The Perception
of Surface Blacks and Whites

What shade of gray a surface appears is related to the perceived distribution of light and shadow, which in turn depends on the perceived spatial relation between the surface and its neighbors

by Alan L. Gilchrist

The lens of the human eye projects onto the retina a two-dimensional image of the three-dimensional physical world. Partly because the retinal image is two-dimensional most investigators of color perception have assumed that depth perception has nothing to do with the process by which the human visual system determines the color of objects. The experiments I shall describe here invalidate that assumption. I have found that a change in the perceived spatial orientation of a surface can change its perceived color from black to white or from white to black.

Traditional explanations of color perception assign no special role to the central nervous system. Because that system governs depth perception my work assigns it a major role. Traditional investigations have also largely ignored the perception of the intensity and color of the light illuminating surfaces whose colors are under scrutiny. In my work I have tried to rectify this imbalance, and for the sake of simplicity I have dealt mainly with what are called neutral colors: white, black and gray.

When white light, which consists of a balanced mixture of all the colors in the spectrum, strikes a colored surface, the surface reflects one wavelength (one color) more than it does the others. This dominant wavelength corresponds to the physical color of the surface. In general the light that illuminates objects in nature is not white but is an unbalanced mixture of various colors. A colored surface reflects some wavelengths in such a mixture more than others, but now the dominant wavelengths do not correspond to the physical color of the surface. The mixture of wavelengths in the light the surface reflects to the eye depends not only on the physical color of the surface but also on the mixture of wavelengths in the light that illuminates the surface. Therefore the visual system should not be considered an instrument that measures the wavelength and the

intensity of reflected light; such measurements would reveal little about surface color. The visual system has the remarkable ability to correctly perceive the physical color of a surface in spite of wide variations in the mixture of wavelengths in the illumination. This is the phenomenon of color constancy.

Neutral surfaces leave unchanged the mixture of wavelengths that illuminates them, but they do alter the intensity of the light. The shades of gray from white to black are all neutral, and so they reflect various degrees of illumination. Surfaces that reflect between about 80 and 90 percent of the illumination are called white, whereas those that reflect between about 3 and 5 percent are called black. In short, the lightness, or perceived grayness, of a neutral surface corresponds to its reflectance: the percentage of illumination it reflects. Again, the visual system has the remarkable ability to determine the lightness of a neutral surface in spite of wide variations in the intensity of the illumination. Such is the phenomenon of lightness constancy, which is analogous to the phenomenon of color constancy for colored surfaces.

A physicist would determine the reflectance of a surface by comparing the intensity of the reflected light with the intensity of the illumination. The human visual system also determines reflectance, but it does so by comparing the intensity of the light reflected from a surface with the intensity of the light reflected from neighboring surfaces. For the visual system to make this comparison there must be constant relative displacement between the retinal image and the retina. This happens naturally because the eye constantly and involuntarily flicks back and forth in tremors with a frequency of between 30 and 150 cycles per second. It has been found that the visual field goes blank in one to three seconds if the relative displacement is

eliminated by artificially stabilizing the image (using a special apparatus that causes the image to remain still on the retina even as the eye moves back and forth). The eye tremors indicate that the receptor cells of the retina function only in the presence of a change of stimulation. Consider what happens to an individual receptor cell during the tremors (assuming that there are no large voluntary eye movements). A cell in the interior of a patch of homogeneous stimulation within the retinal image receives no change of stimulation, whereas a cell at the boundary of a patch does receive such a change.

The relevance of all of this for color vision was elegantly demonstrated by John Krauskopf of Bell Laboratories. He built a display that consisted of a disk of one color, say green, surrounded by an annulus, or ring, of another color, say red. He arranged matters so that the green-red boundary between the disk and the annulus was stabilized, that is, moved in such a way as to follow the tremors. In this way all displacement between the green-red boundary and the receptor cells was eliminated. As a result the retinal boundary disappeared and the observer simply saw a single large red disk.

The most straightforward interpretation of this experiment is that the eye sends the brain only information about changes in light across boundaries, with areas where no change is reported being filled in by the brain as homogeneous. In Krauskopf's display, and in the normal viewing of a large red disk, the eye extracts information from the outer edge. Since no change is reported within the boundary, the brain treats the interior of the disk as homogeneous. Consider what this implies. Krauskopf's observers perceived red even in the center of the display, even though green light was striking the corresponding region of the retina. Hence if the color of the light (green) striking the disk region of the
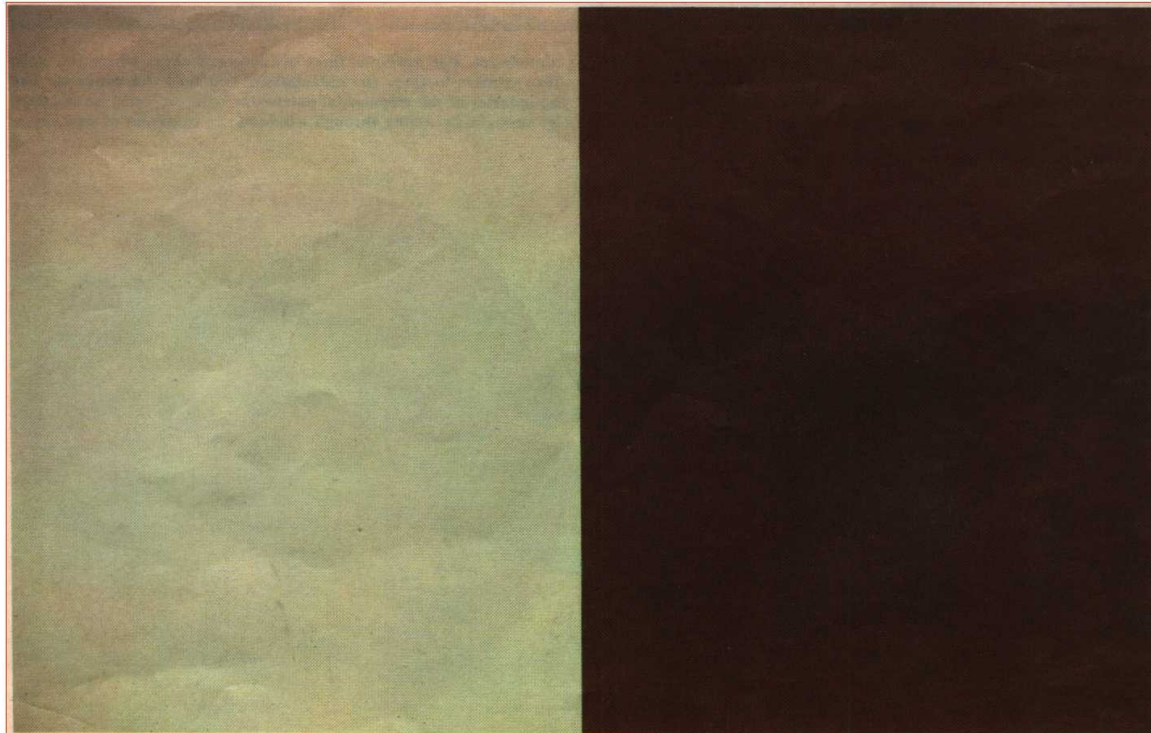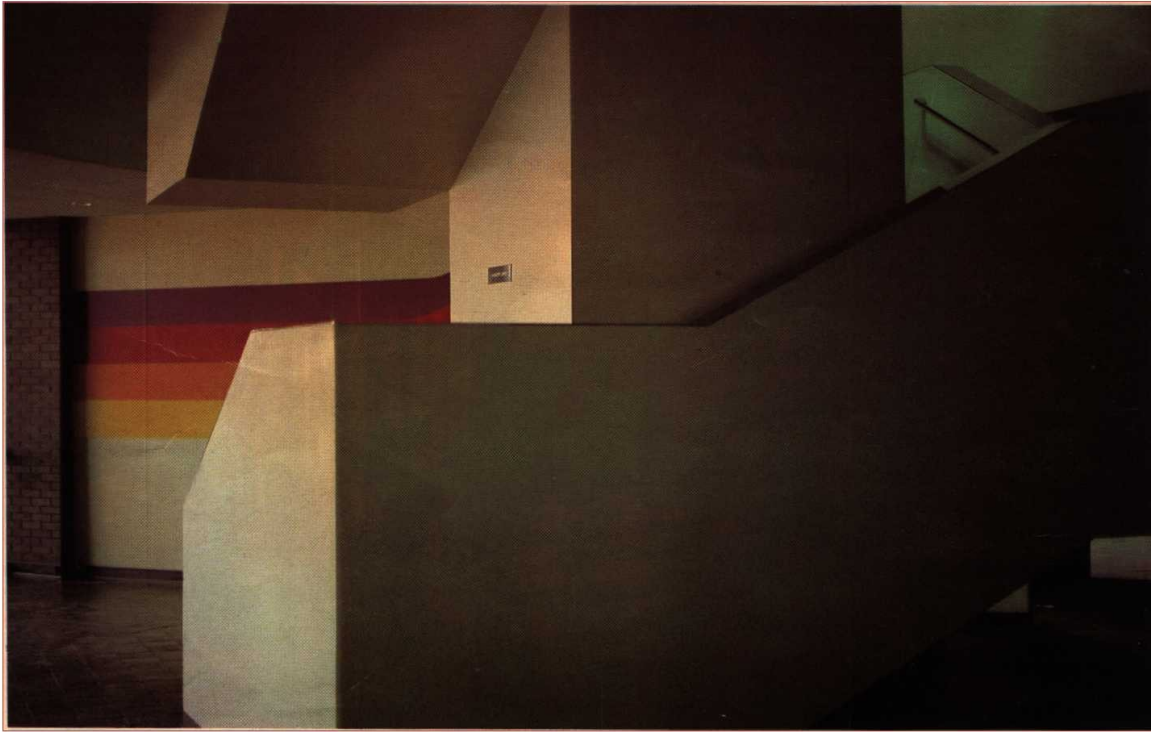
PERCEIVED SPATIAL ORIENTATION affects color perception. In the top photograph it is apparent that the right side of the stairwell is a white surface in shadow and that the left side is an illuminated white surface. Because the visual system recognizes that the stairs turn a corner it correctly perceives that the stairwell is differentially illuminated. The bottom photograph is a closeup view of the stairwell corner. Here the two sides appear to form a plane rather than a corner. The sides are perceived to be illuminated equally, and so the visual system attributes the difference in grayness to the sides themselves. Hence the shadowed side is not correctly perceived as being white.
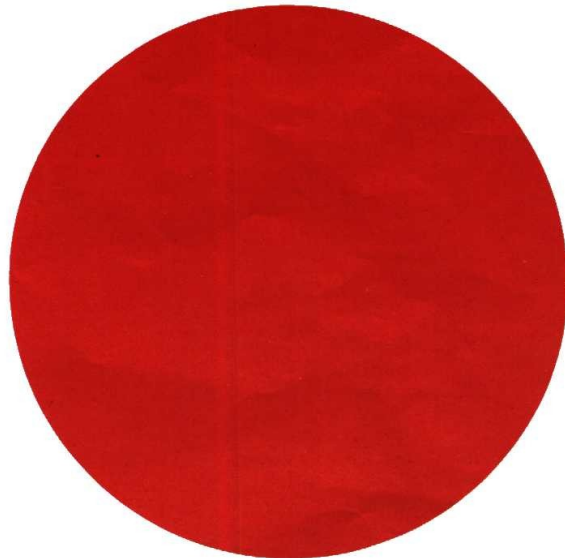
ILLUMINATION EDGES, or illumination differences, are of two kinds: attached and cast (projected). Attached edges are the result of changes in the spatial orientation of surfaces; cast edges are the result of shadows. In this photograph both kinds of edges are present in abundance. For example, there are attached edges where the walls turn corners to form the indentations that hold the windows. The boundaries of the trapezoidal patches of light projected on the floor by sunlight streaming through windows are examples of cast edges.

RETINAL RECEPTOR CELLS gather color information by detecting changes in light. To detect such changes there must be constant relative displacement between the retinal image and the retina, and so the eye scans a surface by constantly flicking back and forth in tremors with a frequency of between 30 and 150 cycles per second. In the display at the left the boundary between the red ring and the green disk was moved in such a way as to follow the tremors. That eliminated all displacement between the boundary and the cells. As a result retinal boundary disappeared and observer saw a red disk (*right*). The experiment was done by John Krauskopf of Bell Laboratories.

retina had nothing to do with the color (red) that was perceived there, then by the same token the color of the light (red) striking the annulus region might have nothing to do with the color (red) that was perceived there.

Krauskopf's display suggests that even under natural conditions the perceived color of a surface depends not on the light emanating from the surface but on the change in the light across the boundary of the surface. The perceived color of a surface cannot depend solely on the change in the light at the edge of the surface because such a change is strictly relative. The change in the light at the edge of a surface depends as much on the color of the background as it does on the color of the surface.

If perceived color depended only on the change in the light at the edge of a surface, the surface would look radically different in color when seen against different backgrounds. Conversely, two surfaces of radically different color could look the same. For example, the change in the intensity of reflected light from a white surface to a middle-gray background is the same as the change in the intensity from a middle-gray surface to a black background. Of course, white and middle-gray surfaces would not look the same when they were seen against these respective backgrounds. Such changes in background color can result in what are called contrast effects, but these effects are not nearly large enough to make the white surface on the middle-gray background look the same as the middle-gray surface on the black background. It is a remarkable empirical fact that perceived surface color remains largely constant in spite of changes in the background, which in turn give rise to changes in the edge information.

The visual system extracts only relative information from an edge, and such information is a small part of what the visual system needs in order to perceive color, a part that can be properly interpreted only in the context of the relative information from many other edges. There is evidence that the visual system integrates neural signals emanating from distinct edges. Recording changes in light and then integrating them is mathematically equivalent to keeping a point-by-point record of the light. Simply recording the changes, however, is more efficient and requires less transmission capacity. Consider an analogous situation in the stock market. To draw a graph showing how the price of a certain stock varies from day to day, one could either find out the price of the stock each day or find out the price only on days when it had changed. The latter method requires less transmission capacity, and so it might be preferred if getting the stock quotation required a transcontinental telephone call.

Yet the information collected by this method could be easily integrated to give the stock prices for each day. I believe something similar to this process of information extraction and integration is going on in the visual system. Such a process has also been suggested by Edwin H. Land of the Polaroid Corporation.

It turns out, however, that the model of reflectance-edge extraction and integration is quite limited in its application. The model works well when the objects whose colors are under scrutiny are uniformly illuminated. It even works when the illumination is changing uniformly over the entire visual field, because changes in the overall illumination have no effect on the relative changes in the light at the boundary between two shades of gray. This is just a special case of the general phenomenon that a change in illumination does not change the ratio of the amount of light reflected off adjacent surfaces of differing reflectance. For example, consider two surfaces whose reflectances are respectively 25 and 50 percent. The ratio of reflected light will remain 1 : 2 regardless of the intensity of the illumination. It has been demonstrated that the neural signal generated at an edge in the retinal image remains the same if the intensity ratio at the edge remains the same.

The model of reflectance-edge extraction and integration runs into trouble when the objects under scrutiny are not uniformly illuminated. The trouble begins with spatial changes in the illumination. Such illumination edges are of two kinds: attached and cast (projected). Attached edges arise from changes in the planarity, or spatial orientation, of surfaces. If two walls of the same color meet at a corner and one wall receives more illumination than the other, then there is an attached illumination edge where they meet. Such edges would also appear in a white plaster-of-Paris sculpture where irregularities in the sculpture's shape are clearly visible because the intensity of the incident illumination varies as the angle between the light source and various areas of the sculpture changes. On the other hand, the boundary of a shadow is a cast illumination edge. So is the boundary of a spotlight projected onto a stage and the boundary of a bright trapezoidal patch of light projected onto a floor by sunlight streaming through a rectangular window.

The presence of illumination edges in the retinal image makes the simple model of reflectance-edge extraction and integration unworkable. If changes in illumination were extracted and integrated right along with reflectance changes, gross errors in perceived lightness would result. For example, a white surface in shadow might appear as a gray surface. Nevertheless, the model can be

further developed to eliminate this difficulty. What is necessary is that before the edges are integrated into an intensity image they must be classified either as reflectance changes or as illumination changes. Once this classification is made, the visual system could separately integrate the two kinds of edge. The result would be two distinct images: one representing surface reflectance and the other representing surface illumination.

My own work has centered on the conditions under which the visual system can distinguish reflectance edges from illumination edges and on how the distinction might be made. I have been guided by the belief that the classification of edges cannot be carried out at the level of the retina. The fact is that spatial changes in illumination and changes in reflectance generate identical edges on the retina. Some workers have suggested that the retina could classify the edges based on their sharpness, with reflectance edges tending to be sharp and illumination edges tending to be gradual. This simple explanation does not work, however, because illumination edges (both cast and attached) are often sharp whereas reflectance edges are sometimes gradual. And yet the visual system can correctly identify the two kinds of edge. Rather than classifying edges solely on the basis of their sharpness on the retina the visual system must classify them on the basis of their relation to all the other edges in the entire retinal image. This means that not only the eye but also the central nervous system play a role in the classification of edges.

That the central nervous system is involved in color perception is suggested by a simple phenomenon Irvin Rock of Rutgers University pointed out to me. He noted that the shadowed side of a corner (for example, where two walls meet) usually seems to be approximately the same shade of gray as the illuminated side. When the corner is artificially made to appear flat by viewing it through a hole, however, the shadowed side seems to be an extremely dark gray, often even black. When the corner is seen this way, no shadow is perceived; the darkness is attributed to the surface itself.

Here is a situation where the edge at the corner can be perceived as either an illumination edge when it is seen correctly in depth or as a reflectance edge when it is made to look flat. This suggests that the visual system relies on depth information to determine whether it is the illumination or the reflectance that is changing at an edge. Moreover, the situation seems to challenge the prevailing view that lightness is determined by intensity ratios between adjacent regions of the retinal image independently of where these regions are perceived to

be in three-dimensional space. To demonstrate conclusively that depth perception (and hence the central nervous system) plays a role in color perception, I tried to reproduce the corner phenomenon under strict laboratory conditions.

I suspended in midair from a hidden support two white surfaces that met to form an outside right-angled corner. Behind the surfaces I put a uniform background of medium intensity. One of the white surfaces received 30 times more light than the other surface. Observers viewed the display and indicated the apparent lightness of each surface by selecting a matching sample from a chart of various shades of gray. A second group of observers viewed the surfaces when they were made to seem flat by being looked at through a hole with one eye.

When the two surfaces looked flat, the difference in intensity was perceived as a difference in reflectance, as I had expected. In other words, the illuminated side (a value of 30) looked white and the shadowed side (a value of 1) looked black. To my surprise, however, the observers saw exactly the same thing when the surfaces appeared to be at right angles to each other. This is essentially what earlier investigators had found in similar circumstances. In a few studies the perceived spatial pattern had been found to have a small effect on perceived lightness, but in most other studies, including my own, there was no such effect.

I was so impressed by the strength of the effect of spatial orientation under natural conditions that I felt something significant about the natural context of the corner must have been lost in my attempt to reproduce the situation in the laboratory. I tried in several ways to enrich the context of the laboratory display. For example, I placed objects in the vicinity of the suspended corner so as to make the lighting conditions manifest to the observers. All these attempts failed to change the results.

Finally I found the effect of spatial orientation when I put a black surface next to and coplanar with each of the original white surfaces. As in the original display, the 30 : 1 ratio in intensity between the white surfaces was seen as a difference in reflectance when they appeared to be coplanar. When the two surfaces were seen at right angles to each other, the same intensity ratio was perceived as an illumination difference. In other words, when the surfaces were seen as being coplanar, the illuminated surface looked white and the shadowed surface looked black, but when the surfaces were seen as an outside corner, the shadowed surface was seen correctly as a white surface in shadow.

A reasonable interpretation of these results is that the visual system operates to account for the entire pattern of retinal intensities as generated by either changes in illumination or changes in reflectance. The original two-surface display could be seen as being uniformly illuminated because the two intensity levels (30 and 1) could be treated as resulting from different shades of gray, namely white (30) and black (1). The more complex four-surface display involves a 900 : 1 range of intensities. (The illuminated white surface is 900 times more intense than the shadowed black surface.) Since the range in intensity of shades of gray that can result solely from reflectance is normally limited to 30 : 1, the four-surface display also requires the perception of some variations in illumination. Which edges will be treated as representing illumination changes and which as representing reflectance changes depends on the overall organization of the scene, particularly the perceived three-dimensional layout.

Next I modified the four-surface display to rule out the possibility that the results were caused by some kind of contrast effect operating between adjacent areas of the two-dimensional retinal image. The horizontal plane of the modified display consisted of a large white square attached to a small black trapezoidal tab that extended outward toward the observer [see illustration on page 118]. The vertical plane consisted of a large black square attached to a small white tab that extended upward. In other words, the tabs extended out into midair from the corner formed by the two large squares, and so each tab appears on three sides against a background that is not in the same plane. As in the original four-surface display, the horizontal plane received 30 times more light than the vertical plane. The tabs are trapezoidal in order to create an illusion in depth perception.

Seen with both eyes the tabs correctly appear as trapezoids lying in their actual planes. In this case the vertical tab appears to be almost white and the horizontal tab almost black. On the other hand, seen through an aperture with only one eye each trapezoidal tab appears to be a small square lying in the same plane as the larger square that surrounds it on three sides. In this case the perceived colors are the reverse: the vertical tab appears to be black and the horizontal tab appears to be white.

In both cases, however, the retinal image is the same: each tab is seen against a square background that surrounds it on three sides. This means that the relation between the target and its background in the retinal image is irrelevant to the target's perceived shade of gray. The shade of gray turns out to depend on the relation between the intensity of the target and the intensity of whatever surface
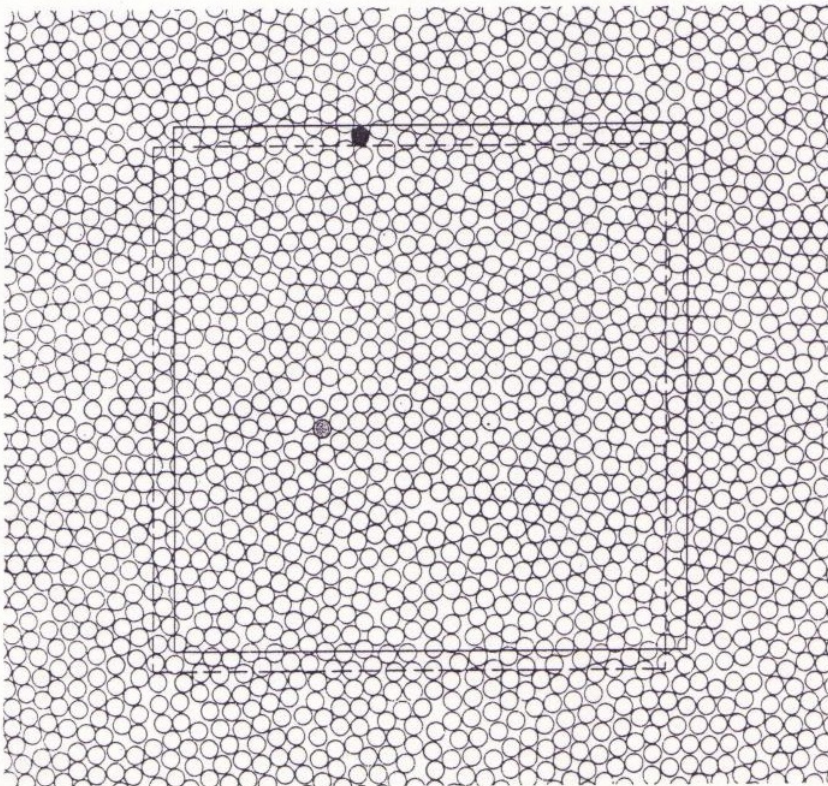

STIMULATION OF A RECEPTOR CELL (*black dot*) changes as the edge of the retinal image moves across the cell when the image shifts from one position (*solid rectangle*) to another (*broken rectangle*) because of the eye tremors. Cells such as the colored one that the boundary of the retinal image does not move across do not receive a change of stimulation. Since the green-red boundary in Krauskopf's display did not cross cells, the display appeared to be red.

seems to lie in the same plane, even if that coplanar surface does not provide the background of the target in the retinal image. Therefore when the display is seen with both eyes, the shadowed white vertical tab appears to be coplanar with the shadowed black vertical background, and so the tab appears to be white. In addition the illuminated black horizontal tab appears to be coplanar with the illuminated white horizontal background, and so the tab looks black. When the display is seen with one eye, the shadowed white vertical tab appears to be coplanar with the illuminated white horizontal background, and so the tab looks black. Moreover, the illuminated black horizontal tab appears to be coplanar with the shadowed black vertical background, and so the tab looks white.

The model of edge extraction, classification and integration was put to a more direct test in a series of experiments designed to probe what observers would see if a reflectance edge were mistakenly perceived as an illumination edge or vice versa. The corner experiments I have described showed that a change in the perceived type of edge results in a change in the perceived lightness of the region bounded by the perceptually altered edge. If perceived lightness depends on the visual system's integration of a series of spatially remote edges, then even stranger perceptual mistakes will manifest themselves. For example, if a single reflectance edge in a complex display is made to appear as an illumination edge, then in certain circumstances this ought to have as great an effect on the perceived lightness of regions that are remote from the altered edge as it has on the regions that are bounded by the altered edge.

At the State University of New York at Stony Brook, Stanley Delman and I designed one such edge-substitution experiment. We studied a familiar display, known as simultaneous lightness contrast, that is often treated in psychology textbooks. Two middle-gray squares are placed respectively on adjacent white and black backgrounds. It turns out that for some reason the gray square on the white background looks slightly darker than the gray square on the black background. This effect, however, was not the subject of our experiment. We wanted to reproduce the same pattern of reflected light that the eye receives from this display when the boundary between the white and black backgrounds was perceived as the boundary between high and low levels of illumination.

To this end we affixed to a wall a large rectangular piece of middle-gray paper that would serve as the immediate background for both target squares. The difference in immediate background intensities was created by casting a beam of light with a sharp edge across half of
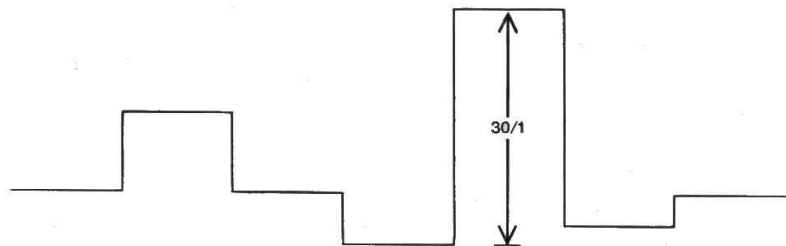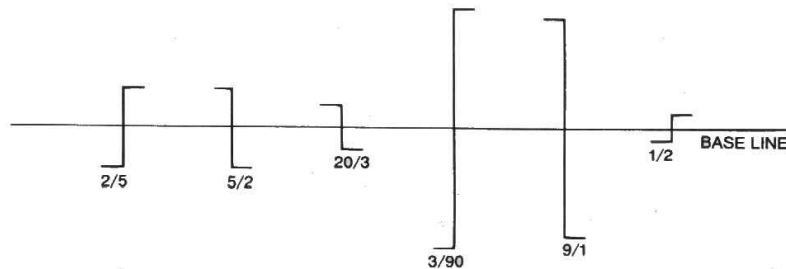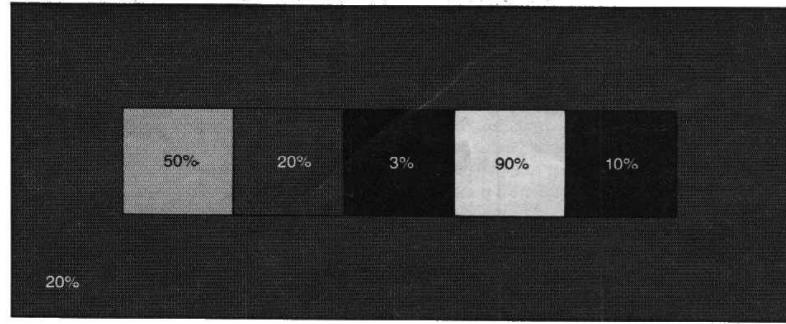
REFLECTANCE OF A SURFACE is the percentage of illumination that the surface reflects. The top part of the illustration indicates the reflectances of several surfaces in a visual scene. The eye responds only to changes in the intensity of the reflected light. The middle part of the illustration indicates the kind of information that is generated at edges in the retinal image. Each vertical line indicates the change in intensity at an edge. The visual system integrates information from all the edges in order to reconstruct the original intensity pattern (bottom).

the gray rectangle. Part of the beam also illuminated the background wall on three sides of the rectangle. The intensity of the beam was set so that the intensity ratio of the illuminated half of the gray rectangle to the shadowed half was 30 : 1, which is equal to the intensity ratio of white paper to black paper, as was the case in the original display. Next we needed to add the target squares. They could not be the same shade of gray (that is, have the same reflectance), because they would then reflect different amounts of light to the eye owing to their unequal illumination.

In the original lightness-contrast display that we were trying to reproduce the two squares reflected exactly the same amount of light. It is also true that in the original display the squares had the same reflectance. That did not con-

cern us, however, because our goal was to reproduce not all the facets of the original display but only the information the visual system received: the pattern of light intensities that reached the eye. To achieve this goal the target square on the shadowed side had to have a reflectance 30 times as great as the other target square, which would be receiving 30 times as much illumination. Therefore we placed a white square on the shadowed side and a black square on the illuminated side.

In this way we thought we had reproduced in a logical manner the light intensities of the original pattern. Now it was time to test our results photometrically and empirically. Photometric measurements showed that the targets reflected equal amounts of light, that the

117

immediate background on one side reflected 30 times more light than the immediate background on the other side and that on a logarithmic scale the target intensities were halfway between the intensities of the two immediate backgrounds. All of this was true of the original display as well. Moreover, the reproduction preserved the geometry of the original.

Empirical considerations also indicated that our reproduction was faithful. Under control conditions observers viewed our reproduction through a rectangular aperture in a cardboard screen. The screen limited the observers' view so that the only parts of the display that were visible were the two square targets and their immediate backgrounds: the illuminated and shadowed halves of the gray rectangle. Under these conditions the display looked just like the original. The immediate backgrounds looked respectively white and black even though they were actually a middle gray that had been differentially illuminated. The illumination edge between the immediate backgrounds was perceived as a reflectance edge, and the square targets looked middle gray, one slightly darker than the other.

All of this served merely to assure us that we had correctly reproduced the pattern of light reflected by the original display. The real test came when we had

observers view the display without looking through the confining cardboard screen. Under these conditions the observers could clearly see that a beam of light was illuminating half of the rectangular piece of paper since the beam of light illuminated the part of the background wall that was behind this half of the rectangle. Now the illuminated target was correctly perceived as black whereas the shadowed target was correctly perceived as white.

The perceived lightness of each target changed radically from the screened display to the unscreened one in spite of the fact that the amount of light each target reflected remained the same, that the amount of light each immediate background reflected remained the same and hence that the neural signal generated at the retinal edge of each immediate background must have remained the same. This clearly means that color perception is not simply a function of the amount of light a surface reflects, of the intensity ratio of a surface to its immediate background or of the neural signal generated at the retinal edge of a surface.

It seems that a much more relativistic process is going on in the visual system. The boundary between each target square and its immediate background gives only the relation between the light reflected by each of these areas. If one target is to be compared with the oth-

er, the relation between the immediate backgrounds must be taken into account. If the immediate backgrounds are perceived to differ radically in reflectance, then this perception, together with the edges between the targets and the backgrounds, suggests that the targets are nearly the same in lightness. If, on the other hand, the immediate backgrounds are perceived to differ radically in illumination, then the targets appear to lie against the same (reflectance) background and hence appear to differ extremely in lightness.

It is now time to consider how the model of edge classification incorporates the above results. In the screened case where the display looked like the original lightness-contrast display the changes at the edges could simply be extracted and integrated to form an intensity profile of the display. In the unscreened case the same intensity profile would be generated if the edges were not classified either as illumination edges or as reflectance edges. If, however, the edges are classified before the integration, and if the integration is done only within each class of edges, then two separate profiles are generated. In that case the reflectance profile of the screened display should be different from the reflectance profile of the unscreened one. The difference is that the middle edge representing the relation between the immediate backgrounds is present only in the reflectance profile of the screened display. This means that in the unscreened display the targets look sharply different in reflectance (namely as white and as black) rather than appearing to be two shades of middle gray. Of course, the edge between the immediate backgrounds that is missing from the reflectance profile would manifest itself in the illumination profile. Thus the illumination profile would show a region of high illumination next to a region of low illumination. In the screened display the illumination profile would simply be uniform.

Alfred Yarbus of the Academy of Sciences of the U.S.S.R. did a similar experiment using the technique by which the image on the retina is stabilized even as the eye moves back and forth. He placed a white and a black region next to each other on a red background [see illustration on page 122]. Disk-shaped targets of the same shade of red as the background were placed in the center of both the white and the black region. Under normal viewing conditions the disk on the white region would look slightly darker than the disk on the black region. Yarbus altered the conditions so as to eliminate the perceptual boundaries between the white and the black region and between these regions and the red background. This he accomplished with an apparatus that caused the physical boundaries to move along with the eye,
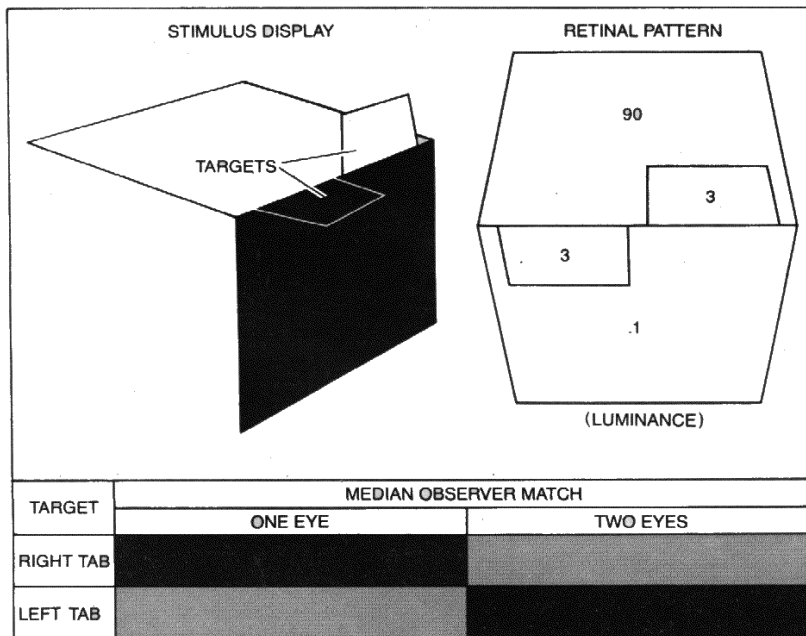
STIMULUS DISPLAY, consisting of an illuminated horizontal plane and a shadowed vertical plane, was viewed by observers who looked down at it from an angle of 45 degrees. Because the targets were trapezoids each seemed to lie in the same plane as one of the larger background squares when the display was viewed with one eye through an aperture. When the targets were viewed with both eyes, they were seen in their actual spatial positions. In both cases the retinal pattern remained the same, although the perceived lightness changed. The chart in this illustration indicates what the observers saw. The perceived shade of gray turned out to depend on the relation between the intensity of the target and the intensity of whatever surface appeared to lie in the same plane, even if that surface was not the background of the target in the retinal image.
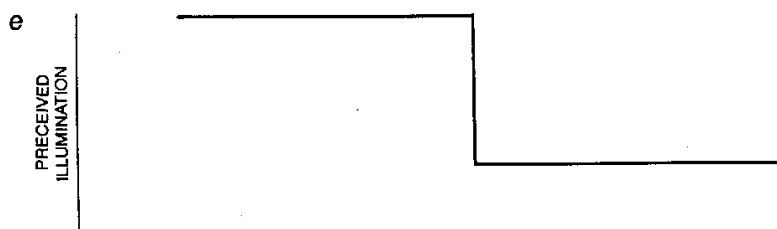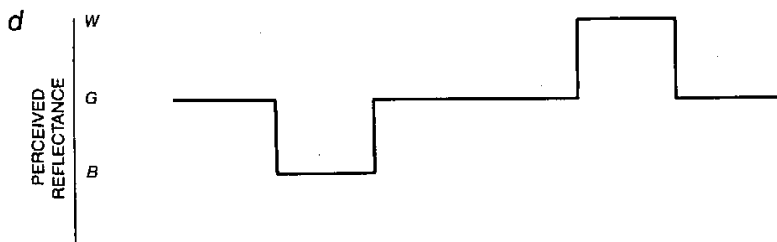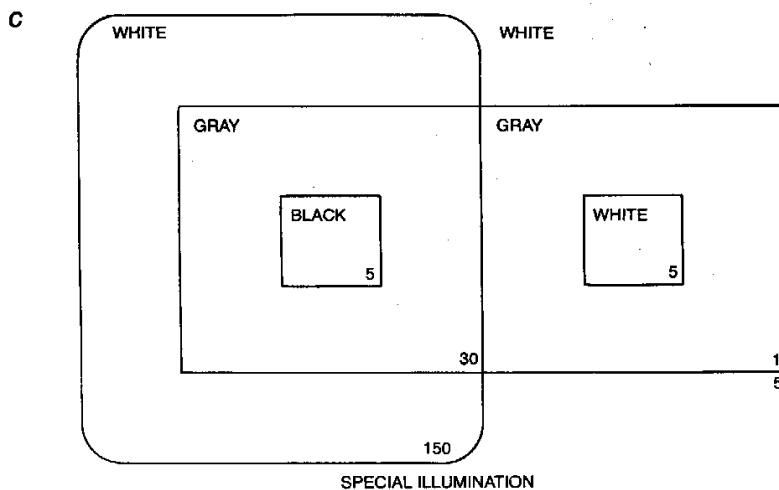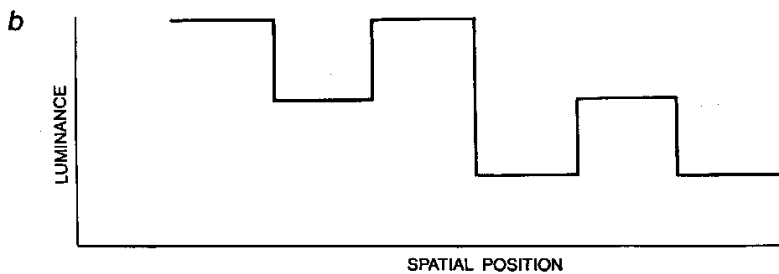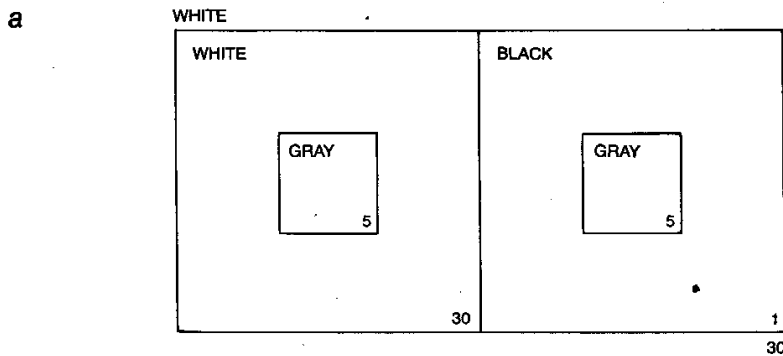
118

as Krauskopf had done with his green disk and red annulus. Since the eye registers only changes in light intensity when there is constant relative displacement between the retinal image and the retina, it could not detect the presence of the white and black regions. As a result the disks appeared to stand on a homogeneous red background. Yarbus did not report his results quantitatively, but he indicated that one disk (the one objectively on the white region) appeared very dark whereas the other disk appeared very light.

Yarbus' results were essentially the same as mine, although there were significant differences in our respective methods that are worth exploring. In Yarbus' experiment and in my unscreened display the targets were made to look as though they had a common background by removing from the reflectance profile the edge that partitioned the background. As far as perceived lightness goes the effects are the same. What is fascinating is that Yarbus removed the edge at the point of extraction, whereas I removed the edge at the point of classification. To put it another way, in Yarbus' experiment the eye never detects the edge, whereas in my unscreened display the eye detects it and then classifies it correctly as an illumination edge. In Yarbus' experiment there is no reason to believe the observers saw the targets as being differentially illuminated.

Yarbus' experiment gives direct evidence for the validity of the concept of edge extraction and integration. The edge of each disk did not change; rather, boundaries that partitioned the background disappeared. As with Krauskopf's green disk and red annulus, the region bounded by the disappearing edge seemed to change color. Here the white and black regions of the background disappeared, leaving only a single uniformly colored background. Yarbus' work, however, reveals an additional phenomenon: the disks appeared to change color even though their boundaries remained constantly visible. The obvious conclusion is that information at one edge (here the edge between a disk and its immediate background) is integrated in some way with the information at a remote edge (here the edge between the immediate background and the surrounding background) before the color of any region is finally perceived.

It would be helpful at this point to discuss why the visual system determines surface colors by comparing intensities of reflected light. As I have mentioned, making such a comparison involves extracting information from edges. This suggests that it is more useful to think of the retinal image as a pattern of edges than it is to think of it as a mosaic of color patches as traditional accounts of color perception would

a

WHITE

WHITE	BLACK
GRAY 5	GRAY 5
30	1
	30

b

LUMINANCE

SPATIAL POSITION

c

WHITE	WHITE
GRAY	GRAY
BLACK 5	WHITE 5
30	1
150	5

SPECIAL ILLUMINATION

d

PERCEIVED REFLECTANCE

W

G

B

e

PERCEIVED ILLUMINATION

have it. The edges are of two kinds, and so the retinal image is actually a dual image.

The retina is a light-sensitive surface designed to register gradients in the optic array. There are basically two kinds of change in the physical world that result in gradients of stimulation on the retina: change in surface reflectance and change in illumination. Each kind of change is capable by itself of producing a complex pattern of stimulation on the retina. From a large expanse of uniformly illuminated wallpaper the retina receives a pattern arising solely from changes in surface reflectance. On the other hand, from the rumpled white sheets of an unmade bed or from a snow-covered landscape the retina receives a pattern of stimulation arising solely out of changes in illumination. Normally these two kinds of change are at work simultaneously. It is as if two separate patterns have been laid on top of each other on the retina. The task of the visual system is to disentangle the two patterns.

The duality of the retinal image implies that each point in the visual system has at least two values: a reflectance value and an illumination value. The perception of transparency involves a similar two-value phenomenon, because a surface of one color is seen behind a surface of another color. The same phenomenon is involved in reflections. For example, when an observer looks through a window, the observer sees two scenes at once: the actual scene outside the window and another scene reflected on the inside of the window. As a result it is possible to see two colors in the same place.

Consider an experience I had recently. A book with a red cover had been left on top of the dashboard of my car in such a way that I could see a red reflection of the book as I looked out through the windshield. I was surprised to find that distant objects retained their normal colors as they were viewed through the red reflection. Even green objects seen through the red reflection looked

green. This interested me because I of course knew that when red and green light are mixed in isolation, they form yellow. Then when I held up my hand to block the rest of the scene and viewed just a patch of the red and green through a small opening between my fingers, I did see yellow. The opening in my hand had imposed the same boundary on the red and the green light, and so they mixed to form yellow. In the original situation, however, the boundary of the red reflection did not coincide with the boundary of the green object. As a result the red and the green light appeared as elements of separate images rather than as yellow. Apparently when one speaks of mixing colors, one should actually speak of mixing edges (changes in the light) rather than of mixing light itself.

Such considerations indicate that the retina does not act as a photocell in measuring the intensity and color at each point in a scene. Man-made measuring devices are designed to respond to only one physical quantity at a time. For example, if a voltmeter is sensitive to resistance or, even worse, to temperature, the meter is considered to be defective. To cite such devices as models of the human sensory system can be extremely misleading. Unlike man-made measuring devices, the human visual system seems capable of processing multiple variables simultaneously. Accounts of color perception relying on inappropriate measuring-device models maintain that the visual system gains information about surface reflectance by sacrificing information about illumination. Such accounts unnecessarily limit the visual system to simple situations, when in fact it can handle complex situations by extracting information about changes in reflectance and illumination.

I have discussed here how the visual system deals with retinal images composed of both reflectance edges and illumination edges. It is also possible to study the perception of images that show only one kind of edge. I wondered how things would look if all the reflectance edges were missing from a situation. Alan Jacobsen and I built two miniature rooms that consisted entirely of illumination edges. We furnished each room identically with the same number of objects of varying size and shape. We painted one room, including all its contents, a matte (nonglossy) white, and we painted the other matte black. Observers viewed each room through an aperture in one of its walls, which prevented them from seeing the light bulb that illuminated the room.

It is important to keep in mind that each of these rooms projects a complex nonuniform pattern of stimulation on the retina of the observer. In the room there are sharp edges at corners and fuzzier ones across walls and curved surfaces. Some of the edges are the result of cast shadows, but all of them are
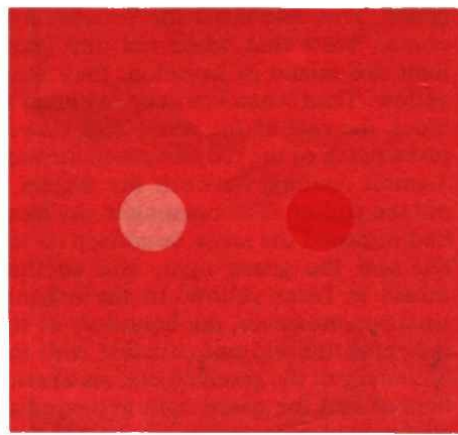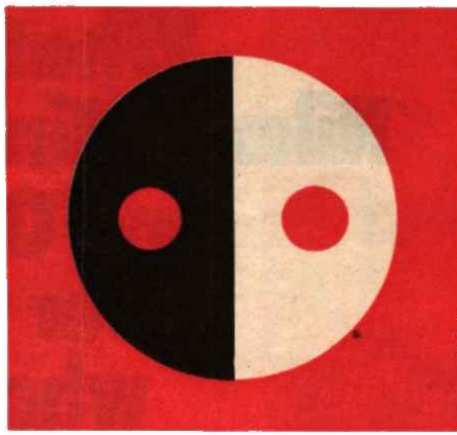
OBSERVERS viewed the display at the left under conditions where the white and black immediate backgrounds were not perceived because the boundary between them and also the boundaries between each of them and the surrounding red background were moved back and forth in such a way that they followed the eye tremors. The target disks then appeared to lie on a homogeneously red background (*right*). The boundary of the disk at the right still carried the information that the disk was darker than its immediate background, and so it appeared a very dark red. The boundary of the disk at the left still carried the information that the disk was lighter than its immediate background, so that it appeared a very light red. This experiment, done by Alfred Yarbus of the Academy of Sciences of the U.S.S.R., indicates that a change in the information at an edge affects regions that are spatially remote from the altered edge.

illumination edges, since we eliminated the reflectance edges by covering everything with paint of a single reflectance.

In presenting these displays to naïve observers we had a number of questions in mind. We wondered whether the edges in the room would be perceived as changes in reflectance, as prevailing theories of lightness perception would predict, or as changes in illumination, as the changes actually were. In other words, would each room be correctly perceived as a single shade of gray? And if it was so perceived, what shade of gray would it be? Would the black room look black and the white room white?

It turns out that 22 of our 24 observers saw each room as being uniform in lightness. They had correctly attributed the variations in intensity to variations in illumination. All of them saw the white room as white. As for the black room, all of them saw it as consisting of only a single shade of gray, although the perceived shade varied from observer to observer. The shades ranged from black to middle gray and averaged to a dark gray.

The model of edge extraction, classification and integration suggested that the illumination edges in the two rooms would be perceived as such and hence that the reflectance of each room would seem uniform. Yet there is nothing in the model so far to predict that observers could identify at least roughly the correct shade of gray. What information makes this identification possible? A plausible hypothesis is that the identification is made on the basis of the intensity of the light. In our experiment the average intensity of the light reflected from the white room is higher than the average intensity of the light reflected from the black room. By modifying our experiment, however, we invalidated

the hypothesis. We lowered the intensity of the light source in the white room and raised the intensity of that in the black room until every point in the black room reflected more light than the corresponding point in the white room. The results were the same: the white room looked white and the black room looked dark gray.

We now have a promising lead to how the visual system determines the shade of gray in these rooms, although we do not yet have a complete explanation. (John Robinson helped me to develop this lead.) We believe the operative factor is the effect of indirect illumination on the retinal pattern. Every point in our rooms received illumination from two sources: directly from the light bulb and indirectly from light reflected from other surfaces in the room. In the black room, which had a reflectance of about 3 percent, there was scarcely any indirect illumination. Direct light accounted for almost all the light shining on the various surfaces. Direct light generates sharp edges, and so the intensity profile of the black room revealed a wild pattern of up-and-down swings. Although the relative pattern of direct illumination turned out to be the same in the two rooms, there was much more indirect illumination in the white room, whose reflectance was about 90 percent. Such abundant indirect light had the effect of smearing out the sharp edges resulting from the direct illumination, so that the intensity profile of the white room revealed a smoother pattern of gradual changes.

When we reduced the illumination in the white room, the shape of the intensity profile did not change. The profile merely showed a uniform decrease in intensity. If the white room were painted

a darker shade of gray, the shape of its profile would change. Our work demonstrates that black rooms have sharper gradients than white rooms. In some way the visual system deciphers the shade of gray of a room from the shape of the intensity pattern of the room. Our work in this area is only at the beginning. As we learn more about the information stored in the intensity pattern, we hope to understand the deciphering process.

My co-workers and I have begun to extend our analysis of shades of gray to the chromatic colors. In one experiment we compared the observations of subjects who viewed a white room illuminated with blue light with the observations of those who viewed a blue room illuminated with white light. All the subjects were able to tell whether the blueness came from the surfaces or from the illumination, in spite of the fact that the light reflected off corresponding patches in the two rooms might be identical. This result clearly indicates that the perception of illumination also plays an important role in the perception of chromatic surface color.

The question of whether or not sur-face-color perception can be explained without reference to the perception of illumination has been debated since the end of the 19th century, when Hermann von Helmholtz first proposed that surface color could be determined only after illumination had been estimated. Helmholtz, however, was unable to specify how the illumination is estimated, so that most accounts of color perception have failed to refer to perceived illumination. The experiments I have described here suggest that Helmholtz was correct in his emphasis on perceived illumination but incorrect in his ideas about its relation to perceived surface color. He assumed incorrectly that the eye measures light. Current work shows that the eye compares light by extracting edge information, and so the visual system acquires information about illumination in exactly the same way that it acquires information about the colors of the surfaces in a scene. It turns out that the perception of illumination and the perception of surface color are parallel processes in the visual system involving the decomposition of the retinal image into separate patterns of illumination and surface color.

POSITION
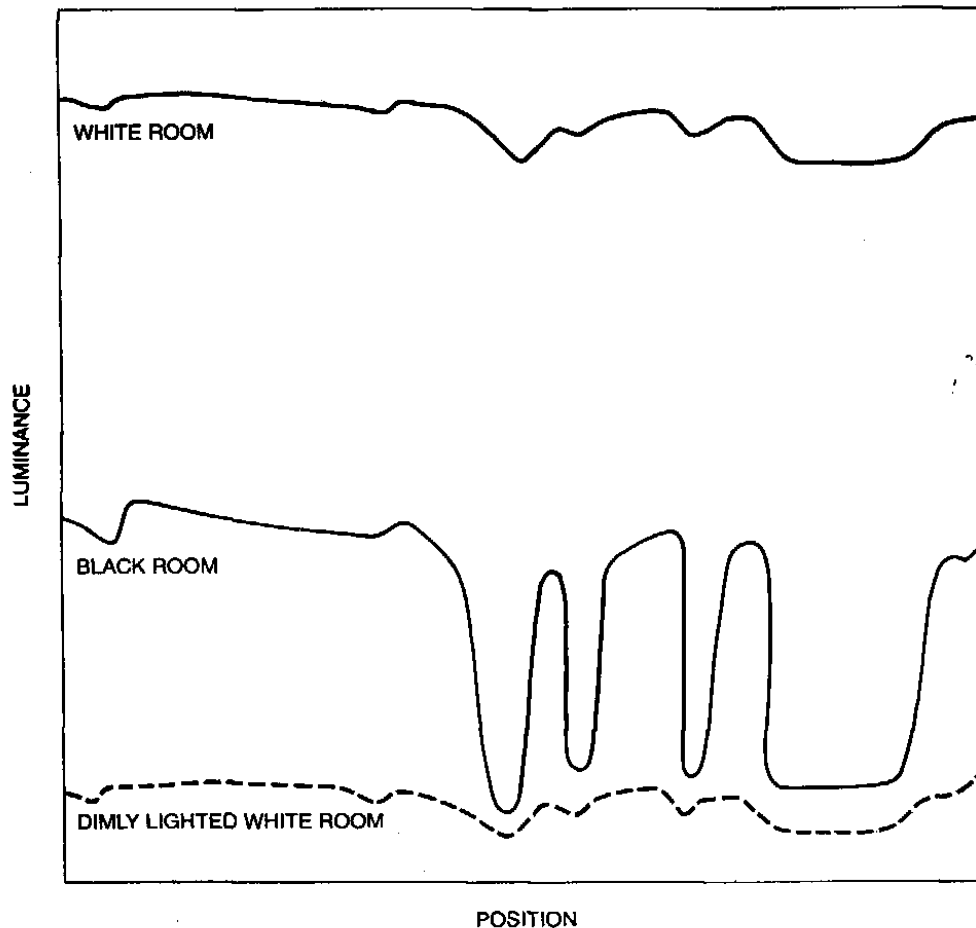
INDIRECT ILLUMINATIONS may play a role in how the visual system determines the grayness of a surface. The top curve depicts the intensity profile of a furnished room that was painted entirely white; the bottom curve, the same room dimly lighted; the middle curve, an identical room painted entirely black. In the white room abundant indirect illumination smeared out the sharp intensity edges resulting from the direct illumination, and so the intensity profiles (*top and bottom*) are comparatively smooth. In the black room there was very little indirect illumination, so that the profile (*middle*) revealed a pattern of up-and-down swings characteristic of edges generated by direct illumination. Perceived grayness does not depend on the intensity of reflected light, since the black room reflected more light than the dimly lighted white room.