CHAPTER 15

Other constancies

Introductory

Colour constancy is by no means the only constancy. There are many others and all are fundamental to the ability of the eye-brain to make practical use of visually acquired information. Paradoxically, although their name suggests stability, they are responsible for the veritable "shifting sands of appearance" which, in its various guises, constitutes one of the main problems for artists seeking to obtain accuracy in drawings or paintings from observation. This is because they ensure that, when we look separately at any two similar features of appearances whether they be whole objects, parts of objects, sections of contour or colours, there is a very strong tendency to see them as being more similar to one another than objective measurement would dictate - often a great deal more so. Our visual systems upset every external relationship by relentlessly forcing everything towards normative dimensions and values. As a result, the constancies involve enlarging and diminishing, squashing and stretching, revolving, darkening and lightening and modifying colour. Any list of the constancies of particular interest to the artist should certainly include (a) ***size constancy,*** (b) ***shape constancy,*** (c) ***orientation constancy*** (d) ***lightness constancy*** and (e) ***colour constancy***.

The approach

Because the constancies provide so many obstacles to achieving accuracy in drawing from observation, drawing teachers are regularly required to help students to cope with their consequences. In the practical situation of the art class, it is seldom feasible or even desirable to go into too much theoretical detail, particularly when the theory is still in the process of being constructed. Long-winded and pedantic clinging to scientific accuracy needs to be replaced by confidently presented, easily understood and inevitably oversimplified ways of explaining the implications of psychological realties and their and neurophysiological correlates.

Before turning to the constancies themselves, it appropriate to set the scene for them by means of a short explanation of where they occur in the analytic cycle and how this effects their influence on visual perception. For this reason I start with sections on "recognition", "long term memory", "analytic looking", and "information storage".

Recognition

As explained in several places, from an evolutionary perspective, the viability of visual perception as a practical tool was dependent on the *eye-brain* finding a solution to the fundamental problem of *recognition*, namely that of being able to classify things that are actually different as being the same. This is an imperative because, in nature no two views of the same object, object-class or abstract quality of appearance are ever identical on two occasions. The eye brain achieves this feat by means of multimodal processing, using the temporal lobe route.

Long term memory

Recognition is not an end in itself, but a means of accessing *Long term Memory* (LTM) and using it to generate instructions as how to react to whatever it is that has been recognised.

Analytic-looking

The organising of these action instructions may not require any further visual information, but it usually it does. If so, the analytic-looking systems will be activated. Their three purposes are:

- To confirm the finding of the recognition systems.
- To decide between alternative possibilities.
- To participate in interactions with the recognised object or situation.

The first two of these can be done, by referring to information stored in short-term visual memory (*STVM*), possibly after further consultation of the exterior visual world. The third involves interactions with the exterior visual world.

Information stored in *STVM*

The analytic-looking stage of visual processing is the only one that involves consciousness. It is this that provides artists with the information they need for drawing from observation. The question is how to access it in a usable form. We know that when it arrives up the optic nerve from the retina:

- The pattern of inputs is constantly changing in step with the ebb and flow of the tapestry of light-borne information entering the eyes.
- The systems that achieve colour-constancy and lightness-constancy have not yet played their part.

We also know:

- That none of the visual systems that mediate *recognition* provide information that could be used directly for drawing from observation.
- That *STVM* contains information temporarily protected from new input from the retina and that it is only because of this protection that we are able to experience a stable visual world.

Making use of the stabilised information.

The next three questions are:

- What visual systems are able to make use of this stability?
- What do they actually do?
- How do they do it?

The only plausible answer to the first question is the *analytic-looking systems*. The answers to the other two require another section.

What do they do?

As a means of getting a better idea of how the sub-systems that make up the analytic-looking system make use of the stabilised information in *STVM*, we need to go to the evidence provided by studying the *constancies*. A rich source of it comes from the experience of drawing from observation.

As indicated above, the *constancies* are a group of phenomena that cause objects or parameters, when viewed separately, to be perceived as being more similar to one another than measurements of differences between them at the level of the retinal image would predict. Clearly, this systematic shift in appearances would make life difficult for artists wishing to achieve accuracy in their drawings from observation.

In order to provide my students with a simple explanation of the nature of the distortions, I cut out complications due to the massive interconnectivity between brain systems and posit a specific region in the parietal area of the brain.[1] It is proposed that this hypothetical region has a number of fixed parameters and that, with every change in the focus of attention or the point of view, everything is adjusted automatically to fit into these. In other words, all parameters are revised towards some norm. Thus:

1. *Axes of symmetry* are tipped towards the vertical or the horizontal
2. Features of appearances are stretched and squashed to make them more similar in *size*.
3. Distortions to shapes due to the angle and distance from which they are viewed are rectified in the direction of their physical shape. For example, rectangular shapes perceived as trapezoids are rectified in the direction of rectangularity and circular shapes perceived as ovals are rectified in the direction of roundness.
4. The range of lightness levels[2] of whatever it is that is being focused upon is stretched or squashed to fit into a standard range.
5. The lightness profiles of surfaces are removed, leaving all regions of similar pigment colour as being uniform in appearance (*spatial colour constancy*).

The constancies in the art class

It is these and other distortions, caused by the fact that the visual world we experience is a creation of the eye/brain, that provide the shifting sands of appearance with which artists have to deal. They lie at the root of a large proportion of the problems that face students in the art class. Attempts at accuracy in drawing from observation are frustrated by the normative tendencies listed above. Thus, depictions of objects undergo various typical distortions. Below is a reprise of the list with examples. Thus, when we look at features of appearances:

1. *Their axes of symmetry, if not already vertical or horizontal, are likely to be depicted as being more so.*

   For example, in the life drawing class, sideways-leaning heads and the features within them are regularly made too vertical and/or too horizontal.

1 Based on the evidence from the unilateral neglect studies that it is here that an axis of symmetry is put in place. The probability that the colour constancy computations take place elsewhere is glossed over.

2 Number of just noticeable differences (JNDs) between the lightest and darkest parts of the scene.

2. *Objects are squashed or stretched in terms of size such that, unless directly compared, objects of different sizes are drawn as being more similar than the extent of their image on the retina would justify. This is the case whether the size difference is due to physical dimensions or the distance of the objects from the viewer.*

   For example, in the figure-drawing class, the eyes, the nose and the mouth tend to be drawn as too big relative to the head, and the head, as too big as compared with the body.

3. *Distortions due to viewing angle and distance are straightened out.*

   For example, the approximately oval appearance of the round top of a mug seen in perspective will tend to be drawn as being more circular than its retinal image. Similarly, the trapezoid shape of the top of a cube seen in perspective will tend to be drawn as being more rectangular than its retinal image.

4. *The lightness range within objects and/or whole scenes is squashed or stretched*

   Thus, when artists set about making an accurate representation of the lightness relationships within an object, their eye/brain systems both (a) extract the object from its context, and (b) stretch or squash its lightness range to fit into the prefixed parameters provided by the system. This causes the artists a problem because, whenever the lightness range of the whole scene, including both object and context, is greater than the lightness range of the object on its own, the darkest and lightest parts of the isolated object will be perceived as darker and lighter respectively than they would be if viewed as part of the whole scene.

   Accordingly, if artists depict what they see when drawing or painting objects of a given lightness-range separately from no matter what context with a greater lightness-range (which will very frequently be the case), they will render (a) the shadows as being too dark and (b) highlights as being too light relative to the darkest and lightest parts of the rest of the scene in which they are situated.

   The fact that this happens so frequently in the drawings and paintings, not only of students but also of professional artists throughout the centuries, provides strong evidence that the strategy of making comprehensive lightness comparisons between the object and context does not come naturally to our species. Likewise, it suggests that artists have either not been as aware of or have not realised the implications of the fact that each new act of visual analysis entails a resetting of the lightness-range. If so, it can be expected that those of them who set themselves to paint or draw what they see by attending to objects in isolation, will find that the lightness relativities in the resulting depiction to be comprehensively awry.

This point is so important that it needs examples. Consider the case of artists working on head and shoulder portraits. One of the requirements is to assess to correct lightness relationships between shadows and highlights on the face. That this may not be as straightforward as it seems is evidenced by the seemingly universal tendency to exaggerate the difference between the two, and to make shadows too dark and highlights too light. The homespun explanation I give to students is that when they are looking at relationships within the face alone, their eye/brain systems give the darkest shadow the lowest possible rating, namely zero, and the brightest highlight, the highest possible rating (for arguments sake, I suggest one hundred or some other arbitrary figure). Accordingly the artists would depict the shadows as very dark and the highlights, as very light. This would be fine so long as the context is not brought into the equation.

However, if the context is taken into consideration, matters can change dramatically. For example, the darkest shadow on the face (being lighter than the darkest shadow in the context) would have to be given a rating of more than zero (for arguments sake, say forty), and the lightest highlight being less light than the lightest region in the context would have to be given a rating of less than one hundred (for arguments sake, say seventy). In this case, the difference between looking at object and context separately and looking at them together would be very dramatic. In virtually all cases in which objects are analysed in natural scenes, the difference will be significant.

The universality of the tendency, either to give scant attention or to completely ignore context when subjecting objects to analysis, can be seen from the regularity with which the shadows and highlights in the depictions of faces appear as both much too dark and much too light relative to the darks and lights in the context. The widespread occurrence of this phenomenon, even in the work of well known artists, provides a strong indication that the long-term memory stores of the artists concerned have no provision for triggering the instructions necessary to solve the problem. What would be needed to have in storage, would be instructions that would cause the artists to make systematic comparisons between the degrees darkness and lightness of all the shadows and highlights on the object with the darkest and lightest regions in the scene.

One way of describing the extent of both within-object and within-whole-scene lightness relationships, is in terms of the *lightness space* they occupy. Thus, objects will almost have a less extensive lightness space than will have the scene within which they are situated. Expressed in these terms, artists who analyse within-object lightness relationships separately from within-whole-scene, lightness relationships, will almost certainly find that they are painting themselves out of the *lightness space* required for depicting the correct *whole-scene lightness relativities*.[3]

5. *The squashing and stretching of the lightness range has another important consequence*:

   - It means that people are more sensitive to colour/lightness differences in scenes, where the range of lightnesses is relatively less, and relatively less sensitive, when the range of lightnesses is relatively more. This is because, as already indicated, whatever the extent of the of the *lightness-continuum* in any scene, the eye/brain systems create a similar number of steps within it. The desire to exploit this extra sensitivity is one of the reasons why Modernist painters have tended to work with more equal lightness colours than their predecessors.[4]
   - Some people might to wonder why paintings of natural scenes, which always have a far smaller intensity range than the scenes themselves, are not perceived as being lacklustre. The reason is that the eye/brain systems give both the painting and the natural scene the same lightness parameters. Thus, both are perceived as containing the same number of lightness steps and, accordingly, as being more similar in terms of their lightness ranges than measurements based on reading from a light meter would predict.

6. *The information-rich modulations of the light reflected from the surface*

3 For a famous example, see "*Cézanne falls short"*<http://www.painting-school.com/cezanne-falls-short/>

4 In 1967, my tutor at art school told me that if I wished to be a serious colourist, I would find it necessary to work with equal lightness colours.

of paintings are separated from the body-colours within them, without ever being made conscious.

As explained in the previous chapter, a major experiential consequence of this separation is s*patial colour-constancy,* the phenomenon by which we see colours as flatter and more fully saturated than the variations and combinations of wavelengths coming into the eyes would predict. This explains why, unless a strategy for avoiding doing so has been learnt, artists will depict the surface as being *too uniform in colour.* But what would such a strategy consist of, if every time we look at colours we are deprived of lightness information relating to them? The answer to this question lies in the fact that the eye/brain's method of acquiring information about the extent of lightness differences is to monitor the changes that occur during the process of transition between sequential acts of analysis. In other words, the only way of making judgements of lightness difference is by making *comparisons*.

Interestingly, the fact that it took the community of artists so long to master chiaroscuro is an indication that the practice of making the necessary whole-field lightness comparisons does not come naturally, but has to be learnt.

7. *The colour of a surface is likely to be painted as being the same when painted on different occasions under different lighting conditions. The cause is temporal colour-constancy.*

With these drawing-class examples to consider and keeping in mind my ad hoc explanation positing a parietal area location where the *constancies* are preconsciously imposed, let us now revisit and elaborate on some of the *constancies* introduced above.

Orientation constancy

If an artist depicts a slightly leaning post with parallel sides as being vertical nobody will be confused if their error is described as due to his or her being influenced by *orientation constancy*. But, this conclusion fails to confront the issue as to what is being oriented. Significantly for artists when drawing from observation there are at least two, seemingly conflicting possibilities. One is that it is the *edge* of the post, the other, that it is its *axis of symmetry*. To decide between them consider the case of an artist drawing a trapezium who tips it away from its orientation on the retinal image, in one direction or another. Luckily the discussion of the *Bakery facade illusion* discussed in *Chapter 10*, provides examples of both possibilities. The illusion that the top edge of the wall slopes down rather than sloping very slightly up as illustrated in the photo and decreed by the laws of linear perspective, depends on seeing it the context of the trapezoid house-front, while seeing it as horizontal requires isolating it from its context. In both cases the distortion is caused by rotating the axis of symmetry towards the horizontal but in the first case the axis is that of the two dimensional trapezium, while in the second case, it is that of the one dimensional edge. Artists need to be aware of the both possibilities and take the necessary precautions against either of them upsetting the relationships in their depictions.

As just explained, *orientation constancy* causes the axes of symmetry of shapes and edges to be perceived as being either more vertical or more horizontal than they would be if compared to a plumb line or a spirit level. A prerequisite of this rotation taking place is the isolation of the object, shape or edge under analysis from its context. Whenever this occurs, all relativities with the context are lost, including orientation relativities. Here are two implications of this loss:

- Once an object of interest has been isolated and rotated, analytic-looking strategies are influenced by *knowledge* of the type of object to be analysed. In other words, it is information, stored in long term memory, that determines the instructions that guide both the *analytic-looking* and the *eye-body coordination* strategies required for interacting usefully with the object concerned.
- A standard way of helping beginners with drawing from observation is to suggest they look for *negative shapes* as a way of bypassing the risk of being influenced familiarity.[5] That this works shows that people have knowledge that can guide the analysis of what are taken to be meaningless configurations. In other words it can be knowledge-guided. Indeed, bringing the unknown into the known is a necessary requirement of all learning, and a primary function of analytic-looking.

To recap, there are three preconscious processes that will always precede the imposition of an axis of symmetry by the eye/brain. These are:

1. Isolation from context.

5 Since the eye/brain can only look as shapes for which it has knowledge, this is a misnomer. See chapter on negative shapes in my book "*Drawing on Both Sides of the Brain*", which can be found at < http://www.painting-school.com/negative-spaces/>

2. Access to relevant knowledge via recognition.
3. A record of location.

These three processes and the putting in place of appropriate axes of symmetry are necessary prerequisites for all analytic-looking tasks.

Size constancy

One way of introducing the concept of *size-constancy* involves looking through the viewfinder of a camera with a telephoto lens. Since this can be adjusted to make distant things look larger and near things to look smaller, it is easy to make the image of any object fill the viewfinder frame from any distance, nearer or further. An analogous result is achieved by the eye/brain systems. In this case, the image of an object, when made larger or smaller as consequence of being perceived from a variety of distances, is stretched or squashed into the same brain-created visual space of an unchanging prefixed size. This is why we tend to perceive the same object as varying less in respect of its relative size than does its image on the retina.

Lightness constancy

The eye/brain achieves lightness constancy in a number of steps:

- It adjusts *pupil-size* according to the amount of light striking the eye at any one time. This enables the regulation of the strength of light that reaches the light-sensitive receptors in the retina.
- The receptors in the retina become more sensitive as the level of illumination decreases and, correspondingly, less sensitive as it increases. Like the control of pupil size this pushes matters in an equalising direction.
- Any multilayered neural system based on *lateral inhibition*, such as those that underpin so much of eye-brain processing, computes relativities rather than absolutes. This produces a tendency in an equalising direction for all modalities of input, including lightness.[6]

However, an equalising tendency is not the same as constancy. Like all the constancies, *lightness-constancy* is associated with conscious, analytic-looking and, like them, it is achieved as a part of the preparation for this key function of visual perception. As with the other constancies, there is a region in the parietal area of

6 A more detailed example of this idea is given in the previous chapter, dealing with the colour constancy algorithm

the brain, that stretches, squashes and/or orients varying inputs into a standardised framework. It is there that lightness inputs are stretched and squashed such that all scenes or objects under analysis, whatever their actual lightness range, are fitted into the same range, this producing *lightness-constancy*.

Spatial and temporal colour constancy

Our ideas about how the eye/brain systems compute these two manifestations of colour constancy are to be found in the previous chapter. In summary, the information concerning the gradations in *reflected-light* is separated out from a remainder that includes information on both *ambient illumination* and on *body-colour*.

The reflected-light profile gives information about four properties of surfaces. These are:

1. *Surface solidity*.
2. *Surface-form profile*.
3. The *borders,* both between juxtaposed regions of body-colour and between cast shadows and adjacent regions of body colour. In all cases, these are revealed by large increments or decrements in the lightness profile.
4. *In front/behind relations*, which are deduced from abrupt and simultaneous increments or decrements in both *lightness-profile* and *body-colour* that occur at the divide between objects and their contexts.

The algorithm that makes use of the variations in reflected-light profile to produce the above information, leaves a remainder. If is allowed to run on it will separate this into two components:

- A non-varying remainder of the reflected light.
- An absorption/reemergence profile of the light that interacts with the pigmentation within the surface.

As explained in the previous chapter, by doing so it provides information that enable us to experience both *ambient illumination* and *body-colour*. As a significant bonus, it does so in a way that ensures that we will experience *body-colours* as remaining stable, despite variations in lighting conditions. In other words, it enables *temporal colour-constancy***.**

Simultaneous lightness contrast and simultaneous colour contrast.

Two other linked factors that regularly influence appearances are:

- *Simultaneous lightness contrast*
- *Simultaneous colour contras*t.

Both are intrinsic to the analytic-looking experience, as well as being the cause of serious problems for artists. Both are consequences of lateral inhibition in the retina and both take the form of an exaggeration of contrast at the borders between adjacent regions of colour.

While both phenomena occur at every border on all occasions, the *colour-contrast* effect is often, to a greater or lesser degree, obscured by the relatively stronger lightness-contrast effect. This is why when artists, who seek to create vivid *colour-contrast effects*, are likely to juxtapose colours that are both approximately equal lightness and well separated on the colour circle, for example, reds and greens or oranges and blues.[7]

Problems for artists due to the simultaneous lightness contrast.

From the above, it is easy to see how *simultaneous lightness contrast* effects could make life difficult for artists when attempting to render the whole-field lightness relations in their paintings. In essence, their problem is that they are confronted by local expansions and contractions of lightness space that are necessarily incompatible with the lightness space of the scene as a whole. Accordingly, if they try to paint what they see, they will find that the parts will not fit convincingly into the whole. As is explained in *"Painting with Light and Colour*", their only recourse is to tone down the lightness differences they see at borders between adjacent regions of different colours and rely on the eye/brains of spectators to recreate the *simultaneous lightness effects* when they view the paintings.

Implications

For scientists of visual perception, who will presumably be well acquainted with the phenomena discussed in this chapter, the main interest of its contents may lie in the examples of how they effect the everyday daily visual world of art-

7 The phenomenon of "*simultaneous colour contrast*" was first described by Chevreul in 1837. He summed up his observations concerning colour contrast effects as follows: "*in the case where the eye sees at the same time two contiguous colours, they will appear as dissimilar as possible, both in their optical composition and in the height of their tone. We have then… simultaneous contrast of colour.*" In particular, he called attention to the degree of mutual enhancement when complementary pairs are juxtaposed.

ists. If so, their treatment in my book "Painting with Light and Colour", will be of even more interest.

For artists, whether they are engaged in painting or drawing from observation or whether their interest is in exploring pictorial dynamics for their own sake, knowledge of how the constancies and the simultaneous contrast effects influence appearances is of evident importance.