magnifies a photograph of a white reflectance facet in HDR that appears white, grey, magenta and yellow; (right) illustrates that the reflectance of the surface is white, and that the illumination has chromatic edges from multiple reflectances.

### Reflectance measurements of an artist’s watercolor rendition

After observers finished the Magnitude Estimations of Munsell designations, we left the pair of 3-D Mondrians in place. One of the authors (Carinna Parraman) painted with watercolors on paper a rendition of a section of both 3-D Mondrians. The painting took a considerable time so as to make the reproduction as close as possible to the appearances in both displays. Figure 8 is a photograph of the watercolor painting.

We made reflectance measurements with a Spectrolino® reflectance meter of selected areas. If constant reflectances in the scene appeared the same to the artist, then all paintings spectra for this area should superimpose. They do not. The artist selected many different spectra to match one scene reflectance (Figure 9). The artist’s selection of paints correlates well with the observers’ selection of Munsell designations. The

artist selected a narrow range of watercolor reflectances to reproduce the LDR scene. Many more paint colors are needed to reproduce the HDR scene.




Figure 8 shows the painted watercolor of portions of the LDR and HDR Mondrians.

Figure 9 shows the watercolor reflectance measurements for two white blocks in the LDR and HDR scenes. The left side plots the Spectrolino® measurements vs. wavelength. The vertical axis is reflectance scaled by  $L^*a^*b^*$  lightness. One block is a tall vertical surface illuminated by the lights and by reflections of other blocks. In the LDR reproduction the reflectances are very similar. In the HDR reproduction the reflectances vary in both lightness and chroma. The second surface with a white wooden block surface is the top of a short block [label 9]. It is well lit in the HDR and has a high reflectance in the watercolor. There is no direct illumination falling on this face in the HDR scene and the watercolor reflectance values are much darker, around  $L^*a^*b^*$  lightness of 40.

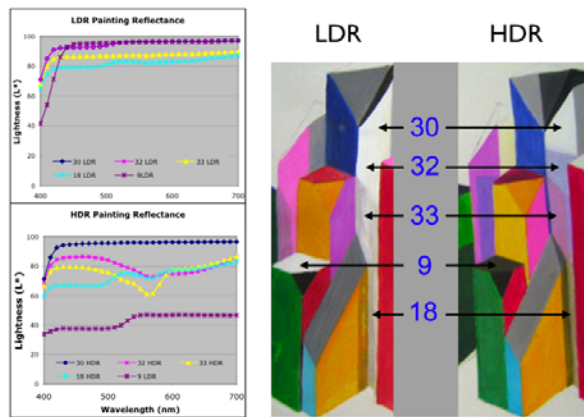


Figure 9 (left) shows the reflectance spectra of the watercolor painting for five areas. Figure 9 (right) shows the location of the measurements with arrows at the ends of the horizontal black lines with numbers to identify each area. The range of watercolor reflectances for LDR is much smaller than that for the HDR reproduction.

The important question addressed by both the magnitude estimation and watercolor reflectance measurement techniques was whether appearance remained the same in LDR and HDR illuminations. Both sets of data show that they do not. Reflectance shows poor to good, correlation with appearance in the LDR case, and poor to bad correlation in the HDR case. We found individual areas with high correlation between reflectance and appearance in both LDR and HDR data. We found that such correlation was frequent in LDR, and rare in HDR. Both sets of measurements give very similar results. Both sets of measurements show that appearance depends on the spatial properties of illumination, as well as reflectance. Edges in illumination cause large changes in appearance, as do edges in reflectance.

## Discussion

Land's Color Mondrian experiment showed that any color sensation can be generated from a single retinal stimulus at a pixel. Changing the radiances from the other pixels in the scene changes the appearance of the pixel of interest. Land's Color Mondrian experiment used an array of different reflectance papers in variable uniform illumination. In Land's experiments appearance were observed to be highly, but not perfectly correlated, with reflectance. [2] Numerous researchers have used vision and artificial intelligence models to predict appearances found in color constancy. Such models fall in to two general groups: *discounting the illumination*, (Table 1 - rows 2-4) and *spatial comparisons* (Table 1 - row 1). In order to *discount the illumination* one must first measure, or calculate, the light falling on the scene. Then, one takes the response to the light coming to the eye and divides it by the estimate of light falling on the object, leaving the object's reflectance. Models that *discount the illumination* assume that appearance correlates with reflectance. The second group of models, *spatial comparisons*, are indifferent to reflectance and illumination. These models synthesize appearance from the relative spatial information in the retinal responses to the scene. It is the visual system's response to edges and gradients that synthesizes appearances. The goal of these *spatial comparisons* models is to treat edges in illumination and edges in reflectance the same. Further, it calculates appearance whether, or not, it correlates with reflectance. [3]

Land's Color Mondrian extended color constancy from a pixel to a complex scene. Since it used a planar array in uniform illumination, it did not measure the appearances of real life 3-D scenes in non-uniform illumination. The experiments in this paper, by simultaneously studying LDR and HDR renditions of the same array of reflectances, extend Land's Mondrian towards real scenes in non-uniform illumination. These experiments show that uniform illumination is nearly impossible outside an integrating sphere. Real

scenes have multiple illuminants with differing amounts of light and different spectral contents. As well, shadows alter luminances and multiple reflections change local spectral content.

Previous experiments [12,13] varied the amounts of long-, middle-, and short-wave illumination (27 different spectra) falling on a flat surface in uniform illumination. These experiments measured the departures from perfect constancy. The results showed small changes in color appearances caused by illumination for highly colored papers and no changes with achromatic papers. These changes did not correlate with the changes in illuminant, rather they correlated changes in edge ratios seen by the broad spectral sensitivity of human cone pigments. These earlier experiments with flat, uniformly-illuminated targets, and the present 3-D Mondrians in nonuniform illumination both show that human color constancy does not work by discounting the illumination.

## Conclusions

Our experiments used two identical arrays of 3-D objects in uniform (LDR) and non-uniform (HDR) illumination. They were viewed in the same room at the same time. All flat facet objects used one of 11 paints. We used two different techniques to measure the appearances to observers of these constant reflectance paints. The first recorded observer estimates of change in Munsell Notation; the second measured an artist's watercolor rendition of both scenes. Both magnitude estimates and watercolor reflectances showed the same results. In nearly uniform illumination many samples appear the same as "ground truth", while few samples do in complex HDR illumination. Even small departures from perfectly uniform illumination generate departures in appearances from reflectance. Vision and artificial intelligence models that attempt to find reflectance and discount illumination do not predict appearance in real life scenes with complex non-uniform illumination. Edges and gradients in illumination behave the same as edges and gradients in reflectance.

## Acknowledgements

We would like to thank all the participants of CREATE, European Union, Framework 6 Marie Curie Conferences & Training Courses (SCF); and staff at the Centre for Fine Print Research University of the West of England Bristol; and support from the PRIN-COFIN 2007E7PHM3-003 project by Ministero dell'Università e della Ricerca, Italy; and the assistance of Alison Davis and Mary McCann.

## References

- [1] E. H. Land and J. J. McCann "Lightness and Retinex Theory", *J. Opt. Soc. Am.* **61** 1-11, 1971.
- [2] J. J. McCann S. P. McKee and T. Taylor "Quantitative studies in Retinex theory: A comparison between theoretical predictions and observer responses to 'Color Mondrian' experiments" *Vision Res.* **16** 445-458, 1976.
- [3] J.J. McCann, "Lessons Learned from Mondrians Applied to Real Images and Color Gamuts", *Proc. IS&T/SID Seventh Color Imaging Conference*, 1-8, 1999.
- [4] H. von Helmholtz, *Physiological Optics*, J. P. G. Southall, ed., Dover, New York, 1962.
- [5] R. W. G. Hunt, *The Reproduction of Color*, Sixth Ed. Chichester, (John Wiley and Sons, 2004).
- [6] M. Ebner, *Color Constancy*, Chapter 6, (Wiley, Chichester, 2007).
- [7] J. J. McCann, "Art Science and Appearance in HDR images", *J. Soc. Information Display*, vol. 15(9), 709-719, 2007.
- [8] C. Parraman, A. Rizzi & J. J. McCann, "Colour Appearance and Colour Rendering of HDR Scenes: An Experiment", *Proc. IS&T/SPIE Electronic Electronic Imaging*, San Jose, 7241-26, 2009.
- [9] J. J. McCann, C. Parraman, & A. Rizzi, "Reflectance, illumination, and edges in 3-D Mondrian Experiments", *Proc. AIC*, Sydney, in press, (2009).
- [10] G. Marcu, "Gamut Mapping in Munsell Constant Hue Sections", in *Proc. 6th IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, p. 159, (1998).
- [11] J.J. McCann, "Color spaces for color mapping", *J. Electronic Imaging*, 8, pp. 354-364, (1999).
- [12] J. J. McCann, "Mechanism of Color Constancy", *Proc. IS&T/SID Color Imaging Conference*, IS&T/SID, Scottsdale, Arizona, **12**, 29-36, 2004.
- [13] J. J. McCann, "Do humans discount the illuminant?", *Proc. SPIE Vol. 5666, Human Vision and Electronic Imaging X*; B. E. Rogowitz, T. N. Pappas, S. J. Daly; Eds., San Jose, p. 9-16, 2005.

## Author Biography

*John McCann (A.B. Biology-Harvard 1964) managed Polaroid's Vision Research Laboratory (1961-1996). He has studied human color vision, digital image processing, large format instant photography and the reproduction of fine art, & HDR imaging. He is an IS&T&OSA Fellow; a past President of IS&T and the Artists Foundation, Boston. He currently consults and continues research on color vision. He is the IS&T/OSA 2002 Edwin H. Land Medalist and IS&T 2005 Honorary Member.*

*Carinna Parraman is senior research fellow at the Centre for Fine Print Research, University of the West of England, Bristol. With a background in pigmented colour mixing for traditional print, her research has developed through long-term usage of digital imaging and printing. Her primary research has included alternative colour methods and digital colour printing for artists and designers. She is lead coordinator of the FP6 European project CREATE (Colour Research for European Advanced Technology Employment).*

*Alessandro Rizzi took the degree in Computer Science at University of Milano and received a PhD in Information Engineering at University of Brescia. Now he is assistant professor and senior research fellow at University of Milano. Since 1990 he is researching in the field of digital imaging and vision. His main research topic is the use of color information in digital images with particular attention to color perception mechanisms. He coordinates the Italian Color Group.*