Reprinted from

PROCEEDINGS OF THE ROYAL INSTITUTION OF GREAT BRITAIN

Volume 47

1974

The Retinex Theory of Colour Vision

EDWIN H. LAND

APPLIED SCIENCE PUBLISHERS LTD

LONDON

The Retinex Theory of Colour Vision

EDWIN H. LAND

We live in a deep sea of radiation that falls on all the objects around us and which is scattered back to us by those objects. The composition and quantity of this radiation is always fluctuating and always unpredictable. The experience of seeing objects cannot depend on the quantity and composition of the radiation scattered by them to our eyes because objects would then vary in appearance as the radiation changed. An object could not have a permanent colour; its colour would be always changing, so that an apple might be the colour of blueberries at one moment and of oranges at another.

This Discourse is about a generally unrecognised animal sense—the ratio-making sense. It is the ratio-making sense which processes the radiation reaching our eyes in such a way as to discover the constant properties of objects in relation to the radiation falling on them. It is the ratio-making sense which keeps an apple looking like an apple in blue sky light when more blue light than red is coming from the skin of the apple to our eyes. It is the ratio-making sense which enables us to see a white paper as a white paper and a dark grey paper as a dark grey paper, even when they are so illuminated that more light may be coming from the dark paper than from the white.

We can best appreciate and understand the ratio-making sense by performing some of the experiments which demonstrate its existence. We note that any piece of paper which has the same efficiency of scattering for all wavelengths in the visible spectrum looks white or grey or black. Therefore, it is convenient to use white, grey and black papers to study the first step in the system within which the ratio-making sense operates. After defining the ratio-making sense by experiment we will proceed step by step to construct a conceptual system based only on experiments presented during the Discourse. At the end of the Discourse we will relate our total concept as it will evolve tonight to several of the principal historical views and conclusions.

We first measure the amount of light scattered back from a series of pieces of paper held successively in one position on a black velvet background and illuminated by a spotlight. Figure 1 shows a piece of paper on the velvet background and a telescopic photometer on a tripod. The slide projected above the speaker shows the view through the telescope. In the slide the black dot at the centre of the crosshairs is the end of a piece of fibre optics attached to the photosensitive element. The meter can read a very small area on the display, about the size of an eraser on a pencil. All the measurements in this Discourse are made with this telephotometer to determine the energy (watts per unit area per unit solid angle) reaching the audience's eyes. Throughout all these demonstrations we use matte papers, that is papers that reflect a constant amount of light in every direction. This ensures that the energy values at the eyes of the audience are the same as they are at the location of the meter. The telescope can be directed, just as one's eyes are directed, to any point in the field of view. The photometer measurement is read on a galvanometer scale projected overhead (see Fig. 1).

We measure the amount of light coming from each of several different papers. We find that a medium grey paper is reflecting a quantity of light that makes the meter read 15 (in arbitrary units). A second paper, dark grey, gives a reading of 5. A third light grey paper reads 35. The black background reads less than 1·0 and a white paper 88.

We take the paper with the highest reading (88), return it to the board and, for easy computation, set the voltage on the spotlight for a photometer reading of 100. 100 is then written on the paper. Without changing the voltage-setting, we remeasure each paper and write 55 on the light grey, 25 on the middle grey, 8 on the dark grey and 1 on the black velvet.

Having measured the energy of the radiation reflected from the papers, it is necessary to measure the efficiency with which the papers reflect the light falling on them. Efficiency is the ratio of two measurements, as illustrated in Fig. 2. First the amount of light coming from a piece of paper is measured with a meter. Second, without altering the illumination we measure the amount of light coming from a standard in the same place as the paper. Usually the highest reflector practically available is used as a standard, for example, magnesium oxide or barium sulphate. This operation characterises the efficiency of a paper as compared to the most efficient reflector of light.

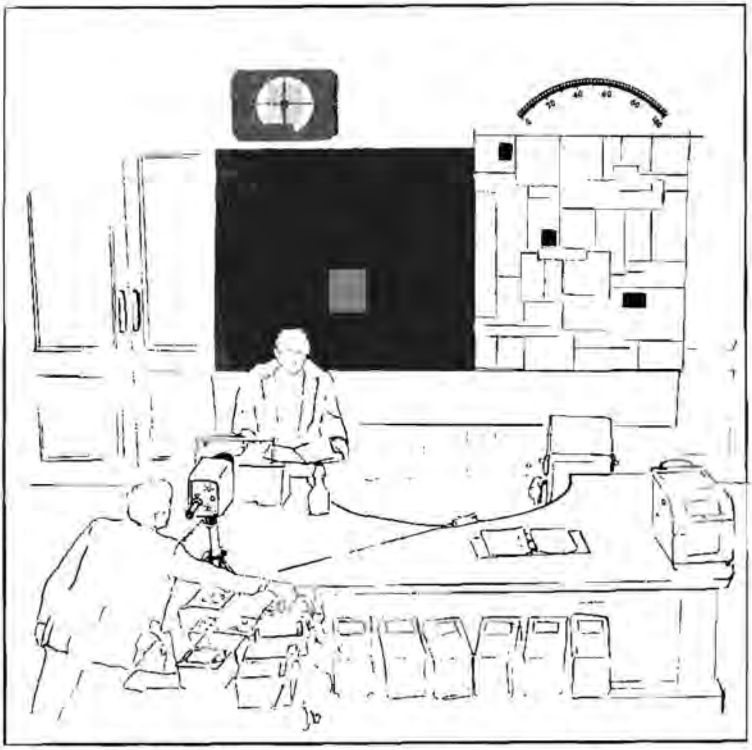

Fig. 1. Sketch of the demonstration apparatus used in the Discourse. Just to the right of the door was a large, black velvet background with a grey paper. This paper was measured with a telescopic photometer mounted on a tripod located near the bottom of the sketch. A slide which had been taken through the photometer was projected above the speaker's head. The scale of the meter and its reading were projected above the Mondrian display on the right.

In tonight's measurements magnesium oxide will be replaced as the standard by the paper with 100 written on it.

We measure each of the papers in constant illumination, holding them one by one at a single position on the board. Their values are then written on them. These are 100/100 on the white paper, 55/100 on the light grey, 25/100 on the middle grey, 8/100 on the dark grey paper and 1/100 on the black velvet. These fractions represent the various efficiencies of the papers and are usually referred to as reflectances. Both the reflectance and the energy numbers are external physical measurements. Reflectance is a property of the objects we look at and energy of the radiation that falls on our retinas. A third concept is needed to designate the sensations from white to black. Let us call those sensations lightnesses. Since each lightness is a position in a *sensation* scale it cannot be measured with a photometer, it must be experienced by a visual system. Furthermore, whereas both energy and reflectance can be measured at any *point* in a display, the

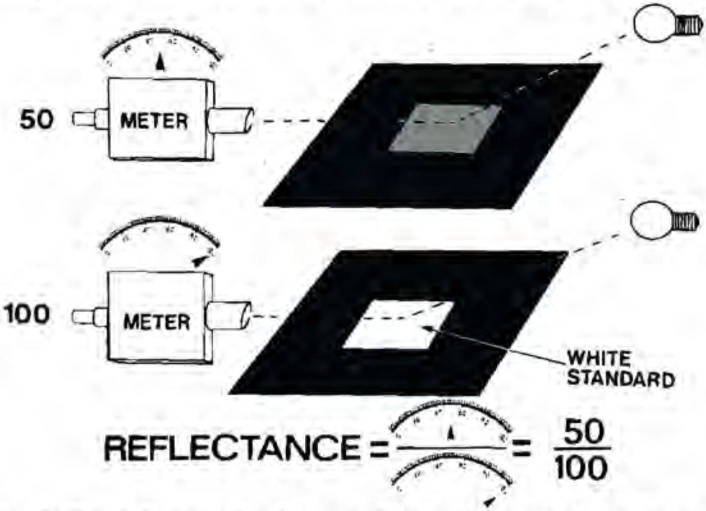

Fig. 2. The definition of reflectance. First we measure with a meter the amount of light coming from a high reflectance standard. Second, in exactly the same conditions, without altering the illumination, we replace the standard with the paper to be tested and measure the amount of light reflected. The ratio of the second measurement to the first is the reflectance of the paper.

biological quantity lightness, in the rôle that it will play tonight, is essentially a characteristic of an area.

We place the papers in the Mondrian-like array shown in Fig. 3. Certainly, as we inspect this array of five papers we would take it for granted that the lighter the paper, the more light is being sent from it to our eyes. This observation leads to a first hypothesis, namely that the amount of light coming to the eye from an area determines its lightness. Thus, by this hypothesis the lightnesses of

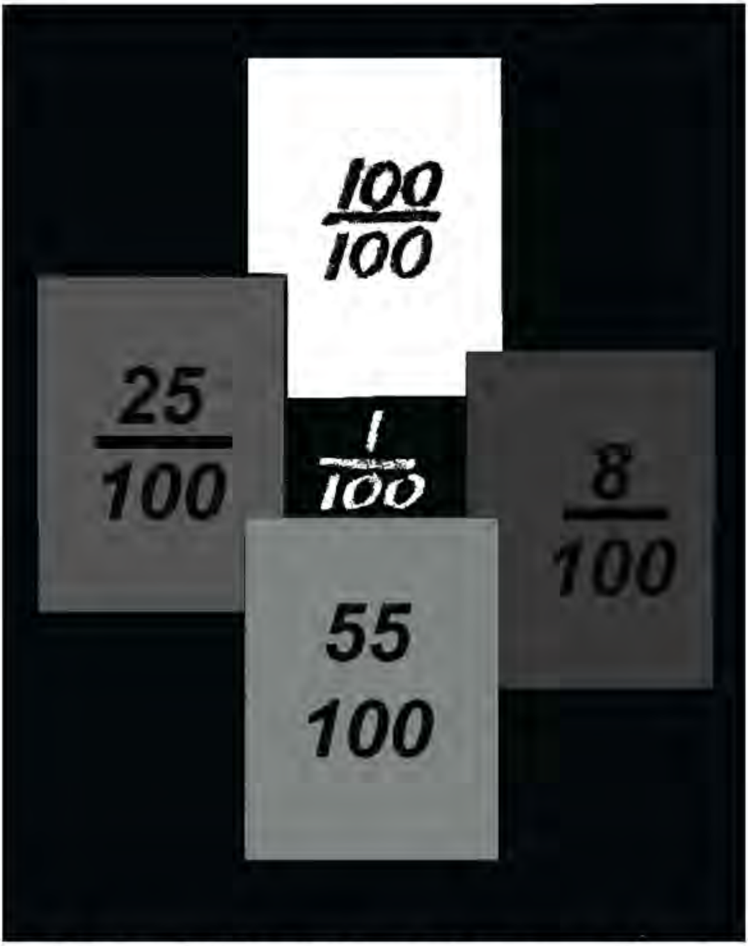

Fig. 3. Photograph of the velvet board with grey papers. The numbers written on each paper are representative of two different hypotheses about how we see light and dark sensations. The numerators are the energies coming from each area to the eye. The ratio is reflectance, namely the relative energy compared to the paper with the highest efficiency for reflecting light. Subsequent experiments show which hypothesis agrees with what we see.

these papers are determined by the values of the numerators, 55, 100, 8, 1, and 25. This is referred to here as the *Energy hypothesis*.

At the same time, however, it can be seen that the lighter the paper, the larger the fraction denoting its efficiency. Thus the lightest paper is characterised by 100/100, a black background by 1/100. This observation leads to a second hypothesis, namely that the fraction of light reflected determines lightness. This is referred to here as the *Reflectance hypothesis*.

Without further experiment we cannot tell which of these two hypotheses provides the dominant reason for the position of the paper in the scale of lightnesses from white to black. If the Energy hypothesis is correct then in three of the four experiments described below the papers should change their lightnesses. If, on the other hand, the Reflectance hypothesis is correct, there should be no significant change in the lightnesses of papers in the display.

When an experiment with the grey paper is carried out in an open room we may tend to over-conceptualise the grey paper as 'the experiment' and to forget that as far as the retina is concerned the display comprises the whole room which forms the pattern on the retina, the grey paper being just one element of thousands composing the pattern. There is no escape from the reality of the pattern on the retina, and the substitution of an infinite void for the room as a surrounding for the grey paper introduces a gravely dangerous artifact, the special case in which the grey paper, whatever its efficiency, is by definition the most efficient scatterer in the field of view. At this stage of the investigation acceptable artifacts can, however, be contrived as substitutes for the multi-coloured room in the form of arrays of white, grey, and black papers, arranged in patterns reminiscent of Mondrian's designs. In these patterns the areas are variegated in size, shape, and greyness to provide enough unpredictability to replace on the retina the pattern of the room and yet enough predictability to enable us to make those measurements necessary for the derivation of laws. For these experiments a five-foot-square collage of matte papers, ranging in reflectance from 4/100 to 100/100 and randomly arranged in a display called the Black and White Mondrian, is used (Fig. 4). This display is located beside the velvet board at the front of the room where it can be illuminated by a projector equipped with a camera shutter, or by a floodlight placed on the floor in front of it.

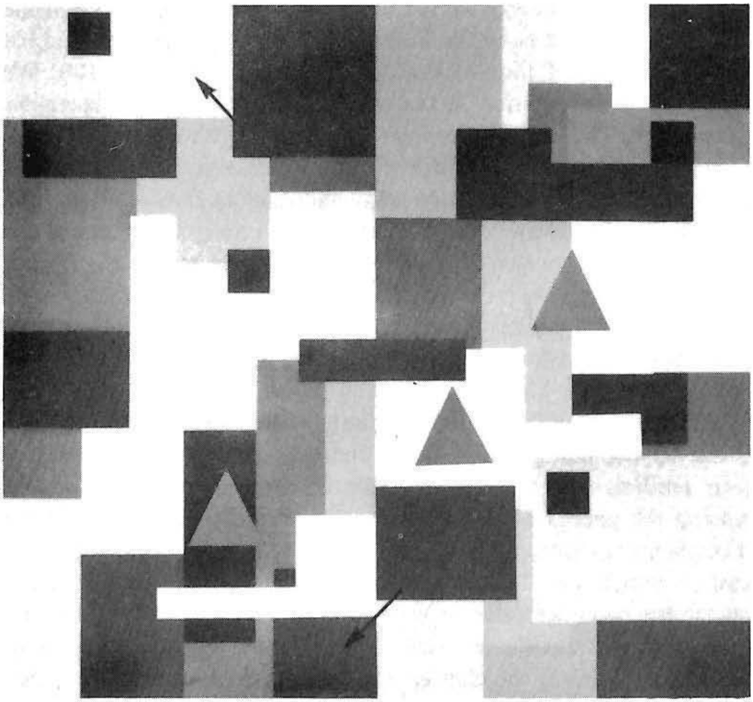

Fig. 4. Photograph of the Black and White Mondrian used in the four experiments that tested whether the lightnesses we see are predicted by the energy hypothesis or the reflectance hypothesis. The three triangular, medium-grey areas were cut out of a single piece of paper. The small difference in lightness shows the extent to which the immediate surround can modify the lightness of areas. Even when one triangle is on a black surround and a second is on a white surround the difference in lightness is small compared to the range of lightness from white to black. The arrows at the top and bottom point to the two areas used in the fourth experiment. The illumination was controlled so that the amount of light falling on each area exactly compensated for the difference in reflectance. The area at the top continued to look light and the area at the bottom continued to look dark even when the same energy came from both. Photographic transparencies which reproduce the experimental stimuli are included in a paper by Land and McCann (1971).

Experiment 1: Cut the illumination by a factor of 20. We point the meter at the white area in the Black and White Mondrian and adjust the overall level of illumination so that the meter reads 100. We then observe the position in the scale of lightness of the rectangles in the array. The energy coming from a nearly black area reads 5. We return the meter to monitor the energy from the white area and change the overall illumination until the reading on the white area drops from 100 down to 5. The measured-energy from the white area has been decreased 20 times so that it now equals the energy that previously came from the black area. All of the areas retain their hierarchical positions and although all the areas in the Mondrian appear slightly darker, the whites are white, the greys are grey and the blacks are black.

The Energy hypothesis predicts that the lightness of areas should change as the energy from them changes. The Reflectance hypothesis predicts that the lightness should remain constant since increasing the energy affects the numerator and denominator of the reflectance fraction equally. From this simple experiment—Experiment 1—which we take for granted whenever we turn on an extra lamp in the room we enter—it can be seen that changing the amount of light does not significantly shift the position of an object in the scale of lightness. In fact, the change in appearance of each paper is surprisingly small when it is considered that we reduce the illumination by a factor of 20. (On the other hand, as we look around the Mondrian, we note that a change of efficiency by a factor of 20 changes lightness the whole way from white to black.) Experiment 1 supports the Reflectance hypothesis.

Experiment 2: Illuminate the Mondrian with an unknown amount of light for 1/10 of a second. To start this second experiment the lens of the illuminator is covered with a camera shutter that can open to pass a pulse of light of one-tenth of a second duration. While the shutter is closed the voltage on the filament—and therefore the level of illumination that falls on the Mondrian—is changed to some unknown new value. The audience keeps their eyes lightly closed for a moment so that they will reach some unknown new state of adaptation to light. They open their eyes—the shutter is snapped. They see a bright display in which the lightness of each area is the same as that seen in continuous illumination. The Energy hypothesis might have been amended by proposing that energy with the aid of an adaptation mechanism might account for the independence of lightness from illumination. The proposal of such an adaptation mechanism would imply that when we decrease the amount of light falling on the whole field of view the light detectors would become less light-adapted and hence more sensitive. Although such an adaptation mechanism might be feasible for the first experiment it is definitely not feasible for the second experiment. For here there is no way the eye can know how to change its sensitivity to compensate for an unknown change in energy.

The Reflectance hypothesis without additional assumptions predicts the lightnesses we see. Changing the overall illumination changes both the numerator and denominator of the reflectance fractions by the same factor. Experiment 2 supports the Reflectance hypothesis.

Experiment 3: Place several identical grey, triangular pieces of paper in various positions on the board, that is, in a variety of surroundings. For this third experiment we take three identical pieces of paper distinguished by their triangular shape and place them on the display to demonstrate that the position in the lightness scale is not a consequence of the triangle's surroundings, which are designed to be arbitrary and variegated. The lightnesses of these three triangular papers are essentially independent of where they are placed. We place them in extreme positions: one in a white area, and a second in a black area (Fig. 4). There is a very small difference between them—the largest change one can get by varying the surround. This change is very small compared to the lightness range from white to black. Position and lateral relationship to the surround have an effect, but that effect is not a significant determinant of lightness. In this experiment both the Energy hypothesis and the Reflectance hypothesis predict that the three triangles should appear the same lightness as each other. Surround, like adaptation, cannot be used in conjunction with the Energy hypothesis as the principal determinant of lightness.

Experiment 4: Illuminate the display so that the energy from a paper with a high reflectance is equal to the energy from a paper with a low reflectance, both being visible simultaneously. In this experiment a paper at the top of the display looks white and another at the bottom nearly black. We show that the energy reaching the audience's eyes from the white paper exactly equals the energy coming to their eyes from the nearly black paper. Thus, this experiment is the *coup de grace* for the Energy hypothesis.

The source of light is placed 3 ft away from the bottom of the 5-ft display. The light source is a little over 9 ft from the top, and this reduces the light falling on the top edge to about 1/9 the light falling on the bottom edge. Much more light falls on the low reflectance paper at the bottom than on the high reflectance paper at the top. We measure the amount of light coming from each area using the photometer. This arrangement of light was chosen to compensate exactly for the difference in reflectance of the two papers. (In Fig. 4 see areas with arrows.) The paper with a high reflectance looks white, and that with a low reflectance dark, even when the same energy is coming from both. The Energy hypothesis predicts that the two areas must appear identical since the light falling on the retina from both is identical. The Reflectance hypothesis predicts that the areas should be at opposite ends of the lightness scale. Experiment 4 supports the Reflectance hypothesis.

In these experiments, we have shown that, for a given piece of paper, we can change the overall *amount* of light falling on the paper, the *duration* of the light, and the *surround* of the paper without substantially changing the lightness sensation associated with the paper. Furthermore, two areas, markedly different in reflectance, were very different lightnesses when the illumination was arranged so that the energy at the eye from one exactly equalled the energy at the eye from the other. In all four experiments the lightnesses seen by the audience were consistent with the Reflectance hypothesis.

Thus we have shown experimentally that lightnesses are not correlated with energies but with reflectances, in an orderly way. Since reflectances are by definition ratios, we can conclude that the lightness-generating sense is a ratio-making sense. It is clear that the ratio-making sense is a powerful and reliable tool. As far as I know, nobody is lightness-blind because we all come to the same conclusions about the hierarchical positions in an array of lightnesses, in spite of the most variegated conditions of illumination.

A METRIC FOR LIGHTNESS

In the next demonstration we can establish the relationship between the *external* ratio, reflectance, and the *internal* sensation, lightness. We make an equal interval scale of lightnesses by offering members of the audience a variety of white, grey, and black papers, and requesting them to pick the extremes, then to pick a paper that lies half-way between the extremes, then two more that lie respectively half-way between the middle and one extreme and between the middle and the other extreme. The process continues in this manner until a full scale has been constructed. The physical reflectance of each paper

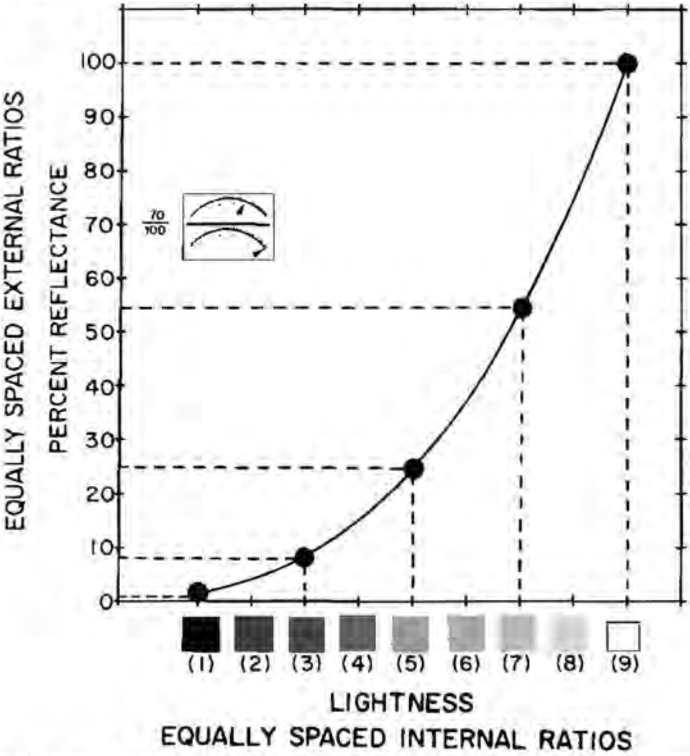

Fig. 5. Graph showing the relationship between reflectance and lightness. We repeat this classical experiment to emphasise the concept we are developing of a correspondence between internal ratios and external ratios. The audience is given paper No. 1 (black) and No. 9 (white) and asked to find a grey paper whose lightness is half-way between Nos. 1 and 9. The chosen paper has a reflectance of 25/100 and is plotted at the lightness 5·0. The audience is then asked to find papers for positions Nos. 3 and 7. The papers chosen have reflectances of 8/100 and 55/100 respectively. The solid line is the Munsell Value curve of equally spaced lightness adapted from Newhall, Nickerson and Judd (1943).

in the series is then measured and a plot made of the position of a paper in the scale of reflectances (Fig. 5). The simple curve on which the points fall for a large number of observers shows that the lightness is the 'pointer' for an internal meter that somehow ascertains the relative efficiencies with which areas in the field of view return whatever illumination is falling on them. What information, then, must the internal ratio-maker utilise to perform its task?

The lightnesses, we see, correspond to the reflectances of the black, white and grey papers. Yet ascertaining reflectance in any of the familiar ways requires an operational step which the eye cannot take. For example, the eye cannot insert a comparison standard next to the object it is regarding. Furthermore, what reaches the eye from each point is clearly the product of the reflectance and the illumination. There remains the circular logical problem that, because the light coming to our eye is the product of the reflectance and illuminance, our eye could not determine reflectance unless illuminance were known and the eye could not determine illuminance unless the reflectance were known. In general, across the field of view, neither reflectance nor illumination is known; and neither is uniform.

My view of evolution is that biological systems do not solve problems but rather evolve in a way such that problems do not arise. The eye–brain system must have evolved in such a way that illumination and its unpredictability do not enter into its determination of lightness. How does the eye–brain do it? In trying to understand how, let us follow Einstein's programme: To understand a natural phenomenon keep inventing a series of machines until you arrive at one with properties that match those of the particular natural domain. Explanation in science can be regarded as the successful search for such a machine—a machine which will serve as a model until it is displaced by requirements imposed by further knowledge.

We can invent for the machine a procedure which determines reflectance without ever knowing or involving the value of illumination, a procedure which permits the illumination to be unknown and unknowable.

We place two grey papers on the board, edge to edge, and take two readings, one on either side of the edge and as close together as possible. The ratio of these readings is then determined. Ideally, the meter would make both readings simultaneously and give only the ratio without bothering to report the two individual readings. We start with an arbitrary selection of two papers and with arbitrary lighting. The two telephotometer readings on either side of the edge are respectively 90 and 30. If we had the meter with a built-in photometer bridge it would read directly the ratio of 30 to 90, that is 1/3. We replace the right-hand paper with another piece of paper and take the readings again. The telephotometer readings on either side of the edge are found to be respectively 90 and 45 and the ratio is 1/2.

We have invented a strategy which our machine uses to determine the reflectance ratio. The key assumptions that make this strategy possible are: (1) changes in reflectance are discontinuous, are abrupt, and form edges; (2) changes in illumination are continuous, slow, and do not form edges, and (3) changes in illumination that are discontinuous and form edges may be regarded as changes in reflectance by both the human visual system and the machine. When two energy readings that form the bridge pair are made at positions closer and closer together, the influence on the ratio exerted by gradually changing illumination decreases. As the ratio of illuminations approaches unity, the ratio of energies approaches the ratio of reflectances.

We can continue in the above manner, replacing the right-hand paper with first one, and then another paper. Because of the bridge structure we are free to alter the illumination in any way we choose. The bridge photometer continues to take readings and make ratios. Whenever we bring to the right-hand position a paper which reflects the same energy as the paper on the left, we will have a ratio of 1·0 across the edge and whenever we bring a paper to the right-hand position which reflects more energy than the one on the left, the bridge-meter will read a fraction larger than 1·0.

When we do find that the paper on the right is higher than the one on the left, we then replace the left-hand paper, which we had been using for a standard of comparison, with the right-hand one. We then remeasure the series of papers. Once again, whenever a reading is obtained that is higher than any previous reading the paper concerned is placed on the left and used for a standard. Thus we have invented for the machine a primitive procedure that, given a variety of choices, determines the highest reflectance among the choices and uses it as standard.

Although this procedure satisfies the requirement for neither knowing nor involving the illumination, it still requires that any paper be placed next to the one with which it is compared. Our machine should have a way of comparing areas that are distant from one another without losing the properties of the bridge system. Some years ago, when we were contemplating an array of black, white, and grey papers and puzzling about the next step, John McCann invented a new Gestalt, a way of looking at the reflectance relationships between widely separated areas. Imagine two pieces of paper that are far apart. A new piece of paper is cut to size and placed so that one edge is adjacent to the first area and the other edge is adjacent to the second distant area. In other words, the paper is as long as the separation between the two areas. Two readings with the bridge photometer are made, the first at the edge between the first area and the long, new paper and the second at the edge between the long, new paper and the second distant area. The bridge photometer reads two reflectance ratios. When these two ratios are multiplied together the product equals the reflectance ratio of the first area to the distant second area. In other words, the product of the ratios across distant edges is the same as the ratio we would get if we placed the two remote areas next to each other. A little thought will show that the new, long paper can be replaced by two or more papers providing that we make a bridge photometer reading at every edge. Therefore by simply proceeding from one area to the next, by multiplying the bridge photometer readings across edges, we can establish a series of sequential products which yields the relative reflectance of any area relative to the first area without knowing or involving the illumination.

Instead of computing reflectances in this manner relative to the *first* area in the row the machine should compute reflectances relative to the *highest reflectance* in the display. If the first of a row of papers has the highest efficiency in that row then all the sequential products as we move along the row will equal the reflectance ratio measured with the highest efficiency paper as standard. If the first area does *not* have the highest efficiency then all the sequential products will be larger than they would be if the first area did have the highest efficiency. Thus areas with higher reflectances than the starting area will have sequential products greater than 1·0. In the earlier procedure one paper was placed next to another and the ratio across the edge read. If the ratio was greater than 1·0, we learned that the paper we had been using as a standard had a lower reflectance than the new paper. We then began anew with the new paper and used it as a standard. Here with a series of papers fixed in place we can also find a way to begin again. We utilise the fact that the sequential product can exceed 1·0 only in an area of higher reflectance than the starting area. Therefore, when we find an area with a sequential product greater than 1·0 we can start all comparisons from there. We have now expanded the ability of our hypothetical machine so that it can establish the reflectance of any area along a row by multiplying the ratios from the bridge photometer readings and by starting over again when the multiplication product exceeds 1·0.

So far the machine has been applied to a single row of papers, and complex displays, such as the Black and White Mondrian, have been ignored. What remains is to expand the model so that it can calculate the reflectance of every area in the Mondrian. Since the Mondrian is a two-dimensional display of papers, and not a simple row, a systematic plan is required for making bridge photometer readings and sequentially multiplying them to calculate reflectances. There are many equivalent ways in which these ratios can be measured and many ways sequential multiplication can be computed. The choice of a particular technique is guided by biological feasibility and personal taste. For simplicity of argument we begin with the assumption that there are many computational paths, each of which can start anywhere and proceed in any direction. At the beginning of the path the sequential product equals the reflectance ratio of two adjacent points, and as we proceed along the path the sequential product equals that ratio multiplied by the ratios from succeeding pairs of points. The value of the sequential product at each pair of points along the path is the partial output of the machine. The total output is the average of all the outputs from all the paths coming from all points of the compass reporting at a single place.

Therefore if we had the biological ability for reading the ratio of energies across edges and for multiplying these ratios, we would arrive at a correct evaluation of the reflectance of each area, all based on this programme of sequential multiplication of ratios of energies across edges. Let us stress again that this technique determines reflectance reliably, irrespective of the level of illumination and the uniformity of the illumination.

The actual biological path, however, must be reading ratios not only across edges between areas of markedly different reflectances, but also as it traverses areas in which the change from point to point is only the small change due to graded illumination. These small changes must not be summated and incorporated in the sequential product since this would raise and lower the lightness of the various areas in response to the vagaries of variable illumination. All of the derivations you will see tonight are based on the presumed availability of biological detectors which cannot read gradients smaller than some threshold value. The bridge photometer cannot detect the small differences caused by gradual changes in illumination. The net result is that the sequential product of bridge photometer ratios for each area is the same as the sequential product of the bridge reading which would have been taken only at the edge.

It must be kept in mind that these operations are the properties of a machine model of the visual system and that such a model is distinct from the actual visual system. A particular way of calculating lightness must be equivalent, but not necessarily identical to, the actual biological processing. There are many different ways of expressing the basis of this model. For example, in this description uni-directional paths have been used as a technique easily adapted for computer simulation (Land and McCann, 1971). The model does not predict that the visual system uses uni-directional paths. It could have used field-effect calculations that process the information for all points simultaneously, or some combination of a variety of computational methods. The model does predict that alternative calculations must give the same results as uni-directional paths. The essence of the model is not affected by the particular choice of computational technique.

EXPERIMENTS WITH COLOUR PAPERS

In this experiment the lights in the room are turned off and two different papers are placed on the Black and White Mondrian. The first paper is easily identifiable by its circular shape, the second by its triangular shape. We illuminate the Mondrian with three projectors, one at a time. Each projector has a different sharp-cut, band-pass interference filter. One filter transmits the long-waves, the second the middle-waves and the third the short-waves (Land, 1964; Land and McCann, 1971). In long-wave light the circular paper is light and the triangular paper dark. In middle-wave light the circular paper is dark and the triangular paper light. In short-wave light the circular paper is dark and the triangular paper is dark, but slightly lighter than the circular paper. All the papers in the Black and White Mondrian remain the same lightness in all three illuminants. It is therefore concluded that the circular and triangular papers have special properties not found in the papers that make up the Black and White Mondrian. The triangular and circular papers each have three different lightnesses, one for each waveband, whereas each black, white and grey paper has the same lightness in every waveband. When all three projectors are turned on simultaneously, the circular paper is red and the triangular paper green, while the Mondrian remains a collection of whites, greys, and blacks. The rest of this Discourse studies the relationship of energy at the eye, reflectance of the object, lightnesses of the object and colour of the object in search of the true determinants of colour.

As in the first part of the Discourse, we can consider two alternative hypotheses for colour. The first is an energy hypothesis, namely that the triplet of energies, one for each waveband, determines the colour we see. According to this hypothesis, the red paper looks red because it sends to the eye larger quantities of long- than middle- or short-wave light; similarly the green paper looks green because it sends more middle-wave light than long- or short-wave light. This hypothesis proposes that colour is determined by the proportions of long-, middle- and short-wave energies. The second hypothesis is that a triplet of lightnesses determines colour, that is, colour is *compared lightness*. With this hypothesis we would characterise the red paper with three lightnesses: light—dark—dark, and the green paper with dark—medium-light—dark. From the results of the experiments we have seen tonight we expect each of these lightnesses to depend on the reflectance of the area in a given waveband.

A strong hypothesis may now be made, namely that every colour is determined by a unique triplet of lightnesses but not by a triplet of energies. In the Colour Mondrian experiment we will test that hypothesis.

THE COLOUR MONDRIAN EXPERIMENT

The Colour Mondrian experiment uses two large four-and-a-half foot square arrangements of coloured papers (see Fig. 6, colour plate). To minimise the rôle of specular reflectance, the papers are not only matte, but are also selected to have a minimum reflectance of at least

10 per cent for any part of the visible spectrum. Each Mondrian in this display is illuminated by the three illuminating projectors described in the previous experiment. The amount of illumination from each projector is controlled by a separate variable transformer. In addition, the projectors have synchronised solenoid-activated shutters to control the duration of the illumination.

With all three illuminating projectors turned on, the variable transformers are set so that the whole array of variegated papers in one Mondrian is deeply coloured and so that, at the same time, the whites are good whites. This is not a critical setting. Then, using one projector at a time, and hence only one waveband at a time, we measure with the telescopic photometer the energy at the eye from any particular area, say a white rectangle. Thus, from a white rectangle, three numbers are obtained that characterise the wavelength-energy distribution of the light at the eye. (The subsequent procedures constitute a null experiment. The particular units of measure, the wavelength sensitivity and the linearity of the meter are not significant in the experiment.) The readings from the white are: 64 $mW.sr^{-1}m^{-2}$ of long-wave light, 31 $mW.sr^{-1}m^{-2}$ of middle-wave light, and 5·2 $mW.sr^{-1}m^{-2}$ of short-wave light.

We have now established the three energies associated with this particular white sensation. With a second identical Colour Mondrian display we can test the hypothesis that the long-, middle- and short-wave energies at a point determine the colour sensation. The three projectors illuminating the first Colour Mondrian on the left are turned off. On the right only the long-wave illumination is turned on. A new area of unknown colour is selected and the long-wave illumination adjusted until the long-wave energy coming to the eye from the selected area is the same as the long-wave energy that, a short while before, came from the white paper, namely 64 $mW.sr^{-1}m^{-2}$. The transformers controlling the middle- and short-wave illumination are then separately adjusted so that the energies sent to the eye from the area are the same as, shortly before, came from the white, namely 31 and 5·2 $mW.sr^{-1}m^{-2}$. Thus, the energies in each waveband from the new area are identical to the three energies which had previously reached the eye from the first white rectangle. Up to this point the audience has been prevented from seeing the second Mondrian in more than one waveband at a time, lest any effects of colour memory enter into the experiment. The observers know nothing about the paper except that it sends to the eye the same energy as a white had previously. The camera shutters on the three projectors are synchronised so that the second Mondrian is illuminated for only 1/10 of a second. The audience responds by stating that the area is 'yellow'. The yellow sensation is associated with the new area despite the fact that the triplet of energies for that area is identical to what, a short while before, came from a white area.

The reason we restricted the light to 1/10 second exposure was to restrict the quantity of light available to bleach visual pigments and hence restrict the light available to alter the state of adaptation of the eye. While one does not see how one could invent a programme whereby adaptation could lead to the result obtained, it was, nevertheless thought desirable to anticipate and dispose of the issue. We now turn on all three projectors and leave them on. This provides a sufficient quantity of light for the eye to reach a new state of adaptation. The audience reports that the area is the same yellow as it was in the 1/10 second flash. If we propose an adaptation mechanism that is relatively insensitive to the quanta caught from the 1/10 second exposure, such as adaptation due to bleaching visual pigment, we cannot explain why the area is yellow in the flash. Alternatively, if we propose an adaptation mechanism that is faster than pigment bleaching and is so sensitive that it adapts sufficiently to explain the yellow colour with less than the 1/10 second exposure, we cannot explain why that process does not continue to change the appearance of areas in continuous illumination. With such fast, sensitive adaptation one would expect to be able to alter significantly the colours of areas in the Mondrian by light adaptation. In this experiment we have been unable to produce significant alteration of colour sensations *in complex, multi-coloured scenes* by changing the state of adaptation.

The three projectors illuminating the Colour Mondrian on the right are now turned off. On the left only the long-wave illumination is turned on and the photometer pointed towards an unknown area on the first Mondrian. Then, in turn, the three projectors for that Mondrian are readjusted separately, so that the three energies that will come to the audience's eyes are again 64, 31, and 5·2 $mW.sr^{-1}m^{-2}$. Again the synchronised shutters are used to view the area for 1/10 of a second. The audience reports green from the same triplet of energies that produced yellow and white.

The next step is to turn on the three projectors on the left illuminating the Mondrian on the left and the three projectors on the right illuminating the Mondrian on the right. It can now be seen that the yellow area and the green area are each simultaneously sending to the observer's eye identical wavelength-energy distributions yet are respectively remaining bright yellow and bright green.

The validity of the photometer measurements can be confirmed by an important control experiment. For this two black viewing tubes, constructed with internal baffles and an adjustable aperture, as shown in Fig. 7, are used. This tube design is important to produce a true optical void so that there is no light reaching the eye, except from the area seen through the aperture. (If the tube does not contain baffles and if the aperture in the end of the tube is not small with respect to the tube's diameter, then the tube becomes a cylinder that integrates the light from surrounding areas and thus forms an 'average surround' around the area seen through the aperture.) The two tubes are positioned so that the light from the middle of the yellow area in the right-hand Mondrian reaches the observer's right eye and the light from the middle of the green area on the left-hand Mondrian reaches his left eye. When the two areas are observed in this manner the sensations are identical to each other and completely different from both the yellow and the green seen in the Mondrians without the tube. The two identical colours are much lighter; they are greyish-white. The observer can move his head and see all of both Mondrians and then move his head back to see only

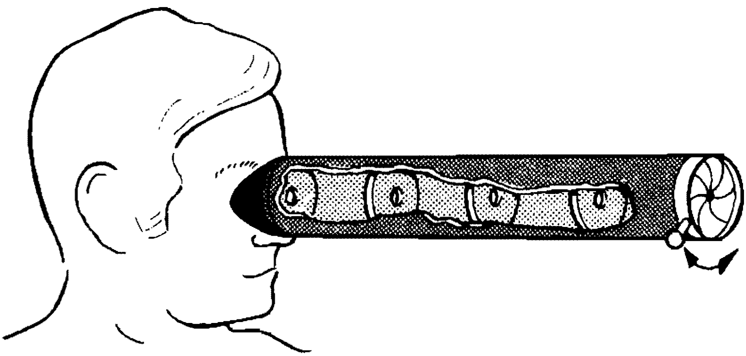

Fig. 7. A black tube.

the areas that send identical wave-length distributions to the eye. As soon as the areas are viewed through the tubes they are identical. As soon as the areas are part of the complex scene they are yellow and green.

Setting aside the black tubes, we can continue to perform the Mondrian experiment with other areas such as blue, grey, brown, red, and so forth. When the variable transformers are changed to send 64, 31, and 5·2 $mW.sr^{-1}m^{-2}$ from any area, all of the areas nevertheless continue to look essentially the same colour that they looked before the transformers were readjusted. Dramatically, the retention of the colour sensations is *not* related to the product of reflectance times illumination, although this product appears to be the only information reaching the eye from each point in the field of view. Therefore, we conclude that the colour sensation at a point in the display has an arbitrary relation to the composition of light at that point. The mystery, then, is how we can all agree wth such precision about blacks, whites, greys, reds, greens, browns or yellows, when there is no obvious physical quantity at a point with which we can specify the colour of an object.

When the triangle and the circle on the Black and White Mondrian were examined it was noted that each of these two papers had a triplet of lightnesses associated with it. Let us look at the two coloured Mondrians. We examine the yellow and green areas first in long-wave, then in middle-wave, then in short-wave illumination. Once again we set the voltages on the two sets of projectors so that the energy from the green paper on the left-hand Mondrian and that from the yellow paper on the right-hand Mondrian are identical. In long-wave illumination the green paper is dark and the yellow paper light. In middle-wave illumination the green paper is medium-light and the yellow paper is light. In short-wave illumination the green paper is dark and the yellow paper also dark. Thus the green paper is characterised by a triplet of lightnesses: dark—medium-light—dark, and the yellow paper by light—light—dark. For each coloured paper in the Mondrian we can similarly find a triplet of lightnesses which characterises that paper, irrespective of the illumination falling on it.

CORRESPONDENCE OF COLOUR SENSATIONS WITH THREE REFLECTANCES

It will be recalled that in the white-grey-black world we found that the lightness corresponded with the reflectance of the objects. The next experiment tests the hypothesis that colour sensation also corresponds with reflectance: in the case of colour, with three reflectances of the object in three wavebands. This experiment has several parts. First, the papers in the Colour Mondrian are measured to determine the triplet of reflectances for each paper. Second, the observer matches each colour in the Mondrian to a colour chip in a standard catalogue of colour sensations. This matching procedure quantifies each sensation. Third, the reflectance number of each paper and the sensation designation of its matching chip are converted to some common scale in order to compare their values.

The first part of the experiment is to measure the triplet of reflectances of each of the papers in the Mondrian. The three reflectances must be measured using the same spectral sensitivity functions that the eye uses. In other words, the photometer that measures the integrated reflectance for long-wave light must have the same spectral sensitivity as that of the long-wave visual pigment in the eye and must integrate over the bandwidth of that pigment. For simplicity we shall refer to all reflectance measurements using the eye's spectral sensitivity curves as integrated reflectance.

Many techniques have been used to measure the sensitivity curves of the visual pigments. The curves we use are the most recent available to us, Paul Brown's measurements (unpublished) of foveal cones, which were obtained by the technique described in the Brown and Wald (1963) paper. These long-, middle-, and short-wave absorption curves are multiplied by the transmission of the eye and the absorption of the macular pigment (Wyszecki and Stiles, 1967) to compute an equivalent pigment sensitivity in the intact eye. We calculated the filter combinations that would alter the spectral sensitivity of a Gamma Scientific S-11 Photomultiplier surface to approximate the spectral sensitivity of the corrected cone pigments. The results are shown in Fig. 8. The best fit for the long-wave pigment curve was provided by the combination of the photomultiplier and Wratten 8 and 106 filters. The middle-wave pigment was matched by the photomultiplier and Wratten 8 and 102 filters and the short-wave pigment by the photomultiplier combined with Wratten 47 and 86A filters.

With these photomultiplier–filter combinations, we are now able to characterise each piece of paper in the Mondrian, whatever the illumination conditions, by a triplet of integrated reflectances measured with detectors whose spectral sensitivities are matched to those of the human visual system.

The second part of the experiment is concerned with the quantifying of each observer sensation. The technique used is based on defining a carefully controlled situation in which there is a unique relationship between the wavelength-energy distribution coming to the eye, the reflectance of the object, and the sensation produced

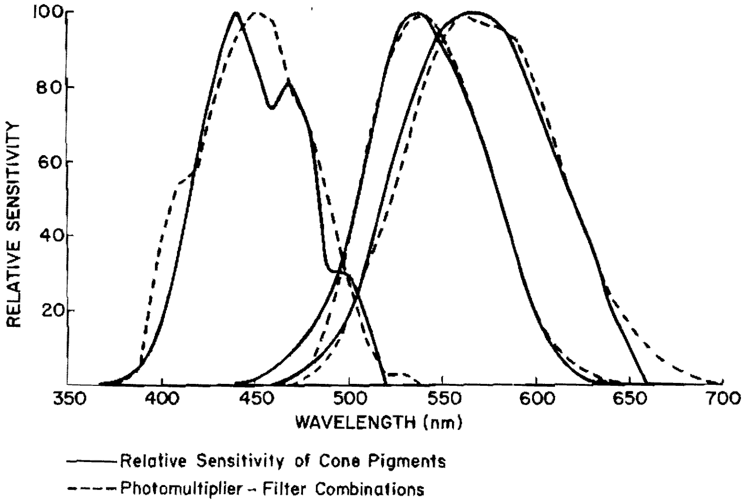

Fig. 8. Graph showing the normalised spectral sensitivity of the three photomultiplier-filter combinations used to match the spectral sensitivities of visual pigments in the retina. The photomultiplier-filter combinations are shown with the dotted lines. The data from Paul Brown's normalised measurements multiplied by the transmission of the eye and macular pigment are shown by the solid lines.

by that object. The controlled conditions are as follows: the illumination must be constant, as must the surround around each area and the state of adaptation. Under these conditions there is a unique sensation associated with each particular colour chip and this sensation can be used to match sensations produced by various wavelength-energy distributions, thus quantifying the second sensation by using quantities measured from the first. Regardless of the conditions producing a colour sensation, as long as it can be matched by a chip in the standard catalogue, it can be quantified by the catalogue specifications for that chip. The observer was given the Munsell Book of Color as the catalogue of colour sensation and was asked to look at a Mondrian in a variety of illuminations and to match each area in the Mondrian to an area in his standard catalogue (see Fig. 9, colour plate).

The third part of the experiment relates the external world of physical measurements to the internal world of sensation. In the first part of the experiment we measured the integrated reflectances of Mondrian areas as a metric of the physical properties of objects in the external world. In the second part we picked colour chips in a controlled environment which matched areas in the Mondrian and thus quantified psychophysical properties of the internal world—sensations. The third and final part of the experiment is to convert the physical measurements and the sensation matches to the same dimension, for example, integrated reflectance. The measurements of the Mondrian are already in terms of integrated reflectance so that we need only measure the integrated reflectance of the matching Munsell chip to have external physical properties and sensations expressed in the same dimension.

As will be appreciated, we cannot conduct this experiment here tonight so let me describe in detail how in our laboratories we carried out this programme to test how good the visual system is as a reflectance measuring device. McCann, McKee and Taylor (in progress) measured the reflectances, integrated under Brown's curves, of various areas in the Mondrian and then had observers choose from the catalogue of Munsell chips the sensations which matched each of those areas in the Mondrian when the Mondrian was illuminated with a variety of widely differing illuminations. In spite of the variegated illumination of the Mondrian, the three integrated reflectances of the chip that the observers chose to match a given area in the Mondrian proved to have a very high degree of correlation with the three integrated reflectances of the area in the Mondrian. These experiments used five different areas in a simplified Mondrian containing 17 areas in all (see Fig. 9).

McCann and his co-workers began with what they called the 'grey experiment', in which they measured the long-, middle- and short-wave energy coming from a grey area. They then performed the 'red experiment', in which a *red* area sent to the observer's eye the same wavelength energy distribution as originally came from the grey area in the 'grey experiment'. Similarly they performed the 'green experiment', the 'blue experiment' and the 'yellow experiment' in which a green, blue, or yellow area sent to the observer's eye the same wavelength-energy distribution as originally came from the grey area in the 'grey experiment'. They duplicated the Colour Mondrian experiment that we demonstrated tonight except for a few procedural differences. First, they used narrow-band instead of broad-band interference filters to illuminate both the Mondrian and the Munsell Book and a much smaller and simpler Mondrian (see right half of Fig. 9). Second, they had the observer match each area in the Mondrian in each of the five experiments to one of the chips in the Munsell Book to quantify each sensation in the manner already described. Third, using the photomultiplier-filter combinations described earlier, they measured the integrated reflectance of each area in the Mondrian in each 'experiment' to find the triplet of integrated reflectances. The triplet of integrated reflectances for each paper was determined by dividing the integrated energy from a paper by the integrated energy reflected from a standard white paper placed on top of the Mondrian. Since the wavebands of integration are broad, in order to duplicate the breadth of sensitivity of the human pigments, each waveband intercepts more than one of the narrow bands of illumination. If we consider the components of each integrated reflectance measurement, we realise that the integrated reflectance of a piece of paper is subject to small changes due to illumination. Integrated reflectance is the ratio of two integrated energy readings, namely, the ratio of the energy from the coloured paper to the energy from the standard white paper. Each integrated energy reading is the sum of three products. Each product is the spectral sensitivity of the meter at the wavelength of one of the narrow-band illuminators times the reflectance of the paper at that wavelength times the amount of illumination at that wavelength. Increases in the illumination term of any one of the products increase the value of that product and hence its contribution to the sum. This change of the weighting of the components of the sum changes the integrated reflectance. It was for this reason that it was necessary to measure the three *integrated* reflectances for each area in each 'experiment'. These measurements make possible a direct comparison of the three integrated reflectances of an area in the Mondrian with the three integrated reflectances of the chip in the Munsell Book which the observer chose to macth it.

A graph of reflectance versus reflectance has the disadvantage that it is not uniformly spaced in terms of observers' sensations. The difference in lightness between a 90/100 and an 80/100 reflector is much smaller than the difference between a 15/100 and 5/100 reflector. Earlier a curve of equally spaced lightnesses as a function of reflectance was described (Fig. 5). Reflectance can be transformed into reflectance-scaled-by-lightness with this curve and a comparison can then be made between the scaled reflectances in the Mondrian and those in the Munsell Book.

The top graph in Fig. 10 shows the results of this comparison for long-wave reflectances. On the horizontal axis the scaled integrated reflectance of each area of the Mondrian is plotted; and on the vertical axis the scaled integrated reflectance of the chip in the Munsell Book chosen by the observer to match the sensations generated by that Mondrian area. The five 'experiments' are identified by different symbols: grey-N, green-G, blue-B, yellow-Y, and red-R. Figure 10 also has similar graphs for middle-wave and short-wave integrated reflectances. We conclude that the triplet of integrated reflectances is a very good predictor of colour sensations.[1]

The success of the Reflectance hypothesis, as compared with the Energy hypothesis, can be demonstrated by plotting a similar graph comparing the energy from an area in the Mondrian with the energy from the observer's choice of chip to match that Mondrian area (Fig. 11). These graphs show that a triplet of energies from a point cannot be used to predict the Munsell chip that the observer will choose as a match. It will be recalled that there are five areas, one for each 'experiment', that sent to the observer's eye a single wavelength-energy distribution. These areas were the grey in the

[1] If, in each waveband, we consider separately the scaled integrated reflectances for each of the five experiments, we observe that there are small systematic differences between the data from the five 'experiments'. For example in the middle-wave graph the R's from the 'red experiment' are systematically slightly higher than the other values. In order to set up the 'red experiment' we had to increase the middle-wave illumination. Just as we observed in the Black and White Mondrian (Experiment I), increases in illumination produce very slight increases in lightness. The small systematic differences found in the five colour 'experiments' correspond to the changes in overall illumination for each of the 'experiments'. The largest average effect on integrated reflectance due to change in overall illumination was less than 10% (McCann, McKee and Taylor, in progress).

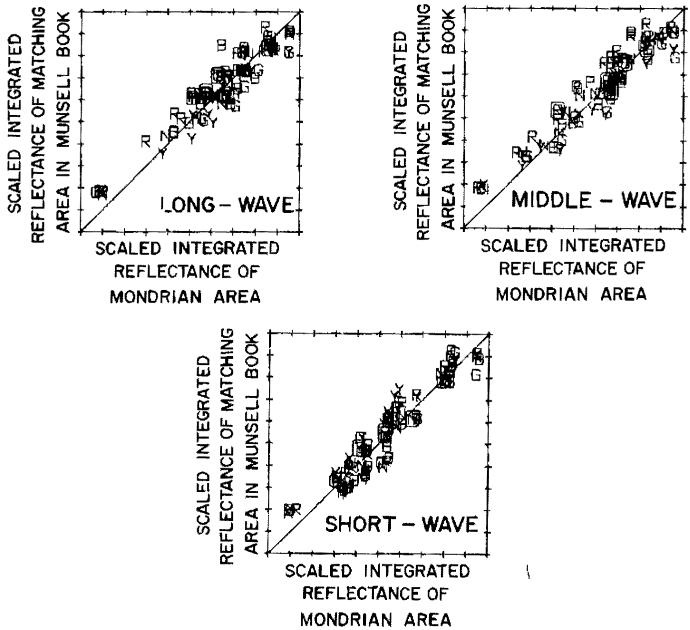

Fig. 10. Comparison of the scaled integrated reflectances of the areas in the Colour Mondrian with the scaled integrated reflectances of the Munsell colour chips chosen by observers to match the areas in the Mondrian. 'Scaled' refers to the fact that the reflectances have been calculated using the curve shown in Fig. 5. 'Integrated' refers to the weighting of the wavelength-energy distribution by the long-, middle-, and short-wave visual pigments and subsequent integration over wavelength. Each letter in a graph corresponds to an area in the Colour Mondrian for the 'experiment' specified by the letter, for example, *all* Mondrian areas in the 'red experiment' are designated by R. See text for further details. The diagonal line from lower left to upper right is where the external measurement, reflectance, is a perfect predictor of sensation.

'grey experiment', the red in the 'red experiment', etc. These areas are identified in Fig. 11 by circled letters. The middle-wave-graph shows that all five Mondrian areas fall on the same value of the horizontal axis because the same energy is coming from each of those Mondrian areas in different 'experiments'. It can also be seen that the chip that matched the yellow area in the 'yellow experiment' reflected much more middle-wave energy than the yellow area. In contrast, the chip that matched the red area in the 'red experiment' reflected

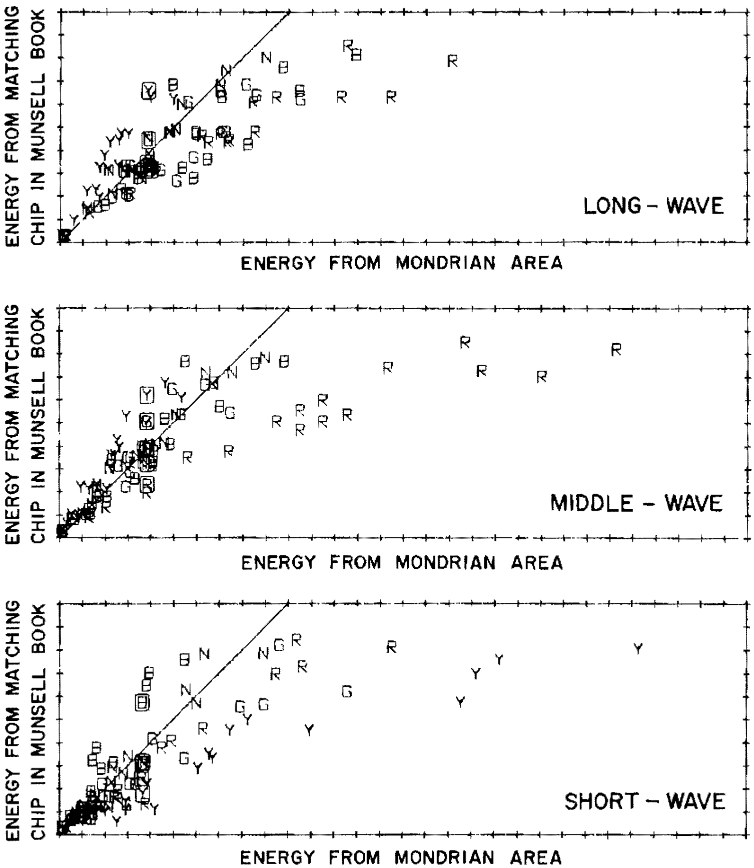

Fig. 11. Comparison of energy coming to the eye from an area in the Colour Mondrian with the energy coming to the eye from the Munsell colour chip selected to match that area of the Mondrian. Each letter in a graph corresponds to an area in the Colour Mondrian for the 'experiment' specified by the letter, for example, *all* Mondrian areas in the 'red experiment' are designated by R. Each of the circled letters corresponds to one of five different areas which is made to reflect a fixed wavelength-energy distribution in the 'experiment' specified by the letter. The diagonal line is where energy from the Mondrian equals energy from the matched Munsell chip. This graph shows that a triplet of energies from a point *cannot* be used to predict the Munsell chip that the observer will choose as a match.

much less middle-wave energy than the red area. In addition, the blue area in the 'blue experiment' and the green area in the 'green experiment' were matched by chips that reflected still different amounts of middle-wave energy. Just as in middle-wave light, in long- and short-wave light the observer also chose chips that reflect a variety of energies to match Mondrian areas that sent a single energy to the eye.

The variety of chips chosen to match the five different areas that sent to the eye the same wavelength-energy distribution illustrates most dramatically the magnitude of the inadequacy of the Energy hypothesis. In Fig. 12 (colour plate) there are printed samples of the five chips chosen by the observer. In fact, in the actual experiment the differences in colour between these samples are even greater because they were viewed in narrow-band illumination which enhances the saturation of these colours.

THE RETINEX

Earlier we proposed a hypothesis, namely that every colour is determined by a unique triplet of lightnesses, the correlates of integrated reflectances. The results in Fig. 10 showed that this hypothesis is correct. This leads us to a further hypothesis, namely that the information from each of several types of receptor in the retina is processed independently of the information from the remaining types. Since we have found that colour sensations correspond with triplets of integrated reflectance, an acceptable theory must contain a mechanism for generating the lightnesses which correspond to the integrated reflectances without the use of external reflectance standards. If we make the simple and fundamental assumption that the long-wave information is independent of the middle- and short-wave information we can sequentially intercompare all the long-wave energies coming to each point in the image and generate reflectance without an external standard as was done in the black, grey, and white experiments described earlier. Similarly, we can intercompare the middle-wave information to form a middle-wave lightness for each area and also intercompare the short-wave information to form short-wave lightnesses. We should state emphatically that if we take the ratio of long-, middle- and short-wave energy at each point, we lose the information necessary to generate the triplet of integrated reflectances and their correlated lightnesses.

The comparison of lightnesses is the determinant of colour. The formation of these lightnesses and their comparison could occur in the retina or in the cortex. Experiments in visual physiology cannot as yet define the location of the interactions that must be occurring. Therefore, I coined the word Retinex (made up of 'retina' and 'cortex') to designate the physiological mechanisms that generate these mathematically independent images. My proposal did not demand that the retinal elements with the same peak sensitivity have to be *directly* connected to each other. Instead, somewhere in the retinal-cerebral structure, elements associated with the same wavelength characteristics co-operate to form independent images in terms of lightness (Land, 1964).

RETINEX RECORDS

We have devised a technique for determining the apparent lightness of objects *as they appear on each independent Retinex*. As has been seen from Brown's curves, the spectral sensitivities of the visual pigments overlap broadly. If a scene were illuminated with the entire range of wavelengths to which a single pigment is sensitive, a large variety of colours would be seen because more than one Retinex would respond. Photographic intermediates called Retinex records permit the isolation of lightnesses that would ordinarily be combined into colour sensations. The technique for seeing long-wave lightnesses is to take photographs of a scene, such as the simplified Mondrian, with a black-and-white film and filter combination. The effective spectral sensitivity of the combination is the same as that of the long-wave visual pigment in the eye. Similarly, film-filter combinations can be found to match the spectral sensitivity of the middle- and short-wave visual pigments. It is the human visual system which converts the energy reflected from the photographic deposit in silver into lightness. Ideally, we would like our observer to examine the resulting black-and-white images with only one set of cones, reporting the lightnesses appropriate to that cone set. However, at any point in the silver pattern the reflectance is essentially the same throughout the visible spectrum. Therefore, with a black-and-white photograph we stimulate all the types of receptors with the same information, in this case the record of energies originally absorbed by a single visual pigment. If we assume that all the Retinexes process the information in an identical manner, then we can propose that sending this identical information to several sets of receptors would yield a pattern of lightnesses like the pattern one would experience from a single Retinex if it were possible to exclude the other two.

We now can see the kind of image that each of our Retinexes sees—whereas, we had not been able to do this without these separate photographs. We can predict the colour sensations by the comparison of the lightnesses we see when viewing these long-, middle-, and short-wave Retinex records.

Four photographs are shown in Fig. 13 (colour plate). The first is a colour reproduction of the simplified Mondrian. The other three are long-, middle-, and short-wave Retinex records of the same subject. Inspecting each of these three photographs with the human visual system generates on all three Retinexes the lightness pattern that would ordinarily be on one Retinex and unavailable.

A COMPUTER MODEL FOR COLOUR

We wished to test whether the quantitative predictions of the Retinex model matched the colour experience of an observer viewing the Mondrian experiment.

The first step in the process is to measure the *energy* of each point in the simplified Mondrian in the long-, middle- and short-wave-bands. For each waveband, the Mondrian is characterised by 480 energies spaced regularly in a 24×20 array. The computer does not know the position of boundaries of areas, just the energy at all points. We chose to use uni-directional paths of length 200 for all three Retinexes. The origin of the path in the 480 point array and the direction of each path is determined by a random number generator. Each path travels straight in the direction from which it came until it reaches the edge of the display where it could either 'reflect' from the outside boundaries or could travel along them. At each point along the path the energy at the point was divided by the energy at the previous point. The ratio was tested to see whether the difference from 1·0 was significant. If it was, then it was multiplied by the sequential product from the previous point. This sequential product was tested to see if it was *larger* than 1·0. If so, it was set equal to 1·0, initiating the new standard. If not it was sent on unchanged. This sequential product was used twice; first it was held to be averaged with all the other outputs that had reached this point in the geography, and second it was sent on to be multiplied into the next sequential product. The sum of the sequential products at each point was then divided by the number of times paths went through that point. This average sequential product is the reflectance calculation or prediction of the model. The triplet of reflectance calculations, one for each waveband, is the prediction of the colour of each area.

We can more readily compare the accuracy of the computer's reflectance calculations by separately comparing long-wave computed lightness against long-wave observed lightness, etc. This comparison is done with graphs similar to those shown in Fig. 10. Figure 14 shows the long-, middle- and short-wave plots of computed

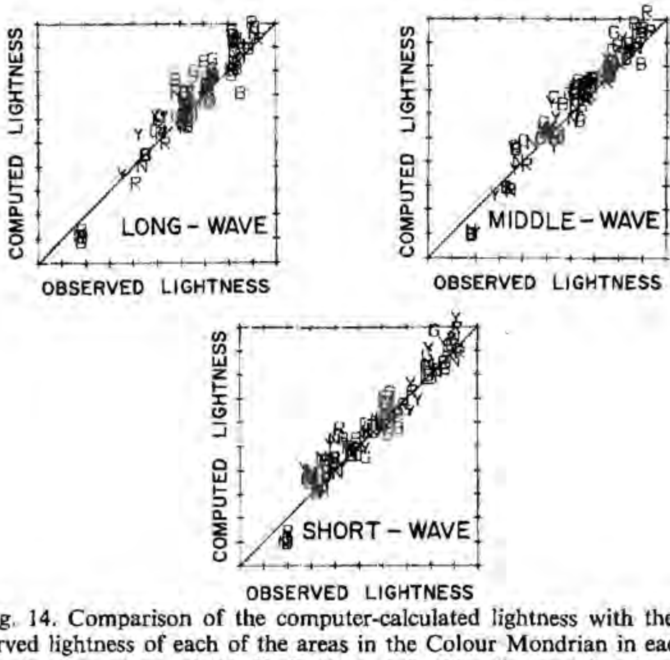

Fig. 14. Comparison of the computer-calculated lightness with the observed lightness of each of the areas in the Colour Mondrian in each of the five 'experiments', N, G, B, Y, and R. The diagonal line is where computer lightness equals observed lightness. These graphs show that the computer program, which is designed around the conceptualised internal procedures, yields as good a fit of computed lightness to observed lightness as the very good fit obtained when reflectance itself is compared with observed lightness (Fig. 10).

versus observed lightness. These graphs show that the computer program, which is designed around the conceptualised internal procedures, yields as good a fit of computed lightness to observed lightness as the very good fit obtained when reflectance itself is compared with observed lightness (Fig. 10).

Thus we have come from the description of a primative property reflectance, via a series of experiments to a consistent theory of colour vision. We have shown that the eye derives reflectance without an external standard; that reflectance is read by an internal meter, the pointer of which is the sensation lightness; and that lightness depends on reflectance, *not* on energy. We have shown that the flux from a patch of one colour can comprise the same ratio of energies as the flux from a patch of an entirely different colour—indeed the second flux may be in every respect identical to the first. We have shown that colour is the function of three lightnesses.

HISTORICAL REVIEW

We can now turn to a brief historical review. Newton stated that 'Every Body reflects the Rays of its own Colour more copiously than the rest, and from their excess and predominance in the reflected Light has its Colour' (Newton, 1730). Today we might say that Newton believed that the wavelength-energy distribution determined colour sensations. Thomas Young (1802), while teaching at the Royal Institution, proposed the next idea in colour vision, namely that there were three kinds of receptors with three different spectral sensitivities and that they worked together at a point. Young's theory might be regarded as a triple Energy hypothesis: the amount of light in each waveband coming from an object determines the response of the different light detectors; the relative response of the three detectors at a point determines colour at that point. The first idea is correct: there are three kinds of receptors, as shown directly by Brown and Wald (1963) and by Marks, Dobelle and MacNichol (1964). The second idea is incorrect: these receptors do not simply act together at a point. As we showed earlier, this simple interaction would lose information that is necessary to form colour sensations.

James Clerk Maxwell (1855) performed experiments that showed that Young's idea that the eye must have three kinds of light-sensitive elements was correct. Nevertheless, his experiments did not reveal to him the paradox that sensations correspond to the reflectances of objects and not to the response of Young's light detectors at each point. Along with Maxwell, Hermann von Helmholtz (1866) helped establish Young's ideas as a basis for colour vision. Helmholtz held firm to the idea that three energies at a point determine colour at that point. At the same time Helmholtz's observations persuaded him that: 'we are accustomed and trained to form a judgement of colours of bodies by eliminating the different brightness of illumination under which we view them' and 'we eliminate the colour of the illumination as well'. Helmholtz made the valid observation that we see the correct colours of objects in spite of variable illumination, but he was mistaken in his view that 'we are accustomed and trained to form a judgement of colours' and that 'we get accustomed to subtracting the illuminating colour from coloured surfaces'.

In a paper delivered to the Royal Institution in April, 1961, I reported on our studies of the effects of memory, surround and highlights as sources of information about colour:

> 'We also find that the colour is not determined by the immediate surround. A lemon placed first against a red background, and second against a cool background, does not change its colour significantly. Neither is the colour determined by the average of the total amount of light on each wavelength band: if it were, a change in the background which covers a large portion of the slide would change all the colours over the picture. Neither is colour determined by reference to highlights, white objects, or other evidence of the source of illumination—if these are all removed, the colour remains. Neither is colour determined from memory: oranges may be painted many different hues and the only ones which appear orange will be those falling in the correct area of the coordinate system.'

The Mondrian experiment which we performed here tonight was designed to remove the observation of colour from the realms of judgement, learning and knowledge of the illumination. Memory colours are eliminated by the abstract nature of the rectangles in the display; and clues to the colour of the illumination are removed by having matte papers which reflect a constant amount in all directions, and/or by viewing between crossed polarisers. It is the Mondrian experiment that brings into intellectual focus the fact that colour sensations are essentially independent of energy and dependent on reflectance. In the Young–Maxwell–Helmholtz approach the interaction is between a member of one set of receptors and a member of a second set and a member of a third set, the three members being at one point on the retina.

In the Retinex theory the only interactions required are within a set; the members of a given set interact with each other. When the members of a given set interact with each other there comes into being a scale of lightnesses produced by that set. Every area in a scene will be characterised by three lightnesses (one for each set) and the colour experience in viewing an area is determined by the comparison of these three lightnesses.

Friday Evening Discourse, 2nd November, 1973

Edwin H. Land was born in 1909, and educated at Harvard University. Major contributions in the fields of polarised light, photography, and colour vision. While at Harvard he invented a new polariser for light in the form of an extensive synthetic sheet. In 1947 he described a camera and film for one-step photography, the evolution of which was reported to the Royal Photographic Society in May, 1973. For the last twenty-five years he has studied colour vision. His technological work is described in about 400 US Patents. He has received numerous honorary degrees and medals, including the Presidential Medal of Freedom (1963) and the National Medal of Science (1967). Fellow of the National Academy of Sciences; Past President of the American Academy of Arts and Sciences. He founded the Polaroid Corporation in 1937 where he is President, Chairman of the Board, and Director of Research.

REFERENCES

Brown, P. K. and Wald, G. (1963). Visual pigments in human and monkey retinas, *Nature*, October 5, pp. 37–43.

Helmholtz, H. von (1866). *Physiological Optics*, Vol. II (ed. J. P. C. Southall, Optical Society of America, 1924).

Land, E. H. (1962). Colour in the natural image, *Proceedings of the Royal Institution of Great Britain*, **39** (176).

Land, E. H. (1964). The Retinex, *American Scientist*, June, pp. 247–64.

Land, E. H. and McCann, J. J. (1971). Lightness and Retinex theory, *Journal of the Optical Society of America*, January, pp. 1–11.

Marks, W. B., Dobelle, W. H., and MacNichol, Jr., E. F. (1964). Visual pigments of single primate cones, *Science*, March 13, pp. 1181–1182.

Maxwell, J. C. (1855). *Scientific Papers of James Clerk Maxwell* (ed. W. D. Niven, Dover Publications, Inc., New York, 1965).
Newhall, S. M., Nickerson, D., and Judd, D. B. (1943). Final report of the OSA sub-committee on the spacing of the Munsell Colors, *Journal of the Optical Society of America*, July, pp. 385–418.
Newton, I. (1730). *Opticks*, 4th Ed., Dover Publications, Inc., New York, 1952.
Wyszecki, G. and Stiles, W. S. (1967). *Color Science*, John Wiley & Sons, New York, London, and Sydney.
Young, T. (1802). *Philosophical Transactions*, **92**, pp. 12–48.

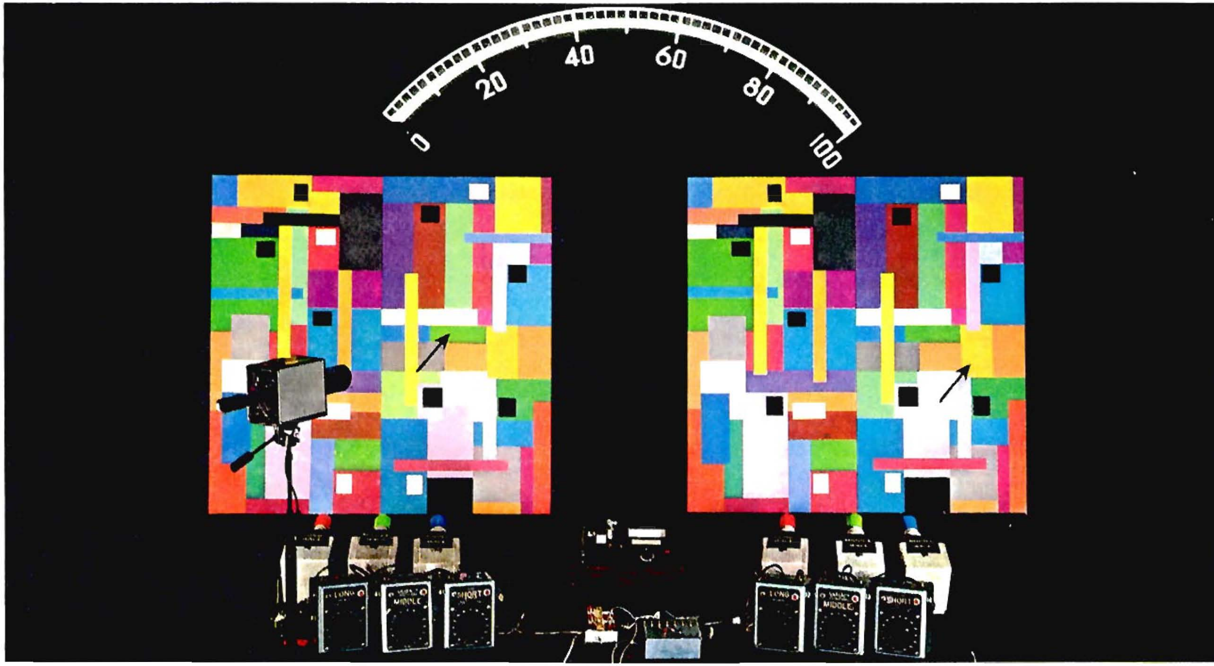

Fig. 6 is a photograph of the equipment in the double Mondrian experiment; it is not intended to reproduce the experiment. Two identical Color Mondrians were each illuminated with three band-pass illuminators with independent brightness controls. Readings from the telescopic photometer on the left were projected on the scale above the Mondrians. In the actual experiment, the three illuminators on the left were set so that the green area indicated by the arrow sent to the eye the same radiation as was sent from the yellow area indicated by the arrow on the right. The green area continued to look green and the yellow area continued to look yellow.

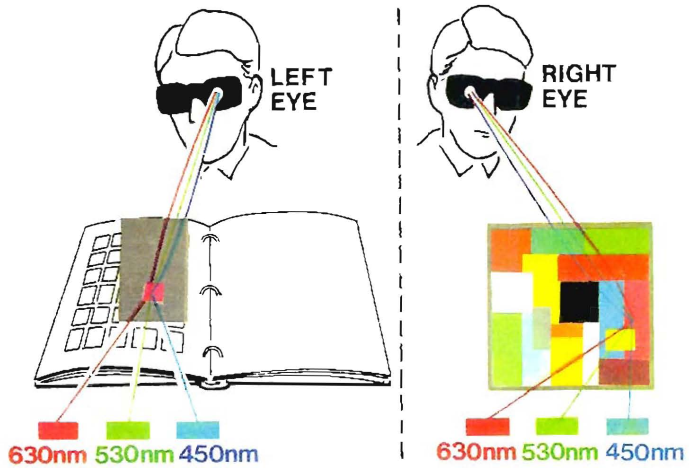

Fig. 9 shows schematically the procedure for quantifying color sensations. The subject sees only the Mondrian with one eye and only the Munsell Book of Color with the other eye. The Munsell color chips are viewed through a hole in a uniform gray piece of paper so that any surrounding effects are constant. The subject's task is to choose the color chip which matches each color area in the Mondrian.

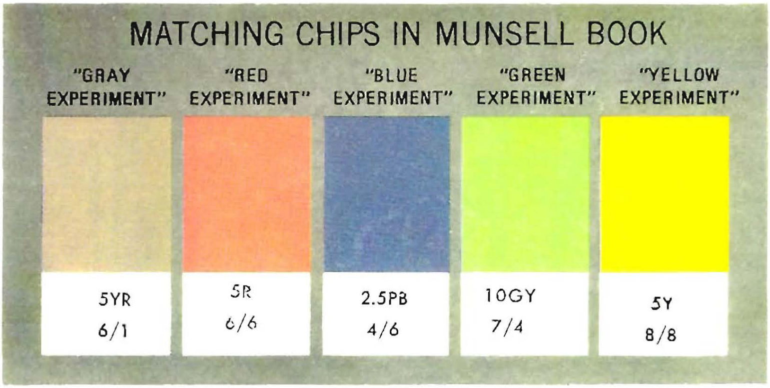

Fig. 12 illustrates the variety of color sensations generated by a single wavelength-energy distribution. The five colors shown are reproductions of the Munsell colors chosen by observers to match the areas which sent identical radiation to the eye. In fact, in the actual experiment the differences in color between these samples are much greater because they were viewed in narrow-band illumination which enhances the saturation of these colors.

A COLOR MONDRIAN

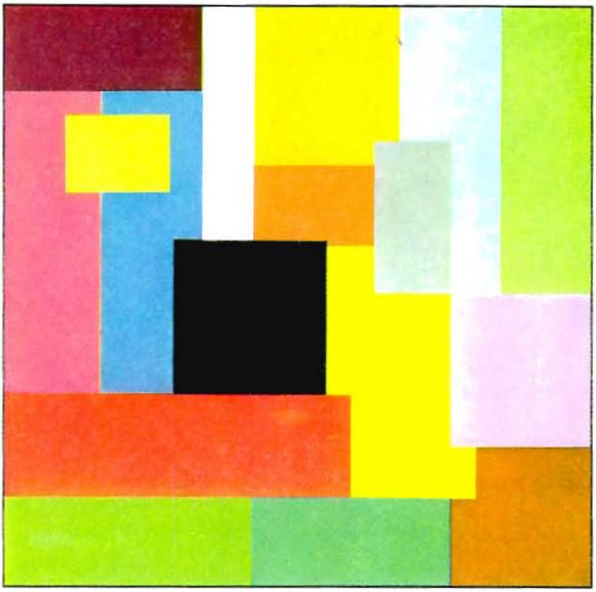

AND ITS RETINEX RECORDS

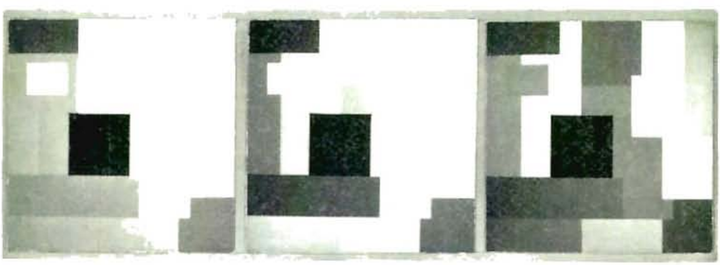

LONG WAVE   MIDDLE WAVE   SHORT WAVE

Fig. 13 contains a color reproduction of the simplified Mondrian and long-, middle-, and short-wave Retinex records of the same subject. A Retinex record is a photograph taken with a film-filter combination having the net sensitivity of one of the pigments of the human eye. These photographic intermediates enable us to isolate lightnesses that would ordinarily be combined into color sensations. Inspecting these three photographs with the human visual system generates the lightnesses that are otherwise unavailable.

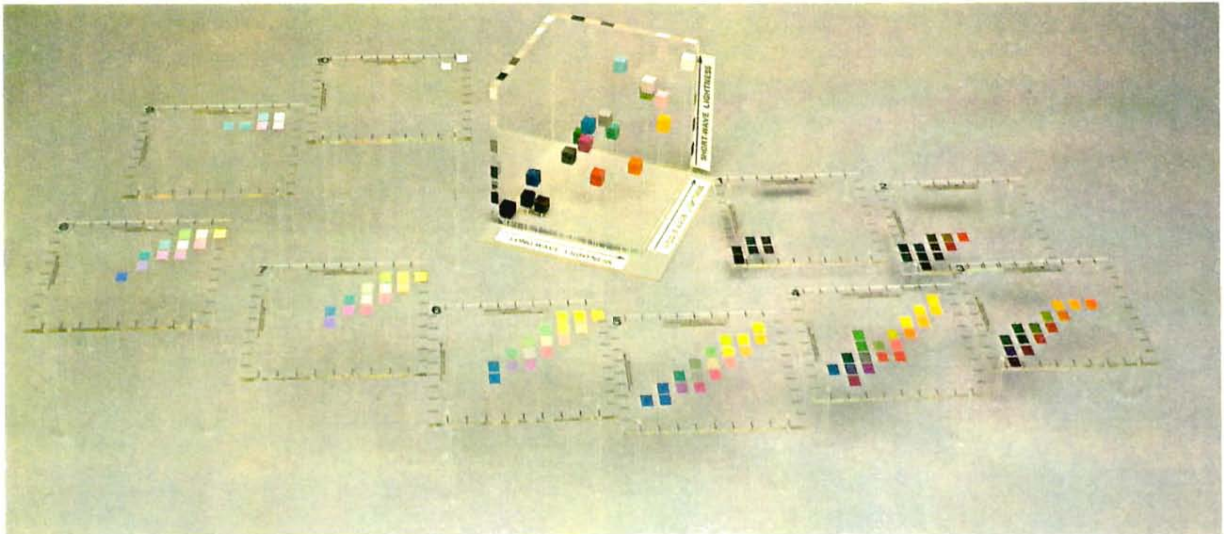

Fig. 15 is a photograph of the model of the three-dimensional color space exhibited in the library. The position of a color in this space is not determined by a triplet of energies at a point, but by the three lightnesses of each area. The three-dimensional model in the top center of the photograph contains representative colors from throughout the space. This composite part of the model is surrounded by ten horizontal planes cut through the three-dimensional space. Each plane is the locus of colors possible with a constant short-wave lightness; for example, the lightness of 5 plane shows the variety of color sensations from all possible long- and middle-wave lightness combinations and a short-wave lightness of 5. Each plane is suspended at a height that corresponds to its height in the composite model. In each plane, the lightness on the middle-wave Retinex is plotted along the edge perpendicular to the bottom of the photograph and the lightness on the long-wave Retinex is plotted parallel to the bottom of the photograph.