

Understanding Natural Vision

H.B. Barlow

Physiological Laboratory, Downing Street,
Cambridge CB2 3EG, United Kingdom

Abstract

Because many of you may not be familiar with the biological structures with which we ourselves process images, I shall start with a brief anatomical tour of the complex and incompletely understood visual system of primates. The biological components which perform the computations may also be unfamiliar, and it is interesting to compare them with physical devices: synapses and neurones are very slow and have a very limited dynamic range, but they are packed at a very high density per unit volume, are very freely interconnected over limited distances, and these interconnections possess self-regulatory properties that can compensate for errors in construction or changes of use.

The problems facing physiologists or psychologists studying natural vision are also very different from yours. First, it is technically very difficult to monitor the activity of just one or two elements at a time in a structure that contains 10^{10} neurones or more; second, with such incomplete information it is hard to interpret the results, and third, there is the constant danger of interfering with normal function. Also, we often fail to understand the goal or purpose of the neuronal interactions that occur in a biological structure; this should not be a problem for physicists because I cannot imagine one of your devices ever doing something clever that its designer had not intended it to do!

The overall achievement of a good visual system is impressive. It does far more than simply provide us with a good picture to look at, for the looker is part of the system. It gives us most of our knowledge of the world, which requires organisation of the results to minimize redundant representation and statistical testing to ensure reliability. All this is done on the spot, in real time, by a system weighing about two kilograms (including power supplies). So far the flow of important concepts has been from physics to biology, but if we knew exactly how a real visual system worked the flow might be reversed.

What is Vision?

One tends to think of natural, biological, vision as the process by which an image of the world around us is formed optically by the eye, transduced into nervous activity in the retina and transmitted to the brain by its optic nerve fibres. But seeing is much more than a matter of getting an image upstairs and into the brain: it includes the process of interpreting the

image and thereby gaining knowledge of the real world that gave rise to the image. It is especially important for me to make this distinction because I am sure most of you are more interested in the second process and that is what you would like me to tell you about; but it is only the first process that we are actually beginning to understand, so that is what I am forced to talk about. With guidance from physicists and engineers we are beginning to see how a real brain tackles a simple task of information engineering, moving an image from one place to another. For understanding the rest we also hope to benefit from those who actually perform the task of gaining knowledge from artificial images and that is one of the points of this symposium.

The following sketchy description of a very complex system will, I hope, give you some feeling for the results we have obtained and the problems that face us.

Macrostructure or Gross Anatomy

Figure 1 is a schematic view of the human visual system from above. The first point of functional interest is the rearrangement of nerve fibres from the two eyes at the optic chiasma; this has the effect of bringing the messages from the right hemifields seen by both eyes to the left cerebral hemisphere, those from the left hemifields of each eye to the right hemisphere. It is quite a shock to realise that the two halves of your visual field are represented in separate structures several centimetres apart. There is, however, a large bundle of nerve fibres called the corpus callosum joining the two halves, and this is in fact the largest tract of fibres in the body.

The structures shown here are composed of nerve cells and nerve fibres which influence each other through synapses. The approximate numbers of elements are indicated, and it also shows the order of the elements, which are synaptically connected in series. In the retina the light is absorbed in receptors (1st order) whose internal potential is thereby affected. These make a synapse with bipolar cells, and by graded release of chemical transmitter substances the internal potential of these bipolar cells is influenced, and they are therefore called 2nd order cells. These in turn influence the retinal ganglion cells whose long processes or axons constitute the optic nerve; these are 3rd order cells. This is much over-simplified, for there are two other types of cell in the retina, horizontal and amacrine cells, that provide other links between ganglion cells and receptors; in some cases amacrine cells may be interposed between bipolar cells and ganglion cells, which would make the ganglion cells and their axons 4th order. There is also some evidence for a direct connection from receptors to ganglion cells, which would make them 2nd order. In the cortex the pathways from cell to cell are so complex that it becomes even less possible to specify the order of a cell, but the numbers I have given do roughly indicate the position in the sequence.

Note that the number of components decreases from above 10^8 in the retina to about 10^6 in each optic nerve. The 10^8 does not represent the number of pixels, for the optical quality of the image in the periphery would not support that number. But it would support more than 10^6 , the number of optic nerve fibres, and it is interesting to see how the retina deliberately degrades the quality of information transmitted from the periphery of the visual field: each fibre is excited by light falling in a large patch of roughly gaussian-shaped sensitivity profile, these patches getting progressively

larger in size and fewer in number towards the periphery, so that resolution is deliberately degraded. The minimum resolvable angle is about 1 minute of arc at the fovea, and increases by about 1 minute for every 2 degrees of eccentricity, so is down to about 6 minutes at 10 degrees. Of course a shift of eye position changes the selection of information that is preserved or discarded, and it is important to realize that the human eye continually makes such a new selection by flicking from one place to another in the visual field. The frequency of these flicks varies greatly, from 2 or 3 times a second if you are reading or keenly interested in what you are inspecting, down to one flick every 2 or 3 seconds if you are in a more leisurely mood; but it is rare for 5 seconds to elapse without a flick. Of course the vast majority of shifts are made automatically without conscious effort; like breathing, in fact, but at a much greater repetition rate.




Fig. 1. Schematic view of the visual system viewed from above, with insets showing retinal neurones (left) and the laminae of visual cortex with a single pyramidal cell (right). The ordinal numbers indicate the serial order of the cell, those in parentheses the total number of cells of that type

In the cerebral cortex the numbers of elements increase to approximately 10^9 in the primary visual area (V 1) (also called the striate cortex or area 17) and another 10^9 (or possibly much more) in the surrounding secondary visual areas (also called parastriate, peristriate, V 2 to V 4 or areas 18, 19, 20). In the fourth layer of V 1 we know that there is a very precise representation of the visual field. It is grossly distorted because cortical area is apportioned according to the number of entering nerve fibres, rather than according to degrees of solid angle in the visual field; as we saw the resolution and density of ganglion cells decline rapidly away from the fovea. But it is precise in the sense that activity at one point corresponds unambiguously to something happening at one point in the visual field, and consequently it is possible to read backwards from the map to the real world: if there is activity at a particular spot in the map, this implies something happening at a particular point in the visual field of the eyes. Note, however, that there is more to be said about what that something is, and also that eye and head positions need to be known if you wish to point your finger to the position of the something in the field. Quite accurate topographical mapping also exists in the other layers of V 1, and there is evidence for quite good mapping in V 2, but in V 3 and V 4 it is very much less precise and appears to lack this unambiguous character that allows one to read back reliably from the map to the real world.

Retinal Components and what they do

Figure 1 also has a diagram of the retina, where it is possible to record from the nerve cells as they interact with each other [1]. Note first their size; though some are smaller, some larger, $10\ \mu\text{m}$ is a reasonable figure for the size of a cell, and it is therefore possible to pack $10^6/\text{mm}^3$. This would leave little space between them, but packing densities close to this are certainly achieved, and may be exceeded in some places.

Now a word about how they work. The ionic concentrations of the fluid inside the cell are maintained by metabolically driven "pumps" at different values from those in the surrounding fluids. Electro-chemical gradients therefore exist across the cell membranes, and ions will pass down these gradients at a rate that depends upon the voltage difference between inside and outside, the concentration ratios for each particular ion, and the permeabilities of the cell membrane for each particular ion. Changes of these permeabilities in turn change the internal potentials of the cell. Take the receptor cell; without light the membrane is quite freely permeable to Na^+ ions, and since there is a strong inward electrochemical gradient for Na^+ , there is a strong inward current in the resting condition. When light falls on the cell this inward current is reduced (by mechanisms that are not fully known) and this results in the inside of the cell becoming more negative. This change in potential spreads to the synapse, at the end of the cell remote from where light is absorbed, and there it causes a reduction in the release of a transmitter substance, whose nature is again unknown. This transmitter influences the permeability to ions of the next cell in the chain, the bipolar cell, with consequent effects on its own internal potential and the amounts of transmitter it releases at its own synapses.

And so it might continue, except for this difficulty: the potential at one end of a cell is not necessarily the same as that at the other, because if you do a simple calculation from the known ionic concentrations you will find that the resistance of the cytoplasm is considerable, and since that of the membrane is not infinite there will be substantial decrements of voltage

along a cell. If the distances are only fractions of a millimetre this decay is tolerable, but passive conduction would be quite inadequate to convey potentials from retina to brain, let alone from your brain to your foot. Hence it is no surprise that in nerves longer than a millimetre or so information is transmitted by a different means: it is done by varying the number of all-or-nothing, self-sustaining, pulses transmitted along the axon. This is much more reliable, for the pulses almost always seem to reach their destination, but there is probably a considerable reduction in the dynamic range of signals. The maximum instantaneous rate of firing of impulses is about 1000/sec, but average rates rarely exceed 200/sec, so in a short period of, say, 1/20 sec only about 3 bits could be transmitted. This is obviously only a very small fraction of what you might expect from an electronic line.

Note that the potentials are slow also in their rise and fall times. The time to peak of a receptor in response to a brief flash is about 30 msec, during which time a fast-bowled cricket ball will have travelled over a yard. In spite of the slowness of the receptors the precision of judging the positions of moving objects corresponds to a time of about 1 msec or better, and it would be interesting to know how this is achieved.

So far I have outlined the biophysics of the retina, but what does it achieve in the way of information processing? Some of these steps are illustrated in Fig. 2, and they can be described very easily, though perhaps too simply, as 1) automatic gain control; 2) signalling deviations above and below the mean instead of absolute values; 3) high-pass temporal filtering; 4) high-pass spatial filtering; 5) conversion of graded signals into varying numbers of output pulses. The first occurs largely in the receptors: when the mean light level changes up or down, the input-output characteristic shifts bodily to higher or lower levels. The second occurs at the bipolar cells and is illustrated in the waveforms of Fig. 2. This also shows the transient character of the discharges from amacrine and ganglion cells. The fourth is brought about by a laterally connecting pathway through horizontal cells, which antagonizes the direct pathway through bipolar cells; the result is that the bipolars signal something approximating to the second spatial derivative of image luminance. The production of pulses is first seen in amacrine and ganglion cells.

Thus a network of elements such as exists in the retina can do tasks equivalent to spatial and temporal filtering, as well as automatic gain control and generation of a pulse-frequency modulated train of pulses. It is not too hard to see how this can be achieved by the interaction of graded signals, together with a threshold device for pulse generation, but we know from studies in other animals that such a network can achieve more. In the rabbit there are retinal ganglion cells that signal the direction of movement of the image over the retina, as well as the presence of oriented objects in the image [2, 3]. I shall not go into these here because there is no evidence for such pattern selectivity in the eyes of man or primates, but their presence in other animals raises an important point. So far I have spoken as if the cell was the important element, and mentioned that they could be packed at $10^6/\text{mm}^3$. But the evidence suggests that the crucial non-linear interaction that gives rise to pattern selectivity can occur at each individual synapse [1,4]. In the layer where this occurs electron-micrographs show [1] that there are some $3 \times 10^8/\text{mm}^3$, and double this figure has been reported in the human cerebral cortex [5]. Synapses may be slow, but they are certainly compact.

There is one other point to discuss in connection with the retina: What is the need for the five operations I said it did? The need for pulse gener-




Fig. 2. Waveforms that are recorded from various cells in the vertebrate retina when a small stimulus spot is shone on to the retina centred on the cell, and when a larger spot that includes surrounding elements is used. The stimuli last about 1 second, and the responses are up to 30 mV in amplitude; in ganglion cells and amacrine the pulses would be larger if recorded with a full frequency response. Depolarization (i.e. reduction of intracellular negativity) is plotted upwards. (Adapted from ref.[2])

ation is clear enough, for graded potentials would decay to nothing before reaching the brain. The need for automatic gain control obviously arises from the enormous ($>10^{10}$) dynamic range of the input signals which have somehow to be passed down output channels of miserably narrow dynamic range, and signalling deviations from the mean also helps. The need for the other two can perhaps be explained in the same way, for it would be the low frequency signals that would tend to cause overload and saturation. But an explanation that is superficially quite different is sometimes given, namely the need to emphasize the borders of objects, or the temporal changes resulting from image movement, because these are "important" characteristics of images. Which of these explanations is correct? Personally I think that the two explanations are related: if one regards compression of the message while preserving its information, that is the reduction of redundancy, as the prime task, this will both tend to reduce the load on optic nerve fibres, and may also explain the psychological importance of transients. But we cannot go into predictive coding and such at this point.

Cortex

The retina performs an intelligible task, and we understand quite a lot about how it does it, but the situation is very different in the cortex, for it is an exceedingly complex structure and also very large. In the Rhesus monkey 60% of the whole surface area of the cerebral hemispheres is thought to be intimately concerned with vision [6]. One can divide these visual areas in the cortex into the primary area also called V1, or striate area, and other

areas called V 2, V 3, V 4, peristriate, parastriate, and so forth. We do not know for certain what tasks these areas perform, even in the case of V 1. Instead of burdening you with a mass of uninterpreted information I shall present a story which is fairly coherent, as far as it goes, but be warned that I have made a biased selection of facts. What I shall say is certainly grossly incomplete and could be wildly wrong.

First look at the primary visual cortex (V 1), about which we know most, thanks mainly to the work of HUBEL and WIESEL [7]. The majority of the fibres from the retina, after relaying in the lateral geniculate nucleus, terminate on cells in layer IV. As you will see from the average separation of two fibres I have indicated, there are a vast number of cells available for each incoming fibre, and there is some evidence that the first task is to recreate among these closely packed cells in layer IV a very precise representation of the 2nd spatial derivative of luminance in the image [8]. This is a much simplified view, for I have ignored evidence about the different terminations of different classes of fibre, about the segregation of the endings from the two eyes in this layer, about colour, and so on, but this idea of the re-creation of the differentiated image, though oversimplified, will have to do for the present.

Now what is this map used for? The first step towards finding an answer is to look at the anatomy: Where does each cell pick up its information from? The answer for the cells representing the map is of course "from the incoming fibres of the optic radiation". But the cells in other layers must also be influenced predominantly from what is mapped locally, because the lateral connections to each cell are most dense in the immediate vicinity and are restricted to a total range of a few millimetres. We do not know the full story about the distribution of lateral spread of connections to a cell, but the total distance across V 1 is several cms, and few would deny that any given cell must be influenced predominantly by what is represented in the local parts of the map close to where it lies.

Next we need to see where the information is transmitted. It has emerged over the last ten years or so that the cells in different layers of the cortex send their messages to different parts of the brain [9]. The axons from the two lowest layers, VI and V, go back towards the region of the LGN relay and to a lower visual centre that I have not mentioned, the superior colliculus. The existence of a large population of fibres feeding back to the LGN is certainly interesting, but so far no-one has succeeded in finding out what they do or in making any theoretical sense of it. The fibres going to the superior colliculus make more sense, for this structure is concerned with the control of eye movements [10] and layer V is presumably important in helping to decide where the eye shall next fixate. However from the viewpoint of image processing the cells in layers II and III are most interesting, for they have the task of relaying information to the other visual areas. Thus the anatomy seems to be telling us that the first major step in image processing (after gain control, and spatial and temporal filtering) is to perform some kind of local analysis on each separate patch of image; this information is then passed on to other parts of the brain.

Now let us look at some physiological evidence (a film produced by HUBEL and WIESEL was shown at this point). Figure 3 shows diagrammatically what you find if you record from cells at various positions in the visual system and map what is known as their "receptive fields". This is the region of the visual field in front of the eye from which responses from a particular cell can be obtained. Initially one uses a small spot of light and records whether you get a response when it is turned on, when it is turned off, or




Fig. 3. Maps of the receptive fields that can be obtained from neurones at different stages in the visual pathway. (Adapted from ref. [2])

not at all. This figure shows the "centre-surround" arrangement found in the retina and LGN, which if you think it through you will see is what you would expect from a cell performing high-pass spatial filtering, or taking (on a rather coarse scale) the second spatial derivative of local light intensities. What is new in the cortex is the fact that the orientation of the stimulus becomes critically important, and this was one of the main findings made by HUBEL and WIESEL when they first recorded from cortical neurones [1] .

What cannot be shown in this figure are the results of more quantitative and extensive investigations of the properties of cells in V1. They are selective for size as well as orientation, and some people [12, 13] believe that the activity of each cell in a set of cells represents the coefficient of one term in a local Fourier transform of the luminance distribution in a local region. We shall be hearing more about this view later on in this conference.

In addition to selectivity for orientation and size one finds that many cells require the stimulus to move in one direction rather than its opposite. Colour is also an important parameter in many cases [14], and finally so is binocular disparity, that is the precise relative alignment of the images in

the two eyes [15]; this is obviously correlated with the distance of a stimulus object from the eyes.

This list of the local properties of the image for which neurones in V1 show selective sensitivity is interesting because it correlates quite well with the characteristics that Gestalt psychologists thought were responsible for causing segregation of the image into parts, and especially the separation of "figure" from "ground". The writings of this school are not easy to follow, but the principle was very clearly expressed in GUZMAN's work [16] in Artificial Intelligence at MIT. He devised a computer programme to derive an account of the positions and dimensions of a set of wooden blocks, given a single line drawing of them from one viewpoint. Since many of them were partially obscured by other blocks he required a means of identifying whether an edge seen to the left of an obstruction belonged to the same wooden block as an edge seen to the right of the obstruction. Co-linearity of the edges was the obvious characteristic to test for, and he named this test a search for a "linking principle" because, if positive, it enabled his programme to link two features as belonging to the same block. Similarity of colour, disparity, motion, and texture are other local characteristics of an image that can be used to establish a high a priori probability that particular parts belong, in some sense, to the same object, and it seems that this is likely to be why these characteristics are detected in V1 and signalled elsewhere [17].

Our knowledge of the other visual areas is still rudimentary but there is evidence that particular areas specialize in colour or motion [18]; perhaps there is one such area specialized for each linking principle. This seems to make sense, and one is tempted to predict that the variables by which a particular linking feature is specified will turn out to be important in the anatomical arrangement of these other areas. The nerve cells of V1 require an accurate topographic map in order to have access to the information that enables them to detect a local linking feature, but to detect the coincident occurrence of the same linking feature in separate parts of the visual field one needs to rearrange the information according to new principles: one needs to collect together information about particular characteristics regardless of whereabouts in the visual field they occur. At the moment there is little evidence that the brain makes such non-topographic maps [17], but time will tell.

Properties of Cortical Components

It is not easy to record the intracellular potentials of cortical neurones, but as far as is known they interconnect and communicate with each other in much the same way as the retinal neurones. However experiments have shown one important property, namely that these connections are not fixed solely by the mechanisms responsible for the development of the brain, but are subject to a limited amount of modification by the use, or lack of it, to which the connections are put during an early phase of maturation [19]. This is not too surprising, for something in the brain must be modifiable because we all learn and remember. Also, it has long been recognised that in young animals a function normally performed in one part of the brain can be taken over by another part following injury or experimental removal of the first part. This happens even with a very complex function such as speech, for injury to the dominant left hemisphere in a child does not prevent the acquisition of good capacity for language in the non-dominant hemisphere, whereas the equivalent damage to an adult would cause profound and permanent defects.

We have no idea how such things happen, but a simpler form of plasticity can be demonstrated in the neurones of the visual system by rearing animals in abnormal visual environments. I cannot go into these experiments here, but I think they require for their explanation a mechanism which was postulated by HEBB many years ago [20, 21]. If activation of a synapse is followed by activity of the postsynaptic neurone, then the effectiveness of that synapse increases. If the synapse fails to activate the neurone, then not only does it become less effective, but also the whole of the presynaptic neurone seems to be adversely affected: its cell body shrinks, and its terminal arborisations wither and become less extensive. Somewhat similar positive and negative effects are produced in peripheral nerves of the sympathetic nervous system by excess or deficiency of a well-characterised substance called nerve growth factor. It is therefore tempting to postulate that "synaptic rewarding factors" are released by activated neurones and picked up by the activating synapses; these would regulate the effectiveness of central synapses and promote growth of the presynaptic cell and sprouting of the successful terminal branches. Components with such adaptable properties would, I imagine, be exciting for electronic engineers, though also disconcerting.

The role of Psychophysics

What I have said so far may have given the impression that what we know about the visual system has been derived almost entirely from anatomical and physiological investigation. Nothing could be more misleading, for all the important properties of the visual system were first established by psychophysical and psychological observations made on the system working as a whole. For example, more than 100 years passed between the establishment of trichomacy as a fact of human colour vision, and the detection within the last fifteen years of three different cone photoreceptors containing three different photosensitive pigments. It is the same with sensitivity, for it was established forty years ago that a rod photoreceptor must signal the absorption of a single quantum, but these signals have only been directly recorded within the last couple of years. Anatomists and physiologists need to be told what the visual system does before they can set about the difficult task of finding out how it does it.

Unique Features of Natural Vision

Table 1 attempts to summarise what goes on in the early stages of natural vision, but what makes it different from artificial image processing? Table 2 attempts to summarize the unique features of natural image processing. I have already mentioned some properties of the components: they are very slow and have a narrow dynamic range, but they can be packed at very high density and have an interesting capacity for self-regulation. The parallel organization of elements is also unusual, and I am sure most of you would agree that nature must have some interesting programmes up her sleeve for handling images, but I do not think anatomists and physiologists are likely to discover what they are without the help of those who handle artificial images. DAVID MARR [22] has argued persuasively that the "computational theory of vision" must be understood before much progress can be made, and I think there is a possibility of fruitful interaction here, provided that we listen to each other. But I think there is a major difference between the aim or purpose of our own visual system and that of most artificial image processing, and I want to enlarge on this briefly.

Table 1. Suggested first steps in natural image processing

Retina	Lateral Geniculate Nucleus	Primary visual Cortex (V 1)	Other cortical areas
Automatic gain control	Aligns, but does not combine inputs from the two eyes	Detects local features of image that are likely to belong to same object —	? Specialize in individual linking features
Signals deviations from mean not absolute luminance	More high-pass filtering	linking features:- Orientation Movement Colour Texture Disparity	? Mapping according to non-topographic variables
High-pass temporal filtering	Controls input to cortex		
High-pass spatial filtering			

Table 2. Unique features of natural image processing

Components	Organization	Purpose
Very slow (<1000 Hz)	Junctions arranged in parallel	To represent visual scenes by activity of a sparse selection of reliable and non-redundant (i.e. independent) elements
Narrow range (<8 bits)	Pattern of connections determines program	
Very many very small junctions ($10^8/\text{mm}^2$)	Many aspects not understood	
Self-regulatory property		

What our own visual system does for us can be crudely described as follows. From a view lasting a second or two it provides reliable knowledge of the positions in space of some dozens of familiar or semi-familiar objects. Let us do a very approximate calculation on the amount of information this represents. Suppose there are 100 objects, each requiring two 16-bit words to identify in our internal catalogue of known objects. For each we might need 3 words for specifying position, and another 3 for orientation. Even with these quite generous figures we only seem to require about 800 16-bit words to specify a particular visual scene, though this would also require a very large store of information about familiar objects. A few more words would be needed if any objects were moving or had an unusual colour, and also for such items as illumination conditions, but in spite of the crudity of this estimate it really seems that such a representation of a whole visual scene would only require about 1000 words. This is quite remarkable, for it is two orders of magnitude below the requirements of a single raw image of any quality. Of course the representation is incomplete and much information will have been lost, but surely the main economy must be achieved by exploiting the redundancy of the population of images that our eyes interpret for us.

This is an idea with an ancient ancestry [23,24] and ATTNEAVE [25] applied it to the psychology of vision 3 decades ago, while I have argued [26] that the physiological facts fit the idea. I think it makes good sense to interpret the early transformations in the visual pathway as measures to reduce redundancy, or attempts to find symbols that are used with high relative entropy and are adequate to represent the image [27]. Gain control, subtracting off the mean value, frequency filtering in the temporal and spatial domains, all fit this interpretation, as does segmentation into sub-regions with more uniform properties. There must be many more stages we do not yet understand, such as the formation of a catalogue of familiar objects together with an adequate means of referring to it: however I think we can see the direction the route is heading even though we have not gone far along it.

One important point that arises from this description of the overall goal of perception is that it gives a means to test each step to see whether progress towards that goal is actually being made; indeed this yard-stick may be what enables the self-regulating properties I mentioned earlier to be effectively utilized. Perhaps if you were to set yourselves the goal of representing images as economically as our brains seem to, and if you were to reach it, then you might achieve a result that would be of very wide interest and importance.

Since this talk was supposed to be about understanding natural vision I must conclude by saying what should by now be obvious:—We do not understand much of it. But we are optimistic that we shall go on finding out more because we have a working system always available to study and we have the physical image processors to tell us what to look for.

References

- 1 J.E. Dowling, *Investigative Ophthal.* 9, 655 (1970).
- 2 H.B. Barlow and J.D. Mollon (Eds), *The Senses* (Cambridge University Press, 1982).
- 3 W.R. Levick, *J. Physiol.* 188, 285 (1967).
- 4 V. Torre and T. Poggio, *Proc. Roy. Soc. B.* 202, 409 (1978).
- 5 M. Collonnier, "The electron-microscopic analysis of the neuronal organization of the cerebral cortex", in *The Organization of the Cerebral Cortex*, F.O. Schmitt, F.G. Worden, G. Adelman and S.G. Dennis, eds. (MIT Press, Cambridge, Massachusetts, 1981).
- 6 D.C. Van Essen and J.H.R. Maunsell, *J. Comp. Neurol.* 191, 255 (1980).
- 7 D.H. Hubel and T.N. Wiesel, *Proc. Roy. Soc. B.* 198, 1 (1977).
- 8 H.B. Barlow, *Nature* 279, 189 (1979).
- 9 C.D. Gilbert and T.N. Wiesel, "Laminar specialization and intracortical connections in cat primary visual cortex", in *The Organization of the Cerebral Cortex*, F.O. Schmitt, F.G. Worden, G. Adelman and S.G. Dennis eds (MIT Press, Cambridge, Massachusetts, 1981).
- 10 R.H. Wurtz and J.E. Albano, *Ann Rev. Neurosci.* 3, 189 (1980).
- 11 D.H. Hubel and T.N. Wiesel, *J. Physiol.* 160, 106 (1962).
- 12 J.G. Robson, "Neural images: The physiological basis of spatial vision", in *Visual Coding and Adaptability*, C.S. Harris, ed. (L. Erlbaum and Associates, Hillsborough, New Jersey, 1980).
- 13 B. Sakitt and H.B. Barlow, *Biol. Cybernetics* 43, 97 (1982).
- 14 C.R. Michael, *J. Neurophysiol.* 41, 572 (1978).
- 15 H.B. Barlow, C. Blakemore and J.D. Pettigrew, *J. Physiol.* 193, 327 (1967).
- 16 A. Guzman, *Proc. Fall Joint Computer Conf. (San Francisco)* Vol. 33, Pt 1, 291 (1968).
- 17 H.B. Barlow, *Proc. Roy. Soc. B.* 212, 1 (1981).

- 18 S.M. Zeki, *Nature* 274, 423 (1978).
- 19 D.H. Hubel and T.N. Wiesel, *J. Physiol.* 206, 419 (1970).
- 20 D.O. Hebb, *The organization of behaviour*, (Wiley, New York, 1949).
- 21 J.P. Rauschecker and W. Singer, *Nature* 280, 58 (1979).
- 22 D. Marr, *Vision*, (W.H. Freeman and Co., San Francisco, 1982).
- 23 E. Mach, *The analysis of sensations, and the relation of the physical to the psychical*, (Dover Publications Inc, New York, 1959; 1st edition 1886).
- 24 K. Pearson, *The Grammar of Science*, (Walter Scott, London, 1892).
- 25 F. Attneave, *Psychol. Rev.* 61, 183 (1954).
- 26 H.B. Barlow, "Possible principles underlying the transformations of sensory messages, in *Sensory Communication*, W. Rosenblith, ed. (MIT Press and John Wiley, New York, 1961).
- 27 H.B. Barlow, "Perception: What quantitative laws govern the acquisition of knowledge from the senses?", in *Functions of the Brain, Wolfson College Lectures 1982*, C. Coen, ed. (Oxford University Press, 1983).