Adaptivity and 'Per learning'

Joseph Rushton Wakeling[a,b,*]

[a]*Institut de Physique Théorique, Université de Fribourg, Pérolles, CH-1700 Fribourg, Switzerland*
[b]*Niels Bohr Institute, Blegdamsvej 17, DK-2100 Copenhagen Ø, Denmark*

Abstract

One of the key points addressed by Per Bak in his models of brain function was that biological neural systems must be able not just to learn, but also to *adapt*—to quickly change their behaviour in response to a changing environment. I discuss this in the context of various simple learning rules and adaptive problems, centred around the Chialvo-Bak 'minibrain' model (Neurosci. 90 (1999) 1137).

1. Introduction

When attempting to model biological learning, what factors should we take into account, and what sort of problems should we expect our models to solve? One of the things which Per Bak always emphasized was that it was not enough to simply learn one task fast: biological neural dynamics had to be able to *adapt*, to *un*learn patterns of behaviour that were no longer working and find new ones. For example, the important early work on 'reinforcement learning' by Barto and colleagues [1–3], which produced much more biologically plausible learning rules than those previously considered, still foundered on this problem, with networks having to be completely reset in order to learn a new problem.

---

* Institut de Physique Théorique, Université de Fribourg, Pérolles, CH-1700 Fribourg, Switzerland.
E-mail address: joe@nbi.dk (J.R. Wakeling).
URL: http://neuro.webdrake.net/.

Particular progress in this regard was made by the work of Dmitris Stassinopoulos, in collaboration with Preben Alstrøm [4] and Per himself [5]. However, it was a few years later that an especially elegant model was developed by Per in collaboration with Dante Chialvo [6]. This model, which I rather cheekily dubbed the 'minibrain'[1] [7], addressed the problem of adaptation by assuming that learning was by only *long-term synaptic depression* (LTD), the weakening of connections. Synapses involved in bad decisions were suppressed, but only enough to render them inactive. Meanwhile, because synapses were not strengthened or reinforced in any way—there was a complete absence of long-term potentiation (LTP)—the strengths of active connections remained barely greater than those of the inactive ones. Thus, the network could easily switch to using a different set of connections if the need arose. It was suggested that it was the absence of any strengthening of connections in the model that was the key to its adaptive ability. In the present work I illustrate this by investigating the adaptive ability of the minibrain when simple forms of LTP are included, compared to the original model and the 'selective punishment' extension later proposed [8].

2. The model

For simplicity, I consider here only the basic feedforward minibrain: 3 layers of neurons ('input', 'intermediary' and 'output') of size n_{ip}, n_{im} and n_{op} respectively. Each neuron in the input layer has a one-way connection to every neuron in the intermediary layer, and similarly each intermediary neuron has a connection to every output. Each connection is assigned a strength value, initially evenly distributed in the interval $[0, W]$ (here $W = 1$). Activity propagates according to *extremal dynamics*: if we stimulate a neuron, the signal travels along the single strongest outgoing connection, and the neuron at the end of that connection then fires, and so on until an output neuron fires.

Should this output be incorrect, a negative feedback signal is sent to the system and the connections responsible are punished by having their strengths reduced by a random amount in the interval $[0, \delta]$ (here $\delta = 1$). Learning efficiency is measured by the total number ϕ_1 of such signals required for complete learning. Depending on a control parameter $\zeta = n_{im}/(n_{ip}n_{op})$, two phases of behaviour are identifiable [9]: for $\zeta < 1$ the network is in the *disordered* phase where complete learning is impossible and $\langle \phi_1 \rangle = \infty$. For $\zeta > 1$ complete learning becomes possible with $\langle \phi_1 \rangle \approx n_{ip}n_{op}$. In the present work, $n_{ip} = 10$, $n_{im} = 100$ and $n_{op} = 10$, with the n_{im} value picked to ensure the network is in the ordered phase so complete learning will always be possible.

---

[1] Much to my embarrassment, after Per and I published our collaboration using this name, Per told me that he had always called it the 'Dante brain', and Dante called it 'Per learning'. Readers are invited to draw their own conclusions but should not read anything whatsoever into the title of the present paper... ;).
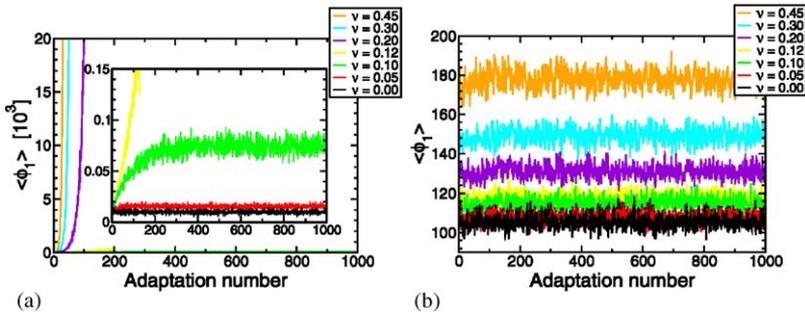
Fig. 1. Adaptive ability of the minibrain model when unbounded LTP is included, compared to the traditional version with LTD only. (a) At each adaptation, one input–output association is randomly reselected ('slow change'). The inset shows more clearly the results for smaller v. Note the astonishing increase in the required value of ϕ_1 for $v > 0$: only for the smallest values, $v \leqslant 0.1 = 1/n_{\mathrm{ip}}$, is this bounded, and then performance is still significantly worse than for the LTD-only network. (b) The network is required to adapt successively back and forth between two different input-output maps (the 'flip-flop' problem). $\langle \phi_1 \rangle$ is bounded in all cases, but nevertheless $v = 0$ (LTD only) provides the best performance. Data averaged over 128 realizations.

3. Unbounded potentiation

The simplest manner of including LTP is symmetric with the negative feedback: in the event of a successful decision, the connections responsible can be rewarded by having their strengths increased by a random amount in the interval $[0, v]$. How does this affect the network's adaptive ability?

Suppose that we present a network with an input–output map to learn and, each time learning is completed, randomly reroll one of the input–output associations. We can call this the 'slow change' problem. Naturally we want to know what value of ϕ_1 the network will require to adapt to the new maps.[2] We stimulate each of the inputs in turn and apply LTP or LTD as necessary; learning is deemed complete when we can run through all the inputs without error. Fig. 1a shows how the average $\langle \phi_1 \rangle$ varies with successive adaptations. Most of the networks with $v > 0$ very quickly become hopelessly addicted, requiring huge amounts of negative feedback to adapt to new input–output maps; this is escaped only with the smallest values of $v > 0$. Indeed, we can observe two distinct phases of behaviour: for $v \leqslant 0.10$, $\langle \phi_1 \rangle$ is bounded (though always worse than the $v = 0$ case), while for $v > 0.10$ we have the real addiction with $\langle \phi_1 \rangle$ growing exponentially.

We can explain this as follows. Consider one input neuron. On average, n_{ip} adaptations will pass before its associated output is reselected. Since it is receiving positive

---

[2] Including LTP in this way might lead one to question whether it is still appropriate to use this measure of learning efficiency. In fact this is a non-issue: LTP can only be applied to active connections, and a connection only becomes active through the depression of connections stronger than itself. Therefore, even with positive feedback, it is still entirely appropriate to use ϕ_1, the number of applications of negative feedback required for complete learning, as a measure of learning efficiency. LTP in this context is not so much a learning mechanism as an 'anti-forgetting' mechanism.
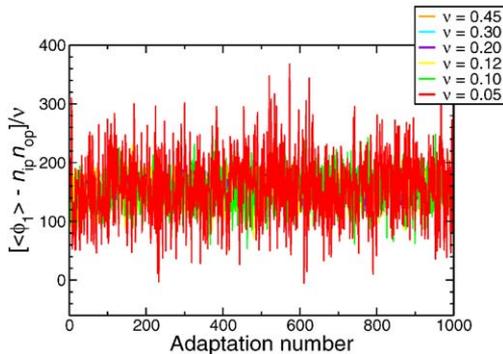
Fig. 2. Collapse of data in Fig. 1b for the flip-flop problem and unbounded LTP, according to the prediction of Eq. (1).

feedback all the while, the amount of potentiation given to its active outgoing connections will be proportional to n_{ip}. Thus, if the ratio v/δ is greater than $1/n_{\text{ip}}$, there will be a divergence between the strengths of the active and inactive connections, leading to the observed addiction. More generally, for any $v > 0$, it is possible to think of a rate of change slow enough that addiction will result.

Fig. 1b shows the results for a different problem, the 'flip-flop' problem. This time, the network starts by having to learn the map $1 \rightarrow 1$, $2 \rightarrow 2$, $3 \rightarrow 3,\ldots$, and, once this has been learned, the map required is switched to $1 \rightarrow n_{\text{ip}}$, $2 \rightarrow (n_{\text{ip}} - 1)$, $3 \rightarrow (n_{\text{ip}}-2),\ldots$; and we continue switching back and forth between these two inverse input–output maps. Again, the learning process consists of repeatedly running through the cycle of inputs until the complete cycle can be run through without error. The slowness of change of the input–output map is no longer an issue, since *all* of the input–output mappings are changed at each adaptation. Thus, as one should expect, the values of $\langle \phi_1 \rangle$ required to adapt remain bounded for all values of v. Nevertheless, performance is still observably worse in all cases than the $v = 0$ case with LTD only.

To explain this, consider again a single input. Once it has been correctly wired up to its associated output, it will have a window of time (while the rest of the network is still learning) in which its active outgoing synapses will be potentiated. The average for this time window will be controlled by the system size, i.e., a constant $f(n_{\text{ip}}, n_{\text{op}})$ for any given network, and the total divergence between active and inactive connections will be proportional to v; thus the gain in the amount of negative feedback required to adapt will be given by v/δ. Mathematically speaking, we have

$$\langle \phi_1 \rangle = n_{\text{ip}} n_{\text{op}} + \frac{v f(n_{\text{ip}}, n_{\text{op}})}{\delta} \tag{1}$$

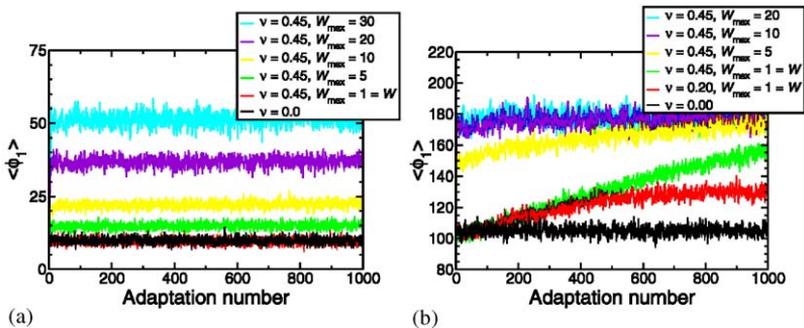which is confirmed by the data collapse achieved in Fig. 2.

Fig. 3. Adaptive ability when bounded LTP is included. (a) Slow change problem. Addiction is prevented but for $W_{max} > 1 = W$, adaptive ability is worse than in the LTD-only ($v = 0$) version. (b) Flip–flop problem. $\langle \phi_1 \rangle$ increases to a maximum value determined by v (see Fig. 1b). Data averaged over 128 realizations.

4. Bounded potentiation

A method to avoid addiction while maintaining potentiation was proposed by Parisi [10] with respect to the Hopfield neural network model. By placing an upper bound on synaptic strength, he was able to construct a 'memory which forgets' and thus avoid the state of total confusion observed if the network were overloaded. This can be easily applied to the minibrain model, requiring synaptic strengths to be bounded in the interval $(-\infty, W_{max}]$, with $W_{max} \geqslant W$. In this bounded case, should potentiation cause a synaptic weight w_{ij} to go over the limit, we simply reset it to W_{max}.

Fig. 3a shows the performance in the slow change problem of different networks with $v = 0.45$ and varying values of W_{max}, as compared to a network with $v = 0$. While the presence of the bound W_{max} has prevented the runaway addiction seen in Fig. 1a, there is no improvement over the simple LTD-only case. In general, the network performs worse, and it is only with $W_{max} = W$ that the network matches the performance of the standard LTD-only minibrain. This is natural when one thinks that inactive connections will always have strength of o(W) whereas active connections will have strength o(W_{max}).

By contrast a curious behaviour is observed in the 'flip-flop' problem of switching back and forth between two different (non-overlapping) input–output maps (Fig. 3b). In all cases the value of $\langle \phi_1 \rangle$ increases with successive adaptations to a maximum value controlled, not by W_{max}, but by v (see also Fig. 1b). How can we explain this? Each time the input–output map switches, the system must suppress the active connections and then search among all connections for the required correct outputs. What makes this different from the slow change problem is that here the majority of possible input–output connections are *always* incorrect. Therefore, the majority of connections will be continually weakened, never strengthened, and the gap between the average synapse strength and the maximum synapse strength W_{max} will diverge. This is equivalent to continually increasing the value of W_{max}, meaning that in the long run the system will
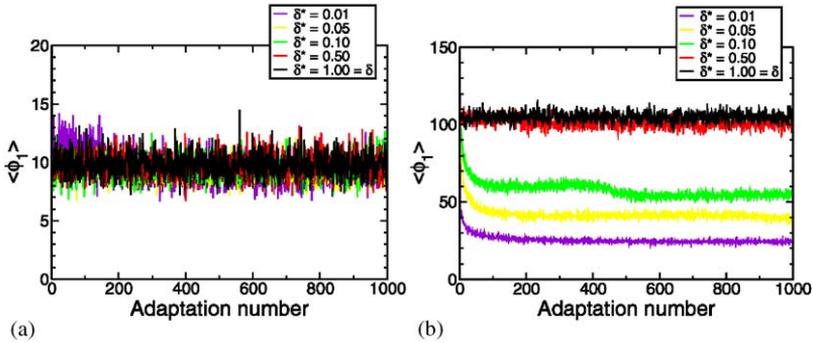
Fig. 4. Adaptive ability with LTD only, but with selective punishment of previously successful connections. (a) Slow change problem. No significant difference is observed for different values of δ^*. (b) Flip-flop problem. When lesser rates of punishment are applied to 'good' connections, $\langle\phi_1\rangle$ *decreases* with successive adaptations: the system has a memory of previously good responses. Data averaged over 128 realizations.

behave as if this limit does not exist, reproducing the behaviour observed in the case of *unbounded* LTP.

5. Selective punishment

Finally, let us consider the case of *selective punishment*. Here a synapse involved in a successful decision becomes permanently marked as 'good'; should it later be punished, it is by an amount in the interval $[0, \delta^*]$ with $\delta^* < \delta$. Thus, a previously good connection that has been suppressed is easier to reactivate than a connection that has never been good.

As Fig. 4a shows, this makes no significant difference to the network's ability to adapt in the slow change problem.[3] Should n_{im} take a smaller value, the effects of the selective punishment become more pronounced and adaptation is initially slower, but this effect vanishes as the network gains 'good' connections to all possible outputs.

However, in the second problem—switching back and forth between two distinct input–output maps—selective punishment proves a considerable benefit (Fig. 4b). While for $\delta^* = 1.0 = \delta$ the value of $\langle\phi_1\rangle$ remains constant with a value of $n_{ip}n_{op}$, as one would expect, values of $\delta^* < \delta$ see a *decrease* in $\langle\phi_1\rangle$ with successive adaptations, towards a minimum value much lower than without the selective punishment. Recall that the selective punishment favours previously good connections over others. Since in this case the number of good responses for each input is small by comparison to the total number of responses, this has the effect of drastically cutting learning times.

---

[3] Different results are observed for different network topologies. For example, on a random network, selective punishment proves very effective at enabling the system to distinguish between those paths that go nowhere or terminate in endless loops, and those that actually lead to output neurons [8].

6. Conclusions

Perhaps the key result of the present work has been to observe that, in the setup considered here, strengthening of synapses *always* carries within it the potential for divergence between the strengths of active and inactive connections. This divergence is governed not merely by the level of potentiation but also by the system size, increasing as the network becomes larger.

The minibrain is a 'toy' model, but it is nevertheless instructive to consider it in the light of biological results. Both LTP and LTD are well-observed in biological neuronal systems but their precise functions remain unclear [11,12]. A variety of different points of view can be found in the literature, with a number of authors explicitly endorsing a selectionist picture of neuronal dynamics [13] where learning is by either elimination or depression of connections. Thinking along these lines one might want to seek other means of positive feedback than LTP, such as the 'synaptic forgiveness' proposed by Klemm et al. [14].

Other authors have suggested that learning may result from a balance of LTP and LTD with a global feedback mechanism to prevent runaway strengthening or weakening of synapses [15]. A modification to the minibrain along these lines has recently been proposed by Bosman et al. [16]. The present results suggest that such a global mechanism may not just be useful, but *vital*, if LTP is to be an effective part of learning.

Acknowledgements

References

[1] A.G. Barto, R.S. Sutton, C.W. Anderson, IEEE Trans. Systems Man Cybern. 13 (1983) 834–846.
[2] A.G. Barto, Hum. Neurobiol. 4 (1985) 229–256.
[3] P. Mazzoni, R.A. Andersen, M.I. Jordan, Proc. Natl. Acad. Sci. USA 88 (1991) 4433–4437.
[4] P. Alstrøm, D. Stassinopoulos, Phys. Rev. E 51 (1995) 5027–5032.
[5] D. Stassinopoulos, P. Bak, Phys. Rev. E 51 (1995) 5033–5039.
[6] D.R. Chialvo, P. Bak, Neuroscience 90 (1999) 1137–1148.
[7] J. Wakeling, P. Bak, Phys. Rev. E 64 (2001) 051920.
[8] P. Bak, D.R. Chialvo, Phys. Rev. E 63 (2001) 031912.
[9] J. Wakeling, Physica A 325 (2003) 561–569.

[10] G. Parisi, J. Phys. A 19 (1986) L617–L620.

[11] C.F. Stevens, Neuron 20 (1998) 1–2.

[12] S.J. Martin, P.D. Grimwood, R.G.M. Morris, Annu. Rev. Neurosci. 23 (2000) 649–711.

[13] O. Sporns, G. Tononi, Selectionism and the Brain, Academic Press, New York, 1994.

[14] K. Klemm, S. Bornholdt, H.G. Schuster, Phys. Rev. Lett. 84 (2000) 3013–3016.

[15] S. Royer, D. Paré, Nature 422 (2003) 518–522.

[16] R.J.C. Bosman, W.A. van Leeuwen, B. Wemmenhove, Neural Networks 17 (2004) 29–36.