Opinion TRENDS in Neurosciences Vol.28 No.2 February 2005

Memory retention – the synaptic stability versus plasticity dilemma

Wickliffe C. Abraham and Anthony Robins

Departments of Psychology and Computer Science, University of Otago, Dunedin, New Zealand

Memory maintenance is widely believed to involve long-term retention of the synaptic weights that are set within relevant neural circuits during learning. However, despite recent exciting technical advances, it has not yet proved possible to confirm experimentally this intuitively appealing hypothesis. Artificial neural networks offer an alternative methodology as they permit continuous monitoring of individual connection weights during learning and retention. In such models, ongoing alterations in connection weights are required if a network is to retain previously stored material while learning new information. Thus, the duration of synaptic change does not necessarily define the persistence of a memory; rather, it is likely that a regulated balance of synaptic stability and synaptic plasticity is required for optimal memory retention in real neuronal circuits.

Introduction

In Hebb's theory [1], the neurobiological basis of any psychological event is the pattern of neuronal activity within cell assemblies. To remember a previously experienced event, its activity pattern needs to be re-established within the same cell assemblies. The encoding that allows long-term memory for an activity pattern is widely assumed to be the setting of synaptic connection weights, which we call here a synaptic configuration. But does this synaptic configuration need to be preserved for the memory to be retained? We suggest that this relationship can hold only for networks that do not participate in further learning. In dynamic learning systems, the continued ability to establish new synaptic configurations is fundamental for minimizing memory loss.

Competing hypotheses for the synaptic basis of memory retention

When considering the likely neural mechanism for the setting of connection weights during storage of long-term memory (LTM), most authors turn to phenomena such as long-term potentiation (LTP) and long-term depression (LTD) [2]. Here we focus on LTP, which exhibits an array of features (e.g. rapid induction, activity-dependence, input specificity, associativity and persistence) that are analogous to the behavioral properties of associative LTM. One reason for the early and continuing interest in LTP is its unusual persistence [3,4], founded on the intuitive rationale that the *sine qua non* feature of a memory storage mechanism should be an ability to persist as long as its referent memory [5]. We refer to this commonly, if perhaps implicitly, held position as the 'synaptic stability hypothesis' for LTM. A strong version of this hypothesis assumes a one-to-one correspondence between a memory activity pattern and a fixed synaptic weight configuration in the circuits that encode it (Figure 1a,i). A more realistic interpretation would allow for some weakening and strengthening of weights over time (Figure 1a,ii), to account for the effects of passive decay and homeostatic adjustments that are nonetheless recoverable through rehearsal or additional learning episodes. However, a fundamental corollary of the stability hypothesis is that when synaptic weights change significantly or revert to baseline, the stored information becomes degraded and memory performance is impaired in line with changes to the functionality of the neural circuit.

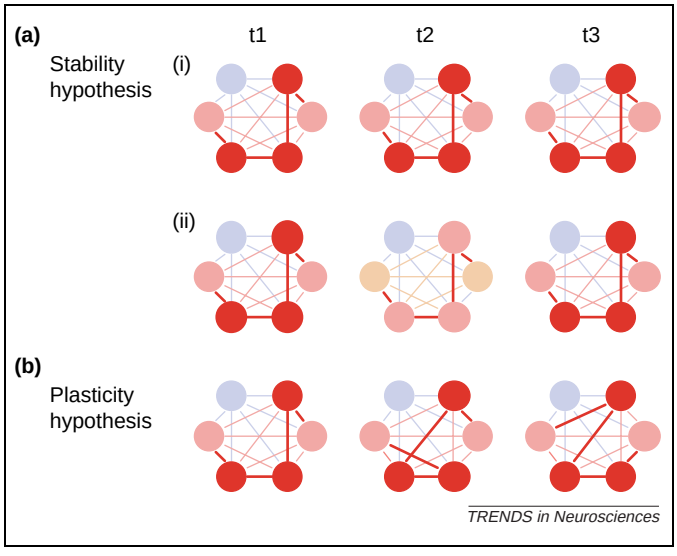

Figure 1. Three distinct models describing the stability or plasticity of synaptic connection weights underlying a stored activity pattern. Neural units are represented by circles and their connections by lines. 'Hotter' colors represent more active units or (with thicker lines) stronger connection weights between units. Encoding of a memory occurs at time t1, and connection weights are modified (i.e. a synaptic configuration is created) to generate the appropriate activity pattern. Retention is measured at two later times, t2 and t3. **(a)** (i) In the strong version of the synaptic stability hypothesis, weights encoding a memory remain permanently fixed for times t1, t2 and t3. (ii) In a weaker version of the hypothesis, the encoding weights vary in strength (e.g. weaken at time t2) but when this happens memory retrieval is impaired. Full memory retrieval requires the original weight configuration to be restored, as in t3. **(b)** In the synaptic plasticity hypothesis, weights encoding a memory can change but the memory is retained as long as the functionality (i.e. the ability to recreate the activity pattern) is preserved.

Corresponding author: Abraham, W.C. (cabraham@psy.otago.ac.nz).
Available online 22 December 2004

 0166-2236/$ - see front matter doi:10.1016/j.tins.2004.12.003

Despite the intuitive appeal of the stability hypothesis, it is conceivable – especially given the well-known capacity that synapses have for plasticity in structure and function [6–11] – that memories can outlast the lifetime of the specific synaptic configurations that originally encoded them. In this way of thinking, a given synaptic configuration does not uniquely define a stored memory [12]; rather, ongoing synaptic plasticity is in fact vital for memory retention. We refer to this general conceptual approach as the 'synaptic plasticity hypothesis' for memory retention (Figure 1b).

In the following sections we will review recent experimental and modeling research on the stability of connection weights, and re-evaluate the basic premise that the persistence of LTM needs to equate with the persistence of a synaptic weight configuration. We propose that memory retention actually requires a regulated balance between stability and plasticity to solve the trade-off between the stability required to retain information and the plasticity required for new learning within the network – a dilemma faced by all learning systems [13].

The synaptic stability hypothesis

As already mentioned, the conventional wisdom underpinning experimental models ranging from simple invertebrates to higher mammals is that LTM storage and retention entail stable, activity-dependent changes in synaptic transmission within the relevant neural circuits. Is it possible to confirm this hypothesis experimentally? Good progress has been made in invertebrate systems, such as sensorimotor synapses of the sea slug *Aplysia*, where the close relationship of synaptic morphology, physiology and behavior has permitted demonstration that facilitation and depression at a population of these synapses closely parallels the corresponding learned reflex changes [14–16]. Nonetheless, the time-course of changes at individual synapses remains an enigma even for this simple neural circuit.

The mammalian hippocampus offers a more complex system, and here it has not yet been possible to identify the precise role played by any subset of synapses in a particular behavioral response. Nonetheless, the fact that LTP in the hippocampus can reliably persist for days or weeks has provided sufficient grounds for hypothesizing that LTP is a mechanism underlying hippocampus-dependent memories [3,5]. In most studies LTP eventually declines to baseline [17–19] and, although this fact has worried some theorists about its relationship to LTM generally [20], it is consistent with theories that the hippocampus has a time-limited role in memory consolidation [21–23]. This begs the question, however, whether LTP can last long enough to support very long-term memory (vLTM) as stored in either neocortical or subcortical brain regions [23,24]. Indeed, combined imaging and recording from single synapses in organotypic hippocampal slice cultures has revealed a set of large synapses that appear resistant to LTP induction, a property interpreted to be ideal for LTM [25]. Recently, we have shown *in vivo* that LTP in the rat hippocampal dentate gyrus does in fact have the capacity to last stably for many months or longer [26,27] (Figure 2a), extending a previous report of stable LTP lasting weeks in area CA1 [28]. Thus, even hippocampal LTP appears to have the capacity for the extreme persistence necessary to store hippocampus-dependent vLTM [29] and, if this property is shared by neocortical and other synapses, a key premise of the synaptic stability hypothesis would be experimentally confirmed.

Recent studies utilizing *in vivo* imaging of spine morphology now offer a second line of evidence in support of the synaptic stability hypothesis. Two-photon laser-scanning imaging in transgenic mice expressing fluorescent proteins in principal neurons has demonstrated remarkable stability of a considerable proportion of spines in neocortex. In the barrel cortex of adult mice, ~50% of the imaged spines were shown to be stably retained for at least a month, with a mean lifetime calculated to be ~120 days [30]. Large mushroom-shaped spines were the most stable. Even more strikingly, an estimated 96% of layer 1 and layer 2 spines of adult visual cortical pyramidal cells were retained over a one-month interval, and spines overall persisted with a half-life >13 months [31] (Figure 2c). These findings, together with evidence for stable synaptic density, dendritic arborization and axonal projections [30,31], are indicative of a general long-term stability of cortical wiring in adults, but leave open the question of the functional stability of the imaged synapses.

If long-term memories are stored as stable synaptic weight changes, it follows that disrupting those configurations should cause loss of memories previously formed and stored in that neural circuit. Indeed, retrograde amnesia for spatial memory can be evoked by induction of saturated LTP in the dentate gyrus [32]. Interestingly, the converse is also true. Exposure of animals to a novel or enriched environment, and the presumed associated new learning, can cause a persistent reversal of stable LTP that had been previously induced [26,33,34] (Figure 2b). Taken together, these data support the synaptic stability hypothesis that memory storage and retention involve enduring changes in transmission at specific synapses.

The synaptic plasticity hypothesis

The synaptic stability hypothesis implies a kind of synaptic phrenology, in the sense that memory activity patterns are encoded by fixed synaptic weight configurations in specific neural circuits. However, it is certainly conceivable that information storage could outlast specific synaptic weight changes. For example, reactivation and rehearsal of previously acquired information could update what would otherwise be decaying synaptic weight changes [35,36] (c.f. Figure 1a,ii). This view nonetheless equates a specific configuration of synaptic weights with a specific memory. Alternatively, one common system theory for memory consolidation is that information moves from a temporary holding store to an anatomically separate permanent store [37–39]. If this is plausible, further shifts during LTM retention could also occur and the association between specific weights and a specific memory is broken down. The latter view implies that widespread synaptic plasticity is vital to memory retention. But is there evidence to support the synaptic plasticity hypothesis?

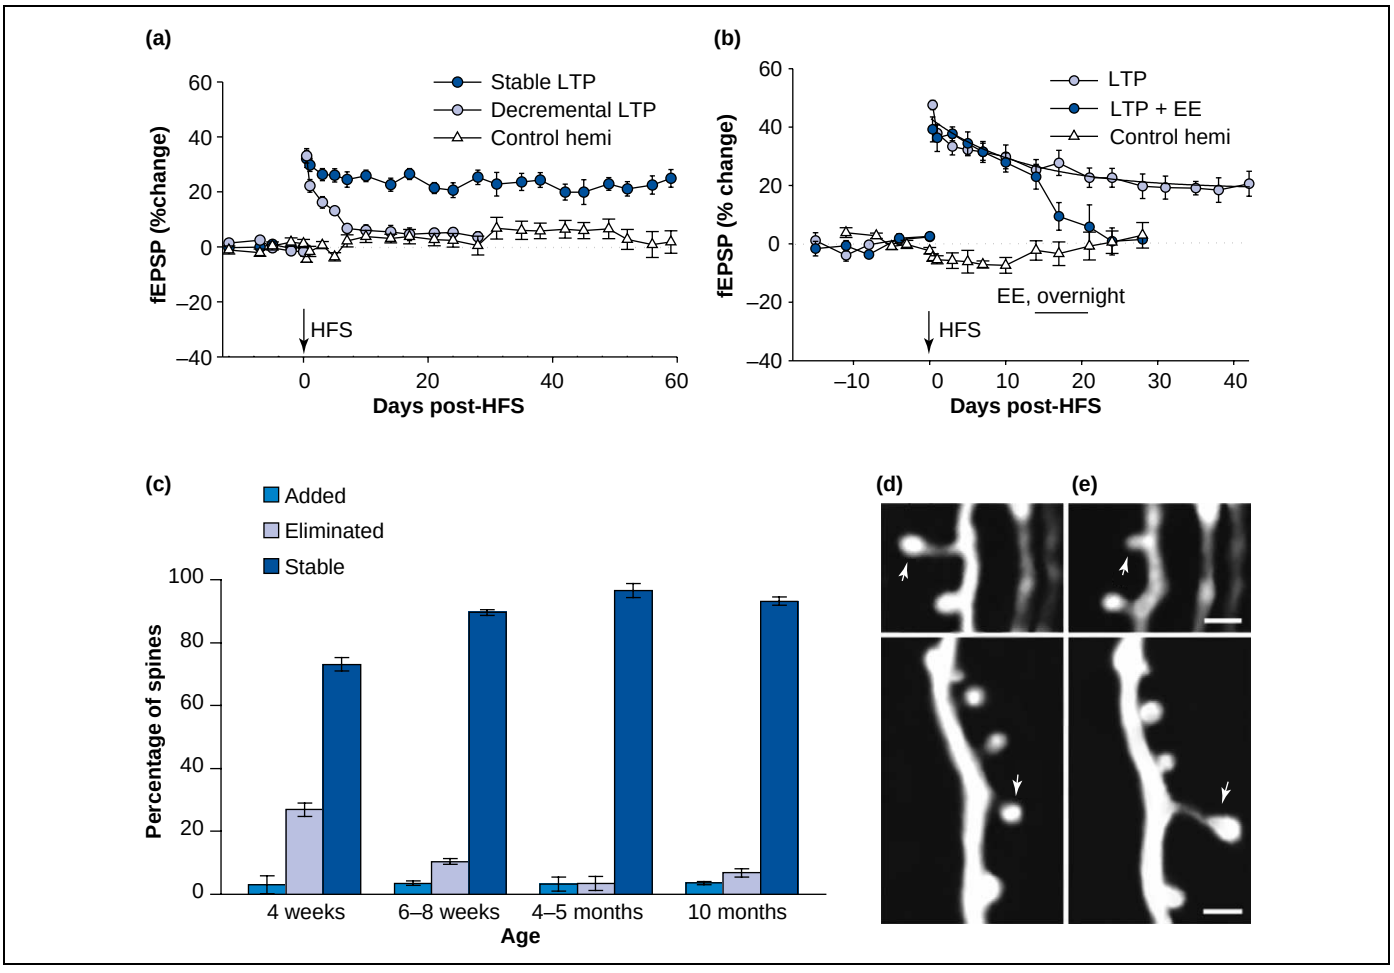

Figure 2. Electrophysiological and morphological examples of synaptic stability and plasticity in whole animals. **(a)** Both decremental and stable LTP lasting months can be induced in the dentate gyrus of awake freely moving rats. The persistence of LTP depends on the number of stimulus trains delivered, the pattern of train delivery and the degree of animal handling before and after LTP induction [26]. 'Control hemi' refers to recordings taken from non-tetanized hemispheres contralateral to the tetanized hemispheres. Abbreviations: fEPSP, field excitatory postsynaptic potential; HFS, high-frequency stimulation. **(b)** Stable LTP in the dentate gyrus is reversible if animals are given opportunities for new experiences through repeated overnight exposures to an enriched environment (EE) containing novel objects, a novel food and conspecific animals [26]. **(c)** Results from two-photon imaging of layer 1 and layer 2 visual cortical spines in transgenic mice expressing yellow fluorescent protein [31]. For spines imaged across a one-month interval, there is a remarkable stability in spine presence, regardless of the starting age, with the exception of one-month-old animals where spine elimination is somewhat more prominent. **(d,e)** Despite stability of their presence, there is ongoing plasticity in the shape of adult spines, suggesting possible functional changes [31]. Arrows point to adult spines repeatedly imaged before (d) and after (e) a one-month period. Scale bars, 1 μm. Panels (a,b) are reproduced, with permission, from Ref. [26] © (2002) the Society of Neuroscience. Panels (c–e) are reproduced, with permission, from Ref. [31].

One of the abiding features of consolidated memory that has been a fundamental element of the synaptic stability hypothesis is its imperviousness to amnestic treatments such as electroconvulsive shock, protein synthesis inhibitors, cooling and a variety of receptor agonists and antagonists. Yet it is well known that experimental reactivation of a memory renders it labile to disruption [40,41] or, alternatively, to further consolidation. This implies that reactivation can institute a new wave of synaptic plasticity. Thus, during the naturally occurring memory retrieval that can occur during related behaviors or sleep [42,43], it is possible that significant plasticity updates or reorganizes synaptic weights and thus retains the memory trace.

Elegant experiments using conditional knockout strains of transgenic mice have revealed that ongoing NMDA-receptor-mediated processes do indeed operate to consolidate and retain memories over weeks and months. Conditional knockout of the NR1 subunit during the first week following contextual fear conditioning or water-maze training caused retrograde amnesia when measured approximately two weeks post-training [44]. Even more strikingly, NR1 knockout for 30 days beginning six months post-training impaired the retention of cued and contextual fear conditioning measured at nine months [36]. These intriguing experiments provide strong evidence that both the consolidation of memory in the early post-training period and the stable maintenance of remote memory require continued neural activity and NMDA receptor activation. Insofar as this is indicative of ongoing synaptic plasticity, these findings appear to support the synaptic plasticity hypothesis for memory retention. However, they do not answer the question of whether the ongoing plasticity merely refreshes the weights for those synapses that were originally modified during initial memory storage [45] (cf. Figure 1a,ii), or whether there are more widespread changes in synaptic structure and function (cf. Figure 1b).

The long-term two-photon imaging experiments previously described, although providing evidence for

prolonged retention of a high percentage of spines and a stability of overall spine density in adult animals, also demonstrated a significant capacity for ongoing synaptic change. Time-lapse imaging in somatosensory cortex revealed that ~17% of spines had lifetimes <1 day, and 23% had a mean lifetime of 2–3 days [30], indicating significant spine turnover. Turnover increased substantially if cortical reorganization was induced by whisker trimming. In visual cortex, where 96% of spines had estimated half-lives of 13 months, there were nonetheless changes in length or diameter of a significant proportion of these spines over a one-month interval [31] (Figure 2d,e). Thus, even in these sensory cortical areas where one might expect considerable wiring stability under normal conditions, there appears to be a clear capacity for structural change that could underpin the synaptic plasticity hypothesis of memory retention. Unfortunately, although the capacity for ongoing synaptic change appears to be present, it has not yet proved possible to record from or image single synapses during memory storage and retention to document their functional state across time. Improvements in the imaging techniques could yet offer ways forward in addressing this key issue but, in the meantime, the plasticity–stability dilemma remains an experimentally unresolved issue for neurobiologists.

Insights from artificial neural networks

Artificial neural networks (ANNs) provide an alternative methodology for addressing the plasticity–stability issue. To the extent that these models capture the computationally significant aspects of the underlying biology, the weighted connections between units in ANNs serve the same function as the synaptic connections between neurons, with the advantage that they can be monitored continuously during learning and retention. In accord with neurobiological theories [1,12], memory retention in ANNs entails preservation of the functionality of the network such that the correct output pattern can be generated in response to its input cue.

Most ANN learning algorithms (e.g. backpropagation [46]) implement a concurrent learning process. Information (i.e. a set of activity patterns) is presented as a whole and processed in one session until a correct output pattern is given for each input. Training is then regarded as finished. After such training the stability of the learned activity patterns rests simply and directly on the stability of the weights. However, this is stability only in the trivial sense that the network remains forever 'frozen' in its current state and no further learning is attempted.

Any useful model of memory needs to account for sequential learning, where new information can be integrated with old information at any time. Standard models typically cannot cope with this task demand, being subject to a rapid loss or 'catastrophic forgetting' of previously stored information as new patterns are learned [47–49]. One effective means of modifying standard algorithms to allow sequential learning is to incorporate partial rehearsal of old patterns while learning new ones. Rehearsal was first explored by Ratcliff [50] and Murre [51]. Robins [52] described a particularly effective 'sweep' rehearsal regime, and showed that rehearsal can use internally generated approximations of patterns instead of using the actual old patterns themselves. Importantly, rehearsal was found to be effective because it maintains the function (input–output behavior) of the network everywhere except the local region of the function within which change is necessary to accommodate a new pattern. This preserves previously learned patterns while incorporating new ones, but the complexity of the task requires a high degree of flexibility and significant change in the weights of the network [49,53] (Box 1).

In short, ANN simulations demonstrate that preserving the contents of memory during ongoing learning requires an active maintenance process (e.g. via rehearsal), which in turn requires considerable flexibility in the connection weights. From these simulations (Box 1), we predict that specific memory patterns in the brain can be encoded by different synaptic weight configurations that are computed as required to accommodate new learning.

Synthesis of the synaptic stability and plasticity hypotheses

Memory storage in the mammalian brain is generally thought to involve synaptic plasticity in extensive neural networks, cortical and subcortical. Despite the size and complexity of the brain, it seems highly likely that an individual cell or network participates in storing many activity patterns. Because memories are formed sequentially rather than all at once, it is difficult to imagine how the synaptic stability hypothesis can by itself account for memory retention in real life unless a separate network were to be used for each memory. Although this could happen in regions such as the olfactory bulb, higher centers do not have this luxury. Indeed, as reviewed here, the neocortex clearly retains substantial capacity for ongoing synaptic plasticity, and ANNs must significantly adapt their connection weights to retain old information in the face of new learning. However, it is obvious that synapses cannot randomly reset their weights on any large scale and yet preserve information. Thus, there is a clear need for stability in the weights throughout the memory network during the intervals between learning or rehearsal episodes, coupled with a capacity for further plasticity to accommodate learning of new information. Previous models of cortical function have incorporated such stability–plasticity considerations either explicitly or implicitly. Hebb [1], for example, viewed activity within the cell assembly as the defining event for memory expression, not the specific configuration of synaptic weights between the individual neurons. Although not discussed by Hebb, it is certainly conceivable that there are multiple solutions for the pattern of synaptic weights that will reinstate the key pattern of activity within a cell assembly. This has been confirmed in the ANN modeling (Box 1). Other theories have emphasized the need for the neocortex to be a slow learner so that information can be incorporated into the existing structure without causing catastrophic loss of earlier learned information [22,37,54]. This implies the need for iterative processing, as is characteristic of ANNs, so that a new synaptic weight configuration can be resolved that retains both old and

Box 1. Memory retention and weight plasticity in ANN simulations

The first formal method for setting the weights of a network was Rosenblatt's 'perceptron convergence procedure' [55]. Minsky and Papet's [56] analysis of this algorithm, however, exposed such serious limitations that research in the newly emerging field of ANNs was largely stalled for decades. The re-emergence of ANNs in the late 1980s was triggered in part by the publication of the backpropagation algorithm [46]. Backpropagation addressed the limitations of Rosenblatt's original method by enabling learning in powerful multilayer networks. A typical example is shown in Figure Ia. Such networks can serve as classifiers, function approximators or content-addressable memories, and exhibit desirable properties such as good generalization and robustness to noise or damage. Today, backpropagation is the most widely studied, extended and practically applied ANN learning algorithm.

We can use a typical backpropagation ANN to illustrate the relationship between the retention of memory patterns and the plasticity of weights. The network described in Figure Ia was trained on a base population of patterns, and then subsequently on a sequence of new patterns. If the new patterns are trained using standard backpropagation, then we observe a sharp rise in error, that is catastrophic forgetting of the base population [47–50] (Figure Ib). Note that the weights, which change only a little to accommodate each most recently trained pattern, remain stable (Figure Ic). By contrast, learning new patterns using backpropagation with rehearsal maintains good retention of all patterns, whereas the underlying weights change significantly [53] (Figure Ib,c). In short, for this typical ANN, relative weight stability is associated with memory loss in the face of new learning, whereas weight plasticity is required for memory retention.

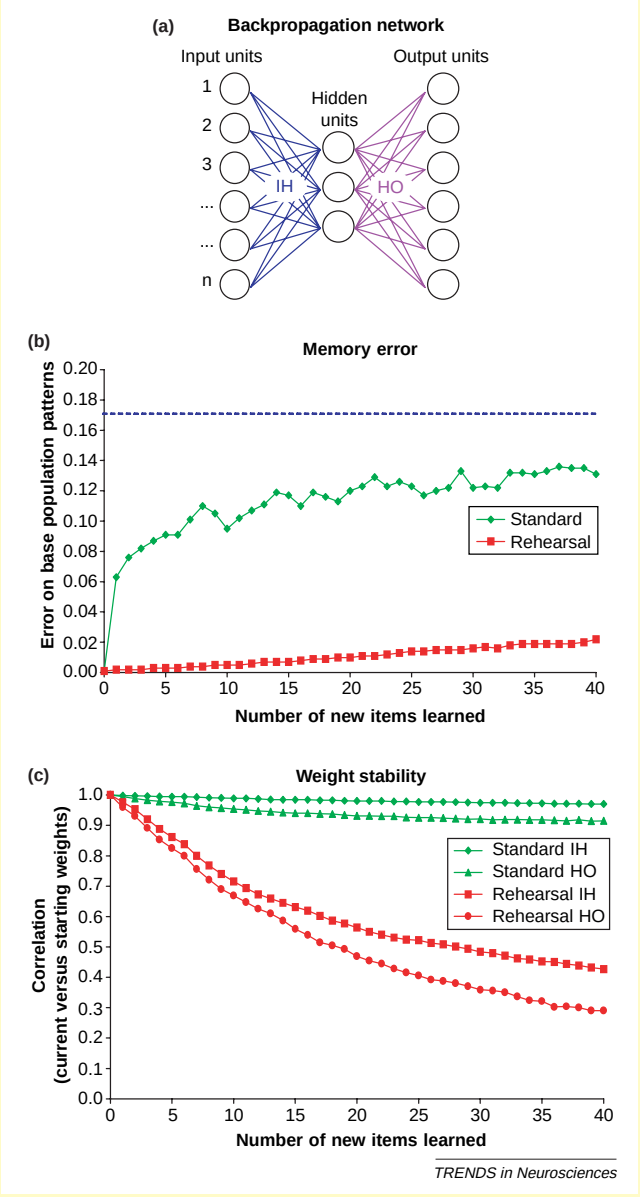

Figure I. Memory versus connection weight stability during sequential learning. **(a)** A typical backpropagation network consisting of input units, hidden units and output units. Units in each layer have weighted feedforward connections to units in the subsequent layer; that is there are input–hidden layer connections (IH) and hidden–output layer (HO) connections. The particular network used in this simulation had 32 input units, 16 hidden units and 32 output units (learning constant 0.3, momentum 0.5, and a normalized summed squared error measure). **(b)** The network was trained on a base population of ten activity patterns (i.e. randomly generated real valued inputs and associated outputs). From this starting point (0 on the *x*-axis), 40 new patterns were trained one at a time. New patterns were trained using either standard backpropagation methods or backpropagation with sweep rehearsal of old patterns. The graphs show the average error of the base population patterns as each new pattern was trained. A plot of the average error for the patterns learned previously in the sequence would be almost identical (i.e. there is no effective difference in performance for the different kinds of old pattern). The blue broken line represents random performance. **(c)** During this learning sequence, the weights of the network change from the starting point established by the base population. The graphs illustrate the stability of IH and HO weights (i.e. their correlation with the weights established after initial learning). Note the high correlation, that is weight stability, for the standard no-rehearsal method, for which catastrophic forgetting nonetheless occurs. Note that the connection weights changed dramatically (i.e. showed a sharply decreasing correlation) for the rehearsal paradigm that preserves old information. Thus, there are multiple synaptic weight configurations that can generate the required output activity pattern, and the network must compute a new one each time a new item is learned. Results shown in (b) and (c) are averages of ten runs of the simulation, using different random patterns for each run.

new information. If synaptic phrenology were true, and a network of synapses and neurons were indeed committed to only one memory, such iterative processing would be unnecessary.

Concluding remarks

There is growing evidence from empirical, theoretical and modeling approaches that memory retention does not require the maintenance of a specific configuration of synaptic weights for the lifetime of the memory. Rather, retention is likely to be an active process and memory-storing synapses must retain the capacity for ongoing plasticity if old information is to be preserved in the face of new learning. Experimental confirmation of this view, however, awaits techniques that can monitor over time both memory strength and the strengths of those synapses mediating memory storage.

Acknowledgements

We thank D. Ireland for assistance with the figures, and K. Harris, K. Kaila, D. Ireland and M. Eckert for comments on previous drafts of the manuscript. This review and related research was supported by grants

from the University of Otago and the Health Research Council of New Zealand.

References

1 Hebb, D.O. (1949) *The Organization of Behavior*, John Wiley & Sons
2 Martin, S.J. *et al*. (2000) Synaptic plasticity and memory: An evaluation of the hypothesis. *Annu. Rev. Neurosci.* 23, 649–711
3 Bliss, T.V.P. and Gardner-Medwin, A.R. (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 357–374
4 Bliss, T.V.P. and Lømo, T. (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356
5 Abraham, W.C. (2000) Persisting with LTP as a memory mechanism: Clues from variations in LTP maintenance. In *Neuronal Mechanisms of Memory Formation: Concepts of Long-Term Potentiation and Beyond* (Hölscher, C., ed.), pp. 37–57, Cambridge University Press
6 Fifkova, E. and van Harreveld, A. (1977) Long-lasting morphological changes in dendritic spines of dentate granular cells following stimulation of the entorhinal area. *J. Neurocytol.* 6, 211–230
7 Markram, H. and Tsodyks, M. (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382, 807–810
8 Harris, K.M. and Kater, S.B. (1994) Dendritic spines: Cellular compartmentalizations imparting both stability and flexibility to synaptic function. *Annu. Rev. Neurosci.* 17, 341–371
9 Maletic-Savatic, M. *et al*. (1999) Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science* 283, 1923–1927
10 Harris, K.M. *et al*. (2003) Structural changes at dendritic spine synapses during long-term potentiation. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358, 745–748
11 Kasai, H. *et al*. (2003) Structure–stability–function relationships of dendritic spines. *Trends Neurosci.* 26, 360–368
12 Dudai, Y. (1989) *The Neurobiology of Memory. Concepts, Findings, Trends*, Oxford University Press
13 Grossberg, S. (1987) Competitive learning: From interactive activation to adaptive resonance. *Cogn. Sci.* 11, 23–63
14 Castellucci, V.F. *et al*. (1978) Cellular analysis of long-term habituation of the gill-withdrawal reflex of *Aplysia californica*. *Science* 202, 1306–1308
15 Carew, T. *et al*. (1979) Sensitization in *Aplysia*: Restoration of transmission in synapses inactivated by long-term habituation. *Science* 205, 417–419
16 Bailey, C.H. and Chen, M. (1989) Time course of structural changes at identified sensory neuron synapses during long-term sensitization in *Aplysia*. *J. Neurosci.* 9, 1774–1780
17 Racine, R.J. *et al*. (1983) LTP phenomena in the rat limbic forebrain. *Brain Res.* 260, 217–231
18 Abraham, W.C. and Otani, S. (1991) Macromolecules and the maintenance of long-term potentiation. In *Kindling and Synaptic Plasticity* (Morrell, F., ed.), pp. 92–109, Birkhäuser
19 Abraham, W.C. *et al*. (1994) Immediate early gene expression associated with the persistence of heterosynaptic long-term depression in the hippocampus. *Proc. Natl. Acad. Sci. U. S. A.* 91, 10049–10053
20 Shors, T.J. and Matzel, L.D. (1997) Long-term potentiation: what's learning got to do with it? *Behav. Brain Sci.* 20, 597–655
21 Squire, L.R. (1986) Mechanisms of memory. *Science* 232, 1612–1619
22 McClelland, J.L. *et al*. (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psych. Rev.* 102, 419–457
23 Teng, E. and Squire, L.R. (1999) Memory for places learned long ago is intact after hippocampal damage. *Nature* 400, 675–677
24 Gale, G.D. *et al*. (2004) Role of the basolateral amygdala in the storage of fear memories across the adult lifetime of rats. *J. Neurosci.* 24, 3810–3815
25 Matsuzaki, M. *et al*. (2004) Structural basis of long-term potentiation in single dendritic spines. *Nature* 429, 761–766
26 Abraham, W.C. *et al*. (2002) Induction and experience-dependent reversal of stable LTP lasting months in the hippocampus. *J. Neurosci.* 22, 9626–9634
27 Abraham, W.C. (2003) How long will long-term potentiation last? *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358, 735–744
28 Staubli, U. and Lynch, G. (1987) Stable hippocampal long-term potentiation elicited by 'theta' pattern stimulation. *Brain Res.* 435, 227–234
29 Nadel, L. and Moscovitch, M. (1997) Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr. Opin. Neurobiol.* 7, 217–227
30 Trachtenberg, J.T. *et al*. (2002) Long-term *in vivo* imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* 420, 788–794
31 Grutzendler, J. *et al*. (2002) Long-term dendritic spine stability in the adult cortex. *Nature* 420, 812–816
32 Brun, V.H. *et al*. (2001) Retrograde amnesia for spatial memory induced by NMDA receptor-mediated long-term potentiation. *J. Neurosci.* 21, 356–362
33 Xu, L. *et al*. (1998) Spatial exploration induces a persistent reversal of long-term potentiation in rat hippocampus. *Nature* 394, 891–894
34 Manahan-Vaughan, D. and Braunewell, K.H. (1999) Novelty acquisition is associated with induction of long-term depression. *Proc. Natl. Acad. Sci. U. S. A.* 96, 8739–8744
35 Wittenberg, G.M. *et al*. (2002) Synaptic reentry reinforcement based network model for long-term memory consolidation. *Hippocampus* 12, 637–647
36 Cui, Z. *et al*. (2004) Inducible and reversible NR1 knockout reveals crucial role of the NMDA receptor in preserving remote memories in the brain. *Neuron* 41, 781–793
37 Marr, D. (1971) Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 262, 23–81
38 Bontempi, B. *et al*. (1999) Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature* 400, 671–675
39 Frankland, P.W. *et al*. (2001) α-CaMKII-dependent plasticity in the cortex is required for permanent memory. *Nature* 411, 309–313
40 Misanin, J.R. *et al*. (1968) Retrograde amnesia produced by electroconvulsive shock after reactivation of a consolidated memory trace. *Science* 160, 554–555
41 Nader, K. *et al*. (2000) Fear memories require protein synthesis in the amygdala for reconsolidation and retrieval. *Nature* 406, 722–726
42 Wilson, M.A. and McNaughton, B.L. (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679
43 Huber, R. *et al*. (2004) Local sleep and learning. *Nature* 430, 78–81
44 Shimizu, E. *et al*. (2000) NMDA receptor-dependent synaptic reinforcement as a crucial process for memory consolidation. *Science* 290, 1170–1174
45 Wittenberg, G.M. and Tsien, J.Z. (2002) An emerging molecular and cellular framework for memory processing by the hippocampus. *Trends Neurosci.* 25, 501–505
46 Rumelhart, D.E. *et al*. (1986) Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition (Vol. 1)* (Rumelhart, D.E., McClelland, J.L. and the PDP Research Group, eds) pp. 318–362, MIT Press
47 Sharkey, N. and Sharkey, A. (1995) An analysis of catastrophic interference. *Connection Sci.* 7, 301–330
48 French, R.N. (1999) Catastrophic forgetting in connectionist networks: causes, consequences and solutions. *Trends Cogn. Sci.* 3, 128–135
49 Robins, A. (2004) Sequential learning in neural networks: a review and a discussion of pseudorehearsal based methods. *Intell. Data Anal.* 8, 301–322
50 Ratcliff, R. (1990) Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychol. Rev.* 97, 285–308
51 Murre, J. (1992) *Learning and Categorization in Modular Neural Networks*, Earlbaum
52 Robins, A. (1995) Catastrophic forgetting, rehearsal, and pseudorehearsal. *Connection Sci.* 7, 123–146
53 Robins, A. and Frean, M. (1998) Local learning algorithms for sequential tasks in neural networks. *Adv. Comput. Intell.* 2, 221–227
54 Milner, P. (1989) A cell assembly theory of hippocampal amnesia. *Neuropsychologia* 27, 23–30
55 Rosenblatt, F. (1962) *Principles of Neurodynamics*, Spartan
56 Minsky, M. and Papert, S. (1969) *Perceptrons*, MIT Press