

Guest editorial

Cerebral predictions

Occasionally, as you go into a restaurant, you pick up a message that says "Wow, this is going to be good", and as you open this number of *Perception* you may get a whiff of the tasty dish of Cortical Neurophysiology Peter Lennie (1998) has cooked for you. This note is intended as an appetising cocktail to prepare you for it, but first a word about the dish itself.

It's a review, and reviews fill me with apprehension because they are so often no more than a boring list of papers you feel you ought to have read yourself. But this one is different, for it has new ideas and tells us to change the way we think about the visual cortex. It is not, Lennie says, composed of parallel pathways and specialised areas providing information to be recombined into our unified perception of the world: instead, each area retains its topographical arrangement of the visual field, makes nearly autonomous decisions, and what we perceive consists of these decisions experienced in parallel. Furthermore, the computational elements that do the work, the cortical neurons, should not be regarded as having specialised selectivities for particular image qualities like colour, movement, depth, or texture: instead they should be regarded as vectors pointing to regions in a space whose many dimensions are these image qualities. In his final words he says "Single-unit recording on its own is a weak instrument for discovering what cortex really does, but when harnessed to a theory it is immensely valuable". I think that's almost exactly right, but we need to try harnessing experimental observations to many

theories, not a single one: we fail to make optimum use of our experimental observations because we don't have enough theories to choose from, and I'll now try to put that right by advancing more of them.

What does cortex do?

It's 70 years since Judson Herrick (1928) said that the cerebral cortex finds and stores knowledge of the world, that it forms the "filing cabinets of the central executive". That judgement may not have been based on much exact experimental or observational evidence, but it has stuck, and I think it is widely accepted as a general explanation for why animals with a large cortex behave differently from those with a small one, and generally have a selective advantage over them (see for example Jerison 1991). But why is knowledge of the world so useful? It's not because it earns us degrees or enables us to write books, but for a much more down-to-earth reason.

It enables us to predict: nothing is more useful than knowing how things will pan out, and this is true in an amazing variety of circumstances over time scales from milliseconds to months and years. It's true of the tennis player who has to judge, from the opponent's movements, past habits, and perhaps the first few milliseconds of the ball's flight path, exactly where the return shot is going, to the manufacturer planning what product to make and how to market it. It's not just true in urban and suburban jungles, but also in real ones, for escaping a predator, planning a foraging trip, or handling social relations in a group of monkeys depend upon correctly foretelling aggressive movements, fruiting habits of trees, or behavioural peculiarities of other monkeys. All this requires extensive knowledge of the world and its ways.

Prediction is a matter of identifying a spatiotemporal pattern at an early stage, and assuming that it will run to completion. You need good knowledge of the commonly occurring spatiotemporal

patterns in the environment, you need to have classified them, you need the ability to discriminate between them, and you need skill in correctly identifying at the earliest possible moment when one of them has started. It could not be done at all if the environment we live in did not have a great deal of statistical structure and regularity - that is, if it did not have spatiotemporal patterns waiting to be classified and identified by suitable neural mechanisms. Doing all this well must place heavy demands on the computational machinery that does it, and it is worth looking at the cortex from the point of view that it is specialised for this purpose.

Visual cortex seems to do otherwise

We naturally look first at the visual cortex, and, from a naive first look, what has been found by recording from its neurons does not seem to lend much support to Herrick's view of what the cortex does, let alone the suggestion that it is involved in prediction. If the cortical cells form the filing cabinets of the central executive, they should be empty in the inexperienced cortex, whereas there are at least a few neurons in the totally inexperienced primary visual cortex that seem to have properties remarkably similar to those of the adult (see Hubel and Wiesel 1970; Freeman and Ohzawa 1992). And if cortex discovers the regularities in the statistical structure of the incoming messages, then the properties of the inexperienced neurons should be plastic and adaptable to the characteristics of the incoming messages--but this is not the picture that emerges from studies of the sensitive period. Instead, it seems that neurons start with strong prejudices about what they should respond to selectively, and, if they are stimulated by their favoured stimuli, then this pattern of responsiveness is maintained, and perhaps they become a little more responsive and selective; but the striking change occurs when they are not so stimulated, for then they sulk and become unresponsive to all stimuli. After the sensitive period, changed experience does not seem to have any permanent effect on the response properties of the neurons.

The initial impression is thus hostile to Herrick's view, but on the other side are the results of recent analyses linking the statistical properties of natural images to the properties of simple cells in V1 (Bell and Sejnowski 1997; Olshausen and Field 1997; van Hateren and van der Schaaf 1998). These results support the idea that V1 cells are adapted to the statistics and form of a set of filters that perform an independent components analysis (ICA) of the images. There are still problems with this so I shall not go into it in more detail, but if it is right then the properties of these cells do reflect the statistics of the world, and therefore do constitute knowledge of it. It is true that this is not explicit knowledge as laid out in a textbook, but it is in an even more useful form, for the receptive fields instantly apply the knowledge to produce a representation in the form dictated by ICA. Why is there this big discrepancy between this latest analysis and the disappointing earlier impression from single-unit recording?

Reconciling the two

I think there are three factors. The first is that, as hard-nosed neurophysiologists, we have tended to ignore or not take seriously enough what the comparative anatomists have been telling us for years about the role of cerebral cortex in behaviour, and we have therefore not even tried to reconcile them. The visual cortex is of course different from the other parts of the cortex, but I refuse to believe that it has no lessons for us in trying to understand what the rest of it does, and the reverse must also be true!

The second reason for being slow in reconciling the two results concerns the weakness of positive instructional effects of experience on the properties of V1 neurons. That does not at first seem to fit. Herrick's idea of learning as a positive process, but in fact much learning is of a negative kind: learning the ways of the world consists of learning what does NOT happen as much or more than learning what does. Consider learning to spell: the facts here are that if you take five-letter words, there are 265 or nearly

12 million possible combinations, but there are less than 50000 actual five-letter words in the English language. We tend to assume that the brain must take the easy course of learning the correct spellings, but it must also learn to reject the much larger number of incorrect ones, and the facts of the sensitive period rather suggest that learning to reject is the crucial task. There is more to it than this of course, but let us go on to the third way our initial expectations may have been incorrectly biased.

If the cortex has acquired enough knowledge of the world to predict well, we tend to think it must have done so through its own individual experience, but its evolutionary history should warn us that this may not be the whole story. Over the past million years or so the human brain has more than doubled in relative size. As far as we know, a striking change like this can only occur through the positive selection of those heritable variations that make some brains larger than others. Of course, the positive selection will not have been for overall brain size as such, but for something like a change in a receptor molecule, a change in connectivity instructions, or the addition or duplication of a brain area, that leads to a behavioural advantage and also entails an increase in brain size. We know none of the details, and the increase in relative brain size is only an indication that inherited changes have occurred, but isn't it obvious that our brains differ from those of chimpanzees, that these differences are genetic, and that they have been selected for because of their behavioural effects?

Applying this to our problem in reconciling the neurophysiology of V1 with Herrick's idea of the overall role of the cortex, we see that it was probably wrong to expect his 'knowledge of the world' to be knowledge acquired through individual experience: much is likely to be knowledge acquired through evolutionary experience. I think this gives a new insight into the whole problem of the neocortex and its evolutionary explosion: it has the best of both worlds, for it learns both ways. I know that this view of the cortex will cause

pain to those who favour a strongly nurturist view of cognition, but Lennie's review should be compulsory reading for those who have put their name to the claim that there is no compelling evidence for innate representations in the cortex" (Elman et al 1996).

More theory

One of the points that rightly concerns Lennie is the vast increase in the number of neurons devoted to the representation of the visual image in the cortex. At first this seems to deny strongly the suggestion that sensory coding reduces the redundancy present in the physical messages transduced by our sense organs, producing an economically encoded representation, for the employment of all those neurons certainly does not look like an exercise in economy. But hang on a minute: Do we know what constitutes an economical representation for the brain? We tend to assume it is the same as an economical representation for British Telecom or AT&T, but this is likely to be wrong for two reasons. The first is that the brain has to do more than simply transmit the representation elsewhere; as Lennie says, it has to make decisions about it, or rather about what it represents in the outer world, and decision-making is a statistical operation that requires counting the numbers of occurrences. Highly compressed representations are very unsuitable for counting, as you would find out if you tried to estimate the frequencies of occurrence of the letters in a passage of prose from the frequencies of the bits in its ASCII representation. This is a problem Gardner--Medwin and I are pursuing elsewhere, but there is another aspect of brain-efficient representation.

One task the visual system does very well is to detect a large variety of useful signals at signal-to-noise ratios that computer vision systems cannot emulate - not yet at any rate - and objective evidence for this is given by the statistical efficiencies with which such tasks are performed. It is therefore enlightening to ask what kind of representations would make such efficient detection possible for a wide variety of tasks. The answer is: Ideally you

need a matched filter for each target you want to detect. You do pretty well with combinations, for example detecting a sine wave of any phase by having a pair of detectors in quadrature, but if each filter adds a bit of noise of its own, then you will always do better by having a filter matched to the particular pattern you want to detect.

Of course, you cannot have a separate detector for every possible pattern, but the desirability of low signal-to-noise ratios is an argument for having as large a number of matched detectors as you can, and if you examine in detail a task like detecting coherent motion, which MT is good at, its neurons begin to look like an enormous array of matched filters, each optimally adapted to detect a patch of coherent motion with particular parameters of position, size, direction of motion, velocity of motion, and depth (Barlow and Tripathy 1997). It's clearly sensible to select detectors for those patterns that occur frequently among its inputs (ie suspicious coincidences), and you also need some mechanism for ensuring that the outputs at a given level do not relay the same inputs but are as different from each other as possible. This is working round to saying that the different outputs classify the different spatiotemporal patterns that occur and signal when each one occurs, which is what you need for prediction. Perhaps we are reaching a job description for visual cortex that is not too different from what Herrick proposed as a general job description for all the cortex; there are many spatiotemporal sequences around, and it would be highly advantageous in a wide variety of circumstances to know about them and recognise at the earliest moment when one of them is about to be completed.

This concludes my cocktail of theory: now enjoy Pete Lennie's meal, and let us test the facts he summarises against as many sensible theories as possible. Sooner or later we'll find the right one.

Horace Barlow

Physiological Laboratory, University of Cambridge, Downing St,
Cambridge CB23EG, UK; e-mail: hbb10@cam.ac.uk

References

Barlow H B, Tripathy S P, 1997 ``Correspondence noise and signal pooling as factors determining the detectability of coherent visual motion" *Journal of Neuroscience* **17** 7954--7966

Bell A J, Sejnowski T J, 1997 ``The independent components of natural scenes are edge filters" *Vision Research* **37** 3327--3338

Elman J L, Bates E A, Johnson M H, Karmeloff-Smith A, Parisi D, Plunkett K, 1996 *Rethinking Innateness: A Connectionist Perspective on Development* (Cambridge, MA: MIT Press)

Freeman R D, Ohzawa I, 1992 ``Development of binocular vision in the kitten's striate cortex" *Journal of Neuroscience* **12** 4721--4736

Hateren F H van, Schaaf A van der, 1998 ``Independent component filters of natural images compared with simple cells in primary visual cortex" *Proceedings of the Royal Society of London, Series B* **265** 359--366

Herrick C J, 1928 *An Introduction to Neurology* (Philadelphia, PA: W B Saunders)

Hubel D H, Wiesel T N, 1970 ``The period of susceptibility to the physiological effects of unilateral eye closure in kittens" *Journal of Physiology (London)* **206** 419--436

Jerison H J, 1991 "Brain size and the evolution of mind" (The fifty-ninth James Arthur Lecture on the Evolution of the Human Mind, 1989) (New York: American Museum of Natural History)

Lennie P, 1998 "Single units and cortical organization" *Perception* **27** 889--935 (this issue)

Olshausen B A, Field D J, 1997 "Sparse coding with an overcomplete basis set: a strategy employed by V1?" *Vision Research* **37** 3311--3325