# Using Color Constancy to Advantage in Color Gamut Calculations

*John McCann*
*McCann Imaging*
*Belmont, Massachusetts, USA*

## Abstract

The human color constancy uses spatial comparisons. The relationships among neighboring pixels are far more important than the absolute differences between the colorimetric values of an original and its gamut-limited reproduction.

If all the pixels in an image have a reproduction error in the same direction (red, green, blue, lightness, hue, chroma), then our color constancy mechanism helps to make large errors appear small. However, if all the errors are randomly distributed, then small errors appear large. This paper will describe experiments using constant errors to produce variable apparent errors and describe a technique of calculating the best appearance image using spatial comparisons. This calculation will be applied to color-gamut problems.

## Introduction

The Retinex model for estimating apparent lightness was proposed by Land and McCann in 1967.[1] It was applied to color constancy by McCann, McKee and Taylor in 1976.[2] In the early 1980's Frankle and McCann[3] extended the ration-product-reset-average operation to highly efficient multi-resolution image processing. Recently, Retinex has been extended for use in calculating the closest color appearance in situations in which the reproduction is made with media having a smaller color gamut than the original.[4]

This paper describes the experiments that lead up to the new gamut Retinex calculation and discusses the results of a sample calculation. The central theme here is that the underlying mechanisms that control color constancy can be used to advantage to make images in a small color gamut resemble images in a large color gamut. These color constancy mechanisms are spatial comparisons between pixels. They show almost no dependence on the L, M, S triplet of radiance at a pixel. It follows that color gamut transformations calculated one pixel at a time produce poor color reproductions, while those done using spatial comparisons give far better results.

## Color Gamut Calculations using a Two-Area Mondrian

Figure 1 illustrates the choices of fundamental processes used in calculating the best color compromise for a limited gamut reproduction.[5] For simplicity we will study only the red record printed in cyan ink. In this scheme, red radiance = 100 is the absence of cyan dye (reflecting the most red light), while red radiance = 0 is maximum cyan dye (reflecting a minimum of red light). The first column shows the original with radiance A=89 from the left area and radiance B=95 from the right area. The second column shows the case in which the color gamut of the reproduction is as large as the original. In this fairly rare case, the reproduction radiances Ar = A and Br = B. That also means that Ar/Br =A/B.

The third column illustrates usual color-gamut transformations. Here the dye sets are such that reproduction radiance Ar =89= A is in-gamut. However, reproduction radiance Br=85 is the best possible compared to the original radiance, B=95. In this example we substitute the closest value, namely 85. This approach conserves the colorimetric values X, Y, Z for area A. It leaves all in-gamut pixels unchanged, thus minimizing the cumulative color distance between original and reproduction for all areas. This choice has a highly adverse effect on the ratios. The original ratio is A/B=89/95. The reproduction ratio has the value Ar/Br=89/85. The reproduction now reports that A is lighter than B, while the original reported that B is lighter than A. Such reproductions with distorted edge ratios make poor reproductions.

The fourth column illustrates the principle of conserving spatial ratios. Here we have the same limitations on gamut. However, the limited gamut of B = 85 causes a shift of A to 80, so as to conserve spatial ratios. Here Ar/Br = A/B is the controlling principle. This mechanism requires an adjustment for area Ar from 89 to 80 even though the 89 is in gamut. This strategy increases the cumulative colorimetric distance for all areas between the original and the reproduction, yet looks better.

## Color Appearance vs. Color Distance

In order to illustrate this last point that colorimetric distance is a poor predictor of best color appearance we can

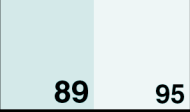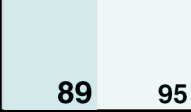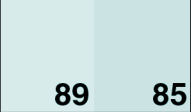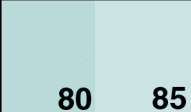## Should reproductions conserve XYZ or spatial ratios?

| Original | Perfect Reproduction | Conserve XYZ | Conserve Spatial Ratios |
|---|---|---|---|
| 89    95 | 89    95 | 89    85 | 80    85 |
| | Ar = A<br>Br = B<br>Ar/Br = A/B | Ar = A<br>Br = .9B<br>Ar/Br ≠ A/.9B | Ar = .9A<br>Br = .9B<br>Ar/Br = A/B |

*Figure 1 illustrates three different approaches to making reproductions. The first column illustrates the original image with two areas with radiances A and B. The second column illustrates the rare occation in which the reproduction media has the same color gamut as the original. Here the the two reproduction radiances Ar and Br equal those (A and B) of the Original. Both conservation of XYZ and conservation of spatial ratios are successful. The third column illustrates the approach of Colorimetry. Here each area is treated separately. The reproduction of the left area equals the original (Ar =A). However, the gamut limit found in this example restricts Br to 90% of B. By selecting the best fit for the right area and the left area independently we alter the ratio Ar/Br compared to A/B. Now the reproduction reports that the left area is lighter that the right. The fourth column on the right illustrates the conservation of spatial ratios. Here the limit of 90% on area B is also applied to the area A. The effect of keeping ratios constant, is that we have made both the areas darker. Despite this change in radiance, the relationship of areas A and B are the same.*

study Figure 2. Here we have made a set of nine different "Umbrella" displays. We call the one in the middle [E] the Original. All the others are reproductions. The observers' task is to judge which of the 8 different reproductions are acceptable and which are not.

**Design of Umbrella Targets**

The specifications for all areas are in Munsell notation. This avoids the problems introduced by the inconsistencies found in L*a*b* spacing.[6] The "Original" [E] has 10 triangular areas. The top area has the same L*a*b* values as Munsell paper 5R 7/6. The remaining 9 areas are all 7/6, each of them are 4 pages apart [5R, 5YR, 5Y, 5GY, 5G, 5BG, 5B, 5PB, 5P, 5RP]. Thus we have placed the 10 areas in a circle of constant lightness and chroma in color space.

All 10 colors in all the remaining 8 targets differ from the original by only one chip in the Munsell Book. The [DEF] series varies in lightness. Each corresponding area in [D] is 8/, and is 6/ in [F]. This change in stimulus is a global shift. It produces noticeable, but acceptable changes in appearance. The changes for all 10 areas are about the same magnitude.

The [AEI] series varies in chroma. Here all 10 papers have /4 in [A] and /8 in [I]. The chances are noticeable, but regarded as acceptable reproductions.

The [BEG] series varies in hue. Here 5R is replaced with 2.5R in [B] and 7.5R in [H]. The same shift was ap-

plied to all 10 areas. The chroma /6 was chosen so that one hue page in the Munsell book equals one chip in lightness value and one chip in chroma. Here the change in color appearance is small, but noticeable and acceptable as good reproductions.

So far all the changes have been global shifts. All 10 papers have moved in the same direction in color space. Such uniform movement is reminiscent of the changes found in color constancy experiments in which the paper display is constant, while the intensity and color of the illumination changes. In color constancy experiments we are familiar with this kind of result, namely that colors change very little with large global shifts in illumination.

The final [CEG] series shows the effect of pseudo-random color shifts. As above, each color patch is one Munsell chip different from the Original [E]. Nevertheless, independent changes produce both large and small visual changes in appearance. These local area changes make for poor reproductions. They distort the appearance of the original in way that make them unacceptable reproductions.

The analogy to color constancy is very compelling. Changing the colors independent of the neighbors disrupts the spatial ratios (Fig 1). Changing the local rations in the context of a color constancy experiment is the same as changing the reflectances of the areas. Changing reflectances of individual papers (local shifts) causes big changes in appearances, while changes in illumination (global shifts) cause small changes. It should be noted that

*Figure 2 shows an "Original" display [E] surrounded by 8 reproductions (see CD-Rom for color images). [Because the Procedings are printed in grayscale, we have included the r, g, b separations. In grayscale the three separationa add to a nearly uniform gray image.] In all cases there are 10 pie-shaped color patches. For all 8 reproductions each individual patch differs from the original by one chip in the Munsell Book of Color. That means that, regardless of the color the difference in color appearances, each of individual chip is a constant distance from the original in the Munsell uniform color space. Thus, changes in lightness, chroma and hue are equal. Along the left-to-right axis [D-E-F] the Original and reproductions vary only in lightness. All 10 patches in D are one Munsell chip lighter than the original. All 10 patches in F are one Munsell chip darker than the original. Along the top-to-bottom axis [B-E-H] the Original and reproductions vary only in hue. All 10 patches in B are shifted counter-clockwise one Munsell chip from the original. All 10 patches in H are shifted clockwise one Munsell chip from the original. Along the upper-left to bottom-right axis [A-E-I] the Original and reproductions vary only in chroma. All 10 patches in A are one Munsell chip less saturated than the original. All 10 patches in I are one Munsell chip more saturated than the original. All of the above reproductions [ABDEFHI] are still reasonable reproductions despite the one chip color shifts. The remaining two reproductions C and G are examples of individual color shifts in hue, lightness and saturation. Unlike systematic shifts, individual shifts create unacceptable distortions of the original. An error of 1 Munsell chip is acceptable if it is global, but not if it is local.*

changing the spectral character of the illumination is not the direct analog of changing lightness, hue and chroma, as we did in the above umbrella experiment. The direct analog of color constancy experiments would be to uniformly lower or raise the long-wave reflectances of all the papers, and do analogous things to the middle-, and short-wave reflectances. However, we began this experiment with the design of substituting paper of known color difference. That has been provided by Munsell Book data. Shifting long- middle- and short-wave reflectances may be a better experiment, but equal color difference data is not available.

The suggestion from these experiments is that color gamut calculations, using spatial comparisons can lead to better in-gamut reproductions. Colorimetrically they will have larger color difference errors, but they will look better. The same color constancy mechanism that reduces large physical shifts in illumination to small, but noticeable appearances can be employed to make gamut-limited reproductions better.

## Today's Retinex Model

There has been remarkably little change in the fundamental operation of Retinex model (Figure 3) 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.

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 experiments[7] showed that there is no single threshold 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. More details on Retinex are available in reference 4. In 1980 Frankle and McCann introduced the multi-resolution version that made real-time image processing possible. It is illustrated in Figure 3 (right).

In reviewing the operation shown in the description of the Retinex model (Figure 3), we see that there are only four 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 operations are simplified to subtraction and addition. Reset is a simple logical operator.

## 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, starting with McCann, McKee and Taylor. 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. Recent images include 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. The process has left the sun image essentially the same: black in the sun has only moved from digit 119 to 126. However, the white in the shade has moved from 119 up to 175.

Recent experiments by Alexander Logvinenko[8] illustrate experiments being studied with a revitalized interest in Gestalt visual phenomena. The input digits are the same for apparently light and the dark diamonds. The calculated lightnesses 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 models evolved from the study of Mondrians can as well calculate appearances of both real life scenes or Gestalt phenomena. Examples of these images can be found in a recent summary publication.[4]

## Color Gamut

The hypothesis connecting 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.[9] In fact good reproductions make all their errors in similar directions.[5]

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 4 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. If the shape of the limited gamut is complicated, we may substitute a three-dimensional LUT for the Best image. 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 Goalin ratio. This New Product is reset to the Bestin 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.
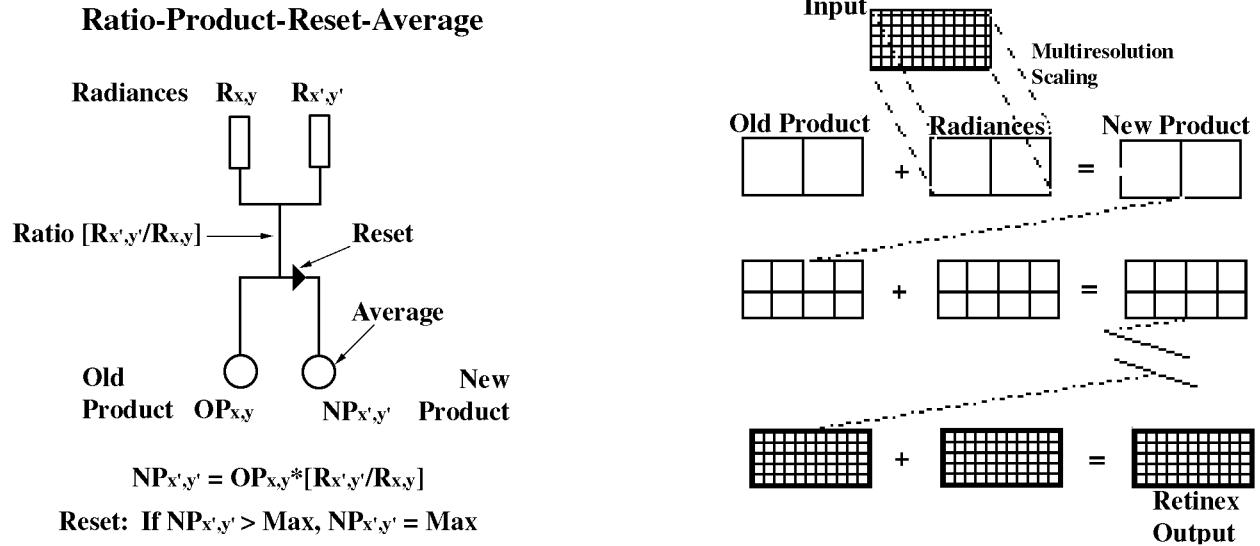
## Ratio-Product-Reset-Average

$$NP_{x',y'} = OP_{x,y} * [R_{x',y'}/R_{x,y}]$$

**Reset: If $NP_{x',y'} > Max$, $NP_{x',y'} = Max$**

*Figure 3 (LeftDiagram). 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 value in the Best image 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. Figure 3 (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 Best image. 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 Old 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 3-Left are repeated. The Process continues until New Product at full-resolution is complete and is used as Retinex Output.*

## Goal       Best

**Ratio**
$$\frac{Goal\ x',y'}{Goal\ x,y}$$

**Product**
$$NP_{x',y'} = OP_{x,y} * \left[\frac{Goal\ x',y'}{Goal\ x,y}\right]$$

**Reset**
$$NP_{x',y'} = \min(NP_{x',y'}, B_{x',y'})$$
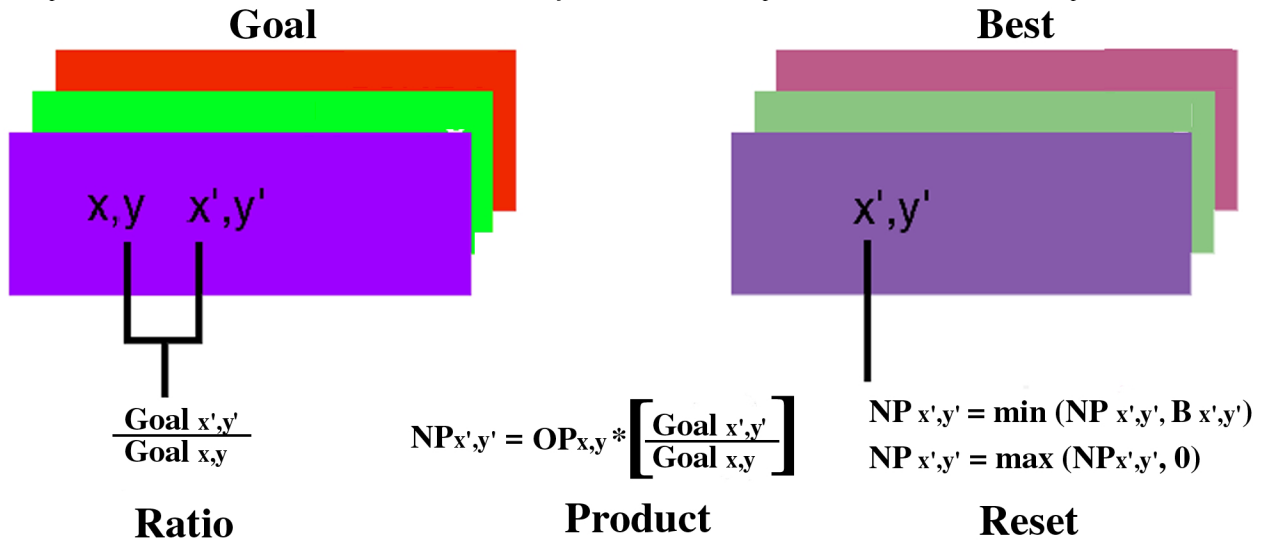$$NP_{x',y'} = \max(NP_{x',y'}, 0)$$

*Figure 4 shows the 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. The specific calculations used to make Fig5 used a 384 by 256 image that was averaged down to 3 by 2 pixels. The Ratio, Product, Reset, Average process (Fig3 left) was used for 4 cylce of comparisons in 8 directions. The interpolation process (Fig3 right) was used to make 6 by 4 images. The process was repeated up to 384 by356 NewProductOut image.*
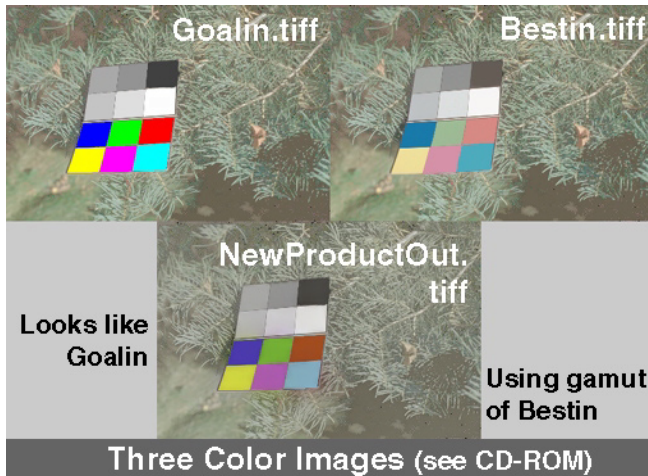
*Figure 5 shows the Goalin, Bestin and NewProuctOut color images. Fig 5a shows the 3 color images(See CD-ROM). Here we see that the NewProuctOut image looks much closer to the Goalin image than the Bestin. Both Bestin and NewProductOut are within the limted gamut of theToyo inks (uncoated). Figures 5b,5c & 5d show the blue, green and red color separation images. By comparing the NewProductOut image with the Goalin image we can see that the gamut retinex product preseved the spatial relationships in each separation. The gamut mapping process that modified the Goalin image to make the Bestin image, adjusted each pixel independent of all the others ( See Fig1b). The NewProuductOut image used spatial comparisons and gamut information to calculate the value at each pixel (Fig1c). The improvement in maching appearances can be seen by studying the three set of color separations. In each image the 12 patch Jobo target shows us the rendition of 6 gray areas in the top two rows and the colors blue, green and red above yellow, magenta and cyan in the bottom two rows. In Fig 5b the blue separation the Goalin colors of blue, green, red are white, black, black; while yellow, magenta, cyan are black, white, white. The spatial gamut retinex process finds in-gamut combinations that maintain, as well as possible, the lightness differences in each color separation. In this record the Goalin whites have been limited to NewProductOut light grays. However, the blacks associated with red, green and yellow have remained black. The comparison with the Bestin image is very interesting. The non-spatial, one pixel at a time process found blue separation values much lower in contrast, in fact all the lighness are close to middle gray. In Fig 5c the green separation the Goalin colors of blue, green, red are black, white, black; while yellow, magenta, cyan are white, black, white. Again the spatial gamut retinex process finds in-gamut combinations that maintain, as well as possible, the lightness differences in each color separation. In this record the Goalin whites have been limited to NewProductOut light grays for green and cyan. However, the blacks associated with red, blue and magenta have remained very close to black. The comparison with the Bestin image shows that the non-spatial, one pixel at a time process found green separation values much lower in contrast. In Fig 5d the Goalin whites have been limited to NewProductOut light grays for yellow and light-middle gray for red and magenta. The rendition of the blacks have been imited to dark gray for both green and cyan. The comparison with the Bestin image shows that the one pixel process found green separation values much lighter than desired. The red separation is the worst of the three, but the the NewProductOut image distorts the appearance of the green much less than the Bestin.*

*Figure 5b shows the blue sepatations. Note that the processed New ProductOut has lightnesses that resembles the Goalin.*

*Figure 5c shows the green sepatations. Note that the processed New ProductOut has lightnesses that resembles the Goalin.*

*Figure 5d shows the red sepatations, the weakest of the three. Note that the pattern of NewProductOut lightnesses resembles the Goalin image more than the Bestin. This is true even though the Bestin has 5 individual areas that are closer to Goalin than NewProductOut.*

This process takes the spatial comparisons from the Goalin and limits the product by the Bestin image (Fig5). 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 resulting image NewProductOut shows a big improvement in appearance compared to the Bestin (Fig 5). The data shows that this new image is in-gamut.

[See Retinex-Gamut color images on CD-ROM version]

Figure 6 plots the Goalin and NewProductOut values of the six colored areas in L*a*b* space. The Goalin values are plotted as large solid squares; the NewProductOut are plotted as solid circles. The tthick solid lines show the distance between Goalin in and the NewProductOut values. The thin black line with arrow head shows the projected (Goalin-NewproductOut) vector. Fig 6a plots the data in the a*b* plane. Here the vectors pass near the central point (a* = b* = 0), but do not intersect there. The area swept out by these vector is shown in gray. Fig 6b plots the data in the L*a* plane. Here the vectors pass near the central point (L* =50, a* = 0), but do not intersect there. The area swept out by these vector is shown in red. Fig 6c plots the data in the L*b* plane. Here the vectors pass well below the central point (L* = 50, b* = 0), sweeping out the area shown in yellow. These graphs show the improved color seen in NewProductOut image are not caused by simple color space projections as: reduction in chroma (Fig 6a), projecting toward middle gray (Figs 6b & 6c). The NewProductOut image was created by optimizing spatial comparisons. Such a process is somewhat similar in that the results project in the vacinity of middle gray (L* = 50, a* = 0, b* = 0). However, their projections sweep out a substantial, nonsymetrical volume. This is an argument that the Gamut Retinex process is fundamentally different from a single, 3D color space transormation.

## Best Color Gamut Compromise

The familiar process, of evaluating the absolute coloimetry of a pixel to see if it is in-gamut, and then replacing it with the nearest in-gamut color, distorts color appearance. 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

Retinex calculations extended to the problem of gamut limited reproductions show promise. The argument developed with the aid of color displays in Figs. 1 & 2 states that global shifts in color similar to those found in color constancy



*Figure 6a plots the 6 Goalin to NewProductOut vectors in a* vs. b* plane.*



*Figure 6b plots the 6 Goalin to NewProductOut vectors in L* vs. a* plane.*



*Figure 6c plots the 6 Goalin to NewProductOut vectors in L* vs. b* plane.*

produce much smaller changes in appearance than local, individual color shifts.  Further, this paper argues that color gamut transformations using spatial comparisons can generate in-gamut reproductions thatlook more like the original, because it employs the benefits of color-constancy processing.  These  reproductions have a greater cumulative difference between original and reproduction, but look better.  Color is a spatial calculation in humans.

## Acknowledgments

## References

1.  E. H. Land and J. J. McCann, Lightness and Retinex Theory,  *J. opt. Soc. Am.*  **61,** p. 1, (1971).

*2.*  J. J. McCann, S. P. McKee and T. Taylor, Quantitative Studies in Retinex theory, A comparison between theoretical predictions and observer responses to Color Mondrian experiments, *Vision Res.,* **16**, pp. 445-458, (1976).

3.  J. Frankle and J.J. McCann,  Method and apparatus of lightness imaging, U.S. Patent 4,384,336, (1983)

4.  J.J. McCann, Lessons Learned from Mondrians Applied to Real Images and Color Gamuts, *Proc. IS&T/SID Seventh Color Imaging Conference, pp. 1-8, (1999).*

5.  J.J. McCann, Making Reproductions of Masterpieces: Managing all the Colors within the Gamut, *Proc. International  Symposium on Multispectral Imaging and Color Reproduction for Digital Archives ,Chiba,  pp. , (1999).*

6.  J.J. McCann, Uniform Color Spaces: 3D LUTs vs, Algorithms, *Proc. Second IS&T PICS Conference, pp. 204-208, (*1999).

7 .  J.J. McCann, R. L. Savoy, and J.. Hall, Jr., Visibility of Continuous Luminance Gradients, *Vision Research,* **18,** pp. 891-894 (1978).

8.  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).

9.  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).

10.  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).

11.  J.J. McCann, *J. Photogr. Sci.Eng.,* **22,** 64-68, (1978).

12.  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.  A.D. Logvinenko, *Perception,  in press, (1999).*

*14.*  J.J. McCann, "Color Theory and Color Imaging Systems: Past, Present and Future", *J. Imaging. Sci. and Technol., 42, 70, (1998).*

## Biography

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.