

Computational Modelling

Peter Dayan

Department of Computer Science, University of Toronto

Computational modelling is playing an increasingly accepted and important role in neuroscience. It is not a unitary enterprise, though, and the distinction between two different sorts of modelling, one interested in description and the other also in function, is illustrated in their application to activity-dependent developmental plasticity and adult conditioning.

1 Introduction

There are two distinct sorts of computational modelling that are commonplace in neuroscience and which are often confused. One is the natural offspring of the coupling of mathematical modelling with technology and concentrates on descriptions of phenomena. Computers permit calculation to be used to probe the quantitative consequences of models that are too complicated to analyse, the complications often coming in virtue of their being quite tightly coupled to experimental data. The classic examples of this are the detailed simulations of the flow of current in single cells with complex geometries, including versions of channel kinetics and the like [1]. The examples we discuss here are algorithms for vertebrate activity-dependent neural development which are based on versions of the Hebb rule for synaptic change. Operation of such rules can convert structure in input activity (*eg* nearby mechanoreceptors having closely related or correlated firing patterns) into structure in the organisation of parts of cortex (*eg* neighbouring cells in somatosensory cortex responding to touch in neighbouring parts of the body). The mathematical models address characteristics of this conversion in a quantitative manner. We will call this form of computational modelling *mathematical modelling*.

The other form of modelling, for which we retain the given name, starts from the (not quite universally accepted [2]) premise that brains are kinds of computers, and asks what computational functions bits of them subserve. This also requires quantitative descriptions, just like mathematical modelling, but it goes beyond them – in the example above not asking just how input correlations get turned into cortical structure by a Hebb rule, but also what this implies about the way information in those inputs is processed; not merely describing the process of Hebbian synaptic change and its structural consequences, but ascribing to it specific computational functions. Notions of the abstraction of tasks from their implementations are pervasive throughout computer science, and the question of what function some structure or mechanism for synaptic change or whatever has, has different answers at different levels of inquiry [3, 4, 5]. This form of modelling is somewhat more controversial, since the basic premise is not universally accepted, the more abstract computational models can appear to be very far removed from their experimental foundations, and the whole enterprise is too teleological for the tremulous.

To see that there is a difference, consider the extremely simplified example in the figure in which

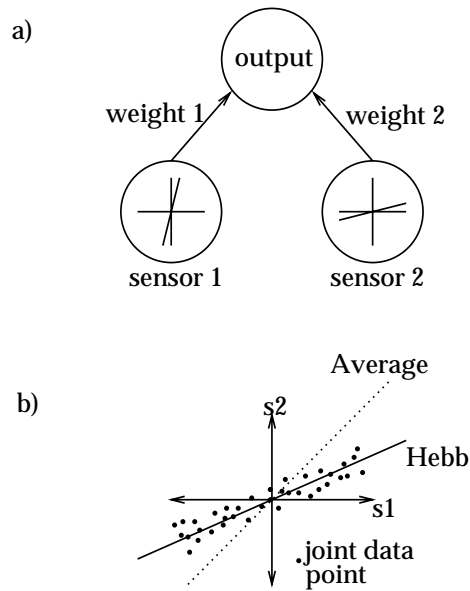


Figure 1: Simple example. a) Two sensors that look at the same quantity. Sensor 1 (s_1) responds more steeply than sensor 2 (s_2). b) Graph of the joint and slightly noisy activities of s_1 and s_2 showing the effect of the different sensitivities in the imbalance of the data points in the two directions. Hebbian learning for the weights can make them line up with the data (solid line). This is as far as mathematical modelling will go – it gives an accurate description of what happens. Computational modelling goes one step further and asks what this way of lining up with the data implies for the way these inputs get processed – in this case more information is extracted than would be from, say, straight averaging. See, for example, Linsker [6].

two sensors with different sensitivities report noisily on the same underlying quantity in a linear manner through a summarising unit with weights that develop according to Hebbian principles. In something of a caricature, mathematical modelling would just describe the relationship between the structure of the input (in this case the correlations between the sensor readings) and the final weights (which line up with the direction in which the inputs jointly vary the most). Computational modelling, however, would reveal that this extracts the maximum amount of information from these inputs – taking a weighted average is better than taking the straight average since it leads to a higher signal/noise ratio. Computational modelling asks more questions and can therefore be more revealing.

Although these two activities may appear to have a lot in common – for example, mathematical modelling also requires abstraction from underlying detail, particular instances of computational models are mathematical models, and the two are usually intimately intertwined in particular pieces of research – there are important differences. We shall try to show this by describing computational modelling in adult conditioning learning and contrasting the two in activity-dependent development. These areas share deep roots, but historically mathematical modelling has played a greater part in the latter, and a restricted form of computational modelling in the former.

Mathematical and computational modelling are pervasive throughout neurobiology. This review therefore focuses only on the distinction drawn above in these two example areas. The references include pointers to some recent work not explicitly discussed here (see particularly [7]. The controversy over the use of abstract algorithms (particularly backpropagation [8]) is not engaged. They mostly share only a certain computational flavour with the brain, however their merits as some of the few ways to explore the class of processing that moderate numbers of simple computational units can perform, are independent of the fact that they are clearly not abstractions over any neural process [9]. Also note that the distinction between the two forms of modelling is not cleanly separable from other common distinctions such as realistic/simplifying [4], or ones based exclusively on computational levels of analysis [3].

The most important area not covered is representation – the codings that are inferred from recordings in various neural systems, particularly from awake and behaving animals. The very notion of representation is itself computational, and abstract connectionist models have been used to suggest various ways in which cells might jointly represent some aspect of a stimulus (in so-called distributed codes) rather than doing it singly (in punctate or grandmother-cell codes). Computational modelling licences investigating these other forms of representation. Just as one example, Georgopoulos and his colleagues (see [10] for a recent review and [11] for some computational analysis) have studied the motor cortex, showing the relationship between the firing of a whole population of cells before a movement and the direction of that movement. The grip of the grandmother-cell has been weakened as investigators have found effects from outside the classical receptive field that is at the heart of this theory, and will be more so as simultaneous recording from many cells becomes possible [12].

Robinson [13] suggests that to the extent that these connectionist-inspired theories of distributed codes are actually borne out, it will be almost impossible to work out in detail how the brain processes particular sorts of information, rather than just the general strategies it employs. This is on the grounds that it is often almost impossible to take a connectionist system that is performing some task using abstract models of neurons, and reverse-engineer just how it does so.

2 Conditioning

Hebb's famous postulate [14] about the circumstances in which the efficacy of a connection between two cells would change can be seen as almost a cellular implementation of part of philosophical associationism. Hebbian learning rules quickly became important in theoretical models of memory and activity-dependent plasticity. However, animal conditioning experiments showed that they could not, at least on their own, account for the psychological data on the circumstances under which animals learn about the associations between stimuli and reinforcers [15]. Conditioned stimuli such as lights or tones generally have to precede delivery of unconditioned stimuli such as shocks or food for animals to learn the association between them; simultaneous presentation is not nearly so effective. Likewise, in a phenomenon called blocking, stimuli that are presented in perfectly appropriate temporal relationships with rewards can also fail to enter into associations if the delivery of reward is signaled by some other, previously conditioned, stimulus.

These breaks with the former doctrine were followed by a separation of most of the research in the neurobiological basis of plasticity, which was conducted along more strictly Hebbian lines, and modelling and experimentation in conditioning. The separation allowed a greater focus on the computational nature of conditioning, albeit divorced from the neural mechanisms by which it occurs, barring some notable and debatable exceptions in invertebrate conditioning and the cerebellum [16].

There are many strands of conditioning research and also theoretical conflicts that have yet to be resolved; however one main focus has been on how animals learn about the causal texture of their environments [15] and learn to act, presumably on the basis of this knowledge. When the task is posed in this way, insights from the study of learned prediction and action in other systems, including artificial ones, can be applied. To put it somewhat starkly, if plasticity in the brain is partly directed at learning predictions, then the way this is done is a member of the set of possible algorithms for learning to make predictions, and so shares any universal features of this set, and likely also features of other and possibly known algorithms. Using insights in this way is exactly what Churchland [19] calls a coevolutionary approach to the computational understanding of the brain, since the theories at all levels evolve in step.

Some of this interplay is evident in the evolutionary path of the influential Rescorla-Wagner [20] learning rule in conditioning. This explains in informational terms the blocking phenomenon described above, suggesting that since the previously conditioned stimulus predicts the delivery of the reward, there is nothing left unpredicted to cause learning for the new stimuli – it is an error correcting rule in the sense that learning is driven by prediction errors and acts to reduce them. Sutton and Barto [21] pointed out that this rule is equivalent to an algorithm called the delta-rule [22], which is well known in engineering as a way of allowing predictions to be made of some quantity on the basis of the values of other quantities to which it is related, and whose computational properties have been extensively explored. All this analysis applies directly to the Rescorla-Wagner rule. Sutton and Barto [21] went on to suggest a modification to this rule to take account of temporal effects in conditioning such as the utility of having the conditioned stimuli precede the unconditioned ones, and this led to a different algorithm for prediction which McLaren and Dickinson [23] pointed out had originally been suggested for animal learning by Konorski [24], the year before Hebb's own suggestion.

Sutton and Barto subsequently suggested a further modification to their own algorithm [25], making it consistent both with more conditioning experiments and with another engineering technique called dynamic programming (DP, [26]), in which prediction and optimal action are intimately bound [27]. DP is a quite general technique for systems to work out how to optimise some measure of return given certain conditions about the world and about the information from the world that is available to them, and so it permits a link to be made with more ethological theories of animal behaviour for which optimalities of various sorts are the starting point. In fact the algorithms turn out to be novel contributions to engineering too [28].

Of course, these conditioning theories are all incomplete – Mackintosh [29] and Pearce and Hall [30] point out the importance of attentional effects, and the true relationship between classical and instrumental conditioning is still unclear – however the interplay licenced in the framework of computational modelling has been productive.

The computational modelling is also incomplete to the extent that it is unclear what the neural mechanisms are over which these algorithms are abstractions. This is of coevolutionary importance since there are various confusing phenomena in conditioning such as the difference between the learning behaviour of conditioned excitors (the stimuli that come to predict the presence of an unconditioned stimulus) and inhibitors (which are described as coming to predict its absence), which are suggestive of neural constraints. The neural basis has attracted more interest in the invertebrate [31, 32] and cerebellar literature (see [16]) than that on cortical learning, although there are some pointers.

As an example of taking computational modelling one stage further, recent evidence from the vertebrate dopaminergic system, which has previously been implicated in the processing of rewarding events on the basis of self-stimulation and drug addiction studies [33], suggests that the firing of the dopaminergic output cells of the ventral tegmental area is consistent with the hypothesis [34] that it reports to its targets the information about prediction errors that would be expected from a rule such as Sutton and Barto's [25, 35]. Importantly for this suggestion, dopamine is believed to modulate synaptic plasticity [36].

Note that this explanatory process is neither top-down nor bottom-up. There are many possible algorithms for prediction, most of which are inconsistent with evidence about what animals actually do or how they do it; there are many facts about conditioning, particularly to do with conditioned inhibitors, whose natural explanation probably lies at the level of their neural underpinning; and yet the organising thrust of the computational modelling comes from the notion of prediction, which is a significant abstraction from any neuron firings.

3 Activity-dependent Development

Research into activity-dependent development has a venerable history. For a system in which the construction of representations at multiple levels *is* a main means of mediating computation, it is an attractive theoretical notion that the representations that form are sensitive to the stimuli that are actually observed. Although significant aspects of development are not dependent on input but are presumably genetically pre-specified, there is substantial evidence for input-driven components and it is those that have been the main focus of theories (but see [37]). Unfortunately, the equivalent of prediction as an organising computational principle has so far eluded the field, and the interplay of mathematical and computational modelling has been less productive.

Mathematical modelling had an early start, as algorithms based on Hebb's hypothesis were shown to be adequate for the development of the orientation selective cells [38] and topographically ordered mappings between sheets of sensory and central neurons [39] as in the retino-tectal projection in frogs or the retino-thalamo-cortical projection in cats and monkeys. Hebb's suggestion is about how innervating axons might cooperate – some form of competition is also required to prevent the outcome of development from being trivial [40]. Different ways of implementing the Hebbian or the competitive parts of the rule, different models for the activity itself (either random or somehow patterned [41]) from which the development arises, and different degrees of abstraction from the neural basis, all lead to apparently different algorithms (*eg* [42, 43, 44, 45, 46]), and

there is as yet not enough experimental data to distinguish between those that compete. Also, they have almost exclusively been confined to the primary sensory areas; cortical mechanisms are not apparently different deeper in, but it is not clear that these algorithms can be sensibly applied in an hierarchical manner.

From the perspective of computational modelling, there seems so far to be lacking an organising principle with the appeal that the notion of prediction had in conditioning theory. Purves, Riddle and LaMantia [47] highlight the extent to which this is true in suggesting that the modular structure in cortex that they all produce are effectively epiphenomenal byproducts of the underlying mechanisms governing synaptic plasticity. This issue is outside the purview of these models, since it is an appeal for a computational explanation of what is going on.

The obvious computational principle (*eg* [48] and many others) would be couched in terms of gleaning information from the input in such a way that it can more easily be used [49, 50] for subsequent processing – indeed it is known that algorithms based on Hebbian learning rules can perform a sort of principal components analysis [51, 52], a standard technique for extracting information from a set of inputs. It is important to bear in mind that there is no sentient ‘user’ but only further sub-cortical areas and also cortical areas presumably organising themselves in similar manners. The most serious concern about this principle, which is why it has been hard to codify, is that it is unclear what it is to be more easily usable – data that from the perspective of one system are information, can easily be noise to another. Various suggestions have been made such as that the outputs of a single step in the process should be more independent than the inputs, and yet convey the same information [49, 53]; however mechanisms for achieving this on a large scale are as yet not known. Related principles such as removing redundancy from inputs [54], have also been suggested; however these work better nearer the sensory periphery and are replaced by less powerful and more system-specific computational principles in modelling cortical organisation [55].

4 Conclusions

This review has pointed out some of the differences between two sorts of computational modelling, one which is an extension of mathematical modelling, and the other which is based on a qualitatively different approach. The two are particularly intertwined in the evolution of conditioning theory and activity-dependent synaptic plasticity, and the dearth of organising computational principles in the latter is thrown into sharp relief. Of course, the two have been somewhat caricatured to reveal their differences. Although mathematical modelling is typically easier to justify, closer to obvious characteristics of the underlying neural substrate, and more amenable to making predictions that can simply be tested, it is in the latter that lie multi-level explanations for neural computation.

Acknowledgements

I am very grateful to Read Montague for discussions and Sue Becker, Geoff Hinton, David Raizen and Chris Williams for comments on earlier drafts. Support was from the Canadian NSERC.

References

- [1] Bower, J, editor. *Trends in Neurosciences*, 1992 **15**:411-64.

This special issue focuses on what we have called mathematical modelling, and presents a number of examples, including the asymmetric coupling between segments of the lamprey spinal cord that was predicted by mathematical modelling (Cohen *et al*), the network controlling the pumping of the heart of the medicinal leech (Calabrese and De Schutter). Kawato and Gomi's article reviews some more computational models of the cerebellar control of the vestibular-ocular reflex.

- [2] Globus, G. Toward a noncomputational cognitive neuroscience. *Journal of Cognitive Neuroscience*, 1992 **4**: 299-310.

An interesting misunderstanding of the nature of computational theories in neuroscience. The paper attempts to knock down a chimerical straw man theory of representation and to justify an supposedly noncomputational alternative in which data from the sensory periphery provides constraints on a neural chemical and electrical dynamic self-organising process. Formalising this notion would seem to lead one straight back to the language and methods of computational analysis in the context of the interpretation of how non-linear dynamical systems compute.

- [3] Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: WH Freeman, 1982.

- [4] Churchland, PS and Sejnowski TJ. Perspectives on cognitive neuroscience. *Science*, 1988 **242**:741-5.

- [5] Churchland, PS & Sejnowski, TJ. *The Computational Brain*. Cambridge, MA: MIT Press, 1992.

An extensive survey of the state of the art in mathematical and computational modelling across a wide range of neural systems. It describes most of the techniques that practitioners use as well as examples of their application, and discusses many of the meta-issues in computational modelling, such as levels of abstraction, as well as the actual issues that the modelling has raised, such as representation in the brain, coordinate systems and adaptation.

- [6] Linsker, R. Self-organisation in a perceptual network. *Computer*, 1988 **21**:105-117.

- [7] Egelhaaf, M and Borst, A. A look into the cockpit of the fly: Visual orientation, algorithms, and identified neurons. *Journal of Neuroscience*, 1993 **13**:4563-4574.

Reichardt pioneered the motion processing system of the fly as a paradigm, and it has slowly succumbed to ever more sophisticated theoretical and experimental techniques. This is a review of what is the nearest thing to a classic example of coevolutionary computational modelling.

- [8] Rumelhart, DE, Hinton, GE and Williams, RJ. Learning representations by back-propagating errors. *Nature*, 1986 **323**:533-6.
- [9] Crick F. The recent excitement about neural networks. *Nature*, 1989 **337**:129-32.
- [10] Georgopoulos, AP, Taira, M & Lukashin, A. Cognitive neurophysiology of the motor cortex. *Science*, 1993 **260**:47-52.
 A review of some of the extensive experimental investigations by Georgopoulos and others in motor cortex looking at codings involving populations of cells rather than single ones. It raises points such as whether the coordinate system for a movement changes as a function of the starting position, shows how the representations observed can also arise in a simple connectionist network, and provides a handy guide to the process of cognitive neurophysiology.
- [11] Sanger, TD. Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation*, 1994 **6**:29-37.
- [12] Wilson, MA and McNaughton, BL. Dynamics of the hippocampal ensemble code for space. *Science*, 1993 **261**:1055-8.
- [13] Robinson, DA. Implications of neural networks for how we think about brain function. *Behavioral and Brain Sciences*, 1992 **15**:644-655.
 If it were the case that simple connectionist systems could be informative about neural processing, what would that imply? The coevolutionary message would be that the computational models should be used to suggest the neural representations and circuits involved in particular processing tasks and the experiments that would confirm or deny these and reveal what changes would be necessary to the computational models. Robinson agrees with the premise, but concludes that since we generally cannot work out in detail how connectionist systems perform a task, we are not going to be able to work out how the brain does so either. He concludes “if we know the learning rules, we may have to accept the inexplicable nature of mature networks.”
- [14] Hebb, DO. *The Organization of Behavior: A Neuropsychological Theory*. New York, NY: Wiley, 1949.
- [15] Dickinson, A. *Contemporary Animal Learning Theory*. Cambridge, England: Cambridge University Press, 1980.
- [16] Gluck MA and Granger R. Computational models of the neural bases of learning and memory. *Annual Review of Neuroscience*, 1993 **16**:667-706.
 This paper reviews the computational modelling of plasticity and its neural underpinnings in olfactory paleocortex, the cerebellum and the hippocampus. The cerebellar theories focus on conditioning in the rabbit eyeblink reflex [17], to which alternative theories have also been turned [18] which are more closely related to those discussed in the section above on conditioning. Error information can plausibly get to the cerebellum through the climbing fibres,

providing a substrate for error-correcting learning rules such as the delta rule. The hippocampus is a tantalising area for modelling because of suggestive behavioural evidence for its role in learning and memory and because of its anatomically regular structure. The difficulty in defining exactly what roles it plays has made many a computational theory flounder.

- [17] Thompson, RF and Gluck, MA. Brain substrates of basic associative learning and memory. In HJ Weingartner and RF Lister, editors, *Cognitive Neuroscience*. New York, Oxford: OUP, 1990.
- [18] Moore, JW, Desmond, JE and Berthier, NE. Adaptively timed conditioned responses and the cerebellum: a neural network approach. *Biological Cybernetics*, 1989 **62**:17-28.
- [19] Churchland, PS. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press, 1986.
- [20] Rescorla, RA and Wagner, AR. A theory of Pavlovian conditioning: The effectiveness of reinforcement and non-reinforcement. In AH Black & WF Prokasy, editors, *Classical Conditioning II: Current Research and Theory*, 64-69. New York, NY: Appleton-Century-Crofts, 1972.
- [21] Sutton, RS and Barto, AG. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 1981 **88**:135-170.
- [22] Widrow, B and Stearns, SD. *Adaptive Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [23] McLaren, IPL & Dickinson, A (1990). The conditioning connection. *Philosophical Transactions of the Royal Society of London, B*, **329**, 179-186.
- [24] Konorski, J. *Conditioned Reflexes and Neuron Organisation*. Cambridge, CUP, 1948.
- [25] Sutton, RS & Barto, AG (1987). A temporal-difference model of classical conditioning. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Seattle, WA.
- [26] Bellman, RE. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957.
- [27] Barto, AG, Sutton, RS and Watkins, CJCH. *Learning and Sequential Decision Making*. Technical Report 89-95, Computer and Information Science, University of Massachusetts, Amherst, MA, 1989.
- [28] Barto, AG, Bradtke, SJ and Singh, SP. *Real-Time Learning and Control using Asynchronous Dynamic Programming*. TR 91-57, Department of Computer Science, University of Amherst, MA. 1992.
- [29] Mackintosh, NJ. A theory of attention: Variations of the associability of stimuli with reinforcement. *Psychological Review*, 1975 **82**:276-98.
- [30] Pearce, JM and Hall, G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 1980 **87**:532-52.
- [31] Hawkins, RD, Abrams, TW, Carew, TJ and Kandel, ER. A cellular mechanism of classical conditioning in Aplysia : Activity-dependent amplification of presynaptic facilitation. *Science*, 1983 **219**: 400-5.

- [32] Hammer, M. An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature*, 1993 **336**:59-63.
 - [33] Wise, RA. Neuroleptics and operant behavior: The anhedonia hypothesis. *Behavioral and Brain Sciences*, 1982 **5**:39-87.
 - [34] Montague, PR and Sejnowski, TJ. The predictive brain: From temporal coincidence to temporal order. *Learning and Memory*, in press.
 - [35] Ljungberg, T, Apicella, P and Schultz, W. Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, 1992 **67**:145-163.
 - [36] Miller, R & Wickens, JR. Corticostriatal cell assemblies in selective attention and in representation of predictable and controllable events. *Concepts in Neuroscience*, 1991 **2**:65.
 - [37] Willshaw, DJ and von der Malsburg, C. A marker induction mechanism for the establishment of ordered neural mappings: its application to the retinotectal problem. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 1979 **287**:203-43.
 - [38] von der Malsburg, C. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 1973 **14**:85-100.
 - [39] Willshaw, DJ & von der Malsburg, C. How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 1976 **194**:431-45.
 - [40] Grossberg, S. How does the brain build a cognitive code? *Psychological Review*, 1980 **87**:1-51.
 - [41] Meister, M, Wong, RO, Baylor, DA and Shatz CJ. Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science*, 1991 **252**:939-43.
 - [42] Miller, KD, Keller, JB and Stryker, MP. Ocular dominance column development: Analysis and simulation. *Science*, 1989 **245**:605-15.
 - [43] Goodhill, GJ and Willshaw, DJ. Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network*, 1990 **1**:41-59.
 - [44] Montague, PR, Gally, JA & Edelman, GM. Spatial signaling in the development and function of neural connections. *Cerebral Cortex*, 1991 **1**:199-220.
 - [45] Durbin, R and Mitchison, G. A dimension reduction framework for understanding cortical maps. *Nature*, 1990 **343**:644-7.
 - [46] Obermayer, K, Ritter, H and Schulten, K. A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy Of Sciences*, 1990 **87**:8345-9.
 - [47] Purves D, Riddle DR & LaMantia AS. Iterated patterns of brain circuitry (or how the cortex gets its spots). *Trends in Neurosciences*, 1992 **15**:362-8.
- Most of the Hebbian-inspired algorithms that turn patterned input into modelled cortex produce (or would produce) modular structure like ocular dominance columns or the whisker

barrels seen in certain somatosensory cortices. Since various species do not exhibit these structures and yet seem unimpaired in the relevant aspects this paper argues that they are epiphenomenal products of the rules for synaptic change. This is an issue that mathematical models are in principle incapable of resolving.

- [48] Linsker, R. How to generate ordered maps by maximising the mutual information between input and output signals. *Neural Computation*, 1989 **1**:402-411.
- [49] Barlow, HB. Unsupervised learning. *Neural Computation*, 1989 **1**:295-311.
- [50] Becker, S and Hinton, GE. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 1992 **355**:161-3.
- [51] Oja, E. A simplified neuron model as a principal component analyser. *Journal of Mathematical Biology*, 1982 **15**:267-73.
- [52] Hertz, JA, Krogh, AS and Palmer, RG. *Introduction to the Theory of Neural Computation*. Addison Wesley, California, 1991.
- [53] Földiák, P. *Models of Sensory Coding*. Thesis, University of Cambridge, England, 1992.
- [54] Atick, JJ and Redlich, AN. Towards a theory of early visual processing. *Neural Computation*, 1990 **2**:308-320.
- [55] Li, Z and Atick, JJ (1994). Towards a theory of striate cortex. *Neural Computation*, 1994 **6**:127-146.

Peter Dayan
University of Toronto
6 King's College Road
Toronto, Ontario. M5S 1A4
Canada