Proc. of SPIE Vol. 1077, Human Vision, Visual Processing, and Digital Display, ed. B E Rogowitz (Aug 1989) Copyright SPIE

The role of simple nonlinear operations in modeling human lightness and color sensations

John J. McCann

Vision Research Laboratory, Polaroid
750 Main Street, Cambridge, Massachusetts, 02139

ABSTRACT

The purpose of this study is: first, to measure the properties of human vision under conditions that are subject to "chromatic adaptation"; second, to compare the human visual properties to a number of different Retinex models. In particular, the experiments study Gray-world properties of Mondrians. Models using averages of radiance or Gray-world assumptions do not show good correlation with observers. Nonlinear, reset models of lightness normalize each pixel to the maxima in each waveband. Triplets of normalized lightness do show good correlation with observer color matches.

1. INTRODUCTION

There have been a number of different analyses of Retinex models of human vision.[1,2] It is important to maintain the many distinctions among the many different calculations and objectives that are proposed as a Retinex model. Often, a conclusion about one of the Retinex implementations is true for the calculation described, but untrue for many other calculations in that class. The purpose of this paper is to discuss the goals, properties, and appropriate analyses of different Retinex models. In particular, the paper will discuss the properties of a nonlinear reset version of Retinex[3,4,5] that has been the subject of the majority of work comparing observer matches with calculated appearance.[6,7,8]

2. WHAT DOES A RETINEX MODEL CALCULATE?

The original work with "Mondrian" displays emphasized that models of vision should begin with the physical record of light everywhere in the field of view. It was not appropriate to be given the spectral properties of the illuminant if one was trying to model color constancy. The computer vision problem is to calculate the human response despite the variability in color and uniformity resulting from various illuminants. One of the important distinctions we need to make early in the modeling process is the exact goal of the calculation. Consider a typical outdoor scene. Fig. 1 shows a swimming float in a lake in New Hampshire, early in the morning. We will discuss three different characteristics of the side faces of the float. First, the physics of the light coming from the float to our eyes. Second, the sensation or appearance of the two sides; and third, the perception or recognition of the paint on surface of the wood.

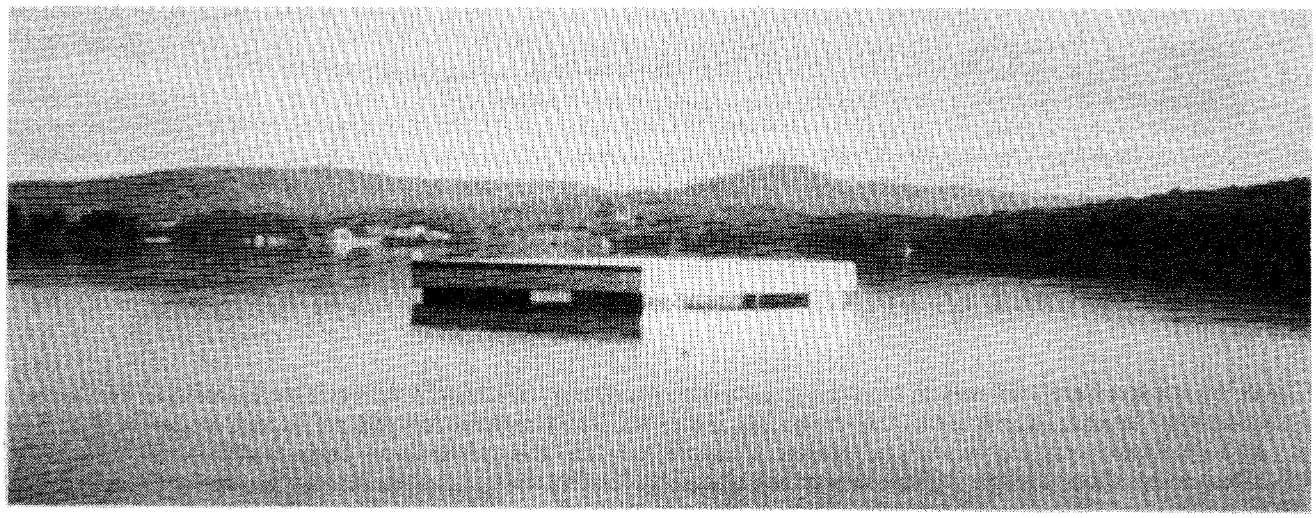

FIGURE 1. The New Hampshire swimming float discussed in the text. The color and amount of light coming from the two sides of the float is very different. The two sides appear nearly the same color. A human observer can recognize that the paint on the float is the same for both sides, even though the faces look somewhat different. Should a retinex model calculate the appearance (Sensation) or the reflectance (Perception)?

The first characteristic is the physical measurements. The sun is low in the sky and illuminates only one face of the float. The other face is illuminated by very blue skylight. The input to any models is the radiance at each wavelength from each pixel. The sunlit side is bright and has a color temperature of about 4000° K. The sky lit face is darker and is 20,000° K. The two faces of the float are very different physical stimuli.

To measure the second characteristic of the sides of the float we need to measure sensation. We can ask people on the beach two different questions about the float. First, we ask them to imagine that they are visual artists-painters. We ask them to pick, from a catalog of color mixtures, a sample to match each of the two faces of the float. Observers select a mixture of white with yellow for the sunlit face and a white with gray and blue for the sky lit face. In this case they have matched the sensation; they have matched the appearance.

To measure the perception we ask the observers to match the color of the paint on the float. In other words, what color paint is on the float? The observers pick the brightest, pure-white paint. In this case, the observers are matching perception—the recognition of the object despite the fact that the two faces appear to be different colors and have different sensations.

A model that calculates sensation must report that the two faces of the float are different. A successful sensation model must render differences in hue and visible gradients due to illumination. A successful model of perception must report that the two faces of the float are identical. Perception models have the goal of calculating the reflectance of the object and should not report appearances due to either illumination or visual phenomena, such as simultaneous contrast. These two goals—calculate appearance and calculate reflectance—are very different. Appropriate models for each need to have different properties to arrive at different solutions.

There has been an interest in algorithms that set out to calculate the actual physical reflectances of objects. There are many different approaches reported in the literature.[9,10] Some investigators developed algorithms that calculated the illuminant, then divided the scene radiances by the calculated illuminant to get the reflectances of objects in the scene.[11] Others[12] sought to solve the problem by assuming that all reflectances can be fit well by a small number of basis functions. Hurlbert and Poggio[1] have developed learning algorithms that seek to solve for the illuminant. Brou, Sciascia, Linden and Lettvin,[13] and Brill[14] have used overall normalization factors. Worthey[15] has used opponent processes. Lee[16] has written programs to exploit information in specular highlights. Each of these approaches represents an interesting formulation of the problem of calculating the reflectances of objects in the scene. This is an important computational problem, but an inappropriate objective for modeling human color appearance matches in complex images and for nonuniform illumination. If the algorithm succeeds in calculating reflectance, then all traces of the illumination will be removed from the computed image. If the objective is defined to be a computed version of what a painter would do, namely record the object's appearance in the illumination environment, then an image of calculated reflectances is undesirable. What is needed is an image that is made up of matches to what human observers see—color sensations. The computed image must include the information derived from the illumination environment.

3. WHAT ARE THE TOOLS OF A RETINEX MODEL?

In the 1971 *Lightness and Retinex Theory* paper, Land and McCann introduced the first computed lightness model. It used five computational steps applied along a one-dimensional path that wandered from pixel to pixel in a two dimensional array. The five components were ratio, product, threshold, reset, and average. Horn later proposed a two dimensional laplacian variation. Frankle and McCann[5] proposed a multi-resolution computation that was extremely efficient. In 1986 Land introduced a designator model that simplified the computation.[17] Brainard and Wandell,[2] and Hurlbert[1] have analyzed the properties of "Retinex" models of lightness.

3.1. Ratio-product

The ratio is the measure of the change in radiance between two pixels. Wallach[18] showed that the ratio of radiances correlated with the change in appearance of two adjacent areas. Wallach's experiments did not include complex images. Imagine a 50% reflectance center on a 100% reflectance surround. As well, in the same scene, imagine a 5% center on a 10% surround. The ratio at adjacent pixels is the same for 50/100 and 5/10. If all four papers are in close proximity to each other the pairs with the same ratio do not appear the same. In order to predict appearance in complex scenes we needed a mechanism that propagated the relative lightness information across the entire image. The product of ratios provided a mechanism that established the relationship of the 5%, 10%, 50% and 100% reflectances.

3.2. Threshold

A ratio-product calculation could establish the appearance of papers in uniform illumination. It, however, could not establish the reflectance relationship in nonuniform illumination. In the 1971 paper Land and McCann described a threshold mechanism. The idea was that there is a rate of gradual change on the retina that cannot be detected. If the smallest detectable edge[19] between two areas has a ratio of 1.003, then the ratio operation should ignore ratios smaller than 1.003. If all computations are performed in the model with pixels that correspond to cones with foveal spacing, then this threshold mechanism is quite powerful. If pixels one minute of arc apart cannot be distinguished for ratios smaller that 1.003 to 1.0, then the product of many such ratios over a thirty degree display cannot be detected for ratio of $(1.003)^{1800}$, or, a ratio smaller than 220 to 1.0. A small threshold propagating across a wide angle can remove a considerable gradient.

3.3. Nonlinear reset

The nonlinear reset introduced by Land and McCann had the function of normalizing each pixel in the image to the maximum. This property is an important one for modeling the color Mondrian experiments. The McCann, McKee and Taylor[6] (MMT) color matches show excellent correlation with the scaled integrated reflectance - a physical measurement that divides the radiance at each pixel by the maximum radiance in the image for each waveband. The 1971 Land and McCann reset mechanism does not search the image for the maximum and use that value as the denominator for all pixels, rather it resets the value of the product when it encounters a product greater than 1.0. This is an important property for modeling targets which exhibit local simultaneous contrast.

3.4. Average

The final step in the original process was the average of different paths that had different histories. It is the combination of nonlinear reset and average that allows the model to predict appearances of test targets demonstrating simultaneous contrast.

4. EXPERIMENTS TO MEASURE THE INFLUENCE OF AVERAGE RADIANCE

The present study, based on earlier work, tests the effect on changes in Gray-world conditions on Color Mondrians and Retinex predictions. It used five different Mondrian test targets. As with previous quantitative studies,[6] it used five different colors of papers, one for each Mondrian. These were gray, red, green, yellow, and blue. Except for the gray, all were near the color gamut of Munsell papers. Their Munsell notations are: N 6.75/, 10 RP 6/10, 5 Y 8.5/10, 2.5 G 7/6, and 2.5 PB 6/8. These displays were illuminated uniformly with three different narrow band lights; 630 nm (long), 530 nm (middle), and 450 nm (short). Next, the illuminant intensity for each of the five displays were chosen to compensate for the paper selected. The amount of long-, middle-, and short-wave illuminants were chosen such that the same triplet of radiances (L, M, S) comes from the gray paper in the initial illuminant; the red paper in illuminant 1; the green paper in illuminant 2; the yellow paper in illuminant 3; and the blue paper in illuminant 4. Finally, five different surround papers (See Fig. 2. - Surround Area A in **Surround Target**) were chosen for each of the five displays to compensate for the illuminant. The long-, middle-, and short-wave reflectances were chosen so that the same triplet of average radiances (AVL, AVM, AVS) came from the global average of all five targets. Thus we have constructed a set of five displays that have the same average over the entire field of view. Any measure of a global average or Gray-world average calculates the displays to be identical. In addition, we have a particular paper in each display that sends to the eye identical triplets of radiances. Just as in the earlier Mondrian experiments, the papers have different reflectances and compensating illuminants. The question is: to what extent does the equivalence of the Total Average Radiance (TAR) for all five targets influence the observers choice in matching color sensations?

A simple, global Gray-world model predicts that the five papers will appear the same because the radiance from each is identical and the average radiance from the entire field of view is the same.

A Ratio-Product-Reset model predicts that the average radiance will have a small effect, but basically the papers will appear very similar to their appearance in the initial illuminant and initial surround. Both surround and illuminant will change the appearance of the five selected papers, but as measured in previous quantitative experiments,[6] the magnitude of these effects is small. The nonlinear reset is the underlying operation that causes the Ratio-Product-Reset model to behave independently of the average properties of the entire field of view. It normalizes to the maximum in each waveband and is only secondarily responsive to the average properties of the image. The Ratio-Product-Reset model has usually been optimized for calculating sensations.. The goal is to calculate color appearance. As will be described later, it is the nonlinear reset along with the a low number of iterations that combine to enable the model to be responsive to local contrast phenomena.

4.1. Surround Targets Experiments

Table 1 describes the experiment. The first column identifies the area of the Mondrian used to measure the radiance. The second column identifies the reflectance of that area. The third column shows that the long-, middle-. and short-wave irradiances on the entire Mondrian were adjusted so that the radiances (column 4) are identical. The fifth column identifies the surround paper that was calculated to compensate for the different irradiances in column 3. The final column shows that the TARs for all five targets are the same.

Target	Reflectance (Munsell Notation)	Irradiance	Radiance	Surround	Total Average Radiance
TAR		IL0	L	Munsell	AVL
GRAY	N 6.75/	IM0	M	N 6.26/	AVM
[Area P]		IS0	S		AVS
TAR		IL1	L	Color-Aid	AVL
RED	10.0 RP 6/10	IM1	M	RVR	AVM
[Area G]		IS1	S	HUE	AVS
TAR		IL2	L	Munsell	AVL
YELLOW	5.O Y 8.5/10	IM2	M	5 Y 8/14	AVM
[Area C]		IS2	S	(Glossy)	AVS
TAR		IL3	L	Munsell	AVL
GREEN	2.5 G 7/6	IM3	M	7.5 G 6/8	AVM
[Area R]		IS3	S	(Glossy)	AVS
TAR		IL4	L	Color-Aid	AVL
BLUE	2.5 PB 6/8	IM4	M	B	AVM
[Area H]		IS4	S	Tint 2	AVS

Table 1.Papers, irradiances, radiances and total average radiances for both the Surround and the Local Surround targets.

We measured the integrated radiance from each area in the Mondrian, in each illuminant. We measured the integrated reflectance of many candidate papers in the appropriate triplet of illuminations. A computer program calculated the TAR for a large number of different papers. We normalized the calculation to the TAR of the gray surround target (TAR GRAY in Table 1) setting this value to 100%. The values for all the surround papers used in these experiments are listed in Table 2.

		Total Average Radiances in percent		
TARGET	SURROUND PAPER	LONG	MIDDLE	SHORT
TAR GRAY	N 6.26/	100.0	100.0	100.0
TAR RED	Color-Aid RVR Hue #2	100.0	106.4	98.2
TAR YELLOW	5 Y 8/14 Glossy	97.9	98.0	97.5
TAR GREEN	7.5 G 6/8 Glossy	95.6	97.5	98.6
TAR BLUE	Color-Aid B Tint 2	102.3	100.0	103.0

Table 2. Total Average Radiance for each waveband.

4.2. Local Surround Targets Experiments

One could argue that the human visual system uses a local average mechanism rather than a global one. This assumption is more difficult to test because a local average hypothesis requires a specification of a size, shape, and weighting function based on radius. One can develop alternative techniques to test the influence of local averages. Fig. 2 shows the **Local Surround Target**. It is made up of the same papers used in the **Surround targets** The papers are simply placed on the surround in a different location. Each area in the Mondrian is surrounded by the Surround A paper . Corresponding papers in all targets have the same shape and size. These new targets do not alter the physical measurements listed in Table 2; they only change the target in its local properties. Instead of surrounding the seventeen papers in the Mondrian with the paper, we changed the surround around each area of the Mondrian. This was done by simply placing each Mondrian paper in the center of an imaginary Mondrian twice as big.

We repeated the experiment described above and had observers match each area in each of the five **Local Surround** Mondrians. The five matches for the area in each target that sent the same radiance (L, M,S) to the eye are listed in Table 3 beside the matches from the Surround Targets.

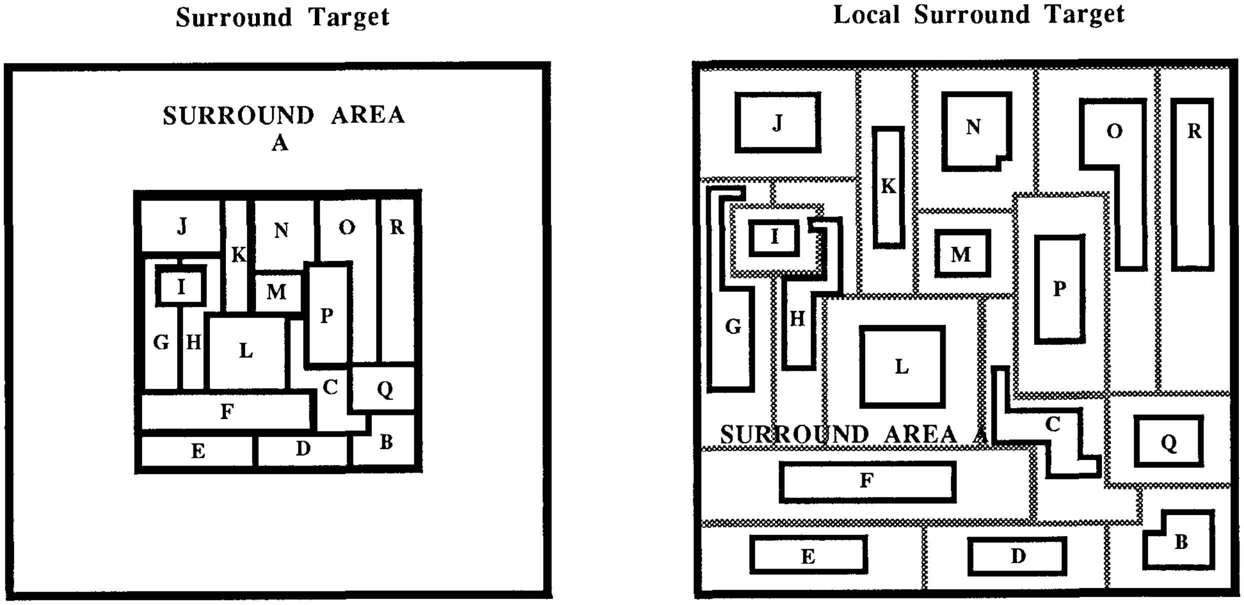

Figure 2. The two types of targets used in these experiments are diagramed in Figure 2. The left figure shows the **Surround Target**. Here the surround area is placed around the Mondrian. The right target is the **Local Surround Target**. Here the surround is placed around each individual area of the Mondrian. Solid lines represent the edges of papers. Dotted lines represent imaginary Mondrian of twice the dimensions.

4.3. Results

The experiments described below measure the influence of Total Average Radiance. The methods and procedures were the same as those described by McCann, McKee and Taylor.[6] Table 3 describes the results. The first column identifies the area of the Mondrian to be matched. The second column identifies the reflectance (Munsell designation) of that area. The third column shows the Munsell designation of the average chip in the Munsell book chosen to match the Mondrian area for the original MMT target. The fourth column shows the Munsell designation of the average chip in the Munsell book chosen to match the Mondrian area (column 1) for the **Surround** target.

These experiments test two simple hypotheses. First, can the color constancy found in the Mondrian experiments be explained by a mechanism that divides the radiance at each pixel by the average radiance of the total field of view? If this hypothesis is correct, then all five patches in Table 3, column 4 should appear identical. The radiances at each patch listed in Table 1 (fourth column) are all L,M,S. The total average radiances of the each display listed in Table1 (last column) are all AVL,AVM,AVS. All ratios are identical; all color matches should be identical. This hypothesis is clearly not correct. The observers chose 5 very different colors: gray N/6, red 2.5R7/4, yellow 5Y8.5/8, green 7.5 GY 7/4 and blue 10B 6/2. The second hypothesis is whether the color can be predicted by a three ratios; each ratio is the radiance at a pixel, in a given waveband,divided by the maximum radiance in that waveband. This hypothesis predicts that each patch listed in Table 1 will be different. Further it predicts that there will be no difference between the McCann, McKee, and Taylor experiments, the **Surround Targets** and the **Local Surround Targets**. The influence of the surrounds will have no effect on a ratio of a pixel to the maximum pixel.

There is very little difference between the matches for the McCann, McKee, and Taylor experiments (Table 3,column 3) and the **Surround Targets** (Table 3, column 4). MMT had large changes in total average radiance, while the **Surround Targets** had the same average. The observer chose chips the differed: for gray by 1 chip in chroma; for red by 1 chip in hue, value and chroma; for yellow by 1 chip in value; for green by 1 chip in hue ; for blue by 1 chip in hue and 2 chips in value and chroma. The above data supports the second hypothesis that observers pick papers consistent with normalization by the maximum in each waveband. It is difficult to asses whether the simple hypothesis is absolutely correct since the uncertainty of the experimental matches (1 chip in hue,value and chroma) is very close to difference found between experiments. Other experiments described later provide evidence that more than a simple normalization is required. Nevertheless, these experiments fail to support a Gray-world normalization and do support a maxima normalization.

Returning to Table 3, the fifth column shows the Munsell designation of the average chip in the Munsell book chosen to match the Mondrian area for the **Local Surround** target. The Surround A papers are the most saturated papers we could find. We placed them so as to completely surround each area of the Mondrian. Even with changing the surround as much as possible with papers, there was no significant change in the observers color matches. The observer matches for the **Surround Targets** and the **Local Surround Targets** differ: for gray by 0.25 chips in value, 1 chip in chroma; for red by 3 chip in hue, value and chroma; for yellow by 0.5 chip in value; for green by by 1 chip in hue, value and chroma; for blue by 1 chip in chroma.

The final column in Table 3 shows the Munsell designation of the chip in the Munsell book calculated to match the Mondrian area for the original target.[6] The nonlinear reset is the underlying operation that causes the Ratio-Product-Reset model to behave independently of the average properties of the entire field of view. It normalizes to the maximum in each waveband and is only secondarily responsive to the average properties of the image. The Ratio-Product-Reset model is able to calculate satisfactory predictions of observer color matches.

Target	**Actual Reflectance**	**Matching Chip MMT**	**Matching Chip SURROUND**	**Matching Chip LOCAL SURROUND**	**Computed Match MMT**
GRAY [Area P]	N 6.75/	5 YR 6/1	N 6.0/	N 5.75/	N 6/
RED [Area G]	10.0 RP 6/10	5 R 6/6	2.5 R 7/4	5 RP 7/4	5 RP 5/4
YELLOW [Area C]	5.O Y 8.5/10	5 Y 8/8	5 Y 8.5/8	7.5 Y 8.5/12	7.5 Y 7/8
GREEN [Area R]	2.5 G 7/6	10 G 7/4	7.5 GY 7/4	5 GY 8/2	5 G 6/6
BLUE	2.5 PB 6/8	2.5 PB 4/6	10 B 6/2	10 B 6/4	7.5 B 6/6

Table 3. Papers in the Mondrians chosen to match in the Surround Targets, the Local Surround Targets, the original McCann, McKee, and Taylor experiments, and the computed predictions made by McCann, McKee, and Taylor[Area H]

5. WHAT DOES PSYCHOPHYSICS REVEAL ABOUT HUMAN VISION?

5.1. Gray-world Assumption

The results in Table 3 show very little change in appearance due to the presence of a **Surround** or the presence of a **Local Surround.** Recall the uncertainty of the experimental matches is roughly 1 chip in hue,value and chroma. For all three versions of the targets(no surround, Surround and Local Surround) the observers' matches for the gray patch span a range of 0.25 units in value and 1 unit in chroma. The observers' Munsell chip matches for the red patch span a range of 4 pages in hue, 1 unit in value and 2 units in chroma. The observers' matches for the yellow patch span a range of 1 page in hue, 0.5 units in value and 4 units in chroma. The observers' matches for the green patch span a range of 2 chips in hue, 1 unit in value and 2 units in chroma. The observers' matches for the blue patch span a range of 1 chips in hue, 2 units in value and 4 units in chroma. These matches do not show significant changes in appearance due to changes in Gray-world properties.

In the MMT experiment the Gray-world averages were all as different as the illuminants. In the **Surround Targets** they were all the same. These experiments did not show a significant dependence on the average of the target even with the **Local Surround Targets**. These papers are the most saturated papers available. We placed them so as to completely surround each area of the Mondrian. Even then, there was no significant change in the observers color matches. Table 3 is strong evidence that a model whose goal is to calculate color sensation be essentially independent of averages. The ratio-product-reset model has this property because of its nonlinear reset. The model normalizes the image with respect to its maximum value, not an average.

5.2. Simultaneous contrast

Why is it that changes in local surround had such a small effect on these Mondrian observations? All of us have seen simultaneous contrast demonstrations in which a change of a background has produced large changes in the sensation of a center patch. First, the most dramatic departures from constancy are due to global normalization. Gelb's classic experiment is the most dramatic example of simultaneous contrast.[20] Here a black piece of paper is the only object in the field of view, by itself it looks a dim white, but unquestionably white. When a white paper is put beside the black paper the white looks white and the black looks black. This is an example of a global normalization process—one that changes a single object from white to black. Another familiar spatial experiment gives us important information about the limits of global normalization. Consider the gray-square-on-white and gray-square-on-black demonstration of simultaneous contrast. If global normalization of the entire field of view were complete, we would expect that observers would report the two gray squares with identical reflectances would have the same appearance. If local mechanisms were the only consideration, then the gray square in the black surround should mimic the results found in the Gelb experiment, and should appear a white, since it is the maximum intensity in the local area. Observer results give important information about the relative importance of global and local interactions. The gray square in the black surround is one value unit out of nine lighter than the same gray in the white surround. If local spatial calculations were the only consideration, the gray in black should appear with a value of 9.0. If global spatial considerations were the only consideration, the gray in black should appear with a value of 5.0 The observer matched the gray in black to a value 6.0 In other words, the spatial normalization mechanism is an imperfect global mechanism. Alternatively, it is a local mechanism that is significantly influenced by information from the entire image. What one draws from both sets of experiments is that the most powerful examples of local influence of a surround are found in situations in which the extent of the surround is large enough to influence the global normalizing mechanisms. The best examples of simultaneous contrast, such as Gelb's experiment, involve changes in the maxima in the image. Experiments that change only the properties of local portions of the image produce small changes in an observer's match—only one or two chips in the Munsell Book.

5.3. Overall Brightness Effects

McCann, McKee and Taylor showed a small but consistent shift in the color matches due to changes in overall illumination. Their data showed that a change in intensity by a factor of 4 caused a change in lightness (on a scale of 0 to 10) of 0.8 units for long and middle-wave lights and 0.6 for short-wave light. As well, numerous other experiments show a corresponding shift in lightness as a function of overall illumination [20-21]. Compared to the color shifts created by color constancy mechanisms, those created by overall brightness are small. Nevertheless, compared to the imperfections in color constancy they can be significant

6. DISCUSSION

Hurlbert[1] has reviewed many of the retinex algorithms. Hurlbert assumes that the average reflectance of each scene in each waveband is the same; that is, she has assumed the Gray-world hypothesis. These Gray-world models are designed to calculate the reflectance, rather than the sensation. If human vision used a Grayworld mechanism, then all the color matches in Figure 3, column 4 and column 5 would be identical The above experiments showed that the human observer does not exhibit Gray-world properties when matching Color Mondrians.

Brainard and Wandell[2] criticize retinex models because they are too sensitive to changes in the color of nearby objects. They argue that the new designator calculation described by Land is equivalent to the energy at a pixel divided by the average of all pixels. They argue that the calculated value of the designator model for each pixel will be influenced by the papers in the scene. As mentioned above a local average model is very difficult to test without quantitatively evaluating the specific properties of the local mechanism. It is beyond the scope of this paper to evaluate Land's designator model.[17] In their Fig 2, Brainard and Wandell vary the reflectance of papers in Mondrians. They used an average-sensitive retinex algorithm and predicted large changes in color. As in the color measurements described above they report virtually no change in color appearance. When discussing this evaluation with Brainard and Wandell, we provided actual calculations of their proposed display targets because we did not know how to write an analytic equation for the nonlinear, image dependent calculation. We measured the integrated reflectance of their papers. We calculated the integrated reflectance (Table 4) using the Ratio-Product-Reset model described by McCann and Houston.[7] These calculations use the nonlinear reset. The Ratio-Product-Reset model predictions are in agreement with the observer's results suggested by Brainard and Wandell. In other words, the Ratio-Product-Reset Retinex model's predictions are relatively insensitive to changes in averages and the choice of papers in the Mondrian.

	R 8/4			B 8/2			G 8/2			10 YR 5/6			Y 8/2			P 6/8		
	L	M	S	L	M	S	L	M	S	L	M	S	L	M	S	L	M	S
B	90	88	87	85	93	91	90	93	84	45	39	17	90	91	77	65	58	90
G	90	88	87	85	93	91	90	93	84	40	35	15	90	90	75	57	48	90
Y	90	88	87	85	93	91	90	93	84	40	35	16	90	91	76	56	50	90
P	90	88	87	85	93	91	90	93	84	39	34	13	90	91	67	60	53	84

Table 4. Nearly constant Integrated Reflectances predicted by the Ratio-Product-Reset model for Brainard and Wandell "Fig 2" experiment.

In the same paper Brainard and Wandell suggested a second set of experiments. Here, they designed 3x3 Mondrians with 7 low-reflectance papers. They then introduced 2 different highly colored, high-reflectance papers. They normalized by the maximum in the field of view. They reported predictions that changed from Mondrian to Mondrian.

Two important questions remain unanswered in the analysis of this second experiment. The first question is that there are no experimental measures of the appearance of the patches in these Mondrians. As described above in section 5.2 the most dramatic departures from constancy are due to global normalization. These proposed Mondrians introduce new global maxima. Further, the proposed second set of 3x3 Mondrians will have different overall brightnesses. As described in 5.3 these brightness changes will cause small changes in appearance. Quantitative measurements of these Mondrians are needed determine the degree of color constancy.

The second question is that Brainard and Wandell "Reset Retinex" model represents accurately the original reset models. Brainard and Wandell chose to analyze the Ratio-Product-Reset model under conditions significantly different from those tested as predictions for color matching experiments.[6,7,8] Brainard and Wandell "assumed the comparison list is infinitely long and that all locations are repeatedly compared with location x [ref. 2, page 1658]." This choice of parameters was motivated by a desire to express the process as an analytical formula. Despite the mathematical advantage of assuming infinite iterations, the real disadvantage is that it does not represent the intended model under discussion. There are two important reasons why the actual Ratio-Product-Reset model uses very few iterations. First, the process was designed to be highly efficient and to use as few iterations as possible. In McCann and Houston there are only eight iterations at each resolution two each in four directions. The values computed by this intentionally nonlinear model are very different with 56 iterations compared to infinite iterations. The second reason for using few iterations is that it installs in the model the dependence on local characteristics of images. As described in MMT, the goal of being able to predict simultaneous contrast targets requires a strong, but not dominant, global influence. Very long paths or many iterations of a multi-resolution process are not appropriate when modeling human color sensations.[6] Finally, corrections for overall brightness need to be added to the calculation to fairly represent the original model.

The goal of calculating sensation is clearly different from the goal of finding the reflectance of objects.[22] If one chooses the goal of calculating the physical reflectance there is not enough information to arrive at a proper solution.[1] If one states that the goal is to calculate the appearance of an object, then there is adequate information in the two-dimensional array of radiances. The sensation problem is simpler. One no longer has to differentiate gradients in illumination from gradients in reflectance or edges in illumination from edges in reflectance; they are treated the same. Both illuminants and reflectances are incorporated in color matches. There is sufficient information to calculate a color match (color sensation).

7. ACKNOWLEDGMENTS

The author wishes to thank John A. Hall and Tom Taylor for their effort in calculating, making and testing these displays.

8. REFERENCES

1. A. Hurlbert and T. Poggio, "Formal connections between lightness algorithms," J. Opt. Soc. Am. A **3**, 1684-1693 (1986).
2. D. H. Brainard and B. A. Wandell, "Analysis of Retinex theory of color vision," J. Opt. Soc. Am. A **3**, 1651-1661 (1986).
3. E. H. Land and J. J. McCann, "Lightness and Retinex Theory", J. Opt. Soc. Am., **61**, 1-11 (1971).
4. E. H. Land, L. Ferrari, S. Kagen, and J. J. McCann, "Methods and apparatus of lightness imaging", U.S. Patent 3,651,252.
5. J. Frankle and J. J. McCann, "Method and apparatus of lightness imaging," U.S. Patent 4,384,336.
6. 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).
7. J. J. McCann and K. L. Houston, "Calculating color sensation from arrays of physical stimuli," IEEE **SMC-13**, 1000-1007 (1983).
8. J. J. McCann, Hall, and E. H. Land, "Color mondrian experiments:the study of average spectral distributions", J. Opt. Soc. Am. (A) **67**, 1380 (1977).
9. D. Marr, "The computation of lightness by the primate retina," Vision Res. **14**, 1377-1388 (1974).
10. B. K. P. Horn, "Determining lightness from an image," Comp. Gr. Img. Proc. **3**, 277-299 (1974).
11. G. Buchsbaum, "The retina as a two dimensional detector array in the context of color vision theories and signal detection theories," Proc. Inst. Elect. Electronic. Engrs. **69**, 772-786 (1981).
12. L. T. Maloney and B. A. Wandell, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," J. Opt. Soc. Am. A **76,** 1673-1683 (1986).
13. P. Brou, T. R. Sciascia, L. Linden, and J. Lettvin, "The color of things," Sci.Am. **255**, 84-91 (1986).
14. M. H. Brill, "Further features of the illuminant-invariant trichromatic photosensor," J. Theor. Biol. **78**, 305-308 (1979).
15. J. A. Worthey, "Limitations of color constancy," J. Opt. Soc. Am. A **2**, 1014-1026 (1985).
16. H. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights," J. Opt. Soc. Am. A **3**, 1694-1699 (1986).
17. E. H. Land, "An alternative technique for the computation of the designator in the retinex theory of color vision", Proc. Natl. Acad. Sci. USA, **83**, 3078-3080 (1986).
18. H. Wallach, "Brightness Constancy and the nature of achromatic colors" J. Exptl. Psychol. **38**, 310 (1948).
19. H. R. Blackwell, "Contrast threshold of the human eye", J. Opt. Soc. Am. 56, 624-637 (1946).

20 C.J.Bartleson, ,and E.J.Breneman, Brightness perception in complex fields, J.Opt. Soc. Am.,**57,** 953-957, (1967).

21 D. Jameson and L.M. Hurvich, Complexities of perceived brightness, Science, **133**,174,179, (1961).

22. J. J. McCann and K. L. Houston, "Color Sensation, Color Perception and Mathematical Models of Color Vision," in: Colour Vision , J.D. Mollon, and L.T. Sharpe, ed. , Academic Press, London, 535-544 (1983).