8. COLOR VISION

The hundreds of dollars extra that consumers are willing to pay for color TV in preference to black and white must mean that we take our color sense seriously. A complex apparatus in the eye and brain can discriminate the differences in wavelength content of the things we see, and the advantages of this ability to our ancestors are easy to imagine. One advantage must surely have been the ability to defeat the attempts of other animals to camouflage themselves: it is much harder for a hunted animal to blend in with the surroundings if its predator can discriminate the wavelength as well as the intensity of light. Color must also be important in finding plant food: a bright red berry standing out against green foliage is easily found by a monkey, to his obvious advantage and presumably to the plant's, since the seeds pass unharmed through the monkey's digestive tract and are widely scattered. In some animals color is important in reproduction; examples include the bright red coloration of the perineal region of macaque monkeys and the marvelous plumage of many male birds.

Color serves many purposes in nature, some of which are still mysterious. The blue spots on this Garibaldi fish fade as it grows older, disappering when it matures. What significance they have for older Garibaldi fish is not known.

In humans, evolutionary pressure to preserve or improve color vision would seem to be relaxing, at least to judge from the 7 or 8 percent of human males who are relatively or

completely deficient in color vision but who seem to get along quite well, with their deficit often undiagnosed for years, only to be picked up when they run through red lights. Even those of us who have normal color vision can fully enjoy black-and-white movies, some of which are artistically the best ever made. As I will discuss later, we are all color-blind in dim light.

Among vertebrates, color sense occurs sporadically, probably having been downgraded or even lost and then reinvented many times in the course of evolution. Mammals with poor color vision or none at all include mice, rats, rabbits, cats, dogs, and a species of monkey, the nocturnal owl monkey.

Ground squirrels and primates, including humans, apes, and most old world monkeys, all have well-developed color vision. Nocturnal animals whose vision is specialized for dim light seldom have good color vision, which suggests that color discrimination and capabilities for handling dim light are somehow not compatible. Among lower vertebrates, color vision is well developed in many species of fish and birds but is probably absent or poorly developed in reptiles and amphibia. Many insects, including flies and bees, have color vision. We do not know the exact color-handling capabilities of the overwhelming majority of animal species, perhaps because behavioral or physiological tests for color vision are not easy to do.

The subject of color vision, out of all proportion to its biologic importance to man, has occupied an amazing array of brilliant minds, including Newton, Goethe (whose strength seems not to have been science), and Helmholtz. Nevertheless color is still often poorly understood even by artists, physicists, and biologists. The problem starts in childhood, when we are given our first box of paints and then told that yellow, blue, and red are the primary colors and that yellow plus blue equals green. Most of us are then surprised when, in apparent contradiction of that experience, we shine a yellow spot and a blue spot on a screen with a pair of slide projectors, overlap them, and see in the overlapping region a beautiful snow white. The result of mixing paints is mainly a matter of physics; mixing light beams is mainly biology.

In thinking about color, it is useful to keep separate in our minds these different components: physics and biology. The physics that we need to know is limited to a few facts about light waves. The biology consists of psychophysics, a discipline concerned with examining our capabilities as instruments for detecting information from the outside world, and physiology, which examines the detecting instrument, our visual system, by looking inside it to learn how it works. We know a lot about the physics and psychophysics of color, but the physiology is still in a relatively primitive state, largely because the necessary tools have been available for only a few decades.

THE NATURE OF LIGHT

Light consists of particles called photons, each one of which can be regarded as a packet of electromagnetic waves. For a beam of electromagnetic energy to be light, and not X-rays or radio waves, is a matter of the wave-length—the distance from one wave crest to the next—and in the case of light this distance is about 5 X 10 to the -7 meters, or 0.0005 millimeter, or 0.5 micrometer, or 500 nanometers.

Light is defined as what we can see. Our eyes can detect electromagnetic energy at

wavelengths between 400 and 700 nanometers. Most light reaching our eyes consists of a relatively even mixture of energy at different wavelengths and is loosely called *white light*. To assess the wavelength content of a beam of light we measure how much light energy it contains in each of a series of small intervals, for example, between 400 and 410 nanometers, between 410 and 420 nanometers, and so on, and then draw a graph of energy against wavelength. For light coming from the sun, the graph looks like the left illustration on this page. The shape of the curve is broad and smooth, with no very sudden ups or downs, just a gentle peak around 600 nanometers. Such a broad curve is typical for an incandescent source. The position of the peak depends on the source's temperature: the graph for the sun has its peak around 600 nanometers; for a star hotter than our sun, it would have its peak displaced toward the shorter wavelengths—toward the blue end of the spectrum, or the left in the graph—indicating that a higher proportion of the light is of shorter wavelength. (The artist's idea that reds, oranges, and yellows are warm colors and that blues and greens are cold is related to our emotions and associations, and has nothing to do with the spectral content of incandescent light as related to temperature, or what the physicists call color temperature.) If by some means we filter white light so as to remove everything but a narrow band of wavelengths, the resulting light is termed *monochromatic* (see the graph at the right on this page).

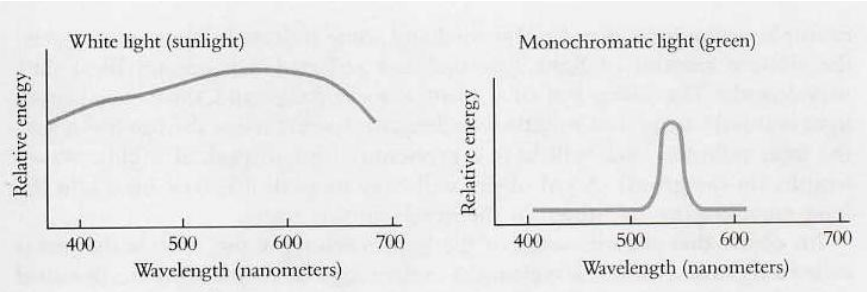

Left: The energy in a beam of light such as sunlight contains a broad distribution of wavelengths, from 400 or less to about 700 nanometers. The gentle peak is a function of the temperature of the source: the hotter the source the more the peak is displaced towards the blue, or short-wave-length, end. *Right*: Monochromatic light is light whose energy is mostly at or near one wavelength. It can be produced with various kinds of filters, with a spectroscope containing a prism or a grating, or with a laser.

PIGMENTS

When light hits an object, one of three things can happen: the light can be absorbed and the energy converted to heat, as when the sun warms something; it can pass through the object, as when the sun's rays hit water or glass; or it can be reflected, as in the case of a mirror or any light-colored object, such as a piece of chalk. Often two or all three of these happen; for example, some light may be absorbed and some reflected. For many objects, the relative amount of light absorbed and reflected depends on the light's wavelength. The green leaf of a plant absorbs long- and short-wavelength light and reflects light of middle wavelengths, so that when the sun hits a leaf, the light reflected back will have a pronounced broad peak at middle wave-lengths (in the green). A red object will have its peak, likewise broad, in the long wavelengths, as shown in the graph on the next page.

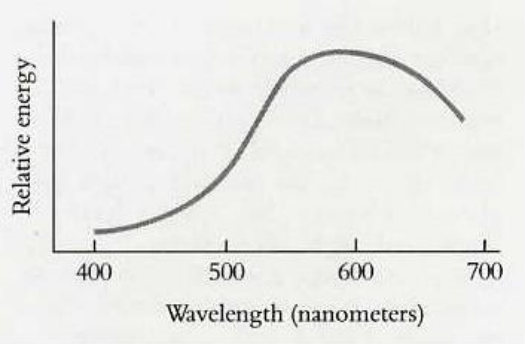

Most colored objects reflect light that is generally richer in some parts of the visible spectrum than in others. The distribution of wavelengths is much broader than that for monochromatic light, however. This graph shows the spectral content of light that would be reflected from a red object, using a broad-band (white) light source.

An object that absorbs some of the light reaching it and reflects the rest is called a *pigment*. If some wavelengths in the range of visible light are absorbed more than others, the pigment appears to us to be colored. *What* color we see, I should quickly add, is not simply a matter of wavelengths; it depends on wavelength content and on the properties of our visual system. It involves both physics and biology.

VISUAL RECEPTORS

Each rod or cone in our retina contains a pigment that absorbs some wavelengths better than others. The pigments, if we were able to get enough of them to look at, would therefore be colored. A visual pigment has the special property that when it absorbs a photon of light, it changes its molecular shape and at the same time releases energy. The release sets off a chain of chemical events in the cell, described in Chapter 3, leading ultimately to an electrical signal and secretion of chemical transmitter at the synapse. The pigment molecule in its new shape will generally have quite different light-absorbing properties, and if, as is usually the case, it absorbs light less well than it did before the light hit it, we say it is bleached by the light. A complex chemical machinery in the eye then restores the pigment to its original conformation; otherwise, we would soon run out of pigment.

Our retinas contain a mosaic of four types of receptors: rods and three types of cones, as shown in the illustration at the top of the next page. Each of these four kinds of receptors contains a different pigment. The pigments differ slightly in their chemistry and consequently in their relative ability to absorb light of different wavelengths. Rods are responsible for our ability to see in dim light, a kind of vision that is relatively crude and completely lacks color. Rod pigment, or *rhodopsin*, has a peak sensitivity at about 510 nanometers, in the green part of the spectrum. Rods differ from cones in many ways: they are smaller and have a somewhat different structure; they differ from cones in their relative numbers in different parts of the retina and in the connections they make with subsequent stages in the visual pathway. And finally, in the light-sensitive pigments they contain, the three types of cones themselves differ from each other and from rods.

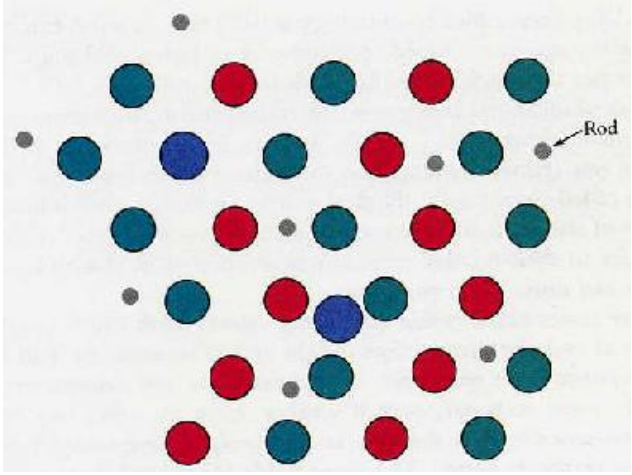

Retinal receptors form a mosaic consisting of rods and the three types of cones. This diagram might represent a part of the retina a few degrees from the fovea, where cones outnumber rods.

The pigments in the three cone types have their peak absorptions at about 430, 530, and 560 nanometers, as shown in the graph below; the cones are consequently loosely called "blue", "green", and "red", "loosely" because (1) the names refer to peak sensitivities (which in turn are related to ability toabsorb light) rather than to the way the pigments would appear if we were to look at them; (2) monochromatic lights whose wavelengths are 430, 530, and 560 nanometers are not blue, green, and red but violet, blue-green, and yellow-green; and (3) if we were to stimulate cones ofjust one type, we would see not blue, green, or red but probably violet, green, and yellowish-red instead.

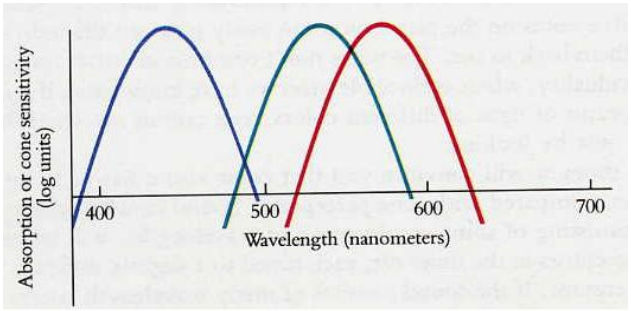

Absorption spectra (or sensitivity curves) differ for the three types of cones. (Spectral-energy curves and absorption curves such as these have their y axes in log units because they operate over such a wide range. The up-and-down position of the x-axis is therefore arbitrary and does not represent zero absorption.)

However unfortunate the terminology is, it is now widely used, and efforts to change embedded terminology usually fail. To substitute terms such as *long*, *middle*, and *short* would be more correct but would put a burden on those of us not thoroughly familiar with the spectrum. With peak absorption in the green, the rod pigment, rhodopsin, reflects blue and red and therefore looks purple. Because it is present in large enough amounts in our retinas that chemists can extract it and look at it, it long ago came to be called *visual purple*. Illogical as it is, "visual purple" is named for the appearance of the pigment,

whereas the terms for cones, "red", "green", and "blue", refer to their relative sensitivities or abilities to absorb light. Not to realize this can cause great confusion. The three cones show broad sensitivity curves with much overlap, especially the red and the green cones. Light at 600 nanometers will evoke the greatest response from red cones, those peaking at 560 nanometers, but will likely evoke some response, even if weaker, from the other two cone types.

Thus the red-sensitive cone does not respond *only* to long-wavelength, or red, light; it just responds better. The same holds for the other two cones. So far I have been dealing with physical concepts: the nature of light and pigments, the qualities of the pigments that reflect light to our eyes, and the qualities of the rod and cone pigments that translate the incoming light into electrical signals. It is the brain that interprets these initial signals as colors. In conveying some feel for the subject, I find it easiest to outline the elementary facts about color vision at the outset, leaving aside for the moment the three-century history of how these facts were established or how the brain handles color.

GENERAL COMMENTS ON COLOR

It may be useful to begin by *comparing* the way our auditory systems and our visual systems deal with wavelength. One system leads to tone and the other to color, but the two are profoundly different. When I play a chord of five notes on the piano, you can easily pick out the individual notes and sing them back to me. The notes don't combine in our brain but preserve their individuality, whereas since Newton we have known that if you mix two or more beams of light of different colors, you cannot say what the components are, just by looking. A little thought will convince you that color vision has to be an impoverished sense, compared with tone perception. Sound coming to one ear at any instant, consisting of some combination of wavelengths, will influence thousands of receptors in the inner ear, each tuned to a slightly different pitch than the next receptor. If the sound consists of many wavelength components, the information will affect many receptors, all of whose outputs are sent to our brains. The richness of auditory information comes from the brain's ability to analyze such combinations of sounds.

Vision is utterly different. Its information-handling capacity resides largely in the image's being captured by an array of millions receptors, at every instant. We take in the complex scene in a flash. If we wanted in addition to handle wavelength the way the ear does, the retina would need not only to have an array of receptors covering its surface, but to have, say, one thousand receptors for each point on the retina, each one with maximum sensitivity to a different wavelength. But to squeeze in a thousand receptors at each point is physically not possible. Instead, the retina compromises. At each of a very large number of points it has three different receptor types, with three different wavelength sensitivities. Thus with just a small sacrifice in resolution we end up with some rudimentary wavelength-handling ability over most of our retina. We see seven colors, not eighty-eight (both figures should be much higher!), but in a single scene each point of the many thousands will have a color assigned to it. The retina cannot have both the spatial capabilities that it has and also have the wavelength-handling capacity of the auditory system.

The next thing is to get some feel for what it means for our color vision to have three visual receptors. First, you might ask, if a given cone works better at some wavelengths

than at others, why not simply measure that cone's out-put and deduce what the color is? Why not have one cone type, instead of three? It is easy to see why. With one cone, say the red, you wouldn't be able to tell the difference between light at the most effective wavelength, about 560 nanometers, from a brighter light at a less effective wavelength. You need to be able to distinguish variations in brightness from variations in wavelength. But suppose you have two kinds of cones, with overlapping spectral sensitivities—say, the red cone and the green cone. Now you can determine wavelength simply by comparing the outputs of the cones. For short wavelengths, the green cone will fire better; at longer and longer wavelengths, the outputs ' will become closer and closer to equal; at about 580 nanometers the red surpasses the green, and does progressively better relative to it as wavelengths get still longer. If we subtract the sensitivity curves of the two cones (they are logarithmic curves, so we are really taking quotients), we get a curve that is independent of intensity. So the two cones together now constitute a device that measures wavelength.

Then why are not two receptors all we need to account for the color vision that we have? Two would indeed be enough if all we were concerned with was monochromatic light—if we were willing to give up such things as our ability to discriminate colored light from white light. Our vision is such that no monochromatic light, at any wavelength, looks white. That could not be true if we had only two cone types. In the case of red and green cones, by progressing from short to long wavelengths, we go continuously from stimulating just the green cone to stimulating just the red, through all possible green-to-red response ratios. White light, consisting as it does of a mixture of all wavelengths, has to stimulate the two cones in some ratio. Whatever monochromatic wavelength happens to give that same ratio will thus be indistinguishable from white. This is exactly the situation in a common kind of color blindness in which the person has only two kinds of cones: regardless of which one of the three pigments is missing there is always some wavelength of light that the person cannot distinguish from white. (Such subjects are color defective, but certainly not color-blind.)

To have color vision like ours, we need three and only three cone types. The conclusion that we indeed have just three cone types was first realized by examining the peculiarities of human color vision and then making a set of deductions that are a credit to the human intellect.

We are now in a better position to understand why the rods do not mediate color. At intermediate levels of light intensity, rods and cones can both be functioning, but except in rare and artificial circumstances the nervous system seems not to subtract rod influences from cone influences. The cones are compared with one another; the rods work alone. To satisfy yourself that rods do not mediate color, get up on a dark moonlit night and look around. Although you can see shapes fairly well, colors are completely absent. Given the simplicity of this experiment it is remarkable how few people realize that they do without color vision in dim light.

Whether we see an object as white or colored depends primarily (not entirely) on which of the three cone types are activated. Color is the consequence of unequal stimulation of the three types of cones. Light with a broad spectral curve, as from the sun or a candle, will obviously stimulate all three kinds of cones, perhaps about equally, and the resulting sensation turns out to be lack of color, or "white". If we could stimulate one kind of cone by itself (something that we cannot easily do with light because of the overlap of the

absorption curves), the result, as already mentioned, would be vivid color—violet, green, or red, depending on the cone stimulated. That the peak sensitivity of what we call the "red cone" is at a wavelength (560 nanometers) that appears to us greenish-yellow is probably because light at 560 nanometers excites both the green-sensitive cone and the red-sensitive cone, owing to the overlap in the green- and red-cone curves. By using longer wavelength light we can stimulate the red cone, relative to the green one, more effectively.

The graphs on this page sum up the color sensations that result when various combinations of cones are activated by light of various wavelength compositions. The first example and the last two should make it clear that the sensation "white"—the result of approximately equal stimulation of the three cones— can be brought about in many different ways: by using broad-band light or by using a mixture of narrow-band lights, such as yellow and blue or red and blue-green. Two beams of light are called *complementary* if their wavelength content and intensities are selected so that when mixed they produce the sensation "white". In the last two examples, blue and yellow are complementary, as are red at 640 nanometers and blue-green.

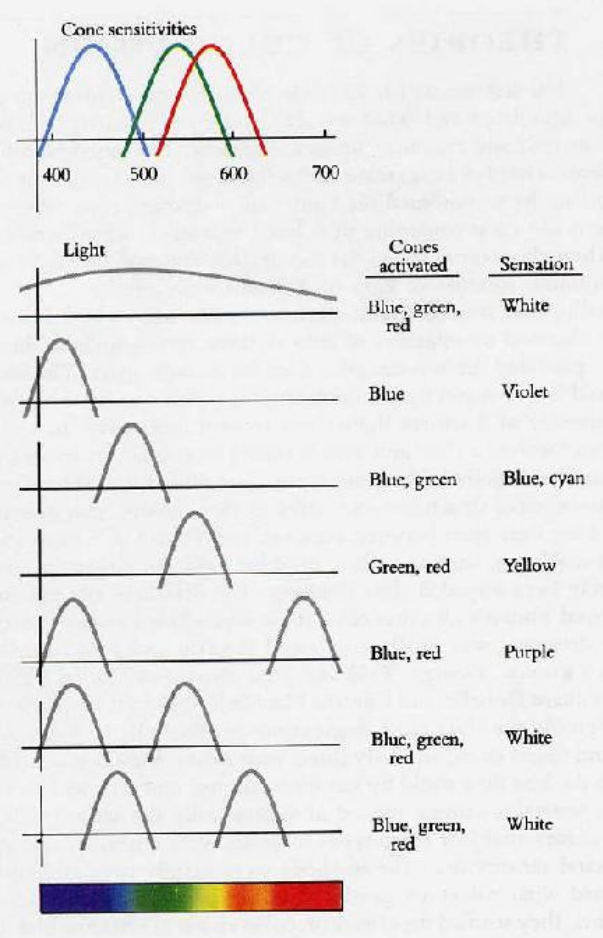

The top graph, "cone sensitivities", repeats the graph on page 5. The rest of the figure suggests which cones will be activated by various mixtures of colored light and what the resulting sensations will be.

THEORIES OF COLOR VISION

The statements I have made about the relationship between what cones are stimulated and what we see depend on research that began with Newton in 1704 and continues up to the present. The ingenuity of Newton's experiments is hard to exaggerate: in his work on color, he split up white light with a prism; he recombined the light with a second prism, obtaining white again; he made a top consisting of colored segments, which when spun gave white. These discoveries led to the recognition that ordinary light is made up of a continuous mixture of light of different wavelengths. Gradually, over the eighteenth century, it came to be realized that any color could be obtained by mixtures of light of three wavelengths in the right proportions, provided the wavelengths were far enough apart. The idea that any color could be produced by manipulating three controls (in this case, controls of the intensity of the three lights) was termed *trichromacy*. In 1802 Thomas Young put forward a clear and simple theory to explain trichromacy: he proposed that at each point in the retina there must exist at least three "particles"— tiny light-sensitive structures—sensitive to three colors, red, green, and violet. The long time span between Newton and Young is hard to explain, but various roadblocks, such as yellow and blue paints mixing to produce green, must surely have impeded clear thinking. The definitive experiments that finally proved Young's idea that color must depend on a retinal mosaic of three kinds of detectors was finally confirmed directly and conclusively in 1959, when two groups, George Wald and Paul Brown at Harvard and William Marks, William Dobelle, and Edward MacNichol at Johns Hopkins, examined microscopically the abilities of single cones to absorb light of different wavelengths and found three, and only three, cone types. Meanwhile, scientists had had to do the best they could by less direct means, and they had, in fact, in the course of several centuries arrived at substantially the same result, proving Young's theory that just three types of cones were necessary and estimating their spectral sensitivities. The methods were mainly psychophy steal: scientists learned what colors are produced with various mixtures of monochromatic lights, they studied the effects on color vision of selective bleaching with monochromatic lights, and they studied color blindness.

Studies of color mixing are fascinating, partly because the results are so surprising and counterintuitive. No one without prior knowledge would ever guess the various results shown in the illustration on page 9 in Chapter 3—for example, that two spots, one vivid blue and the other bright yellow, when overlapped would mix to produce a white indistinguishable to the eye from the color of chalk or that spectral green and red would combine to give a yellow almost indistinguishable from monochromatic yellow.

Before discussing other theories of color, I should perhaps say more about the variety of colors that theories must account for. What colors are there besides the colors in the rainbow? I can think of three. One kind is the purples, which we don't find in rainbows, but which result from stimulating the red and blue cones, that is, from adding long- and short-wavelength light, or, loosely, red and blue light. If to a mixture of spectral red and blue lights—to purple—we add the right amount of the appropriate green, we get white, and so we say that the green and purple are complementary. You can, if you like, imagine a circular dial that gives all the spectral colors from red through yellow and green to blue and violet, then purples, first bluish-purple and then reddish-purple, and finally back to red. You can even arrange these hues so that complements arc opposite each other.

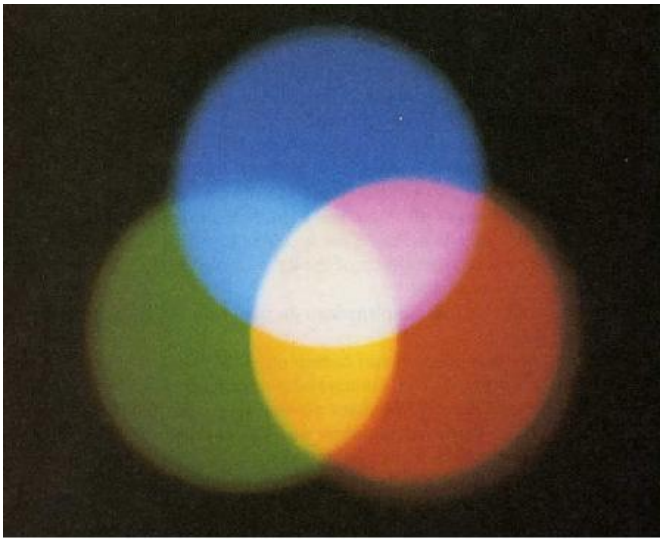

With three slide projectors and three filters, three overlapping spots (red, green, and blue) are projected onto a screen so that they overlap. Red and green give yellow, blue and green give turquoise, red and blue give purple, and all three—red, blue, and green—give white.

The concept of *primary colors* does not even enter this scheme: if we think of primaries in terms of the three receptor types, we have greenish-yellow, green, and violet, shades hardly consistent with the idea of three pure, basic colors. But if by primary we mean three colors from which any other hues can be generated, these three will do, as will any other three that are far enough apart. Thus nothing I have said so far gives any justification for the idea of three unique primary colors.

A second kind of color results from adding white to any spectral color or to purple; we say that the white "washes out" the color, or makes it paler—the technical term is that it *desaturates* it. To match any two colors, we have to make their hues and saturations the same (for example, by selecting the appropriate position on the circle of colors and then adding the right amount of white), and then we need to equate the intensities. Thus we can specify a color by giving the wavelength of the color (or in the case of purple, its complement), the relative content of white, and a single number specifying intensity. A mathematically equivalent option for specifying color is to give three numbers representing the relative effects of the light on the three cone types. Either way, it takes three numbers. A third kind of color these explanations do not cover is typified by brown. I will come to it later. Young's theory was adopted and championed by Hermann von Helmholtz and came to be known as the Young-Helmholtz theory. It was Helmholtz, incidentally, who finally explained the phenomenon mentioned at the beginning of this chapter, that mixing yellow and blue paints gives green. You can easily see how this differs from adding yellow and blue light by doing the following experiment, for which you need only two slide projectors and some yellow and blue cellophane. First, put the yellow cellophane over the lens of one projector and the blue over the other and then overlap the projected images. If you adjust the relative intensities, you will get a pure white in the area of overlap. This is the kind of color mixing we have been talking about, and we have said that the white arises because the combined yellow and blue light manages to activate all three of our cones with the same relative effectiveness that broad-

band, or white, light does. Now turn off one projector and put *both* filters in front of the other one, and you will get green. To understand what is happening we need to know that the blue cellophane absorbs long-wavelength light, the yellows and reds, from the white and lets through the rest, which looks blue, and that the yellow filter absorbs mainly blue and lets through the rest, which looks yellow. The diagram on this page shows the spectral composition of the light each filter passes. Note that in both cases the light that gets through is far from monochromatic, the yellow light is not narrow-band spectral yellow but a mixture of spectral yellow and shorter wavelengths, greens, and longer wavelengths, oranges and reds. Similarly, the blue is spectral blue plus greens and violet. Why don't we see more than just yellow or just blue? Yellow is the result of equal stimulation of the red and the green cones, with no stimulation of the blue cone; this stimulation can be accomplished with spectral yellow (monochromatic light at 580 nanometers) or with a broader smear of wavelengths, such as we typically get with pigments, as long as the breadth is not so great as to include short wavelengths and thereby stimulate the blue cone. Similarly, as far as our three cones are concerned, spectral blue light has about the same impact as blue plus green plus violet. Now, when we use the two filters, one in front of the other, what we get is what both filters let through, namely, just the greens. This is where the graphs shown on this page, for broad-band blue and yellow, overlap. The same thing happens with paints: yellow and blue paints together absorb everything in the light except greens, which are reflected. Note that if we used monochromatic yellow and blue filters in our experiment, putting one in front of the other would result in nothing getting through. The mixing works only because the colors produced by pigments have a broad spectral content. Why discuss this phenomenon here? I do so partly because it is gratifying to explain the dramatic and startling result of mixing yellow and blue paint to get green, and the even more startling result—because it is so unfamiliar to most people—of mixing yellow and blue light to get *white*. (In a chapter on color theory in a book on weaving, I found the statement that if you mix yellow and blue threads, as in warp and weft, you get green. What you do get is gray— for biological reasons.) The artificial results of mixing paints is doubtless what has led to the idea of "primary colors," such as red, yellow, and blue. If any special set of colors deserves to be called primary, it is the set of red, blue, yellow, and green. As we will see in the section on Hering's color theory, what justification all four have as candidates for primaries has little to do with the three cones and much to do with the subsequent wiring in the retina and brain.

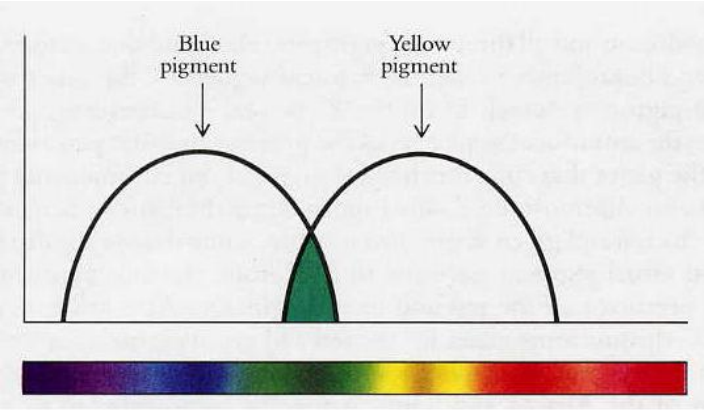

The blue filter passes a fairly broad band of wavelengths centered about 480 nanometers. The yellow filter passes a fairly broad band of wavelengths centered about 580 nanometers. Both together pass only wavelengths common to the two—light at a fairly broad band of wavelengths centered about 530, which gives a green.

THE GENETICS OF VISUAL PIGMENTS

In the early 1980s Jeremy Nathans, while still an MD-Ph D student at Stanford, managed to clone the genes for the protein portions of human rhodopsin and all three cone pigments. He found that all four pigments show strong homologies in their amino acid sequences: the genes for the red and green pigments, which lie on the X, or sex, chromosome, are virtually identical—the amino acid sequences of the proteins show 96 percent identity— whereas the genes that code for the blue pigment, on chromosome 7, and for rhodopsin, on chromosome 3, show much larger differences, from each other and from the red and green genes. Presumably, some time in the distant past, a primordial visual pigment gave rise to rhodopsin, the blue pigment, and the common precursor of the red and green pigments. At a much more recent time the X-chromosome genes for the red and green pigments arose from this precursor by a process of duplication. Possibly this occurred after the time of separation of the African and South American continents, 30 to 40 million years ago, since old world primates all exhibit this duplication of cone pigment genes on the X-chromosome, whereas new world primates do not. Cloning the genes has led to a spectacular improvement in our understanding of the various forms of color blindness. It had long been known that most forms of color-vision deficiency are caused by the absence or abnormality of one or more of the three cone pigments. The most frequent abnormalities occur in the red and green pigments and affect about 8 percent of males. Because of the wide range of these abnormalities the subject is complex, but given our molecular-level understanding, it is fortunately no longer bewildering.

Very rarely, destruction to certain cortical areas can cause color blindness. Most often this occurs as the result of a stroke.

THE HERING THEORY

In the second half of the nineteenth century, a second school of thought arose parallel to, but until recently seemingly irreconcilable with, the Young-Helmholtz color theory. Ewald Hering (1834-1918) interpreted the results of color mixing by proposing the existence, in the eye, brain, or both, of three opponent processes, one for red-green sensation, one for yellow-blue, and a third, qualitatively different from the first two, for black-white. Hering was impressed by the nonexistence of any colors—and the impossibility of even imagining any colors—that could be described as yellowish-blue or red-dish-green and by the apparent mutual canceling of blue and yellow or of red and green when they are added together in the right proportions, with complete elimination of hue—that is, with the production of white. Hering envisioned the red-green and yellow-

blue processes as independent, in that blue and red do add to give bluish-red, or purple; similarly red added to yellow gives orange; green added to blue gives bluish-green; green and yellow gives greenish-yellow. In Hering's system, yellow, blue, red, and green could be thought of as "primary" colors. Anyone looking at orange can imagine it to be the result of mixing red and yellow, but no one looking at red or blue would see it as the result of mixing any other colors. (The feeling that some people have that green looks like yellow added to blue is probably related to their childhood experience with paint boxes.) Hering's notions of red-green and yellow-blue processes seemed to many to be disconcertingly dependent on intuitive impressions about color, but it is surprising how good the agreement is among people asked to select the point on the spectrum where blue is represented, untainted by any subjective trace of green or yellow. The same is so for yellow and green. With red, subjects again agree, this time insisting that some violet by added to counteract the slight yellowish appearance of long-wavelength light. (It is this subjective red that when added to green gives white; ordinary (spectral) red added to green gives yellow.)

We can think of Hering's yellow-blue and red-green processes as separate channels in the nervous system, whose outputs can be represented as two meters, like old-fashioned voltmeters, with the indicator of one meter swinging to the left of zero to register yellow and to the right to register blue and the other meter doing the same for red versus green. The color of an object can then be described in terms of the two readings. Hering's third antagonistic process (you can think of it as a third voltmeter) registered black versus white. He realized that black and gray are not produced simply by absence of light coming from an object or surface but arise when and only when the light from the object is less than the average of the light coming from the surrounding regions. White arises only when the surround is darker and when no hue is present. (I have already discussed this in Chapter 3, with examples such as the turned-off television set.) In Hering's theory, the black-white process requires a spatial comparison, or subtraction of reflectances, whereas his yellow-blue and red-green processes represent something occurring in one particular place t in the visual field, without regard to the surrounds. (Hering was certainly aware that neighboring colors interact, but his color theory as enunciated in his latest work does not encompass those phenomena.) We have already seen that black versus white is indeed represented in the retina and brain by spatially opponent excitatory and inhibitory (on versus off) processes that are literally antagonistic.

Hering's theory could account not only for all hues and levels of saturation, but also for colors such as brown and olive green, which are not only absent I from any rainbow, but cannot be produced in any of the classical psychophysical color-mixing procedures, in which we shine spots of light on a dark screen with a slide projector. We get brown only when a yellow or orange spot of ' light is surrounded by light that on the average is brighter. Take any brown and exclude all the surround by looking at it through a tube, a black piece of rolled up paper, and you will see yellow or orange. We can regard brown as a mixture of black—which is obtainable only by spatial contrasts—and orange or yellow. In Hering's terms, at least two of the systems are at work, the black-white and the yellow-blue.

Hering's theory of three opponent systems, for red-green, yellow-blue, and black-white, was regarded in his time and for the next half-century as rivaling and contradicting the Young-Helmholtz three-pigment (red, green, and blue) theory: the proponents of each

were usually strongly partisan and often emotional. Physicists generally sided with the Young-Helmholtz camp, perhaps because their hearts were warmed by the quantitative arguments—by such things as linear simultaneous equations—and turned or cooled off by arguments about purity of colors. Psychologists often sided with Hering, perhaps because they had experience with a wider variety of psychophysical phenomena. Hering's theory seemed to be arguing for either four receptor types (red, green, yellow, and blue) or worse, for three (one subserving black-white, one yellow-blue, and one red-green), all in the face of mounting evidence for the original Young hypothesis. In retrospect, as the contemporary psychophysicists Leo Hurvich and Dorothea Jameson have pointed out, it seems that one difficulty many people had with the Hering theory was the lack, until the 1950s, of any direct physiological evidence for inhibitory mechanisms in sensory systems. Such evidence became available only a half-century later, with single-unit recordings.

By imagining the voltmeters to be measuring positive to the right and negative to the left, you can see why Hering's work suggested inhibitory mechanisms. In a sense, the colors yellow and blue are mutually antagonistic; together they cancel each other, and if the red-green system also reads zero, we have no color. Hering in some ways was fifty years ahead of his time. As has happened before in science, two theories, for decades seemingly irreconcilable, both turned out to be correct. In the late 1800s, nobody could have guessed that at one and the same time the Young-Helmholtz notions of color would turn out to be correct at the receptor level, whereas Hering's ideas of opponent processes would be correct for subsequent stages in the visual path. It is now clear that the two formulations are not mutually exclusive: both propose a three-variable system: the three cones in the Young-Helmholtz and the three meters, or processes, in the Hering theory. What amazes us today is that with so little to go on, Hering's formulation turned out to describe cell-level central-nervous-system color mechanisms so well. Nevertheless, color-vision experts are still polarized into those who feel Hering was a prophet and those who feel that his theories represent a fortuitous fluke. To the extent that I am slightly to the Hering side of center, I will doubtless make enemies of all the experts.

COLOR AND THE SPATIAL VARIABLE

We saw in Chapter 3 that an object's whiteness, blackness, or grayness depends on the light that the object reflects from some source, relative to the light reflected by the other objects in the scene, and that broad-band cells at an early stage in the visual pathway—retinal ganglion cells or geniculate cells—can go far to account for this perception of black and white and gray: they make just this kind of comparison with their center-surround receptive fields.

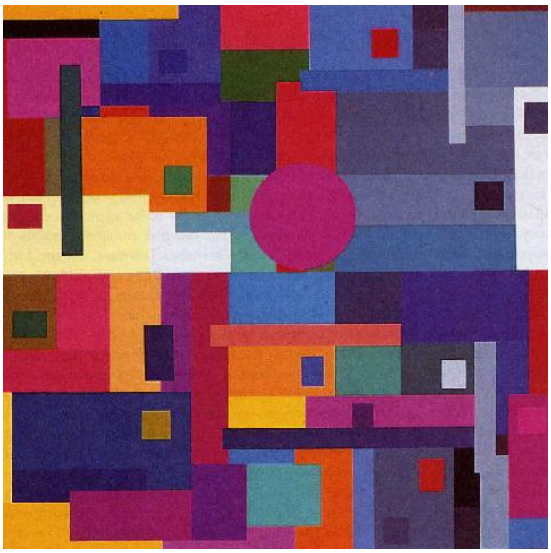

In many of his experiments Edwin Land used a Mondrian-like patchwork of colored papers. The experiments were designed to prove that the colors remain remarkably constant despite marked variations in the relative intensities of the red, green, and blue lights used to illuminate the display.

This is surely Hering's third, spatially opponent black-white process. That the spatial variable is also important for color first began to be appreciated a century ago. It was tackled analytically only in the last few decades, notably by psychophysicists such as Leo Hurvich and Dorothea Jameson, Deane Judd, and Edwin Land. Land, with a consuming interest in light and photography, was naturally impressed by a camera's failure to compensate for differences in light sources. If a film is balanced so that a picture of a white shirt looks white in tungsten light, the same shirt under a blue sky will be light blue; if the film is manufactured to work in natural light, the shirt under tungsten light will be pink. To take a good color picture we have to take into account not only light intensity, but also the spectral content of the light source, whether it is bluish or reddish. If we have that information, we can set the shutter speed and the lens opening to take care of the intensity and select the film or filters to take care of color balance. Unlike the camera, our visual system does all this automatically, and it solves the problem so well that we are generally not aware that a problem exists.

A white shirt thus continues to look white in spite of large shifts in the spectral content of the light source, as in going from overhead sun to setting sun, to tungsten light, or to fluorescent light. The same constancy holds for colored objects, and the phenomenon, as applied to color and white, is called color constancy. Even though color constancy had been recognized for many years, Land's demonstrations in the 1950s came as a great surprise, even to neurophysiologists, physicists, and most psychologists.

What were these demonstrations? In a typical experiment, a patchwork of rectangular papers of various colors resembling a Mondrian painting is illuminated with three slide projectors, one equipped with a red, the second with a green, the third with a blue filter. Each projector is powered by a variable electric source so that its light can be adjusted over a wide range of intensities. The rest of the room must be completely dark. With all three projectors set at moderate intensities, the colors look much as they do in daylight. The surprising thing is that the exact settings do not seem to matter. Suppose we select a

green patch and with a photometer precisely measure the intensity of the light coming from that patch when only one projector is turned on. We then repeat the measurement, first with the second projector and then with the third. That gives us three numbers, representing the light coming to us when we turn on all three projectors. Now we select a different patch, say orange, and readjust each projector's intensity in turn so that the readings we now get from the orange patch are the same as those we got before from the green one. Thus with the three projectors turned on, the composition of light now coming from the orange patch is identical to the composition of light that a moment ago came from the green. What do we expect to see? Naively, we would expect the orange patch to look green. But it still looks orange—indeed, its color has not changed at all. We can repeat this experiment with any two patches. The conclusion is that it doesn't much matter at what intensities the three projectors are set, as long as some light comes from each. In a vivid example of color constancy, we see that twisting the intensity dials for the three projectors to almost any position makes very little difference in the colors of the patches.

Such experiments showed convincingly that the sensation produced in one part of the visual field depends on the light coming from that place and on the light coming from everywhere else in the visual field. Otherwise, how could the same light composition give rise at one time to green and at another to orange? The principle that applies in the domain of black, white, and gray, stated so clearly by Hering, thus applies to color as well. For color, we have an opponency not only locally, in red versus green and yellow versus blue, but also spatially: center red-greenness versus surround red-greenness, and the same opponency for yellow-blueness.

In 1985, in Land's laboratory, David Ingle managed to train goldfish to swim to a patch of some preassigned color in an underwater Mondrian display. He found that a fish goes to the same color, say blue, regardless of wavelength content: it selects a blue patch, as we do, even when the light from it is identical in composition to the light that, in a previous trial and under a different light source, came from a yellow patch, which the fish had rejected. Thus the fish, too, selects the patch for its color, not for the wavelength content of the light it reflects. This means that the phenomenon of color constancy cannot be regarded as some kind of embellishment recently added by evolution to the color sense of certain higher mammals like ourselves; finding it in a fish suggests that it is a primitive, very basic aspect of color vision. It would be fascinating (and fairly easy) to test and see whether insects with color vision also have the same capability. I would guess that they do.

Land and his group (among others, John McCann, Nigel Daw, Michael Burns, and Hollis Perry) have developed several procedures for predicting the color of an object, given the spectral-energy content of light from each point in the visual field but given no information on the light source. The computation amounts to taking, for each of the three separate projectors, the ratio of the light coming from the spot whose color is to be predicted to the average light coming from the surround. (How much surround should be included has varied in different versions of the theory: in Land's most recent version, the surround effects are assumed to fall off with distance.) The resulting triplet of three numbers—the ratios taken with each projector—uniquely defines the color at that spot. Any color can thus be thought of as corresponding to a point in three-dimensional space whose coordinate axes are the three ratios, taken with red light, green light, and blue

light. To make the formulation as realistic as possible, the three lights are chosen to match the spectral sensitivities of the three human cone types.

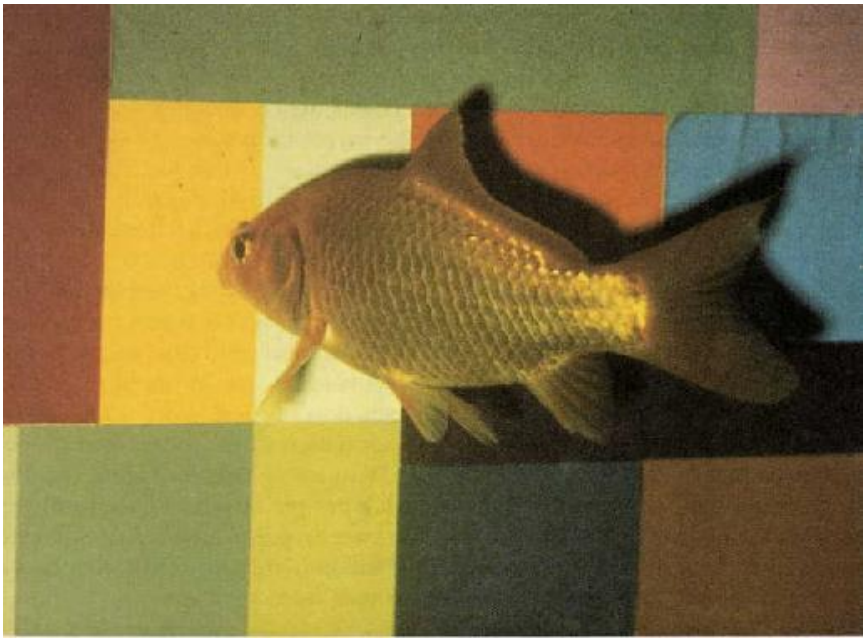

In David Ingle's experiment, a goldfish has been trained to swim to a patch of a given color for a reward—a piece of liver. It swims to the green patch regardless of the exact setting of the three projectors' intensities. The behavior is strikingly similar to the perceptual result in humans.

That color can be so computed predicts color constancy because what counts for each projector are the *ratios* of light from one region to light from the average surround. The exact intensity settings of the projectors no longer matter: the only stipulation is that we have to have some light from each projector; otherwise no ratio can be taken. One consequence of all this is that to have color at all, we need to have variation in the wavelength content of light across the visual field. We require color borders for color, just as we require luminance borders for black and white. You can easily satisfy yourself that this is true, again using two slide projectors. With a red filter (red cellophane works well) in front of one of the projectors, illuminate any set of objects. My favorite is a white or yellow shirt and a bright red tie. When so lit, neither the shirt nor the tie looks convincingly red: both look pinkish and washed out. Now you illuminate the same combination with the second projector, which is covered with blue cellophane. The shirt looks a washed-out, sickly blue, and the tie looks black: it's a red tie, and red objects don't reflect short wavelengths.

Go back to the red projector, confirming that with it alone, the tie doesn't look especially red. Now add in the blue one. You know that in adding the blue light, you will not get anything more back from the tie—you have just demonstrated that—but when you turn on the blue projector, the red tie suddenly blazes forth with a good bright red. This will convince you that what makes the tie red is not just the light coming to you from the tie. Experiments with stabilized color borders are consistent with the idea that differences across borders are necessary for color to be seen at all. Alfred Yarbus, whose name came

up in the context of eye movements in Chapter 4, showed in 1962 that if you look at a blue patch surrounded by a red background, stabilizing the border of the patch on the retina will cause it to disappear: the blue melts away, and all you see is the red background. Stabilizing the borders on the retina apparently renders them ineffective, and without them, we have no color.

These psychophysical demonstrations that differences in the spectral content of light across the visual field are necessary to perceive color suggest that in our retinas or brains we should find cells sensitive to such borders. The argument is similar to the one we made in Chapter 4, about the perception of black or white objects (such as kidney beans). If at some stage in our visual path color is signaled entirely at color-contrast borders, cells whose receptive fields are entirely within areas of uniform color will be idle. The result is economy in handling the information. We thus find ourselves with two advantages to having color signaled at borders: first, color is unchanged despite changes in the light source, so that our vision tells us about properties of the objects we view, uncontaminated by information about the light source; second, the information handling is economical.

Now we can ask why the system evolved the way it did. Are we to argue that the need for color constancy led to the system's evolving and that an unexpected bonus was the economy—or the reverse, that economy was paramount and the color constancy a bonus? Some would argue that the economy argument is more compelling: evolution can hardly have anticipated tungsten or fluorescent lights, and until the advent of supersuds, our shirts were not all that white anyway.

THE PHYSIOLOGY OF COLOR VISION: EARLY RESULTS

The first cell-level physiological information came 250 years after Newton from the studies of the Swedish-Finnish-Venezuelan physiologist Gunnar Svaetichin, who in 1956 recorded intracellularly in teleost fish from what he thought were cones but turned out later to be horizontal cells. These cells responded with slow potentials only (no action potentials) when light was directed on the retina. He found three types of cells, as illustrated on the next page: the first, which he called L cells, were hyperpolarized by light stimulation regardless of the light's wavelength composition; the second, called r-g cells, were hyperpolarized by short wavelengths, with a maximum response to green light, and depolarized by long wavelengths, with a maximum response to red; the third, which with Hering in mind he called y-b cells, responded like r-g cells but with maximal hyperpolarization to blue and maximal depolarization to yellow. For r-g and y-b cells, white light gave only weak and transient responses, as would be expected from white's broad spectral energy content. Moreover, for both types of cell, which we can call *opponent-color cells*, some intermediate wavelength of light, the *crossover point*, failed to evoke a response. Because these cells react to colored light but not to white light, they are probably concerned with the sensation of color.

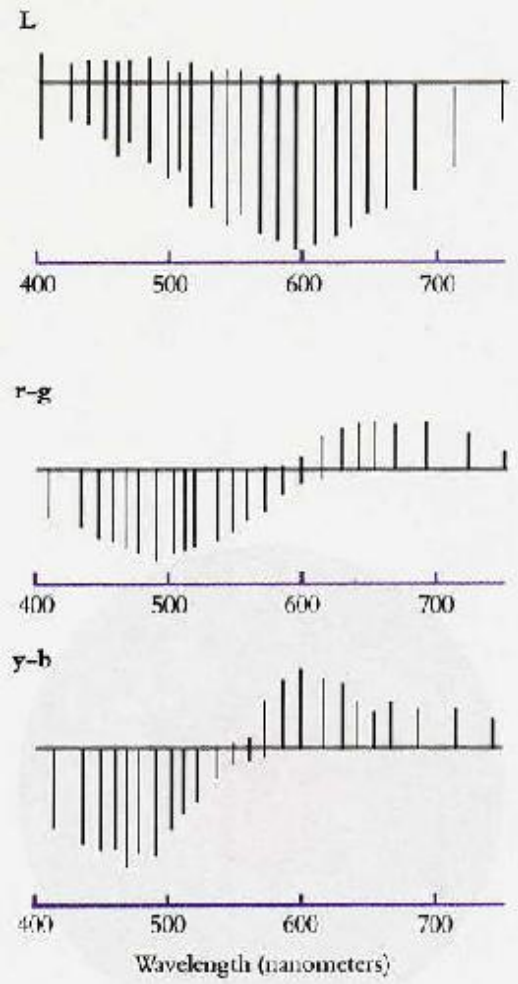

Gunnar Svactichin and Edward MacNichol recorded the responses to color of horizontal cells in the teleost fish. Deflections pointing downward from the gray line indicate hypcrpolarization; those pointing upward indicate depolarization.

In 1958, Russell De Valois (rhymes with *hoi polloi*) and his colleagues recorded responses strikingly similar to Svaetichin's from cells in the lateral geniculate body of macaque monkeys. De Valois had previously shown by behavioral testing that color vision in macaque monkeys is almost identical to color vision in humans; for example, the amounts of two colored lights that have to be combined to match a third light are almost identical in the two species. It is therefore likely that macaques and humans have similar machinery in the early stages of their visual pathways, and we are probably justified in comparing human color psychophysics with macaque physiology. De Valois found many geniculate cells that were activated by diffuse monochromatic light at wavelengths ranging from one end of the spectrum to a crossover point, where there was no response, and were suppressed by light over a second range of wavelengths from the crossover point to the other end. Again the analogy to Hering's color processes was compelling: De Valois tound op-ponent-color cells of two types, red-green and yellow-blue; for each type, combining two lights whose wavelengths were on opposite sides of some crossover point led to mutual cancellation of responses, just as, perceptually,

adding blue to yellow or adding green to red produced white. De Valois' results were especially reminiscent of Hering's formulations in that his two classes of color cells had response maxima and crossover points in just the appropriate places along the spectrum for one group to be judging the yellow-blueness of the light and the other, red-greenness. The next step was to look at the receptive fields of these cells by using small spots of colored light, as Torsten Wiesel and I did in 1966, instead of diffuse light. For most of De Valois' opponent-color cells, the receptive fields had a surprising organization, one that still puzzles us. The cells, like Kufflcr's in the cat, had fields divided into antagonistic centers and surrounds; the center could be "on" or "off". In a typical example, the field center is fed exclusively by red cones and the inhibitory surround exclusively by green cones. Consequently, with red light both a small spot and a large spot give brisk responses, because the center is selectively sensitive to long-wavelength light and the surround virtually insensitive; with short-wavelength light, small spots give little or no response and large spots produce strong inhibition with off responses. With white light, containing short and long wavelengths, small spots evoke on responses and large spots produce no response.

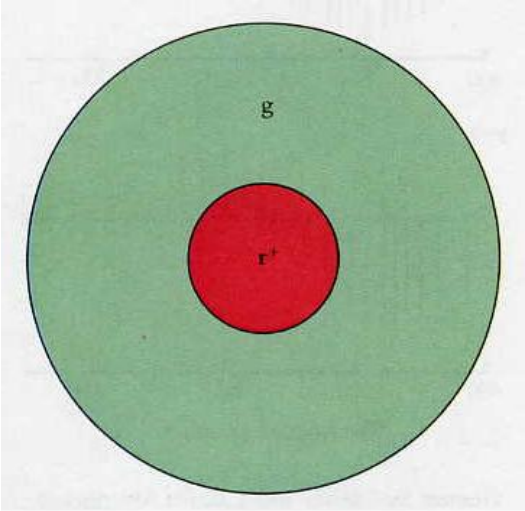

In a typical type *I* receptive field, the center receives excitatory input from red cones; the surround, inhibitory input from green cones.

Although our first impression was that such a cell must be getting input from red cones in the center region and green cones in the surround, it now seems probable that the total receptive field is a combination of two overlapping processes, as illustrated in the figure on the next page. Both the red cones and the green cones feed in from a fairly wide circular area, in numbers that are maximal in the center and fall off with distance from the center. In the center, the red cones strongly predominate, and with distance their effects fall off much more rapidly than those of the green cones. A long-wavelength small spot shining in the center will consequently be a very powerful stimulus to the red system; even if it also stimulates green cones, the number, relative to the total number of green cones feeding in, will be too small to give the red system any competition. different-shaped sensitivity-versus-position curves.

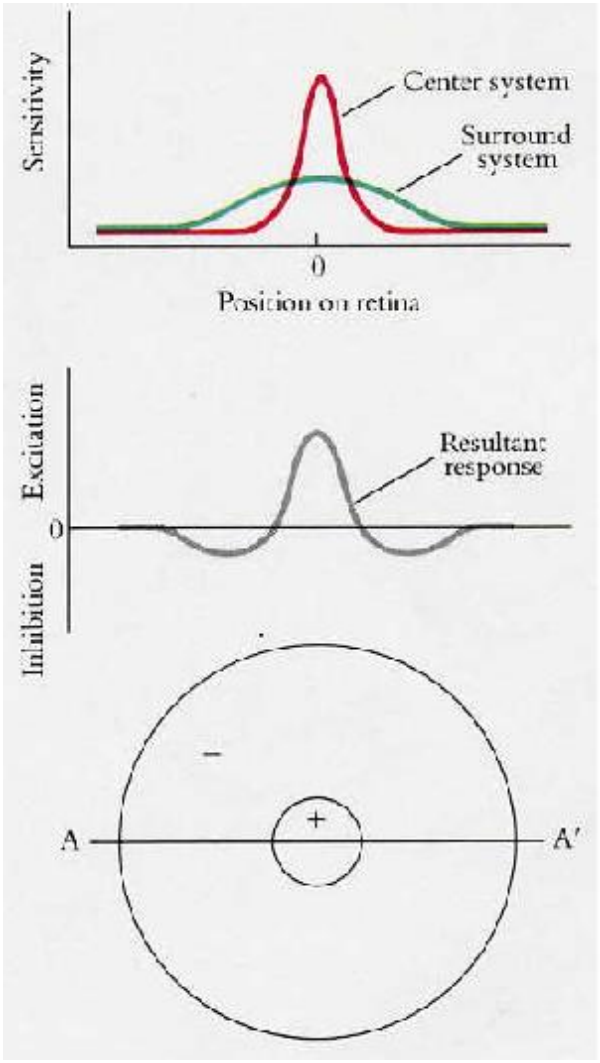

These graphs plot the sensitivity of a cell (measured, for example, by the response to a constant very small spot of light) against retinal position along a line AA' passing through the receptive-field center. For an r+ center-g-surround cell, a small red spot gives a narrow curve and a small green spot, a much broader one. The lower graph plots the responses to light such as white or yellow that stimulates both of the opponent systems, so that the two systems subtract. Thus the red cones dominate in the center, which gives on responses, whereas the green cones dominate in the surround, which yields off responses.

The same argument applies to the center-surround cells described in Chapter 3, whose receptive fields similarly must consist of two opponent circular overlapping areas having Thus the surround is probably not annular, or donut shaped, as was originally supposed, but filled. With these opponent-color cells in monkeys, it is supposed—without evidence so far—that the surrounds represent the contributions of horizontal cells.
The responses to diffuse light—in this case, on to red, off to blue or green, and no response to white—make it clear that such a cell must be registering information about color. But the responses to appropriate white borders and the lack of response to diffuse light make it clear that the cell is also concerned with black-and-white shapes. We call

these center-surround color-opponent cells "type 1".
The lateral geniculate body of the monkey, we recall from Chapter 4, consists of six layers, the upper four heavily populated with small cells and the lower two sparsely populated with large cells. We find cells of the type just described in the upper, or parvocellular, layers. Type 1 cells differ one from the next in the types of cone that feed the center and surround systems and in the nature of the center, whether it is excitatory or inhibitory. We can designate the example in the diagram on page 20 as "r+g-". Of the possible subtypes of cells that receive input from these two cone types, we find all four: r+g-, r-g+, g+r-, g-r+. A second group of cells receives input from the blue cone, supplying the center, and from a combination of red and green cones (or perhaps just the green cone), supplying the surround. We call these "blue-yellow", with "yellow" a shorthand way of saying "red-plus-green".
We find two other types of cells in the four dorsal layers. Type 2 cells make up about 10 percent of the population and have receptive fields consisting of a center only. Throughout this center, we find red-green opponency in some gives little or no response, because one effect cancels the other—the reverse of what would seem to be required for color contrast.
What we can say of type i cells is that they are likely to play an important part in high-precision form perception, given their tiny field centers and their responsiveness to black-and-white contours. As we saw in Chapter 6, we have several ways to measure visual acuity, the ability of our visual system to discriminate small objects; these include the smallest separation between two dots that can just be discriminated and the smallest detectable gap in a circle (called the Landolt C). Acuity measured in either of these ways turns out, for the fovea, to be about 0.5 minute of arc, or about i millimeter at a distance of 8 meters. This corresponds well with the distance between two cones in the fovea. Type 1 geniculate cells that get their input from near the fovea have receptive-field centers as small as about 2 minutes of arc in diameter. It seems likely that in the fovea one cone only contributes to a field center. So we find a reasonable fit between acuity and smallest field-center sizes of lateral geniculate cells. The ventral pair of geniculate layers differs from the dorsal four in being made up entirely of cells whose field centers are broad-band. The cells do show a curious form of color opponency that no one understands and that I will say no more about. Most people assume that these cells are color-blind. Their field centers are several times larger than centers of parvocellular cells, and they differ in several other interesting ways. We presently suspect that these cells feed into parts of the brain that subserve depth and movement perception. To elaborate further would take us far from color and require another book. Most of the cells I have been describing for the lateral geniculate have also been observed in the retina. They are more conveniently segregated in the geniculate and are easier to study there. We do not know what the geniculate contributes to the analysis of visual information in the monkey, besides its obvious function of handing on to the cortex the information it receives from the eyes.

THE NEURAL BASIS OF COLOR CONSTANCY

Since type 1 cells in the lateral geniculate body seem not to be geared to make color-spatial comparisons, we probably have to look beyond the retina and geniculate. To test the idea that such computations might go on in the cortex. Land's group and Margaret Livingstone and I examined a man who had had his corpus callosum severed surgically to treat epilepsy. Spatial-color interactions did not take place across the visual-field midline, that is, the color of a spot just to the left of the point at which the subject was looking was not affected by drastic changes in the colors in the right visual field, whereas normal subjects observed marked differences with such changes. This suggests that the retina by itself cannot mediate the color-spatial interactions. Although no one had seriously claimed that it could, the question continued to be debated, and it was satisfying to have some experimental evidence. The experimental results are consistent with our failure to find retinal ganglion cells that could plausibly be involved in color-spatial interactions. The goldfish, which makes spatial comparisons very much like ours, has virtually no cerebral cortex. Perhaps the fish, unlike us, does make such computations with its retina. Nigel Daws' discovery in 1968 of *double opponent cells* in the fish retina seems to bear this out. In the monkey, as I describe in the next section, we find such cells in the cortex but not in the lateral geniculate or the retina.

BLOBS

By about 1978, the monkey's primary visual cortex, with its simple, complex, and end-stopped cells and its ocular-dominance columns and orientation columns, seemed reasonably well understood. But an unexpected feature of the physiology was that so few of the cells seemed to be interested in color. If we mapped a simple or complex cell's receptive field using white light and then repeated the mapping with colored spots or slits, the results as a rule were the same. A few cells, perhaps as many as a 10 percent of cortical upper-layer cells, did show unmistakable color preferences—with excellent responses to oriented slits of some color, most often red, and virtually no response to other wavelengths or even to white light. The orientation selectivity of these cells was just as high as that of cells lacking color selectivity. But most cells in the visual cortex did not care about color. This was all the more surprising because such a high proportion of cells in the lateral geniculate body are color coded, and the geniculate forms the main input to the visual cortex. It was hard to see what could have happened to this color information in the cortex.

Suddenly, in 1978, all this changed. Margaret Wong-Riley, at the University of California in San Francisco, discovered that when she stained the cortex for the enzyme cytochrome oxidase, the upper layers showed an unheard of inhomogeneity, with periodic dark-staining regions, pufflike in transverse cross section, about one-quarter millimeter wide and one-half millimeter apart. All cells contain cytochrome oxidase, an enzyme involved in metabolism, and no one had ever imagined that a stain for such an enzyme would show anything interesting in the cortex. When Wong-Riley sent us pictures, Torsten Wiesel and I suspected that we were seeing ocular-dominance slabs cut in cross section and that the most monocular cells were for some reason metabolically more active than binocular cells. We put the pictures in a drawer and tried to forget them.

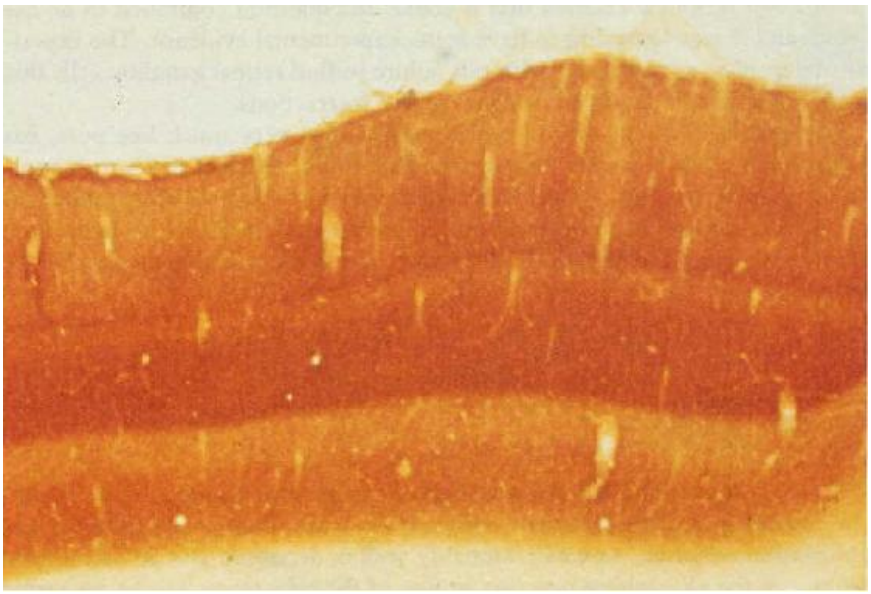

This cross section through the striate cortex shows the layers stained for the enzyme cytochrome oxidase. The darker zones in the upper third of the section are the blobs.

Several years elapsed before it occurred to us or anyone else to examine the primary visual cortex with this stain in sections cut parallel to the surface. When that was finally done, roughly simultaneously by two groups (Anita Hendrickson and Alan Humphrey in Seattle and Jonathan Horton and myself in Boston), a polka-dot pattern appeared—to everyone's complete surprise. An example is shown in the photograph on this page. Instead of stripes, we saw an array of bloblike structures for which no known correlates existed. Wong-Riley's inhomogeneities have been called by almost every imaginable name: dots, puffs, patches, and spots. We call them "blobs" because the word is graphic, legitimate (appearing even in the Oxford English Dictionary), and seems to annoy our competitors.

The next task was obvious: we had to record from the striate cortex again, monitoring the experiments histologically with the cytochrome-oxidase stain to see if we could find anything different about the cells in the blobs. Margaret Livingstone and I set out to do this in 1981. The result was quite unexpected. In traversing a distance of one-quarter millimeter, the diameter of a blob, it is possible to record from roughly five or six cells. Each time we crossed a blob, the cells we saw completely lacked orientation selectivity, in marked contrast to the high orientation selectivity shown by the cells outside the blobs.

One might explain this absence of orientation selectivity in either of two ways. First, these cells might receive their input unselectively from oriented cells in the nonblob neighborhood and consequently still respond specifically to lines (slits and so forth)—but by pooling all the various orientations, still end up with no preference. Second, they could resemble geniculate cells or cells in layer 4C and thus be simpler than the nonblob orientation-selective cells. The question was quickly settled: the cells were mostly center-surround. A few more experiments were enough to convince us that many of them were color coded.

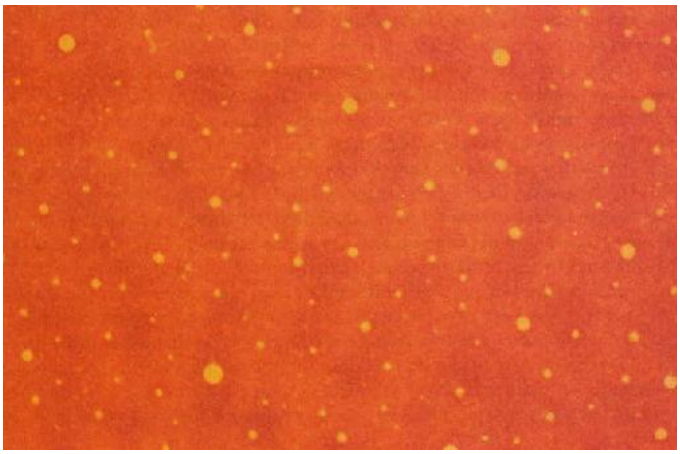

The dark areas are blobs seen tace on, about 50 of which form a polka-dot pattern. This section, through layer 3 of area 17, is parallel to the cortical surface and about 0.5 millimeter beneath it. (The yellow circles are blood vessels cut transversely.)

Over half the blob cells had opponent-color, center-surround receptive fields, but they behaved in a decidedly more complicated way than type 1 cells in the lateral geniculate body.in the lateral geniculate body. They gave virtually no responses to white spots of any size or shape. But to small colored spots shone in the center of the receptive field they responded vigorously over one range of wavelengths and were suppressed over another range: some were activated by long wavelengths (reds) and suppressed by short (greens and blues): others behaved in the reverse way. As with geniculate cells, we could distinguish two classes, red-green and blue-yellow, according to the position of the maximum responses. (Here, as before, red, green, and blue stand for the respective cone types, and yellow implies an input from the red and green in parallel.) So far, then, these cells closely resembled opponent-color, center-only geniculate cells (type 2). Their field centers, like centers of type 2 cells, were large—several times the size of type 1cell centers. They were unresponsive to small shone anywhere in their receptive fields. Most surprising was the finding that these color-coded blob cells, unlike type 2 cells, were mostly unresponsive to large colored spots, regardless of wavelength content. They behaved as though each center system was surrounded by a ring ofopponency. To take the commonest type, the r+g- center seemed to be surrounded by a ring that was r-g+. Margaret Livingstone and I have called these cells *double-opponent* because of the red-green or yellow-blue opponency in the center and the antagonism of the surround to any center response, whether "on" or "off". They are therefore unresponsive not only to white light in any geometric form but also to large spots, regardless of wavelength content. As already mentioned, Nigel Daw coined the term *double-opponent* for cells he saw in the retina of the goldfish. Daw suggested that cells like these might be involved in color-spatial interactions in man, and a few years later, with Alan Pearlman, he searched carefully in the macaque monkey lateral geniculate for such cells, without success. From the late 1960s on, double-opponent cells had occasionally been observed in the monkey cortex, but they were not clearly associated with any anatomical structure. We still do not understand some things about these cells. For example, in the r+g- just described, a red spot surrounded by green often gives a poor response, not the vigorous

one we might expect. Mixed with the two classes of double-opponent cells (red-green and yellow-blue) were ordinary broad-band, center-surround cells. Again, these broadband cells differed from cells in the upper geniculate layers and from cells in 4C Bata in having several times larger center sizes. Blobs also contain cells that are indistinguishable from geniculate type 2 cells, resembling double-opponent cells but lacking the receptive-field surround.

Margaret Livingstone and I have proposed that the blobs represent a branch of the visual pathway that is devoted to "color", using the word broadly to include blacks, whites, and grays. This system seems to separate off from the rest of the visual path either in the lateral geniculate body or in layer 4 of the striate cortex. (The geniculate probably projects directly but weakly to the blobs. It seems likely that layer 4C beta also projects to them, and it may well form their main input. Whether 4C alpha: projects to them is not clear.) Most blob cells seem to require border contrast in order to give responses at all: either luminous-intensity borders, in the case of the broad-band, center-surround cells or color-contrast borders, in the case of the double-opponent cells would respond. As I argued earlier, this amounts to saying that these cells play a part in color constancy.

If blob cells are involved in color constancy, they cannot be carrying out the computation exactly as Land first envisioned it, by making a separate comparison between a region and its surround for each of the cone wavebands. Instead ^ they would seem to be doing a Hering-like comparison: of red-greenness in one region with red-greenness in the surround, and the same for yellow-blueness and for intensity. But the two ways of handling color—r, g, and b on the one hand and b-w, r-g, and y-b on the other—are really equivalent. Color requires our specifying three variables; to any color there corresponds a triplet of numbers, and we can think of any color as occupying a point in three-dimensional space. We can plot points in such a space in more than one way. The coordinate system can be Cartesian, with the three axes orthogonal and oriented in any direction or we can use polar or cylindrical coordinates. The Hering theory (and apparently the retina and brain) simply employ a different set of axes to plot the same space. This is doubtless an oversimplification because the blob cells making up the three classes are not like peas in pods but vary among themselves in the relative strengths of surrounds and centers, in their perfections in the balance between opponent colors, and in other characteristics, some still not understood. At the moment, we can only say that the physiology has a striking affinity with the psychophysics. You may ask why the brain should go to the trouble to plot color with these seemingly weird axes rather than with the more straightforward r, g, and b axes, the way the receptor layer of the retina does. Presumably, color vision was added in evolution to the colorless vision characteristic of lower mammals. For such animals, color space was one-dimensional, with all cone types (if the animal had more than one) pooled. When color vision evolved, two more axes were added to the one already present. It would make more sense to do that than to throw out the pooled system already present for black-white and then have to erect three new ones. When we adapt to the dark and are using only our rods, our vision becomes colorless and is again plotted along one axis, to which the rods evidently contribute. That would not be easy to do with r, g, and b axes.

At present we can only guess how double-opponent cells are wired up. For several reasons we suspect that they receive their inputs from type 2 cells: their field centers are about the same size as type 2 field centers and much larger than centers of type 1 cells;

they are intermixed with type 2 cells in the blobs; and finally, as already pointed out, type 1cells, with their opponent-color antagonistic surrounds, seem especially inappropriate as building-blocks for double-opponent receptive fields. For a red-on center double-opponent cell, the simplest arrangement, as illustrated on the next page, would be to have excitation from one or a few red-on green-off type 2 cells whose field centers were contained within the double-opponent cell's center, and excitatory inputs from red-off green-on type 2 cells whose centers were scattered through-out the double-opponent cell's periphery. Or the surround could be formed by inhibitory inputs from red-on green-off type 2 cells. (Originally we favored a scheme in which the inputs were made up of type 1 cells. Logically such an arrangement is possible, but it seems much more awkward.)
This leaves unsettled the part that type 1 cells play in color vision—if they play any part at all. These are the most common cells in the lateral geniculate body, and they supply the lion's share of the input to the visual cortex. Their obvious color coding makes it easy to lose sight of the fact that they are beautifully organized to respond to light-dark contours, which they do with great precision. Indeed, in the fovea, where their centers are fed by one cone only, they have no choice but to be color-coded. (The mystery is why the surround should be supplied by a single, different, cone type; it would seem more reasonable for the surrounds to be broad-band.) Given this massively color-coded input it is astonishing that interblob cells in the cortex show so little interest in color. The few exceptions respond to red slits but not to white ones, and are thus clearly color coded. For the most part it would seem that the information on wavelength carried by type 1 cells is pooled, and the information about color lost.
In one sense, however, it is not discarded completely. In Freiburg, in 1979, Jiirgen Kriiger and Peter Gouras showed that cortical cells often respond to lines formed by appropriately oriented red-green (or orange-green) borders at all relative intensities of red and green. A truly color blind cell, like a color blind person, should be insensitive to the border at the ratio of intensities to which the cones respond equally. These cells presumably use the type i color information to allow contours of equal luminance to be visible by virtue of wavelength differences alone—of obvious value in defeating attempts at camouflage by predators or prey. The recognition of colors as such would thus seem to be an ability distinct from the ability to detect color borders, and to require a separate pathway consisting of type 2 cells and color-opponent blob cells.
Our tendency to think of color and form as separate aspects of perception thus has its counterpart in the physical segregation of blobs and nonblob regions in the primary visual cortex. Beyond the seriate cortex the segregation is perpetuated, in visual area 2 and even beyond that. We do not know where, or if, they combine.

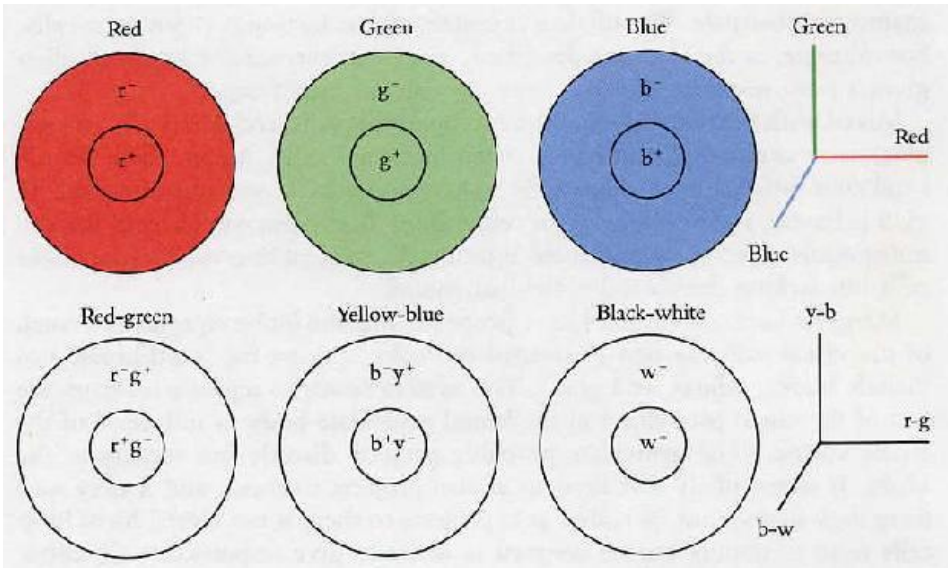

Top: Land's original formulation of the color-constancy problem seems to call for three kinds of cells, which compare the activation of a given set of cones (red, green, or blue) in one region of retina with the average activation of the same set in the surround. The result is three numbers, which specify the color at the region. Thus yellow, brown, dark gray, and olive green each has a corresponding triplet of numbers. We can therefore plot colors in a color space specified by three axes, for red, green, and blue. *Bottom*: A mathematically equivalent system also gives three numbers, and is probably closer to the way the brain specifies color. At any point on the retina, we can speak of red-greenness, the reading an instrument would give if it were to record the relative stimulation of red and green cones (zero for yellow or white). This value is determined for a particular region, and an average value is determined for the surround; then the ratio is taken. The process is repeated for yellow-blueness and black-whiteness. These three figures together are enough to specify any color.

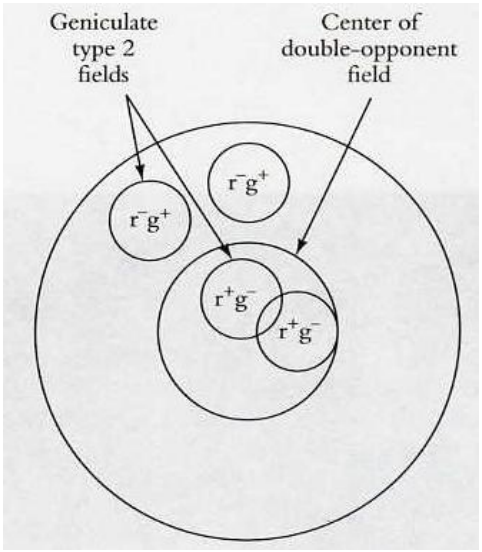

A double-opponent cell could be built up from many geniculate type 2 cells. If the cell is r+'g- center, r-g+ surround, then its inputs could be a large number of r g type 2 cells with fields scattered throughout the cell's receptive field center, and r g type 2 cells with fields scattered throughout the cell's receptive field surround, all making excitatory contacts with the double-opponent cell. Alternatively, the surround could be formed from r+g- type 2 cells, making inhibitory contacts.

CONCLUSION

The subject of color vision illustrates so well the possibilities of understanding otherwise quite mysterious phenomena—the results of color mixing or the constancy of colors despite changes in the light source—by using a combination of psychophysical and neurophysiological methods. For all their complexity, the problems presented by color are probably simpler than those presented by form. Despite all the orientation-specific and end-stopped cells, we are still a long way from understanding our ability to recognize shapes, to distinguish shapes from their background, or to interpret three dimensions from the flat pictures presented to each of our eyes. To compare the modalities of color and form at all may itself be misleading: remember that differences in color at borders without any differences in luminous intensity, can lead to perception of shapes. Thus color, like black and white, is just one means by which shapes manifest themselves.