Brain Theory - Biological Basis and Computational Principles
A. Aertsen and V. Braitenberg (Editors)

Cell assemblies versus single cells

Horace Barlow

Physiological Laboratory, University of Cambridge,
Cambridge CB2 3EG, England

In the report of the 1990 Brain Theory meeting I gave several reasons for the superiority of single cell descriptions of sensory function over cell assembly descriptions (Barlow 1992), and I have recently given a more complete argument (Barlow 1994) for believing that perception is best explained at the single neuron level. It is not of course claimed that descriptions at this level are appropriate for other problems in neuroscience, such as the actions of drugs, or explaining social behaviour; these require one to consider membrane receptor molecules in the one case, or how individual humans interact in the other, and the behaviour of single neurons would not be illuminating for either. I shall not repeat all this, but want to make three points. The first is to rebut an argument one frequently hears *against* single neuron explanations because it fails to understand what is claimed for single neurons and ignores the behavioural facts that their properties can account for. The second is to attack some of the arguments *for* cell assemblies that seem to me unfounded. Finally I shall list some evidence for assemblies that I am beginning to find convincing.

1 MISLEADING TESTS OF THE SINGLE NEURON HYPOTHESIS

An example of the misleading argument says that cells in V1 are only crudely tuned to orientation, so that the response of a single such cell would only change when the orientation of a stimulus line was changed by several degrees. It is then claimed that orientation changes well under one degree can be discriminated psychophysically; therefore, it is said, psychophysical orientation discrimination must be based on more than a single cell. This is wrong because the wrong psychophysical results have been selected for the comparison. As you increase the length of a line, orientation discrimination improves, and the high discrimination ability quoted above is only obtained for lines many times as long as the extent of the receptive field of the V1 cells (see Andrews 1967). One cannot expect a cortical neuron to be influenced by the parts of the line outside its receptive field, so the appropriate comparison to make for a cortical neuron with a receptive field, say, 9 minutes long is the psychophysical orientation discrimination for 9 minute long lines, and this is only a small fraction of the figure for long

lines. To make a fair comparison between psychophysical performance and that of single neurons one must use a task for which the neuron is adapted, and compare it with psychophysical performance at the identical task.

It has been clear for a long time that known classes of neurons (i.e. ones whose activities are readily isolated) do not explain all aspects of psychophysical performance, nor would one expect them to. For instance single retinal ganglion cells in the cat appear to be sensitive enough to account for the ability of the intact animal to detect small, brief, flashes of light containing only a few quanta (Barlow, Levick and Yoon 1971), but no single retinal ganglion cell could account for the threshold for large, long duration, stimuli; it is quite obvious that no retinal ganglion cell can integrate information over very large extents of the visual field, but such integration is necessary to detect the small flux of quanta that the intact animal can detect; presumably there are cells at more central locations in the visual pathway that integrate information from many retinal ganglion cells, but we do not know where they are so we cannot confirm or refute this. Likewise there are presumably cells that combine information over much more extended regions than the typical V1 neuron, and which are thereby able to discriminate orientations of a fraction of a degree, but again we do not know where they are so we cannot confirm or refute this speculation either. But we do know that the intact animal can make such discriminations, and the only mechanism we know of that could improve the signal noise ratio sufficiently is integration by more centrally located neurons over larger areas, and perhaps longer times. This is the reason for suggesting that "Whenever two stimuli can be distinguished reliably, then some analysis of the physiological messages they cause *in some single neuron* would enable them to be distinguished with equal or greater reliability" (Barlow 1994).

Comparisons between behavioural and single unit performance have been made for tactile discriminations in monkey cortex (Talbot et al 1968), wavelength discrimination in monkey LGN (Devalois, Abramov and Mead 1967), sensitivity to light in retinal ganglion cells of the cat (Barlow, Levick and Yoon 1967), responses of single tactile fibres from the hand in humans (Vallbo 1989), spatial resolution (Parker and Hawken 1985) and contrast discrimination (Barlow, Kaushal, Hawken and Parker 1987) of cortical neurons in monkey, ability to distinguish coherent motion of crossed gratings in monkey V5 (MT) (Movshon, Adelson, Gizzi & Newsome 1985), the ability to detect coherent motion of random dots by neurons of monkey MT (Newsome, Britten & Movshon 1989), and many others that I do not know about or cannot now recall.

There are still problems in explaining psychophysical performance: 1) Often we do not know where to find the cells that seem to be required for particular tasks. 2) We do not know what is the best measure of the response of a single neuron to use - peak impulses per second, the number in some time interval, or perhaps a measure involving synchrony or rhythmicity. 3) The dynamic range of single neurons often seems inadequate to account for psychophysical performance. 4) The general principle that emerges is that psychophysical sensitivity for a particular task parallels that of the most sensitive of the neurons, but how does the brain succeed in ignoring the noise contributed by the other neurons that are less sensitive?

In spite of these difficulties single neurons are way ahead of cell assemblies when it comes to accounting for psychophysical behaviour, and I hope no-one will continue to be misled by foolish comparisons between psychophysical performance at one task and that of single neurons at another.

2 IMAGINARY VIRTUES OF CELL ASSEMBLIES

Hebb (1949) attached great importance to cell assemblies in his attempt to develop a conceptual system " . . . which relates the individual nerve cell to psychological phenomena", and it is an idea that has had enormous appeal to others. I understand this appeal myself, for neurons come in such vast numbers, possess such vast variety, and behave with such vast irregularity, that at first they do not seem to offer a promising basis for explaining behaviour. But on the other hand some of the virtues of cell assemblies are imaginary.

2.1 Time scale

Hebb introduced reverberating cell assemblies as the means of extending the time scale of nerve excitation to match that of the psychological mechanisms involved in learning and perception, but we now know that the supposed short time scale of nerve excitation processes resulted from ignorance; in Hebb's time it was not even universally accepted that synaptic transmission was chemical, or that it could be inhibitory as well as excitatory. We now realize that there are many intracellular processes that occur on quite a slow enough time scale to match the psychological processes that Hebb was concerned with, for instance changes of intracellular Ca^{++} levels, the actions of neuromodulators and second messengers, the restitution of ionic levels by pumps, and probably the slow diffusion and removal of external transmitters such as NO. The fact that cell assemblies are unnecessary for the main purpose for which Hebb introduced them does not of course prove that they do not exist, but it should make one cautious about them.

2.2 Reducing errors

Another imaginary virtue of cell assemblies is that they are less error-prone than single cells. Whether this is true or not depends upon the conditions that have to be met for the assembly to work correctly. If, for instance, each cell in the assembly has to function correctly for the whole assembly to function correctly, then it is easy to see that the assembly misfunctions when <u>any</u> of its component cells misfunctions, so the overall error rate is greater than that of the individual cells. Single cell representations can in fact be robust: if, for instance, you wish to represent a complex event or concept in a way that would resist the death or malfunction of individual neurons, then quite a good way of doing so would be to represent that event or concept by a single cell, and then make a few spare copies of that cell; the event or concept would then survive until all copies were lost.

2.3 Minimizing effects of cell death etc

Actually the argument from cell death is almost always bogus; it depends upon quoting the daily number of neurons that die, without pointing out that this is a minute fraction of the total number of cells. Similarly arguments from the resistance of the brain to mutilation are also weak because a) careful tests often reveal effects of mutilation that superficial examination had missed, b) from failure to realize that cell assemblies and distributed representations (like computers) are not inherently resistant to damage, c) because a limited amount of reduplication of neurons will minimize the effects of injury.

2.4 Exploiting combinations

The fact that the number of combinations of cells is much greater than the number of cells is another imaginary virtue of cell assemblies. The idea is that if important external objects or events (a cow or a catastrophe) are represented by the joint activity of n of the N neurons, rather than by a single one, then it is possible to represent a much larger number of such events. But no-one has ever suggested that we use a mutually exclusive representation, in which only a single cell is active, and in any plausible distributed representation each of the individual cells corresponds to some definable set of objects or events in the field of things represented. Ideas differ about the *sparseness* of the representation (what proportion of the cells are typically active), and about the nature of the subsets of objects or events that cause a single cell to be active. At one extreme people believe that a high proportion of cells are active and that this subset is arbitrary and meaningless in the way that the bits of the ASCII code define almost arbitrary subsets of keyboard characters; at the other it is thought that the subsets are usually themselves meaningful and are selected to make it possible to represent naturally occurring events sparsely – the idea of *cardinal* cells (Barlow 1972, 1994). One of the merits of representing meaningful and useful objects and events by single neurons is that it improves the efficiency of learning.

2.5 Learning in distributed representations

One of the important things one has to do with a representation is to learn associations with particular classes of the things represented: one learns that certain sounds signify dinner, or that certain sights signify danger. Now learning such associations is essentially a matter of determining that the entries in a 2x2 contingency table are non-random, and any mechanism that can do this well requires access to all four entries. This is the main reason for believing that learning must occur at Hebbian synapses, for one cannot easily point to any other position in the brain which has available the information required to determine whether a pre- and post-synaptic neuron fire in an associated way.

Now consider the task of learning an association with a class of objects or events represented by a *combination* of active cells in a cell assembly. To find out whether the class is present or not one must check that all the relevant cells in the assembly fire, except in the trivial case when *all* the cells fire *only* for that class (i.e. for non-overlapping assemblies, which are essentially reduplicated single cell representations). Now there is no single location where all the necessary information is available; one might form associations with each cell of the

assembly separately, and in some cases this might work reasonably well, but there are obvious problems in doing this. Imagine trying to form an association with a keyboard character by measuring the associations with each of the bits that represent it; random errors would arise from the fact that each bit is active for other letters than the one of interest, and there is also the possibility of consistent errors. Of course you could resolve the character from the bits, but this would mean creating a single neuron representation for the character, which is just what cell assemblies were supposed to avoid. If you want to form an association with a class of events you must surely find a location where the occurrence of that class is signalled; nerve cells have the characteristics required to do this while I know of nothing else in the brain that has.

3 BUT ALL THE SAME . . .

Cell assemblies do have some real virtues, and quite recently anatomical and physiological facts have emerged that I agree suggest that they are important in the cerebral cortex.

3.1 Intermediate level descriptions

The first point that carries weight with me is not actually new. It is obviously desirable to find a level of description intermediate between the whole brain and the single neuron, and I am impressed by the success of invertebrate neurophysiologists in accounting for the way small assemblies of cells generate rhythms and complex sequences of movements (see for instance Getting 1989). There must be similar cooperative interactions involved in generating directional selectivity and other forms of pattern selectivity in sensory pathways. Understanding how a group of half-a-dozen, or perhaps several hundred, cells cooperate to perform some task would obviously constitute an enormous step on the path of explaining how the 10^{10} cells of the cerebral cortex control behaviour.

3.2 Connectivity of cortical neurons

Second, the facts elucidated by Braitenberg and Schüz (1991), Douglas and Martin (1991), and others about the connectivity of cortical neurons seem to point to assembly-like bahaviour. What makes me hesitate here is that these anatomical and physiological facts do not provide an intuitively satisfactory account of the known physiological properties of cortical neurons. How is the receptive field of a cortical neuron formed? If the firing of a neuron depends upon synchronous input from a group of other neurons, should not these other neurons all have identical receptive fields? Why then are the receptive fields of cortical neurons always found to be different from each other? I shall be happier when the anatomy and physiology of cortical neurons fit together better.

3.3 Well-timed impulses

Third, well-timed impulses (Legéndy & Salcman 1985; Abeles et al 1993; Abeles, this volume) seem to me difficult, but perhaps not impossible, to account

for *except* by something like synfire chains. Nerve impulses can recur at astonishingly regular intervals in some preparations, it's just that this behaviour does not seem typical of single cortical neurons. And although single-cell events might recur with a timing accuracy of a few percent up to intervals of several hundred milliseconds, I don't think one would expect the absolute accuracy to be independent of interval. In one sense their discovery fits my arguments for the importance of single cells very well, for if the well-timed behaviour is widespread, neurons are not as noisy and irregularly behaved as they seem to be when one only attends to one cell at a time. The existence of well-timed impulses is certainly an important new fact about the cortex.

For these three reasons I'm much more impressed by cell assemblies now than I was even a few years ago, and I'll be watching them with interest.

REFERENCES

Abeles, M., Prut, Y., Bergman, H., Vaadia, E. & Aertsen, A. (1993). Integration, synchronicity and Periodicity. In A. Aertsen & W. v. Seelen (Eds.), *Spatio-temporal aspects of brain function*, Elsevier Science Publications.

Andrews, D.P. (1967) Perception of contour orientation in the central fovea Part II: spatial integration. *Vision Research*, **7** 999-1013.

Barlow, H.B. (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, **1** 371-394.

Barlow, H. B. (1992). Single cells versus Neuronal Assemblies, pp 169-173 in *Information Processing in the Cortex*. Eds A Aertsen and V Braitenberg. Berlin: Springer-Verlag.

Barlow, H.B. (1994). The Neuron Doctrine in Perception. In M. Gazzaniga. (Eds.), *The Cognitive Neurosciences*. Cambridge, Mass: MIT Press.

Barlow, H.B., Levick, W.R. & Yoon, M. (1971) Responses to single quanta of light in retinal ganglion cells of the cat. *Vision Research*, **11** (supplement No 3) 87-101.

Barlow, H.B., Kaushal, T.P., Hawken, M. & Parker, A.J. (1987) Human contrast discrimination and the contrast discrimination of cortical neurons. *Journal of the Optical Society of America*, **A 4** 2366-2371.

Braitenberg, V. & Schüz, A. (1991). *Anatomy of the Cortex: statistics and Geometry*. Berlin: Springer-Verlag.

Devalois, R.L., Abramov, I. & Mead, W.R. (1967) Single cell analysis of wavelength discrimination at the lateral geniculate nucleus in the macaque. *Journal of Neurophysiology*, **30** 415-433.

Douglas, R.J. & Martin, K.A.C. (1991) A functional microcircuit for cat visual cortex. *Journal of Physiology, London*, **440** 735-769.

Getting, P.A. (1989) Emerging principles governing the operation of neural networks. *Annual Review of Neuroscience*, **12** 185-204.

Hebb, D.O. (1949). *The Organization of Behaviour*. New York.: Wiley.

Legéndy, C.R. & Salcman, M. (1985) Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *Journal of Neurophysiology*, **53** 926-939.

Movshon, J.A., Adelson, E.H., Gizzi, M.S. & Newsome, W.T. (1985). The analysis of moving visual patterns. In C. Chagas, R. Gattass & C. Gross (Eds.), *Pattern Recognition Mechanisms* pp. 117-151. Rome: Vatican Press. (Reprinted in *Experimental Brain Research, Supplementum* **11** 117-151, 1986).

Newsome, W.T., Britten, K.H. & Movshon, J.A. (1989) Neuronal correlates of a perceptual decision. *Nature*, **341** 52-54.

Parker, A. & Hawken, M. (1985) The capabilities of monkey cortical cells in spatial resolution tasks. *Journal of the Optical Society of America*, **A 2** 1101-1114.

Talbot, W.P., Darian-Smith, I., Kornhuber, H.H. & Mountcastle, V.B. (1968) The sense of flutter -vibration: Comparison of human capacity with response patterns of mechano-receptive afferents from the monkey hand. *Journal of Neurophysiology*, **31** 301-334.

Vallbo, Å.B. (1989). Single fibre microneurography and sensation. In C. Kennard & M. Swash (Eds.), *Hierarchies in Neurology: a reappraisal of a Jacksonian concept* (pp. 93-109). London: Springer.