

Implications of Synaptic Digitisation and Error For Neocortical Function

Recent work in hippocampus (1,2) has shown that correlation-based synaptic strengthening occurs digitally and may spread to nearby, inactive, connections. We are exploring a model in which co-activity strengthens connections by adding synapses. The added synapses usually appear at the co-active connection, but with some low probability E they appear at neighboring cells (3). Regular Hebbian learning is thus a special case of our model, with $E = 0$. We assume that the total number of synapses formed by a given presynaptic cell upon a row of postsynaptic cells is constant. With $E = 0$, the equation that describes our model is $dy/dt = (w - \langle w \rangle)y$, where y is the number of synapses at a connection, w is the "fitness" of a synapse (which depends on the degree of correlation across it), and $\langle w \rangle$ the mean fitness of all synapses. Allowing errors in placing new synapses, the equation becomes $\partial y/\partial t = (w - \langle w \rangle)y + 0.5 w E \partial^2 y/\partial x^2$, where x is position along the postsynaptic row, and w is a function of x . This equation represents the random left/right placement of erroneous synapses as a diffusion process. If the presynaptic neuron is well correlated ($w = w_m$) with a small mesa of neurons, and poorly correlated ($w = w_p$) with a large surrounding plateau of neurons, then in the steady state most of the synapses will accumulate in the high fitness zone, but there will be an exponential fringe of synapses trailing off into the low fitness surround. This fringe of synapses will lower the accuracy of the connection compared to the zero error case, to an extent that depends on E and w_m/w_p . We have verified this using a computer simulation in which synapses appear according to a mesa-plateau fitness function, with a probability E of misplacement to the right or left (Fig 1a). As expected we found that increasing E or decreasing (w_m/w_p) increased the spread of the connection beyond the mesa. We also simulated the transfer of synapses from an initial high fitness region at one end of the row to a new high fitness region at the other. As expected we observed (Fig 1b) a variable delay in the transfer process that was most evident with small numbers of synapses.

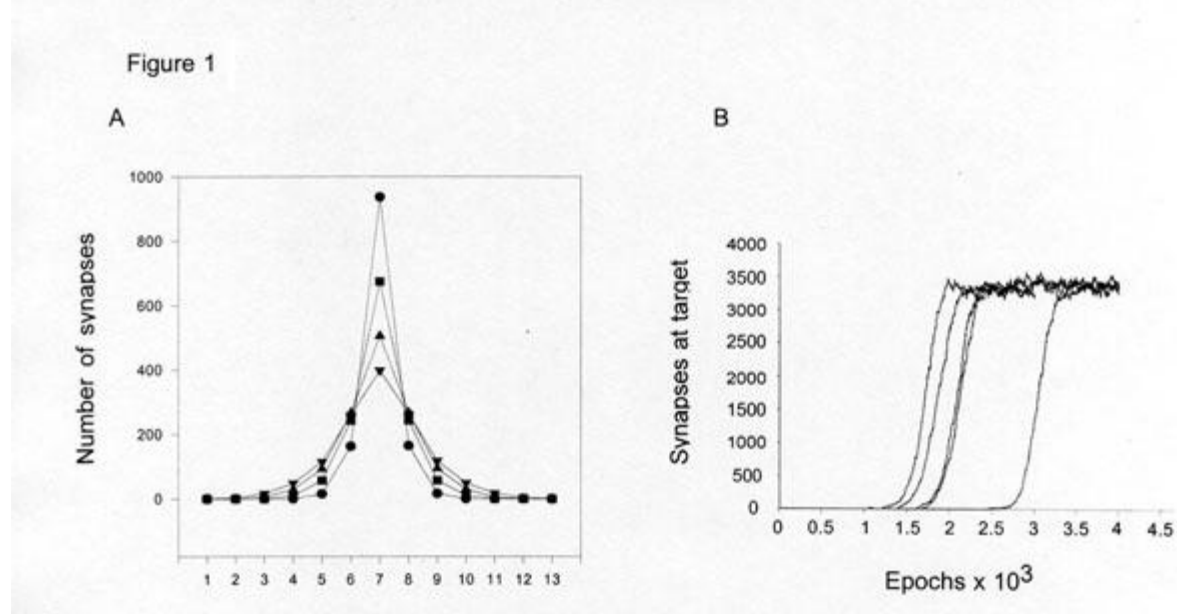


Fig 1. (A) Steady state synapse spread beyond the target (cell 7) increases with E (circles: 0.1, squares: 0.2, up triangles: 0.3, down triangles: 0.4; $w_m/w_p = 1.4$). (B) 13,000 synapses were

placed initially on cell 1 and the number of synapses at the target (cell 13) is plotted for 5 repeated runs with $E = 0.1$, $(w_m/w_p) = 1.05$.

The conclusion that the accuracy of correlation-dependent wiring depends on both the fidelity with which new synapses are placed and upon the sharpness of the correlation profile is not surprising but it has interesting implications for the formation of large networks such as neocortex. The useful size of networks may be limited by the accuracy with which they can be wired. Evolution has presumably lowered error rates as far as possible, by placing synapses on spines isolated from neighbors. However, errors ultimately reflect molecular noise, and cannot be eliminated without making synapses bulky and sparse. It is doubtful if neocortex has appreciably more reliable synapses than other brain regions. The other obvious way to improve the accuracy of cortical wiring is to enhance the sharpness of correlations. This strategy is also unavailable to cortex (in the short term) because correlations depends on patterns in the real world and on analysis by previous layers. However, although a given cortical area cannot improve the presented correlations, at least it can measure them. If it determines that the profile of correlations is sharp (w_m/w_p exceeds some threshold), then Hebbian learning can proceed. If however the profile is blunt, learning should be disabled until changes in the world, or in preceding layers, produce clearer patterns. Fig 2 shows a circuit that ensures any desired degree of wiring accuracy. A top layer of cells projects to a middle, output, layer, and the matrix of digital synapses between these layers transforms input vectors to output vectors in the standard connectionist manner. Consider an input cell which is currently wired to an output cell (shown as solid circles), but which could become wired (open circles) to the neighbors of the output cell as a result of errors in the strengthening of the existing connection. The bottom cells act as correlation detectors. Thus the bottom central cell is excited when it receives nearly simultaneous spikes from the currently connected pair of cells, but is inhibited by its neighbors if they receive nearly simultaneous spikes from the top neuron and the neighbors of the output neuron. The depicted distal placement of inputs from the top layer and the proximal placement of inputs from the middle layer will assist in such coincidence detection, but could be supplemented by various nonlinearities. If the difference in the correlation signals measured by the central bottom neuron and its flanking neighbors is sufficiently great that accidental formation of erroneous connections will not spread, then the central bottom cell fires and signals the top layer cell that all its synaptic connections within the middle layer can be plastic. This could be achieved if the bottom cells send axons back to the top cells to control their firing mode, and if the firing mode controls the plasticity of the synapses formed by top layer cells. If the correlation profile measured by the bottom cell is too low, then it remains silent, and the top-to-middle synapses remain implastic. This arrangement hinges on the distinction between "connected" and "unconnected" inherent in synaptic digitisation. Only if connections are sparse is it feasible to use a layer of correlation detectors to control learning. In the simplest case, there need only be $n \times m$ bottom cells, where n is the number of cells in the top or middle rows, and m is the number of neighbors. The arrangement shown in Fig 2 is similar to that in neocortex, with the top row being the thalamic relay cells, the middle row, layer 4, and the bottom row, layer 6. Firing of layer 6 cells switches relay cells from burst to tonic mode. A similar solution should also apply to the subsequent layers of cortical processing, from 4 to 2/3 and from 2/3 to 5. However now the feedforward connections whose plasticity must be regulated are entirely intracortical, so that the

burst/tonic transition of thalamic relay cells cannot be directly used. Instead the plasticity of these connections must be regulated postsynaptically. It is likely this is done by layer 1 inputs to apical tufts of layer 2,3 and 5 pyramidal cells, which set up local calcium spikes, which facilitate backpropagation along the apical dendrite. These arrangements also agree well with the observation that there are, in cat visual cortex, "simple" layer 6 cells (which would compute the T-4 correlations) and "complex" layer 6 cells (which would compute the 4-2/3 and 2/3-5 correlations). Interestingly these simple and complex layer 6 cells send collaterals to layers 4 and 2/3/5 respectively, where they make shaft "drumstick" synapses like those formed on the distal dendrites of relay cells. These synapses could use the same correlation sharpness signals to modulate within-layer recurrent excitation.

Figure 2




Fig 2. Proposed thalamocorticothalamic circuitry to contain error. Vertical arrow controls presynaptic plasticity. Horizontal arrows show interactions computing correlation ratios.

This solution to the problem posed by synaptic error carries 2 interesting drawbacks. First, neocortex will learn more slowly than if it were a fully-connected error-free network. Second, the solution does not prevent error, but merely contains it. If the connections are plastic, errant synapses can still occur (though they will likely be eliminated by competition with "correct" synapses). Errant synapses can serve as seeds of future learning, so that a "mutating" partly-connected network can eventually attain the weights achievable directly in a fully-connected perfect network. Consider the extreme case where, as a result of daytime learning, an errant

synapse flourishes and a relay cell transfers its allegiance to a neighbor of the original layer 4 cell (Fig 3). Before any further plasticity is allowed, the connections of the layer 6 cells must be updated, so they can continue to limit error spread. This updating, which involves breaking the dotted connections and making the dashed connections in Fig 3, must be done offline, with the feedforward T-to-4 connections rendered implastic. Layer 4 to layer 6 connections, which define the columns of cortex, are permanent. The connections from layer T to layer 6 can be updated if the T cell fires in bursts, since this will fire the newly connected layer 4 cell, which will fire the correct layer 6 cell. The connection from the new layer 4 cell to its partner layer 6 cell will undergo Hebbian strengthening, and errant synapses will form onto the neighbors of that layer 6 cell, as required. If this bursting activity advances as a calibrating wave across thalamic nuclei, it will automatically update all the T-to-6 connections. Updating the return pathway to thalamus is trickier, because the new layer 6 cell must find all the relay cells that comprise the receptive field of its partner layer 4 cell. Because the layer 6 cells act as correlation detectors, they can do reverse correlation analysis to determine this receptive field. White noise must be played into thalamus to perform offline updating of the corticothalamic connections. Now it is required that a group of relay cells fire the new layer 4 cell, so the calibrating, white noise, input should be in tonic mode. These requirements match the features of slow wave and paradoxical sleep. In the former case traveling bursts are imposed on thalamus by its reticular nucleus, in the latter, irregularly discharging brainstem cholinergic neurons bombard relay cells with brief nicotinic epsps.

Fig 3 shows the start and end of the allegiance transfer, but not intermediate points, when the thalamic cell makes synapses on both layer 4 cells, which at the end of the day comprise an extended high fitness zone. Now it is the average correlation in the high fitness zone that should be compared to the flanking correlations, and used to control the plasticity of the thalamic neuron. This can be achieved if there is also offline updating of the lateral interactions between neighboring layer 6 cells, so that during the transfer the layer 6 cells marked 0 and 1 act as a unit, comparing their average activity to that of the flanking cells (marked -1 and 2), and both feeding back to the relay cell. This (and similar updating in layers 2, 3 and 5) can again be accomplished in slow wave sleep. Because the connections of the layer 6 complex-type cells that control postsynaptic plasticity must also be updated offline in a similar manner during sleep, and because the required calibration signals (burst waves and white noise) are generated thalamically, layer 6 control of postsynaptic plasticity must also loop back through thalamus, via matrix relay cells that synapse in layer 1. In this case white noise input must be played into the postsynaptic cells, presumably by random matrix cell spikes fired into the apical tufts under conditions where apical epsps initiate somatic spikes.

Figure 3

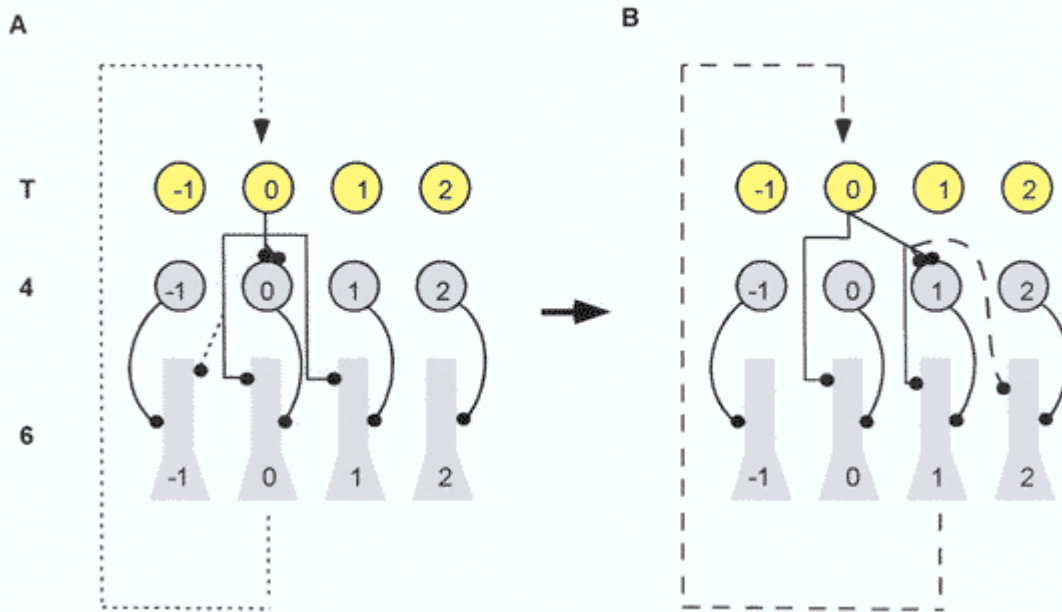


Fig. 3. Rewiring after an allegiance shift. Dotted (A) connections change to dashed (B).

There are gaps in the above reasoning. First, new synapses may be placed with perfect fidelity. Under these conditions there is no way to explore new connections, and by shrinking synapses to the point where errors occur, great increases in computational power could be obtained. Second, implicit in the notion that large networks need accurate wiring is an assumption of nonlinearity. Our simple model of synapse formation is linear, and wiring degrades gracefully with increasing error rates or decreasing correlation sharpness. However, real computation requires nonlinearities, as does slow learning in backpropagation networks. We are looking for wiring catastrophes in nonlinear models. Third, the exact mechanisms of presynaptic plasticity control are unclear. Fourth, although the proposed detailed wiring is consistent with known anatomy, it goes slightly beyond it. Nevertheless, the principle of error containment by correlation ratio measurement offers an attractive new view of cortical function.

References

1. C.C.H. Petersen, R.C. Malenka, R.A. Nicoll and J.J. Hopfield. All-or-none potentiation at CA3-CA1 synapses. *Proc. Natl. Acad. Sci.* 95 (1998), 4732-4737
2. F. Engert and T. Bonhoeffer. Synapse specificity of long-term potentiation breaks down at short distances. *Nature.* 388 (1997), 279-284.
3. P.R. Adams. Hebb and Darwin. *J.Theor.Biol.*195(1998), 419-438.

BIOGRAPHIES Kingsley Cox did his PhD at Sussex University cloning peptide receptors in snail neurons. At Stony Brook he has done postdoctoral work on calcium imaging in zebra fish, and most recently on modeling synaptic darwinism. Paul Adams is a Professor in the Neurobiology Department. His early work was on the biophysics of synaptic transmission, spiking and calcium signalling. More recently he has worked on cellular aspects of thalamus. He has a PhD from London University.