

From synaptic errors to thalamocortical circuitry

Terry Elliott

Recent data indicate that newly grown synapses in the brain are not guaranteed to innervate their desired target, but can form instead on nearby targets. Such 'errors' introduce representational inaccuracies but improve representational flexibility. Optimizing accuracy and flexibility requires detecting correlated activity and disabling plasticity, explaining the structure of the thalamocortical circuit.

How does the brain learn from experience? How does it maintain the integrity of learned experiences while retaining the capacity to learn new ones? And how does it discard, or forget, experiences that have ceased to be of value? These are all questions of central importance to neuroscience. The ability of synapses to undergo change, called synaptic plasticity, resulting in the strengthening or weakening of the influence of one neuron over another, is believed to constitute a major component in the biological basis of learning. Broadly speaking, synaptic plasticity can be either physiological (e.g. a change in the level of neurotransmitter released at a synapse) without any associated change in the morphology of the network of synaptic connections, or it can be anatomical (e.g. growth of more synapses or removal of existing ones). Frequently, physiological plasticity precedes anatomical plasticity [1].

Conventional models of synaptic plasticity
Mathematical and computational modellers of synaptic plasticity have traditionally ignored anatomical plasticity, focusing instead on physiological plasticity in an anatomically fixed, completely connected network in which all pre-synaptic or afferent neurons synapse on all post-synaptic or target neurons. Following Hebb's postulate [2], the rule for synaptic plasticity in these networks is that a synapse is strengthened if the afferent and target cells' firing patterns are correlated, but weakened if their firing patterns are anti-correlated: cells that 'fire together wire together', as the cliché

runs. Physiological models for the Hebb postulate are long-term potentiation (LTP) [3] and long-term depression (LTD) [4]. In LTP, a high-frequency sequence of afferent spikes coupled with a post-synaptic response elicits a strengthening, or potentiation, of the synapses involved. In LTD, a low-frequency series of afferent spikes coupled with a small or no post-synaptic response elicits a weakening, or depression, of the synapses involved.

'With synaptic errors a model can [...] dig itself out of a no-longer-required pattern of connectivity...'

Experimental data indicate, however, that LTP is not specific to the stimulated pathway but can spread to neighbouring, unstimulated pathways [5–8]. Furthermore, LTP appears to occur in a digital, all-or-none fashion: investigation of pairs of neurons coupled by only one synapse in the hippocampus reveals that synaptic strength either doubles or remains unchanged during LTP [9]. One interpretation of this latter result is that LTP, when it induces a change at a synapse, leads to synaptic doubling, perhaps involving the growth of a new synapse. The former results, however, suggest that this new synapse is not guaranteed to form on the activated post-synaptic cell, but can form on nearby, inactive cells too. But if LTP is a biological model for experience-dependent learning, then can these 'errors' in synaptic growth actually disrupt the laying down of appropriate, new patterns of synaptic connectivity or corrupt existing ones? This is the question addressed in a recent paper by Paul Adams and Kingsley Cox [10].

A mathematical model of anatomical plasticity

Adams and Cox write a simple mathematical model of anatomical plasticity in which the possibility of synaptic errors, or 'mutations' as they call them, is explicitly included. Their synaptic growth rule consists of two pieces. One piece corresponds to a

standard Hebb rule, although rather unconventionally formulated. Instead of expressing the Hebb rule in terms of correlations between pre- and post-synaptic firing patterns, they express it in terms of a 'fitness function'. Essentially, where the 'synaptic fitness' exceeds average fitness, synapses grow; but where synaptic fitness is lower than average, synapses are removed. In conventional Hebbian language, above-average fitness corresponds to correlations in pre- and post-synaptic activity higher than average, whereas below-average fitness corresponds to correlations in activity lower than average.

The second piece of their synaptic growth rule models synaptic errors, and takes the mathematical form of a diffusion term with a characteristic diffusion constant. The size of this constant or error rate determines how rapidly errant synapses can spread or 'diffuse' from a region of high fitness into a region of low fitness. In the steady state of a dynamic equilibrium, the spread of errant synapses into a region of low fitness is exactly matched by their removal there. The profile of errant synapses in the low fitness region is that of a decaying exponential, with some characteristic space constant. This space constant depends on the error rate and on the ratio of fitnesses in the high and low fitness regions. The larger the error rate or the smaller the fitness ratio, the larger the space constant, leading to a broader errant synaptic profile in the low fitness region.

Dynamic flexibility of anatomical models of plasticity

It is precisely the error rate that allows the model studied by Adams and Cox to remain 'open to new ideas'. In conventional models, given a fixed set of experiences, final states of connectivity are typically reached in which the system becomes 'stuck': the presentation of a new set of experiences, which might require the system to adopt a different pattern of connectivity, does not induce any change

in the existing network. With synaptic errors, however, a model can, albeit perhaps rather slowly, dig itself out of a no-longer-required pattern of connectivity and fall into a different pattern. Synaptic errors therefore admit of dynamic flexibility in the network's architecture, but they also reduce the accuracy of the resulting representations.

Optimizing accuracy and flexibility: the thalamocortical circuit

Finding the optimal trade-off between flexibility and accuracy is critical. What is this optimum? Adams and Cox have no definitive answer, but suggest that the brain might set its own optimum by switching off plasticity in some synapses when they threaten to reduce accuracy unacceptably. They argue that there is likely to be an absolute minimum, biophysically realizable synaptic error rate below which a system cannot fall, so that the brain cannot simply rely on reducing the error rate to zero when necessary. Suppose, instead, that there is a minimum-acceptable level of accuracy, or, because the accuracy is set by the spread of errant synapses, a maximum-acceptable space constant. This maximum space constant sets a minimum-acceptable level for the fitness ratio. Adams and Cox thus suggest that when the system detects that the fitness ratio is about to fall below this minimum-acceptable level, it should disable the plasticity of those synapses involved.

So far so good. But how is the brain to detect when the fitness ratio is at this minimum level, and how does it switch off the plasticity of afferent synapses? Adams and Cox suggest that neocortical circuits and their thalamic inputs appear ideally suited for these purposes [11,12]. Layers 4 and 6 of the neocortex are the principal recipient layers for input from the thalamus, with layer 6 also receiving input from layer 4. Layer 6 cells therefore receive input from both the thalamus and layer 4, and could potentially calculate the ratio in correlations between pre- and post-synaptic activity in adjacent pairs of layer 4 cells required by their scheme. Furthermore, layer 6 cells send outputs back to the thalamus, changing the firing mode of thalamic cells from burst to tonic [13]. Adams and Cox argue that tonic mode enables the plasticity of thalamic inputs onto layer 4 cells, and burst mode switches off plasticity. They therefore propose that

layer 6 cells fire only when the correlation ratio exceeds the minimum-acceptable level, and that this spike output alters the pattern of thalamic firing, enabling the plasticity of thalamic inputs to layer 4.

Evaluation and conclusion

From small beginnings, in the form of a consideration of errant synapses, Adams and Cox arrive at a rather elegant account of the structure and function of the thalamocortical circuit. Such an account is sorely needed: the paucity of decent theories in the neurosciences is such that the thalamus is commonly regarded as a mere 'relay' of sensory input to the neocortex. However, the history of science is replete with elegant but wrong theories. How seriously should we take the ideas of Adams and Cox? It is important to separate the two parts of their approach.

...the paucity of decent theories in the neurosciences is such that the thalamus is commonly regarded as a mere "relay"...

The first part, based on an analysis of a model of synaptic errors, is a much clearer contribution. There is no doubt that the modelling community has largely neglected anatomical plasticity, and has ignored the impact of synaptic errors on the formation and resulting accuracy of neuronal representations of experience. Synaptic errors, although reducing representational accuracy, have the considerable virtue of increasing representational flexibility, allowing a neuronal network's synaptic resources to be redeployed in the event that existing representations are rendered obsolete by changing experiences. But synaptic errors also introduce potentially fatal inaccuracies into these neuronal representations.

Their explanation of the thalamocortical circuit, in an attempt to balance neuronal flexibility and accuracy, is, on the other hand, much more difficult to assess. It is possible that other, simpler mechanisms might achieve the reduction in synaptic errors that their scheme requires. For example, their model assumes that the space constant characterizing the spread of errant synapses can become infinitely large. However, there is certainly a limit on the size of axonal arborization, and such a limit will have the effect of imposing a

maximum-achievable space constant. If this maximum is near the maximum-acceptable space constant discussed earlier, then no machinery at all is necessary to halt errant synaptic spread. Of course, appealing to axonal arbors of limited size does potentially limit network flexibility, but then so does imposing a maximum acceptable space constant. To be sure, theirs is an elegant and courageous attempt at explaining the structure and function of the thalamocortical circuit, but whether or not their account is correct is ultimately an experimental matter.

Acknowledgement

I thank the Royal Society for the support of a University Research Fellowship.

References

- Colman, H. *et al.* (1997) Alterations in synaptic strength preceding axon withdrawal. *Science* 275, 356–361
- Hebb, D.O. (1949) *The Organization of Behavior*, John Wiley & Sons
- Bliss, T.V.P. and Lomo, T. (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356
- Lynch, G.S. *et al.* (1977) Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature* 266, 737–739
- Bonhoeffer, T. *et al.* (1989) Synaptic plasticity in rat hippocampal slice cultures: local 'Hebbian' conjunction of pre- and postsynaptic stimulation leads to distributed synaptic enhancement. *Proc. Natl. Acad. Sci. U. S. A.* 86, 8113–8117
- Kossel, A. *et al.* (1990) Non-Hebbian synapses in rat visual cortex. *NeuroReport* 1, 115–118
- Schuman, E.M. and Madison, D.V. (1994) Locally distributed synaptic potentiation in the hippocampus. *Science* 263, 281–284
- Engert, F. and Bonhoeffer, T. (1997) Synapse specificity of long-term potentiation breaks down at short distances. *Nature* 388, 279–284
- Petersen, C.C.H. *et al.* (1998) All-or-none potentiation at CA3–CA1 synapses. *Proc. Natl. Acad. Sci. U. S. A.* 95, 4732–4737
- Adams, P.R. and Cox, K.J.A. (2002) Synaptic Darwinism and neocortical function. *Neurocomputing* 42, 197–214
- Gilbert, C.D. (1983) Microcircuitry of the visual cortex. *Annu. Rev. Neurosci.* 6, 217–247
- Callaway, E.M. (1998) Local circuits in primary visual cortex of the macaque monkey. *Annu. Rev. Neurosci.* 21, 47–74
- Sherman, S.M. and Guillery, R.W. (1996) The functional organization of thalamocortical relays. *J. Neurophysiol.* 76, 1367–1395

Terry Elliott

Dept of Electronics and Computer Science, University of Southampton, Highfield, Southampton, UK SO17 1BJ.
e-mail: te@ecs.soton.ac.uk