THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Computational modelling of memory retention from synapse to behaviour

Citation for published version:

van Rossum, MCW & Shippi, M 2013, 'Computational modelling of memory retention from synapse to behaviour', *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2013, no. 03, P03007.
<https://doi.org/10.1088/1742-5468/2013/03/P03007>

Digital Object Identifier (DOI):

[10.1088/1742-5468/2013/03/P03007](https://doi.org/10.1088/1742-5468/2013/03/P03007)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Statistical Mechanics: Theory and Experiment

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

TOPICAL REVIEW

Computational modelling of memory retention from synapse to behaviour

Mark C W van Rossum and Maria Shippi

Institute for Adaptive and Neural Computation

School of Informatics

University of Edinburgh

Edinburgh EH8 9AB, UK

E-mail: mvanross@inf.ed.ac.uk

Abstract. One of our most intriguing mental abilities is the capacity to store information and recall it from memory. Computational neuroscience has been influential in developing models and concepts of learning and memory. In this tutorial review we focus on the interplay between learning and forgetting. We discuss recent advances in the computational description of the learning and forgetting processes on synaptic, neuronal, and systems levels, as well as recent data that open up new challenges for statistical physicists.

Introduction

Memory provides the brain with one of its most fascinating aspects: its ability to acquire knowledge, learn new skills, and change the computations it performs. Biological memory has been studied extensively by philosophers, psychologists, experimental and computational neuroscientists. From this research it is well-known that our memory is not one giant general purpose system, but that several memory sub-systems exist, such as episodic, semantic, familiarity, procedural, and working memory [1]. Here we focus on long term memory for facts and events, so called explicit memory, which best corresponds to the common use of the term memory. There is not always an anatomical segregation of the sub-systems. To learn a given task, plasticity in multiple brain areas might be required, while a given brain region might support multiple memory systems.

Synaptic plasticity and memory

Learning, the process of acquiring memories, has in many cases been observed to be correlated with changes in the connection strength (synaptic weight) between neurons, known as *synaptic plasticity*. See [2] for a critical discussion of the link between synaptic plasticity and memory). Computational neuroscientists like to (over)simplify and often assume that learning has a 1-to-1 correspondence to synaptic plasticity. Although it is known that other properties of neurons, such as their input-output relation, e.g. [3], also change during learning, we also make that assumption here.

Given the correlation between synaptic changes and learning, an enormous research effort has been dedicated to long-term synaptic plasticity. The experimental discovery of Long-Term Potentiation (LTP) in 1973 by Bliss and Lomo is often regarded as the starting point of modern synaptic plasticity research [4]. In their experiments a stimulation and a recording electrode were placed in the hippocampus of a life rabbit. When a brief voltage pulse was given with the stimulation electrode, the resulting electric field caused substantial activity in nearby neurons. The recording electrode, placed away from the stimulating electrode, picked up activity of a population of neurons that received input via the synaptic connections from the stimulated population. When 200 pulses were applied with a pulse interval of 50ms (20Hz), an increase in the response of the downstream population was observed. Notably this change in the response strength persisted for hours, even days, after the induction, hence the name long term potentiation. Although this induction protocol was highly artificial and unlike naturally occurring neural activity, it did suggest that these long-lasting changes might be similar to processes involved in memory formation. Subsequent experiments discovered that synapses can also be weakened, so called Long-term Depression (LTD), using a low frequency (1Hz) stimulation.

In the decades that followed, revolutions in experimental techniques have significantly enriched our understanding of synaptic plasticity (as well as increased the complexity and diversity of the phenomena to almost unmanageable levels). For example, patch-clamp recordings have lead to the characterization of the effect of precise action potential timing on synaptic plasticity (see below). Genetic manipulations have revealed proteins crucial for plasticity, and allowed for imaging involved proteins. More recently it has become possible to image synaptic dynamics in vivo, highlighting that not only the strength of synapses changes, but that synapses are created and deleted even in adult animals [5, 6].

Hebb's hypothesis has been a common framework in the interpretation of these studies. It is a specific form of correlation learning that states that when a neuron participates in activating a target neuron, the connection should be strengthened. It is believed to be one of the important principles in plasticity (see [7] for a historic perspective).

Figure 1. Schematic diagram showing the 3 different levels for memory stability considered here: the synaptic level, the neural level, and the systems level. Below: Depending of the description level, the model of a typical stimulus and the memory strength measure differ, as will the determinants of the stability. (Synapse cartoon from [12]).

Computational models of memory and plasticity While some computational neuroscience models are highly accurate, the biophysical model of the action potential comes to mind, in the field of plasticity such accuracy is thus far unattainable. Rather plasticity models approach learning and memory from a number of different angles: A normative model one might optimize a learning rule with respect to a certain criterion, such as storage capacity or error on a task. A mechanistic model integrates the known underlying components and interactions to understand the larger system. Finally, in a phenomenological, descriptive approach one constructs compact models that best fit and unify highly diverse experimental data. In this tutorial review, we will see examples of all three approaches.

In the statistical physics community the Hopfield model of memory is the best known [8]. In this network model binary neurons are connected via a symmetric weight matrix, causing the network to always settle in one of a number of attractor states. The weight matrix can be set such that the attractor states of the network correspond to desired memory patterns. The attractor dynamics ensure that associative properties of this memory model. Statistical physics tools such as spinglas theory have been very fruitful in calculating the capacity of the network and stability of the attractor states [9]. However, also in other models tools from statistical physics are extremely useful.

A short bibliography for a neuroscience-naïve reader might be in order: Ref. [10] is a textbook on theoretical neuroscience. Ref. [11] provides a older but good account of neural networks from a computational point of view. More biological detail can be found in various textbooks.

Stability and memory

This review focusses on the dynamics of learning and forgetting and the determinants of memory retention on the synaptic, neural and systems level, Fig.1A. Although these levels are related, at each level there are typical approaches on how stimuli are constructed, and how memory strength is measured.

We consider on-line learning in which new inputs are continuously presented and learned by the system. This should be contrasted to the case in which the input data-set is limited and fixed, as happens in many practical applications of neural nets, but which is perhaps less relevant for biology. From a practical point of view, if the system is much older than any time-scale in the system and the input is stationary, we may assume that equilibrium has been reached, which simplifies analysis considerably and helps the use of tools from statistical physics. On-line learning highlights that any system has only a finite memory capacity; some information has to be forgotten

in order for new information to be learned [13, 14, 15, 16, 17, 18, 19]. As a result the memory strength, the definition of which may vary with description level, becomes a dynamical quantity.

It is interesting in this context to note that the brain does not appear to have an active forgetting mechanism. ‡ Instead, it appears that forgetting is the result of biophysical decay processes, or the result of overwriting.

Stability on the neural level

We start with the description of learning and forgetting on the level of single neurons and small networks, Fig.1 (middle column). This is the best studied level of description in computational neuroscience, and the most amenable to methods from physics [11]. The neural activities are modelled as continuous variables, corresponding to the neuron's firing rate averaged over an unspecified time interval. The synaptic strength is described as a single continuous variable, also known as the synaptic weight, labelled w_i . The synaptic change Δw_i at synapse i is typically modelled as function of the pre-synaptic input x_i at that synapse, the post-synaptic output activity y and the weight itself. A generic description of plasticity might thus be

$$\Delta w_i = f(x_i, y, w_i, w_j) \quad (1)$$

where $f()$ is some arbitrary function. The weights of other inputs, w_j (with $j \neq i$) are allowed to affect the plasticity as well, allowing for competitive mechanisms (see below). If we concentrate on the first bi-linear term that requires pre- and post-synaptic activity and does not depend on synaptic weight, we have

$$\Delta w_i = \epsilon x_i y \quad (2)$$

where ϵ is typically a small parameter called the learning rate, which determines the speed with which synapse is updated. Sometimes this equation is justified as being a Taylor-approximation of Eq. 1, but as there is no obvious small parameter, this is not really justified. Nevertheless, Eq. 2 highlights some basic properties and problems. The rule is Hebbian, that is, when both pre- and post-synaptic neurons are highly active, the synapse between them strengthens. Hebb put down this idea in his book as a way to create cell assemblies of neurons representing associations [20].

We consider the effect of this plasticity rule in an unsupervised scenario in which the goal of the plasticity is to learn the statistics of the input set, such as happens for instance during development of the visual cortex when the neurons in the visual cortex become selective to oriented bars [21]. In an unsupervised model, the post-synaptic activity is the result of all the inputs it receives. In a linear approximation one can write $y = \sum_i w_i x_i$. If the learning rate ϵ is small, one can simplify to $\frac{d\mathbf{w}(t)}{dt} = \epsilon Q \mathbf{w}(t)$, where $Q = \langle x_i x_j \rangle$ is the input correlation matrix and $\mathbf{w}$ is the weight written as a vector. As Q is a positive definite matrix, the weights will diverge, and moreover, the activity y will diverge or saturate, which is clearly unbiological. One might note that if x and y are firing rates, and thus positive, then always $\Delta w \geq 0$, hence the synapse never depresses. Including linear depression terms, however, does not solve the instability.

Various methods have been introduced to constrain the unlimited growth of the synapses: A simple solution is to impose hard bounds (or weight clipping) that constrain the weight, say, within the interval $[0, 1]$. More commonly, one fixes the length of the weight vector, e.g. $\sum_i w_i^2$, forcing competition between the weights. Alternatively, one can include weight dependence in the learning rule to prevent further strengthening of already large weights. A particular example of this is Oja's rule, in which the weights equilibrate to the normalized PCA vector of the input correlation matrix [22]. In Bienenstock-Cooper-Munro (BCM) theory the problem is solved by adjusting the plasticity threshold that determines whether a certain stimulus strength leads to

‡ Nor do computers: deletion typically means that it becomes allowed to overwrite memory blocks previously reserved for the deleted item.

Figure 2. The effect of the details of the plasticity rule on synaptic retention. Adapted from [26].

A. Left: The synaptic strengths of a neuron with 800 Poisson stimulated inputs after reaching equilibrium. With standard STDP without weight dependence the synaptic weights are close the minimal and maximal boundaries (set at 0 and 200 pS). Right: The decay of a selected synaptic weight after it has been artificially set to a high value (larger black dot left plot), due to random pre- and post-synaptic activity resulting in LTP and LTD. The decay is slow due to the bistable dynamics.

B. Left: In weight-dependent STDP the synaptic weights have similar values in equilibrium, but they have a central rather than bi-modal distribution. Right: In response to perturbation, the synaptic weight decays much more quickly back to equilibrium. (note difference in x-axis scale).

C. In a single layer network simulation with weight-dependent STDP, lateral inhibition stabilizes otherwise unstable receptive fields; the retention time increases as inhibition increases.

LTD or LTP. When post-synaptic activity is for instance too large, the threshold shifts slowly up so that some activation pattern lead to depression [23].

Importantly, the equilibrium weights will depend not only on the input statistics, but also strongly on the chosen regularisation method, as can be shown using fixed points and stability analysis [24]. Yet, despite its importance, there is little conclusive experimental data on this issue. For instance, evidence for synaptic competition has been sporadic. A practical problem is that the mechanism that prevents unlimited weight growth typically acts on a much slower time-scale than the plasticity, which hinders experimental characterization. For instance, synaptic homeostasis, which likely plays an role in setting the synaptic weight, is known to have a time-scale of about a day [25]. Similarly, the BCM threshold is thought to shift slowly. Apart from that plasticity of inhibitory connections (i.e. from inhibitory neurons) as well as onto inhibitory neurons might help prevent run-away plasticity, however, such plasticity is poorly characterized.

The learning rule does not only affect the equilibrium synaptic weights. The precise learning rule has also a profound effect on memory retention time. This was studied in the context of Spike Timing Dependent Plasticity (STDP), although similar observations have been made earlier [17]. In STDP the synaptic update is formulated in terms of the spike time of pre- and post-synaptic spikes: if a presynaptic spikes precedes a post-synaptic one within some 10ms the synapse is strengthened, if the reverse happens the synapse is weakened [27, 7]. Erasure of memory traces is prominent in STDP, as in its simplest form any pre/post spike pair will modify the synapse. In the simulation the synapses of a neuron were stimulated with Poisson inputs, and after a while, equilibrium established. To measure the synaptic retention time, one arbitrary synapse was set to a high value and the decay back to equilibrium was measured. Standard STDP without weight dependence except for hard bounds was compared to weight-dependent (soft-bound) STDP in which potentiation becomes gradually harder for large synapses. Using hard-bound STDP plasticity rules, the retention time was orders of magnitude larger than for soft-bound learning, Fig 2A,B. The results from the simulation are accurately described using diffusion theory and Kramer’s escape theory [26]. The intuition is that with hard-bound STDP the synaptic weight distribution is segregated by a large region of low probability, therefore it takes an exponentially long time for synaptic weights to equilibrate, while in the soft-bound case the net drift determines the equilibration time.

The situation gets dramatically more complicated when recurrent networks with plastic synapses are considered. As now the activity, which leads to plasticity, becomes dependent on the synaptic weights themselves. Only in a few cases analytical solutions are known [28]. The

network interactions not only change the synaptic weight equilibria, but can also alter the forgetting dynamics. In particular, lateral inhibitory connections, prevalent in cortex, can dramatically increase the synaptic retention time, Fig. 2C. This opens the possibility that memory retention is regulated by inhibition; it is not proven that such a phenomenon exists in biology (but see [29]). Furthermore, in a network setting it is less clear how information is to read out, using Hopfield networks would be one option. Thus, it is clear that much work remains to be done on the interaction between networks and plasticity.

State based models An alternative to models where the synaptic weight is a continuous variable, one can consider state-based models in which each synapse is described with a state-diagram and each state has an associated synaptic weight. The limited number of states immediately prevents any runaway plasticity. The simplest case is a binary synapse in which synapses are either on or off. Multiple states in these models can have the same synaptic weight emphasizing that to fully characterize the state of the synapse and its response to induction protocols, only knowing its weight might not be sufficient. Finally, state diagrams also provide an easy way to introduce phenomenological state-dependent transitions or interactions between synapses [30]. Interestingly, discretization of synaptic weight does not necessarily reduce the storage capacity of synapses much. In principle a continuous noise-free synapse can store infinitely many bits, but under ideal circumstances the maximum storage capacity in most models is between 0.1 and 1 bit/synapse for both continuous and binary synapses [31, 32].

To describe memory retention a Markov description can be used when potentiation and depression events occur randomly. For a binary synapse the memory will then decay exponentially. Once more states are introduced, the decay of the weight becomes a superposition of the exponentials decays, with decay times given by the eigenvalues of the Markov transition matrix. These transition matrices can be engineered to give power-law decay [33, 34]. This is relevant as the decay of human memories has been well fitted with a power law with a power close to one [35] (usual caveats about power-laws apply). However, it is not known whether the power law decay is attributable to the synaptic level.

Stability on the synaptic level

Stabilization and de-stabilization of memory is also linked to biophysical stability, Fig.1 (left column). Biophysically, LTP (and LTD) induction protocols lead to calcium entry into the synapse. The locally elevated calcium kicks off a complicated molecular cascade that finally, along with other changes including pre-synaptic ones, results in the insertion of more post-synaptic receptors in the cell membrane, so that a subsequent stimulus will yield a stronger response. This so called early phase of LTP will decay back to baseline in a few hours. Under the right circumstances the strengthened synapse stabilizes (late-phase LTP), a process that is dependent on protein synthesis [36]. Strong extra-cellular stimulation, as well as reward signals promote this process of *synaptic consolidation*. Moreover, late-LTP of a synapse 15 minutes before or after can stabilize a synapse that only received weak stimulation. According to the *synaptic tagging* hypothesis, the weakly stimulated synapse piggy backs on the protein synthesis triggered by the late-LTP induction at the other synapse [37].

The observed stability of LTP is particularly remarkable given that the post-synaptic receptors are continuously replaced by new ones; their lifetime is probably only a few hours [38]. This would seem an odd system to underlie synaptic modifications that are known to last for months [39]. Instead, it is thought that another substrate underlies biophysical stability. Biochemical cascades with positive feedback loops have been proposed to create bi-stable dynamics, which can be made very robust to fluctuations [40, 41]. Specific candidates for such a mechanism are the

auto-phosphorylation of the CaMKII molecule [42, 43], and sustained up-regulation of insertion of new AMPA receptors via PKM ζ [44]. Note that mechanisms based on bistable or multi-stable dynamics predict step-like rather than continuous changes in synaptic strength when inducing plasticity, for which some evidence has been found [45, 46]. On first sight it would also lead to discrete synaptic strengths, however, this does not necessarily follow if additional processes, such as homeostasis, co-determine synaptic strength.

The sub-cellular cascade that determines synaptic strength is complicated and contains hundreds of components. Yet more quantitative models of the sub-cellular cascade are now appearing [47]. The amount of LTP and its persistence in these models is in general non-linear in the number of induction pulses and their temporal spacing, Fig. 1 (middle). This implies that some electro-physiological protocol will be more efficient in inducing long lasting LTP than another despite having the same number of pulses [39, 48, 49]. It has been shown that this effect can contribute to the behavioural observation that learning with spaced repetition is more efficient than massed learning [50] (see also [51]).

Still most models on this level are however far from complete. This is partly due to a lack of knowledge of the biology and the many kinetic parameters involved. On the theoretical side there are also challenges, such as understanding of the role of the stochasticity - there are sometimes only dozens of copies of a certain protein in a synaptic spine [52], and understanding the effect of spatial inhomogeneity. Furthermore, the spine morphology, which is highly plastic as well, has been hypothesized to interact with the synaptic dynamics [53].

Stability on the systems level

A systems level description of learning and memory typically involves *in vivo* studies and tests memory retention with behavioural tasks possibly combined with lesions and/or pharmacological manipulations, Fig.1 (right column). The systems level description of learning and forgetting is the most challenging for computational models, as multiple brain regions interact, while knowledge about the underlying neural activity is limited. Moreover, it is complicated to interpret behaviour using computational models.

The state of the animal seems to play an important role in stabilization: In rats it was found that *in vivo* LTP persists longer when paired with either punishment or reward as compared to unpaired LTP [54]. Similarly, novel environments can both erase previously induced LTP, as well as increase the efficiency of subsequent LTP protocols [39, 55]. These effects appear to be partly mediated by dopamine, which is thought to signal reward and novelty throughout the brain by volume transmission. In humans strong emotional stimuli can boost the storage of irrelevant information (so called flash-bulb memory), which has been linked to the synaptic tagging described above [56].

Yet, it is important to realize that, in analogy with computers, stability of memory does not necessarily require that synapses are immutable, as memories can be moved to a different network [57]. A well-known example of the transfer of memory happens between hippocampus and neocortex, a process called *systems consolidation*. It is generally believed that new episodic memories are initially stored in the hippocampus, while subsequently, over a periods of days or weeks, memories are moved to cortex. The hippocampus is ideally placed as a temporary memory buffer: it receives input from many brain regions and it is highly plastic. The crucial role of the hippocampus has been shown in both humans and animal lesion studies. Patient H.M. is the best known for his antero-grade amnesia: after the removal of his hippocampi he lost recently formed memory as well as the ability to store new information. Strikingly, he retained old memories formed long before the operation. Animal lesion studies confirm this basic picture.

The duality of the cortical and hippocampal system is a subject of intense study and debate. It

has been argued, using data like H.M.'s, that the cortex is thought to code for life-long accumulated knowledge about the world. In contrast, the hippocampus acts as a temporary memory buffer. As the cortex learning rate is limited, too rapid learning in the cortex would lead to interference with pre-existing knowledge. The hippocampus replays the new information to allow for slow integration of new information [58]. However, recent experimental data suggest that the cortex can learn rapidly if organized relevant prior knowledge exists [59]. In the above view it is also likely that information is re-arranged when being transferred to the cortex and integrated with existing cortical information. As Grossberg highlighted with his stability-plasticity dilemma, it is highly problematic to simply merge all incoming data [60]. His Adaptive Resonance Theory solves this by a vigilance signal that triggers the recruitment of new nodes when sufficiently new situations are encountered.

A modern probabilistic view of the consolidation process is that the cortex codes the statistical prior knowledge about the world. Recent advances in machine learning show that algorithms such as Boltzmann machines [61], can build useful high level representations from data without supervision [62, 63]. However, the relation to biology is thus far obscure.

Finally, even when they are stored in the cortex, memories never become immutable. Interestingly, recall of consolidated memories can make them vulnerable to alterations and deletion. That is, applying drugs that block memory formation during recall of a memory stored previously, will erase that memory [64, 65], suggesting the re-opening of a window of plasticity to update cortical representation, so called re-consolidation.

Discussion and outlook

In this review we have considered memory storage and forgetting at the synaptic, neural, and systems level, and we have seen that there are important determinants of stability of memory at all these levels. Functionally, the task facing the brain might not be to remember as much as possible for as long as possible, but rather to decide what to remember and what to forget [37]. Support for this view comes from estimate of the typical learning rate of the order of 1 bit/s to yield about 10^9 bits of information of explicit memory for adults [66]. Meanwhile the world record memorizing digits of π , held by Akira Haraguchi, stands at 100,000 digits (41 kbytes). It takes about one year to learn such large arrays. This shows that the human memory-writing rate is very slow, much slower than the information rate from sensory input. It suggests that information is forgotten, unless deemed important enough to remember (likely by a sophisticated neural algorithm).

Computational models of memory and the above selection process might obviously inspire machine learning, but there might also be benefits for health. Numerous brain disorders, including Down syndrome, fragile X and Alzheimer's, are correlated to defects in synaptic plasticity. Although the causal link between the defect in plasticity and the disorder is often weak, a better understanding of plasticity might still help to develop cures or treatments. Oppositely, recently a number of genetically modified mice have been engineered that are smarter than regular mice on standard memory tests [67]. But even without resorting to genetics, plasticity models might help educators, or, god forbid, advertisers to optimize memory formation by changing the way information is presented and spaced. In contrast to enhanced memory, 'enhanced forgetting' would only seem a nuisance, but also here applications exist. Selective forgetting might help people with post traumatic stress disorder, as well as former drug addicts to fall back in old routines [68].

The rate of new experimental discoveries in neuroscience on this topic is high and shows no signs of slowing. For computational neuroscientists this is sometimes frustrating, but mostly it provides an exciting field of research, evoking memories of the golden era of physics.

Acknowledgments

MS was supported by the Erasmus Mundus Eurospin programme. We thank Richard Morris for discussion.

References

- [1] L. R. Squire. Memory systems of the brain: a brief history and current perspective. *Neurobiol Learn Mem*, 82(3):171–177, Nov 2004.
- [2] S. J. Martin and R. G. Morris. New life in an old idea: the synaptic plasticity and memory hypothesis revisited. *Hippocampus*, 12:609–36, 2002.
- [3] J. R. Moyer, L. T. Thompson, and J. F. Disterhoft. Trace Eyeblink Conditioning Increases CA1 Excitability in a Transient and Learning-Specific Manner. *J. Neurosci.*, 16:5536–5546, 1996.
- [4] T. Bliss and T. Lomo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path,. *J. Physiol.*, 232:331–56, 1973.
- [5] D. B. Chklovskii, B. W. Mel, and K. Svoboda. Cortical rewiring and information storage. *Nature*, 431(7010):782–788, Oct 2004.
- [6] A. Holtmaat and K. Svoboda. Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat Rev Neurosci*, 10(9):647–658, Sep 2009.
- [7] H. Markram, W. Gerstner, and P. J. Sjöström. A history of spike-timing-dependent plasticity. *Front Synaptic Neurosci*, 3:4, 2011.
- [8] J. J. Hopfield. Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proc. Natl. Acad. Sci.*, 79:2554–2558, 1982.
- [9] D. Amit. *Modeling brain function: The world of attractor neural networks*. Cambridge University Press, 1989.
- [10] P. Dayan and L. F. Abbott. *Theoretical Neuroscience*. MIT press, Cambridge, MA, 2002.
- [11] J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the theory of neural computation*. Perseus, Reading, MA, 1991.
- [12] M. Graupner and N. Brunel. Mechanisms of induction and maintenance of spike-timing dependent plasticity in biophysical synapse models. *Front Comput Neurosci*, 4, 2010.
- [13] M. Mézard, J. Nadal, and G. Toulouse. Solvable models of working memories. *J. Phys.*, 47:1457–1462, 1986.
- [14] G. Parisi. A memory which forgets. *J. Phys. A: Math. Gen*, 19:L617–M20, 1986.
- [15] J. Nadal, G. Toulouse, J. Changeux, and S. Dehaene. Networks of Formal Neurons and Memory Palimpsests. *Europhysics Letters (EPL)*, 1:535–542, 1986.
- [16] D. Amit and S. Fusi. Constraints on learning in dynamic synapses. *Network: Computation in Neural Systems*, 3(4):443–464, 1992.
- [17] D. Amit and S. Fusi. Learning in neural networks with material synapses. *Neural Computation*, 6(5):957–982, 1994.
- [18] S. Fusi. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biological Cybernetics*, 87(5):459–470, 2002.
- [19] M. van Rossum, M. Shippi, and A. Barrett. Soft-bound synaptic plasticity outperforms hard-bound plasticity. submitted PlosCB, 2012.
- [20] D. Hebb. *The organization of behavior*. New York: Wiley, 1949.
- [21] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. *Computational Maps in the Visual Cortex*. Springer, 2005.
- [22] E. Oja. A simplified neuron model as a principal component analyzer. *J. Math. Biol.*, 15:267–273, 1982.
- [23] E. L. Bienenstock, L. N. Cooper, and P. W. Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, 2:32–48, 1982.
- [24] K. D. Miller and D. J. C. MacKay. The role of constraints in Hebbian learning. *Neural Comp.*, 6:100–126, 1994.
- [25] G. G. Turrigiano, K. R. Leslie, N. S. Desai, L. C. Rutherford, and S. B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391:892–896, 1998.
- [26] G. Billings and M. C. W. van Rossum. Memory retention and spike-timing-dependent plasticity. *J Neurophysiol*, 101(6):2775–2788, Jun 2009.
- [27] A. Morrison, M. Diesmann, and W. Gerstner. Phenomenological models of synaptic plasticity based on spike timing. *Biol Cybern*, 98(6):459–478, Jun 2008.
- [28] M. Gilson, A. Burkitt, D. Grayden, D. Thomas, and J. van Hemmen. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. I. input selectivity–strengthening correlated input pathways. *Biological cybernetics*, 101(2):81–102, 2009.
- [29] F. Gambino and A. Holtmaat. Spike-timing-dependent potentiation of sensory surround in the somatosensory cortex is facilitated by deprivation-mediated disinhibition. *Neuron*, 75(3):490–502, Aug 2012.

- [30] A. B. Barrett, G. O. Billings, R. G. M. Morris, and M. C. W. van Rossum. State based model of long-term potentiation and synaptic tagging and capture. *PLoS Comput Biol*, 5(1):e1000259, Jan 2009.
- [31] C. Meunier and J.-P. Nadal. Sparsely coded neural networks. In M. A. Arbib, editor, *The handbook of Brain theory, 1st edition*. MIT press, Cambridge, MA, 1995.
- [32] A. B. Barrett and M. C. W. van Rossum. Optimal learning rules for discrete synapses. *PLoS Comput Biol*, 4(11):e1000230, Nov 2008.
- [33] S. Fusi, P. J. Drew, and L. F. Abbott. Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611, 2005.
- [34] A. Mehta and J. Luck. Power-law forgetting in synapses with metaplasticity. *Journal of Statistical Mechanics: Theory and Experiment*, 2011:P09025, 2011.
- [35] D. Rubin and A. Wenzel. One hundred years of forgetting: A quantitative description of retention. *Psychological Review; Psychological Review*, 103(4):734, 1996.
- [36] K. Reymann and J. Frey. The late maintenance of hippocampal LTP: Requirements, phases synaptic tagging, late-associativity and implications. *Neuropharmacology*, 52(1):24–40, 2007.
- [37] R. L. Redondo and R. G. M. Morris. Making memories last: the synaptic tagging and capture hypothesis. *Nat Rev Neurosci*, 12(1):17–30, Jan 2011.
- [38] R. Malinow and R. C. Malenka. AMPA receptor trafficking and synaptic plasticity. *Annu Rev Neurosci*, 25:103–126, 2002.
- [39] W. C. Abraham, B. Logan, J. M. Greenwood, and M. Dragunow. Induction and Experience-Dependent Consolidation of Stable Long-Term Potentiation Lasting Months in the Hippocampus. *J. Neurosci.*, 22:9626–9634, 2002.
- [40] F. Crick. Memory and molecular turnover. *Nature*, 312(5990):101, 1984.
- [41] A. Hayer and U. S. Bhalla. Molecular switches at the synapse emerge from receptor and kinase traffic. *PLoS Comput Biol*, 1(2):137–154, 2005.
- [42] J. Lisman, H. Schulman, and H. Cline. The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat Rev Neurosci*, 3(3):175–190, Mar 2002.
- [43] M. Graupner and N. Brunel. STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Comput Biol*, 3(11):e221, Nov 2007.
- [44] T. C. Sacktor. How does PKM ζ maintain long-term memory? *Nat Rev Neurosci*, 12(1):9–15, Jan 2011.
- [45] 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.*, 95:4732–4737, 1998.
- [46] D. H. O’Connor, G. M. Wittenberg, and S. S.-H. Wang. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc Natl Acad Sci U S A*, 102(27):9679–9684, Jul 2005.
- [47] J. H. Kotaleski and K. T. Blackwell. Modelling the molecular mechanisms of synaptic plasticity using systems biology approaches. *Nat Rev Neurosci*, 11(4):239–251, Apr 2010.
- [48] Q. Zhou, H. W. Tao, and M.-m. Poo. Reversal and stabilization of synaptic modifications in a developing visual system. *Science*, 300(5627):1953–1957, Jun 2003.
- [49] S. Ajay and U. Bhalla. A role for ERKII in synaptic pattern selectivity on the time-scale of minutes. *European Journal of Neuroscience*, 20(10):2671–2680, 2004.
- [50] Y. Zhang, R.-Y. Liu, G. A. Heberton, P. Smolen, D. A. Baxter, L. J. Cleary, and J. H. Byrne. Computational design of enhanced learning protocols. *Nat Neurosci*, 15(2):294–297, Feb 2012.
- [51] C. Beste, E. Wascher, O. G. Enten, and H. R. Dinse. Improvement and impairment of visually guided behavior through LTP- and LTD-like exposure-based visual learning. *Curr Biol*, 21(10):876–882, May 2011.
- [52] G. Antunes and E. De Schutter. A stochastic signaling network mediates the probabilistic induction of cerebellar long-term depression. *J Neurosci*, 32(27):9288–9300, Jul 2012.
- [53] C. O’Donnell, M. F. Nolan, and M. C. W. van Rossum. Dendritic spine dynamics regulate the long-term stability of synaptic plasticity. *J Neurosci*, 31(45):16142–16156, Nov 2011.
- [54] T. Seidenbecher, K. G. Reymann, and D. Balschun. A post-tetanic time window for the reinforcement of long-term potentiation by appetitive and aversive stimuli. *Proc Natl Acad Sci U S A*, 94(4):1494–1499, Feb 1997.
- [55] S. Li, W. K. Cullen, R. Anwyl, and M. J. Rowan. Dopamine-dependent facilitation of ltp induction in hippocampal ca1 by exposure to spatial novelty. *Nat Neurosci*, 6(5):526–531, May 2003.
- [56] D. Moncada and H. Viola. Induction of long-term memory by exposure to novelty requires protein synthesis: evidence for a behavioral tagging. *The Journal of neuroscience*, 27(28):7476–7481, 2007.
- [57] W. C. Abraham and A. Robins. Memory retention—the synaptic stability versus plasticity dilemma. *Trends Neurosci*, 28(2):73–78, Feb 2005.
- [58] J. McClelland, B. McNaughton, and R. O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995.
- [59] D. Tse, R. F. Langston, M. Kakeyama, I. Bethus, P. A. Spooner, E. R. Wood, M. P. Witter, and R. G. M. Morris. Schemas and memory consolidation. *Science*, 316(5821):76–82, Apr 2007.

- [60] S. Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cog. Sci.*, 11:237-63, 1987.
- [61] Hinton and Sejnowski. Learning and relearning in Boltzmann machines. In *PDP*. MIT Press, Cambridge, 1986.
- [62] S. Káli and P. Dayan. Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. *Nat Neurosci*, 7(3):286-294, Mar 2004.
- [63] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504-507, Jul 2006.
- [64] K. Nader, G. Schafe, and J. Le Doux. Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, 406(6797):722-726, 2000.
- [65] K. Nader and O. Hardt. A single standard for memory: the case for reconsolidation. *Nat Rev Neurosci*, 10(3):224-234, Mar 2009.
- [66] T. Landauer. How much do people remember? some estimates of the quantity of learned information in long-term memory. *Cognitive Science*, 10(4):477-493, 1986.
- [67] Y. Lee and A. Silva. The molecular and cellular biology of enhanced cognition. *Nature Reviews Neuroscience*, 10(2):126-140, 2009.
- [68] Y. Xue, Y. Luo, P. Wu, H. Shi, L. Xue, C. Chen, W. Zhu, Z. Ding, Y. Bao, J. Shi, et al. A memory retrieval-extinction procedure to prevent drug craving and relapse. *Science*, 336(6078):241-245, 2012.