CHAPTER 13

Edwin Land's colour-constancy demonstration

Introductory

As explained earlier, a key event in my life was the encounter with Professor Marian Bohusz-Szyszko. The ideas he shared set me off on a lifelong journey of discovery. My first step was to set about testing his seemingly extravagant assertion that it is only necessary to follow two rules to guaruntee a good painting:

- *There must be no repetition of colour on the same picture surface.*
- *All the colours used must be mixtures containing at least a trace of complementary.*

After four years of experimenting, I proved, at least to my own satisfaction, that there is a special quality in all paintings that abide by these two rules. It is difficult to describe, but it involves the creation of a sense of pictorial space and harmony.

Fortunately, a troubling paradox arose that would eventually have a profound effect on the development of the ideas presented in this book. It concerned the Professor's physics-based proof of the invariable variability of colours in nature. This asserted that no two parts of any surface will reflect exactly the same wavelength combinations into our eyes due to:

- *The complexity produced by the inter-reflecting surfaces*
- *Variations in viewing angles and distances*
- *Atmospheric filtering*

The paradox is that, if the light reflecting from two parts of a surface can never be characterised by the same wavlength combination, how could artists repeat colours on a picture surface? Even if two regions were painted with exactly the same pigment-colour, how could these appear as the same?

Other people might already have known how to resolve this mystery, but for many years I had no idea how to do so. My first inkling of a solution came after *many years, as a result of reading a paper by Edwin Land,*[1] *the inventor of the Polaroid camera. In it was a powerful demonstration of the phenomenon of "colour constancy" and an attempt to explaining it. What the demonstration showed was a region of colour within a multicoloured display (henceforth referred to as the MCD) being perceived as remaining the same, even when the experimenter changed the combination of wavelengths being reflected from it. I was excited because here were two colours being perceived as the same despite reflecting different wavelength combinations into they eyes? It was a eurika moment. However a big problem emerged for it was soon clear to me that the explanation suggested by Land was not neurophysiologically plausible. Another explanation had to be found. I could never have guessed at the treasure trove of discoveries that came out of my struggles to find it. This chapter describes Land's demonstrations in the context of an earlier attempt at explaining colour constancy. The next chapter introduces our neurophysiologically plausible colour constancy algorithm.*

A precursor

Towards the end of the eighteenth century, in 1789, the French mathematician and scientist, Gaspard Monge (1746-1818) gave a demonstration to the French *Académie Royale des Sciences*.[2] It was very simple. He asked his audience to look at piece of red card through a red piece of glass. Much to their surprise the red became "*white*". Monge related this counter-intuitive phenomenon to another in which he was interested, which can be demonstrated as follows. First focus two beams of white light from two separate projectors on the same white-coloured screen. Place a red filter in one of the beams so that it filters out all but the long wavelengths of light. In this *red* beam place a ping pong bat or some other opaque object so that it casts a shadow on the screen. Since the shaded region is illuminated only by the white beam, common sense would indicate that it would appear to be white or grey. But it will not do so. Rather, it will be seen as very pure looking *green*. It is an example of *induced colour.*

Taking these two phenomena together Monge was able to make the imaginative leap required to associate them with a third, now known as *colour-constancy*, and to produce a plausible sounding explanation for it. Incidentally, it was the same as the one Edwin Land came up with nearly two hundred years later

1 Land, E.H., 1979, The Retinex Theory of Colour Vision. Scientific American

2 Described in John Mollon, 1995, Seeing Colour. In: *Colour Art and Science*, Eds.. Lamb and Boreau, C.U.P.

in the *Scienitic American*. Both concluded that the evidence was consistent with the hypothesis that the visual system can compute *ratios* between the colours in a MCD. One important implication of their speculation was that it meant that *colour-constancy* will only occur when colours are seen in the context of other colours within a MCD.

However, there is a problem with the ratio-based hypothesis: Though both fitting the evidence and being mathematically sound, it lacks *neurophysiological plausibility*. Calculating ratios is not the kind of thing that the neural networks used in the first stages of visual processing can do. However, since it had to be these that mediate *colour constancy*, there must be a way of coming to the same conclusion by means of neurophysiologically plausible mathematics. In the next chapter an *algorithm* is described that fits these two requirements. Before describing this, it is appropriate to look, first, at some earlier ideas on the subject of colour constancy and, then, at Land's demonstrations of the phenomena which show just how powerful it is.

Colour-constancy

If there actually were a direct correspondence between wavelength of light and colour, then an analysis of the wavelength composition of the light entering the eye would always predict the colour experience. In some cases, such as those of the rainbow or individually illuminated patches of pigment, such a prediction will be fulfilled. However, in the majority of cases it will not be. Thus, in complex and everyday scenes, the colour of an object tends to stay the same regardless of the light illuminating it. For example, a red apple seen in daylight will continue to appear as being the same colour throughout the day, despite the fact that the conditions of illumination are constantly changing. It even seems to be much the same when viewed under tungsten light. This phenomenon is known as "colour-constancy".

As mentioned above, colour-constancy had already been a subject of interest at the end of the eighteenth century, when Gaspard Monge gave his talk. Various other scientists added their ideas in the nineteenth century. These included Hermann von Helmholtz (1821-1894) and Ewald Hering (1834-1918).

Helmholtz attempted an explanation of colour-constancy pointing out that, if the effect of changes in the source of illumination could be neutralised or discounted, the colours of surfaces would not vary. He further suggested that such a neutralisation could be achieve by shifting what he called the "*achromatic point*" in colour-space to correspond to the average of all the colours in the scene. He thought that this process would require some form of unconscious inference (in modern times this might be translated as *higher level processing*). Significantly, he suggested no mechanism by which this calculation could be accomplished.

The name of Hering is associated with the discovery of *opponent-colour cells* which work on the principle of *lateral inhibition*. In a paper published in 1874 (coincidentally, the same year as the first 'Impressionist' exhibition) he argued that colour-constancy could be explained by the same sensory processes that produce the phenomena of *simultaneous contrast* and *sensory adaptation*. Although he could not offer a complete explanation, he saw that the way forward would require knowledge of neurophysiology.

There is no need to understand the ideas of Helmholtz and Hering in detail. What is important is that these early speculators on the subject of colour-constancy, following in the footsteps of Monge, had already understood that any explanation for the phenomenon would lie in activity within the neural systems of the brain. However, since neurophysiology was in its infancy, they were in no position to check out any speculations about mode of functioning of the neural systems they posited.

During the next hundred years, despite various efforts by scientists, little significant progress was made. Two possible explanations for this tardiness are worth mentioning.

- The first is that mainstream of colour-perception scientists were working within the prevailing tradition of studying colour responses via small patches of colour or light either individually or as adjacent pairs. This means that their experimental design cut out the contextual effects that makes colour constancy possible.
- The second is the low status of colour in visual perception theory. Until very recently and perhaps to this day there has been an overwhelming tendency to assume that edges and shading are the key variables and that colour is an almost expendable, *"last minute"*, evolutionary bolt-on. After all, we can make perfect sense of black and white films.

This being the case, the possibility that the understanding of colour-constancy, and the way it is derived, might be of fundamental importance to the understanding of human visual perception had been largely overlooked.

While the ratio-based proposals of Edwin Land were quickly abandonned, his two experiments were to wake up the scientific world. One reason for their great influence was the clarity of their outcomes. Though it could be argued that they showed nothing new, what the did provide was evidence that could not be gainsayd. The first experiment produced the data upon which his *Retinex Theory* was based. The second, which might more legitimately be called a "*demonstration*", produced unequivocal proof that wavelength and colour should not be confused.

The light 'primaries'

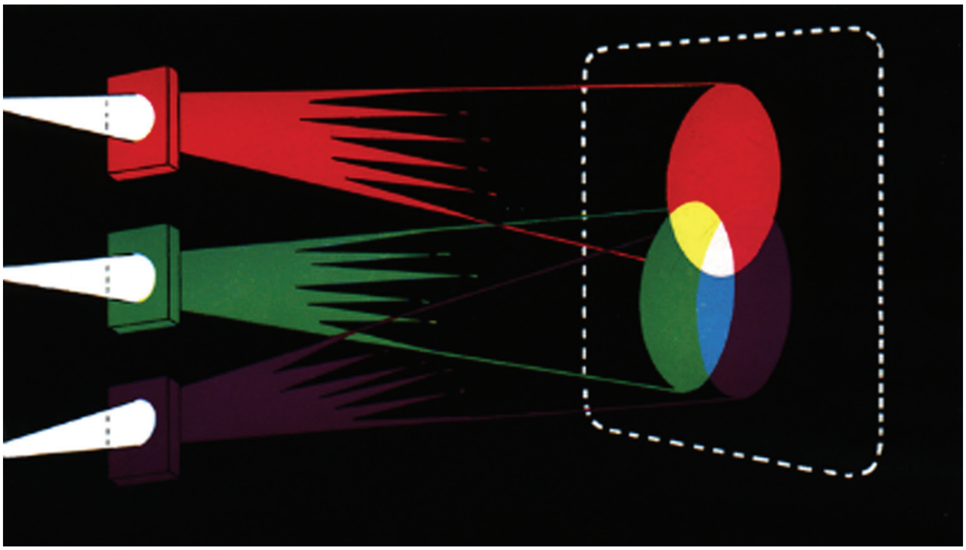

Figure 1 : Three projectors creating overlapping discs of the three primary colours.

From my students I have learnt that a great many people are extremely vague on the subject of the light primaries and their mixtures. For this reason, when I explain Land's demonstration to my students, I start with basic knowledge. *Figure 1* shows three projectors, one with a red filter, one with a blue filter and one with a green filter creating three discs of coloured light on a white painted screen, in a dark room. Colour mixtures occur where they overlap. Thus: blue and red make violet; blue and green make blue-green; and, red and green make yellow. Where all three discs overlap, white is created.

Many people are surprised to find that the three light primaries do not include yellow. To help them get a feel for the situation, I explain that adding two light sources always results in a beam of increased intensity. It follows from this that the most intense or brightest colours will necessarily be mixtures. Thus: white is a mixture of all three primaries, and, therefore, is the brightest. Similarly, violet, blue-green and yellow are brighter than the colours from which they were mixed (henceforth, referred to as "*parent colours*"). In short, parent colours cannot be the brightest. Looked at this way it becomes intuitively understandable that yellow, the experientially the brightest colour, must be a secondary.

Coloured lights making the full gamut of colours

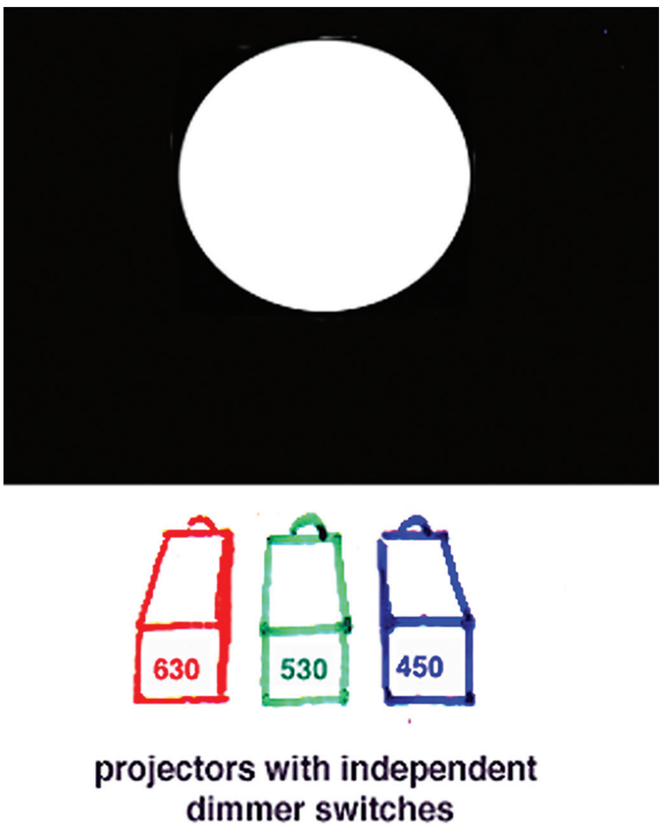

Figure 2 : Projecting three discs of light onto the same region of a white board. a rheostat (dimmer switch) is attached to each projector.

It is easy to see that, if the beams of light from the three projectors in *Figure 1* are moved inwards so that they overlap progressively more and more, the area of white will become larger and larger. When the beams overlap completely, as in *Figure 2,* the effect will be the same as that produced by one projector projecting white light. The result will be as illustrated, namely one disc of white on the screen. With this arrangement it is easy to demonstrate light-mixing in a different way. First, each projector can be switched on independently, thus providing the three light-primaries, red, green and blue. Alternatively, any two of the three of projectors can be switched on (that is to say, blue and red, blue and green and, red and green) producing the light secondaries, violet, blue-green and yellow.

More pertinent to the ideas about to be presented is the fact that the possible number of mixtures can be enormously increased by adding a dimmer switch (rheostat) to each of the projectors. In this way, in principle at least, the intensity of each light-source can be varied infinitely and, if so, an infinite number of wave-length combinations can be projected at the screen, producing a very large number of different colours.

My students often seem to think that this number should be infinitely large, but this reflects the confusion between the physical properties of light and the psychological phenomenon of colour. The students forget that colour is a product of the visual system and that the sensitivity of that system is limited. However, though not infinitely large, the number of different colours to which the eye-brain is sensitive is pretty impressive. Estimates suggest that it is in the low millions.

The Experiment

Land's experiment is very simple. The set-up was as just described, using three projectors and a screen, but with two crucial differences. The first is that the screen was made into an *MCD* by covering it with various sized rectangles of coloured paper (as illustrated in *Figure 3*) and, the second, is that a telescopic photometer (i.e. a light meter that can take readings of small areas of the screen from a distance) has been introduced to provide a means of monitoring changes in lighting conditions.

From the scientific point of view, the arrangement of coloured papers could have been arbitrary. All that was required was a representative range of colours, including the primaries, the secondaries, black and white. However, for idiosyncratic aesthetic reasons, Land took a great deal of trouble over the appearance of his *MCD* and he was proud enough of the result to feel justified in calling it a "*Colour Mondrian*", in homage to the Dutch Constructivist painter of that name. Though the attribution may have made the fastidious Mondrian turn in his grave, the relevance of Land's gesture in the present context follows from the painting-like properties of the *MCD*. Like them, it is an "*arrangement of colours on a flat surface*".[3] Clearly the findings of the experiment, if they turned out to be interesting, would be relevant to the experience of looking at paintings.

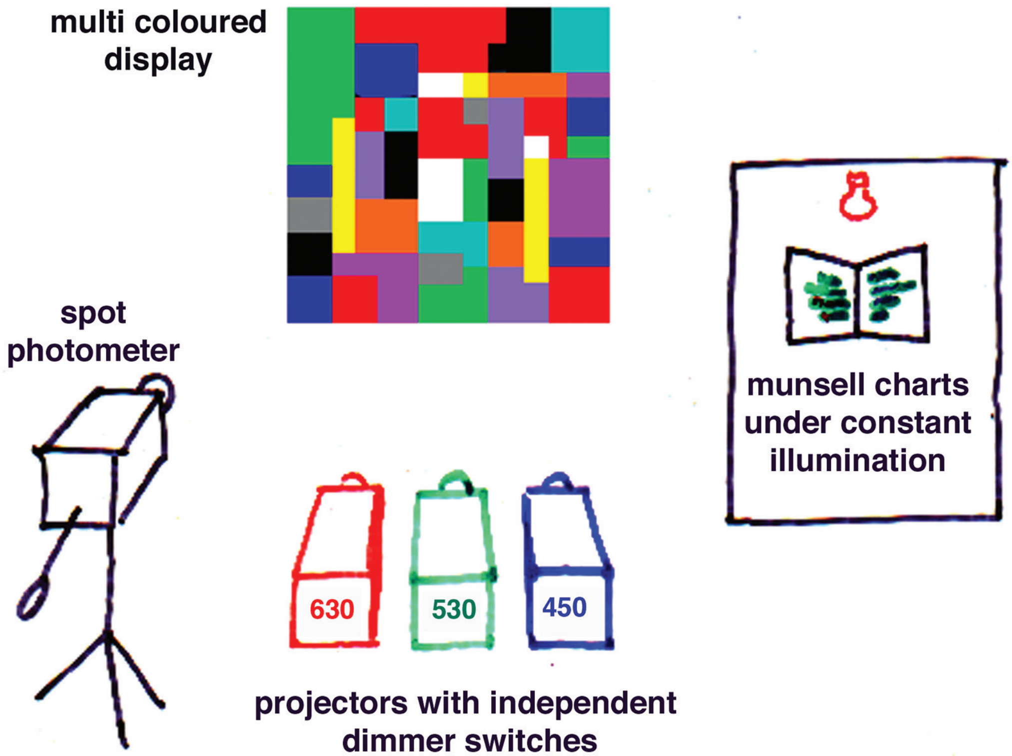

Figure 3 : The set-up for land's first demonstration. The numbers written on the back of the projectors correspond to the wavelength used.

Land's experiment also required a means of measuring what people were seeing. This explains the presence of the Munsell Book of Colour[4] illuminated by a constant white light source. For those who have not come across this book, it is the generally accepted scientific reference for the classification of coloured surfaces. It consists of seventy-two pages of samples (usually referred to as "*colour chips*"), with each page dedicated to a different hue. The chips in any

3 The phrase famously used by Maurice Denis.

4 Munsell Color, 1976, Macbeth Division of Kollmorgen Corporation, Baltimore.

vertical column vary in respect of "*lightness*" (what Munsell calls "*value*") and, in any horizontal column, in terms of "*saturation*" (what Munsell calls "*chroma*"). *Figure 4*, shows a small selection of these chips laid out to form a *"Munsell Color Tree*", according to the scheme found in the *Munsell Book of Colour*.

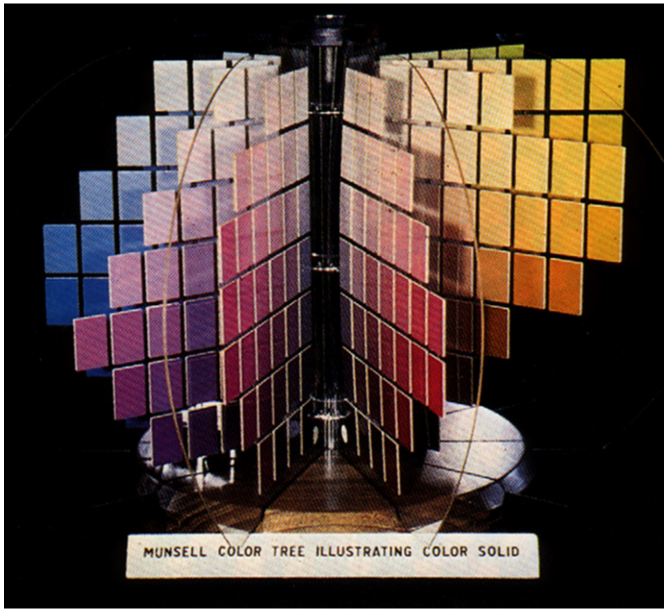

Figure 4: A munsell solid

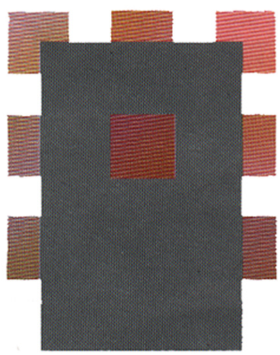

Figure 5: A part of a page from the munsell book of colour, with grey card to surround each colour tried out

The experimental task

For his experiment, Land's experimental subjects were asked to compare colours in the MCD with chips in the Munsell Book of Colour, using one eye to look at the former and the other eye, to look at the latter. Their task was to find the Munsell chip which most closely corresponded to a selected colour in the MCD. After every judgement, the *telescopic photometer* was used to take a measurement of the light reflected from the selected colour when illuminated separately by each of the three coloured light sources. The upshot was three reflectance readings, which Land termed *reflectance-triplets*.

This first trial having been completed, the rheostat settings for the three projectors were changed such that they would produce different combination of photometer readings. When the adjustments had been made, the experimental subjects were asked to repeat the same procedure as before. Thus, they were shown an identical sequence of colours and asked to make another set of matches between the *MCD* and the Munsell chips. In this way, the set of *reflectance triplets* was recorded for each new combination of light sources. Sufficient experimental subjects were used to give statistical validity to the findings.

The main result of the demonstration was extremely simple. As long as all three of the light sources were contributing to the illumination of the MCD, then, no matter what the reflectance triplets, the experimental subjects *matched the same display colour with the same Munsell chip*. In other words, no changes occurred in the appearance of the variously illuminated colours. They were perceived as remaining absolutely constant.

Neutralising painted colours

We are now ready for Land's second and experientially more powerful way of showing the same phenomenon. For this purpose, it will help to explain how a slightly modified version of the equipment just described can be used to neutralise brightly coloured surfaces (in other words, to make them look grey). This takes us back to fundamentals. The green appearance of green paint is due to the fact it absorbs all but the predominantly green wave-lengths which give rise to the colour impression it generates. Roughly speaking, this is the same as saying that it absorbs the red light. It follows that if there is only red light illuminating a region of complementary green paint, all of it will be absorbed and the surface will be deprived of its colourfulness. Likewise, a red surface illuminated by a complementary green light or a blue surface illuminated by a complimentary yellow light will also be perceived as achromatic. In all cases the outcome will tend towards an absence of light, which is somewhat prosaic way of describing "*black*". This being the case, it is easy to see that just the right mixture of white and red light could be used to neutralise a region of green, while leaving some white light still reflecting from it. The result would be a perception of *dim white* or, as it is usually called, *grey*. Similarly, it should be evident that with appropriate manipulations of the dimmer switches, the equipment described is capable of neutralising any other painted colour.

Land's second demonstration

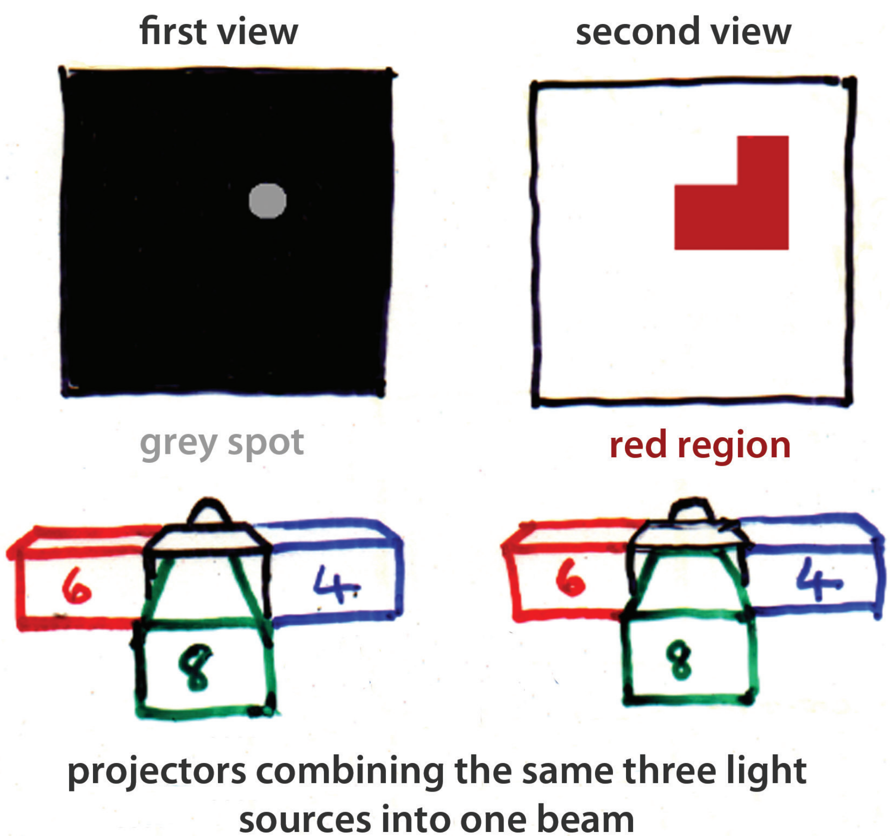

Figure 6: The set-up for Land's second demonstration. Here the red region is isolated, whereas for the experiment it was part of a MCD.

Whilst not having had the luck to be one of the subjects in Land's experiment and, therefore, only knowing of it from the *Scientific American* article, I did have the opportunity to experience his second demonstration. The reason was that my interest in the phenomenon of colour-constancy led me to visit Semir Zeki, *Professor of Neurobiology* at the *University of London*. Professor Zeki had been using Land's MCDs to investigate the neurophysiological correlates of colour-constancy and by 1980 he was publishing interesting results, most notably, the disovery of colour coded cells in a speciific part the visual cortex (known as Area V4). In his laboratory he showed me the set-up illustrated in *Figure 6*. It is essentially the same as the previous one, except for two variations:

- The projectors had been arranged in such a way that, by means of a prism, their light could be fused into a common beam.
- A diaphragm had been positioned such that the beam could be narrowed down or opened up at will.

After explaining this set up to me, Professor Zeki focused a narrowed-down beam so that it shone on one and only one of the colour areas in the MCD. It happened that, on this occasion, he had chosen a bright red area. As the light was now illuminating a uniform region of colour, that part of the MCD became equivalent to the uniformly coloured blank screen, described above. This meant that manipulations of the three rheostats would now bring about changes in colour appearance. In particular, as shown in the illustration and in accordance with the explanation given above, the red patch could transformed into a neutral grey.

The miracle occurred when the diaphragm was opened out for, although the wave-length combination of the illuminating light remained unchanged, the colour appearance of the patch, which I had just been seeing as grey, did not. I ccould not help being astonished when before my very eyes, it was transformed into the colour it had been when illuminated by white light. In other words, it had become as brilliant a red as ever it had been.

Professor Zeki told me that equivalent seemingly miraculous transformation can be wrought for all the different colours on the MCD, and that it is experienced by one hundred per cent of people with normal colour vision.

So what had brought about this seeming miracle? Since the transformation occurred at the moment when all the colours of the MCD became simultaneously visible, there was no escaping the conclusion that it was due to interactions that occur only when numbers of colours are viewed at the same time. In other words, colour constancy is a context-dependent phenomenon.

Implications

Why was I so excited by all this? The answer is quite simple. What Land's experiment clearly shows and what Zeki's demonstration so vividly confirms is that regions of surfaces of identical pigmentation can be perceived as the same, even when reflecting different (sometimes very different) wave-length combinations into the eye of the viewer. Though, in this case, the composition of the reflected light was varying over time rather than over space, was this not exactly the same paradox that had troubled me within the dogmas of Professor Bohusz-Szyszko? From this moment on, I knew that, if ever I could find a way, I had to get to the bottom of the mystery of how the eye/brain combination enables colour constancy.