Calculating Lightnesses in a Single Depth Plane

John McCann

McCann Imaging
Belmont, MA/USA

Abstract

The experiments in this paper demonstrate that “Early Vision” mechanisms can account for the appearance of three “Diamond Wall” demonstrations, without reliance on apparent illumination, transparency, apparent depth and junctions.

Part 1 demonstrates that simultaneous contrast and the Katz-Albers effect can explain the appearance of a“ Diamond Wall” display.

Part 2 reviews the “Straight Edge” experiments designed to show changes in lightnesses consistent with perceived illumination. These experiments showed that both this effect, and White’s effect, are caused by very-low-spatial frequency sampling.

Part 3 applies complex “Early Vision” models to images associated with the “High Vision” lightness hypotheses. The argument is that flat displays, which are perceived as flat, require a quite complex visual mechanism just to account for the properties of flat lightnesses. Any experimental verification of the existence of “High Vision” lightness mechanisms should be tested first with realistic complex “Early Vision” models.

The results show that “Early Vision” mechanisms can account for appearances in “Diamond Wall” experiments. If “Early Vision” mechanisms can explain these results, then these experiments cannot be used as evidence for the existence of “High Vision” mechanisms.

Key words: lightness, brightness, Early Vision, High Vision, models of vision, sensation, perception.

Introduction

Lightness, the appearance of objects between white and black, is more complicated for humans than for image sensing technologies. All cameras, whether film, vidicon or charge coupled device, record images by counting photons. Humans generate images in which the lightness is independent of the quanta count.[1] A wide variety of different mechanisms have been proposed to account for what we see. Land and McCann[2] proposed a spatial comparison extending Wallach’s[3] observation about the importance of edge ratios. Stockham[4] suggested low-spatial frequency filtering. Gilchrist[5], Adelson[6] and Logvinenko[7] all suggested mechanisms in which apparent depth and/or apparent illumination altered the appearance of lightness.

The problem raised by Land and McCann was whether spatial comparisons could be used to separate reflectance from illumination given that the only input data was the radiance, (reflectance x illumination), at each pixel in the image. The initial work suggested that reflectances changed abruptly and that illuminations changed gradually, so that reflectance and illumination could be teased apart by the different rates of change on the retina.

This mechanism was shown to be incorrect for human vision for two reasons. First, our measurements of real life images showed that there are as many gradual changes in reflectance and abrupt changes in illumination as the converse. Second, extensive quantitative experiments[8] 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 *rate of change on the retina* threshold mechanism. This created a redefinition the goal of our model’s calculations. If humans see identical lightnesses from both reflectance displays and transparency displays having identical radiances, then we cannot really isolate actual reflectance from actual illumination.

The distinction between sensation and perception[11] becomes important. In this paper, *lightness sensations* are the appearances of objects between light and dark. *Brightnesses* are the sensations of how intense the illumination appears. *Perceived lightnesses* are sensations of lightness that have been modified by past experience. We can easily identify that a piece of paper has a uniform reflectance (*perceived lightness*), even though we see that there is a shadow falling across it. We see that the paper in the sun has higher *brightness sensations* and higher *lightness sensations*. The operational definition of the *sensation lightness* is the paint selected by a *fine-arts* painter to make a shadow. He adds black paint to the shadow to recreate the *sensation*. The job of the *house* painter is recognize the color of the paint on the house, perform the *perception* task, and ignore the illumination. Here the house painter reports that the two halves of the white paper are the same, regardless of the shadow. People can perform the tasks assigned to both types of painters. Scientists have to be particularly careful about the questions they ask observers. In this paper we will discuss *lightness sensations*, asking observers to perform exclusively the *fine arts* painter task.

Adelson has used the nomenclature of Early, Mid, and High Vision for different types of models (Figure 1). In this framework Early Vision is often described as quanta catch of the receptors. However, more important is the idea that lightness is calculated before depth. The sensation of lightness is an input to mechanisms that calculate apparent depth. The shading of objects is useful information in computing the shape in apparent depth of that object. High Vision is the equivalent of the OSA's definition of Perception, involving cognition. Here, apparent depth, and/or apparent illumination information feeds back into the human's computation of lightness. Mid Vision is the most practical because it assumes a simultaneous solution of lightness, depth, apparent illumination, and apparent reflectance. However, if your goal is to write equations that model visual response given the quanta-catch array on the retina, it is the most difficult to implement. There is no obvious place to start the process.

The terms Early, Mid and High, should not be used to infer locations in the visual pathway such as retina, lateral geniculate and cortex. The term *Retinex* was coined to emphasize the fact that we do not know where lightnesses are generated. Although we know much more about the neural process in the visual system, we still know very little about the neural locations of sensations and perceptions. In this paper Early Vision is analogous to *shape from shading*. Namely, it is a set of hypotheses that model lightness sensations using simple calculations, but excludes cognitive feedback. Lightness is available to be used as an input component of an apparent depth mechanism[10]. High Vision is *shading from shape*; namely, it is the set of visual mechanisms that uses apparent depth and/or apparent illumination feedback in the calculation of lightness.

The B&W Mondrian[2] had a white area and a black area sending the same radiances to the eye. Experiments with computational models of the sensation lightness showed that the reset, "normalization" process was the mechanism predicting appearance in B&W Mondrians. As well, it predicts the different appearance of grays on white and black backgrounds (*Simultaneous Contrast*). These displays have been successfully modeled by many different generations of Retinex models, starting with McCann, McKee and Taylor[9]. Later experiments with real life images, 20 years ago, used 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, a Retinex model created a new low-dynamic-range image displaying details of both sun and shadow areas. Recent experiments have used Retinex processing to generate improved color images in reproductions made on small gamut media.[10]

Experiments by Alexander Logvinenko[7] provide quantitative measurements of experiments being studied with a revitalized interest in Gestalt visual phenomena[5,6]. The experiments use "Diamond Wall" displays introduced by Adelson. The input digits are the same for apparently light and the dark diamonds. Adelson's and Logvinenko's analysis centers on the apparent depth, apparent illumination and apparent transparency observed from these flat displays. In their analysis observers see planes in the image, differing in depth, that appear to be in bright illumination (light areas) and other planes that appear in shadow (dark areas). They propose the perception of different depth planes, or shadow planes, causes the change in lightness. The analytical problem is that the displays use light and dark surrounds as well. We have to tease apart the effects due to depth from those due to simultaneous contrast [Part 1]. Further, we need to identify visual effects that are controlled by low-frequency sampling effects [Part 2]. Finally, we need to give "Early Vision" a fair chance, as to whether it can account for complex targets. Since Leonardo da Vinci, we have known that lightness is more complicated than the input value at a pixel. We should use a complex "Early Vision" lightness model that can account for a variety of flat, planar displays which appear flat, to see if they can account for lightness without invoking a depth hypothesis. Displays that lack both apparent depth and apparent shadows would still require a complex lightness mechanism to explain simultaneous contrast, the Black and White Mondrian, the Color Mondrian, and other complex phenomena. What happens when a complex "Early Vision" lightness model analyzes a display used to establish the case for "High Vision" [Part 3]?

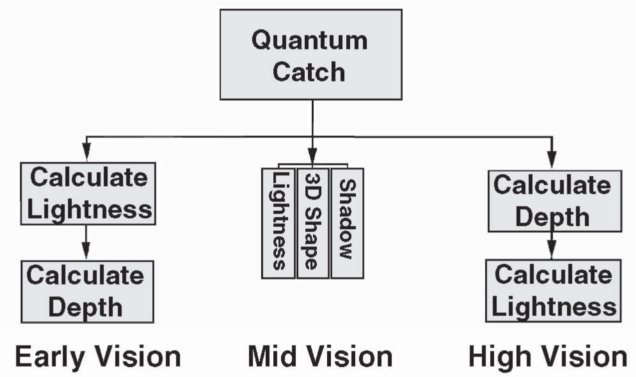

Figure 1. A diagram of alternative theories of lightness mechanisms. "Early Vision" suggests that lightness sensations come before depth perceptions. "High Vision" suggests apparent depth comes before lightness. "Mid Vision" suggests that lightness, shape and shadow are all calculated in parallel.

These "Diamond Wall" experiments center on whether the underlying mechanisms are the result of "Early Vision", implying that lightness is calculated before apparent depth and apparent illumination are undertaken, or "Mid Vision", implying the simultaneous solutions of both depth and lightness, or "High Vision" mechanisms that imply depth information is used to estimate lightness sensations. It is beyond the scope of this paper to describe in detail a visual model that will predict matching lightnesses for all 38 experiments presented. Rather, the focus of this paper is whether Early

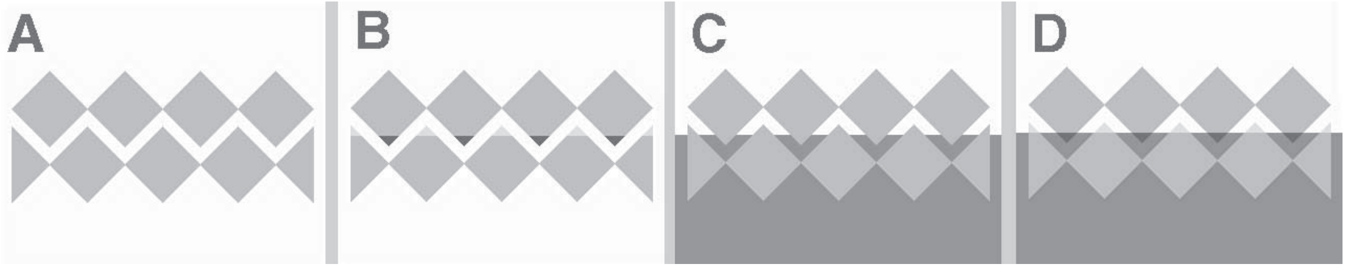

Figure 2. These four displays support a "High Vision" mechanism. A is a control using identical gray diamonds. B adds a dark tips to the top diamonds and light tips to the bottom diamonds to show that the tips do not change the appearance of the grays. C places a dark surround around the bottom row, while leaving a light surround around the top row. Surprisingly, the change of the grays is small. D is the combination of the tips from B and the surround from C. Here the rows of gray diamonds surrounded by white look darker than those surrounded by black. The "High Vision" hypothesis is that the human visual system interprets the grays on white as grays in bright light and the grays on black as shadow. The dark tips on the upper diamonds and the light tips on the bottom ones provide consistent evidence of the bright light/shade hypothesis. In the "High Vision" explanation, humans perceive the lightness of the lower diamonds to be lighter because they appear to be in a shadow.

Vision constructs can account for High Vision demonstrations.

Part 1: Simultaneous Contrast

The Diamond Wall experiments can be explained a variety of ways. They were originally designed as demonstrations of perceived illumination (Figure 2). They also have an explanation rooted in simultaneous contrast (Figure 3).

Experimental Evidence for High Vision

Experiments reported by Adelson and reprinted by Fairchild[12] (Figure 2) argue that lightness is controlled by the perception of a horizontal shadow. A, B, and C are controls. A shows that identical grays on white appear equal. The addition of tips in B does not change the appearance. Surprisingly, the addition of white and near-black surrounds in C to the rows of diamond also has little effect. The addition of both tips and surrounds in D makes the gray on white darker than the gray on black. The tips were intentionally chosen to create the perception of a high illumination horizontal stripe above a shadow. The tips of the gray diamond are light in the bright light and dark in the shadow. The appearance of the gray in the shadow is lighter. By this argument the "High Vision" perception of light and shadow is the controlling mechanism.

The experiments reproduced in this article have been compressed in size and placed in juxtaposition for economy of printing. The reader may want to place uniform surround around these images or reproduce the figures for themselves.

Experimental Evidence for Simultaneous Contrast

A second set of similar experiments can be used to make the case for the familiar "Early Vision" simultaneous contrast mechanism[13,17]. Figure 3 shows darker grays in white surrounds in E, F and D. In all three examples the middle-gray diamonds appear darker in the white surround. The unusual result is C. Here the absence of tips, or the presence of continuous gray diamonds have shut off the simultaneous contrast mechanism. How can we sort out the two opposing explanations? The solution centers on C. Is it a control proving that surround does not matter, or is it a curious anomaly?

Early vs. High Vision

The really interesting display is C. For some reason, the fact that the diamonds extend over the horizontal black-white line shuts off the usual simultaneous contrast effect. This result is the opposite of the role of edges found in the O'Brien-

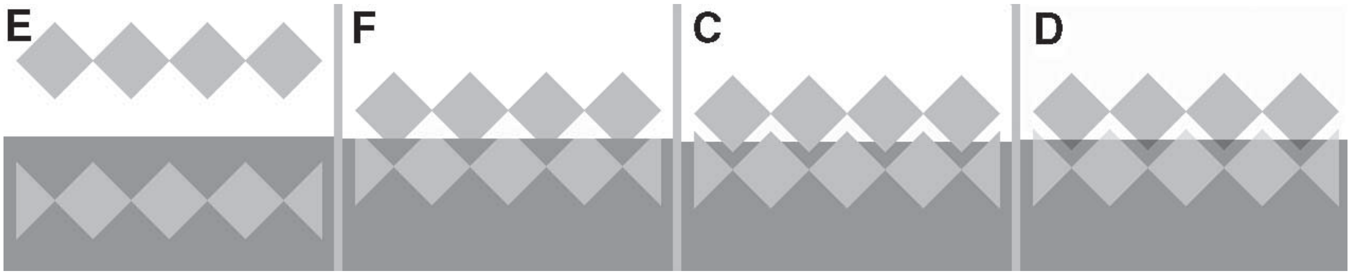

Figure 3. These four displays support simultaneous contrast. E is the familiar grays-on-white and grays-on-black. Grays-on-white looks darker. In F the diamond are moved to the white-black edge, with tips removed. The grays-on-white still look darker. C and D are the same as Figure 2. C exhibits anomalous behavior. It is different from all the rest. D behaves just like E and F, exhibiting the familiar simultaneous contrast effect. Display C has shut off simultaneous contrast.

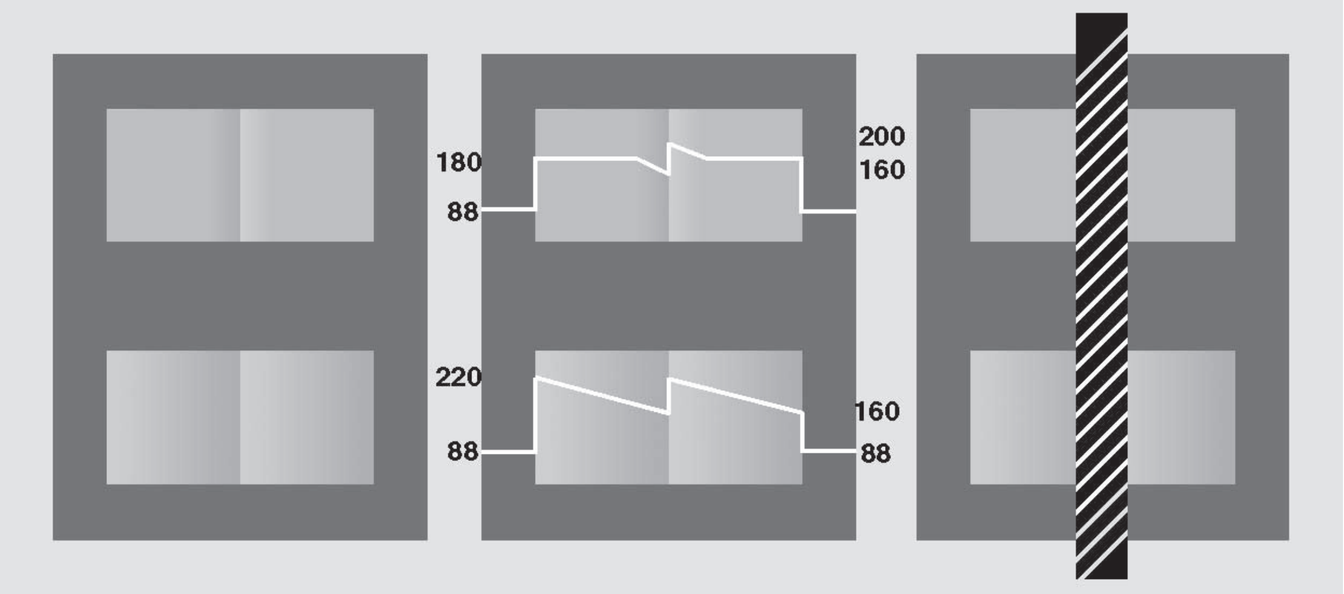

Figure 4. Demonstrations of the O'Brien, Craik, Cornsweet effect (top) and Land and McCann's "Two Squares and a Happening"(bottom). The left figure shows the displays. The central figure superimposes plots of pixel values on top of the displays. The surround of both has a value of 88. The Cornsweet is flat at 180 until it starts a gradient down to 160 near the center. It has an abrupt increase to 200 and then a gradual decrease back to 180. The "Two Squares and a Happening" starts at 220 and gradually decreases to 160 at the center. Then it jumps to 220 and again gradually decreases to 160. Although most of Cornsweet and all of "Two Squares and a Happening" are identical the edge is powerful enough to make the areas look different. The right figure adds the "Happening", namely a new object that obscures the boundaries between the two squares. Obscuring the edge releases the two squares to appear the same.

Craik-Cornsweet effect[14] and Land and McCann's "Two Squares and a Happening"[2]. Figure 4 shows these experiments, as well as plotting the pixel values of the displays. These effects show pairs of nearly identical squares (top) and identical squares (bottom) that look different. In the right figure we see that obscuring the edge releases identical squares to look the same.

Figure 5 shows the Katz-Albers Effect[14] which is the same phenomenon we see in Figure 2C. Figure 5 left is a large gray area on a black and white surrounds. The gray appears uniform. When a "Happening" separates the gray into two independent parts these two identical parts appear different. All the examples in Figures 4 and 5 demonstrate the powerful influences of edges and "Happenings" that obscure edges. Could these effects explain the behavior of C and D?

The Role of Edges

Figure 6 shows that the addition of a variety of lines across the gray diamonds restores simultaneous contrast. In G- a pair of thin light and dark lines, in H - an intermittent light and dark dots, and in Display I - a thick black line re-

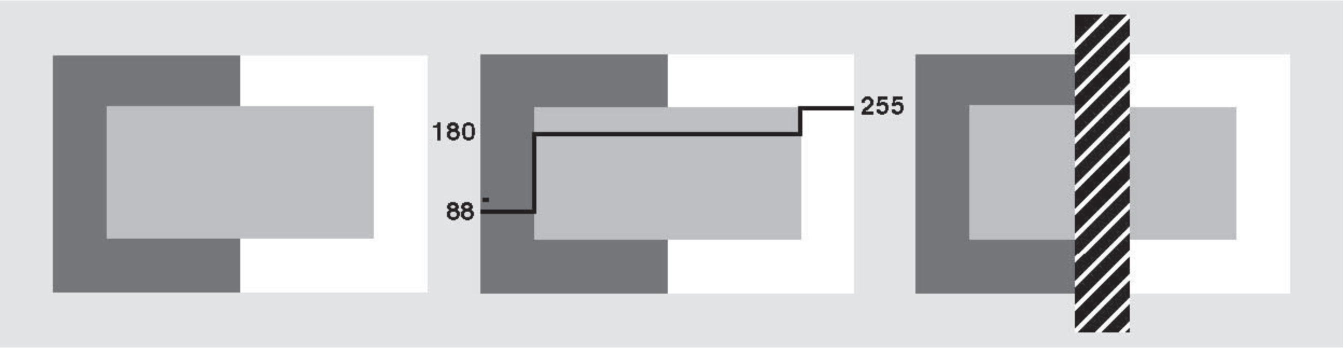

Figure 5. Demonstration of Katz-Albers Effect. The left figure shows the display. The central figure superimposes a plots of pixel value on top of the display. The surround on the left has a value of 88. The pixel value is flat at 180 until it reaches the white surround of 255. Although identical grays in white and black surrounds usually look different, the lack of a separation is powerful enough to make the whole gray area look uniform. The right figure adds the "Happening": namely a new object that creates a boundaries between the two identical halves. Obscuring the lack of an edge releases the two squares to appear the different.

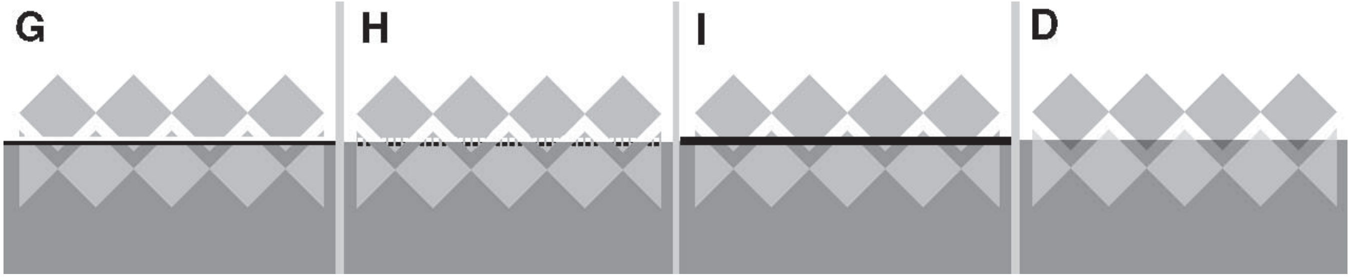

Figure 6. Four displays used to test the role of edges in displays C & D. G shows that the introduction of a light and dark line across the gray diamonds along the white-black border restores simultaneous contrast. H shows that dots work as well; I shows that a simple black line does better. By this hypothesis, the anomalous behavior of display C is controlled by the uniform gray diamonds. Any contour that breaks up the uniformity of the diamonds will restore simultaneous contrast.

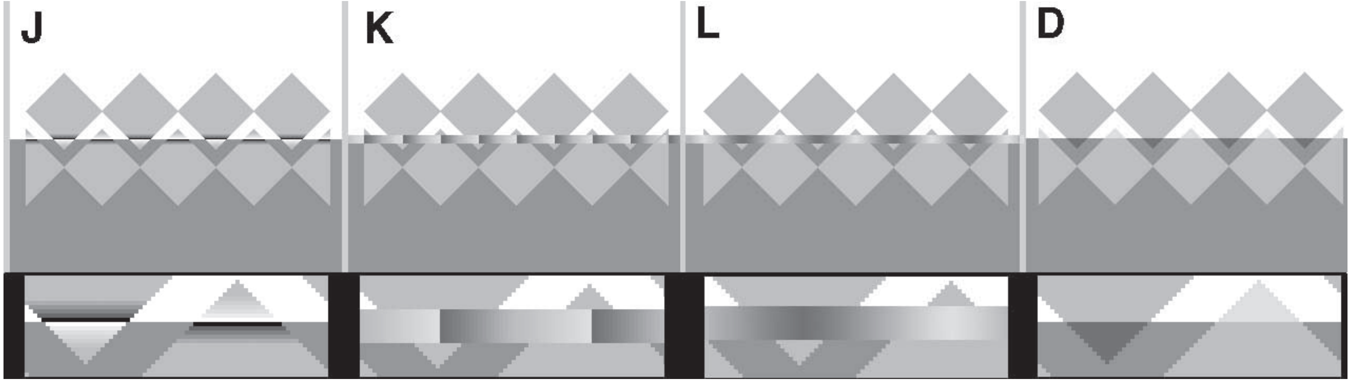

Figure 7. Four additional displays testing gradient edges. Below each display, is a 4x magnification of 2 tips. This will allow you to see the construction of the gradients which are hard to see in the displays. Display J uses a Cornsweet edge in the gray tips. It is a gradient that equals the gray diamond at the top, as it moved down it gets darker, until the level of the white-black edge, it jumps lighter and then gets darker until it is again equal to the gray diamond. Display K uses horizontal gradients. In the center of each diamond tip there is a vertical edge. The light side of the edge is equal to the light tip in D. The dark edge is equal to the dark tip in D. These lines are connected by gradients of values. The center of the gradient has the same value as the gray diamonds. Display L has a gradient oscillating from the value of the light tip in D to the value of the dark tip in D. All of these insertions near the light/dark horizontal edge release the diamonds to look different, thus re-enabling simultaneous contrast.

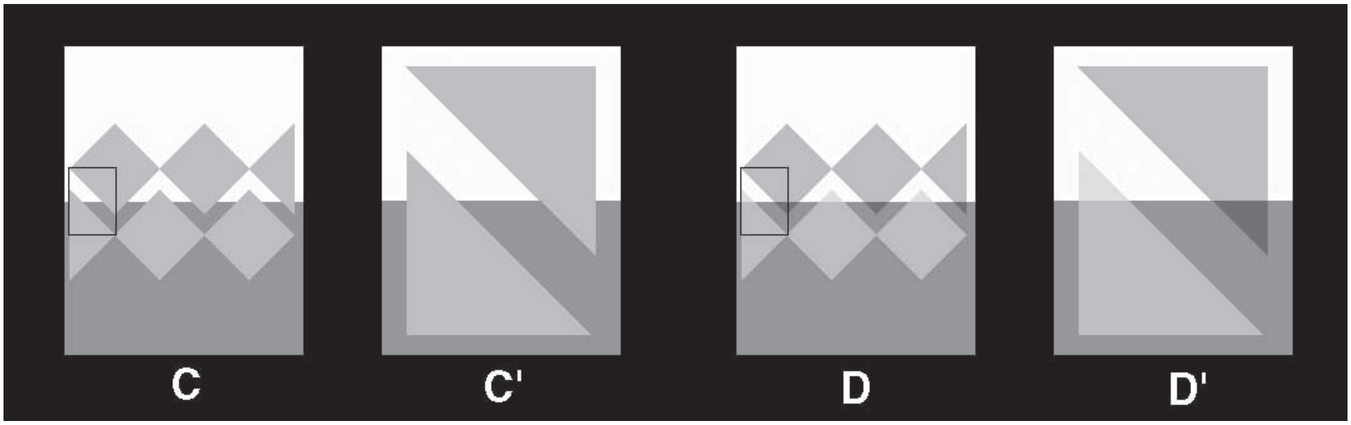

Figure 8. These displays repeat C and D while introducing simplified versions C' and D'. The same lightnesses are seen in displays C and a simplified one-quarter diamond pattern C'. The same is true for D and D' with tips. There is no difference in lightness between the complex displays with implied illumination and the simple ones without any implied illumination.

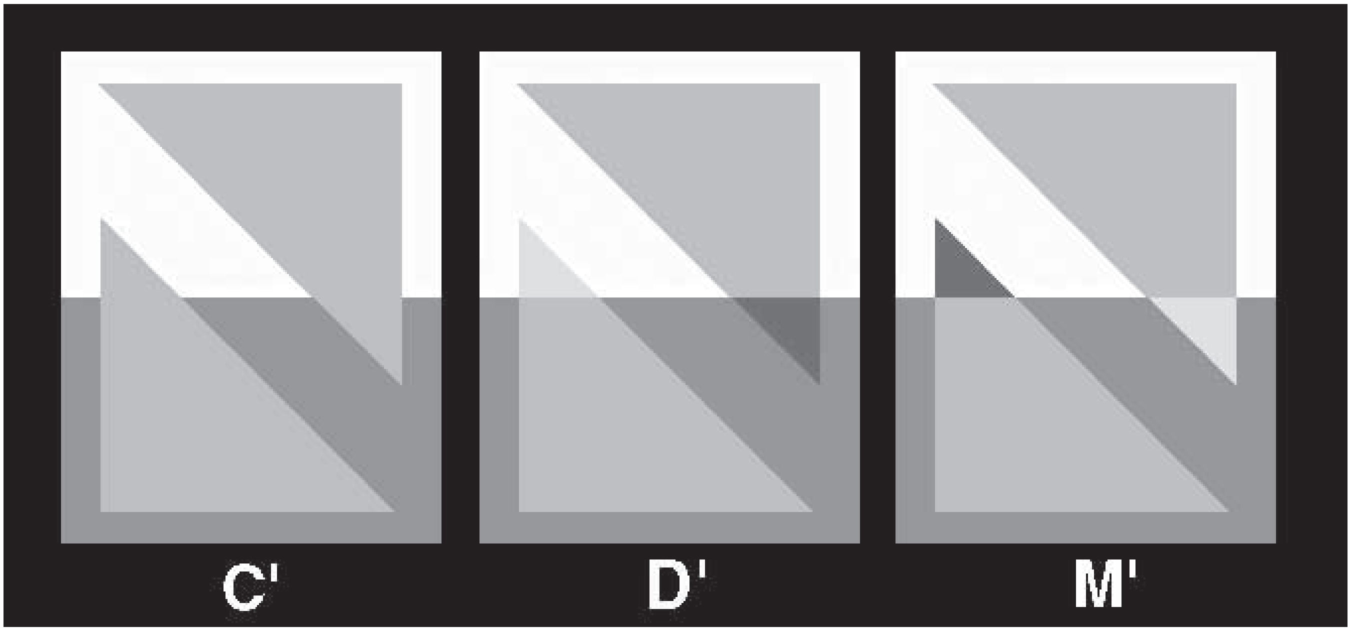

Figure 9. Display M' shows the reversal of the tips, so that they are inconsistent with the illumination hypothesis. The difference in lightness is slightly smaller than that found in D', however it is larger that found in C'. These inconsistent tips act to restore Simultaneous Contrast.

lease the two rows of diamonds to look different. It is interesting to note that the lines in G, H and I are not equally good at releasing the grays. Display I looks more like D than G and H.

Figure 7 pursues the idea that any edge introduced into the gray tip area so as to extend the black white edge across the tip will restore simultaneous contrast. Display J uses Cornsweet edges, K and L use horizontal rather than vertical gradients. In each case, we see that these objects restore simultaneous contrast.

We then made a series of different colored tips all with the same lightness value as the diamonds. Only hue and saturation were different from Display C. These colored tips restored simultaneous contrast as well.

Simplified Displays

Figure 8 shows a portion of Display C, along with a simplified, one-quarter diamond segment. The black lines identify the part of C that was magnified to make C'. On the right side of Figure 8 is the same comparison of D and D'. C' and D' show the same contrast behavior as C and D. The effects we see in C and D do not require long edges intersecting many areas. They do not require either apparent depth or apparent shadow. The contrast behavior is the same in C'/D' as in C/D, because the mechanism controlling their lightness is local.[17]

Inconsistent Tips

Figure 9 shows three "Quarter Diamond" displays. On the left is C', in which the central grays look most like each other. In the middle is D' which adds tips consistent with a shadow interpretation. If we visualize that the entire lower triangle is made of the same paper, and the top triangle is a different darker paper, and if we visualize the horizontal level of the white-black edge in the surround as the edge of a cube with two visible faces, and we visualize that the top face of the cube is in bright light and the bottom is in a shadow, we can see that the change in appearance of the chips is consistent with a shadow making the bottom triangle tip lighter in the bright light and the top triangle tip darker in the shade. In Display M' we reverse the tips. Now, the lighter tip is on the bottom and the darker is on top. It is still possible to visualize the white-black edge as the edge of a cube, but the tips are clearly different patches than the rest of the triangles. M' shows that reversed tips act to restore simultaneous contrast.[17] Illumination-consistent tips are not required for the gray area surrounded by white to look darker than the gray area surrounded by black.

In summary, there are three arguments for simultaneous contrast or "Early Vision" as the controlling mechanism for lightness sensations in this type of "Diamond Wall" experiment (Figure 2). First, any edge restores simultaneous contrast (Figures 6 &7). Second, simple patterns have the same effect on lightnesses as complex ones designed to enhance apparent depth and illumination. (Figure 8). Third, the tips used in D to make a shadow-consistent image restores simultaneous contrast. The same is true if we reverse the tips and make a shadow-inconsistent image (Figure 9).

Simultaneous contrast is all we need to understand the lightness effects in these diamond experiments. That is not to say that these displays all have the same apparent depth. They do not. It is just to say that we do not need apparent depth and illumination effects to explain the lightness differences. These experiments are arguments for "Early-Vision" lightness mechanisms because simultaneous contrast can account for lightness appearance without the need for feed-

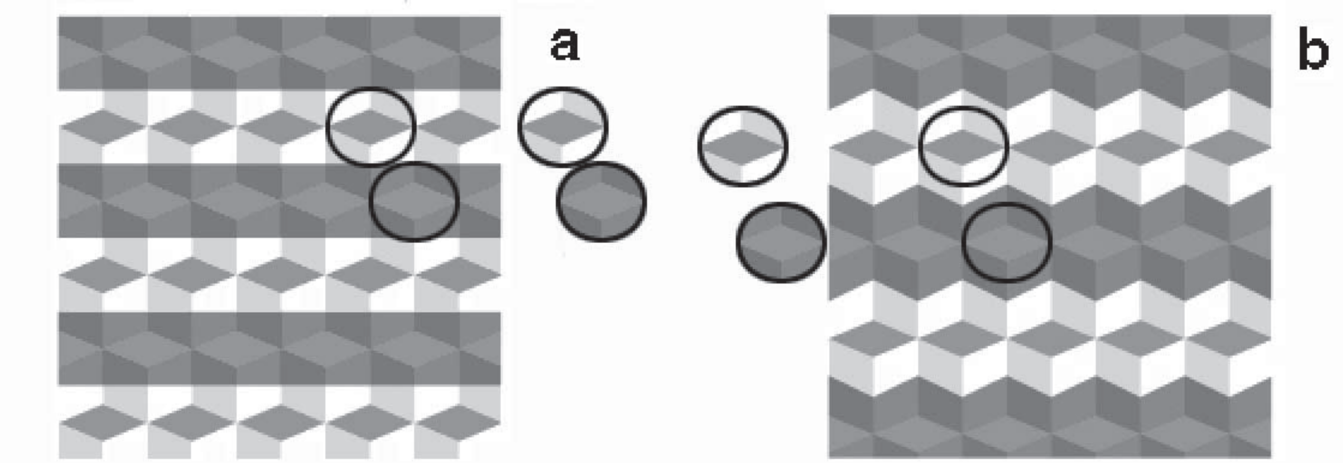

Figure 10. Adelson effect illustrating the difference in lightness with straight light-dark edges and angular borders. Logvinenko's argument is that the straight edges are associated with a perceived illumination edges, and are different from grays that are not associated with perceived illumination. This pair of displays have the same gray diamonds and the same surrounding areas. The black ellipses identify the areas around gray diamonds that are identical in both displays. Copies of the identical segments are placed in the margin, at the same relative spacings, for comparison . The gray diamonds have the same horizontal separations, but have different vertical separations.

back from higher-level mental processes.

Part 2: Low-Spatial Frequency Sampling

In this section we study a different "Diamond Wall" experiment described by Adelson in 1993 and measured using the Munsell Scale by Logvinenko. (Logvinenko's Figure 1)[7]. The "Straight Edge" effect is shown here as Figure 10. Logvinenko measured the differences in lightness between the gray diamonds with light surrounds and those with dark surrounds for both left display **a** and right display **b**. His observers reported that the difference in diamonds in **a** was 1.0 Munsell unit greater than in **b.** Logvinenko argued that perceived illumination is responsible for this effect. The pixel values of the four adjacent areas in a both **a** and **b** are the same. [White = 253, light gray = 195, gray diamonds = 128, dark gray = 113, very dark gray = 93]. How can a low level mechanism account for the observed lightness differences when all the adjacent areas have the same pixel values? Figure 10 also identifies the identical areas and their relative placement on the page.

The "Straight Edge" effect in Display **a** creates both apparent depth and apparent shadows. The question we study here is whether the *apparent* properties are influencing the lightnesses, or whether the difference in spatial composition of the rest of Figure 10 can account for the difference in lightnesses between **a** and **b.**

Figure 11 shows that the "Straight Edge" effect has different visual properties than the diamonds in Part 1. There, the lightness differences were the same in the simple vs. repeating patterns. In Part 2, Figure 11, the lightness differences are different in the simple vs. repeating patterns Here we reduced 10a and 10b to their simplest components. The light and dark grays in **a** (left) look the same, or nearly the same, as those in **b** (right), and their differences in lightness are equal. Thus, we believe that the "Straight Edge" effect is due to a different underlying visual mechanism than the experiments in Part 1.

White's Effect

White's Effect[15], the inverse simultaneous contrast effect, is shown in Figure 12 right. The gray bars on the left are darker (wih more adjacent black) than those on the right (wih more adjacent white).

Recent papers[17] has provided a very wide range of interpretations of the effect. Spehar et. al. look to High Vision mechanisms. White's 1981 and Moulden & Kingdom papers demonstrated the need for two different low-level mechanisms for simultaneous contrast and White's effect. Blakeslee and McCourt's 1999 paper shows that an oriented difference of gaussian (DOG) model can account for both White's experiments and the White and White 1985 experiments. McCann argues that the gray bars in White's effect

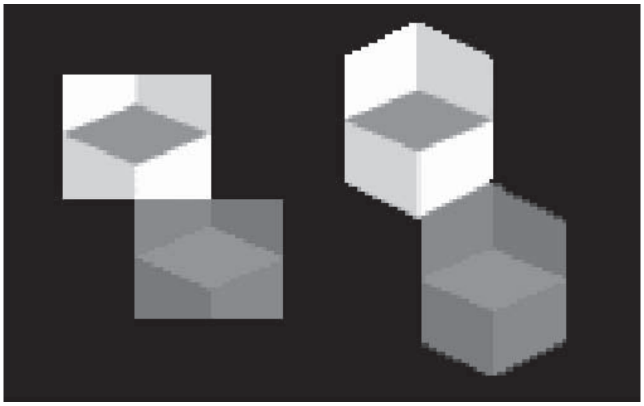

Figure 11 shows that the elements that make up Figure 10 **a** and **b** do not exhibit the "Straight Edge" effect. The difference in the diamonds' lightnesses is the same on both sides.

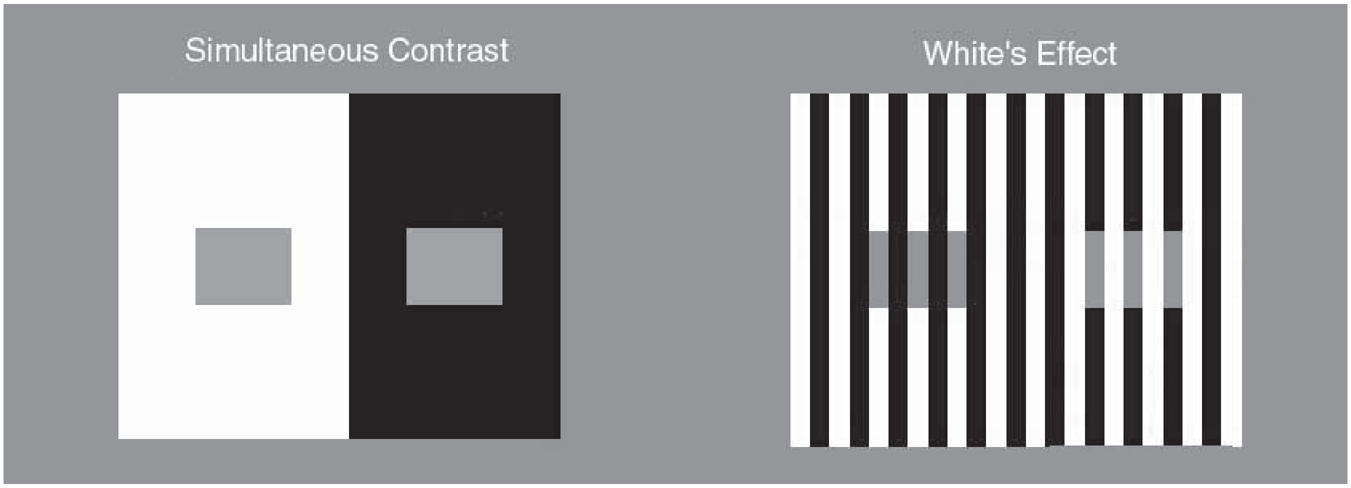

Figure 12. The comparison of simultaneous contrast with White's effect. The gray on right in White's effect is lighter than the left. In White's effect there is more white adjacent to the gray than black. (The effect gets much stronger if you squint your eyes. Somehow low-spatial frequency information is controlling the effect.)

Figure 13 demonstrates the influence of spatial averaging by large receptive pools. White/black circles indicate the area in the underlying White's display that is averaged. The left and right columns of digit are the average value for each area. All the pixels in the circles have been replaced by the average values. Single receptors in the gray bars and pools up to the width of the gray stripes give equal responses (144). As the pool get larger the black stripes on the left lower the pool response to as low as 73, while the white stripes raise the pool on the right to a high as 198. As the pool gets larger, the pool response returns to equal values. The large pools demonstrate that White's lower lightness on the left correlates with lower average digits. The gray on the left is darker, because the stimulus is lower in the very-low-spatial frequencies.

are controlled by the values in very-low-spatial frequent channels. Bars look lighter because they have higher values in the very-low-spatial frequency channels.

Figure 13 illustrates how multiresolution sampling can account for the lightnesses on the right bars being higher than those on the left. The gray stripes have the same digital value [144] in both halves of the display. If we measure the average response from a circle of pixels (white/black outline) with the diameter equal to the width of the stripe, both average values are equal (Figure 13-top row). When the diameter of the area increases to two times the width of the gray stripe, the average (gray surrounded by black) get smaller [104], while the average (gray surrounded by white) gets larger [162] (Figure 13-second row). These coarse visual receptors are sensing both the gray stripe and the white or black surrounds to make the averages different. At four times the gray stripe width, the average (gray surrounded by black) get smaller [73], while the average (gray surrounded by white) gets larger [138] (Figure 13-third row). At eight times the gray width, the averages both equal 108. The size of the receptor pool has grown to a size so as to integrate all the details into a equal averages (Figure 13-bottom row).

Full-resolution models that can successfully predict lightness in simultaneous contrast cannot predict White's effect. Nevertheless, almost any multiresolution model can account for these results by sampling the output from each resolution independent of the others. Figure 14 illustrates the differences between unsuccessful (left) and successful (right) White's effect models. On the left the model averages the input image to form a series of smaller, lower resolution images. The processing begins with the smallest image. The results are interpolated up to make an output image the size of the next larger input image. The interpolated precessed image and the averaged input image are then processed and interpolated up to the next size. This operation is iterated until all sizes are processed and a full-resolution processed image is produced. Using a Retinex model[10] with a very wide range of parameters we failed to replicate Whites effect data.[16]

Figure 14-(right) illustrate how spatial summation of

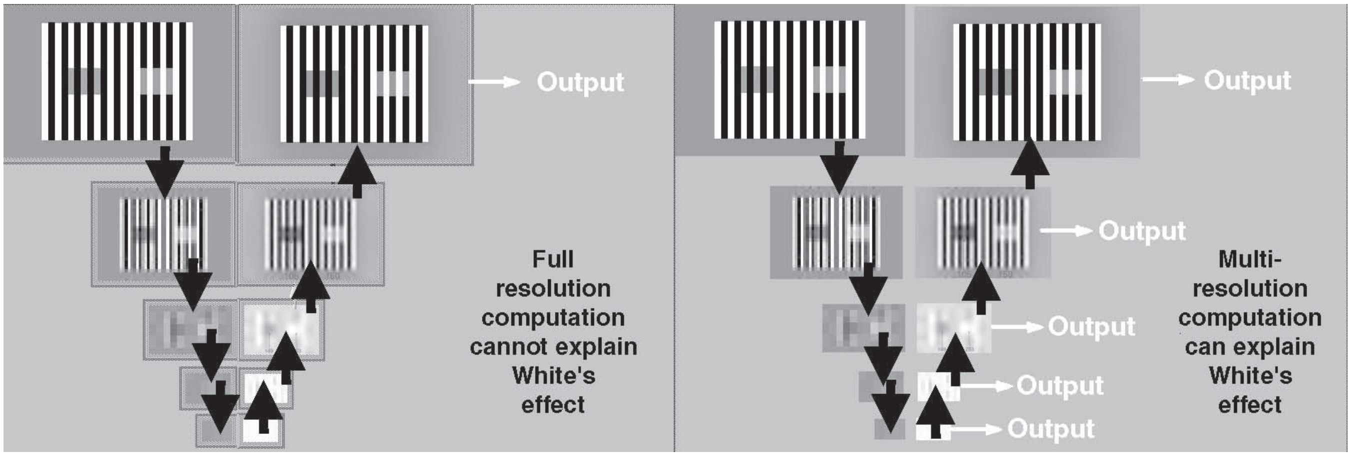

Figure 14. The left side shows that multiresolution computations that averages the input image into a series of smaller, lower-resolution images. Beginning at the smallest image and applying models such as ratio- product-rest-average processes do not give results consistent with White's effect. The right side shows that models that combine low-spatial resolution partial outputs with mid- and high-spatial frequency outputs give results consistent White's effect.

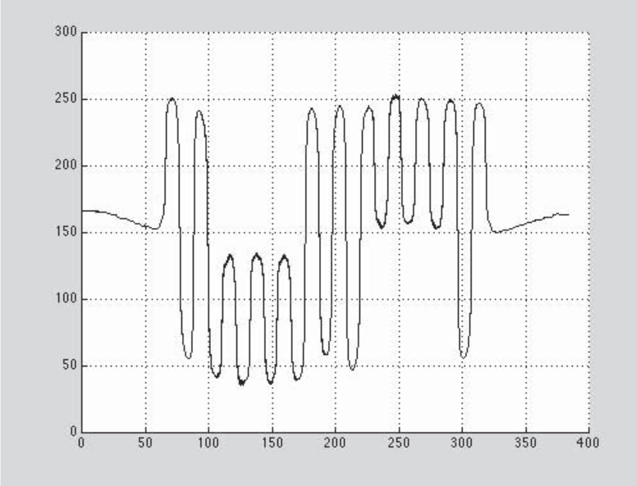

Figure 15 shows the plot of output digits from a combined image of different averaged layers. Figure 12 (right) was the input. A series of 8 averaged layers were made and rescaled using bicubic expansion. The proportional summation of all

independent channels can explain Whites effect. Instead of generating a single output from the highest resolution image, we can extract a partial output from each resolution. Figure 15 plots the calculated output using a primitive model that combines input data from different average layers. Starting with the full resolution image, make a series of 8 smaller image layers, each half the size of the previous layer. A pixel in the smaller image is the average of four pixels in the next larger one. For didactic purposes, we used Photoshop to rescale and recombine the different average images. We used bicubic interpolation to create a family of spatially averaged images all the same size. Second, we pasted each averaged image into a layered image with 30% transparency for each layer. In the combined image the grays on the left are darker than those on the right, thus showing that the combination of image made by low-spatial frequency sampling is all that is needed to explain White's effect. Figure 15 is a plot of output digital values vs. horizontal distance. The output associated with the gray bars on the left adjacent to black bars is less than the output of the same gray bars on the right. Spatial sampling can account for the difference in gray bar appearance in Whites effect.

This simple demonstration and spatial filtering experiments by White, Moulden & Kingdom, Blakeslee & McCourt all provide evidence of the important role of very-low-spatial frequencies in the appearance of equal grays in different surrounds.

Straight Edge Effect

Let us return to the "Straight Edge" effect in Figure 10. Can low-spatial frequency sampling, used as an explanation for White's effect, explain the difference in lightness between gray diamonds in Figure 10**a** and 10**b**? Or, as Adelson and Logvinenko asserts, the lightnesses are different because the long light-dark straight edges (10**a**) can be interpreted as an illumination edge, while the sawtooth light-dark edges (10**b**) cannot.

The vertical spacing between corresponding gray diamonds is smaller in Fig 10a than in 10b. The display in Figure 10 is 256 pixels in height. If we compare the averages both figures we find they are somewhat similar for 2x2, 4x4, 8x8, 16x16. However, they are markedly different for 32x32. Fig 16a (left) shows the 32x32 averaged down and bicubically interpolated back up. It is uniform. Fig 16b (right) shows the 32x32 averaged down and bicubically interpolated back up. It is far from uniform. It shows sinusoidal variations in intensity in the vertical direction. The black diamonds show the corresponding locations of the gray diamonds in Figure 10.

We can remove this low-spatial frequency artifact from Figure 10 by reconstructing it with the same vertical distances in **a** and in **b**. It is possible that matching the displays at the average 32 level might make the "Straight Edge" disappear?

Figure 17 is the reconstruction of Figure 10 with equal vertical distance between gray diamonds in **a** and **b**. The lightness differences between 17a with straight edges are still

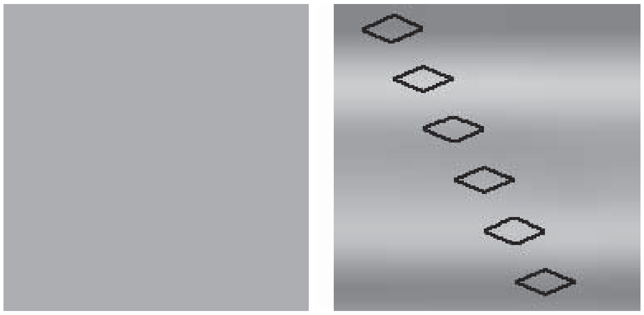

averaged layers make the gray bars 135 on the left, and 155 on the right (15% higher).

Figure 16 shows the average digits in Figure 10 at the 32 x32 level. The average for Fig 10a (left) is uniform. The

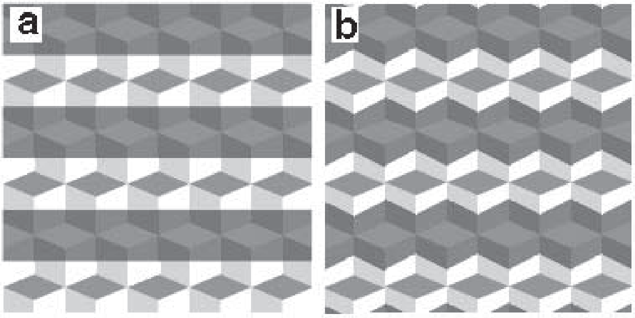

outlines show the locations of the gray diamonds at full resolution.

Figure 17 shows the revised versions of Logvinenko displays

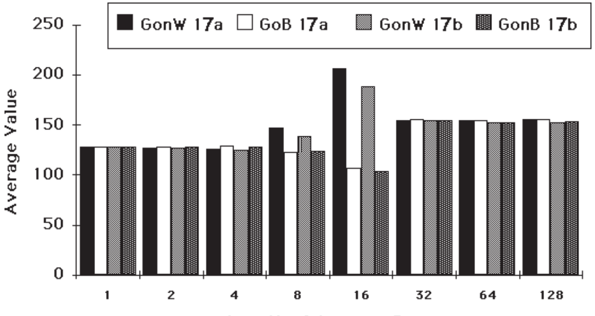

with equal vertical spacing of diamond in **a** and **b**. The lightnesses of the diamonds surrounded white and black is still larger in Figure 17 a.

Figure 18 shows the histogram of average low-spatial frequency input. The horizontal axis plots the length of each side of the averaging box. The value 2 means that a 2 by 2 area of 4 pixels were averaged. After averaging a full size image was reconstructed using bicubic expansion. The vertical axis plots the average digital value of 120 pixel samples from the middle of the gray diamonds in Fig 17. Data is shown for gray diamonds surrounded by white areas in Figure 17a with legend title [GonW 17a], gray diamonds surrounded by black areas in Figure 17a [GonB 17a], gray diamonds surrounded by white areas in Figure 17b [GonW 17b], gray diamonds surrounded by black areas in Figure 17b [GonB17b]. Averages with box length 8 and 16 show grays that have unequal average values. The difference in lightness between [(GonW 17a-GonB 17a) - (GonW 17 - GonB 17b)] larger than 17b. Although this reconstruction of Figure 10 made the two displays more like each other, it did not remove the lightness differences. It seems that the 32x32 average was below the low-spatial frequency visual response.

Again, we took the digital data input for Figure 17 and averaged it to model lower-spatial frequency responses. We wrote a program that calculated the average values in a box with lengths 2, 4, 8, 16, 32, 64, and 128. We used bicubic expansion to regenerate images the size of the original. We averaged the values for 120 pixels corresponding to the center gray diamond surrounded by whites [GonW] and 120 pixels from the gray diamond on black [GonB]. These values show the average input values for low-spatial frequency input information.

Figure 18 shows the histogram of these averages. The data shows the average values of input data for different low-spatial frequency images. Sampling using boxes of length 1, 2, and 4 show equal average values of 128. When the box has sides of lengths 8 and 16 the average values include the higher, white surround pixel values and the lower, black surround values in the average. When the box length reaches 32, 64 and 128 the data incorporates several rows of data and averages approach 155 with slightly different averages of the entire images.

The data for length 8 and 16 show that the difference in lightness for 17**a** is slightly lighter that for 17**b**. Nevertheless, this small difference is enough to account for the small difference in lightness reported by Logvinenko.

If a spatial sampling hypothesis is correct, then making Figure 17 taller should make the effect go away. Figure 19 is a new elongated display that magnifies Figure 17 four times (vertical only). These displays still use the same pixel values in adjacent areas. Figure 19 still has straight edges in **a,** and corrugated edges in **b**. With this version of “Diamond Wall” the grays look the same, regardless of the nature of the light/dark edge. This behavior reinforces the sampling explanation of the “Straight Edge” effect.

In summary, there are three consistent experiments that support the sampling hypothesis. First, the images have different low-frequency energy distributions (Figure 18), second, the phenomenon disappears with a constant surround (Figure 11) and third, it goes away with vertical magnification (Figure 19). We see that two different “Diamond Wall” experiments have very different visual properties. Part 1 has the same lightness effect in complex and simple displays (Figure 8), while Straight Edge effect is found only in complex displays (Figures 10 &17). Each can be explained by different “Early Vision” mechanisms suitable for their different image construction.

Part 3: Complex “Early Vision” Models

We should use a complex “Early Vision” lightness model that can account for a variety of flat displays, that appear flat, to see if they can account for lightness without invoking a depth hypothesis.[17] Displays that lack both depth and apparent depth still require a complex lightness mechanism to explain simultaneous contrast, the Black and White Mondrian, the Color Mondrian and real life scenes. These experiments

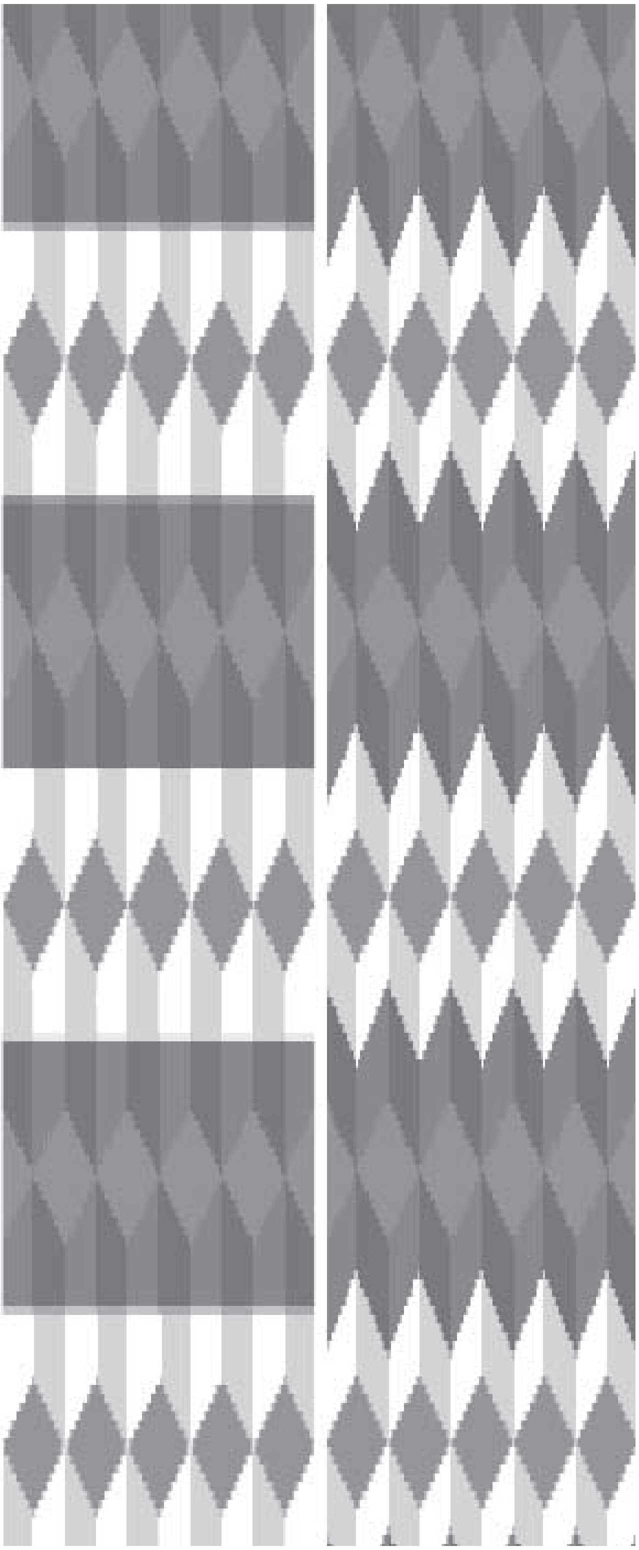

is 9 for length 8 and 15 for length 16. This shows that low-spatial frequency can account for the fact that the difference in lightness in Figure 17a is larger than Figure 17b.

Figure 19 is the same as Figure 17 in the horizontal dimension. In the vertical dimension each line is repeated four times. The "Straight Line" effect has been removed, even though all the perceptual arguments about apparent illumination shadows still hold. The fact that the difference in gray sensations are the same is an argument for the low-spatial frequency sampling explanation of Figure 17.

produce lightnesses that are significantly different from the quantum catch of the input. Retinex models have been developed that can calculate lightness that match observer data. What happens when a complex "Early Vision" lightness model analyzes a display used to establish the existence for "High Vision"?

Flat Lightness Phenomena

In simultaneous contrast (Figure 12 left) grays appear 1.0 Munsell lightness unit different in different surrounds. There is no real, or apparent depth in this display. Many hypothetical mechanisms can account for the lightnesses in Figure 12-left. Many are models that just use local interactions; other are global with imperfect normalization to the maxima in each waveband.[10] The term *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. 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 surrounds, show that the word "normalize" must be used with caution. Figure 20 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 on the left surrounded the gray with an equal width of white. Observers matched the gray to Munsell value 1.5. The next experiment cut the white into four pieces and placed them 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 experiments 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 whites[18] make the same point. Spatial positions of test patches relative to whites have large effects on lightness. Matches varied from 1.5 to 3.5 Munsell units. That is 28% of the lightness difference between white and black. This variability was controlled by

Figure 20. These six experiments illustrate that the

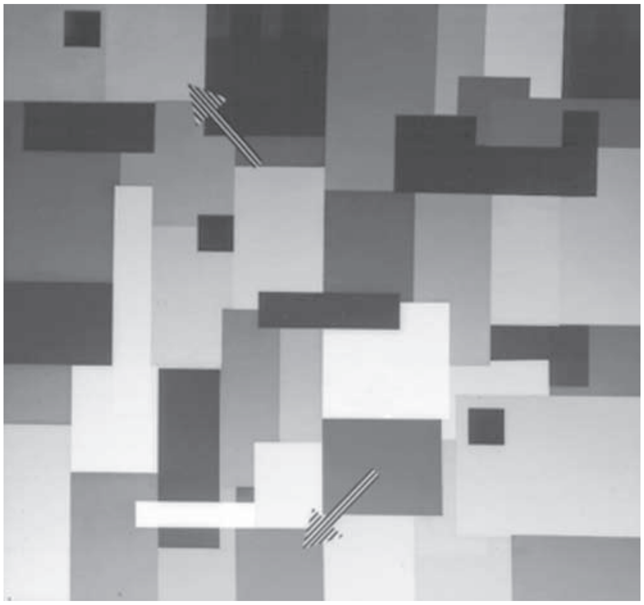

appearance of the central gray area depends on the location of the white. This points out the importance of enclosure as a part of color appearance.

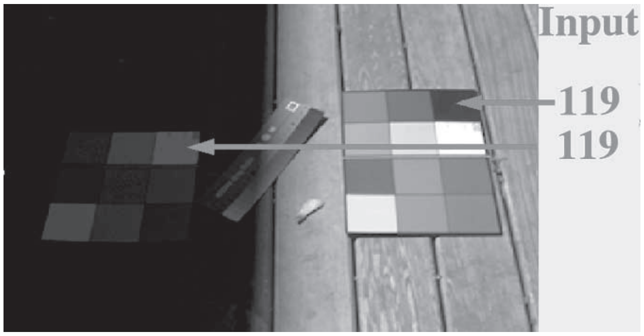

Figure 21 shows the Black and White Mondrian. The tips of arrows show light areas at the top and dark areas at the bottom with equal radiance at the eye.

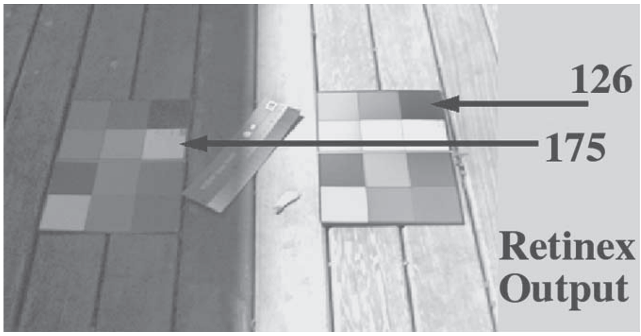

Figure 22. The scanned photograph of the sun/shade image. The white in the shade has the same input digital value (119) as the black in the sun.

Figure 23. The Retinex Output photograph of the sun/shade

the enclosure and separation of the area white area. A successful lightness model must mimic these human visual properties for flat displays, that look flat. In the theoretical evaluation of complex displays we should turn first to models that can account for a large variety of flat lightnesses. Such models can provide an "Early Vision" alternative explanation to "High Vision" phenomenon.

Ratio-Product-Reset-Average models have been used calculate the observed lightness in simultaneous contrast (Figure 12- left) since 1970.[19] Black and White Mondrian had a white patch and black patch sending the same radiances to the eye. (Figure 21). It was modeled in the original Land and McCann [2] article. Color Mondrians were modeled by the same technique.[9] A variety of outdoor images including "John at Yosemite", a real-life equivalent to the B&W Mondrian. It consisted a black card in the sun having the same radiance as a white card in the shade.[1] Figure 22 is a recent photograph of two Jobo reflectance 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. The shadow was 32 times darker than the sun. The black reflectance in sun and the white reflectance in the shade both have 119 as the scanned input digit. Figure 23 shows the processed image[10] in which the sun image is essentially unaltered. The black reflectance in the sun has only moved from digit 119 to 126. However, the white reflectance in the shade has moved from 119 up to 175. This display is a flat recreation of "John at Yosemite" without depth and apparent depth.

Many different experiments have been modeled using the same Ratio-Product-Reset-Average model[10]. It takes the information from the entire field of view and calculates an imperfect global normalization of the image. It is sensitive to the separation and enclosure of this maximum. It works with quanta catch at the retina as the only input. It can account for the observed lightness in all the above flat images without depth or apparent depth information. Can it account for 'Diamond Wall" apparent depth experiments?

Lightness Models and Gestalt Displays

Figure 24 shows Logvinenko's gradient experiment (Logvinenko's Figure 4)[7] generates the largest lightness change between diamonds. The digital values for all pixels in a 384 by 256 display was the input. The input image was averaged down to a 2 by 3 pixel image (See Figure 25-right). This image was processed using Ratio-Product reset-Average steps (See Figure 25-left). This result was interpolated to the next size larger image and the process was repeated. The processed result from the 384 by 256 image is the output of the model. (A detailed discussion of the steps in this calculation, including MatLab code can be fount in a recent review paper[21].)

The left border of Figure 24 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.

image. The white in the shade has a processed value of 175 compared the black in the sun with 126.

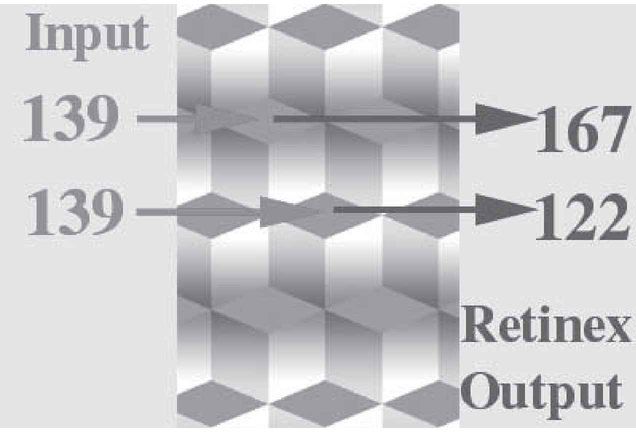

Figure 24. The input and Retinex output for Logvinenko's Gradient experiment. Despite constant input (139), a Ratio-Product-Reset-Average model can calculate lightness outputs consistent with observer values. The model designed for the B&W Mondrian accurately predicts the lightness of Logvinenko's diamonds.

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, including the successful prediction of Logvinenko gradient "Diamond Wall".

Discussion

It should be noted that the combination of Adelson's and Logvinenko's ideas need to be handled carefully. Both authors have experimented extensively with Adelson's "Diamond Wall" patterns. That is not to mean that their analysis of human vision are the same. Adelson has restated some of his earlier hypotheses in a "Mid Vision" framework using the categorization of junctions. It is reminiscent of the Lettvin's triads and the Todorovic and Zaidi et. al. T junctions[20] However, Adelson's intersections deal with apparent lightness, while Lettvin's dealt with the color of the illuminant. Logvinenko builds his arguments around perceived illumination, rather than perceived depth.

If we review Part 3, this time evaluating apparent illumination, we arrive at similar conclusion, but with an altered list of experiments. Both the Black and White Mondrian and Figure 22 have apparent illuminations - one is gradual, and one is abrupt. In principle, perceived illumination could be used to account for these experiments, if we only knew how to calculate this perception. As described earlier, such a calculation requires that we that know how to extract both reflectance and illumination from radiance. As discussed earlier, humans see the same lightness sensations from identical radiances made by either reflectances and transparencies. As the problem stands today, we can calculate lightness and color

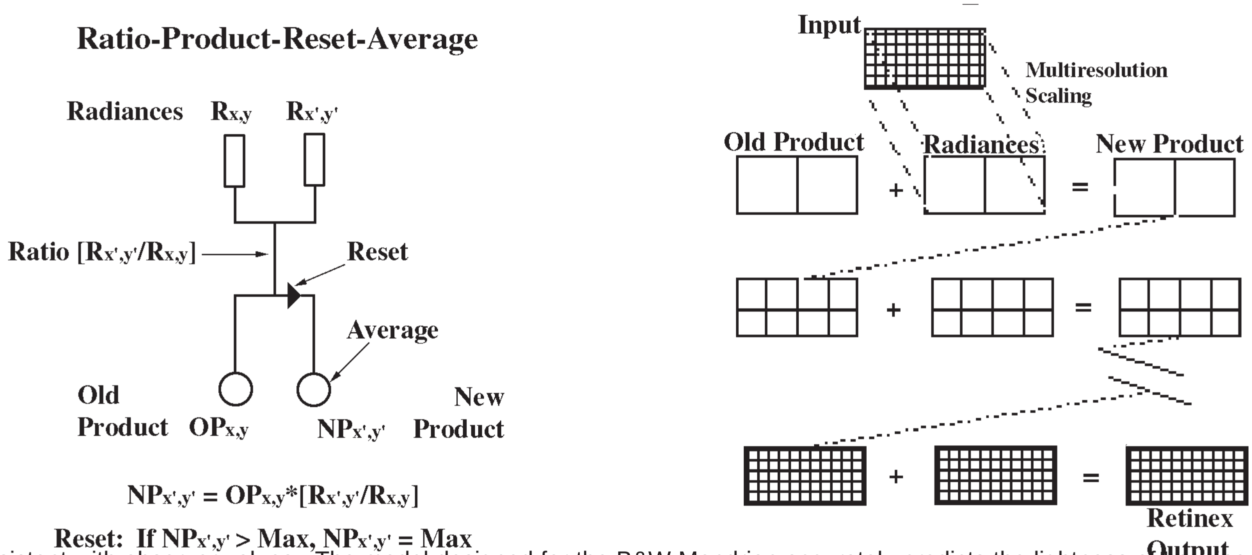

Figure 25 (left). The explanation of Ratio-Product-Reset-Average operation. Here we calculate the New Product (NP) for the output pixel x',y'. We begin at the starting pixel x,y using the Old Product (OP). All OP's are initialized with the maximum value. The product of the radiance Ratios times the Old Product is reset if greater than the maximum and averaged with the previous New Products. This operation simulates the imperfect normalization to white. Figure 25 (right). An illustration of the Multiresolution aspect of the Retinex calculation. The calculation uses three data planes. The Old Product is initialized to the maximum value. 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 to the maximum 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 25-Left are repeated. The Process continues until New Product at full-resolution is complete and is used as Retinex Output.

sensations using "Early Vision" mechanisms limited to simple computational steps (sum, difference, reset, average). These lightness sensations are then available to calculate perceptions of depth and illumination. We can model lightness sensations using only the radiance on the retina. So far, we can account for the lightness in "Diamond Wall" experiments with "Early Vision" mechanisms. It makes calculating the perceived depth and illumination easier to calculate, if lightness sensations are out of the way. We can use calculated lightness as the input to shape from shading to help calculate apparent depth.

A model for "Diamond Wall" experiments

The three parts of this paper break a single intellectual framework (diamond walls) into three familiar frameworks (simultaneous contrast, receptor pooling, and B&W Mondrians). These constructs do not require three independent submodels of vision. Although quantitatively predicting the lightnesses all 30 experiments described above is beyond the scope of this paper, a single model is not unreasonable. The original Retinex model set out to predict B&W Mondrians. Early results[9] showed that appropriate processing parameters enabled the model output to be predict simultaneous contrast. Recent experiments[16-17] showed that a multiresolution Retinex model will account for low-spatial frequency sampling effects. By parallel, instead of sequential, combination of each multiresolution output, the model can account for observer data in Part 2. The incorporation of simultaneous contrast and receptor pooling does not affect the model's ability to predict Mondrians. A completely successful model of vision would quantitatively predict all three intellectual frameworks. The discussion of three frameworks does not imply three models, just three familiar ways to think about our very complex visual system.

Conclusions

This paper has reviewed a wide variety of visual displays all dealing with the "High Vision" and "Early Vision" theories.

In Part 1, Adelson's "Diamond Rows" can have either a depth / shadow perception explanation, or a simultaneous contrast sensation explanation. The experiments here argue for the sensation mechanism for three reasons. First, as with the Katz-Albers's effect, any edge added to the anomalous target C restores simultaneous contrast. Second, the long rows of diamonds, helpful in creating apparent illumination, are not necessary. The same lightness effects are found in simple, one quarter diamond displays. Third, simultaneous contrast is restored when the tips in Display D are reversed. The tips, now inconsistent with a depth/shadow perception, perform the same "releasing" function .

In Part 2, we found the "Diamond Wall" displays in the "Straight Edge" effect were very different in the very-low-spatial frequency components. We made an improved display that still exhibited the effect. We showed that this improved image did not remove all the differences in low-spatial frequency input and that sampling can the cause of the effect.

In Part 3, we showed that the largest lightness shifts from the Logvinenko Gradient display are similar to B&W Mondrians and can be explained by an "Early Vision" calculation.

All of these experiments share the same categorization. They are all explicable by "Early Vision" mechanisms, since simultaneous contrast, or multiresolution sampling or a Retinex model can predict their appearance. The corollary of that statement is that these experiments cannot be used as evidence of the existence of "High Vision" lightness mechanisms.

Acknowledgments

The author wants to thank Alexander Logvinenko and Ted Adelson for very helpful discussions and Alan Gilchrist, Michael Brill and Mary McCann for their comments and criticisms of the manuscript.

References

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

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

[3] H. Wallach, "Brightness constancy and the nature of achromatic colors", J. *Exptl. Psychol*, 38, 310-324, (1948).,

[4] T.P. Stockham, "Image Processing in the Context of a Visual Model", *Proc. IEEE,* 60, 828-284, (1972).

[5] A. Gilchrist, "Perceived lightness depends on perceived spatial arrangement", *Science,* 195, 185-187, (1975).

A. Gilchrist, "The perception of surface blacks and whites", *Scientific American*, 24,(3), 88-97, (1979).

A. Gilchrist, S. Delman, and A. Jacobsen, "The classification and integration of edges as critical to the perception of reflectance and illumination", *Perception and Psychophysics,* 33(5), 425-436, (1983).

A. Gilchrist, "Lightness contrast and failures of lightness constancy: a common explanation.", *Perception and Psychophysics*, 43 (5), 415-424, (1988).

A. Gilchrist, "*Lightness, Brightness and Transparency*, Hillsdale: Lawrence Erlbaum Associates, (1994).

J. Cataliotti, & A. Gilchrist, "Local and global processes in lightness perception" *Perception and Psychophysics*, 57(2), 125-135, (1995).

[6] E. H. Adelson and A. P. Pentland, "The perception of shading and reflectance, in Channels" in *Visual Nervous System: Neurophysiology, Psychophysics and Models*, Ed.. Blum, London: Freund, 195-208, (1991).

E. H. Adelson, "Perceptual Organization and the Judgment of Brightness", *Science*, 262:2042-2044 (1993).

P. Sinha, and E. H. Adelson, "Recovering Reflectance in a World of Painted Polyhedra", *Proc. of Fourth International Conference on Computer Vision* , 4, 156-163, Berlin; May 11-14 (1993).

P. Sinha, and E. H. Adelson, "Verifying the Consistency of Shading Patterns and 3-D Structures", *Proc. IEEE Workshop on Qualitative Vision*, 71-81, New York, NY; June 14 (1993).

E. H. Adelson, "Lightness Perception and Lightness Illusions" in ., *The Cognitive Neurosciences*, M. Gazzaniga, ed., Cambridge, MA:

MIT Press, 339-351 (1999).

[7] A.D. Logvinenko, "Lightness Induction Revisited", *Perception*, 28, 803-816, (1999).

[8] J.J. McCann, R. L. Savoy, and J.. Hall, Jr., "Visibility of Continuous Luminance Gradients", Vision Research, 18, 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(2), 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(7), 891-894, (1978).

J.J. McCann, "Visibility of gradients and low-spatial frequency sinusoids: Evidence for a distance constancy mechanism", *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(2), 76-79, (1978)

[9] 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).

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

[11] J .J, .McCann & K. L .Houston, "Color Sensation, Color Perception and Mathematical Models of Color Vision," in: *Colour Vision*, J.D. Mollon, & Sharpe, ed., Academic Press, London, 891-894 (1983).

J. J. McCann, "Simultaneous Contrast and Color Constancy: Signatures of Human Image Processing", in Proc of 9 th. Annual Vancouver Cognitive Science Conference, *Colour Perception: Philosophical, Psychological, Artistic and Computational Aspects*, Oxford Press, 88-101, (2000).

[12] M. D Fairchild, *Color Appearance Models* , Addison-Wesley Pub. Co., Reading, p. 153, (1997).

[13] J. J. McCann and A. Logvinenko, "Lightness from Early Vision Spatial Comparisons, Mid Vision Junctions or High Vision Perceptions : The Debate Part 2", *Perception*, Proc. ECVP, Trieste, (1999).

[14] T. N. Cornsweet, *Visual Perception*, Academic Press, New York, 270,279, (1970);

V. O'Brien, "Contour Perception, Illusion and Reality", *J. opt. Soc. Am.*, 48, 112-119, (1958).

K. J W. Craik, *The Nature of Psychology*, Cambridge University Press, London, (1966).

J, Albers, Interaction of Color, Yale University Press, New Haven, 1963.

[15] M. White, "A new effect of pattern lightness", *Perception*, 8, 413-416, (1979);

M. White, "The effect of the nature of the surround on the perceived lightness", Perception, 10, 215-230, (1981).

M. White and T. White, "Counterphase lightness induction", *Vision Res*., 25, 1331-1335, (1999);

[16] J. J. McCann, "Spatial Contrast and Scatter: Opposing Partners in Sensations", B. Rogowitz and T. Pappas, Eds., *SPIE Proc.*, 3644, 97-104, (1999).

[17] J. J. McCann, "Calculating The Lightness Of Areas In A Single Plane", B. Rogowitz and T. Pappas, Eds., *SPIE Proc.*, 3959, 243-253, (2000);

B. Blakeslee and M. E. McCourt, "A muliscale spatial filtering account of White effect, simultaneous brightness contrast and grating induction", *Vision Res.*, 4361-2377, (1999);

B. Blakeslee and M. E. McCourt, "Similar mechanisms underlie simultaneous brightness contrast and grating induction", Vision Research, 37, 2849-2869, (1997);

B. Moulden and F. A. A. Kingdom, "White's effect: a dual mechanism", *Vision Res*., 29, (1989);

F. A. A. Kingdom and B. Moulden, "A multi-channel approach to brightness coding", *Vision Res.*, 32,1565-1582, (1992);

F. A. A. Kingdom, "Old wine in new bottles? Some thoughts on Logvinenko's Lightness induction revisited", *Perception*, 28, 929-934, (1999);

J. M. Hans du Buf and S. Fischer, "Modeling brightness perception and syntactical image coding", *Optical Eng.*, 1900-1900, (1995);

B. Spehar, A, Gilchrist, and L. Arend, The critical role of relative luminance relations in White's and grating induction", *Vision Res.*, 35, 2603-2614, (1995).

[18] J. J. McCann and R. L. Savoy, "Measurement of Lightness: Dependence on the Position of a White in the Field of View ", B. Rogowitz, Ed., *SPIE Proc*., 1453, 402-411, (1991).

[19] J. J. McCann, E. H. Land and S. M. V. Tatnall, , "A Technique for Comparing Human Visual Responses with a Mathematical Model for Lightness", *Am. J. Optometry and Archives of Am. Acad. Optometry*, 47(11), 845-855, (1970).

[20] J. Lettvin, "The Colors of Colored Things", *MIT Research Laboratory of Electronics Quarterly Progress Reports*, 87, 193-225, (1967).

D. Todorovic, Lightness and Junctions, *Perception*, 26, 379-395, (1097).

Q. Zaida, B. Spehar, and M. Shy, Induced effects of backgrounds in three-dimensional configurations: the role of T-junctions", *Perception*, 26, 395-408, (1997).

[21] B. Funt and F. C. and J. J. McCann, "Retinex in Matlab", *Proc. IS&T/SID Eigthth Color Imaging Conference*, Scottsdale, in press, (2000).

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. Along with Giordano Berreta, he co-chaired the IS&T /SPIE Electronic Imaging 2000 Meeting. He is currently consulting and continuing his research on color vision. mccanns@tiac.net