The knowledge used in vision and where it comes from

HORACE B. BARLOW

Physiological Laboratory, University of Cambridge, Downing Street, Cambridge CB2 3EG, UK

SUMMARY

Knowledge is often thought to be something brought from outside to act upon the visual messages received from the eye in a 'top-down' fashion, but this is a misleadingly narrow view. First, the visual system is a multilevel heterarchy with connections acting in all directions so it has no 'top'; and second, knowledge is provided through innately determined structure and by analysis of the redundancy in sensory messages themselves, as well as from outside. This paper gives evidence about mechanisms analysing sensory redundancy in biological vision. Automatic gain controls for luminance and contrast depend upon feedback from the input, and there are strong indications that the autocorrelation function, and other associations between input variables, affect the contrast sensitivity function and our subjective experience of the world. The associative structure of sensory message can provide much knowledge about the world we live in, and neural mechanisms that discount established associative structure in the input messages by recoding them can improve survival by making new structure more easily detectable. These mechanisms may be responsible for illusions, such as those produced by a concave face-mask, that are classically attributed to top-down influences.

1. INTRODUCTION

In discussions among the organizers before this Discussion Meeting, Geoff Sullivan defined the knowledge used in knowledge-based vision as follows: 'It is a store of extra facts and relevant information that is brought to the task of interpreting the image by the perceptual or computational system'.

On this view knowledge-based vision is practically synonymous with 'top-down' vision, but this article argues that we should take a broader view of the nature of knowledge and where it comes from. To be useful in vision, knowledge must tell us how objects, events and transformations in the outside world influence visual images. Such knowledge can be derived from different sources and introduced into the system in different ways.

(i) The knowledge can be obtained by experience of 'what works' and embodied in the structure of the visual system: in a biological system, evolution selects genes that form a well-functioning system that aids survival.

(ii) The knowledge can be obtained from a library or another source of communal knowledge and can be used to modify how the machinery works: this could take the form of debugging a program or entering a set of required parameters, and it can be likened to teaching or training a biological system.

(iii) The knowledge can be acquired from the images by an analysis of their redundancy and used to change the way the incoming messages are coded. This possibility arises from the fact that objects, events and transformations in the outside world influence images by modifying their structure or redundancy, so some at least of the requisite knowledge can in principle be derived from analysing the images themselves.

Of course it remains an open question which is the best way of obtaining and using knowledge, but the history of photography illustrates how information obtained and used in the third way can supplant the second method.

2. KNOWLEDGE-BASED PHOTOGRAPHY

Until quite recently photographers needed to bring a large amount of knowledge to the task before they could take good photographs in a wide variety of conditions. They had to know which film to load for the photographs they planned, then they had to take into account time of day, season of the year, amount of cloud cover, type of scene, and apply all this knowledge to decide the aperture and exposure. Of course they also had to know how to mix developer and fixative and how to use them, but just consider the first step—taking the photograph—for this aspect of the photographer's knowledge has now been built into the camera and split in two: part of it is embodied in a mechanism within the camera itself, and the other part uses 'bottom-up' information provided by photocells. It took a few years for modern automatic exposure systems to evolve, but they now incorporate much of the knowledge the photographer himself formerly brought to the task, and they are able to do this because they use input signals about the luminance and contrast of the scene derived from the light striking the camera. Measurements made on the sensory messages themselves can provide much knowledge that formerly had to come 'from the top'.

This paper presents evidence from human vision and the neurophysiology of visual pathways which strongly suggests that sensory messages are indeed analysed in the way required to provide this knowledge. Visual messages are not random and incoherent, that is they have a great deal of structure or redundancy, and if this redundancy is taken account of and used properly visual sensitivity will be improved and visual performance will be quicker and more reliable.

Some of the structure in visual images is almost always present and can be handled by mechanisms that are permanently in place: such genetic or inherited knowledge corresponds to the first way of introducing knowledge defined at the start of this article. Examples of the second, taught or learned knowledge, might be finding out what snakes are poisonous by referring to a book, or predicting the perspective transformations of a three-dimensional object by applying the appropriate equations, but notice that it is not always easy to decide which type is involved in a particular task. There are probably innate mechanisms making us afraid of snakes and spiders, and it is very likely that our visual system is aided in recognizing and applying perspective transformations by innately wired connections.

The third way of introducing knowledge, through analysis of the incoming messages themselves, is the main focus of this article. In mammalian vision the coding of incoming messages is modified according to properties of the images the visual system receives during an early 'sensitive period' of development (Movshon & Van Sluyters 1981). In later life there must also be modifications when a single glimpse of a pattern allows us to recognize a subsequent appearance, or when we are able to remember a tune after a single hearing, but there are much simpler mechanisms at earlier levels in the visual pathway, so let us start with these.

3. ADAPTIVE MECHANISMS IN VISION

Figure 1 (Barlow 1969) shows how retinal ganglion cells adjust their response characteristics when the mean luminance to which they are exposed is changed. Essentially they normalize their response to the local mean luminance of the image, and the on-centre units signal increases from this mean while the off-centre units signal decreases. This exemplifies how the system discounts an elementary piece of knowledge (local mean luminance) in order to make better use of the dynamic range of the neuroneal signalling system. Similar gain control systems are of course universal in video cameras, and physiologists are hot on the trail of the biophysical mechanisms involved. The principle is so well known that one hesitates to call it 'knowledge-based vision', and we shall not be concerned with it any further here, but it is definitely an example where 'relevant information . . . is brought to the task of interpreting the image', and the required information is not 'top-down' but is right there in the input.

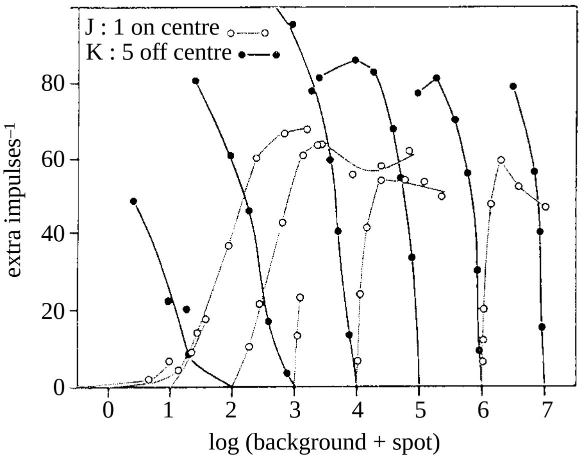

Figure 1. How mean luminance is discounted in the retina. This figure shows the response characteristics of an on-centre and an off-centre ganglion cell from the retina of a cat adapted to different mean background luminances. Log (mean luminance) was set to one of the seven values shown on the abscissa, then the luminance of a one degree central spot was changed to the value whose log (luminance) corresponds to the abscissa of one of the dots on the lines connected to the adaptation abscissa value. A small increment in luminance evokes extra impulses from the on-centre cell and a small decrement evokes them from the off-centre cell. The cell's whole characteristic is shifted according to the mean luminance it has adapted to (from Barlow 1969).

In addition to mechanisms adjusting for luminance variations, and discounting them in the output, there are mechanisms in the primary visual cortex (Ohzawa *et al.* 1985) and to some extent earlier in visual pathways (Smirnakis *et al.* 1997) that discount changes in the range of contrast of the incoming signals; again the neurophysiological mechanisms are being uncovered (Carandini & Ferster 1997). There may be more to this process than is fully recognized (see below and Geisler & Albrecht (1995)), but a feedback signal for contrast is perhaps again too simple to qualify as 'knowledge'.

Now consider the form of adaptive change in coding shown in figure 2. What is shown to change here is not just the sensitivity to increments or decrements of light, but the form of the contrast sensitivity curve. It changes from a low-pass form at low mean luminance to a band-pass form at high mean luminance, and this change is thought to result from the greater effectiveness of lateral inhibition at higher background luminances (Barlow *et al.* 1957). It is generally supposed that the added low-frequency attenuation compensates for the greater amplitude of low spatial frequencies in natural images (Field 1987; Tolhurst *et al.* 1992; Ruderman 1994); so what the visual system seems to be doing here, as with luminance and contrast adaptation, is to discount automatically a constant feature in the pattern of sensory stimulation. However, here the feature discounted is a more interesting property of the input: the shape of the power spectrum depends upon the autocorrelation function and is not a simple mean,

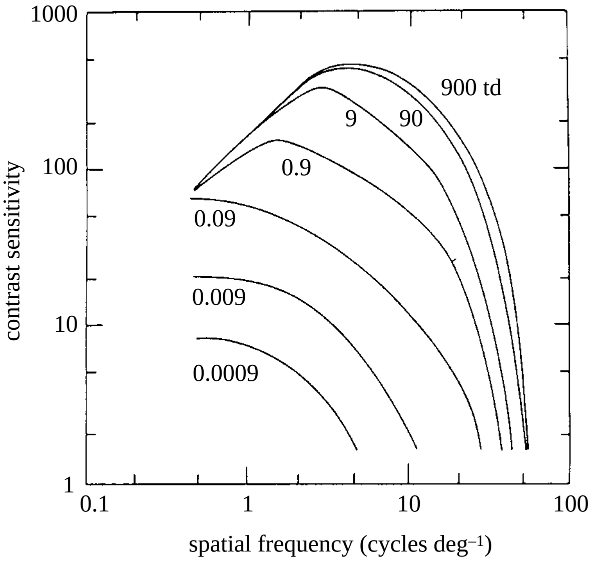

Figure 2. The shape of the contrast sensitivity curve of human subjects changes with mean luminance. Contrast sensitivity is the reciprocal of the minimum detectable modulation of a spatial sine wave whose frequency is plotted as abscissa and whose mean luminance is given on each curve. The change in shape is thought to compensate for the excess of low spatial frequencies present in the power spectrum of all natural images, and since the power spectrum is determined by the autocorrelation function, this suggests that the visual system takes account of an associative property of the images it receives (from Van Nes & Bouman 1967).

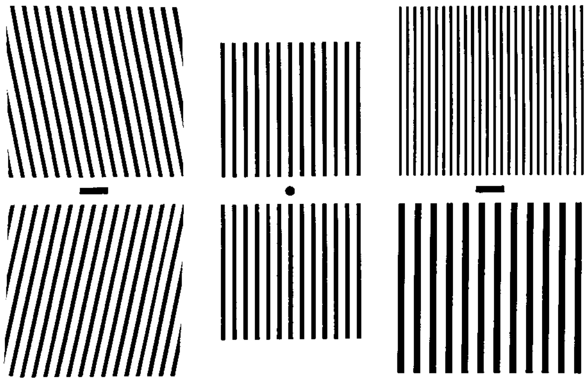

Figure 3. First look at the dot between the central pair of gratings and confirm that these are both vertical and have the same spacing. Then place the figure at a distance of about 30 cm and fix your eyes for at least 30 s on the black strip between the left-hand pair of gratings, moving them along the strip but not outside it. Next transfer your gaze to the central black dot and compare the upper and lower gratings; for a brief period the upper should appear tilted to the right and the lower to the left. To produce an illusion of size, repeat the procedure by adapting to the right pair of gratings. These after-effects persist much longer if the adaptation period is prolonged (from Blakemore 1973).

as with luminance or contrast. Can the system adapt to associations or contingencies?

That may seem unlikely, but look at the well-known illusion shown in figure 3. This shows that exposure to a grating pattern for a minute or so desensitizes the system to that particular spatial frequency and orientation, thus shifting the position of maximum excitation when the eye views a pattern of slightly different frequency or orientation. Viewing the grating for a minute or two would change the autocorrelation function in the part of the visual field concerned, so we should expect such after-effects if it is continuously being estimated and used to discount established associations. A map of the autocorrelation function at different positions in the visual field in the recent past, which is what seems to provide the feedback for this type of adaptation, perhaps does constitute worthwhile 'knowledge'.

How universal are associative or contingent adaptive effects? One of the best known is the 'waterfall phenomenon' that follows prolonged inspection of a moving region of the visual field, and one of the most dramatic is the effect described by McCollough (1965). To generate this one alternates one's gaze, preferably for a period of several minutes, between a vertical grating of one colour and a horizontal grating of a complementary colour; after this one looks at plain black-and-white vertical or horizontal gratings, and observes that each is tinged with a colour complementary to the one it had received during the adaptation period. The contingency adapted to here is more complicated than that implied by a simple change in power spectrum or autocorrelation function, and since McCollough's pioneering work in 1965 adaptation to many other contingencies has been tested. Not all of these work equally well, but Harris (1980) summarized the situation by saying that, in every case reported, some degree of contingent adaptation had been found.

Originally it was thought that adaptation that was selective for a particular contingency implied the existence of cells selectively sensitive to that particular contingency. This was the motivation for McCollough's original experiment, which was done soon after Hubel and Wiesel had discovered that there are neurones in the primary visual cortex of primates that are selective for both colour and orientation (Hubel & Wiesel 1968). The demonstration of contingent adaptation to colour and orientation seemed a triumphant vindication of this idea, and a period of enthusiasm for the non-invasive 'psycho-anatomical' exploration of the properties of neurones in the visual system followed, but the universality of contingent adaptation casts grave doubt on the whole enterprise. As Harris (1980) pointed out, it is reasonable to suppose that there are units for some contingencies, but surely not for all contingencies!

4. CONTINGENT ADAPTATION IN SINGLE NEURONS

The output activity of a neurone depends upon a single variable representing the combined effects of all its synaptic inputs, but adaptation could be occurring at individual synapses and could therefore be much finer-grained: the range of adaptive states of a single cell may be multidimensional and enormously wider than the range of excitatory states. Could the adaptive state of an individual neurone be different following adaptation to different combinations of its inputs, combinations that do not necessarily produce different outputs?

Contingent adaptation is a special case where a component of the adaptive state produced by joint exposure to two stimuli A and B specifically affects responses to that conjunction and differs from the adaptive state produced by exposure to either of them alone. We know that it occurs in the brain as a whole from the McCollough and numerous other effects, but we do not know if it occurs in single cells. The paper in this volume by Carandini *et al.* (1997) tests experimentally for the occurrence of contingent adaptation in single neurones of the primary visual cortex of macaques, and provides evidence that it does occur in some circumstances. The purpose of the next few paragraphs is to give a feeling first for the range of phenomena in which it may be involved and then for the functional significance of such a mechanism.

5. THE LAW OF REPULSION

Let us start with a 'law of repulsion', which summarizes a lot of empirical psychological observations (Barlow 1990). It states that when two stimuli frequently occur together their representations in the brain repel each other, meaning that the representations discourage or inhibit each other so that each is weaker than it is when each stimulus is presented alone. It is supposed that the degree of repulsion is constantly modifiable and therefore represents the average strength of the association between the two stimuli over some period in the past. Repulsion frequently occurs between closely related pairs of stimuli, for example those in the same modality of sensation and in similar regions of the sensory field, but of course it cannot occur between all possible pairs of stimuli.

Lateral inhibition in sensory pathways may be a simple example. There is certainly a strong tendency for neighbouring points on a sensory surface to be excited together, and the very frequent occurrence of mutual inhibition is an example of repulsion. In spite of the short-term variability implied by figure 2, this mechanism certainly has a strong genetic component, so the 'period in the past' referred to above must, in this case, involve an evolutionary time scale.

The 'waterfall effect' (the well-known after-effect of adapting to motion) can be explained when one realizes that even a small moving stimulus excites many units. During adaptation these will respond together, so mutual inhibition between them will develop and their responses will decline, thus accounting for the decrease in the perception of movement during the adaptation period. When the movement ceases the persistent increased strength of the mutual inhibition between the neurones tuned to the adapted direction of movement will make them less sensitive than those tuned to the opposite direction, unbalancing the system and causing the perception of movement in the reverse direction. A similar explanation can account for after-effects such as those shown in figure 3, for the McCollough effect, and for the numerous other forms of contingent adaptation mentioned above.

6. ILLUSIONS PRODUCED BY THE ABSENCE OF HABITUAL ASSOCIATIONS

We experience a world with a very complex structure in which there are numerous associations between the sensory stimuli we receive, and the strengths of the repulsions between these stimuli should, according to the law, reflect the usual strengths of these associations. Many well-known illusions occur when such habitual associations are prevented from occurring, and the law of repulsion neatly explains these. For instance motion parallax and stereo-parallax provide evidence for the distance of objects that is usually strongly correlated, but when viewing a stereo-pair the cue from motion parallax is absent. The law tells us that strong repulsion must have developed between the representations of signals from stereo-disparity and those from motion parallax, so if, while looking at a stereo-pair, one moves one's head in a way that would induce motion parallax in a real three-dimensional scene, one expects one's perception of the scene to be perturbed by the absence of the habitual motion parallax. This happens: one perceives the objects in the scene actively moving by amounts that would precisely remove the motion parallax signals. Those who have never seen this illusion should seize the next opportunity to experience it.

Judgements of size and distance provide many examples of illusions that might result from repulsion induced by past experience of associations, and some of these include: (i) impression of distance from haze and blue tinge; (ii) micropsia induced with a prism or concave lens; (iii) the 'toy-town' illusion from an excess range of disparities; (iv) illusions of depth and slant from vertical disparities; (v) false rotation of a 'concave' face; (vi) many effects described by Kersten (1997). One should always consider contingent adaptation before accepting innate mechanisms or top-down influences. I certainly do not claim that contingent adaptation is their proven cause, and even if it is there is much that remains obscure, such as the locus of the changes associated with repulsion and the duration over which repulsion is established. But it is a possible cause of a wide range of phenomena and needs to be considered in each case. Let us examine the role it might play in the paradoxical motions of a concave face-mask that Gregory (1970, 1997) described, for this can be taken as a canonical example of a so-called top-down influence.

As you walk past a normal stationary head, the tip of the nose, the eyes and the ears have successively lower angular velocities because of their successively greater distances from you. This is the normal pattern of motion parallax for faces, and since it has been experienced countless times it should have established the appropriate pattern of 'repulsions' between the representation of your own voluntary motion and the motions of the components of a normal face. With a normal face, the patterns you receive as you move past are those for a face moving its direction of gaze in the direction opposite to the direction you are moving, but you do not see this: you see a fixed head, not such a rotating head. Presumably the repulsions, in conjunction with your voluntary movement, counteract the expected signals for contrary rotation of the head and prevent you making that interpretation.

When you walk past a concave version of a face instead of a normal one the pattern of motion parallax experienced is that appropriate for a real (convex) face rotating faster in the same direction as your motion. The repulsions now work in the wrong direction to compensate for the expected contrary rotation and you have a strong experience of the head following you and even turning faster than you are moving. In passing it should be noted that this explanation would lead one to expect that a face painted on a flat surface should appear to follow you as you walk past it. I think this effect can be seen, but it is equivocal and certainly not nearly as vivid as the experience with the concave face; perhaps it would work better for individuals who have never previously seen a portrait and have not adapted to that experience.

There must be innate features involved in this illusion, such, for instance, as the connections that set up 'face cells' in the infero-temporal cortex, and the mechanisms that allow signals for your own motion to be correlated with the motion parallax signals. But simply labelling it 'top-down' is too simple and draws attention away from both the innate factors, and the key role that may be played by the past experience of correlated input signals.

7. FUNCTIONAL ROLE OF CONTINGENT ADAPTATION

One might express the functional role of the repulsions in the above example in the form of the following matrix equation:

subjective experience = input image + input image × [pattern of repulsions induced by past experience of faces and the current voluntary movement]

Since the repulsions are negative, the second term on the right-hand side subtracts the expected distortions resulting from the voluntary motion in a manner reminiscent of Donald MacKay's 'matching response' envisaged 40 years ago (MacKay 1955). The same basic idea has been given greater precision in many forms recently and may come closest to realization in the work described by Hinton & Ghahramani (1997) in this volume. But the functional role of contingent adaptation can also be understood in simpler examples.

The luminance gain control mechanism illustrated in figure 1 shifted the response characteristic of the ganglion cells to suit the range of luminances in the incoming image. Figure 4 illustrates how repulsion between two habitually correlated input signals can improve their representation. The top left panel illustrates how the responses (black dots) of two uncorrelated variables are distributed over the whole response plane, for dots occur in all the large squares, which are supposed to represent the very coarsely resolved gradations of the input signals. If they are strongly associated, as shown in the top right panel, the responses are clustered along the diagonal; many of the squares are unoccupied and the representation is obviously less effective. Adaptation of the type shown in figure 1 cannot improve the situation, for all it can do is to redistribute the responses along the horizontal and vertical axes separately, and it can take no account of their joint distribution. However, if the axes themselves are skewed, as shown at lower left, the resolvable states of the system also become clustered along the diagonal and the responses now occupy all the resolvable areas. Such a transformation of coordinates corresponds to forming new variables by subtracting a portion of each of the two original input variables from the other one, the fraction subtracted corresponding to the repulsion coefficient. When the responses of these new variables are plotted orthogonally they are distributed among all the resolvable spaces, as shown in the bottom right diagram.

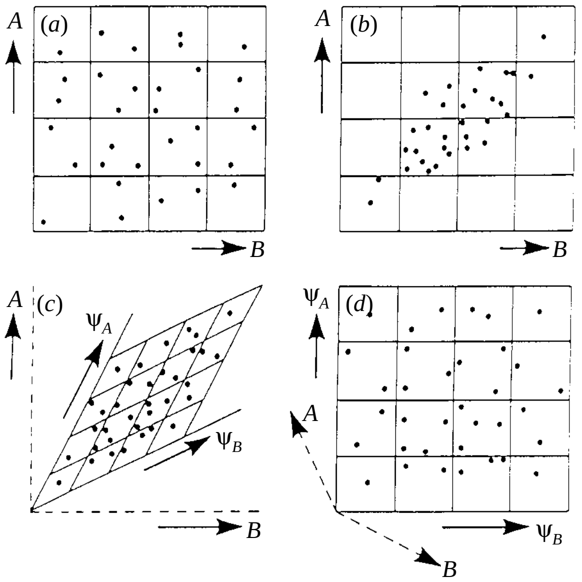

Figure 4. Repulsion (see text) creates a skewed coordinate system that improves the representation of associated stimuli. Two stimuli A and B that occur independently of each other are distributed over the whole A–B plane (*a*), but if A and B are positively associated the responses are clustered along the diagonal (*b*). If repulsion between A and B develops, then new variables ψ_A and ψ_B are created with $\psi_A = A - kB$ and $\psi_B = B - kA$. When the original coordinates A and B are used (*c*) the distinguishable states of ψ_A and ψ_B are now clustered along the diagonal, matching the distribution of the values of A and B, and if the new variables ψ_A and ψ_B are used (*d*) instead of A and B the points are distributed over the whole plane again. Note the relation between 'positive repulsion' in this representation and 'negative shear' in the paper by Anderson (1997).

The role of decorrelation or repulsion can be looked at qualitatively in yet another way. Anybody familiar with the analysis of variance will appreciate how much the sensitivity of a statistical test is improved by removing variance caused by known associations. Thus if the effect of age on the height of school children is removed, the residual variance becomes much less than the overall variance and the effect of another variable, such as the sex of the children, their home location, or their parental income, can be detected much more easily.

The visual messages aroused by the world around us are full of very complicated associative structure, and the more that can be removed by a mechanism such as repulsion, the more sensitive will our perception be to other sources of variation in the environment. It goes without saying that the ability to detect such other factors is likely to be a powerful promoter of survival.

8. LIMITATIONS

Notice that the bottom-up influence of associative structure goes far beyond what was suggested in the initial example of an automatic exposure system in photography, for instead of simple signals for mean luminance and contrast, the associative structure of the input is used. Földiák (1990) devised a simple network for bringing about pairwise decorrelation, and the need for it was argued by Barlow & Földiák (1989), but the brain may well use higher-order associations as well as the pairwise ones implied so far in this discussion.

One must also take account of the fact that the visual system is not a straightforward hierarchy but consists of a network of areas with strong interconnections in all directions. This organization by itself makes one suspect that it is simplistic to think only in terms of top-down and bottom-up influences; the transformations in each area are open to modification according to the activity in other areas, and may be temporally or permanently modified by what is going on in these other areas. In such a complex system it is particularly important to realize that the structure of the messages being handled provides a source of knowledge.

But deriving and using knowledge from the redundancy of the input has limitations. First, the associative structure of sensory messages is enormously complicated, and to work it all out from first principles would be a horrendous task. Short cuts and hints provided by genetic prewiring must be a great help to a biological brain, and for the same reason it is almost certainly advantageous to impose as much high-level knowledge as possible on to an engineered visual system. At the start of this article three ways in which knowledge could be derived and introduced into a visual system were mentioned, but although there is little doubt that there is a role for all of them the proportions in which they should be used is still not clear.

One also needs to be reminded that the ideas of this article are not new, but have been around for at least 40 years (Attneave 1954; Barlow 1959; Watanabe 1960) and the principle goes back to Locke, Helmholtz and Mach. In this volume these ideas are particularly important for the papers by Carandini *et al.* (1997) and Hinton & Ghahramani (1997), and some of the nicest examples of perceptual effects that may result from redundancy analysing mechanisms are to be found in the paper by Kersten (1997). The prominent use of the Karhunen–Loéve transform (KLT) or principal components analysis (PCA) in the papers by Edelman (1997), Bobick (1997), Taylor (1997) and Pentland (1997) show that the benefits of analysing covariance are recognized in these areas of computer vision. But even though Attneave introduced the idea of redundancy into psychology as early as 1954, the idea does not seem to have taken hold there.

The problem of knowledge-based vision never looks the same once it has been recognized that much that is simplistically attributed to top-down influences is either provided by the genetic structure of the visual system, or obtained by analysis of the redundancy present in the inputs themselves. To avoid the top-down/bottom-up dichotomy blinding us to more important questions, the term 'top-down' should perhaps be challenged whenever it is used, not least because the visual system has no 'top'.

REFERENCES

Anderson, J. A. D. W. 1997 Representing geometrical knowledge. *Phil. Trans. R. Soc. Lond.* B **352**, 1131–1141. (This volume.)

Attneave, F. 1954 Informational aspects of visual perception. *Psychol. Rev.* **61**, 183–93.

Barlow, H. B. 1959 Sensory mechanisms, the reduction of redundancy, and intelligence. In *The mechanisation of thought processes*, pp. 535–539. London: Her Majesty's Stationery Office.

Barlow, H. B. 1969 Pattern recognition and the responses of sensory neuronees. *Ann. NY Acad. Sci.* **156**, 872–881.

Barlow, H. B. 1990 A theory about the functional role and synaptic mechanism of visual after-effects. In *Vision: coding and efficiency* (ed. C. B. Blakemore). Cambridge University Press.

Barlow, H. B. & Földiák, P. 1989 Adaptation and decorrelation in the cortex. In *The computing neurone* (ed. R. Durbin, C. Miall & G. Mitchison), pp. 54–72. Wokingham: Addison-Wesley.

Barlow, H. B., Fitzhugh, R. & Kuffler, S. W. 1957 Change in organisation in the receptive fields of the cat's retina during dark adaptation. *J. Physiol.* **137**, 338–354.

Blakemore, C. 1973 The baffled brain. In *Illusion in nature and art* (ed. R. L. Gregory & E. H. Gombrich), pp. 847. London: Duckworth.

Bobick, A. F. 1997 Action, activity, and movement: the role of knowledge in the perception of motion. *Phil. Trans. R. Soc. Lond.* B **352**, 1259–1267. (This volume.)

Carandini, M. & Ferster, D. 1997 A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* **276**, 949.

Carandini, M., Barlow, H. B., O'Keefe, L. P., Poirson, A. B. & Movshon, J. A. 1997 Adaptation to contingencies in macaque primary visual cortex. *Phil. Trans. R. Soc. Lond.* B **352**, 1151–1156. (This volume.)

Edelman, S. 1997 A model of visual recognition and categorization. *Phil. Trans. R. Soc. Lond.* B **352**, 1193–1204. (This volume.)

Field, D. J. 1987 Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am.* A **4**, 2379–2394.

Földiák, P. 1990 Forming sparse representations by local anti-Hebbian learning. *Biol. Cybern.* **64**, 165–170.

Geisler, W. S. & Albrecht, D. G. 1995 Bayesian analysis of identification performance in monkey visual cortex: non-linear mechanisms and stimulus certainty. *Vision Res.* **35**, 2723–2730.

Gregory, R. L. 1970 *The intelligent eye.* London: Wiedenfeld& Nicolson.

Gregory, R. L. 1997 Knowledge in perception and illusion. *Phil. Trans. R. Soc. Lond.* B **352**, 1123–1129. (This volume.)

Harris, C. S. 1980 Insight or out of sight? Two examples of perceptual plasticity in the human adult. In *Visual coding and adaptability* (ed. C. S. Harris), pp. 95–149. New Jersey: Laurence Erlbaum Associates.

Hinton, G. E. & Ghahramani, Z. 1997 Generative models for discovering sparse distributed representations. *Phil. Trans. R. Soc. Lond.* B **352**, 1179–1192. (This volume.)

Hubel, D. H. & Wiesel, T. N. 1968 Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* **195**, 215–243.

Kersten, D. 1997 Perceptual categories for spatial layout. *Phil. Trans. R. Soc. Lond.* B **352**, 1157–1165. (This volume.)

MacKay, D. M. 1955 The epistemological problem for automata. In *Automata studies* (ed. C. E. Shannon & J. McCarthy), pp. 235–251. Princeton University Press.

McCollough, C. 1965 Color adaptation of edge-detectors in the human visual system. *Science* 1**49**, 1115–1116.

Movshon, J. A. & Van Sluyters, R. C. 1981 Visual neural development. *A. Rev. Psychol*. **32**, 477–522.

Ohzawa, I., Sclar, G. & Freeman, R. D. 1985 Contrast gain control in the cat's visual system. *J. Neurophysiol*. **54**, 651–667.

Pentland, A. 1997 Content-based indexing of images and video. *Phil. Trans. R. Soc. Lond.* B **352**, 1285–1292. (This volume.)

Ruderman, D. 1994 Statistics of natural images. *Network* **5**, 517–548.

Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W. & Meister, M. 1997 Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**, 69–73.

Taylor, C. J., Cootes, T. F., Lanitis, A., Edwards, G., Smyth, P. & Kotcheff, A. C. W. 1997 Model-based interpretation of complex and variable images. *Phil. Trans. R. Soc. Lond.* B **352**, 1269–1276. (This volume.)

Tolhurst, D. J., Tadmoor, Y. & Chao, T. 1992 Amplitude spectra of natural images. *Ophthalm. Physiol. Opt.* **12**, 229–232.

Van Nes, F. L. & Bouman, M. A. 1967 Spatial and temporal modulation transfer in the human eye. *J. Opt. Soc. Am.* **57**, 410–406.

Watanabe, S. 1960 Information-theoretical aspects of inductive and deductive inference. *IBM J. Res. Dev.* **4**, 208–231.