

WHY HAVE MULTIPLE CORTICAL AREAS?

H. B. BARLOW

Physiological Laboratory, Downing Street, Cambridge CB2 3EG, England

Abstract—Image processing requires free access to information about all parts of an image, but a nerve cell in V1 can only interact directly with a tiny fraction of the other cells in V1. The problem this poses might be alleviated by forming secondary "neural images" in which information is re-arranged, and some possible rules of projection for forming such images are explored. It is also suggested that all parts of the cerebral cortex detect, and subsequently signal, suspicious coincidences in their inputs. Acquiring knowledge of the associative structure of sensory messages, in the form of the unexpected coincidences that occur, may represent the beginning of the formation of a working model, or cognitive map, of the environment.

Visual cortex Image-processing Computer-vision Segregation Mapping Flow-field Neural images

INTRODUCTION

Many of you may have noticed that this is a silly title, since it implies that the problem lies solely in the multiplicity of areas. Of course we do not really understand what any part of the cortex does, but David Van Essen's talk forces the question about multiplicity on us, and I first want to suggest an answer based upon a limitation of neurons as computational elements. But for this to make much sense one also needs ideas on what the whole cortex does and how its single neurons might do it, so I shall include a few brief paragraphs on the function of the cortex as a whole.

WHY NOT ONE AREA?

I think the answer is to be found by considering something that systems theory tends to draw one's attention away from. It is happiest dealing with linear quantities and well-behaved components, but as everyone knows the brain is composed of difficult materials—white and grey matter do not behave at all as systems' theorists would like. The way the brain sets about its tasks must be very much determined by the limitations of its components, as with any other complex system, so let us start out by making a comparison.

Table 1 lists some of the properties of a fairly modern computer and the brain. There is the well known difference between them in the slow, parallel plan of the brain's computations and

the rapid serial operation of most computers, but I want to draw attention to two other differences. First there is the accuracy, depending primarily on word length in the computer and the dynamic range of nerves in the brain. Here the computer can easily manage 8 decimal digits with its 32 bit words, but a nerve fibre can only transmit 2, perhaps only 1 or less, in a time of 1/10 sec or so. There is a remote but exciting chance that we have under-estimated the potentialities of nerves if, as Abeles (1982) and von der Malsberg (1981) suggest, the precise timing of impulses is important in the cortex, but even if they are right there is no way that I can see whereby nerves could approach the accuracy that is commonplace in computers.

I am not going to spend time on the consequences of the limited dynamic range of neurons, but I believe that it certainly enables one to understand, and might have enabled one to predict, much of the neurophysiology of the retina. For example any communications engineer could have suggested adaptive mechanisms, the use of differential inputs in lateral inhibition and colour opponency, and parallel pathways such as the "on" and "off" systems and X and Y channels, as ways of alleviating the problem posed by the limited dynamic range of nerve fibres. Perhaps one might gain a similar understanding of the cortex if one could put one's finger on the weakness of neurons as components for image processing, so let us go back to Table 1.

Table 1. Comparison between computer and brain

	Motorola VME/10	Human brain
Architecture	Serial	Parallel
Speed—Hz	10^7	30
Accuracy—digits	> 8	1 or 2
Address range	10^7 Words	10^4 Neurons
Number of elements	10^7	10^{11} Cells
		10^{14} Synapses
Weight—kg	25	2
Volume—m ³	0.2	0.002
Working life—Years	< 10	> 10
Maintenance	Contract	Self-repairing Self-adjusting
Programming	Needs it	Does it

THE PROBLEM OF LIMITED INTERCONNECTIONS

Look at the row which I have called "Address range". For the computer this is well over 10^7 words directly addressable without any fuss. It is not quite clear what the appropriate analogy for address range is in the brain, but the number of other nerve cells from which a given nerve cell receives messages, and the number to which it directly transmits messages, must be related. I'm not of course happy that this is a close analogy, but the range of connectivity of neurons is certainly important, and the generous figure of

10^4 that I have given here is three orders of magnitude below that for the computer. But it really needs to be read in conjunction with the next line of the table—the total number of cells. You will see that each neurone makes connections with only one ten millionth of all the cells in the brain. What this means is that access of information between one part of the brain and another is not by any means direct and simple. Both Allman and Kaas (1974) and Cowey (1979) have suggested that shortening the interconnections between cells is an important consideration in topographic maps and their relationships to each other. However I think the limited number of cells with which any given cell can connect adds a new dimension to this argument; it makes one realise the importance of facilitating the interconnection of those particular cells that, for functional reasons, must be able to interconnect. Furthermore the ability to find out about other parts of a picture is of special importance in image processing, so overcoming restrictions on interconnection is likely to be particularly important in the functioning of the visual cortex.

To illustrate this consider the classic problem dealt with by Guzmán (1968). His task was to program a computer to describe blocks from a




Fig. 1. A line drawing of a jumbled pile of solid blocks is quite easily sorted out by the human eye: regions 34 and 20, for instance, obviously belong to the same thin block, even though in human V1 they might be represented 5 cm or more apart. Guzmán (1978) wrote a computer program to do this and emphasized the importance of "linking principles"—in this case the orientation of edges—in facilitating the establishment of relations between different parts of the same block (from Guzmán, 1968).

line drawing of them such as Fig. 1, and assign them their 3-D positions. Notice first that one's own visual system knows immediately when the two halves of a block belong to each other, even when another block occludes a large part of it. For the computer the key to recognising that, for example, 20 and 34 in Fig. 1 are the same face of the same block was the fact that their edges are co-linear. The obscured part of these edges spans about 1/4 the width of the whole picture, so if the whole subtended 12° at your eye the obscured part would subtend 3° . The magnification factor in the human V1 is more than 15 mm/deg in the foveal region so the two parts of this line are represented nearly 5 cm apart! Considering that there are some 20 million cells per cm^2 in the cortex one can see that there is little chance that a cell representing one part of the edge will connect with a cell representing another part of that edge across the gap. Furthermore they are unlikely both to make a connection to a third cell unless there are specific provisions for this to happen.

SHORTENING THE SEARCH

Now for the computer the importance of appreciating that the edges of the face of a block are co-linear lines in the fact that it is the key enabling the computer memory to be searched for common parts of a block; that was why Guzmán named orientation a "linking feature", for it enables links to be found between one part of an object and another part elsewhere in the image. It also enables that search to be narrowed and thereby shortened, for instead of scanning the whole symbolic representation of the picture for the entry of an edge that is nearly co-linear, one might arrange for lists of edges to be ordered according to their orientation, and search that part which is nearly the same as the edge of interest.

Better than relying on orientation alone would be to use the "Hough transform", in which edges are represented in a 2-D matrix with orientation as one dimension and distance from an origin measured normal to that orientation as the other dimension; in such a representation co-linear parts coincide, so to match up two edges of the same face you only have to search a small part of the whole matrix. That is the sort of thing that might be possible for a neuron in spite of its miserable capacity to interconnect with other neurons. The principle is illustrated in Fig. 2 (top).

COMMON FATE

Segregation by motion reveals another capacity of the human visual system to recognise similarities over considerable distances in the visual field; a figure can suddenly be made to appear in a random array if a subset of marks delineating a cat or other figure starts to move coherently. Gestalt psychologists would, I think, have said that it is the "common fate" of the parts of the cat that enables one to segregate it so readily from the noisy background, but for us it is perhaps easier to regard direction of motion as a "linking feature" analogous to orientation and co-linearity in Guzmán's task. As in that case, its importance must surely be that it enables a directed search to be made for the representation of other parts of the same object, thus avoiding an exhaustive search through the whole representation of the image. As with the Hough transform, a 2-D representation of motion can be found that would greatly simplify the search for objects having the same direction and velocity of motion, and if motion in the image was re-mapped in this way one can imagine a single cell having access to information about the motion of all parts of an object, even when these occupied quite widely separate parts of the visual field. This motion map is illustrated in Fig. 2 (middle). If there were reciprocal connections from such a map back to a topographical map in V1 or V2 it might be possible for them to do the cortical equivalent of making the parts of the cat stand out by "highlighting" or "flashing". Such a map would also be very useful for the control of eye movements, for a concentration of activity at a particular point indicates the direction and velocity of the eye-movement required to reduce the retinal drift of the parts of the image that gave rise to those points, thereby enabling them to be seen better.

Similarly an object can be distinguished from its background by virtue of its colour, and Fig. 2 (bottom) shows how remapping according to colour might enable this to be done. It also suggests that crude information about topographical position might be retained within each colour region in a manner analogous to that whereby orientation is positionally mapped in V1.

These examples may have shown you that rearrangement of information by non-topographic mapping is potentially useful, and the possibility that it occurs should be taken

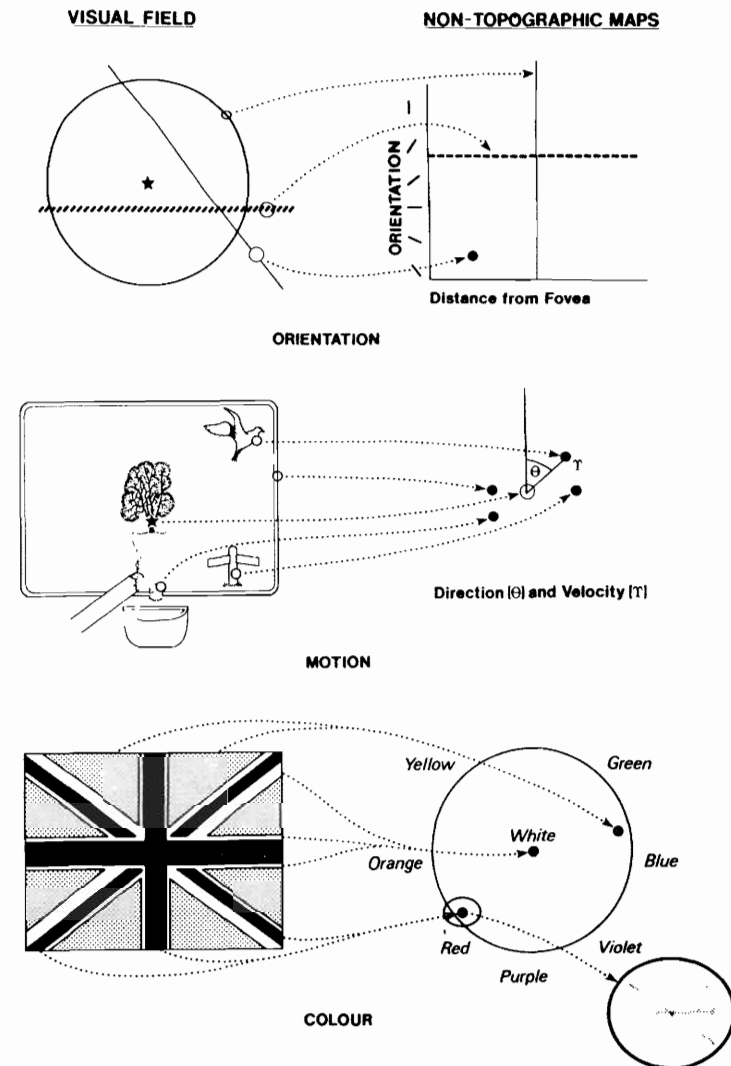


Fig. 2. Transformations, or non-topographic maps, of a visual image can bring close together items of information that are represented far apart in the original image or a topographic map of it. At the top the Hough transform is represented; here the ordinate gives the orientation of a line segment, the abscissa the shortest distance of a line of that orientation from the origin. Thus a straight line in the image maps to a point in the transform, a row of parallel, oblique, line segments maps to a horizontal row of dots, and a circle concentric with the fovea maps to a vertical line. The middle section shows a non-topographic map of motion. The observer fixates a tree seen through the window of a moving vehicle; this point is stationary on the retina and is mapped at the origin of the motion map. The frame of the window moves to the left relative to the tree, and is represented at a distance to the left of the origin proportional to the velocity of its image's movement. Similarly for other points. The bottom section shows how colours might be mapped on a colour circle, thus making it easy to establish links between regions of the same colour in the original image. The enlarged portion of the red region of the colour map suggests that crude topographical information might be retained within each small part of a non-topographic map: in V1 orientation is mapped onto position within each hypercolumn, and the cortex is therefore capable of representing more than two variables by position (from Barlow, 1981).




Fig. 3. The ray diagram shows how a simple rule for deviating a ray leads to the formation of an optical image. Neural images would result if the projection position of a nerve fibre was deviated according to its own pattern-selective properties (see Figs 4 and 5).

seriously. The fact that nerve fibres can topographically map one array of cells (such as the retinal ganglion cells) onto another array (such as LGN or cortex) proves that they are experts at finding their way to their appropriate termination sites; all that non-topographic mapping requires is that the termination site should depend upon the pattern selective property of a neuron as well as upon the position of its dendritic field. An analogy with the ray diagram that shows the formation of an optical image may make this clear. The central ray in Fig. 3 is undeviated, but other rays are deviated by an amount proportional to their eccentricity, and the consequence of this very simple "mapping rule" is the formation of an optical image. In the same way, if a nerve fibre can project, not to its

normal place in a topographic map, but to a position deviated according to its own pattern selective properties, then interesting "neural images" can be created. I shall consider two problems to show how modifications to straightforward topographical mapping can help to solve problems concerned with motion.

FOCUS OF EXPANSION

Since Gibson (1950), it has been recognized that a "focus of expansion" is a potentially useful point to determine in the flow field resulting from image movement; for instance it can indicate the expected point of impact when an observer moves in a straight line towards a plane surface. The heavy arrows in Fig. 4 show a set of velocity vectors, and it is easy to see that the way to find the focus of expansion is to extend them backwards until they meet. But it is not so easy to see how the brain can perform an operation like this, so it is worth exploring the possibility that it forms a "neural image" according to the following rules.

Each point in the primary image has neurons optimally sensitive to a wide variety of directions and velocities of motion, and for a straightforward topographic map all of these would project to the same position in the secondary map. Now suppose that each of them projects to a position deviated according to the

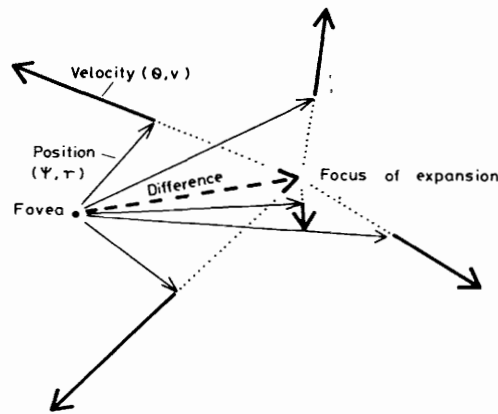


Fig. 4. Determination of the focus of expansion from a neural image. The thin vectors (ψ, r) indicate the positions of five objects relative to the fovea. The heavy vectors (θ, v) show the direction and velocity of movement of these objects, which are moving away from a single focus of expansion. Mathematically this can be determined from the difference vector $(\psi, r) - K(\theta, v)$, and this is shown (with appropriate selection of K) as a heavy dashed vector. Neurally this would be realised if direction and velocity selective neurons project, not to their exact topographic position in a secondary map, but to a position deviated according to their own direction and velocity selective properties.

following scheme. In Fig. 4 position in the visual field is represented by the thin vector (ψ, r) , and the motion at that point is represented by the thick vector (θ, v) . Now consider the new vector representing the difference between position and motion vectors, $(\psi, r) - K(\theta, v)$. The thick dashed line shows this new vector and it will be seen that, for this particular value of K , it is the same for all the expanding points, and that its termination is at the focus of expansion. Thus if there was a secondary projection in which each directionally selective neurone projected to a position modified according to the direction and velocity of movement to which it was itself most sensitive, there would tend to be a maximum of activity at the focus of expansion. The rule for modifying a neuron's projection position is implied by the construction of Fig. 4 and the expression for the difference vector $(\psi, r) - K(\theta, v)$.

It is interesting to look at this in slightly more detail, for a map of difference vectors has interesting properties: if everything is stationary, it will reduce to a topographical map, but otherwise the representation of moving objects is displaced towards or beyond the point from which they are diverging. If maps of the difference vector are made using different magnitudes of the constant K , there will be one in which all vectors coincide, as shown in Fig. 4, provided that there is a genuine focus of expansion. Not only can this indicate the expected point of impact with a plane, but the value of K required to make the difference vectors meet at a point is proportional to the time delay before the moment of impact (Hoyle, 1957). It seems to me that such combinations of motion and topographic maps have properties that are interesting and potentially useful, but above all one now sees ways that the brain might perform a new range of tasks: it is just a matter of each neuron modifying its propensity to map topographically in accordance with its own pattern selective properties.

THE APERTURE PROBLEM

Movement in an image is signalled mainly by the displacement of lines or edges, and the only information that can be extracted locally from such a displacement is the component of the motion that is perpendicular to the orientation of the edge or line. This problem has been elegantly analysed both experimentally by Movshon *et al.* (1984), and theoretically by

Hildreth (1984). To obtain the true motion of a solid body it is necessary to combine the information from two or more non-parallel edges. Clearly this problem must be solved before a true motion map such as that of Figs 2 or 4 could be formed, and a way that nerves might do this is shown in Fig. 5.

Start by mapping motion vectors (superimposed as in Fig. 2) from an area of the visual field large enough to contain most parts of a rigid moving object, but in this case consider that we are dealing with vectors that correspond to the component of motion perpendicular to the edges of the object. To make such a map, each motion selective cell would first project to the position in the motion map corresponding to the direction and velocity of a moving edge that stimulates it optimally. Now such an edge motion can result from a whole range of true motions, and the locus of this family of true motions is obtained by drawing a tangent through the position corresponding to the perpendicular motion [see stippled areas in Fig. 5(a)]. Thus if the terminal field of a particular cell extends along this tangent, when it is activated it will create activity at all positions in the map corresponding to true motions that would excite that cell optimally. When a rigid object with many edges moves (without rotation) through the visual field many neurons will be excited, but the terminal fields of all these neurons will intersect at one point, which corresponds to the true velocity and direction of the moving object. In practice such a map would be blurred and indistinct, for direction and velocity tuning are usually broad and imprecise, but the point of maximum activity might still be located by strong mutual inhibition between neurons that summate local activity over regions in the motion map.

There are other ways of solving the problem, of course. For instance the reciprocal of velocity, "slowness", might be mapped, as in Fig. 5(b). In such a map, the normal components of edge-motions of differing orientations but the same true velocity map to points on a straight line; one can imagine a "simple cell" in such a map with a dendritic field that summates activity along a line and thus responds selectively to a particular true motion.

The message so far is: "To detect similarities over more than a fraction of a degree in the visual field you must rearrange the information in the image so that similar events at separate points are represented close to each other." I




Fig. 5. Two neural images that could solve the aperture problem. At the *top* objects moving in the visual scene (left) are mapped as in Fig. 2 (middle), but in this case the component of motion perpendicular to edges is (more realistically) considered. Now suppose that the terminal fields of neurons are not confined to regions at the tips of the vectors corresponding to the motions to which they are most sensitive, but are extended along a tangent at that point, as shown by the stippled strips; for linear motions of the three edges of the triangle, these strips intersect at the position corresponding to the true motion of the triangle, and there will tend to be a focus of maximum activity at this point. The *lower* diagram shows a similar construction for "slowness" vectors inversely proportional to velocity. In the motion map the terminations of these vectors lie on a straight line tangential to a circle round the origin at the position of the true slowness vector of the moving triangle; a cell at this point with a linearly extended dendritic field would summate excitation from all the moving edges.

think the brain *must* use this principle in order to narrow the search for parts of the image that have some property in common, and it makes sense to suppose that reprojection to another area where the magnification factor is less, and where information about one modality is concentrated, might be a first step towards achieving this. Rearrangement according to non-topographic principles would also help a great deal, and it should be noted that this need not necessarily be done on the whole field simultaneously; it could for instance be done separately on large overlapping regions, and the secondary area would then retain a coarse topographical map. There are many problems, such as symmetry detection, facial recognition, or recognising similarities after perspective transformations, that require bringing together information over considerable regions of the visual field, and it is a challenging problem to invent maps which might concentrate the information required for their solution. Ballard (1984) has explored some of the computational tasks that can be facilitated by the use of parameter spaces, which are equivalent to non-topographic maps.

The kinds of problem that might be solved by forming non-topographic neural images are those that Gestalt psychologists drew attention to, and it is an exciting prospect that some ancient psychology may be on the point of linking up with modern neuro-anatomy. But we can't claim much progress unless we begin to see the general nature of the task that the cortex performs. We have suggested that the point of forming new neural images by reprojection is to enable similarities to be detected in distant parts of the image, but does this have any significance for the role of the cortex as a whole?

WHAT IS CORTEX FOR?

Van Essen and Maunsell (1980) showed that about 60% of the cortical area in the macaque was directly concerned with vision, but it is obvious that our knowledge of cortical physiology does not go 60% of the way towards explaining that part of a monkey's behaviour which is controlled by its cortex. It would be easier to approach the problem this poses if one knew what behaviour the cortex *does* control, but unfortunately experimental psychology gives one very little help on the general aspects of this problem.

Without trying to justify this here I shall

assume that the answer is: "Behaviour that requires accurate knowledge and understanding of an individual's normal environment" (see Barlow, 1985). Thus I am saying that the cortex forms and stores what Craik (1943) called a working model of the environment and would nowadays probably be called a cognitive map of it. The very large amount of information required for such a model or map, and its importance in the production of useful behaviour, has become very obvious in the design of artificially intelligent expert systems (Feigenbaum and McCorduck, 1983). Does our knowledge of the physiology of the cortex and its anatomical organization fit in with the notion that its role is to acquire and store this knowledge?

Maybe it does, better than might be supposed, but the first question is how the knowledge required for this model can be obtained. Some of it must be genetic in origin, for animals are well-adapted to their environments, often exquisitely well-adapted. But I think Craik regarded the knowledge that an individual acquires during its lifetime as more interesting and important than this innate store, and his followers might agree. Circus animals and humans acquire much from deliberate education, but this can only be old knowledge, and new knowledge must be gained by some process of induction or statistical inference. This may seem a lot to expect of cortical neurons, but I think a plausible case can be made that they are doing it.

New knowledge comes from the rejection of null hypotheses, and it actually makes quite good sense to regard a peripheral nerve fibre, such as one in the optic nerve, as signalling when it is unusually strongly stimulated; one might say that its own null hypothesis is that the sensory stimulus is not unusual in some particular way, and it signals when this hypothesis is violated (Barlow, 1969). But the range of hypotheses that could be represented by single afferent fibres is small and restricted compared with the range of hypotheses concerned with associations among events signalled in these afferents, and the important knowledge required for the working model or cognitive map must lie in the patterns of associations that occur in the sensory stimuli that an animal receives. It is the detection and signalling of these associations that is suggested as the special task of the cortex. This is not a new idea, for the earliest comparative anatomists thought the cortex must increase the range of associations an ani-

mal could detect, Hebb (1949) wrote much about identifying correlated inputs, and the role of the cortex in identifying covariation has been emphasized by MacKay (1978). But the form in which the notion is now advanced is much more specific (Phillips, Zeki and Barlow, 1984; Barlow, 1985).

Suppose that cortical neurons, during the "sensitive period" discovered by Hubel and Wiesel (1970), adjust their pattern selectivity so that they respond to the patterns of association in their inputs which occur most often; these patterns may be called "suspicious coincidences", for that is what they are: joint occurrences that happen more often than would be expected if there was no causal connection between them. After the sensitive period is over these neurons will continue to respond selectively to the "suspicious coincidence" to which they first became attuned, and will thus have incorporated knowledge of the causal connection in the environment that was responsible for the coincidence in the first place. It is not clear how far such a mechanism would go towards the acquisition of a cognitive map or working model, but at least it might be a beginning.

To test the plausibility of this idea it must be discussed in relation to the pattern selectivity of cortical neurons, their modifiability by experience, and the facts about cortical organization summarised by Van Essen.

PATTERN SELECTIVITY AND MODIFIABILITY

Cortical neurons in V1 and V2 are selectively sensitive to the orientation of bars and edges, their size or spatial frequency content, their binocular disparity, and their colour. It is generally agreed that rudiments of these forms of selectivity are present before an animal has acquired any experience of sensory stimuli in its normal environment, but that experience is required to develop this selectivity to the extent that is found in the adult, and even to retain the innate rudiments. Is it plausible to claim that this development occurs because of suspicious coincidences in the sensory input?

At a common sense level the answer must be yes: high-pass filtered images, such as are produced by lateral inhibition in the retina and lateral geniculate, possess plentiful edges—far more than would occur in a random, unstructured image. Similarly binocular excitation occurs frequently simply because the eyes point

almost in the same direction, so the same object will excite through both eyes. In both these cases the evidence that deprivation disrupts the development of the adult degree of selectivity is now good (Movshon and Van Sluyters, 1981). For spatial frequency selectivity the same is true, and indeed the inexperienced visual system does not appear to respond to high spatial frequencies nearly as well as the experienced adult, whether the inexperience is due to youth or deprivation (Derrington, 1979; Blakemore and Vital-Durand, 1984). Thus the adult visual cortex can be said to incorporate some knowledge of the usual structure of visual images, a kind of working model of them, in the forms of selectivity of its neurons.

MULTIPLE MAPS AND CORTICAL ORGANIZATION

Here we come back to the beginning, where it was suggested that multiple maps were necessitated by the limited capacity of neurones to interconnect as widely as is convenient for image-processing operations. The detection of similarities in different parts of an image that result from the presence of an object, which is the basis for segregation of figure from ground, is an example of detecting suspicious coincidences, presumably ones that have occurred before and thus represent part of the stored knowledge, working model, or cognitive map, of the environment.

The unified point of view that is sketched in this paper explains topographic mapping in the primary projection areas along associationist lines. Events occurring close together in the visual field are *a priori* likely to be associated, just as are events that occur close together in time. It is therefore desirable to represent them close together in the brain in order to bring them within the limited connectivity range of the cortical neurons. The special property of a topographic map that makes it necessary for this purpose is the fact that it preserves neighbourhood relations, which no other mapping does. For other areas, other non-topographic principles are likely to be involved in order to enable associations that do not depend upon simple proximity to be detected—for example associations in depth, motion, texture, or colour. It may be too much to expect topographic information to be entirely discarded in the early stages after V1, because in the real world one would not want to ignore proximity entirely when faced with evidence of similarities of

motion, depth, or colour; thus evidence of topographic order in secondary areas is not inconsistent with the idea, but in this paper I have, for obvious reasons, emphasised the new principles and what they may achieve.

As well as assigning to topographic maps a function that used to be attributed exclusively to so-called association areas, this point of view also allows plenty of room for developmental "hard-wiring" of cortical connections. The task of determining the associational structure of the sensory messages that impinge upon the cortex is so vast that there is plenty of need for ontogenetic guidance of connectivity in order to preselect the types of coincidence that can be detected. But of course the test of these ideas is not whether they provide a satisfying framework for the facts we already know, but whether they lead to the right experimental questions to ask in the future.

REFERENCES

- Abeles M. (1982) *Local Cortical Circuits*. Springer, Berlin.
- Allman J. M. and Kaas J. H. (1974) The organisation of the second visual area (V II) in the owl monkey. A second order transformation in the visual hemifield. *Brain Res.* **76**, 247–265.
- Ballard D. H. (1984) Parameter networks. *Artificial Intell.* **22**, 235–267.
- Barlow H. B. (1969) Pattern recognition and the responses of sensory neurons. *Ann. N.Y. Acad. Sci.* **156**, 872–881.
- Barlow H. B. (1981) The Ferrier Lecture: Critical limiting factors in the design of the eye and visual cortex. *Proc. Roy. Soc., Lond.* **B212**, 1–34.
- Barlow H. B. (1985) Cerebral cortex as model builder. In *Models of the Visual Cortex* (Edited by Rose D. and Dobson V.). Wiley, New York.
- Blakemore C. and Vital-Durand F. (1984) Development of the monkey's geniculocortical system. In *Development of the Visual Pathways in Mammals* (Edited by Stone J., Dreher B. and Rapaport D. H.). Liss, New York.
- Cowey A. (1979) Cortical maps and visual perception: the Grindley memorial lecture. *Q. J. exp. Psychol.* **31**, 1–17.
- Craik K. J. W. (1943) *The Nature of Explanation*. Cambridge Univ. Press.
- Derrington A. (1979) Effect of visual deprivation on the development of spatial frequency selectivity in kitten striate cortex. *J. Physiol.* **300**, 62P.
- Feigenbaum and McCorduck P. (1984) *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*. Addison-Wesley, Reading, Mass. (1983) and Pan Books, London.
- Gibson J. J. (1950) *The Perception of the Visual World*. Houghton Mifflin, Boston.
- Guzmann A. (1968) Decomposition of a visual scene into three-dimensional bodies. *Proc. Fall Joint Computer Conf.*, San Francisco, 1968, Vol. 33, Part I, pp. 291–304.
- Hebb D. O. (1949) *The Organisation of Behaviour*. Wiley, New York.
- Hildreth E. C. (1984) The computation of the velocity field. *Proc. R. Soc. Lond.* **B 221**, 189–220.
- Hoyle F. (1957) *The Black Cloud*. Harper & Row.
- Hubel D. H. and Wiesel T. N. (1970) The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *J. Physiol.* **206**, 419–436.
- Malsburg C. von der (1981) Internal report 81-2, Department of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Goettingen, F.R.G.
- Movshon J. A. and Van Sluyters R. C. (1981) Visual neural development. *Ann. Rev. Psychol.* **32**, 477–522.
- Movshon J. A., Adelson E. H., Gizzi M. and Newsome W. T. (1984) The analysis of moving visual patterns. *Expl. Brain Res.*
- Phillips C. G., Zeki S. and Barlow H. B. (1984) Localization of function in the cerebral cortex: Past, Present and Future. *Brain* **107**, 327–361.
- Van Essen D. C. and Maunsell J. H. (1980) Two-dimensional maps of the cerebral cortex. *J. comp. Neurol.* **191**, 255–281.