CHAPTER 29

More on experienced reality

Introductory

*The purpose of this chapter is to complete the preparation for the exercises, found in the next chapter. While doing them, it will be interesting to keep in mind one of the leitmotivs of both volumes in this series of books ,[1] namely the question of how the word "accuracy" should be defined. Are we going to relate our efforts to what I have been calling "**measured reality**" or to what I have been calling "**experienced reality**" and what do these phrases mean in practical terms? In the interests of furthering the discussion, other questions need asking. Perhaps the most fundamental of these is:*

- *"What information is required to persuade the eye/brain to perceive a depiction of an object or scene as bearing a resemblance to its source?"*

Since the answer can be "surprisingly little", we find ourselves faced with two further questions, namely:

- *"Which aspects of appearances are essential for recognition and which can be given low priority or even left out."*
- *"Where do exaggeration, distortion and abstraction fit in?"*

Recognition

At various points in this book the importance of *recognition* has been emphasised,[2] with special reference to its supreme achievement of being able to classify objects as being the same when they are actually different. Much has already been said about how each particular object (for example, a particular face, a particular tree or a particular chair) will always vary according to viewing conditions. Its size changes with differences in viewing distance, its shape alters with viewing angle, its colour is transformed by lighting conditions and almost

1 *"Painting with Light"* and *"Painting with Light and Colour"*..

2 And in much greater depth in *"What `Scientists can Learn from Artists"*.

everything varies with the passage of time. However, even though the changes in the appearance of any particular object can be large, they are small when compared with the those that occur between different examples of the same object-type, such as the faces of different people, different species of tree, different chair designs, etc..

Although nobody has a comprehensive answer to the question as to how this seeming miracle is achieved, it is possible to come to certain strong conclusions, most of which can be based on personal experience and common sense. For example, it is clear that recognition can be achieved at a glance, in a fraction of a second and, therefore, that it does not depend on long drawn out processes. Also, we can work out that it cannot be based on the perception of object contours since they vary with every change in viewing conditions.[3] To make use of these, the eye/brain would have to memorise an infinite number of different contour representations for every single recognisable object, which is clearly impossible.

Nor can it be based on any aspect of surface-appearance since all of these likewise vary continuously due to the ever changing direction, strength and relative influence of primary and secondary light sources. In other words, the information contained in the light entering the eyes that enables us to see arrangements of colours, lightnesses, shadows and shading is infinitely variable. Accordingly, as with contour, it would be impossible for the eye/brain to use it as a basis for recognition, unless by means of processes that simplify.

Another feature of recognition-systems is that they are able to function perfectly well when confronted by arrangements of colour/lightness variables that they have never come across before. We can be sure of this, if for no other reason than that successful likenesses in paintings are always made using colour/lightness relativities that can never have existed in the natural world.[4]

We can also learn from the fact that drawings and paintings which are accepted as likenesses can contain contours and shapes that do not correspond precisely to contours and shapes in the real world. Matisse famously made a group of four self-portraits using line-drawing in which his project was to maintain a high degree of resemblance despite making each version as different as possible from the others.[5]

This lack of need for precise information can also be deduced from the fact

3 Although there are limits imposed by the sensitivity characteristics of visual systems.

4 The physics of absorption/reflection properties of surfaces alone makes this clear.

5 Theoretically one of these could be accurate but not all four. In practice none of them were.

that it is possible to recognise likenesses in highly abstract portraits, such as Picasso's cubist portraits of Gertrude Stein and Ambrose Vollard, or in newspaper cartoon drawings containing hugely distorted features.

So what can artists learn from the realisation that neither contour nor lightness relations can be used as the primary basis for recognition? First and foremost, they discover that the details of the shape and colour aspects of appearances which so many of those who aspire to literal accuracy labour so hard to capture, can be of no more than of secondary importance with respect to triggering recognition.

If it is not contours or surface properties, upon what aspect of appearances does recognition depend? To answer this question we need to find what can remain the same under a variety of different viewing conditions. More specifically what kind of description is capable of accommodating a possibly infinite numbers of differences of contour and surface-form into a unified description?

Although the fact that we all have the ability to recognise makes it clear that evolution has arrived at an answer, it took scientists a long time to puzzle out the main characteristics of what it might be. The reason for the delay was that the solution that eventually emerged depended on a conception of how the eye/brain works that could hardly have been imagined before recent advances in neurophysiology and computer modelling. The breakthrough had to wait until after the identification of visual-systems capable of separating out different components of the information contained in the light coming into our eyes.[6] Once this step had been taken, the solution was not long in coming. It comprises three parts:

- Multimodal processing.
- Analytic-looking cycles.
- Context.

Each needs a short section to itself.

Multimodal processing

As a sort of informal experiment I have asked many people over a number of years what the words "*round, red and slightly squishy*" first bring to their mind. The answer given by a huge proportion of them (I would guess at well over 90%) has been "*a tomato*". Easily the second most popular reply was "*a red rubber ball*". Other answers usually feature red fruit, but there are very few of these.

6 "*Introduction to the science*"

If I had asked separately, "*What is round?*" "*What is red?*" Or "*What is slightly squishy?*", I feel certain that tomatoes and red rubber balls would no longer have dominated the answers. Indeed they might not have appeared at all.

The point is that combining the three qualities that separately have minimal classifying power, brings us close to a plausible classification. True there is some ambiguity over what is the correct answer and therefore a need for making checkups that will decide between the different candidates but, as we shall see in the next section, evolution has made provision for that as well. The point being made here is that the round, red and slightly squishy example demonstrates the classifying power of cross-correlations between independently varying modalities of information. More generally, it suggests that unambiguous classifications can be arrived at on the basis of ambiguous information sources. If, to emphasise this point, we compare a large number of red tomatoes we will find that they vary considerably in shape, colour appearance and degree of squishiness. Thus, in the "*round, red and slightly squishy*" example we find that:

- The word "*round*" is being used to represent many possible shapes, none of which will be actually round.
- The word "*squishy*" to represent many possible degrees of squishiness.
- The word "*red*" is being used to represent a gamut of different reds.
- All three words are being used flexibly to cover many possibly examples of tomato sizes, shapes and other characteristics.

What this illustrates is that words that, when *used independently*, are vague enough to allow different views, colours and tactile experiences of objects to be taken as being the same, when *cross-correlated*, can provide relatively, or even totally, unambiguous classifications. Accordingly, the strategy of using *cross-correlations between independently varying modalities of vague information* meets the main requirement for achieving *recognition.* It is also difficult to see how this key achievement of our visual systems could be arrived at in any other way.

The usefulness of information provided by a nonvisual sensory system ("*squishiness*") in the process of arriving at the tentative "*tomato*" classification reminds us that, even for visual recognition, clinching cross-correlations can be provided by nonvisual sensory modalities of information (sound, smell, taste, touch, body temperature, heartbeat rate, etc.). In addition, recognition systems are able to make use of information coming from various kinds of memory store, cross-correlations with these can be used in refining classifications.

The analytic-looking cycle: What it does

It is important to remember that *recognition* is not a goal in itself, rather it is a staging post in a process of triggering *learnt actions*, whether they be described as "*habits", "skills"* or *"ways of thinking*". One of the most regularly used habits/skills are those that guide the eye, head and body movements that enable the focus of attention required for *analytic looking*.

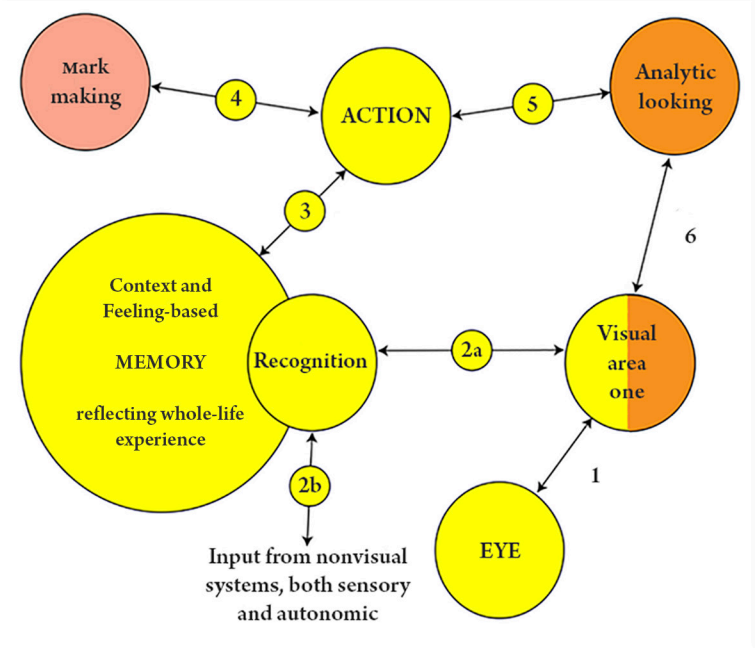

Figure 1 : The analytic-looking cycle.

The neural systems and linkages that enable ***analytic-looking*** are far too complicated to represent in a diagram. However matters can be simplified by concentrating on ***function*** as I have done my best to do in *Figure 1*, which shows how the processes that enable ***recognition*** tap into ***action instructions***, which have been previously stored in ***long-term memory,*** and that it is these that enable ***conscious visual analysis***.

It shows two pathways, both of which extract information from *Visual Area 1*, the first destination of the information coming up the optic nerve from the eyes. The lower of pathway (the one that enables *recognition*) is totally ***subconscious***. It makes use of multimodal processing to access relevant ***action-instructions***

stored in "***long-term memory***", or some other, more temporary task-specific form of memory.[7] The higher pathway mediates *visually guided actions*, including those required for all analytic looking.

Despite the enormously impressive classifying power of multimodal processing, as the "*round, red and slightly squishy*" example suggests, it invariably leads to some degree of ambiguity and can result in a great deal of it. This is a inevitable outcome of the dependence of our "*recognition systems"* on cross-correlations between crude descriptions.

Remember also that:

- The cruder the descriptions, the more different versions of the same object they can stand for.
- The greater the number of such descriptions that are being correlated, the cruder each of them can be to obtain the correct classification.

There is no need for artists to know exactly how these descriptions are structured, to realise that, even if they were to be made conscious (which, as already explained, they are not), they would provide little if any information in a form that would of direct use to the when drawing or painting from observation.

The loss of information due to the necessity of relying on crude descriptions could be catastrophic if it were not for the existence of a independent means of accessing information coming from the retina. That is why *Figure 1* is divided into two halves: One supplying the ***recognition systems*** and the other linking to the ***analytic looking systems***. This second pathway is needed for two reasons:

- It enables the confirmation or completion of classifications. Thus, if we take the example of the tomato that might alternatively be a red rubber ball, it is necessary to take a further look at the information coming from *visual area 1* to decide between the interpretations. This could be done on the basis of the degree of roundness (tomatoes are never exactly round), the surface-texture (rubber balls tend to be matt while tomatoes have a glossy sheen) or on colour (despite the many possible variations in their colour appearance, the pigmentation of a tomato can never be the same as that of a rubber ball). Once the tomato alternative has been selected, it may be useful to find out more about it. For example the analysis may be targeted via questions such as: "*How red*?", "*How round*?" or "*How squishy?"*. Any of

7 The memory stores known as *short-term visual memor*y and *working memory*.

these could trigger a secondary analysis by means of some combination of visual, touch and feel systems.

- It is necessary for guiding actions, such as the ones used for picking tomatoes off the vine, selecting them in a shop, preparing them for use in a salad or drawing them from observation.

The analytic-looking cycle: How it works

We are not finished with the ability of the eye/brain to loose information. Just as preconscious recognition processes are accompanied by necessary and significant information-loss, so are the analytic-looking ones.

This is because their functioning depends on the preparatory regularisation of the size, orientation, shape, lightness and colour of the recognised objects that is necessary for providing a degree of stability to the ***analytic looking process***. What is ***lost*** in the process of providing these ***constancies of visual perception*** is the ***actual relativities*** in all the above listed dimensions, providing a universally experienced cause of frustration for artists seeking accuracy in the representation of objects.

Fortunately, just as the two pathways system enables the eye/brain to keep hold of and make later use of the visually derived information that it must lose during the preconscious stages of recognition, it can also keep hold of the information about size, orientation, shape, lightness and colour that has to be lost in order to compute the constancies. That is can do so can be deduced from the precision with which it can guide the actions used for:

- Locating, grasping and manipulating objects.
- The comparative looking activity necessary for making judgements of relative lightness and colour.

It can be hypothesised that the computational basis of these skills makes use of the information provided by the degree of distortion required to achieve the required uniformity of size, orientation, shape lightness and colour. If so, it must be available fractionally before analytic looking can take place.

In addition to all this, there is the question of *directional-indicators*. Because multimodal processing flourishes on vagueness and ambiguity, information concerning what is *top, bottom, left and right* interferes with the efficiency of recognition processes and is, consequently, bypassed by them. Since, on the

contrary, knowledge of these properties of appearance is essential for analytic-looking, the eye/brain systems have evolved their way of inserting *directional indictors* so that they are available for use at the analytic looking stage of visual processing.

Context

All neural processing is ***context-based***, whether one looks at the basic mechanism or at their high level achievements. For example:

- The firing of any one neuron only has meaning in the context of the firing pattern of its neighbours and other interconnected neurons.
- A person that is easily recognised in a place that is associated with him/her, might not be recognised, or only recognised with difficulty, when encountered by chance in a foreign country.

The second example the facilitatory function of ***context*** comes in the form of the information stored in ***long-term memory*** that has been built up over a lifetime, from birth to a fraction of a moment ago. Without this resource and the ability to retrieve its contents, recognition could not happen. With it, the level of efficiency can be truly astonishing.

Elsewhere in this series I have used the example of a woman going to her wardrobe to find a red party dress which she knows is hanging in it. As she possesses no other red dresses, she does not need multimodal processing to find it for one modality alone (colour) would be enough to identify both it and its location. From this point onwards, the woman could almost certainly walk over to the dress, take it out of the wardrobe and put it on, all with her eyes shut.

In practice, she will keep them open for they can still be useful in finding answers to a number of practical questions. For example:

- *Which is the front and which the back?*
- *Do any buttons need undoing to allow the dress to be put on, etc.*?
- *Has the material acquired any untoward creases that will need ironing?*

Since each of these is asked in the context of existing knowledge of the dress, the search for an answer can be precisely targeted. Accordingly, in this situation, a minimum of cues will provide the information required to activate, first, ***recognition*** and, then, the ***action-instructions*** required for fulfilling the task in hand.

However, since this efficiency is based on a store of knowledge which is only capable of dictating strategies of looking that have been relevant to past contexts, there is a flip side to its value. If a new situation arises, the same knowledge can no longer help. For example, if the rejected dress was one of several dresses of the same red, the cue of redness alone would no longer enable certainty. More modalities of information would be required.

The efficiency trap and how to circumvent it

As already indicated, the triple strategy of ***multimodal processing*** for primary recognition, ***context-based memory*** for limiting potential alternatives and the ***analytic-looking cycle*** for confirmation or completion of object classification enables the human visual system to perform a vast panoply of tasks with formidable efficiency. However, even the highest levels of efficiency have intrinsic limits. The explanation why this is so is based on two interrelated facts, namely:

- That efficiency is the product of the "*useful coordinations of habits*" normally referred to as "*skills*".
- That all habits, whether described as "*good*" or "*bad*" or whether associated with "*action*" or with "*thought*" can only be based on knowledge deriving from past experience.

From this we can conclude that an intrinsic property of ***skill*** is that of not being well adapted for coping with new, ***unfamiliar*** situations. This must be true even for the skills that the eye/brain has developed for dealing with novelty, including those recommended in this book as aids to personal expression and creativity. No wonder *Modernist Artists* have seen skill as a major enemy.

Luckily for the human race, the processes of evolution ensured that brain systems have a way of responding fruitfully to *novelty*, all of which entail the use of *same/difference judgements* between similar but different visual inputs. Most of these are generated in one of two ways:

- Situations in which performance does not match up with intentions, as when a drawing that is intended to be an accurate copy fails to achieve its goal. The luck lies in the fact that, by their nature, the mismatches, commonly referred to as *"mistakes",* reveal *differences,* and that the differences are never predictable.
- Movements of body, head or eye relative to an object being analysed which call attention to any variations in appearance within a region under

analysis. Without either mismatches or movement, the feedback necessary to trigger the processes which underpin learning and creativity would not be available.

Another approach to demonstrating the intrinsic limits of ***efficiency*** starts with an attempt to define the word. How about, "*Completing a job to the required standard in the least possible time and with the least possible effort*?"

When applied to ***analytic-looking*** this definition can be rephrased as "*reducing the time spent on analysis to a minimum*" or, more colloquially, as ***looking as little as possible***. For example, the woman finding the red dress in the wardrobe could manage perfectly well on the one brief glance needed for establishing its position. No need to expend any more effort. A further example would be eight year old children who, when asked to draw familiar objects from observation, limit themselves to one or two brief glances.[8] Once they have recognised what it is they are looking at, they can complete their drawing without further reference to it.

Nor is this essentially lazy form of looking confined to children. Skilled artists too can use their knowledge-base to minimize the necessity for active analysis. In the case of the children, it is usually obvious to viewers that their drawings fall woefully short of literal accuracy. In contrast, some skilled artists using a similar minimal looking strategy are able to deceive people into thinking that they are drawing what they see in front of them. However, for a number of reasons, they will almost certainly be deceiving themselves. All of these are related to the fact that no matter how familiar an object:

- The image of it arriving at the eye is always unique.
- Since edges have no width, the lines used to represent them in depictions can never replicate real-world appearances
- Since no colour in nature can be matched by mixture of paint colours, these can never replicate real-world appearances

All analytic-looking is knowledge-driven and the all knowledge is based on generalisations that can never match the uniqueness of the object being analysed since they are derived from past experience of it. In as far as artists can provide the kind of visual equivalence to this uniqueness that we call *"accuracy"*, they can only do so by means of assiduous, time-consuming use of the analytic-look-

8 Pratt, 1983, " *A perspective on traditional artistic practices*" in Freeman and Cox (Eds) "*Visual Order*"?

ing cycle and using trial and error to match of the contents of their painting with the scene it is supposed to represent..

Artists need to keep this truth in mind at all times, particularly if they are experiencing no difficulty in achieving what they believe to be accuracy. As Degas asserted, "*It is necessary to assume that I know nothing for it is the only way to make progress*". He also said enough about the limiting effects of skill to make it clear that, when he made this statement, he was making a plea against blind reliance on any seeming efficiency of knowledge-driven looking strategies.

However what Degas does not seem to have realised is that *"knowing nothing"* is impossible: Artists making drawings from observation have no option but to use knowledge-driven looking strategies, since no targeted looking activity can take place unless planned in advance.

This is why, if we seek to approach as near as possible to accuracy, we need to focus attention on the aspect of visual perception where knowledge and appearances come nearest to corresponding to one another. It should be no surprise to find that this necessitates looking in terms of parts (the role of analytic-looking) rather than wholes (requirement for recognition). More precisely, it means going down to the level of abstract relations until the necessary correspondence is found, and then making comparisons.

The problem that then arises is to work out how something that has been broken down into a collection of simple elements can be reconstructed into never-before-seen configurations and relationships. Ways of finding solutions has been a main subject of the two books on the practice of drawing and painting.[9]

The music of colour and line relations

The skill of being able to deconstruct what we see into simple parts and relations between them, and then to reconstruct them into accurate wholes has value beyond that of achieving high levels of realism. It is only when we concentrate on abstract relations that we can become aware of what has been described as "*music of colour, texture and shape*", in which visual journeys from region to region across the picture surface provide the melodies and whole-field relations provide the chords. No wonder so many *Modernist Painters,* being fascinated by these new dimensions of experience, turned to abstraction.

9 *"Drawing on Both Sides of the Brain"* and *"Painting with Light and Colour"*.

Implications

How far does a combination of this chapter and earlier ones help us to pin down the nature of what I have been calling "experienced reality"? I am afraid the answer has to be,"very little", for what we mean when we use the word "see", will always remain inherently elusive. There is no escaping the fact that this intimate aspect of our daily lives is in constant flux. First, we are confronted with one of the most fundamental truths of visual processing, namely that, in order to look at anything, we have to take it out of context. Second, we find that, when we do so, the object of our attention is subject to the 'constancies', a state of affairs which catastrophically disrupts whole-scene-relations. Third, the only way of approaching the task of finding out what these relativities actually are is by means comparative looking. True, this provides us with a succession of instances of what, in some senses, can be called reliable information about differences. But it only does so at the expense of finding a new starting point for each comparison, and this is a manoeuvre that entails resetting the whole system to what is almost certain to be a different lightness range.

One thing this inherent instability of visual perception means is that, when we try to analyse a scene (including if it is a painting), we will be facing a problem in some ways analogous to representing running water with a still image. No doubt this is what made Cézanne describe painting as "So damned difficult".

But, perhaps finding a definitive solution is not what we really want. As Robert Browning wrote, "Man's reach should exceed his grasp or what's a heaven for?" The advantage to artists of the ideas presented in these books is that they can at least help us struggle fruitfully towards this goal. It may be unattainable but, as many of us can testify, groping our way towards it can be a richly rewarding experience. Maybe, this is why Cézanne followed his statement about the difficulty of painting with the heart felt assertion, "I want to die painting."

Differences relative to "measured reality"

Despite the elusiveness of "experienced reality", we can nevertheless identify some of its main characteristics. Not the least important of these is the degree to which it diverges from "measured reality". Three reasons for this divergence are fundamental to the ideas presented in this book. Thus:

- *Differences in appearances due to variations in viewing conditions are ironed out by the various constancies.*

- *Lightness-contrast and colour-contrast effects exaggerate differences at the common borders between regions of different colours.*
- *Despite its enormous importance with respect to everyday visual experience, the part of the light that is reflected from surfaces without entering into them is never available for direct analysis, except as the false body-colour we describe as shadows, shading and highlights.*[10] *If we want to give drawings and paintings the qualities of surface, space and light that experienced reality provides, we have no option but to turn to theory and experiment.*[11]

But, perhaps even more significant than any of these ever fascinating and elusive features of appearance is the influence of the processes that enable "recognition". These take us a long way indeed from "measured reality". In doing so they liberate artists to explore not only exaggeration, distortion and abstraction, but also the value of inaccuracy as a route to artistic creativity.

Matisse's self-portrait project (described earlier) confronts us with two questions of fundamental importance for any artist interested in achieving the goal of creating what people will describe as a "likeness". These are:

- *What are the minimal cues necessary to provide it?*
- *To what extent and in what ways can inaccuracy help or hinder its realisation?*

Both are of particular importance in the next chapter.

10 See"*Painting with Light*", PART ONE of this volume.

11 Or, blindly follow the rules of Marian Bohusz-Szyszko.