

J. D. Mollon (1977), Chapter 4 of "The Perceptual World," edited by K. von Fieandt and I. K. Moustgaard, London, Academic Press


4 | Neural Analysis

The Feature Detector


The last two decades have seen the doctrine of specific nerve energies elaborated in the concept of the "feature detector". The inquisitive micro-electrode has shown that vertebrate sensory systems contain neurons apparently designed to detect those features of the physical environment that either are of immediate biological significance to the animal, or else are especially informative for the identification of objects. Above all, we find neurons that are alert for a particular kind of change, the temporal change that may reveal prey or predator, the spatial change that may mark the contour of an object.¹

The detection of movement, for example, is of universal and primitive importance and it has long been suspected from phenomenological evidence that in man's visual system there survive specialized mechanisms that detect movement. "It is experimentally certain", wrote William James, "that the feeling of motion is given to us as a direct and simple sensation." In the periphery of our visual field we can detect slight movements even when we cannot spatially resolve the starting and ending points of the movement. In unusual circumstances, movement may be phenomenally dissociated from change of position or of shape (Bowditch and Hall, 1880;

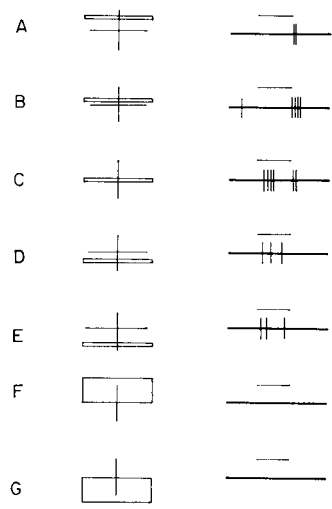
1. For further discussions of the matters covered in this chapter the reader may turn to Barlow (1972), Blakemore (1974), Dodwell (1970), Hubel (1963), Konorski (1967), Lettvin, Maturana, McCulloch and Pitts (1959), Virsu (1973) and Exner (1894). The neuro-physiology of auditory pattern recognition is reviewed by Evans (1974b). Much useful material is to be found in the successive volumes of the *Handbook of Sensory Physiology*, published by Springer-Verlag, and of the *Neurosciences Study Program*, published by the Rockefeller University Press (Vols I and II) and the MIT Press (Vol. III).



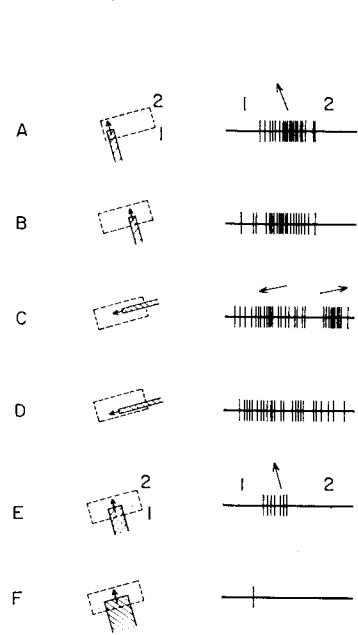
a. Simple



c. Lower-order hypercomplex



b. Complex



d. Higher-order hypercomplex

Wertheimer, 1912; Graybiel *et al.* 1946; see also chapter 16 p. 416); and movement, like other sensory qualities such as brightness and colour, can give rise to simultaneous and successive contrast. As early as 1894, considerations of this kind led the prescient psychophysicologist Exner to suggest that our visual system contains single cells sensitive to the direction of movement (see Fig. 4.5). Neurons of this kind, Exner argued, must occur at a stage before the signals from the two eyes are combined, for, if corresponding parts of the two retinæ are simultaneously stimulated by movement in opposite directions, the direction of the illusory after-effect depends on which eye is tested. In 1953 Barlow published the first report of a class of ganglion cells in the retina of the frog that were particularly sensitive to a small object moving within their receptive field.² Such cells, he noted,

2. See also chapter 16 p. 415. The receptive-field of a cell is that area of the receptive surface, of the retina, say, or the skin, within which stimulation can affect the response of the cell concerned.

Fig. 4.1 Hubel and Wiesel's classification of the types of neuron found in the visual cortex of the cat.

In each case typical stimuli are shown to the left and the neural responses they characteristically evoke are shown to the right. In (a) and (b) the short lines above the neural records represent the period for which the stimulus was present and in (c) and (d) the arrows represent the direction of stimulus movement.

(a) *Simple cell*. The maximum response was elicited by a vertical slit of light that passed through the centre of the receptive field (E). The responses was reduced if the orientation or the position of the slit was changed. (From Hubel and Wiesel, 1959.)

(b) *Complex cell*. A horizontal slit evoked a response anywhere within a receptive field that had a diameter of three degrees of visual angle. When the slit was shone anywhere within the upper half of the receptive field the response was at the offset of the stimulus (A, B); in the lower part of the field the response was to the presence of a stimulus (D, E). When the slit was in an intermediate position (C), the cell responded at both onset and offset. However, and this is critical to the definition of a complex cell, a summation does not occur within the two parts of the field: large rectangles covering the entire lower or upper halves of the field (F, G) were completely ineffective. (From Hubel and Wiesel, 1962.)

(c) *Lower-order hypercomplex cell*. To evoke a strong response the stimulus had to be of a particular orientation and had not to extend beyond the boundaries shown by broken lines (C). The optimum length was about three degrees. In the case of this cell the left-hand antagonistic region seemed to be more powerful than the right hand one (compare G and H); and some lower-order hypercomplex units are stopped at only one end. (From Hubel and Wiesel, 1965.)

(d) *Higher-order hypercomplex cell*. The cell responded to a narrow tongue introduced into the receptive field in either of two orientations 90 degrees apart (compare A and C). The exact position of the tongue did not matter (compare A and B; and C and D) but if its width was increased the response was reduced (E, F). In A, B, E, F the duration of the samples was 5 s; in C, D it was 10 seconds. (From Hubel and Wiesel, 1965.)

could serve the frog well as “fly detectors”. Cells with these properties, however, have not proved exclusive to the insectivorous frog and have since been found in the retinae of many different vertebrates, including primates. Cells that are sensitive to the actual direction of movement have been discovered in the retinae of rabbits and pigeons; and, despite an earlier belief that the analysis of movement was *encephalized*³ in higher mammals, such cells have now been reported in the retinae of cats (Hoffmann, 1973). A typical unit in the rabbit’s retina responds with a burst of firing to movement in the preferred direction whether the stimulus is a white spot on a dark ground or a black spot on a light ground; it may be inhibited by movement in the opposite, or *null*, direction (Barlow *et al.* 1964).

Cells that are sensitive to the direction in which a delicate tactile stimulus is drawn across a particular area of skin have been found in the cortex (somatosensory Area 1) of the monkey (Whitsel *et al.* 1972). Similarly, cells specific to the direction of movement of a sound source have been reported in the auditory system of the cat (Al’tman and Markovich, 1969); but the truest auditory analogue to movement across the retina is perhaps movement along the basilar membrane (see chapter 7), that is, a change in the pitch of the stimulus. Whitfield and Evans (1965) have, in fact, described cells in the auditory cortex of the cat that respond to only one direction of change of acoustic frequency.

The catalogue of the feature detectors found in the major sensory systems is now a very long one. Figure 4.1 shows the classical types of unit described by Hubel and Wiesel in the visual cortex of the cat and the monkey (see also chapter 5). The *simple cell* has a bar-shaped receptive field that can be divided into discrete excitatory and inhibitory regions. The response to a stationary edge or bar can be predicted from a plot of the field made with flashing spots. Thus the simple cell responds optimally to a bar or edge in a particular orientation and a particular position. The *complex cell* typically has a larger receptive field and responds throughout its field to a bar or edge of a particular orientation; its receptive field, unlike that of a simple cell, cannot be divided into distinct excitatory and inhibitory regions and sustained firing is elicited by the movement of a bar across the field. The *lower-order hypercomplex cell* requires that the bar or edge be *stopped* at one or both ends: if the stimulus is too long, if it extends beyond the excitatory field, the response is inhibited. The *higher-order hypercomplex cell* responds to stimuli in either of two perpendicular orientations and could conceivably measure curvature. Hubel and Wiesel’s classificatory scheme, introduced more than a decade ago, is still in general use, although it needs addition and modification and although, as we shall see, the model

3. Postponed until the cortex.

that it originally inspired is no longer tenable. The student should still read the classical papers (Hubel and Wiesel, 1962; 1965; 1968).

Often the specificity of a cell is to a conjunction or disjunction of stimulus qualities. Thus a particular unit may require that a bar is in a particular region of the visual field, is in a particular orientation, is of a particular colour and is moving in a particular direction (cf. Gouras, 1974, Fig. 10).



The Argument from After-effects

There is every reason to suppose that man's sensory systems are organized in a similar way to those of animals, but here we must rely primarily on psychological evidence. One of several indirect approaches has drawn upon *perceptual after-effects*. If a particular class of feature-detector is present in one of our sensory systems, then, it is argued, we may be able selectively to fatigue these feature-detectors and cause a change either in the detectability of the feature concerned or in its phenomenal appearance.⁴



By adapting to the grating of Fig. 4.2a and then looking at 4.2b the reader can experience one of the most celebrated of all after-effects, the *tilt after-effect*. Long before Hubel and Wiesel had made micro-electrode recordings from the visual cortex and thus before the existence of orientationally selective neurons was known, J. J. Gibson argued from the tilt after-effect that the orientation of a line should be regarded as a simple sensory quality, comparable, in fact, to brightness or to colour (Gibson, 1937).

A relatively recent addition to the catalogue of features extracted by the visual system is *spatial frequency*. The striking after-effect that can be obtained from Fig. 4.2c was first reported in 1969 by Blakemore and Sutton. Some years earlier visual scientists had begun to use stimuli that consisted of gratings in which luminance varied sinusoidally along one axis. Such a stimulus is a visual analogue of a pure tone (see chapter 7 and Fig. 7.1) and its frequency is usually expressed as the number of cycles per degree of visual angle. It initially recommended itself because it was the mathematically simplest of spatial stimuli, but there soon came the suspicion that the visual system contains channels tuned to particular spatial frequencies (Campbell and Robson, 1968) and single neurons with such selectivity have been discovered in the visual systems of cats and primates (see chapter 3 Fig. 3.8b). But what would be their role in perceptual analysis? Some suppose that these analysers of spatial frequency are directly responsible for pattern recognition, that, in fact, the visual system performs a spatial



4. An introductory account of perceptual after-effects is given by Mollon (1974). A comprehensive summary of visual after-effects will be found in Anstis (1975).



(a)



(b)



(c)

at: studies during the pri-
 ds engaged in an Employ-
 l his intense application
 ible of any void space in
 asion for domestick soci-
 He left behind him about
 Leibnitz, his rival, likewi-
 ces, tho' not so rich: But
 money which he had hoard
 inary examples, and both



discovery of them, it is a light entirely new and
 pecked.

What Sir Isaac Newton aims at quite through-
ticks, is the Anatomy of Light; this expression is n
 bold since it is no more than the thing it self.
 experiments, the smallest ray of Light that is co
 into a dark room, and which cannot be so smal
 that it is yet compounded of an infinite number o
 rays, is divided and dissected in such manner, th
 Elementary rays of which it is compoied, are sep
 from each other, and discover themselves ever
 ringed with its particular colour, which after this
 reason can no more be altered. The first total ra
 fore the dissection, is white, and this whiteness
 from all the particular colours of the Primitive
 The separating these rays is so difficult, that wher



(d)

delt, began by taking hold of the
 to climb to unknown principles;
 them only in such manner as they cou
 chain of consequences. The former fe
 clearly understands, to find out the
 sees; the latter sets out from what h
 find out the cause, whether it be c
 self-evident principles of the one
 him to the causes of the phenome
 the phenomena do not always lea
 ciples sufficiently evident. The bou
 two such men in their pursuits thro
 were not the boundaries of Their



delt, began by taking hold of the
 to climb to unknown principles;
 them only in such manner as they cou
 chain of consequences. The former fe
 clearly understands, to find out the
 sees; the latter sets out from what h
 find out the cause, whether it be c
 self-evident principles of the one
 him to the causes of the phenome
 the phenomena do not always lea
 ciples sufficiently evident. The bou
 two such men in their pursuits thro
 were not the boundaries of Their

(e)



(f)

*Fourier analysis*⁵ and that what is recognized is the resulting Fourier spectrum. However, the spatial frequency channels could have other functions. Firstly, they may serve to identify the density of texture of a surface and texture density may be on a par with colour: it allows the separation of figure from ground and of one object from another, a separation that is a preliminary to the recognition of form. In addition, *gradients* of texture density (see chapter 11 p. 310) are important in our judgments of distance (Gibson, 1950). Another function for the spatial-frequency channels might be seen in our ability to attend selectively to global or to local aspects of a visual stimulus: if we are looking at a page of text we can choose to attend to a single word or to a single letter of the word or to a serif on a particular letter.⁶

Classes of After-effect and Types of Code

The spatial-frequency after-effect belongs to a class of after-effects in which the test stimulus is phenomenally displaced away from the adapting stimulus along some perceptual dimension. Figure 4.3, which is based on a hypothesis introduced by von Békésy in 1929,⁷ shows how such after-effects can be explained if the perceptual dimension is coded by place (see chapter 3 p. 55), if, for example, different spatial frequencies are represented by the activity of different members of a set of neurons. We must

5. Fourier analysis is the analysis of a complex waveform into a set of pure sine and cosine waves from which it could be reconstructed by simple addition. See chapter 7 p. 211 and Fig. 7.3.

6. The analysis of spatial frequency by the visual system is reviewed by Campbell (1974), by Campbell and Maffei (1974) and by Robson (1975).

7. An explanation of this kind was first extended to the tilt after-effect by Sutherland (1961) in a paper that has not received the acknowledgement it deserves.

Fig. 4.2 To obtain the tilt after-effect, gaze for about 30 seconds at the central horizontal bar of (a). Keep your eye moving to and fro along the length of the bar to avoid producing an ordinary after-image. Then switch your gaze rapidly to the central spot of (b). For a moment the vertical bars of (b) will appear tilted in directions opposite to those of the adapting grating in (a).

To obtain the spatial-frequency after-effect, move away from the page until the finest grating is still just clearly resolvable and look for about a minute at the central bar of Figure (c), again moving your eye to and fro along the bar. Then look quickly at the central spot of (b). For a moment the lower grating of (b) will look denser while the upper grating will appear coarser.

If you have difficulty securing either effect, try adapting for longer periods or try viewing from a different distance.

We do not have to use gratings. The two further adaptation figures, (d) and (f), and the test figure (e), allow the reader to judge the generality of the spatial frequency after-effect and its transfer from one kind of stimulus to another. (After Campbell and Maffei; Blakemore; Mayhew; Anstis.)

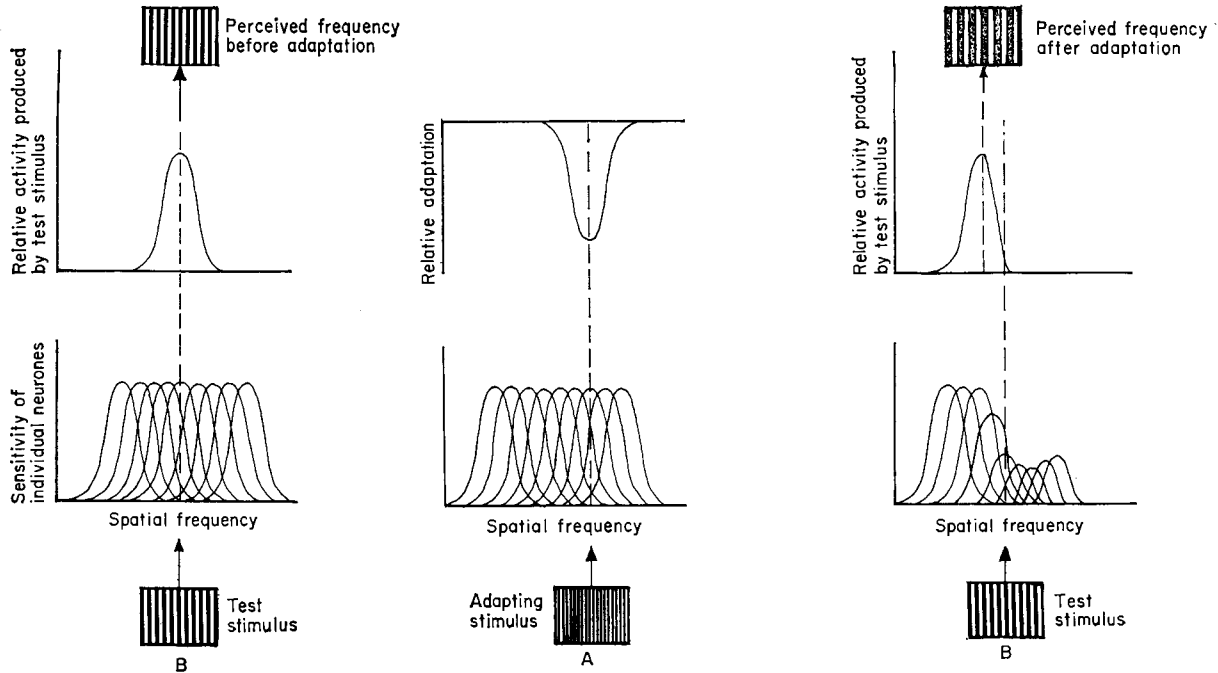


Fig. 4.3 A qualitative model of how after-effects occur when coding is by place. (After Békésy; Osgood; Sutherland; Blakemore.)

first make the assumption that neurons representing different values on the dimension have overlapping sensitivities and that decisions by later mechanisms are based on the position of peak activity in this population of neurons. Prolonged exposure of a particular value on the dimension (A in the figure, for example) will depress the corresponding neurons to an extent that is inversely related to their response to the adapting stimulus. If we now present a test stimulus (B) that lies to one side of A, the activity it produces must be weighted by the adaptation left by A. So the peak of activity produced by B, and thus perhaps its phenomenal appearance, will be displaced from A. If, however, the test stimulus is coincident with A, the peak activity will be reduced but will not be displaced (provided there are no asymmetries in the tuning curves of individual neurons in the series). Thus stimulus dimensions coded by place will show the *distance paradox*: the maximum after-effect will be for stimuli some distance along the dimension from the adapting stimulus. Appropriate measurements are available for the spatial-frequency after-effect and indeed distortion is greatest for test gratings lying one half to one octave⁸ either side of the adapting frequency (Blakemore and Sutton, 1969).

The movement after-effect, which is one of the class of phenomena traditionally called *negative after-effects* (chapter 3 p. 66), seems to differ in a critical way from the Blakemore-Sutton effect that we have just been discussing. The relevant feature of the stimulus appears to change during the adaptation period: the movement comes to appear slower (Gibson, 1937) whereas spatial frequency does not appear to change during adaptation.⁹ This observation may suggest that speed is not coded by place, but rather that all movement in a particular direction excites the same neuron, speed being represented by the frequency of impulses or the magnitude of a graded potential. Figure 4.4, which is based on a hypothesis introduced by Exner, shows schematically how the movement after-effect may arise. Two neurons are shown (L and R), one sensitive to movement to the left, the other sensitive to movement to the right. We have to assume that when their outputs are equal the later stages of visual analysis take the stimulus to be stationary and that we have here an example of an opponent process, a limiting case of coding by place such as discussed in chapter 3 (see p. 67). During prolonged stimulation by stimuli moving to the left, the response of neuron L slowly wanes. When stimulation ends the neuron will be depressed for several seconds and an inequality in the spontaneous activity of L and R may be treated as movement to the right. Alternatively, L and R may be

8. To say that a frequency is increased or decreased by an octave is to say that it is doubled or halved.

9. Since the adapting stimulus is here being regarded as a test stimulus the question is close to that of whether the distance paradox holds.

linked by mutual inhibition; and adaptation of L may transiently release R from inhibition.

Critical to this account of the movement after-effect is the question of what happens when we adapt to relatively slow movement and then look not at a stationary test stimulus but at one that is moving faster than the adapting stimulus. If speed is coded by frequency then the test stimulus should appear slower than normal; if speed is coded by place, the perceived speed should be displaced away from that of the adapting stimulus and the test stimulus should appear faster than it normally would. This crucial experiment has been neglected, but Carlson (1962a) did find one condition in which a test stimulus appeared to move faster after adaptation to slower movement.¹⁰




Fig. 4.4 Exner's explanation of the movement after-effect. Before adaptation both neurons are spontaneously active; after adaptation only R is active. Graded potentials could be substituted for action potentials.

A more clear-cut candidate for coding by frequency might be seen in apparent contrast. The apparent contrast of a grating, unlike apparent spatial frequency, does change during adaptation; and the maximum reduction in apparent contrast is at the adapting frequency (Blakemore and

10. There are other difficulties that face too simple an account of the movement after-effect. If the retina receives no patterned stimulation at the end of adaptation, the after-effect can be stored for several seconds at almost full strength (Wohlgemuth, 1911); adaptation for 15 minutes to a rotating spiral will produce a weak after-effect the next day (Maslin, 1969); and the after-effect of lateral movement can be abolished by the absence of a stationary framework (Day and Strelow, 1971). The present writer confirms these phenomena from his own observations.

Campbell, 1969). There is no distance paradox. We can draw conclusions only about the level of analysis at which adaptation occurs; but at this level, contrast would not seem to be coded by place.

Hierarchies

How do individual neurons gain their specificity to a complex feature of the stimulus? Although we have a good idea of the trigger features for neurons at different levels of different sensory systems, we know very little of how neurons are connected together to achieve these specificities. It is supposed,




Fig. 4.5 Exner's model of directionally-selective movement detectors (Exner, 1894, p. 193). *a-f*, etc. are the projections of discrete retinal points. *S*, *E*, *Jt* and *Jf* are directionally-selective cells. Each retinal point is individually connected with each movement detector, but for simplicity the connections are shown in the diagram as passing through the projections of neighbouring retinal points. The time required for excitation from any particular retinal point to reach a particular movement detector is approximately proportional to the distance given in the diagram.

How does the model deal with dark objects moving on a bright ground? Exner supposes that considerable processing occurs between the retinal receptors and their central projections (*a-f*) and that by the stage shown in the diagram a dark object is represented by a positive signal.

The cells *a1-a4* project to the external eye muscles (only four are shown in the diagram) and to the cortex.

however, that sensory neurons are organized into successive levels and that the specificity of a particular neuron is achieved by selective inputs, excitatory and inhibitory, from the previous level. The successive levels of analysis may reveal themselves in anatomically distinct layers, as in the retina, but elsewhere cells at different functional levels may be intermingled and we can only infer their connections from physiological and anatomical experiments.




Fig. 4.6 Barlow and Levick's model for directional selectivity in the rabbit retina. Excitation passes from the light-sensitive receptors (*R*) via the bipolar cells (*B*) to the ganglion cell (*G*), but the bipolar cells can be inhibited by horizontal cells, which pick up from receptors in an adjacent region of the retina and conduct laterally in the null direction through a teledendron (*Td*) (but cf. Ch. 5 p. 106). Thus responses are inhibited when an image moves in the null direction but not when it moves in the preferred direction. A particular ganglion cell must receive inputs from a set of bipolars all selective for the same direction of movement; for although the receptive field may have a diameter of three degrees of visual angle, the directional selectivity may hold for movements of less than a quarter degree anywhere within a large region of the field. (From Barlow and Levick, 1965, *J. Physiol.*, 178, 477–504.)

The principle of selective summation is explicitly embodied in Exner's model of movement detection (Fig. 4.5). The cell for leftward movement ($\mathfrak{y}t$) receives input from horizontal rows in the retinal array and responds optimally only when it receives simultaneous input from several retinal elements. Exner called such a cell a "Summationszelle"; in modern jargon, it is an AND-gate. When a stimulus moves from *c* to *a* the excitation from later stimulated points travels over shorter paths and so the inputs are temporally coincident at $\mathfrak{y}t$. The excitation reaching *E* on the other hand, is

dispersed in time. We are to imagine that the set of movement-detectors, \mathcal{Y}_t , \mathcal{Y}_f , S and E , is reduplicated many times for different regions of the retina. When it did become possible to record from the movement-detectors that Exner had so remarkably foreseen, it turned out that his coincidence model was not enough and that inhibitory as well as excitatory processes were required. Figure 4.6 shows Barlow and Levick's model of directionally selective units in the rabbit retina (see: structure of retina, chapter 5 p. 101). When a stimulus moves in the *null*, or non-preferred, direction the delayed inhibition reaching the bipolar cell via the horizontal cells cancels the response to direct excitation, but when movement is in the preferred direction the inhibition arrives too late to prevent the excitatory response. The evidence for inhibition lies in the fact that a response to a spot moving in the null direction may be less than that to a stationary spot: a purely excitatory model could not account for this finding.

A model that is very similar in principle has been offered by Bishop *et al.* (1971) to account for the directional selectivity of simple cells in the cat's cortex; and the reader will readily see how an analogous scheme would explain the presence in the auditory system of neurons specific to a particular direction of change of frequency.¹¹

The most celebrated of all hierarchical models is that offered by Hubel and Wiesel for the early stages of visual shape recognition in cat and monkey. Although it increasingly looks as if their account of the visual cortex may be wrong, the student should have some grasp of their model, because of the pervasive influence it has enjoyed for over a decade; and because of its beautiful simplicity.

We begin at the level of the bipolar cell of the retina (cf. Figs. 3.2, 4.6; see also chapter 5 p. 104). A typical unit may respond when light falls on the centre of its receptive field and be inhibited by light falling on a surrounding annulus (Fig. 4.7). Such a cell would be called an *on-centre* unit. It receives an excitatory input from those receptors that lie in the centre of its field and an inhibitory input from those in the surround; the latter input probably reaches it via horizontal cells. Other cells have *off-centres* and *on-surrounds* (see also chapter 6 pp. 168-170). Many retinal ganglion cells reflect the properties of the bipolar cells from which they receive their primary input. Thus already in the retina we have cells that have little interest in steady, diffuse light (which stimulates both *on* and *off* regions of the receptive field), but which respond strongly to local contrast, to edges and to spots. "Die Netzhaut . . . schematisiert und karikiert" (Mach, 1900).

11. Fernald (1971), however, has explained such specificity by a model that is formally similar to Exner's: the synaptic inputs on the dendrites of the summation cell lie at varying distances from the soma and so different sequences of input lead to either temporal coincidence or temporal dispersion of excitation.

The cells at the next level of the hierarchy, at the level of the lateral geniculate nucleus (LGN), have receptive-field properties similar to those of retinal ganglion cells, but are even less sensitive to diffuse light (chapter 6 pp. 111–113). It is now confirmed, by simultaneously recording from retina and LGN, that most LGN units take their excitatory input from a very

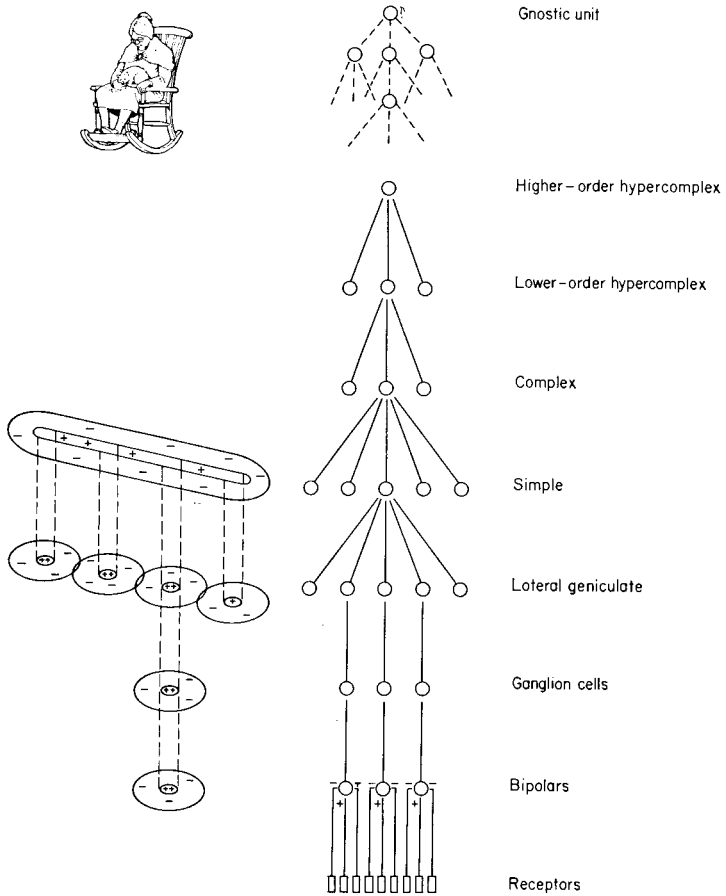


Fig. 4.7 The hierarchical model of the visual system. To the right are shown schematically the classes of neuron and their connections. At the higher levels these connections are strictly hypothetical.

To the left are represented typical receptive fields for cells at early stages of the system. The plus and minus signs mark areas in which a stimulus produces excitation and inhibition respectively. The diagram shows how the receptive field of a simple cell might correspond to the receptive fields of a set of lateral geniculate units.

The hypothetical “gnostic unit” is discussed later in the chapter.

small number of retinal ganglion cells of the same type (Cleland *et al.* 1971).

The next stage in Hubel and Wiesel's hierarchical model is the cortical simple cell. The orientational selectivity of the simple cell is achieved by taking a set of LGN units that have receptive fields lying in a row on the retina and feeding their outputs into the same simple cell (Fig. 4.7).

Simple cells in turn provide the input for complex cells. Each complex cell receives inputs from a set of simple cells that all have the same preferred orientation and so comes to have its characteristic property of responding to a bar of the correct orientation anywhere within the field.

Likewise, the complex cells furnish the input for lower-order hypercomplex cells. Thus a hypercomplex unit may come to be specific to length because it receives excitatory input from one complex cell and inhibitory input from other complex cells with fields lying at the end of its excitatory field. A higher-order hypercomplex cell may draw inputs from lower-order hypercomplex cells that have orthogonal preferred orientations (see chapter 12 pp. 337–340).

Difficulties for Hubel and Wiesel's model. X- Y- and W-cells

Unhappily this account of the sequential analysis of the visual cortex may no longer be tenable. It looks less and less likely that simple cells provide the exclusive input to complex cells. It seems certain that complex cells receive direct ("monosynaptic") input from the axons of LGN units: when the optic pathway is electrically stimulated and recordings are made from cortical units, complex cells respond with such a short latency and with so little variability that it is improbable that extra synapses are interposed (Stone, 1972; Toyama, *et al.* 1973; Singer, *et al.* 1975). Especially persuasive is the finding that some complex cells respond with a shorter latency than any simple cell in the same sample (Hoffman and Stone, 1971). Many complex cells do also receive inputs from other, as yet unidentified, cortical cells and it is an open question whether the direct or indirect inputs are most important in normal visual analysis. Indirect evidence against Hubel and Wiesel's hypothesis comes from a report that simple cells respond best to stimuli moving relatively slowly whereas complex cells often prefer higher velocities to which simple cells hardly respond (Movshon, 1974). Hammond and MacKay (1975) have recorded from the cat's visual cortex while stimulating the eye with the kind of "visual noise" that is obtained when a domestic television set is mistuned. If one area of such noise is moved relative to the rest, observers perceive contours at the boundaries of the moving area, even though a true contour is not present in the stimulus. Most simple cells proved to be indifferent to these "kinetic contours" but many complex cells responded vigorously to them. It is

unclear how this result would arise if complex cells took their main input from simple cells.

One possibility is that simple and complex cells are the respective projections of two different classes of retinal ganglion cell, *X-cells* and *Y-cells*. *X-* and *Y-cells* are primarily distinguished according to whether summation within their receptive field is linear: an *X-cell* does not respond when a sine-wave grating is so positioned that the changes in luminance over one half of the receptive field are the exact inverse of the changes over the other half and when, we may suppose, the inputs from the excitatory and inhibitory areas of the field directly cancel each other (Fig. 4.8); but null positions of this kind cannot be found for *Y-cells*, which respond to any variation of luminance within their field (Enroth-Cugell and Robson, 1966).




Fig. 4.8 The responses of an off-centre *X-cell* (*A*) and an off-centre *Y-cell* (*B*) to the introduction and withdrawal of a sinusoidal grating pattern. The ordinate of each graph shows the number of action potentials per second. The lowermost, rectangular trace shows when the sinusoidal pattern is turned off (downward deflexion) and when it is turned on (upward deflexion). To the right is shown the relation of the grating to the receptive field. (The "phase angle" is the angular position, in degrees, of the grating relative to the centre of the receptive field). Notice especially the difference in the responses of *X* and *Y* cells when the changes in luminance over half of the receptive field are the exact opposite of the changes over the other half, when, that is, the phase angle is 90 degrees or 270 degrees. (From Enroth-Cugell and Robson, 1966.)

There prove to be other differences between the two classes. Y-cells have large cell bodies and large axons and have high conduction velocities. Although their receptive fields can be divided into concentric ON and OFF areas, the centres are relatively large and the surrounds are notably so. X-cells have smaller somata, smaller axons, lower conduction speeds and smaller receptive fields. Y-cells respond only transiently to local illumination and prefer briskly moving stimuli; X-cells give a sustained response to local illumination. The axons of X-cells appear to project primarily to the lateral geniculate nucleus whereas those of Y-cells project both to the LGN and to the superior colliculus in the mid-brain. (Hoffman, 1973; Wright and Ikeda, 1974; Robson, 1975).¹² Are the Y-cells the sentinels that detect changes and events while leaving to the X-system the detailed analysis of form? The reader who knows the literature of the last hundred years will hear many an echo in such a speculation; and will hold his judgment.

LGN units can also be divided into X and Y classes and it has been supposed that the simple cells of the cortex receive excitatory input from X-cells while complex cells receive a parallel excitatory input from Y-cells (Stone, 1972; Stone and Dreher, 1973). Such a model accounts for many of the characteristic properties of simple and complex cells. A major difficulty for the hypothesis is the report of Singer *et al.* (1975) that both X and Y fibres project to both simple and complex cells, but in this study X and Y cells were distinguished only by the indirect criterion of conduction latency.

Hypercomplex cells have also fallen from their place in the hierarchy. Some hypercomplex cells receive input from X-cells of the LGN (Hoffman and Stone, 1971). Dreher (1972) has suggested that there are two types of hypercomplex unit, those with the properties of simple cells and those with the properties of complex cells: the former prefer slow movement and their receptive fields can be divided into distinct ON and OFF regions, whereas the latter respond best to faster movement and their fields cannot be divided into distinct regions. Rose (1974) has gone further and has argued that many of the cells that would conventionally be classified as simple or complex respond more strongly to short bars than to long and that there is no quantitative basis for regarding hypercomplex cells as a distinct class rather than as merely extreme cases of simple or complex cells.

12. A third type of ganglion cells, W-cells, have been described in the cat's retina. As a class, these cells are marked by very low conduction speeds, but they appear to be very heterogeneous in their properties, resembling many of the units that have been catalogued in the retinæ of frogs and pigeons. Many have receptive fields that cannot be divided into centre and surround and some are directionally selective. Their cell bodies are thought to be very tiny and they are believed to project only to the superior colliculus. (Stone and Hoffman, 1973; Hoffman, 1973; Cleland and Levick, 1974).

The Gnostic Unit

It may be time to level Hubel and Wiesel's hierarchy, but whether the hierarchical model is wrong in principle or only in detail we do not know. Some have been ready to extrapolate beyond Hubel and Wiesel's experimental findings and to suppose that the hierarchical analysis was maintained through successively more abstract levels of feature extraction until was reached a single cell that detected a particular object or person or word. Konorski (1967a, b), who has given this view explicit expression, called such a cell a *gnostic unit*; another popular term has been *grandmother cell*. Barlow (1972) has spoken of *cardinal cells*. Although each gnostic unit can be regarded as lying at the top of a pyramidal structure, any particular receptor element lying at the base of the pyramid may contribute to very many such hierarchies and thus we may equally conceive of the pyramids as inverted. Konorski proposes that the only output from a particular perceptual hierarchy is from the gnostic unit at its apex: perceptual systems, he argues, are built to recognize objects and things, not the features or properties of objects. His conjecture is especially plausible in the case of hearing, for the untrained listener cannot identify the individual components that go to make up the sounds of speech. Man's sophisticated capacity to examine individual features of a visual stimulus depends, Konorski suggests, on the formation of special gnostic units.

The gnostic unit brings the doctrine of specific nerve energies to its ultimate development. The hypothesis has appealed because the means and the nature of recognition are made clear. It has been easy to imagine how linkages might be formed between a gnostic unit and an analogous cell that stood at the apex of an outflowing response hierarchy and controlled a spatial and temporal complex of individual muscle movements.

As we mentioned in chapter 3, earlier versions of the doctrine of specific nerve energies allowed that our immediate perceptual experiences depend not only on which sensory fibres are active but also upon our past experience, upon our expectations and upon interactions within and between sensory modalities. However, the hypothesis of the gnostic unit, as expressed in Barlow's Fourth Dogma (Barlow, 1972), supposes that there is a fixed relation between the activity of a cardinal cell and perceptual experience. When our perception of an object is changed by secondary influences, then what has changed is which gnostic unit is most active. We are to imagine that the input patterns that play upon a gnostic unit correspond not only to abstract descriptions of the sensory stimulus but also to the past and present states of other gnostic units.

The slender chances of finding a gnostic unit by electrophysiological

recording in animals would be maximized by using stimuli that were of biological significance to the species. Funkenstein *et al.* (1971), recording from neurons in the primary auditory cortex of unanaesthetized squirrel monkeys, drew their stimuli from the monkey's elaborate repertoire of socially significant vocalizations and reported that some cells were specific to one particular class of call or to calls with very similar acoustic properties. Gross *et al.* (1972) have described a unit in the inferior temporal cortex of the macaque monkey that responded best to a cut-out silhouette of a monkey's hand that was held against an illuminated background in front of the animal.

What are we to make of such findings? It does seem that the very specificity of the response excludes an explanation in terms of general arousal; and Newman and Symmes (1974) have shown that the specificity to vocalizations of cortical units in the squirrel monkey does not change during changes of arousal that occur spontaneously or are produced by direct stimulation of the midbrain reticular formation. However, the electrophysiologist can never know that he has exhausted the set of stimuli that would lead to response in a particular unit: there may always be untried stimuli that would tease an equal, or stronger, response from the cell. Thus a more recent and thorough study of the responses of cortical units in the squirrel monkey shows that more than one class of vocalization will usually evoke a response from a given cell (Newman and Wollberg, 1973). The silhouette of a monkey's hand that was used by Gross *et al.* closely resembles the regular gratings that have more recently been found to evoke strong responses in primary visual areas.

The Association of Perceptual Qualities

As it stands, the doctrine of the gnostic unit presents an immediate difficulty. Within a particular hierarchy, information that is not relevant to the definition of the object is abandoned. In Hubel and Wiesel's model, for example, information about exact retinal position is lost between simple cell and complex cell. One kind of specificity is gained at the expense of others and the signals converging on the postulated gnostic unit cannot carry their ancestry with them. How then can accidental properties, such as position, be associated with the object? A man who looks at a busy expressway can report immediately that it is the Volkswagen that is yellow and is moving slowly towards Brooklyn whereas it is the Citroën that is silver and is moving swiftly towards Nassau County.

How is this perceptual synthesis achieved? Until now in this chapter we have been proposing that each of the major senses should be regarded as several sense-organs rather than as one and we have suggested, for example,

that the several properties of the retinal image (position, brightness, colour, direction and speed of movement, shape and so on) are analysed independently. We must now look more guardedly at the nature and extent of this independence.

The problem is one of the most acute, most central and least answered problems in perceptual theory. At the onset, however, we can exclude the two extreme possibilities. We cannot have a system in which the properties are not separated at all. It would be no good having single units, or cell assemblies, that were specific to yellow Volkswagens moving left at a distance of three metres or to Grandmother wearing her red shawl and bending to pick up a knitting needle. The primary difficulty here is not the numerical one of the number of cells required but the fact that we need to explain how we are able to recognize a Volkswagen for what it is independently of its accidental properties such as colour and position. Equally, however, the dissociation cannot be complete, for, as we have seen, accidental properties can be associated with their objects.

Let us consider briefly the types of psychological and physiological evidence that can be brought to bear on the question of the association of sensory qualities. We shall concentrate, in our illustrative examples, on the degree to which colour and form are dissociated in visual analysis.

Phenomenal Dissociation

In 1934 the British Post Office adopted a new design for its postage stamps. Plate I shows a halfpenny stamp of the old issue and a three-halfpenny stamp of the new issue. The changes from the typographic issue on the left to the photogravure issue on the right were slight ones: the fine etching of the typographic issue was absent on the newer issue; and the newer issue was the more saturated in colour. The reader is invited to try an experiment suggested by Creed (1935). Fuse the two stamps of Plate I by means of a stereoscope (see chapter 11 p. 293). Most observers see a brown stamp, differing only slightly in colour from the three-halfpenny stamp, but they also see the etching that is present on the cheek of King George V in the halfpenny stamp. Thus the hue of the fused image is drawn from one eye and the form is drawn from the other.

Here we have a dissociation of hue and form that is similar to the dissociation of movement from position or shape mentioned at the beginning of this chapter (p. 71). Another example is seen in the demonstration by Treisman (1962) that the colour of a stimulus may be suppressed in binocular rivalry¹³ while the form of the stimulus still contributes to stereopsis.

13. When the two eyes are simultaneously stimulated by stimuli that are very different in form or in colour, the two stimuli often are not perceptually fused but are seen in rapid alternation. This alternation is known as *retinal* or *binocular rivalry*. (See chapter 11 p. 298).

Contingent After-effects

Evidence for the partial association of properties in analysis has been sought in the curious phenomena called *contingent after-effects*. Using Plate II the reader may observe for himself one of the most notable of these, the McCollough effect. This mysterious effect cannot be explained in terms of ordinary after-images (see chapter 6 p. 189), since any particular point on the retina has been stimulated equally by red and green light during the adaptation period. The illusory hues are contingent: they appear only when gratings of the appropriate orientation are present. It is argued that the McCollough effect reveals the existence of neurons that are specific to both colour and orientation. We have to assume that the apparent colour of a grating normally depends on the relative activity of orientation detectors tuned to different colours. During the adaptation period there occurs a selective depression of neurons specific to, say, green bars tilted 45 degrees to the right and of others specific to red bars tilted to the left. A black and white grating tilted right will then come to look pink.

Analogous after-effects of colour have been found that are contingent on direction of movement and on spatial frequency. Equally there is a tilt after-effect that is contingent on colour and there are movement after-effects that are contingent on colour or on spatial frequency. Table 4.1, which has been prepared by Dr. P. Thompson, conveniently summarizes most of the visual contingent after-effects that have so far been described. (For further references the reader may turn to the review by Skowbo *et al.* 1975.) In every case it is possible to argue that there are neurons selective for two, or more, properties of the stimulus.

In some ways, however, the contingent after-effects are more akin to the phenomena of conditioning. Firstly, for example, they persist much longer than sensory after-effects are normally thought to: ten or fifteen minutes of adaptation may produce a McCollough effect that lasts for days, weeks or months. Secondly, a McCollough effect is dissipated much more rapidly if the subject is repeatedly exposed (for example, in repeated testing) to black and white gratings after the end of the adaptation period (Jones and Holding, 1975; Skowbo *et al.* 1975) and this process resembles the extinction of a learned response. Thirdly, contingent movement after-effects have been reported to be stronger half-an-hour after the adaptation period than they were on immediate testing (Mayhew and Anstis, 1972); this is the phenomenon that in learning theory would be called *reminiscence*. But if the McCollough effect is a kind of conditioning why are the illusory colours roughly complementary to those associated with the gratings during adaptation? We have to suppose that the unconditioned response of the visual system to an excess of redness in the world is to turn down the gain

Table 4.1.

		Dimension of After-effect				
		Colour	Movement Direction	Spatial Frequency	Orientation (Tilt)	Brightness
Dimension of Contingency	Colour		Mayhew and Anstis 1972 Favreau, Emerson and Corballis 1972	Virsu and Haapasalo 1973	Held and Shattuck 1973	
	Movement Direction	Hepler 1968 Stromeyer and Mansfield 1970				Mayhew and Anstis 1972
	Movement Velocity	Hepler personal comm. to Stromeyer and Mansfield 1970				
	Spatial Frequency	Harris 1970 Breitmeyer and Cooper 1972 Stromeyer 1972 Leppman 1973	Mayhew and Anstis 1972 Walker 1972		Wyatt 1974	
	Orientation (Tilt)	McCollough 1965	Mayhew and Anstis 1972	Wyatt 1974		Mayhew and Anstis 1972
	Curvature	Riggs 1973 Sigel and Nachmias 1974 Stromeyer and Riggs 1974				

of red-sensitive channels and that when later, in the testing situation, the appropriate grating (the conditioned stimulus) elicits this response the phenomenal result is that the black and white test-stimulus appears greenish. However, this learning process will have to be a rather special and local one: traditionally the McCollough effect has been produced by alternate presentation of red and green stimuli, but both red and green are concurrently present as we adapt to Plate II.

The problem of the association of properties during analysis has also been approached through the conventional, non-contingent after-effects:¹⁴ it is asked whether an after-effect that has been established in the presence of some other sensory quality is still present when this second quality is changed. For example, Virsu and Haapasalo (1973) have shown that the Blakemore-Sutton effect (the spatial-frequency after-effect that we discussed earlier in this chapter) is as strong when adapting and test stimuli are heterochromatic as it is when they are of the same colour; and so they conclude that there are some channels at some stage of the visual system that are specific to spatial frequency and not specific to colour.

Electrophysiology

We have already mentioned electrophysiological evidence for units that are specific to more than one property of the stimulus, but the interesting question is whether this multiple specificity increases or decreases at successive levels of analysis, and here the answer, in the case of our example of colour and form, is contradictory. Hubel and Wiesel (1968) suggested that the analyses of form and colour were progressively dissociated in Area 17 of the monkey's cortex: the ON and OFF areas of the receptive fields of six out of twenty-six simple cells had different spectral sensitivities, whereas only 7 per cent of complex cells and a very few hypercomplex cells were colour-coded. Gouras (1974) has systematically examined the spatial and chromatic properties of cells in Area 17 and has reached a similar conclusion: most of the cells that had concentric receptive fields were opponent-colour cells whereas no hypercomplex cell in the sample was colour-selective. Yet Yates (1974) has reported the very contrary, that the more narrowly tuned is a cell to colour the more likely is it to have complex field properties. Zeki (1973), recording from an area he calls V4 in the prestriate cortex of the monkey, finds colour specificity associated with most of the types of receptive field described by Hubel and Wiesel.

However, Zeki's work addresses our question at a second level. He finds that almost every cell in V4 is specific for colour, whatever its other stimulus requirements. Elsewhere he finds an area of cortex in the rhesus

14. It could be argued that the majority of visual after-effects are contingent, if only upon retinal position.

monkey (the posterior bank of the superior temporal sulcus) where every cell is sensitive to movement and most are directionally specific; again the dominant feature was associated with varying degrees of specificity for form, but specificity for colour was apparently absent (Zeki, 1974).

With this suggestion that there are anatomically distinct centres for different stimulus properties we end our discussion of the evidence for independent analysis. Our purpose has been only to illustrate the types of finding that bear on this question. We have seen good evidence for partial independence; but the best argument against complete independence remains the *a priori* one with which we began. If we wish to retain a hierarchical model we must allow information to flow between hierarchies and in both directions within a hierarchy. There must be access to the Volkswagen's pedigree.

The Development of Neural Specificity

Lurking uncomfortably behind successive versions of the doctrine of specific nerve energies has always been the question of how neurons come to gain their specificity. Are animals equipped genetically with the rich and intricate pattern of feature-detectors that they need in order to face the environment into which they are born; or does the environment itself impress an organization upon an unschooled nervous system? Recent research begins to answer this ancient question. Few branches of sensory physiology are now so active and so intriguing.¹⁵

At first it seemed that much of the apparatus found in the visual cortex of the mature cat was present as soon as the kitten opened its eyes; for Hubel and Wiesel (1963) reported that both simple and complex cells could be found before the kitten had had any visual experience. These cells resembled those of the adult in every way except in their sluggishness and in the readiness with which they habituated to repeated stimulation. The question, however, has become controversial. Barlow and Pettigrew (1971) bluntly denied that there were any cortical neurons in the inexperienced kitten that were truly orientation-specific. A current view allows that some units are orientation-specific (Blakemore, 1974b; Barlow, 1975); but most units are not and it is clear that the inexperienced cortex is very different from that of the adult. Cells may have inputs from both eyes, they are almost always specific to the position of a stimulus, and they may be specific to the direction and velocity of movement, but most units are at once promiscuous and fickle in their response; and fine tuning for orientation and for binocular disparity waits upon visual experience (Pettigrew, 1971). That this is so is not remarkable in the light of anatomical evidence that less than 1 per cent

15. The reviews of this field by Blakemore (1974a, b) and by Barlow (1975) are recommended.

of synapses in the visual cortex develop before eye-opening in the kitten (Cragg, 1975).

The remarkable influence of early experience has been revealed by experiments in which animals have been reared in selective visual environ-




Fig. 4.9 An apparatus for exposing kittens to an environment containing stripes of only one orientation. The kitten wears a black ruff that prevents it seeing its own body. (From Blakemore and Cooper, 1970.)

ments. Hirsch and Spinelli (1970; 1971) brought up kittens in the dark for three weeks after birth and then fitted them with goggles that allowed one eye to see only an array of three vertical bars and the other eye to see only three horizontal bars. When, after several weeks of this restricted visual experience, recordings were made from the visual cortex of such a kitten, it was found that almost every cell was driven from only one eye and that if there was a stimulus orientation that excited the cell preferentially, then it

corresponded to the orientation to which that eye had been exposed. Blakemore and Cooper (1970) reported a slightly different experiment with a similar result. They reared kittens in an environment containing only vertical lines or only horizontal ones. The kitten, fitted with a broad, black collar to prevent its seeing its own body, stood on a clear glass platform inside a tall cylinder covered with stripes in one of the two orientations (Fig. 4.9). When electrophysiological recordings were later made from such kittens it was found that the preferred orientations for cortical cells were clustered around the particular orientation to which the kitten had been exposed (Fig. 4.10). Similarly, when kittens are exposed to an environ-




Fig. 4.10 The distribution of optimal orientations for 52 cortical neurons from a cat brought up in a visual world containing only horizontal stripes (left) and for 72 neurons from a cat exposed to vertical stripes (right). Each line represents the orientation to which one particular cell responded maximally. In the normal cat the optimal orientations are distributed randomly around the clock. (From Blakemore and Cooper, 1970.)

ment containing only coarsely-spaced white spots on a dark background, most cells come to respond as strongly to spots of a corresponding size as they do to the long lines that are normally so effective (Van Sluyters and Blakemore, 1973).

Is there a *critical period* during which the infant visual system can be modified especially readily? For the cat the effects of exposure to a particular orientation are greatest between the ages of 3 weeks and 14 weeks (Blakemore, 1974b). If one of a kitten's eyes is closed between the ages of 3 weeks and 8 weeks, there is a marked decline in the number of cortical cells driven by that eye, but this susceptibility to eye closures declines thereafter and disappears around the end of the third month (Hubel and

Wiesel, 1968). Monocular deprivation of this kind during the critical period has a permanent effect: even if the deprivation is followed by a long period of normal visual experience the properties of cortical cells remain abnormal (Wiesel and Hubel, 1965).

During the critical period an exposure to a biased environment of only a few hours may change the responses of single units. However, the effects may take time to develop: Pettigrew and Garey (1974) exposed kittens to a vertical grating for 5–22 hours and found that if they immediately made electrophysiological recordings from the visual cortex then many cells responded to a broad range of stimuli and some responded to all orientations except the one that had been used during the initial exposure. After a number of hours had elapsed, many more cells responded to vertical stimuli with the fineness of tuning that is characteristic of the adult neuron.

How does selective experience exercise its dramatic effect on the visual cortex? If a kitten is exposed to vertical lines, do cells that might normally have been destined to respond to horizontal lines degenerate through desuetude, or do they instead come to respond to the prevailing stimulus?

The evidence suggests that the latter is the case. Firstly, the electrophysiologists do not find large areas of silent cortex that would correspond to the missing orientations and indeed the density of active cells is similar to that of the cortex of a normal animal (Pettigrew *et al.* 1974). Secondly, an unusual early environment may create units of a type not found in normal animals. Thus Hirsch and Spinelli (1971), in the study described above, found a few cells that had a receptive field consisting of three, parallel, bar-shaped regions and corresponding in configuration and size to the array of bars to which one eye had been exposed. Similarly, alternating exposure to vertical and horizontal bars may produce some units that are excited by both orientations and respond particularly well to a cross moving through the receptive field (Blakemore, 1974b).

The findings of the developmental electrophysiologists would appear to have important clinical implications for man. We should avoid subjecting the infant to a restricted or biased visual environment. It is known, for example, that if an eye is astigmatic in the child no amount of optical correction can later restore normal acuity in the adult, and an explanation of this loss might be seen in the finding that kittens reared wearing cylindrical lenses lack cortical neurons sharply tuned to the orientation that has been blurred by the lenses (Freeman and Pettigrew, 1973).

We may end this chapter by concluding that higher mammals are equipped with a comprehensive armoury of feature-detectors and that, in phylogeny and in ontogeny, these feature-detectors are matched to the environment in which an animal finds itself. Whether the same principles are maintained in the analysis and recognition of objects, we do not know.




Plate I The halfpenny stamp on the left (a) is from the typographic issue of King George V; the three-halfpenny stamp (b) is from the photogravure series introduced in 1934. The stamps above and below b are to allow comparison between the fused image and the two monocular stimuli. If the brown stamp does not immediately dominate when you fuse the two, try attenuating the halfpenny stamp with a neutral filter (e.g., one lens of a pair of sunglasses). Colour reproduction is never perfect and the experiment is best attempted with the original stamps. (The 1934 photogravure issue can be identified from its dimensions, 22.25 mm \times 18.3 mm; later issues were slightly smaller.) To comply with Post Office regulations it has been necessary to insert the oblique bar on the stamps.

Plate II To obtain the McCollough effect, gaze at the red and green pattern (A) for about three minutes. The viewing distance should be about two feet. Do not let your eye rest on any one point for very long and try to look as frequently at red as at green areas. Resist the temptation to tilt your head to one side.


Looking at the uncoloured pattern (B) you should see illusory colours that vary according to the orientation of the tilted lines: where the lines are tilted to the right you will see pinks and where they are tilted to the left you will see pale greens. Don't expect the illusory hues to be very strong: the effect is scientifically, rather than phenomenally, striking. Try tilting the book, or your head, 90 degrees to one side: the apparent colours should exchange positions.

One of the most remarkable aspects of the phenomenon is its persistence. Try testing yourself again after half an hour. If you gaze at the coloured figure for rather longer than three minutes—for, say, quarter of an hour—you may produce an effect that survives for days or weeks.

To obtain the basic effect some readers will need to look at the coloured pattern for a little longer than three minutes, some for a little less. Since the phenomenon, once established, is very stable, it does no harm to glance occasionally at the uncoloured patterns.



(a)



(b)