

Lessons Learned from Mondrians Applied to Real Images and Color Gamuts

John McCann

*McCann Imaging
Belmont, MA/USA
mccanns@tiac.net*

Abstract

This paper is an attempt to integrate a wide variety of psychophysical experiments into a computational model to calculate color appearance. Having described the fundamentals of such a model, we turn to applying this model to printing wide dynamic range, real-life scenes and finding the best reproduction of an image with limited printer gamut.

Key words: Mondrians, Retinex, Complex Images, Models of Vision.

Introduction

In 1963 Edwin Land decided he wanted to take his work on vision in a different direction. He quipped that many critics thought that his conclusions about human vision were suspect, because the experiments were photographic, and that these critics thought that he (Land) could do anything with photographic film. He set out to do experiments using just papers, lights and a very good telephotometer. Red and White experiments had shown the importance of complex images, so Land started with displays using about 100 papers. He asked Lucretia Weed to make a display resembling Mondrian's painting in the Tate Gallery, London. As it turned out Lucretia finished the display before we could find color photographs of that Mondrian painting. Years later, Hank Spekreijse, at a conference in Amsterdam, pointed out that his countryman Mondrian had never used high-chroma greens. So, although these experimental displays are higher in chroma and contrast than the Tate Mondrian painting, they are better visual test targets than the original because they have a much larger range of colors.

Nigel Daw's experiments¹ with afterimages had convinced Land to avoid regular arrays of squares. Daw's experiment had two parts: First, he had observers make a strong color afterimage of an object by fixating at a point in a color image, say a square pillow, that formed a diamond afterimage on the retina. Second, he asked observers to describe the afterimages as the observers moved their gaze to different fixation points on a black and white image of the same scene. When observers fixated on new points in the image, the mis-

match of contours between external scene and the internal afterimage inhibited the visibility of the afterimage. However, when the observers fixated on the original point registering the color afterimage in the black and white image of the pillow, the afterimage became more visible. For several minutes the observers could make the color afterimage of the pillow appear by looking at the original fixation point, and make it disappear by looking at another point in the image. The conclusion is that afterimages are inhibited by different contours on the current image. Hence, Land avoided the problem of afterimages by making each color in the Mondrian a different size or shape. Regular arrays of constant size patches are subject to the problem that the color afterimage of last square fits the contours of the new area of interest.

Land set out to study the appearance of colors by varying the spectra, intensity and duration of illumination falling on a large variety of Mondrians.

Black and White Mondrians

Gelb had shown that the same radiance from the same paper could look white or black, depending on other papers in the image. One of Land's first experiments was to adjust the uniformity of the illumination of the B&W Mondrian (Figure 1) such that a near-white paper far from the illuminant had the same radiance as near-black paper close to the light. Now, the same scene had a white and a black sending to the eye exactly the same radiances. In other words, nearly the entire range of sensations between white and black were generated by identical colorimetric stimuli.²

Lesson: Ratio-Product Required

At this time we began to look for models of how to calculate lightness from radiance. The work of Hans Wallach³ was very important in our thinking. He had shown that constant differences in sensation were associated with constant ratios of luminance. This property is a necessary but not sufficient property for a model of lightness. The missing property is illustrated in the B&W Mondrian shown in Figure 1. Here we have two sets of concentric circles with constant ratios of reflectance, but different reflectance values. The illumination is greater at the bottom than the top. The illumination was




Figure 1. A Black and White Mondrian that includes two pairs of concentric circles. The original Mondrian was used to point out that the grays at the tips of the arrows look different even though the radiances are equal. The gradual change in illumination from high at the bottom to low at the top makes a low reflectance paper at the bottom have the same radiance as a high reflectance paper at the top. The concentric circles are added here to make the point that local ratios cannot account for appearance in complex images. A ratio-product element is required in a model for lightness.

chosen to cancel the differences in absolute reflectance. That means that the small circle at the top is the same radiance as the small circle at the bottom. As well, the large circles have equal radiances. The two sets of circles do not appear the same. The small circular patch at the bottom appears darker than the one at the top. It has identical radiances: the small circles have an identical ratio to the large circles. Nevertheless, the small circles are not the same lightness.

A model of lightness needs to use the ratio of luminances for constant lightness differences. It needs an additional operation to introduce long-distance interaction that can model the fact that the top circle is lighter than the bottom one. That operation in Retinex models is achieved by replacing the ratio with the ratio-product⁴. The calculation starts with the current value at one pixel and multiplies that value by the ratio of the luminance of the comparison pixel divided by the radiance at the starting pixel. This New Product is the current value of the output pixel. Models that iterate this ratio-product operation successfully predict the lightness sensations reported by observer's matches of Figure 1.

Color Mondrians

The Black and White Mondrians showed that a particular luminance at the display and a particular quanta catch at the retina can appear any lightness (white to black). The Color Mondrian experiments showed that a particular triplet of radiances at the display and a particular quanta catch at the retina

can appear any color (nearly all sensations).

The first Color Mondrian experiments⁵ used Mondrians made of color papers, and three different, non-overlapping spectral illuminants (L = long-, M = middle-, S = short-wave visible light). In this experiment to get the same quanta catch from a green paper as red paper, the experimenter measured separately the L, M, S radiances from the red paper that appeared red in ILL, ILM, ILS illumination. (Fig 2, Step 1) The experimenter then moved the telephotometer to measure the radiances from the green paper. He separately adjusted the overall ILL, ILM, ILS illumination falling on the Mondrian until he got the L,M,S radiance values measured from the red paper, thus matching quanta catch (Fig 2, Step 2). The observer reported that the green paper looked green in the new ILL1, ILM1, ILS1 illumination. This is despite the fact that it is the same quanta catch as the red. By repeating this procedure with different papers we showed that a particular quanta catch at the retina can appear any color.

Lesson: Physical Basis of Incomplete Adaptation

The McCann, McKee and Taylor experiments matched each color in the 18 area Mondrian in each triplet of illuminants.⁶ This data showed quantitatively the lack of perfect color constancy because the average observer responses were different for different illuminants. McCann, McKee and Taylor (MMT) reported the measured reflectance of the papers in each illuminant using the cone sensitivity functions as L-, M-, S-wave radiance detectors. The spectra of the MMT illumination are controlled by three narrow-band interference filters (10 nm). The spectra of the matte surface papers are continuous across the visible spectrum.

Reflectance is a physical quantity measured by dividing the radiance of the sample by the radiance of the white for each wavelength. Usually reflectance instruments measure this ratio with a bandwidth of 1 to 10 nm using illuminants that are flat or at least monotonic across the spectrum. In such cases, the reflectance is usually constant with small changes in illumination.

Using broad-band cone-based sensors changes the picture. Here the sensors exhibit crosstalk. The L-cone sensor responds primarily to the L illuminant, but also responds significantly to the M illuminant and responds somewhat to the S illuminant. When the experimenter changes the amounts of the ILL, ILM, ILS illumination, the ratio of radiance from the sample to the white changes. The effect is due to the cross talk components. The L cone radiances are not simply illumination (ILL) times reflectance (Lr), but rather the sum of Lcone responses to three illuminants.

$$L_{cone} = a(Lr*ILL) + b(Mr*ILM) + c(Sr*ILS)$$

So, despite the fact that the paper surface is unchanged, we can measure changes the integrated reflectance values changed with changes in illumination. McCann, McKee and Taylor measured these changes and reported them in detail. The purpose was to try to correlate color matches with scaled integrated reflectance. The data shows that a substantial part of the small, but measurable color shift observed in color constancy experiments correlates with integrated reflectance - a phenomenon in the domain of physics. Most of these shifts

have nothing to do with any of the many adaptation mechanisms attributed to the domain of biology. In other words, the cause of much of the “incomplete adaptation” phenomenon is simply the physics of the broad-band spectral sensors found in cones. If the reflectance ratios change, then appearance changes.

Lesson: No Average Found in Adaptation

The more interesting versions of the Color Mondrian are those done with colored surrounds⁷. The original McCann, McKee and Taylor experiments do not try to discriminate between the various normalization techniques, such as maximum, average or grayworld and local average.

In order to get constant L, M, S radiances from different papers we changed the ILL, ILM, ILS illuminant intensities everywhere in the image. That action changed the average radiance (AVL, AVM, AVS) integrated over the entire image. That leaves the door open to grayworld explanations of color constancy. A better MMT experiment is one that includes a surround around the Mondrian. In Fig 2, Step 2 we change the ILL, ILM, ILS illuminants to get the desired L, M, S radiances at a particular point. The consequence is shifts in the average as big as the illuminant changes. Now we change the color reflectance of the surround area around the Mondrian to move the AVL, AVM, AVS averages back to be equal to the averages in the original illumination (Fig 2, Step 3).

Any grayworld theory now has no change in grayworld at input; any adaptation theory has no change in adapting average stimulus. Observer matches shows that color constancy works as it did in Step 2, with no change in grayworld or adaptation state. The corollary of this is that these experiments designed to look for the influence of grayworld and of adaptation find it has no effect on color constancy experiments.

To further elaborate on this point we performed a third set of Color Mondrian experiments⁸. In the previous experiments we balanced the change in illumination with a canceling change in reflectance of the surround. There grayworld and adaptation are reduced to zero. With new surround papers we introduce a change in grayworld as large as that introduced by a change in illumination, but we use the original illumination constant. As above we change the ILL, ILM, ILS illuminant to get the desired L,M,S radiances at a particular point. The consequences are shifts in the AVL, AVM, AVS averages as big as the illuminant changes. We measure the amount of this shift of AVL, AVM, AVS averages. Now we change the color reflectance of the surround area around the Mondrian to move the AVL, AVM, AVS averages as much as the illumination change would have. In other words we shift the average as much as the MMT experiment would have, but we view the display in the original illumination. Any grayworld and any adapting mechanism must introduce large changes in appearance. The matching data show that appearance does not change. Changes in grayworld and changes in adaptation do not cause changes in color appearance.

We have just tested the effect of grayworld by manipulating the average of the entire field of view. We found that such averages do not control color constancy. There is the possibility that there exists a local adaptation mechanism that uses integration areas smaller than the entire display. For this reason we made additional sets of Mondrians in which each piece

of the Mondrian is surrounded by an average controlling surround. As above we changed the illumination and the surround to reach the zero adaptation. With the spaced Mondrians we also have zero local adaptation. As above this does not matter⁷.

Quanta Catch at the Retina

Illumination




Figure 2. Three steps are necessary to perform color constancy experiments that can separate the underlying mechanism controlling color constancy. The traditional color constancy experiment demonstrates the small effect caused by making large changes in illumination that in turn cause very small changes in color appearance. By adding a third step to the experiment we can show that the controlling mechanism is the spatial comparison between each pixel and the maxima in each waveband. Further, it shows that grayworld and adaptation mechanisms are not involved in color constancy.

As well, we made spaced Mondrians in which the surround changed the local average while viewing in the original illumination. These surrounds had very little effect on color matches, showing that local adaptation is not involved in color constancy⁸.

Experiments looking for the effects of grayworld and adaptation cannot find them.

Maximov's Shoe boxes

Vadim Maximow wanted to test the effects of removing maxima from Mondrians. He devised the elegantly simple shoebox viewer (Figure 3a). The hole cut in the top allows illumination to fall on the small Mondrian at the end of the box. Colored filters controlled the intensity and the spectral distribution of the light. A viewing tube, made out of black plastic 35mm film containers, restricted the angle of view. Maximov's goal was to control the illuminants and the reflectances in the shoebox to overcome color constancy.

In principle it is easy to do (Figure 3b). Imagine two Maximov shoeboxes; one for the top Mondrian and one for the bottom. Select two filters that attenuate the color spectra, but do not reduce the light at any wavelength to zero. We used Wratten CC50R and CC50C. We measured the difference in color between the two filters. Find pairs of colored papers that have exactly the same color shift. It was impossible to find existing papers because the constraints were too difficult. It required that the papers be manufactured to fit the measurements. We used a Canon CLC 500 printer to make two Mondrians with identical spectral shifts.⁹

The result was that we could turn off the color constancy mechanism for complex images using a pair of Maximov shoeboxes. One (Top Fig 3a) had a Mondrian with 5 colors shifted towards red and viewed in pale cyan illumination; the other (Bottom Fig 3a) had 5 colors shifted towards cyan and viewed in pale reddish illumination. Despite the fact that the reflectances were different, the color appearances were the same. Ordinarily illumination has little or no effect. Why did Maximov's boxes turn off color constancy?

The answer is that the two Mondrians have to look identical because every pixel in the entire image had identical quanta catches as the corresponding pixels in the other image. Whenever two images have identical quanta catches everywhere, they look the same. We worked hard to find papers that all shifted the same amount. The reward was that by shutting off color constancy, we could experiment to see what new papers in the field of view would turn it back on.

Lesson: Independent R,G,B Normalizations

Figure 3c illustrated that the introduction of a white rectangle around the central patch revived color constancy. Two careful observations are important here. First, the whites do not look the same in both boxes. It looks reddish in the CC50R box and coolish in the CC50C box. The influence of the illuminant shift is visible. Second, the two sets of 5 original papers look almost the same as they do in the room side by side. They still have a reddish or coolish cast depending on the illumination. Nevertheless, the striking conclusion is that the introduction of white to both displays brought color constancy back to life.

Further experiments to test the hypothesis that humans normalize to the maximum in each waveband were performed in a similar set up.¹⁰ If the information from long-wave cones is normalized to the maximum L cone response, then white, and pure red, yellow and magenta will have the same effect. Our experiments supported this interpretation because all materials that provided a new maximum red restored color constancy.

Figure 3a.




Figure 3b.




Figure 3c.




Figure 3. Maximov shoeboxes are used to demonstrate that color constancy is controlled by the maxima in each waveband. Figure 3a shows the shoebox. Figure 3b illustrates that color constancy can be turned off by matching quanta catches at every pixel in the image. Figure 3c demonstrates that the introduction of a new maxima restores color constancy.

If the information from middle-wave cones is normalized to the maximum M cone response, then white, and pure yellow, green and cyan will have the same effect. Our experiments supported this interpretation because all materials that introduced new maximum green restored color constancy.

If the information from short-wave cones is normalized to the maximum S cone response, then white, and pure cyan, blue and magenta will have the same effect. Our experiments supported this interpretation because all materials that provided a new maximum blue restored color constancy.

The set of colors that did not revive color constancy were: black, gray, dark colors including reds, yellows, greens, cyans, blues and magentas.

Although the Color Mondrians experiments have made strong implications that humans independently normalize L-, M-, S-wave information to their maxima, these experiments provide direct experimental proof.

Separation and Enclosure

Normalize is a mathematical term that means scaling all values in a set by the maximum value. In principle mathematics does not care if the “normalization” process is performed by hand, slide rule, calculator or computer, even though the precision and speed of the calculation may vary. Mathematical normalization does assume that each member of the set be treated exactly the same as all the others. This is where the term normalize as applied to human visual processing gets into trouble. Does a white area have exactly the same effect on every other pixel in the field of view regardless of its location relative to the white?

Experiments measuring the influence of white show that the word “normalize” must be used with caution. Figure 4 illustrates six different displays made of large photographic transparencies. Each display has a 1.5° gray square. In each experiment the observer matches this central square to a standard display. Each display has a very dark background with an optical density of 3.0. The variable in this experiment is the placement of a fixed area of white. The first display at the left surrounded the gray with an equal width of white. Observers matched the gray to 1.5 Munsell units. The next experiment cut the white into four pieces and placed the on each side of the gray. The effect of this was to remove the white from the corners. Measurements showed a small increase in matching lightness to 1.8. The third experiment removed the white from one side. The match moved up to 2.1. The next two experiment moved the whites out to the corners. Lightnesses moved to 2.5 and 3.5. Finally, placing all the white on one side, lightness reached 3.5.

These experiments along with others that measure the effects of angular separation from whites¹¹ make the same point. Spatial positions of test patches relative to whites have large effects. Matches varied from 1.5 to 3.5 Munsell units which is 28% of the lightness difference between white and black. This variability was controlled by the enclosure of the area white area.

Lesson: Properties of Biological Normalization

It appears that for vision, the word “normalize” should never be used without a modifier such as “biological normal-




Figure 4. These six experiments illustrate that the appearance of the central gray area depends on the location of the white. This points out the importance of enclosure as a part of color appearance.

ization”. Although we have seen above there is considerable evidence that humans scale their retinal data relative to the maxima, there is as well considerable evidence that visual normalization does not treat each of the pixels the same. Spatial interactions are an important element in modeling lightness and color constancy.

Today’s Retinex Model

There has been remarkably little change in the fundamental operation of Retinex model (Figure 5) since first proposed in 1967 at Land’s Ives Medal Address to the Optical Society of America. The original proposal used the Ratio Product, Reset and Average. The original proposal also used a threshold operation on the Ratio step. The argument then was that reflectances had sharp edges and illumination edges were gradual. A threshold that removed small gradual changes in radiance would be of great value in modeling the B&W Mondrian (Figure 1).

Extensive psychophysical experiments have shown three important changes in theory. First, in real life scenes, illumination can have sharp edges and gradual changes in reflectance. The original hypothesis that the model could separate illumination from reflectance was wrong. Second, extensive quantitative experiments¹² showed that there is no single rate of change in radiance on the retina at visual threshold. In other words, we could not find psychophysical support for the threshold mechanism. Third, extensive experiments with models showed that the reset, “normalization” process was the mechanism predicting appearance in B&W Mondrians. These experiments showed the importance of controlling the number of interactions at each multiresolution level. If the calculation uses one iteration at each level the results show very strong local interactions. If the calculation uses hundreds of iterations the results show very strong global interactions, stronger than the data on enclosure (Figure 4) and separation require. The key factor in a successful Retinex model is the selection of the number of iterations at each level, which controls the degree of spatial interaction, that, in turn, mimics human vision.¹³

In 1980 Frankle and McCann introduced the multiresolution version that made real-time image processing possible. It is illustrated in Figure 5 (right).

Lesson: A Few Biological Operators

In reviewing the operation shown in the description of the Retinex model (Figure 5), we see that there are only three operations: ratio, product, reset and average. In implementing these calculations we have always converted the input to log radiance. The consequence is that ratio and product opera-

Ratio-Product-Reset-Average





Figure 5 (Left Diagram). The explanation of Ratio-Product-Reset-Average operation. Here we calculate the New Product (NP) for the output pixel x',y' . We begin at the starting pixel x,y using the Old Product (OP). All OP's are initialized with the maximum value for that waveband. The product of the radiance Ratios times the Old Product is reset if greater than the maximum and averaged with the previous New Products. The product of the radiance Ratios times the Old Product is reset if greater than the maximum and averaged with the previous New Products. Figure 5 (Right Diagram). An illustration of the Multiresolution aspect of the Retinex calculation. The calculation uses three data planes. The Old Product is initialized to the maxima. The original full-resolution image is illustrated as Input at the top. The input is averaged down to make a series of multiresolution planes ending with two pixels. This average Radiance image is the second data plane. The third data plane is for the output of each iteration and is called the New Product. Starting with two pixels we multiply the Old Product at the starting pixel and multiply it by the ratio of Radiances for the starting and output pixels. That product is Reset and averages with previous New Products at the output pixel. To get to the next level, the New Product is interpolated to twice the size and placed in the Old Product data plane. The Radiance data plane uses the next larger (8 by 2) average of the Input. The Ratio-Product Reset-Average calculation illustrated in Figure 5-Left are repeated. The Process continues until New Product at full-resolution is complete and is used as Retinex Output.

tions are simplified to subtraction and addition. Reset is another biologically friendly operation. Difficult operations such as multiplication and exponentiation have not been a part of the model. There has been a desire from the beginning in 1963 to restrict the operator to simple biologically friendly mechanisms.

Real Life Images

The B&W Mondrian had a white patch and black patch sending the same radiances to the eye. It was successfully modeled by many different generations of Retinex models.




Figure 6a. The scanned photograph of the sun/shade image. The white in the shade has the same digital value (119) as the black in the sun.

Later experiments with real life images 20 years ago demonstrated a scene with a boy holding a white card in the shade that had the same radiance as the black paper in the sun. Again, Retinex model created a new low-dynamic-range image displaying details of both sun and shadow areas.

Figure 6a (left) shows an input image scanned from a photograph of two Jobo targets: one in sun and one in shade. The photo was taken in Belmont, MA on a cool fall day without a single cloud in the sky. As on that day in Yosemite, the shadow was 32 times darker than the sun. The black in sun and the white in the shade both have 119 as the scanned input digit. Figure 6b (right) is processed by the algorithm described in




Figure 6b. The Retinex Output photograph of the sun/shade image. The white in the shade has a processed value of 175 compared the black in the sun with 126.




Figure 7. Logvinenko cubes pattern. Digital input on the left; Retinex output on the right. Retinex models can predict appearances that were previously attributed to cognitive behavior.

Figure 5. The process has left the sun image essentially the same: black in the sun has only moved from 119 to 126. However, the white in the shade has moved from 119 up to 175.

Lesson: Lightness Models and Gestalt Displays

Figure 7 shows a recent experiment by Alexander Logvinenko¹⁴. It is one of many experiments being studied with a revitalized interest in Gestalt visual phenomena. The left border of the picture shows that the input digits for the light and the dark diamond are the same. The right border shows that the lightness predicted for those diamonds are 122 and 167. When we translate digits to Munsell Values we find that retinex Output predicts a difference of about 2 Munsell lightness Units. Logvinenko measured a difference of 2.2 Lightness units.

The general conclusion is that the model evolved from the study of Mondrians can as well calculate appearances of both real life scenes or Gestalt phenomena.

Color Gamut

The strong argument running through these experiments is that humans calculate color using spatial comparisons. A variety of experiments show that the sum of errors (distances in color space) is a very poor predictor of the quality of a reproduction.¹⁵ In fact good reproductions make all their errors in similar directions.¹⁶

If that premise is true then spatial comparisons could be helpful in finding a set of in-gamut colors that look like the out-of-gamut original. Fig 8 illustrates the color gamut Retinex calculation. We begin with two input images, instead of one. We have the Goal image that has the large gamut. Second, we have the Best image that represents the limited gamut of the reproduction media. In fact we may substitute a three-dimensional LUT for the Best image, if the shape of the limited gamut is complicated. Again, we begin by averaging down each of the R, G, B separations to a small number of pixels for both the Goal and the Best image. We take the Old Product initialized to maximum and multiply by the Goal ratio. This New product is reset to the Best image or the Best data LUT. This process is repeated and the New Product values from this resolution are interpolated up to the next resolution. The process is repeated for R, G, B.

This process takes the spatial comparisons from the goal and limits the product by the Best image. The iterative process keeps reinforcing the ratios found in the goal while the reset forces the New product to migrate toward an image with all the same ratios, regardless of the absolute input values of the Goal image. The results show a big improvement in appearance compared to the Best. The data shows that this new image is in-gamut.

[See Retinex-Gamut calculated images on WWW version]

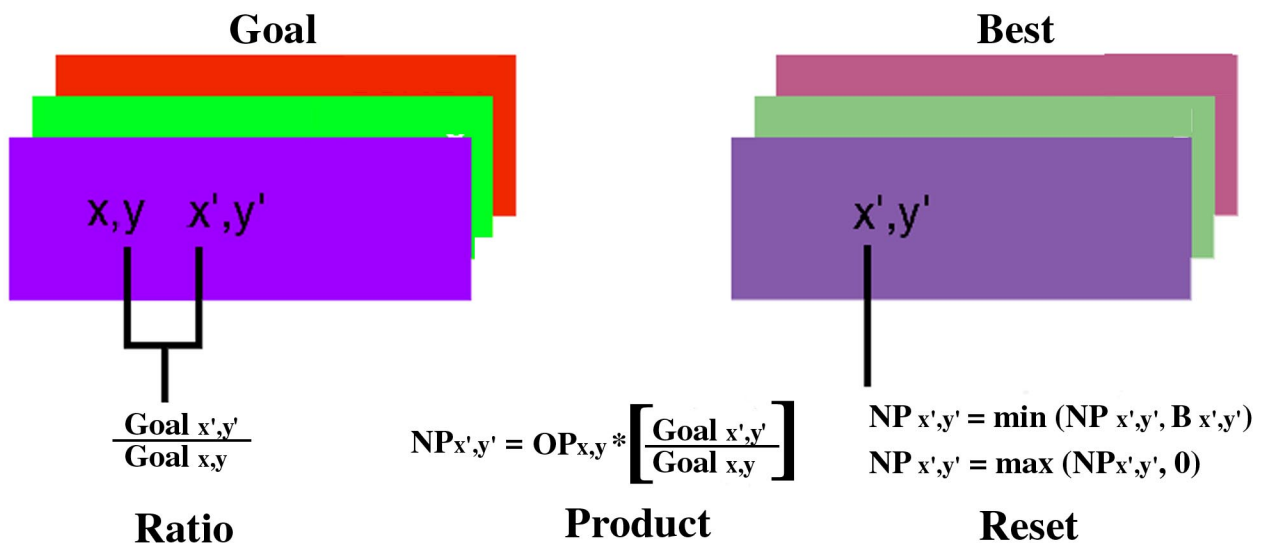


Figure 8. Schematic diagram of the Color Gamut Retinex Calculation. This calculation uses the Ratios from The Goal image and Reset from Best image to put spatial comparisons in the search for best reproduction using a limited gamut.

Lesson: Spatial Comparisons Predict Appearance and Find Best Color Gamut Compromise

The process of evaluating the absolute intensity of a pixel to see if it is in-gamut and replacing it with the nearest in-gamut point distorts images. Take two areas next to each other. Let us assume that one area is in-gamut and the other is not. If we leave the in-gamut pixel value unchanged while changing the out of gamut point, we have replaced the ratio of these two areas with a new ratio and a new color relationship to each other. It is a far better thing to change both pixel values so as to leave the spatial comparisons constant. The best reproduction is the one that reproduces the most spatial comparisons.

Conclusions

In 1861 Maxwell wrote his last paper on color; after that he turned his attention to electromagnetic fields.¹⁷ Land often quipped that the study of color vision would have been very different had Maxwell studied color after electromagnetic fields. For Land it was the other way round. Although his first published paper was on retinal rivalry¹⁸, his scientific interest had been focused on polarized light, polarized images, and instant images until he was 48. He accidentally observed colors in a red and white image that he could not explain. By then, he had several hundred patents on making images. He was fascinated because vision did not work like anything else; it is spatial in nature and cannot be explained one pixel at a time.

Acknowledgments

This work would not have been possible without two remarkable mentors, Edwin Land and Ansel Adams and many exceptional colleagues including Jeanne Benton, Lucretia Weed, Suzanne McKee, Tom Taylor, John Hall, Jon Frankle, Karen Houston, Jim Burkhardt and Jay Thornton and Mary McCann.

References

- ¹ N. Daw, Why After-Images Are Not Seen In Normal Circumstances, *Nature*, 196, No.4860, pp1143-1145, (1962).
- ² E. H. Land and J. J. McCann, Lightness and Retinex Theory, *J. opt. Soc. Am.* **61**, pg. 1, (1971).
- ³ H. Wallach, *J. Exptl. Psychol.* **38**, pg. 310, (1948).
- ⁴ E.H.Land The Retinex Theory of Colour Vision, *Proc. Roy. Institution Gr. Britain*, **47**. pp 23-58, (1974).
- ⁵ E. H. Land, The Retinex, *Am. Scientist*, **52**, pp.247-264, (1964).
- ⁶ 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**, pp. 445-458, (1976).
E.H.Land The Retinex Theory of Color Vision, *Scientific American*, **237**, No. 6 pp 108-128, (1977).
- ⁷ J.J. McCann, The Role of Nonlinear Operations in Modeling Human Color Sensations, *SPIE Proceedings, B. Rogowitz ed.*, **1077**, pg. 355-363, (1989).
J.J. McCann, Color Mondrian Experiments Without Adaptation, *Proceedings AIC, Kyoto*, p159 -162 (1997).

- ⁸ J.J. McCann, Magnitude of Color Shifts from Average-Quanta Catch Adaptations, *Proc. IS&T/SID Fifth Color Imaging Conference*, p215-220, (1997).
- ⁹ J.J. McCann, Rules for Colour Constancy, *Ophthalm. Opt.*, **12**, pp 175-177, (1992).
- ¹⁰ J.J. McCann, Color Constancy, Small Overall and Large Local Changes, *SPIE Proc.*,**1666**, pp.310-321, (1992).
- ¹¹ J.J. McCann and R.L. Savoy, Measurement of Lightness: Dependence on the position of a white in the field of view, *Human Vision, Visual Processing, and Digital Display II*, B. Rogowitz ed., *Proc SPIE*, **1453**, pp. 402-411, 1991. J.J. McCann, Scene Normalization Mechanisms in Humans, *Proc. IS&T/SID Second Color Imaging Conference*, p5, (1994).
- ¹² J.J. McCann, R.L. Savoy, and J. Hall, Jr., Visibility of Continuous Luminance Gradients, *Vision Research*, **18**, pp.891-894 (1978).
R.L. Savoy and J.J. McCann, Visibility of low spatial-frequency sine-wave targets: dependence on number of cycles, *J. opt. Soc. Am.*, **65**, No. 3 pg 343-350, (1975).
J.J. McCann and J.A. Hall, Jr. Effects of average-luminance surrounds on the visibility of sine-wave gratings, *J. opt. Soc. Am.*, **70**, No. 2, pp 212-219, (1980).
J. J. McCann, R. L. Savoy and J. A. Hall, Jr. Visibility of Low Frequency Sine-Wave Targets, Dependence on Number of Cycles and Surround Parameters, *Vision Research*, **18**, No.7, pp 891-894, (1978).
J.J. McCann, *J. Photogr. Sci.Eng.*, **22**, 64-68, (1978).
R.L. Savoy, Low Spatial Frequencies and Low Number of Cycles at Low Luminances, *J. Photogr. Sci Eng.*, **22**, No. 2, pp 76-79, (1978).
- ¹³ E.H. Land, An Alternative Technique for the Computation of the Designator in the Retinex Theory of Color Vision, *Proceedings of the National Academy of Science*, **83**, pp.3078-3080, (1986) *He described a variation on spatial processing using complex spatial ratio detectors. This model did not include a reset.. We prefer the original reset property for the reasons stated in the text.*
- ¹⁴ A.D. Logvinenko, *Perception*, in press, (1999).
- ¹⁵ J.J. McCann, "Color Theory and Color Imaging Systems: Past, Present and Future", *J. Imaging. Sci. and Technol.*, **42**, 70, (1998).
- ¹⁶ J.J. McCann, Making Reproductions of Masterpieces:
Managing All the Colors within the Gamut, *Proc. Multispectral Imaging and Color Reproduction for Digital Archives*, Chiba, (in press), 1999.
- ¹⁷ W. D. Niven, ed., *The Scientific Papers of James Clerk Maxwell*, Dover Publications, Inc., New York, 1965
- ¹⁸ E.H. Land, The Use of Polarized Light in the Simultaneous Comparison of Retinally and Cortically Fused Colors, *Science*, **83**, p. 309, (1936).

Bibliography

John McCann received his B.A. degree in Biology from Harvard University in 1964. He managed the Vision Research Laboratory at Polaroid from 1961 to 1996. His work concentrated on research in human color vision, large format instant photography and the reproduction of fine art. He is a Fellow of the IS&T. He is a past President of IS&T and the Artists Foundation, Boston. He is currently consulting and continuing his research on color vision.

This paper contains a number of color illustrations and photographs. It is also available on the web at

[www.imaging.org/ ?????](http://www.imaging.org/?????)

The web version also includes color samples of the Color Gamut Calculation.