CHAPTER 7

An artist comes to work amongst scientists

Introductory

A main purpose of this Chapter is to give some idea of what was for me a highly creative relationship with the splendid community of scientists into whose lair I innocently wandered soon after I first arrived at the University of Stirling, in 1975. At the time I had minimal knowledge of perceptual and cognitive psychology and I had never met a live psychologist of any kind. Over the months that followed I became friendly with several members of the Psychology Department. When I showed an interest in their subject, they gave freely of their time in explaining things to me, were generous in putting journal papers that might be of interest in my way, welcomed me to join in departmental seminars and, almost immediately, encouraged me to join them in experiment. In short, they nursed me through the learning slopes of becoming an experimental psychologist and over subsequent years were ever on hand if I needed advice or information.

As time went by the network of helpful people spread for I found attentive ears and invaluable sources of information in other departments: Environmental Science, Computing Science and Biology. Later, it widened its scope even further to include researchers from a variety of countries, whom I met at international conferences. The list of well-informed friends grew and grew and the debt I owe to them is great.

So why was I bothering them with questions? The short answer is that, being an artist and art teacher with an artist's and art teacher's point of view, I had an agenda that often saw me going off at tangents relative to the subjects that I could read about in books. I needed help in exploring new territories.

Two other purposes of this chapter are to indicate ways in which my background as an artist contributed to the originality of the findings described and to provide an introduction to some of the ideas that will be dealt with in greater detail in later chapters.

A first encounter

My initiation into scientific experiment came unexpectedly. It occurred very soon after my arrival at the University of Stirling where I had an appointment as an artist with nothing to do but work for three years on a painting project.[1] Whether by chance or design, the coffee room nearest to my workplace was that of the Psychology Department. Timidly entering it for the first time, I found myself confronted across a coffee table by a small bearded man who, knowing hardly anything about painting, but a great deal about the *psychology of perception*, engaged me in conversation on the subject of the development of *children's drawing*. He began telling me what he knew about "*Theory of Intellectual Realism"* which was the at the time in vogue and spawning journal papers on a worldwide scale. As I was to learn later, some of these had been written by distinguished professors, coming from the most prestigious universities. However, I was as yet unaware of all this and, because I was an artist with experience of teaching adults to draw from observation, I could see immediately that the theory was deeply flawed. It is based on the assumption encapsulated in the aphorism: "*Children draw what they know and adults draw what they see*". My problem with this catchy formulation was that my experience as a teacher left me in no doubt that the psychologists who supported this theory could not have looked at many drawings made by adults. For years, I had been battling with the problem of *knowledge-vitiated* adult productions. The psychologists had clearly built their theoretical edifice on foundations of sand.

My new friend, Bill Phillips,[2] whom I later discovered to be a leading expert on *visual memory*, was sufficiently interested in my incredulous response to suggest we do an experiment together. The first palpable result of our cooperation was a publication.[3] A longer-term outcome was the grant and the appointment to a *Senior Research Fellowship* that enabled me to undertake the studies that are described in the following chapters.

There was another longer-term outcome of the first experiment that proved to have great significance, for it provided the conceptual framework which was to constrain my thought when I delved deeper into the world of perceptual and cog-

1 Tom Cottrell Memorial Fellow, 1975-1978

2 Now Professor W. A. Phillips, *Centre for Cognitive and Computational Neuroscience,* at the University of Stirling, and Since June 2005 Fellow of the *Frankfurt Institute of Advanced Studies*.

3 Phillips, W. A., Hobbs, S. B. and Pratt, F. R., 1978, "*Intellectual Realism in Children's' drawings of Cubes*". Cognition, Vol.. 6, pages 15-33.

nitive psychology. The fact that everything started with the flawed *Theory of Intellectual Realism* turned out to be a blessing in disguise. Despite my criticisms, its formulation had the considerable virtue of highlighting the importance of the role of *knowledge* (i.e. Information stored in long-term memory that enables recognition and supports skills) in visual processing. From the very first experiment it was abundantly clear that *what people know* influences the information that they pick up for use in their drawings and how they make use of it. This may not be very surprising, but it is well worth while to keep it firmly in mind.

Once the focus was on *knowledge*, the question of accessing it arose and with it the importance of *recognition.* Add to this the obvious importance of *analytic looking* and it became increasingly evident that, if we were to make progress with our investigation, we would need take a holistic view, one which gave equal importance to the five, main stages of drawing from observation, namely:

1. The *structure of the object* or *scene* being drawn.
2. Accessing *knowledge* relating to it via *recognition*
3. Refining the *knowledge* using *analytic looking*
4. Using the accessed knowledge to guide both the subsequent *analytic looking* and *line production.*
5. Making use of the *emerging drawing* as a source of feedback.

Keeping all five constantly in mind proved to be very rewarding.

It was gradually dawning on me just how ambitious our project was. In the event, the ramifications of our research programme were so extensive that the final outcome proved to be something approaching a *general theory,* not only of the *acquisition and use of visually mediated skills* (with special reference to painting and drawing) but also of *creativity*. But, how this came about will emerge slowly.

Gaining courage

Meanwhile, there were other outcomes of bringing my artistic experience into a *Psychology Department*. Some of these were not very evident to start with. For example, it only gradually dawned on me that the story of my baptism into perceptual and cognitive psychology could be taken as evidence in support of a well-known phenomenon. From books, I already knew of stories of scientists, from one realm of enquiry, being able to see what scientists within another domain had missed, for the very reason that their framework of reference was different.

Bringing my artistic background to bear on the "*Theory of Intellectual Realism*" was a case in point. It provided a fruitful perspective that had previously been lacking. However, at the same time it left me with the foundations for a healthy scepticism concerning the conclusions of psychologists. If a whole community of eminent researchers could be blinkered once, presumably other communities and individuals could be so again.

For these positive and negative reasons and despite my lack of training in the subjects, I found myself progressively open to the possibility that I might be able to contribute something worthwhile to the study of *visual perception* and *skill acquisition*.

At first, I was very unsure of myself on this score, but time was to boost my confidence. On several occasions and in a variety of circumstances my experience as an artist led me to question the work of psychologists and, again and again, the process of questioning was to lead in fruitful directions. For example, J. J. Gibson mounted a formidable assault on the whole basis of the traditional laboratory experiments used in visual perception.[4] He argued that by holding too many variables constant, experimenters had cut themselves off from critically important aspects of visual perception. Gibson's ire had been first aroused because of the traditionalists had overlooked the focus of his own early research findings, namely on the role of *movement* (whether of the viewer or of the object being viewed) as a generator of visual cues. However, once on the warpath, he found much grist to his mill.

It has to be admitted that I found and still find many of Gibson's arguments persuasive with respect both to what the traditionalists had overlooked and to his new approach to seeking out knowledge. For a brief period I risked finding myself amongst the ranks of the Gibsonians. However, my knowledge as an artist was to save me. When reading a book on linear perspective, I noticed a striking resemblance between Gibson's description of traditional laboratory experiments and the use of *perspective frames* by artists of the calibre of Leonardo da Vinci, Albrecht Dürer and Vincent van Gogh. It was clear that extremely similar and equally stringent controls over variables were being used by creative artists of the first rank, presumably enabling them to *see* in ways that would otherwise have been difficult or impossible for them. I could not avoid the conclusion that something did not add up. Could Gibson have got something wrong?

4 J.J. Gibson, 1980, The Ecological Approach to Visual Perception, Houghton Mifflin,

This question haunted me for some days until a thought-provoking insight popped into my head. Perhaps, *constraint* plays a regular and fundamental part in visual perception whenever the eye/brain is searching new information. What Gibson had overlooked is that movement is just as much of a constraint as standing still, looking with one eye or taking in information in at a glance. He was right that a stationary eye/brain cannot use movement through the environment to generate visual cues, but he seemed to have overlooked the possibility that movement might deny it cues that are available to someone who is standing still. Could the simple truth be that the different contingencies require different visual systems and that a large portion, if not all, of these might depend for their efficacy on constraint of some kind? Following this line of thought I made a list of constraints used by artists when working and found considerable support my hypothesis (see *Chapter 9*).[5]

The synthesis of Marian Bohusz-Szyszko

Another fruitful outcome of my artist's perspective related directly to my experience of making paintings. One of my main interests was *whole-field colour relations* and, in particular, the degree to which changing one colour on a picture surface had an effect the appearance of all the others. Where this came from originally, I cannot be certain, but for its extraordinary importance to me I am indebted to the influence of Marian Bohusz-Szyszko and his synthesis of ideas coming from his predecessors (see the *Preface* and *Chapter 6*).

Less positively, I became aware of a serious flaw in the Professor's explanation as to why his ideas worked. The whole edifice depended on the undesirability of seeing separate regions of colour in different parts of the picture surface as being the same. Yet, according to his own argument, the laws of physics as they apply to the effect of the totality of primary and secondary light sources that illuminate very surface in the natural world, ensure that no two parts of a picture surface, even if containing regions of identical pigment colours can ever be reflecting identical wavelength combinations. My question was, "*How could it be possible to see something that is actually different as being the same*?".

As I could think of no way of resolving this conundrum, , for the time being at least, I had to live with my inability to do anything. However, the paradox stayed

5 *Which is an abbreviated and edited version of:* Pratt, Francis, 1985, "*A perspective on traditional artistic practices*." In, "*Visual Order. Studies in the Development of Representational Skills*". Eds.: Freeman, N.H. and Cox, M. Cambridge University Press.

as a niggling presence in the back of my mind. It was still lurking there some ten years later, when, a couple of years after I arrived, at the University of Stirling, a well-meaning psychologist friend placed an article entitled "*The Retinex Theory of Colour Vision*" on my desk.[6] In it the author, Edwin Land, attempts to provide an explanation for *colour constancy*, the phenomenon whereby people see regions of colour as unchanging despite considerable variations in the wavelength composition of the light coming from them.

Land's article starts with a report of a powerful and astonishing demonstration of colour constancy, using a multicoloured display. He is then able to prove that its occurrence depends on two conditions: The targeted colour must be perceived in the context of an array of other colours, and the entire display must be illuminated spectrum of wavelengths broad enough to stimulate all three of cone receptor types.[7, 8]

What Land's demonstration showed conclusively was that the eye-brain can classify a wide range of different wavelength combinations of light reflected from regions of a surface as being the same colour. I immediately realised that this finding had a great deal in common with what I had been looking for, namely a region of colour that is perceived as being the same as another region situated on the same picture surface, even though the two are reflecting different wavelength combinations. From this moment on I was on the alert for any information about *colour constancy* and turned what I found over and over in my mind, trying to see how on earth it could resolve my painting paradox.

The breakthrough came some four years later as a result of an encounter with Dr. Alistair Watson who was working in the *Department of Environmental Science* on the interpretation of satellite images by computer. We swapped ideas and, from a combination of my speculations and his mathematics, Alistair came up with a way of computing colour constancy on the basis of input from a video camera directed at a multicoloured array similar to the one used by Land.[9] We

6 Land E.H., 1977, "*The Retinex Theory of Colour Vision* ", Scientific American, Vol. 237 , p.108

7 In the case of Land's demonstration, the light mixture was created from a blend of three beams each of a different light primary, coming from three projectors. Each of these was equipped with a dimmer switch to allow the experimenter to vary the ratio of between the three light sources.

8 More on this in *Chapter 13*

9 We never published our colour constancy algorithm because we found that several other so called "*lightness algorithms*" had pipped us at the post. For a review of these see: Annie Hulbert, 1986, "*Formal Connections between Lightness Algorithms*" Journal of the Optical Society of

called it the "*Colour Constancy Algorithm*",[10] even though we were soon to realise that it could compute much more than just colour constancy.

It is impossible to overemphasise what a turning point in my intellectual life this constituted. The input of new ideas came just at the right moment to loosen the logjam at which my speculations deriving from the drawing experiments had arrived. It proved to be the first step in a process that has since provided explanations for a wide range of visual phenomena of interest to artists and indeed anyone interested in visually mediated skills. In particular it gave a significant boost to my search for an understanding of the eye-brain processes that underpin such skills.

Alistair and I discovered that for a physicist with serious mathematical skills[11] and an artist with hardly any, we had a surprising degree in common. Immediately, we warmed to the idea of sharing ideas about *vision*, *recognition* and *how brains work*. These gradually proliferated and soon after our first encounter we decided to see if we could set up an interdisciplinary group to study how the eye-brain uses visual acquired information. It seemed evident that the most appropriate way of testing our ideas would be by means of *computer modeling*. This is why the first recruit was Leslie Smith from the *Department of Computing Science*.[12] We called our new baby "*The University of Stirling Vision Group*" (*USVG*) and organised discussions amongst interested people from various Departments.[13] My involvement in these and in the production of working papers was to prove to be the most intellectually exhilarating period of my life.

One of the basic tenets we promoted within the *USVG* members was that computer programmes would only be acceptable if the principles upon which they worked were consistent with all known structural and functional features of

America A, Vol. 3, page 1684

10 "*Algorithms*" are the mathematical process that arrive at solutions one step at a time and for this reason are perfectly suited to providing the basis for computer programmes.

11 An early paper of his was: Watson, A.W., 1974, 'A Reimann Geometrical Explanation of visual illusions'. Maths Psych.

12 Now Professor Leslie Smith: Head of the 'Computational Intelligence Research Group' and Chairperson of the 'IEEE UKRI Neural Networks.

13 Computing Science, Psychology (with inputs on visual perception, visual memory and neurophysiology) Environmental Science, and Biology Departments. Of particular importance to us was the input on neurophysiology from Dr. Karel Gisbers and on biochemical aspects of brain function from Dr. Peter Brophy - now Professor Peter Brophy, Director of the Centre for Neuroscience Research at University of Edinburgh.

the eye and the brain. Our idea was to take into account current knowledge of:

1. The structure of the *retina*.
2. The structure of the *visual cortex* and its functional relation with other parts of the brain.
3. The working principles of *neural systems*.
4. Brain *biochemistry*.

In other words, our aim was *neurophysiological and biochemical plausibility*. We felt that though others had given lip service to this limitation on theorising, too many of them had not yet taken it seriously enough.

In addition we thought it important to take into consideration that the eye/brain has been evolving over millions of years. The lesson we drew from this was that, whereas computer designers can choose the structural features of their machines according to a virtually infinite number of alternative principles and computer programmers can arrange matters according to whim, biological systems are constrained by their origins. They had to evolve step by step from the most primitive organisms in such a way that every advance had to have been plausible. Thus, the fifth limitation that we placed on ourselves was:

5. Evolutionary plausibility

It helped us a lot to keep these five considerations in mind.

The artists got there first

Meanwhile, the subject of *colour constancy* also jogged my thoughts back to the ideas of J. J. Gibson concerning the perception of paintings. The great man seemed to take every opportunity to show his scorn for experiments on visual perception based on using photographs or other images found on flat surfaces. He asserted that, being impoverished sources of visual information, which they indubitably are, studies of them could have little relevance to how people use their eyes in everyday, real life, circumstances.

However, I found the courage to confront this argument thanks to my newly acquired realisation that constraints on the richness of the visual input, while inhibiting certain visual systems, will almost certainly enable others. Rather than worrying about the information loss in two dimensional images of scenes as compared with the scenes themselves in the real world, I began to look for visual systems that might be enabled by images in on flat surfaces in the absence of 3D or motion-generated information.

To understand the way I approached this task, it is necessary to take into account a number of puzzles that were in my artist's mind. They all concern effects of *repeated colours* in paintings. Here is a list of them:

1. Why should the repetition of colours in paintings be so disruptive to the visual experience of the viewer?
2. Why is it so difficult to take one's attention off the repeated colours and give it to other elements on the picture surface?
3. Why do repeated colours so regularly and predictably seem to jump out of their appointed place in illusory pictorial space?
4. Why do paintings with very large numbers of repeated colours become garish?
5. Why do repeated colours destroy any sense of luminosity within the illusory picture space?
6. Why do they undermine perceptions of three dimensional space within it?

In large part, all these questions stem from developments in painting influenced by the work of Seurat and Cézanne and, later, reflected in the propositions of Professor Bohusz-Szyszko.

You will remember that in addition to the first assertion that in any one painting no two colours must ever be the same, there was a second which stated that the colours used in making paintings should always be using paint-mixtures containing some proportion, however small, of pigment colour that is complementary to the predominant pigment colour. Even before meeting Alistair, I had been convinced that it could be no coincidence that this recalled Land's requirements for enabling his colour-constancy demonstration to work, namely that the light illuminating the multicoloured display must contain contributions from all three projectors, each projecting a different one of the three light-primaries. Indeed, as it turned out, sharing this conviction with him was on of the keys that led to his colour constancy algorithm (see below and *Chapter 11*).

The colour constancy algorithm and paintings

The answers to my artist's questions came flooding in from the work on Alistair's colour constancy algorithm. Basically it computes rates of change across the reflected-light profiles of surfaces. Abrupt changes signify borders between regions of colour and the extent of the change provides the basis for body-colour information in such a way that it can be subtracted off, leaving only slow varying gradations that provide information indicating *surface-form profiles.* Since the surface of the multicoloured display is flat, this profile indicates its flatness. Similarly, if we run the algorithm on a curved or domed surface it will tell us that it is curved or domed, whichever is the case.[14]

By using three filters corresponding to three primary colours and running the algorithm three times, it is not only able to provide information about *body-colour,* but also about the wavelength combinations of the residue. Since this corresponds to light being reflected directly from surfaces, it informs us of *ambient illumination*. Accordingly, assuming that the surface is being illuminated by daylight, or some other complex light source containing wavelengths from across the spectrum, the reflected-light from every part of the surface will include wavelengths from across the spectrum.

It follows that if artists want to represent the light reflected from surfaces in nature as they are they will need to use colours that include wavelengths from across the spectrum. None of the artists colours in their tubes or pans have this property, since the reason why the are perceived as having their characteristic hues (cadmium red, ultramarine blue, lemon yellow, etc.) is the selectivity of their absorption characteristics. The only way of rectifying the situation is to mix in other pigment colours that will provide the missing wavelengths.

Eureka! Surely, here was the explanation for Professor Bohusz-Szyszko's requirement that all colours should be mixtures containing at least some proportion of complementary? Clearly, the Professor's rules would be needed if artists were to create illusions light in illusory pictorial space.

From this starting point, other insights proliferated.

Edges of shadows

A seeming problem raised by the algorithm occurs at the borders of shadows, where there is a sudden jump in the lightness profile at borders between different regions of the same pigment-colour. The outcome is an ambiguity that the algorithm cannot resolve. It has no way of deciding whether this jump is due to a change in body-colour or a change in lightness. Placed in this predicament, its only way out of ambiguity is to opt for one of the alternatives. We know the

14 More about how the algorithm works in *Chapter 14.*

choice made by our eye/brain systems. It is to classify sudden changes in lightness at the edges of shadows as sudden changes in pigmentation.

This was another eureka moment. It provided an explanation as to:

- Why Semir Zeki found black-coded cells amongst the colour-coded cells in the area V4 of the visual cortex. What is special about these is that they react to the colours in Land's multicoloured display as we see them and not according to their wavelength profiles. In other words, they are subject to *colour constancy*.
- Why *shadows look to be black*er than their reflective properties alone would predict.
- Why it never occurred to artists to represent shadows in any other way than by adding achromatic black or grey paint, until after Seurat's pointillism, with its complementaries in every group of dots, revolutionised shadow painting.

Edges between overlapping surfaces

The trail was getting more and more intriguing. The next obvious step was to consider what might happen at the edges of objects or surfaces where the neighbouring colour would be some distance behind them. Here also there would be a sudden change in the lightness profile that would be at least partially and, quite likely, completely independent of changes in body colour. Surely, as in the case of the edges of shadows, this also should lead to ambiguity and leave the eye/brain systems with interesting choices. Could it be that we were onto another important discovery?

But the rapid rise in our excitement was abruptly cut short when we discovered that trying to compute such changes would completely screw up the algorithm. We found ourselves halted in our tracks just when everything seemed to be working out so nicely. It was time to stop and assess the situation. What was our achievement so far? We had an algorithm, based on ideas relating to known patterns of neural processing that not only perfectly modelled the outcomes of Land's demonstration but also predicted seeing shadows as achromatic. Its only defect was that it would be screwed up by in front/behind relations. There were two alternatives: abandon the algorithm as a guide to visual perception or work out how it might cope with the edges of objects.

Naturally we opted for the second choice, which meant that we had to put our thinking caps on again. To help ourselves we went back to our knowledge of how the eye picks up light-borne information. Very quickly, something that should have been an obvious leapt out at us. Our algorithm uses serial processing whereas eyes use massive parallel processing which means that they can use simultaneous comparisons between different aspects of the same visual input. Perhaps our problems could be explained by this difference. As we did not have the computing power to programme simulations of parallel processing, we had to resort to theorising. When we did, we concluded that comparing different views of the same layout, does indeed give the possibility of disambiguating sudden changes at borders between pigment-based colours and those at the edge of objects where both *pigment colour* and *lightness profile* change simultaneously.

Moreover, if the image were to be in black and white, texture information could take over the role of colour as the disambiguating agent. Either way, it seemed theoretically possible to compute the information necessary to show the existence of a border at which both colour and lightness (or texture and lightness) were simultaneously changing. From these speculations we concluded that, in principle at least, our algorithm, if running in the context of parallel processing, could compute the fact of spatial separation between an object and its background (henceforward referred to as "*3D spatial separation*").

Simultaneous lightness contrast

We were encouraged further still by an unexpected bonus. It should really have occurred to us beforehand, particularly after our finding with the blackness of shadows, but we realised that the algorithm would produce something very like the *lightness contrast effects* demonstrated by Gilchrist (See *Chapter 12*) and well known to artists as "*simultaneous lightness contras*t".[15, 16]

Testing hypotheses

In sum, the theoretical signs were very encouraging and all the findings and speculations so far had a very good feel about them. However, we still needed to find confirmation for our speculations. Since they could not be tested by means of computer modeling, was there any other way of making progress? The obvious answer was to work out predictions deriving from them and test these experimentally.

15 Artists confuse *lightness* and *brightness*.

16 Much more detail on this in *Chapter 14,*

As an artist, I did not see much of a problem. Indeed, I felt that I had already done the experiments in my years of testing the tenets of Professor Bohusz-Szyszko. However, it was clearly important to redo my experiments, this time with the new ideas in mind. One of my reasons for setting up the *Painting School of Montmiral*, which I have been running since 1987,[17] was to give myself the opportunity to do just this. There, I could test everything. More than twenty years later, I can confirm that the method I adopted of testing the ideas by predicting outcomes in student paintings, has worked a treat, and has taken my understanding far beyond the reach of my imagination at the time.

A description of my findings is to be found in the two volumes on painting associated with this book.[18] What they show is that we had stumbled upon a *comprehensive theory of the role of colour in the perception of light and space in paintings*.

Ideas begin to crystallise suggestively

The discoveries relating to the colour constancy algorithm also gave a boost to the speculations coming from the drawing experiments. Ideas began to crystallise in a most remarkable way and, as they did so, the focus was shifting inexorably in the direction of three related subjects, namely:

1. *Analytic-looking,*
2. *Recognition*
3. *Creativity*.

More about the "*blindness*" of experts

However, before saying anything about these, a short detour seems appropriate. This requires returning to examples of consequences of my perceiving matters from the perspective of an artist. I have chosen two examples that confirmed for me my discovery that even accredited experts can be very blinkered about matters outside their domain of specialist knowledge. A minor one comes first. Thus, on one occasion, I attended a seminar in which tracings, taken from photographs of the faces of famous people currently in the news, were being used for recognition experiments relating the to police enquiry work.[19] The speaker,

17 The Painting School of Montmiral in the Tarn region of S.W. France.

18 "*Painting with Light*" and *"Painting with Colour*"

19 Composite images, such as Identikit and Photofit, issued as means to catch criminals, had

represented a group of researchers working with the police on a big grant. What caught my attention was the speaker assuming that outline drawings traced from photographs must be "*realistic*". He seemed to have no doubt that it is reasonable do so. I was flabbergasted, not only by the assumption but also by the fact that none of the highly qualified psychologists attending the seminar thought it necessary to challenge it.

I am confident that artists reading this book will see why I was surprised. Like me, they will realise immediately just how many awkward decisions have to be made in choosing which features should be traced, exactly where a line should be placed and at precisely what point it should start or come to an end. Artists know that variations in the breadth and intensity of the lines can be critical in creating good likenesses. They also know that assessment of significance is vital is fundamental as are consequent decisions as to what to include and what to leave out: *Putting in everything in is a recipe for disaster*.

What undermined the value of the whole experiment was that lines in the photographs produced for the seminar showed little evidence that the tracer had taken any of the above matters into consideration. I asked myself: "*How can the psychologists so deceive themselves*?" The only answer I could offer relates to the well-attested capacity of human beings for being able to see what they want to see: The psychologists recognised the traced images because they already knew who they were. There was no sign here of any double-blind testing.

Nor should it have been only the artists who called the images into question.? Gibsonians amongst the audience should have been appalled. Of all the visual features that provide the least degree of invariance, outline must be the champion. Every single viewpoint, viewing distance and lighting condition will modify or radically change them. The idea that recognition could have anything to do with classifying in terms of the precise contour-curvatures should have seemed absurd to them.

The purpose of mentioning this example of a group of blinkered perceptual psychologists is to provide a context for suggesting that such people were a common phenomenon. Over and over again, I found myself incredulous the lack of visual sophistication amongst them. Nor did there seem to be any correlation between this phenomenon and the intellectual prowess of the person concerned: Without exception they seemed to me to be extremely able people.

proved unreliable and were being tested experimentally against artists likenesses and other possibilities.

If I search my memory for examples impressive intellects who can be naïve in matters of visual perception, the one that stands out comes from a seminar given by one of the outstanding speakers I had the pleasure to hear. His name was Geoffrey Hinton and for me, and from what they later told me, for other members of the audience, it was a very special experience to witness such clarity of exposition from such an evidently brilliant mind.[20]

However, at one point Hinton was talking about the possible connectivities between the eight corners of a *Necker Cube*.[21] With one part of myself, I was fascinated by the mental gymnastics he displayed, juggling so skilfully and thoroughly with the totality of the two and three dimensional relationships that he was finding. But with another part, I found myself becoming uneasy because he seemed to be under the impression that what he was saying might have something to do with human visual perception. Such an idea fitted very badly with both my experience of analytic looking as an artist and the speculations emerging from my experimental work on drawing skills.[22] From both these perspectives, Hinton's proposals had a completely the wrong feel. In particular, the neural systems involved in the visual systems dealing with early-processing, which I was learning about, seemed poorly adapted for juggling with input in such complex ways. Unless I had was very much mistaken, Hinton's ideas ranked very low on the neurophysiological plausibility test. I could not stop myself from becoming extremely sceptical concerning any connection between his exciting ideas and actual eye-brain function. I could only see Hinton's ideas as another example of the increasingly familiar combination of brilliance and implausibility.

For me the most useful part of Hinton's talk was a mathematical demonstration that a network of large overlapping receptive fields, as found in the human retina, could mediate the capture of information with the same level of precision of an array of individual receptors. This was a breakthrough in my understanding of how what I had previously perceived as the extremely bizarre feature of structure of the retina could actually have a role in the experience of seeing as I

20 I was not at all surprised to find that Hinton was in the course of constructing a distinguished career. An example of the kind of thing he was working on at the time is: Hinton, G. E. and Anderson, J. A. (1981) Parallel Models of Associative Memory, Erlbaum, Hillsdale, NJ..

21 A skeleton representation of a cube consisting of edge lines such that all could be seen simultaneously.

22 Pratt, Francis, 1984, "*A Theoretical Framework for Thinking about Depiction*." In, "*Cognitive Processes in the Perception of Art*". Eds. Crozier, W.R. and Chapman, A.J., North-Holland Press.

knew it from daily life. What Hinton had to say was for me like a distant light at the end of a very long tunnel.

Reaction to the Neural Netters.

Meanwhile the *USVG* was going strong and its activities lead to grant applications and via them to the founding of the *Centre for Cognitive and Computational Neuroscience* (*CCCN*) at the *University of Stirling*. This was very good news for the *University* but for me a setback, since it signalled not only a rather abrupt end for the *USVG* but also and a parting of the ways.

Those who became part of the *CCCN* veered off in the direction of an approach to the computer modeling of neural processing that was currently setting fire to the imagination of droves of cognitive psychologists worldwide.[23] Their interest was to find a learning algorithm that works upon the same principles as those brains used for acquiring knowledge. This meant modeling neural networks as found in the brain and providing them with neurophysiologically plausible learning rules.

Initially my feelings about neural nets were very positive for at least three reasons.

- I had been completely convinced by two of my colleagues, Karel Gisbers and Richard Bambridge, that neurophysiology must have much to teach researchers into brain function.
- The same two colleagues had long before directed me to D.O. Hebb's trail blazing ideas of on neural substrates of *the organization of behaviour*.[24] As the researchers that I came to call the "*Neural Netters"* were clearly taking these into consideration, I felt they were setting off on a good pathway.
- My friend Bill Phillips was extremely excited and bombarded me with excellent sounding arguments about the validity and importance of neural nets and buttressed them with lists of the impressive people who had espoused them. I could hardly follow all his arguments, but

23 The general idea is that since the brain is indubitably a network of interconnected neurons, any attempt to model its way of working, using computers, should consider the structure of this network. Of particular interest was supposed to be the study of how brains learn.

24 Hebb, DO, 1950, The organization of behavior: A neuropsychological theory. New York: Wiley.

surely this international array of top people could not all be wrong?

However, with the passage of time, the initially favourable impressions came to be modified. The first step of my partial disillusionment came when attending seminars by accredited Neural Netters reporting on the progress of their research. Despite the fact that the avowed basis of the neural net initiative was to use brain systems as models for computing ideas, over and over again, at question time, after yet another brilliant, largely mathematical exegesis, a question would come from someone in the biological sciences. It would be of the form: "*What an interesting and impressive talk, but I cannot quite see how your algorithms fit in with what is known about neurobiological (or biochemical) systems*." The speaker's answer would be disarming and along the lines of: "*Well I see what you mean, but don't you agree that what I have worked out is worth following up in its own right*?"

It might well have been, but whether it was or not made no difference to the fact that for me it was a sellout. I could only see it as a betrayal of the principle of using established knowledge of brain function to guide computer modeling. How could an algorithm that brushed this aside be of any value in the context of an initiative based on it?

Alistair claimed mathematical reasons for rejecting the particular approach to learning algorithms currently in favour with the neural net community and adopted at the CCCN. He seemed very sure that computations based on it would always be too slow. I had great respect for him but was not capable of having an opinion on this subject. However, whether he was right or wrong seemed of little importance for I had come to see learning algorithms, despite their obvious potential, as a sidetrack relative to the trail I was now following with increasing excitement. One way or another, the work on drawing skills and the implications of the colour constancy algorithm had focused our interest on the structuring of visually-acquired information.

The way we looked at the situation was as follows. As my psychologist friends had so often emphasised, about one third of the brain is devoted to preparing visually-acquired information before it is ready for *recognition* and *use*. Even without taking into consideration the multiple other systems that contribute to *recognition*, the clear implication is that a massive amount of preparation is necessary if these capacities are to function as they do. Presumably useful learning, no matter the method adopted, would be a great deal more difficult, if not impossible, without it. Accordingly we concluded that whatever the progress made by the *Neural Netters*, it would almost certainly be of restricted usefulness, unless their neural nets were trained on appropriately prepared information, or unless computing power increased exponentially.[25] As far as we were concerned, they were putting the cart before the horse. Learning algorithms, could be left to other people. Hopefully they would find ones that we could use later on.[26]

Making rules and sticking to them

As time went on I became more and more convinced that the rule of giving high priority to neurophysiological plausibility was a good one and worth sticking to ay all costs. Good reasons seemed to be coming at me from all angles. Here are three examples:

1. As already explained, Edwin Land's *retinex theory of colour vision* was a theory concerning how the eye-brain computes *colour-constancy*. The experimental results showed that ratios between triplets of reflectance values derived from a multicoloured display under different wavelength combinations remained constant. Like the brilliant French mathematician Gaspard Monge who had developed similar ideas two hundred years earlier,[27] Land jumped to the conclusion that the brain must be computing these ratios. At this juncture, he seems to have given little consideration to the subject of neurophysiological plausibility.[28]
2. A follower of David Marr produced an article concerning *stereopsis*, in which two possible methods were suggested. One, the author claimed, worked better, the other was more like what happens in the brain. I was astonished to find that he chose to pursue the former.
3. Gibson, in pursuit of his movement-related ideas, found that movement through the environment created invariant patterns in the flow of information entering the eyes. He was in no doubt that these "*flow fields*",

25 Which it did in the next 25 years. "*Deep Mind*" the current computer learning miracle of neural net heritage, requires mind boggling amounts of memory storage.

26 Now that we are well into the 21st century, we can see that the neural netters persevered and that Geoffrey Hinton, armed with access to ever increasing amounts of information storage, went on to be one of the pioneers "*Deep Mind*"

27 Cf. John Mollon, 1986, "*Seeing Colour*". A chapter in: "*Colour Art and Science*", Eds. Lamb and Boreau, C.U.P.

28 However, he clearly soon realised his mistake. Cf.: Edwin H. Land, August 1983, "*Recent advances in retinex theory and some implications for cortical computations: Colour vision and the natural image*". Proc.. Natl. Acad. Sci., USA Vol. 80, pp. 5163-5169.

as he called them, must be used to enable visually guided navigation through the environment. Like Land's ratios, this was all very well, but left open the question as to how the eye-brain could compute the complex, movement-generated patterns.

Computing flow fields

Unlike Gibson, who was against looking inside the head, we thought immediately of looking at neural processing structures. By doing so, we came up with what seemed to us to be a very convincing way of computing flow fields. Our guiding principle of neurophysiological plausibility dictated that the best approach would be to work out how the actual retina, with its multiple overlapping receptive fields of different sizes with their propensity for lateral inhibition[29] would respond to them. Our conclusion was that a structure with these features and properties would come very close to performing the necessary computations and do so in a manner analogous to the procedures used in the colour constancy algorithm. However, there was a missing variable and it was the search for this that led us to see the significance of the known limitations on the time taken for receptors to integrate information. It is because of these that a "*greying-out*" occurs during the rapid eye-movements known as *saccades*.[30]

Likewise, if a stationary eye is exposed to a fast moving target it also greys out (as a film in a camera would blur). What we realised was that information about speed of passage through the environment would automatically be computed from the pattern of *greying-out* within the receptive fields. To give an example that should make this clearer, it is easy to show that as a viewer approaches a surface it expands and that the rate of expansion increases with nearness of the object. Clearly, if the expansion rate is too rapid, greying-out of receptors in the retina will occur. However, since the expansion-rate is greater towards the peripheries of the retina and lesser towards the centre of it which remains stable, this greying-out will not be uniform. Accordingly, the regions at the edges of the retina will grey-out before those at the centre. Similarly, since the receptive-fields in the retina vary greatly in size and, the smaller ones will grey-out more rapidly than the larger ones. Without working out details, it should be evident

29 For more about how neurons inhibit the activity of their neighbours as a means of producing same/different information, see *Chapters 14,16* and *18*.

30 Campbell, F. W. and Wurtz, R.M., 1978, Saccadic omission: why we do not see greying out during saccadic eye movement. Vision Research: Vol. 18 pp. 1297-1303

that either the nearness to a surface or speed of approach can be computed with some accuracy by means of a simple count of which receptive field size-groups are still operational. In this way, that which otherwise would seem to be a complex calculation becomes extremely simple. What is more it does so in a highly neurophysiologically-plausible manner.

Computing colour constancy

The same conclusion can be drawn from an examination of our colour-constancy algorithm. No need to work out Land and Monge's ratios when the need for calculating them could be bypassed by a neurophysiological plausible system working on the simplest of principles. As with the computation of flow fields our theoretical speculations revealed that comparisons between outputs from different sized overlapping receptive fields within the retina can provide a particularly simple method of calculation. Indeed, we concluded that much the same mechanisms that we were proposing for computing the rate of approach to surfaces, would also produce information concerning texture, orientation, edge, surface-form, colour, ambient illumination etc. It was all very exciting. Indeed, with the advantage of hindsight we can now see that we were well on our way to the "*Context-based model of perceptual processing*" which provides the climax to this book.[31]

A fruitful detour

As indicated earlier, after coming to the conclusion that useful *learning*, no matter what the method adopted, might well be impossible without the massive preparation that occurs in the actual brain, we decided to concentrate our attention on *preparation* rather than on *learning* itself. For the time being, we were happy to leave that task to others. However, we found the way the ideas were developing ideas nudged us into an unpremeditated detour.

Once again we started by looking at known eye and brain structures. A key step followed when we had the idea of thinking about the well established research that showed that the early processing parts of the visual system, though usually referred to as "*hard-wired*", have to sensitise themselves in the early days, months and years of life. This they do so as a result of the exposure of the retina to patterns of light entering it from the outside world, as modified by the

31 See *Chapter 27*

baby's body, head and eye movements. In other words, as suggested earlier, the hard-wiring of the visual system is the product of a process of learning.

The conclusion that lower level visual processing systems are able to learn, was exciting for us because it gave direction to our search for a neurophysiologically plausible learning algorithm. Surely it would help our progress if we could make use of the fact that so much more was known about both structure and working principles of the early processing stages of visual perception than about what happens in the further reaches of the brain, where others were seeking their inspiration? Accordingly, we searched out and pondered research findings relating to the component parts of neurons (dendrites, cell body, axons) and the mathematics of their functioning.[32]

We also gave special thought to the structure of the *processing columns* found in the visual areas of the brain. All along we had assumed the importance of *lateral inhibition*. What we had not considered was the importance of the fact that when axons branch, each of the bifurcations outputs the same signal, with the result that they become *signal amplifiers*. The potential significance of this simple discovery became evident when we learnt of a biochemically plausible learning rule, namely that learning (in the form of an increased sensitisation) takes place whenever a neuron is overloaded with input.[33]

With these ideas in mind we hypothesized what would happen if we passed information through a columnar network of neurons capable of amplifying signals by means bifurcating axons with each abiding by the learning rule just mentioned. Much to our delight we found that what we called our "*learning column*" automatically sensitized itself to input, In other words, it enabled learning. Alistair gave our ideas algorithmic expression and wrote a computer programme that worked a treat. We demonstrated its ability to learn in a public display.[34] For this we linked a computer programmed with Alistair's programme to a video camera and demonstrated that the combination could "*learn to see"* in three different ways:

1. When the video camera was directed towards a slideshow of favourite

32 Largely through reference to Shepherd G.S., 1983, *Neurobiology*, O.U.P. Oxford.

33 Lynch G. 1985, "*Synapses, circuits and the beginnings of memory*", a seminar given at the University of St Andrews, Department of Psychology.

34 *Phenomenon*, an exhibition designed by the author and stimulated by ideas coming from the University of Stirling Vision Group, under the auspices of the Stirling District Council. June – September 1987, The Tollbooth, Stirling, Scotland

photographs, a *learning-column* could sensitize itself to different intensities of illumination. In other words, the computer had learnt to be a *light meter*.

2. When exposed to a display of the spectrum of colours, a linked group of three *learning-columns* could sensitize themselves to sixty-four colours which the computer learnt to recognise.
3. When a number of primitive shapes (a square, a circle, a triangle and a rod) were waved in front of an artificial retina, composed of an 18 X 18 matrix of laterally inhibited, light-sensitive sensors attached to a *learning-column,* the computer learnt to recognise them.

The parsimony of the programme was remarkable: Not only was learning very rapid but was achieved by means of BBC computers, using 32K of memory. Notice that our belief in the advantage of preparing information was given support by the efficacy of the artificial retina.

Low priority given to colour

Amongst psychologists of the period when I was at Stirling, there was much resistance to the idea of investigating colour in the context of recognition research. Their prejudice came from two sources. The basis for the first was the fact that human beings (not to mention many other species of animal) are able to make sense of black and white images. Did this not prove that colour is largely redundant? A second reason for the prejudice came from the fact that at the time the neurophysiological research which was exciting everyone was coming from the *single cell recordings* that were revealing edge and orientation detectors in the visual area at the back of the neocortex.[35] All this coincided with the dawn of the era of the computer modeling of visual systems. The acknowledged powerhouse of ideas in the domain of early visual processing was a group working with David Marr at Massachusetts Institute of Technology.[36] The possibilities of gaining information from arrays of edge detectors and interactions between them were many, presumably enough to keep the modelers happy for many years.

Meanwhile, single cell recording of colour sensitive cells was less advanced and for many the modeling of the complexities of colour vision seemed a daunt-

35 Hubel, D.H and Wiesel, T.N., 1977, The Ferrier Lecture/ Functional architecture of macaque monkey visual cortex. Proc. R.Soc. London. B 198, Pages 1-59

36 Marr, D, 1983, Vision, Freeman

ing prospect. Not surprisingly, when the general view was that the role of colour in the recognition of objects wasn't very important, it was the consensus that it could safely be left aside until later.

This was certainly not our view. Many perceptual problems could not be solved without colour. Amongst these were my painting-related problems and Alistair's satellite imagery problems. Yet again, I found myself looking at a situation from a different viewpoint than the main body of psychologists. This time, because of the nature of his work and the ongoing promise of our cooperations, I was joined by Alistair.

Not for the first time, choosing the unconventional route gave us a considerable advantage. Moreover, as should perhaps be apparent by now, I was very far from convinced by edge-detection related approaches to recognition. As just pointed out, edges, being the most variable of features, must surely be amongst the least learnable. Once again, my feeling was that too many of the computer modelers were trying to solve problems in a way that the brain would be incapable of doing.

This is not to say that edges are not important. Of course they are. But, not for the preliminary processes of recognition. Presumably, their day would come when we delved more deeply into *analytic looking* and its role both in the later stages of *building object-descriptions,* in organising *action-instructions* and in the processes that enable *creativity*.[37]

Recognition and analytic looking

It is now the time to go a little more deeply into the subject of *analytic-looking* as it is used in drawing and painting from observation. The structure of inputs, the organisation of outputs and knowledge-acquisition all play essential parts in the implementation of these skills.

Earlier I told how my starting point as an experimental psychologist was an attempt to demonstrate the shortcomings of the *Theory of Intellectual Realism*. I also explained one of its main benefits from my point of view. Entering into the subject of drawing from observation from the perspective of knowledge-driven looking and line production meant that from the start I was pushed into thinking about the main aspects of the drawing cycle as being intimately intertwined. This *all-inclusive mind-set* had important consequences when I came to consider

37 *Chapters 19-21*

recognition and *analytic looking*. One conclusion to which it led me was that the main function of the former is to enable the latter and the organisation of actions in general (including line production and verbalisations)

After Bill and I received our grant and I had been welcomed into the Psychology Department as a *Senior Research Fellow,* our experiments involved people making copies of line drawings. Many of our findings were predictable, but none the less significant for that.[38] Not surprisingly and fortunately, the first experiment showed that *experience has a large effect on the type of drawing produced*. We could not have proceeded with our programme of research if this had not been the case. Further studies produced a second predictable finding, namely that the *presence or absence of the model during drawing activity makes a difference to the result*: People draw more accurately if they can look back and forth between model and copy than when they try to draw it in its absence, from memory. A later and equally anticipated discovery was that the *extent of knowledge* correlates with the *nature of looking strategies*. Thus, in the experiments in which people made copies of cubes, greater experience meant more looking back and forth. Typically, the looking behaviour took the form of a combination of long and short looks.[39] Those with less experience used a minimum of looks and hardly any of these were long ones.

So what have these results got to do with recognition? A great deal, for it is recognition that accesses the knowledge which is the fruit of experience and which guides both the looking behaviour and the line production-strategies. Our simple but important insight was that enabling these two types of action is the totality of what the recognition processes are required to do in order to mediate graphic skills (and perhaps all other skills). Complete knowledge of the object being represented (whatever that might mean) is clearly unnecessary and certainly not available in memory.

If we are not looking for complete knowledge, what are we looking for? Guided by the principle of parsimony, we came to the conclusion that the answer must be the *least information necessary*. To clarify the implications of this last statement, some thought-experiments may help. The fact that these use a very simple set up may seem to weaken the experimental findings, but later it will

38 Pratt, F.R., 1983, "*Intellectual Realism in children's and adults' copies of cubes and straight lines*." In, "*Acquisition of Symbolic Skills*." Eds: Rogers, D.R. and Sloboda, J.A., Plenum Press.

39 Short looks were defined as lasting less than one second (enabling no more than two saccades). Long looks lasted more than a second, enabling three or more saccades.

be suggested that this is not at all necessarily the case. In anticipation of this conclusion, it is worth recalling my criticism of Gibson's idea that experimental controls risk distorting the value of experiments. My argument is that, on the contrary, constraint is of the very essence of everyday visual perception and that therefore provide a very important variable.

The simple set-up for the thought-experiment comprises a computer-video camera combination (henceforward referred to as the "CVCC") learning to recognise two objects placed in a box with a grey interior. Let us suppose that the objects are a red-spotted white cube and a green cone . Further, let us assume that the CVCC is capable of isolating either of the two objects from its background. If so, given the severely limited context, recognition could be achieved on the basis of colour alone (red spots on white versus green), texture alone (dotted versus smooth), surface-form-profile alone (flat versus curved) or, even, position alone (assuming that the objects remain stationary). There would be no need to engage in the more complicated business of analysing shape or working out how the parts fit together. Moreover, in this situation, the analysis needed for achieving recognition could be very crude indeed. Take the example of colour. The input from the whole surface of the cube could be averaged. A way of doing this would be to use a filtering system with filters large enough to blur the distinction between the red dots and the white ground. These would be performing a function analogous to that of large receptive fields in the retina. As a result of the filtering, the red and the white would be fused into a pink and this would be just as useful a cue as the red spots on white. There would be no need for the engaging in the more complex task of identifying the spots and differentiating them from their background.[40]

In some circumstances, the CVCC might be left in doubt about the validity of conclusions derived from colour, texture, surface-form-profile or relative position information. As a result, it might initiate in a targeted search. For example, it could analyse the contours of the objects looking for a curve. Should it find one, it will have resolved the uncertainty of the situation (since cubes have no curves).

If it is to search around the contours of an object, the input device (the "*eye*" of the CVCC) will need to be mobile (rotate on its stand or move from side to side or up and down), as indeed it would have to be for many other tasks (particularly ones that require information about position relativities). This activity

40 Other demonstrations of the potential usefulness of the crude information provided by large receptive fields will be found elsewhere, particularly in *Chapter 18*.

would have an analogous function to that of eye, head and body movements and the need for it is evidence that the visual information required cannot be provided by bottom-up processes alone. A *perceptual-cycle* is necessary in which information provided by a first stage is confirmed or elaborated by means of subsequent stages. Notice that the cycle is being triggered by a need. This same need also determines the subsequent strategy for looking, the one that takes the form of a targeted, top-down search for further information. Another very important point to make is that this type of search is necessary for making available information that cannot be accessed by bottom-up processes alone.

Action control and different forms of memory

The subject of making visual searches brings us naturally to the topic of "*action control*". As an approach to thinking about what this entails, consider the case of programming a CVCC to make an accurate copy of the visible edges of a cube. For this it would be necessary to provide *action instructions* to guide both *looking activity* and *line output.* In order to plan the latter, a number of information-capturing looks would be required for making estimates of the relative positions, length and of orientations of the different edges. Since relativity judgements involve *comparisons*, each and every one of them would require the video camera to change its direction of focus. There would also be a need for some kind of *memory store,* for how can two things be compared unless one of them is being held in store while the other is being attended to?

But what kind of store does this have to be?

Types of memory

Our visual systems use four types of memory. They have been given the names "*iconic"*, "*short-term visual (STVM)*, "*working*" and "*long term*"*(LTM)*. Analytic looking requires all four.

- *Iconic memory* is required for the first step of all visual processing. It is manifest in comparisons, using lateral inhibition, between inputs to linked receptor cells (or linked groups of them) in the retina. Since these are changing continuously with every movement of the eye and with every modification of the environment, the only information available is the differences between inputs. Since differences cannot be computed unless information about what is being compared is retained long enough to allow the comparison to take place, a form of memory is required. This has been given the name of "*iconic memory*". A fundamental feature of this kind of memory is that it is necessarily destroyed by the input that enables the comparison. Consequently, it only endures a matter of milliseconds.

- *Short term visual memory* is required because there is a time lapse between the arrival of information in visual cortex and its use the organisation of actions, including the actions required for guiding appropriate analytic looking. This is necessary to allow time (a) for the processes that enable recognition to take their course, (b) for the recognition systems to access *action-instructions in long term memory,* and (c) for the *organisation of actions* relevant to the analytic-looking task in hand. During the time taken by this process, there is a need (a) to store the information that is to be analysed, and (b) to block disruptive information coming up the optic nerve. However, as soon as the information targeted by the action-instructions has been extracted, the block must be removed. Accordingly, the storage time is severely limited. This is why it is called *"short-term visual memory"*. It is also why, although considerably more robust that iconic memory, that one of its necessary characteristics is extreme fragility.
- *Working memory* is required for the organisation of actions, including those that control the *eye-movements* involved in *analytic-looking*. For this purpose various modalities of information, coming from a variety of sources, located both in current input and in long-term memory, have to be assembled. Clearly, this could only be possible if each of the modalities is held in store while the action is being planned. The required storage facility is called "*working memory*".

  Once the action instruction have been given, it is necessary for the action organising system to be ready to respond to subsequent inputs of information. Accordingly *working memory* is also necessarily fragile. However, the complexity of the task for which it is used means that it is a lot less so than *short-term visual memory*.

  Working memory is also different from *short term visual-memory* in that its organisational strategies can be learnt. Accordingly, it can become more efficient when it is mediating familiar looking or doing skills. One example of this in the context of drawing from observation is that it can learn to increase rapidity of *information pick-up*, allowing, amongst other things of great value to artists, for faster drawing without loss of accuracy.

- *Long-term memory* is required for whole list of purposes, including the organisations of actions, such as those required for *analytic-looking*. By their nature, all skills are built up progressively and must endure over long periods of time. The fact that people can become skilled at copying tasks makes it clear that the knowledge required for enabling them to be so must be available for use whenever the skill in question is required. That this is possible demonstrates that, however ephemeral its contents, there must be something about a scene being drawn that has a degree of permanence (for example (a) the invariants that underpin the rules of linear perspective and anatomy, (b) the fact that no two colours in any one scene are ever the same and (c) the fact that all scenes are made up of a limited number of so called "*visual primitives*").[41] If this degree of permanence in both the external appearances and effective strategies for making use of it did not exist, skill acquisition would be impossible.

The importance of context

Up to now, we have been neglecting an important factor. Hardly a mention has been made of "*context*". It is high time to remedy this state of affairs, since *context* has a role in absolutely all cognitive activity, at all levels. In later chapters, its central importance in eye/brain function will be given greater emphasis. For the time being, it may help to keep three considerations in mind:

1. Firstly, at the most fundamental level, *context* is required for the functioning of the three workhorses of all biological-processing systems, namely *relativity, comparison* and *transformation*. This must be the case since, when any two things are related, compared or transformed, the one necessarily provides the *context* for the other.
2. Secondly, the use of *context* is essential in the key domains of classification and the recognition. To enable these fundamental processes, the eye-brain uses cross-correlation between different modalities of information, a process that depends on each modality standing as contextual

41 Much more on this in my books on the practice of drawing and painting.

information for those to which it is being related. The outcome is a formidably powerful classification tool, which provides the information-base that underpins recognition processes.

3. Thirdly, *context* is of fundamental importance to analytic looking, which depends for its *action-plans* on the eye-brain's *knowledge-base*. This not only determines the direction of gaze but also controls the level and domain of abstraction accessed, which it does in a top-down, context dependent and, therefore, extremely parsimonious manner. For example, once it has recognised what it is looking at, it can decide to concentrate the analytic-looking resources on:
   - Whole objects.
   - Parts or features of objects.
   - How the parts or features fit into each other.
   - How objects relates to other objects
   - How objects relate to their background.

When doing so, it can make use of information coming from, texture, lightness relativities, surface-form profile, ambient illumination, body-colour, surface connectivity, in front/behind information, contour curvature, orientation detectors, directional indicators, etc. In all this, the *knowledge-base* provides the *context* for analytic decisions. In doing so, it both filters out what it has learnt to classify as unnecessary information and directs attention where it can be most useful.

In summary, the lesson to be learnt from the above remarks on *context* is clear. Anyone seeking to understand or model visual-processing must keep it constantly in mind. The need to do so is just as essential as it is for artists when making paintings in which every colour on the picture surface influences the way each and every one of the other colours on it is perceived.

"*Gestalts*" and "*feelings*"

As a conclusion to this chapter, let us return to the theme of the benefits of approaching research from a personal and, therefore, idiosyncratic viewpoint. This time the focus is no longer on my perspectives as an artist but on the influence of my background in science. Although I started this chapter by asserting my "*minimal knowledge of perceptual or cognitive psychology*", the knowledge that I did have played an important role in the development of my ideas. In particular it was responsible for triggering doubts relating to the ideas of the *Cognitive Psychologists*.

A main reason why my extremely sketchy scientific knowledge was able to perform this valuable service was that it was so out of date. This was because it came largely from my student days at *Art School* and *Teacher Training College,* both of which institutions took it for granted that perceptual and developmental psychology were relevant to their educational aims. One of the first things I was asked to do at my *Art School* was to read a book entitled "*Art and Visual Perception*" by Rudolf Arnheim[42] and at *Teacher Training College*, I was subjected to a healthy dose of "*Behaviourist*" theory.

In view of this background, it was quite a shock to learn from my new colleagues at the University of Stirling that, in their view, Arnheim's *Gestalt* ideas were outmoded and that the *Behaviourists* were wrong-thinking enemies of progress. Although prepared to believe that there must be good reasons for their strongly negative views, I could not help feeling that there might also be an element of throwing the baby out with the bath water. Let me explain why this was so, by means of reference to my own experience.

My Art School teachers were all abstract artists. The reason why they encouraged their students to read Arnheim's book was that they saw his *Gestalt* ideas as grist to their mill. In particular, they were enthused by the emphasis on the perceptual importance of *groupings* of shape, colour, texture, etc.. No wonder this was the case for it fitted so well with their experience of making and looking at paintings. As for myself, even before I went to *Art School*, I already had good reasons for a similarly positive reaction. Don't forget that I had already been taught by Professor Bohusz-Szyszko that repetitions of identical colours are to be avoided at all costs if pictorial harmony is to be achieved. Also, from my efforts to follow his rules, I had learnt that groups of almost identical but never repeating colours are essential to the creation of the kind of pictorial dynamics that most excite me. Evidently there is something both about relationships between identical colours situated in different parts of a picture surface and about interactions between members of groups of similar but different colours distributed across a picture surface that contribute to pictorial dynamics. Accordingly, it was natural for me to sympathise with the ideas of the *Gestalt Psychologists*.

Nevertheless felt I should listen to what my Stirling colleagues had to say. I could not help feeling that they might have good reasons for their dismissive

42 Rudolf Arnheim, 1954, Art and Visual Perception, Faber, London

attitude, and indeed they did. But these had nothing to do with *gestalt* groupings. Rather they targeted the *Gestalt* theorists' rather woolly ideas about brain function which by this time were indeed well and truly out of date.

It was a similar story with my psychologist friends attitude towards the Behaviourists stimulus-response model. In their opinion its formulators had ruled the roost for too long. How could anyone, they argued, so blatantly ignore events that take place in the brain between the activation of the sensory receptors and behavioural responses? Surely all those billions of cells and neural processes must have a function? How could anyone doubt that these should be made the main focus of research?

Although persuaded by this argument, I could not help feeling that before condemning the *Behaviourists* out of hand it would be only fair to remind myself of what I had learnt about their experiments. My memory was that, by means of a simple *reward* and *deterrent* regimes, experimenters could reliably train rats to turn a given way at the junction in a T shaped box and pigeons to play ping-pong in hardly any time at all. The rewards used tended to be food and the deterrents, electric shocks. I also remembered sometimes wondering whether the *Behaviourists* were wasting their time training animals to accomplish such trivial tasks? However, I hesitated before coming to such a conclusion because:

1. Amongst them were many very able people.
2. They had been given much support over a long period of time by the most highly placed grant-giving bodies.
3. They had done an incredibly large number of experiments.
4. Their paradigm worked not only with rats and pigeons but also, to a limited extent, with human beings.

Surely something of lasting value must have come from the combined efforts of all these clever, well funded researchers?

Asking myself this same question more recently, I focused on the fact that the main preoccupation of the *Behaviourists* was *learning* and that they saw it in terms of *habit-formation, habit-perpetuation* and *habit-loss*. Nobody could seriously question the importance of any of these four subjects, certainly not anyone interested in the *acquisition of the analytic looking skills used in drawing and painting from observation.* How could they when the development of those skills requires unlearning some habits and substituting them with newly acquired ones?

Desire and fear

More pertinent in terms of my new interest in brain function, neither reward nor deterrent would work in the absence of the basic instincts of *desire* and *fear*. Without these motivating forces neither the rats nor the pigeons would have made any progress. But, as everyone must know, *desire* and *fear,* as modified by *experience,* are fundamental feelings in us all, and it is hard to believe that anyone would deny that they must have a role to play in the acquisition and use of knowledge. This may seem obvious to parents who have tested the efficacity of rewarding and punishing on their own children and likewise to artists whose feelings play such an important role in their thought processes and actions. For myself, it would be hard to persuade me either that the choices I make when working on my own paintings are not influenced by positive and negative feelings telling me what is "*good*" and "*bad*" or that these are anything other than the fruit of the combination of instinct and learning that wheedles its way into every aspect of brain function.

In summary, even though there can be little doubt that the *Cognitive Psychologists* were right to pay more attention to what is happening in the head, they were wrong to ignore the role of feeling in *knowledge-acquisition, analytic looking* (hearing, tasting, sensing, etc.) and *though*t. From the earliest days of my efforts to understand the perceptual and cognitive processes involved in artistic practices, I decided that in any model of them with which I was associated would give a fundamentally important place to the guiding role of the feelings.[43]

Implications

My artist's perspective unremittingly pushed me in the direction of the personal framework for thinking about many of the subjects that most interest me.[44] Looking at visual and cognitive processing in the light of a specific practical task that not only requires analytic-looking and line-production skills but also depends heavily on the use of feedback, suggested to me an abundance of questions and some answers. How could I resist the opportunities that came my way

---

43 See the summarising diagram in: Pratt, Francis, 1984, "*A theoretical framework for thinking about depiction.*" In, "*Cognitive Processes in the Perception of Art*". Eds. Crozier, W.R. and Chapman, A.J., North-Holland Press.

44 For an early stage in this process, guided by the drawing studies, see Pratt, Francis, 1984, "*A theoretical framework for thinking about depiction*." In, "*Cognitive Processes in the Perception of Art*". Eds.: Crozier, W.R. and Chapman, A.J., North-Holland Press.

to investigate further?

The remainder of this book tells the story of what I found with respect to brain function rather than my discoveries relating to the specifics of artistic skills and creativity, which are dealt with in the other volumes in this series. The main conclusions were reached by 1986, when I wrote, with the help of long and lively discussions with Alistair and Leslie, the text to be found in Chapter 21, which presents a speculative model of the "operational principles used by the eye-brain-body when coding and making practical use of the ever-changing array of light-borne information entering the eye". Although written so long ago on the basis of the evidence available at the time, there seems no reason to change it. In general, subsequent research seems to confirm its speculations, particularly those parts of it that emphasise the roles of context and the feeling based systems in visually mediated activity.

Enough has been said. The twin purposes of this chapter have been fulfilled. An explanation has been given of how I found the courage to take the plunge into the domain of scientific research and how the process of furnishing this and elaborating on some issues that arose provided a preview of some of the ideas to be found in the chapters that follow.