REFLECTANCE, ILLUMINATION, AND EDGES IN 3-D MONDRIAN COLOUR-CONSTANCY EXPERIMENTS

John McCann[1], Carinna Parraman[2], Alessandro Rizzi[3]
[1]McCann Imaging, USA.
[2] Centre for Fine Print Research, University of the West of England, UK
[3]Dip. Tecnologie dell'Informazione - Università degli Studi di Milano, Italy

ABSTRACT

Colour constancy remains an important topic in colour research today, as it has for more than a century. Technological advances in digital capture and image processing have expanded studies of this constancy from individual colour patches to entire complex images. In Land's Colour Mondrian experiment he controlled uniform illumination over an array of more than 100 coloured papers to demonstrate and measure variable colour appearances from identical retinal stimuli. Colour appearances remained nearly constant despite large changes in the spectrum of the illumination. As well, all colour appearances were observed from the same R, G, B stimuli at a pixel.

Our new 3-D Colour Mondrian display is made of two identical sets of painted wooden shapes: one viewed in uniform illumination; the other in highly directional non-uniform illumination. There are more than 100 facets in each display. There are, however, only 6 chromatic, and 5 achromatic paints. The three-dimensional nature of these test targets adds shadows and multiple reflections, not found in flat Mondrian targets. These properties enrich the targets and make them more like real scenes. Both 3-D Mondrians, the Low-Dynamic- Range (LDR) in uniform illumination, and the High-Dynamic Range (HDR) in non-uniform illumination, were side-by-side, in the same room, at the same time. We asked observers to compare the appearances of individual three-dimensional surfaces having identical reflectances. These measurements give us another perspective on how reflectance, illumination and image structure affect colour constancy. The same paint does not exhibit perfect constancy, but rather shows significant shifts in lightness, hue and chroma in response to non-uniform illumination. This paper describes the physical measurements of the pair of 3-D Mondrians and their illuminants. Further, it describes observer measurements (magnitude estimates) of appearance changes from identical reflectances. We also made a second set of appearance measurements of both scenes. An author painted a watercolour of both LDR and HDR Mondrians at the same time. The hand-painted watercolour captures the observer appearances and renders them in the watercolour reflectance values. Here we measured appearances by measuring reflectances of the artist's rendering.

There are alternative theories to explain constancy: one discounts the illumination to determine the reflectances of objects; the other synthesizes appearance from edges and gradients in the retinal image. Data from both magnitude estimates and painting reflectance measurements help us to understand just how well humans do at discounting the illuminant. Both measurements find that appearance and reflectance correlate only in limited special cases.

Keywords: Colour Constancy, Discount Illumination, Colour Appearance, GrayWorld,

CONTACT mccanns@tiac.net

INTRODUCTION

Human colour 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. Land's Colour Mondrian experiment[1] used many coloured papers in uniform illumination. There are a number of models of colour constancy used to predict colours from the array of radiances coming to the eye or camera. They not only use a variety of image processing assumptions, they have different goals for the model. The early discussions of constancy began with understanding the appearances of objects in different spectral illuminations. Helmholtz proposed the idea that humans discount the illumination,[2] so that appearances correlated with recognizing the object, namely its reflectance. This principle is incorporated in pixel-based colour appearance models such as CIELAB and CIECAM.[3] 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 colour appearance without measurements of illumination as input.

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.[4] One of the important assumptions used in 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 average 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. Numerous computer vision algorithms use variations on the Gray-World hypothesis. Computer-vision Gray-World models measure success by how well they can calculate an object's reflectance in different spectral illuminants.[4]

Other imaging approaches to colour constancy have different goals.[5] These approaches return to the original goal of modelling appearances for humans.[6,7] The principal difference is that they attempt to calculate the colour that observers see, instead of the object's reflectance. These models do not attempt to solve the ill-posed problem of separating illumination from reflectance. They simply take the spatial distributions of scene radiances in three channels and synthesize an "appearance" image. These models replace a pixel's radiance with spatial comparisons (ratios of radiances) as the primary building block.[1] This approach introduces scene-dependent appearance predictions that do not always equal reflectance. These models do not make any Gray-World assumptions requiring that all scenes must have constant spatial illumination and constant average reflectance.

The experiments in this paper are designed to study the interplay of reflectance, illumination and spatial content in human colour appearance. We replaced the flat array colour papers used in Land's Mondrian with a collection of three-dimensional painted blocks. We replaced the 100 plus colour 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: the first as uniform as possible, and the second as highly directional. The uniform illumination has a low dynamic range (LDR), while the directional one has a high dynamic range (HDR).[8] 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 our complex array of radiances. Here we measure the effects of illumination on constant reflectances. Human vision models need this data to assess how well their predictions match appearances. Since computer-vision colour constancy compares predictions with reflectance, they do not compare results to appearance.

Figure 1(left) shows the low-dynamic-range (LDR) scene; (right) shows the high-dynamic-range (HDR) scene.

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. Each of the flat surfaces had one of eleven different paints applied. Figure 1 shows photographs of the two scenes. Figure 2 (left) shows a circular test target with 11 painted sections. Figure 2 (right) lists the Munsell chip closest to each paint, evaluated in daylight, and the measured Yxy in LDR light.

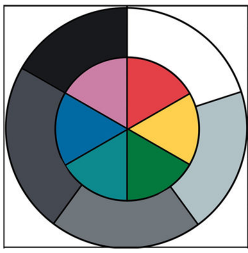

Name	Munsell	Y	x	y
R	2.5R 4/14	43.2	.631	.335
Y	2.5Y 8/14	105.0	.553	.429
G	10GY 5/12	29.7	.370	.547
C	10BG 6/8	39.9	.350	.434
B	5PB 4/12	20.6	.351	.350
M	10P 6/10	65.1	.520	.359
W	N 10/	156	.464	.417
G7.5	N 7.5/	76.3	.456	.418
G6	N 6/	44.5	.456	.418
G4	N 4/	21.2	.452	.416
K	N 1/	9.12	.462	.413

Figure 2 shows painted flat test target; the paint designations, the Munsell designation and Y,x,y values.

Characterization of LDR & HDR Illuminations

Figure 1 (left) shows LDR Mondrian in illumination that was as uniform as possible. The blocks were placed in an illumination cube (Figure 3left). 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. 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 3 left-centre). 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 centre of a perfect integrating sphere.

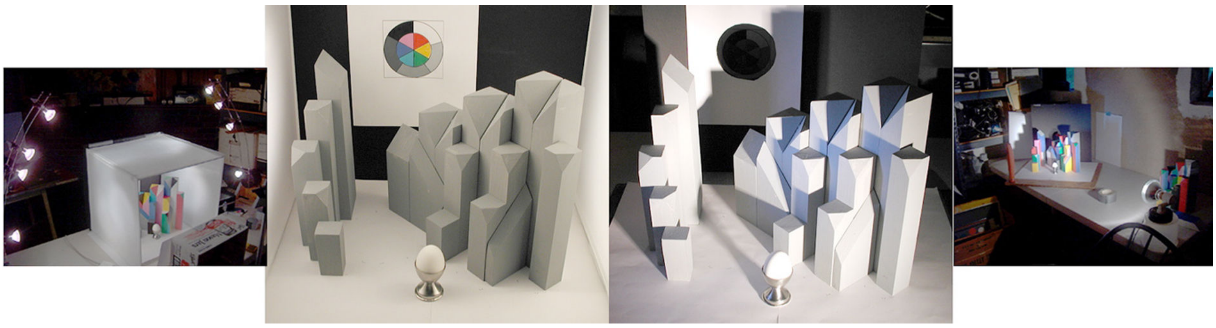

Figure 3 (left) LDR 3-D Colour Mondrian with lights; (left-centre) LDR Gray-World 3-D Mondrian; (right-centre) HDR Gray-World 3-D Mondrian; (right) HDR 3-D Colour Mondrian with lights.

Figure 1(right) is a photograph of the HDR Colour Mondrian illuminated by two different lights. One was a 150W spot light placed to the side of the 3-D Mondrian at the same elevation. It was placed 2 meters from the centre of the target. The second light was an array of WLEDs in a flashlight. It stood vertically and was placed quite close (20 cm) on the left (Figure 3 right). Although both are considered variants of white light they have different colour appearance. The placement of these lamps produced highly non-uniform illumination and increased the dynamic range of the scene. Figure 3 (right-centre) shows that directional illumination can change a pixel's radiance as much as reflectance. Even though 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. 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 this chamber significantly increased the range of the HDR target. However, human observers had no difficulty seeing the circular target.

MAGNITUDE ESTIMATION MEASUREMENTS

The project CREATE, (Colour Research for European Advanced Technology Employment), is funded by the European Union, Framework 6 Marie Curie Conferences & Training Courses (SCF). (www.create.uwe.ac.uk). The program's aim is to develop an international cross-disciplinary community within the areas of science and technology and the arts and commerce that could exchange both practical and theoretical knowledge. The intention is to foster novel ideas within a group of professional scientists and artists. Thirty-three CREATE participants were invited to view and analyse the HDR and LDR Mondrians. 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 2) 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. They 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 pre-printed forms. Observers were asked to estimate hue changes starting from each of the 6 patches of colours in the 'ground truth" standard [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 colour halfway between R and Y (Figure 4 left). 100%Y meant a complete shift of hue to Y. Observers estimated lightness differences on a Munsell-like scale indicating either increments and decrements, or the apparent lightness value (Figure 4 centre). Observers estimated chroma by assigning paint samples as 100% (Figure 4 right).

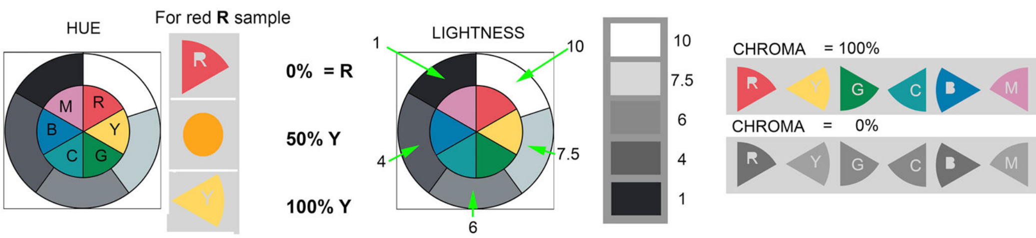

Figure 4- left shows the "ground truth" reflectance samples and illustrates the strategy for magnitude estimation of hue shifts; (centre) lightnesses; (right) chroma.

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 colour space,[9] as the measure of change in appearance due to non-uniform illumination. Figure 5 shows the results of selected areas in the pair of 3-D Mondrians. Figure 5 plots the colour change estimates for one observer, 6 paints and 28 observations. We converted the observer estimates to Munsell chip designation, and then to MLAB colour space. We assumed that the Munsell Book of Colour is, as intended, equally spaced in colour. MLAB converts the Munsell designations to a format similar to CIELAB, but avoids its large departures from being equally spaced in appearance.[10] 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 5 show each magnitude estimate for each colour paint. For example, as shown in the first rows, observers evaluated 5 facets with red paint. 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 1 right), R equals 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 values for the LDR scene are plotted on the left side of Figure 5.

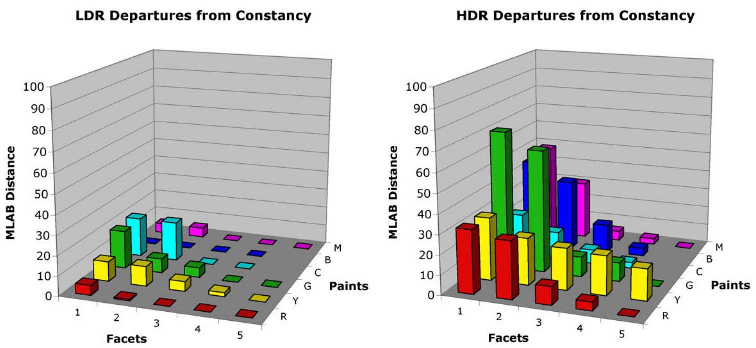

Figure 5 plots the MLAB distance between the ground truth and one observer's colour magnitude estimates for 28 surfaces in each illumination.

As seen above in Figure 5(left), there are small changes in appearance due to the lack of perfect 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 grey samples. We observed similar results in all other observers' data. These 3-D Mondrians provide an excellent test bed for evaluating approaches to colour appearance. Figure 6 left shows a region of the 3-D Mondrian that has white paint. In the HDR illumination observers reported white, grey, pink and yellow appearances from the same white reflectance. The discrepancies from white were caused by shadows and light reflected from coloured surfaces.

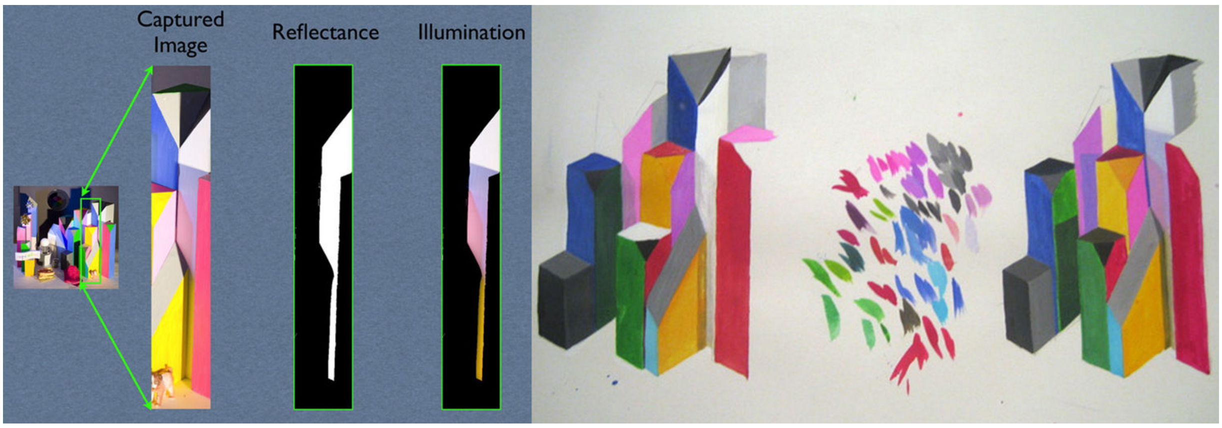

Figure 6 (left) magnifies a photograph of a white reflectance facet that appears white, grey, magenta and yellow; (right) shows the painted watercolour of portions of the LDR and HDR Mondrians.

ARTIST'S SCENE RENDERING

After observers finished the Magnitude Estimations of Munsell designations, we left the pair of 3-D Mondrians in place. One of the authors (CP) painted with watercolours on paper a rendition of a section of both 3-D Mondrians. Figure 6 (right) is a photograph of the watercolour painting. We made reflectance measurements with 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 7 compares the spectral measurements of the LDR and HDR paintings for 5 surfaces painted white. The LDR painting reflectances were similar (all achromatic whites and light greys), varying between L* values of high nineties to low eighties (Figure 7-left). The HDR painting reflectances were very different, varying as low as L* values of high thirties, with very different spectral content (Figure 7-right). The artist's selection of paints correlates well with the observers' selections of Munsell designations.

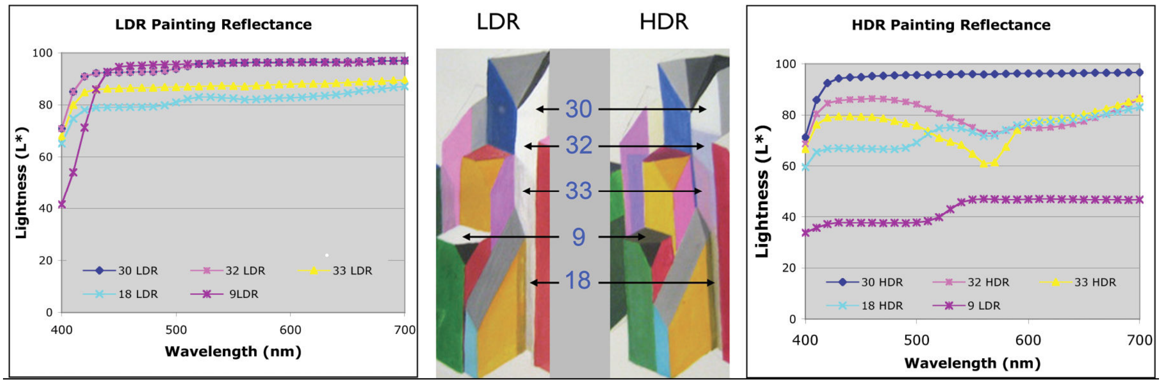

Figure 7 (left) plots the LDR paintings reflectances reproducing the appearances of 5 white reflectance facets; (centre) a photographic map identifying surfaces 9, 18, 18, 30, 33; (right) plots HDR paintings reflectances that appears white, grey, magenta, yellow and dark grey. The plots scale % reflectance using L* on the vertical axis.

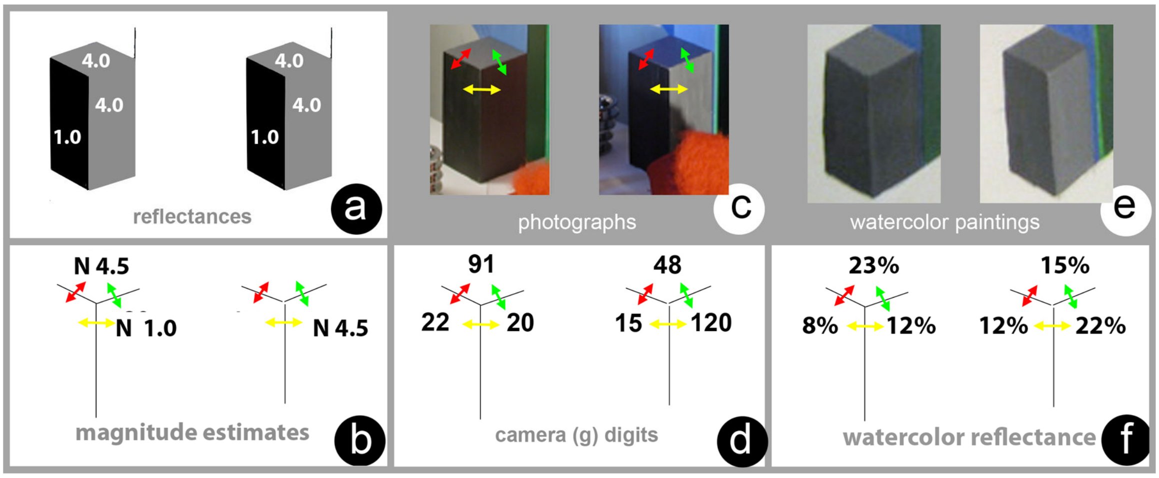

Figure 8 shows six different renditions of a small dark-grey block below the centre of the scenes. (a) shows that the top and right side of the blocks have N4.0 grey paint, while the left side has N1 paint. (b) shows that the magnitude estimates do not correlate with reflectances. (c) shows photographs of the LDR and HDR blocks with coloured arrows to identify edges. (d) shows camera digits from the above photographs. (e) shows photograph of the blocks in the watercolour painting. (f) shows the average reflectance measured from the corresponding areas of the watercolour. The magnitude estimates of appearance, the camera digits and the watercolour reflectances provide very similar renditions that do not correlate with reflectance. The camera digits show luminance changes at edges that are consistent with appearance, as measured by both magnitude estimates and watercolour reflectances.

Figure 8 uses a single block with adjacent faces with the same paint to illustrate the experimental results. First, it shows that the top and right faces have the same reflectance. Second, it shows the magnitude estimates and painting reflectances along with camera digits from photographs. Appearances are consistent with edge information, but not with scene reflectances.

DISCUSSION

Land's Colour Mondrian experiment showed that any colour 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 Colour 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.[11] Numerous researchers have used vision, and artificial intelligence, models to explain these results and the broader context of colour constancy. Such models fall in to two general groups: discounting the illumination, and spatial comparisons. 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 light falling on the object, thus calculating the object's reflectance. Models that discount the illumination assume that appearance correlates with reflectance. The second group of models are indifferent to reflectance and illumination.[6,7] 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 this second group of models is to treat edges in illumination and edges in reflectance in the same manner. Further, it calculates appearances, whether they correlate with reflectance, or they do not.[6,7]

Land's Colour Mondrian extended colour 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 spectral contents. In complex real world illuminations, appearances rarely correlate with reflectance, although some specific examples are found in all scenes.

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 (R, Y, G, C, B, M, W, G1, G2, G3, K). We used two different techniques to measure the appearances of these constant reflectance paints. The first recorded the observer estimates of change in Munsell Notation; the second measured the reflectances in an artist's painted rendition of both scenes. Both sets of data showed the same results. Some reflectance samples in a complex scene appear the same as the "ground truth" reference sample. In nearly uniform illumination (LDR) many samples look the same as the reference. In complex illumination (HDR), few do. Even small departures from perfectly uniform illumination generate departures in appearances from reflectance. Vision and artificial intelligence models that attempt to find reflectance by discounting 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. In order to calculate appearances computational models must use the scene's spatial information.

REFERENCES

[1] E. H. Land and J. J. McCann "Lightness and Retinex Theory", *J. Opt. Soc. Am*. 61 1-11, 1971.

[2] H. von Helmholtz, *Physiological Optics*, J. P. G. Southall, ed., Dover, New York, 1962.

[3] R. W. G. Hunt, *The Reproduction of Colour, Sixth Ed*. Chichester, John Wiley and Sons, 2004.

[4] M. Ebner, *Color Constancy, Chapter 6*, Wiley, Chichester, 2007.

[5] J. J. McCann, ed., *"Retinex at Forty"*, *Journal of Electronic Imaging*, **13**, pp. 1-145, 2004.

[6] J. J. McCann, Black Capturing a black cat in shade: past and present of Retinex color appearance models, Journal of Electronic Imaging, **13**, 36-47, 2004.

[7] A. Rizzi and J. J. McCann, "On the behavior of spatial models of color", in *Color Imaging XII: Processing, Hardcopy, and Applications, Proc. IS&T/SPIE Electronic Imaging*, 6493-02, 2007.

[8] J. J. McCann, "Art Science and Appearance in HDR images", *J. Soc. Information Display*, vol. 15(9), 709-719, 2007.

[9] G. Marcu, "Gamut Mapping in Munsell Constant Hue Sections", in *Proc. 6th IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, p. 159, (1998).

[10] J.J. McCann, "Color spaces for color mapping", *J. Electronic Imaging,* 8, pp. 354-364, (1999).

[11] J. J. McCann S. 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.