A working memory model based on fast Hebbian learning

A Sandberg[1]**, J Tegnér**[2,3] **and A Lansner**[1]

[1] Department of Numerical Analysis and Computer Science, Royal Institute of Technology,
100 44 Stockholm, Sweden
[2] Department of Physics, Computational Biology Unit, Linköpings Institute of Technology,
581 83 Linköping, Sweden
[3] Stockholm Bioinformatics Center, SCFAB, Stockholm University, 106 91 Stockholm, Sweden

E-mail: asa@nada.kth.se, jespert@ifm.liu.se and ala@nada.kth.se

Abstract
Recent models of the oculomotor delayed response task have been based on the
assumption that working memory is stored as a persistent activity state (a 'bump'
state). The delay activity is maintained by a finely tuned synaptic weight
matrix producing a line attractor. Here we present an alternative hypothesis,
that fast Hebbian synaptic plasticity is the mechanism underlying working
memory. A computational model demonstrates a working memory function
that is more resistant to distractors and network inhomogeneity compared to
previous models, and that is also capable of storing multiple memories.

1. Introduction

The mechanisms behind the cognitive functions of the human brain remain enigmatic. Relevant
experimental data at the microscopic level of ion channels, synapses and neural activity as
well as the macroscopic level of psychophysics and cognitive psychology are currently rapidly
accumulating. Computational models provide a way to bridge the gap between these levels of
description and to integrate information from multiple sources into a coherent picture. A good
predictive model allows the experimenter to design maximally informative new experiments.

The memory systems of the brain are key players in cognitive functions. They exist in
several different forms and have been characterized along dimensions like episodic-semantic
and declarative-procedural. Here we focus specifically on working memory, the retention
or maintenance of information for short periods of time, usually linked with an ongoing
behavioural task. The paradigmatic experiment has been the delayed response task (see [18,
p 1]). In the oculomotor version of this task [15], the location in visual space has to be
remembered over a few seconds after which a suitable response should be generated.

For a long time, computational theories and models of memory processes addressing the cellular and network level have focused on long-term memory. The dominating paradigm has been the Hebbian cell assembly theory and its mathematical instantiation as attractor network memories [3, 25, 54]. More recently, working memory processes have attracted the attention of modellers [9, 12, 47, 51]. In contrast to network models of long-term memory, current working memory models rely mainly on reverberatory, localized persistent activity [2, 9, 18] in a network with fixed connectivity. These line attractor models originate from previous models of neural fields [1] and of the visual hypercolumn [5, 21]. They have been used to reproduce a variety of experimental findings from the prefrontal cortex (PFC) in the delayed response oculomotor task.

These models, however, do not take into account important characteristics of the underlying neuronal substrate like synaptic plasticity on the task relevant timescale 0.1–10 s [22, 23] and neuronal adaptation [33], which is known to destabilize line attractors [28]. Their operation is also sensitive to distractors unless additional stabilizing factors are included [9]. In addition, working memory is generally thought to be able to hold several items at the same time [32, 34]. It has been hard to find a robust mechanism for this in the context of persistent activity working memory models.

Here we investigate the alternative hypothesis that short-term Hebbian plasticity is a critical mechanism underlying the phenomenology in WM tasks. The presentation of a stimulus induces the formation of a corresponding attractor state which can later be read out as reverberatory activity. Persistent activity is still an integral part of this hypothesis but it now acts as an indicator of which stored memory is currently relevant and active. Short-term forms of memory based on fast synaptic plasticity have previously been suggested [50]. A unification of different memory mechanisms acting on a range of timescales is an attractive consequence of such a hypothesis.

Since the Hebbian property is crucial in attractor memory models our hypothesis suggests the existence of fast Hebbian synaptic plasticity in the underlying cortical memory networks, e.g. the PFC. The existence of such forms of synaptic plasticity has not yet been experimentally established, but remains an open possibility.

In the following, we investigate this hypothesis in the context of a recently described attractor network model capable of acting as a long-term as well as short-term palimpsest memory [43]. A phenomenological model of cellular firing rate adaptation and synaptic dynamics on a timescale of a few hundred milliseconds is also added in order to examine the stability or memory states and the storage of multiple memory items.

2. The network simulation model

Our simulation model is based on the Bayesian confidence propagation neural network (BCPNN) architecture [24, 29, 31, 43]. The BCPNN model is derived from a statistical view of learning and neural processing, where the activity of units correspond to estimates of the probability of various features. It employs a Bayesian–Hebbian learning rule that reinforces connections between simultaneously active units and weakens or makes connections inhibitory between anti-correlated units. It is here used as an attractor network model of a similar kind as the Hopfield net, but with a high degree of synaptic plasticity, i.e. very fast learning and forgetting.

The BCPNN learning rule has previously been used to set the weights in a cortical network model implemented with biologically detailed compartmental model neurons and with cortical minicolumns as its functional units [14]. That study suggests that it is reasonable to view the units of the BCPNN network as cortical minicolumns.

The learning rule produces both positive and negative weights. A connection from one minicolumn to another in the model lumps together all synaptic connections between the pyramidal cells in the respective minicolumns. Long-range synaptic connections are assumed to be excitatory but they may contact pyramidal cells in the receiving minicolumn either directly or via local inhibitory interneurons (e.g. bipolar cells) that in turn inhibit their companion pyramidal cells. In this manner, either a positive or a negative functional connection is realized. There is experimental support for plasticity in cortical synapses of both types involved in the suggested inhibitory pathway [17, 37].

To model cellular adaptation and synaptic depression/facilitation we use a simple phenomenological model. The adaptation is modelled as fast and temporary 'unlearning' implemented as an associative projection with negative gain and a short time constant [40, 41]. This results in a decrease of the effective synaptic weight between units active together as well as an increasing negative bias, thus removing valleys in the energy landscape while the network state remains in them. When units become inactive the projection returns to a resting value, thus restoring the original synaptic efficacy. This adaptation model was chosen due to its simplicity and symmetry with the learning rule used, and produces the same type of attractor-switching behaviour as other cellular adaptation [7, 30, 55] and synaptic depression based models [36, 48].

The BCPNN acts as a palimpsest memory where new information overwrites old and memories decay at a rate set by a learning time constant modulated by a print-now signal [43]. By temporarily up-regulating the print-now signal it is possible to imprint relevant stimuli while partially over-writing information already stored. While the print-now signal is zero no weight changes occur. We have hypothesized that the print-now signal could correspond to dopaminergic modulation of synaptic plasticity, which is well represented in e.g. the PFC and facilitates synaptic plasticity [8, 11, 35, 53].

This model allows us to simulate long-term as well as intermediate and short-term memories. In a short-term memory of this type the effective capacity is set by how strongly new information is imprinted, at the same time forcing old information to decay. With a print-now signal above a certain level, the memory becomes episodic, i.e. no repetition of the stimulus is necessary.

Details of the learning rule and network update equations used in this paper can be found in the appendix.

3. Set-up of the delayed oculomotor task

As in previous bump state models of working memory, the canonical experiment is based on the oculomotor delayed response task. In this task a monkey is trained to fixate on a central mark on a screen during a brief presentation of a peripheral cue. The gaze remains on the mark during a subsequent delay period until a signal is given for the animal to make a saccade to the cue position.

In our model the network was first subjected to a period of no input corresponding to the pre-trial period. This was followed by a 300 ms cue stimulus and a simultaneous print-now signal which caused an update of synaptic weights. After the cue the network spent 3 s in a delay period with no input and no print-now signals.

After the delay period a reset signal was given in the form of strong stimulation to all neurons in the network together with another print-now signal. While models with persistent activity only need to reset the neural activity itself [20, 26], in this model the reset signal is also assumed to be associated with a print-now signal erasing the changed synaptic weights. This print-now signal could be the one that imprints the next stimulus, acting as a gating signal [11].

The network consisted of 100 units fully connected to each other (this is similar to the set-up in [5]; the hypercolumn in this model is, in fact, analogous to several identical hypercolumns of the type used in [43] connected to each other). Each unit received input with different spatial tuning, corresponding to a projection from a population of location sensitive cells in the parietal lobe. For simplicity of display the units were ordered according to their favoured orientation. Each target angle θ corresponded to an input of the form

$$I_i = Z \exp\left(-\left|i - \frac{\theta N}{2\pi}\right|^2 \Big/ \sigma^2\right) \tag{1}$$

where Z is a normalization constant, N is the number of neurons and $\sigma = 10$. The distance is assumed to wrap around the population, producing a ring-shaped network metric. If not otherwise stated the input gain (g_I) was set to 1, the learning time constant (τ_L) to 7.2 s and the print-now signal to 1 or 0.

In these simulations the time constant of the adaptation projection τ_A has been given a value of 160 ms, corresponding to the decay rate of the action potential related Ca^{2+} pool that contributes to the accumulated after-hyperpolarization in a previous biophysically detailed pyramidal cell model [13]. In the initial experiments reported below $g_A = 0$, while in the subsequent experiments with adaptation in section 4.1 $g_A = 0.35$.

The noise input was Gaussian with mean 0 and variance 1 and a default gain (g_N) of 0.1. In the inhomogeneity experiments synaptic strengths and time constants (which were set individually for each synapse) were subjected to noise. The noise was uniformly rather than normally distributed in order to avoid negative time constants and synaptic sign reversals.

The model parameters have been collected in table A.1 in the appendix.

4. Results

The network was tested in the simulated delayed oculomotor task, both with a single and several cues to remember. Different distractors were added during the delay period and the tolerance to noise and parameter inhomogeneity was investigated.

The network exhibited one-shot learning of a bump shaped attractor state. The bump activity profile can, due to rotation symmetry, be viewed as a tuning curve. When exposed to a cue input consisting of a single bump for 300 ms, the network was able to sustain the bump throughout the delay period due to the change in local excitatory connections and long-range inhibition (figure 1). The increase in lateral inhibition stabilized the location of the bump while the increased local excitation induced a persistent activity state. The reset signal abolished the synaptic changes, thus dissolving the persistent activity state.

The activity of units not corresponding to the cue was reduced during the delay period relative to the spontaneous activity before the cue or after the reset signal.

As a rule, the bump state was resistant even to high levels of noise and did not drift, since its position was determined by the weight matrix and the neuron biases. Since it was the only attractor state of the network the activity profile resumed its shape and location if disturbed or if the activity was reset. Experiments with increasing levels of noise showed a broadening of the bump state, until the noise amplitude was so high that it dominated over the recurrent input. At the same time the population vector remained fixed to the peak of the original cue, i.e. there was no true drift.

An interesting phenomenon was the sharpening of the bump state relative to the cue signal during the delay period (figure 2(a)). This was due to the creation of a weight matrix with a nonlinear relationship to the original bump size. The learning process of attractor networks does not in general guarantee memory states identical to the input causing them, but it tends to
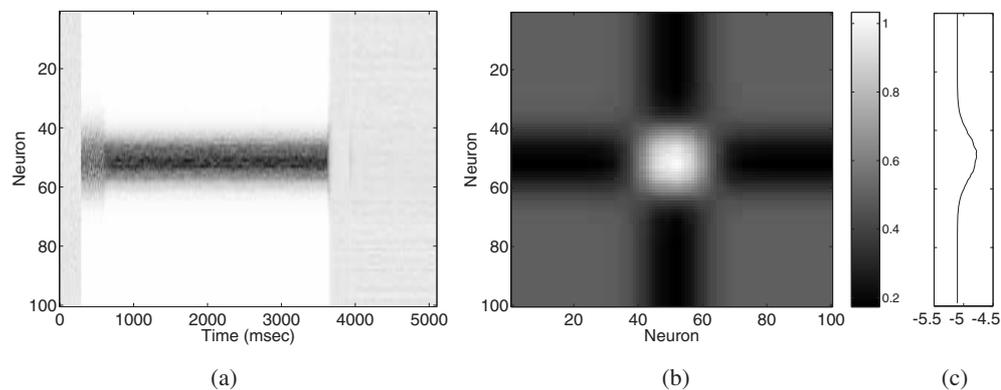
Figure 1. (a) Activity of network during a simulated oculomotor delayed response task. Activation of network units is indicated by grey-scale. The cue exposure caused a broad activation between 300 and 600 ms, followed by a persistent activity state until the reset signal at 3600 ms. (b) Weight matrix after the cue period. Bright areas correspond to excitatory interactions and dark areas to inhibition. The units near the target direction have formed a self-exciting group, with lateral inhibition to and from all other units. (c) Bias of the units. Most units have a very low bias, but those units close to the target direction have become more excitable.

create stable states in the vicinity of the input. The exact relationship between the shape of the original cue signal and the persistent activity pattern will in general depend on details of the learning rule used and its time constants; hence biological observations of these relationships can be used to support or rule out proposed network models.

If several patterns were presented in turn during the cue phase (each pattern for 300 ms), the last one remained active (figure 2(b)). If a sequence of targets were shown during the cue period the weight matrix developed into a band matrix with translation symmetry corresponding to a ring attractor family (there is an asymmetry caused by the aging of the memories, but for the current parameters and cue exposure times this is small). Each unit developed excitatory connections to units with similar spatial tuning and inhibitory connections to remote units (figure 2), similar to the rotation invariant pre-wired synaptic matrices used in previous bump state models.

These multiple bump states were metastable attractors: a sufficiently strong stimulus ('distractor') could shift the network state to one of the other stored patterns (figure 3). However, distractors had only temporary effects when a single target had been stored. For increasing input gain it became possible to temporarily shift the bump towards other positions, but when the input ended the network returned to the stored bump state. The shift appeared to be 'elastic' and non-linear: for small input gains the distance moved under the influence of external input was proportional to the gain, up to a critical level where the bump instead moved directly to the input direction. When several targets had been stored the distractor would cause a bump intermediate between the distractor location and the closest target; the bumps would be strongly attracted to the stored target. For a large number of targets the behaviour was similar to other line-attractor models, including fast virtual rotation between the original location and the stimulated location.

The network functioned without strong inter-trial interference when subjected to several task cycles (cue–delay–reset). It is worth noting that an explicit reset signal was not necessary if new cues were imprinted strongly enough; the new cue would erase the old stored information (figure 4).
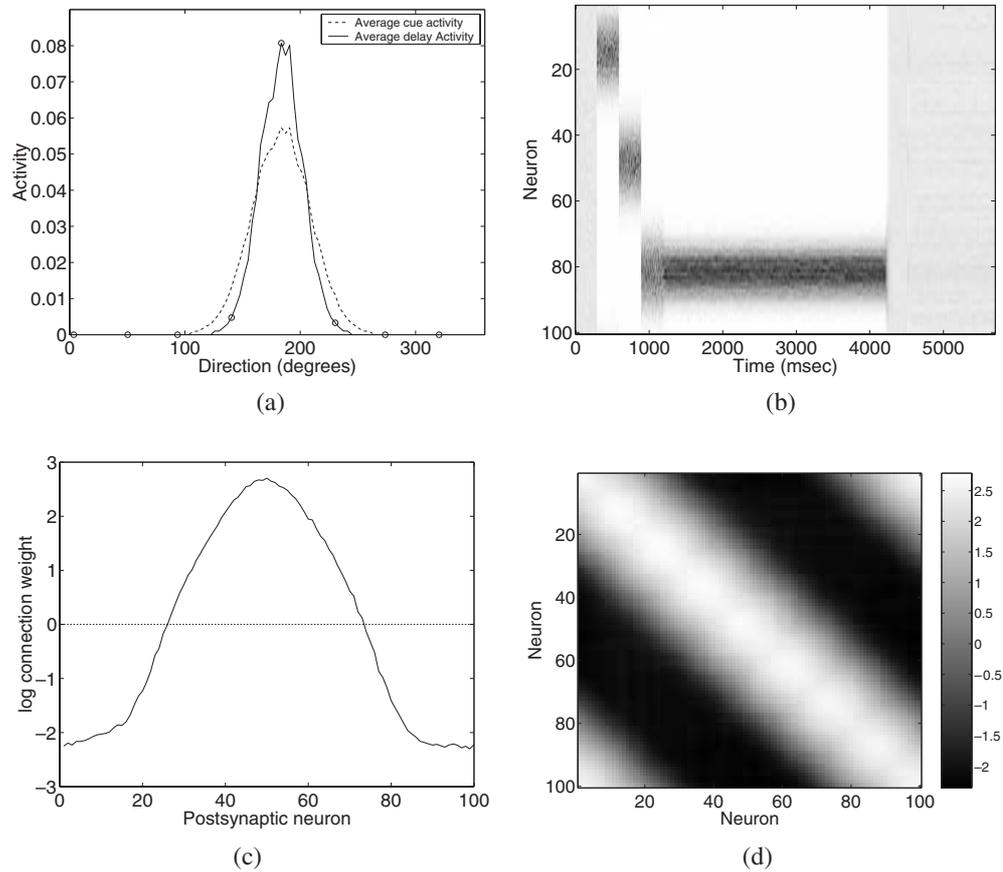
Figure 2. (a) Typical bump activity profile for the network. Mean activity during the cue and delay period for one trial. Eight equidistant points have been marked for comparison with [15, figure 9]. Shape fluctuations due to a noisy cue are imprinted in the delay activity, as can be seen near the top. (b) Activity when the network was exposed to three cues during the cue period. This created a synaptic matrix with three metastable memory states, each individually similar to the state in figure 1. The activity remained in the state corresponding to the last cue until being reset. (c) Plot of weights from a single neuron after learning eight bump attractors. (d) Weight matrix after learning eight targets.

The network was extremely resistant to network inhomogeneity, synaptic noise and sparse connectivity. Stable, if noisy, bump states persisted both for multiplicative synaptic noise where weights were multiplied by uniform random numbers between 0 and 7 (not shown in figure) and noise injected into the units. Similarly, the learning time constants could vary over a large range without the loss of bump states. Dilution of the connections caused a graceful degradation (figure 5).

The network was also resistant to selective removal of connections between units with distant receptive fields or random removal of connections in specific regions, although the loss of mutual inhibition enabled the co-existence of multiple bumps in different regions at sufficient dilution.

An interesting effect was seen when the print-now signal was not completely turned off during the delay phase, but kept on at 0.1 strength. Strongly activated neurons tended to link

(a) One target
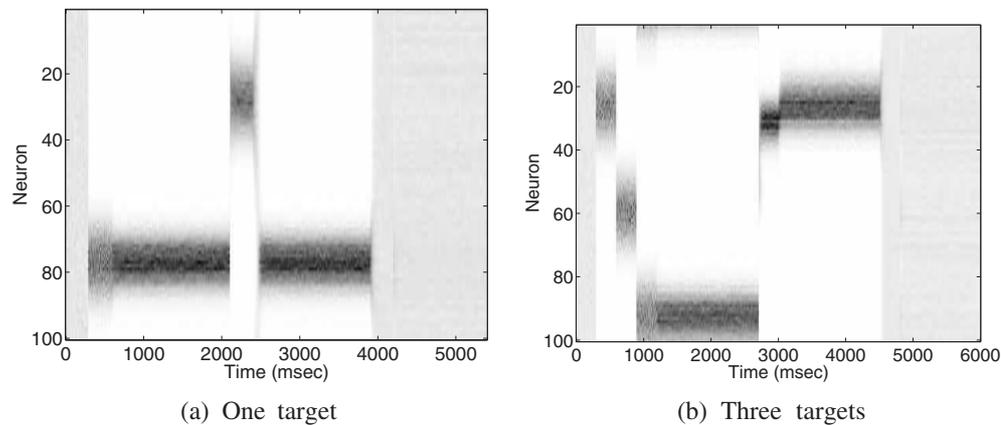
(b) Three targets

Figure 3. (a) If one pattern was stored, distractor input had no lasting effect: the activity returned to the stored target. (b) If several targets were stored sufficiently strong input allowed the network to shift between them. In the three target simulation the distractor occurred at position 40, but the resulting bump was strongly attracted to the memory state at position 25.
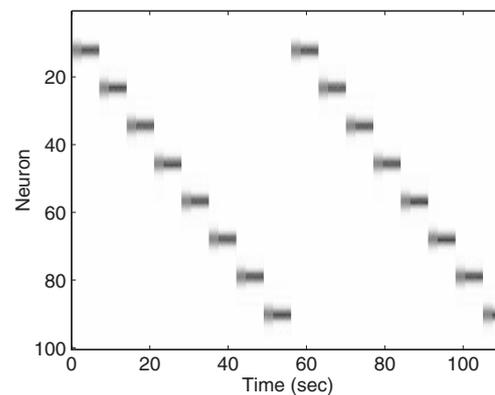
Figure 4. Multiple trial cycles. In each trial one cue is learned, remains sustained and is reset by the arrival of the next cue. The print-now signal is set to $\kappa(t) = 15$ during each cue phase.

to each other, creating an attractor growing narrower and sharper (figure 6). This produced gradual decay of some activities and increase of others, a ramping behaviour similar to what has been observed in unit recordings in working memory experiments [16, 39].

4.1. Adaptation

With adaptation present ($g_A > 0$) the stability of attractor states became time dependent.

For low values of adaptation the bump state remained stable and the system behaved as before. For strong adaptation the bump state instead became oscillatory, first decaying to a state where most neurons were weakly active and then returning to the bump shape again once the initially adapted neurons and synapses had regained their efficacy. This cycle then repeated itself.

For intermediate values of g_A, the network exhibited a stable bump when trained with one target, while moving between two or more bump states when trained with several targets

(a) Noise and dilution   (b) Noise and inhomogenity
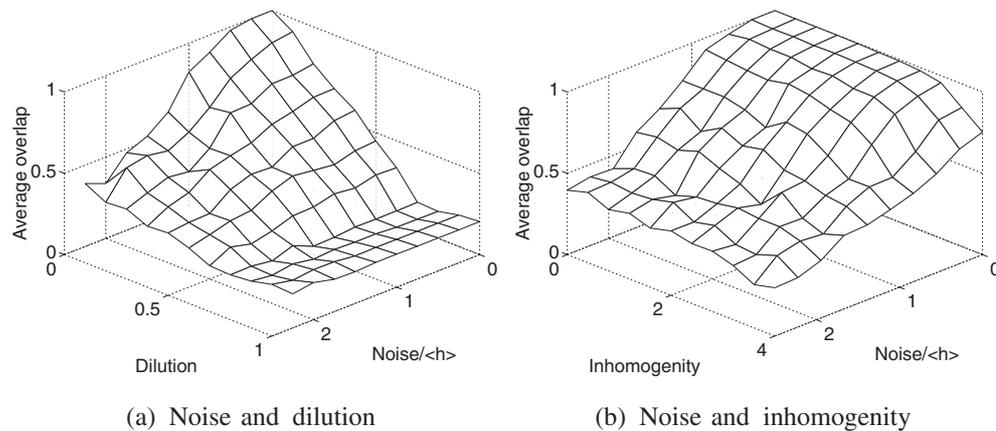
Figure 5. Network performance as a function of noise level, synaptic dilution (a) and standard deviation of time constants (b). The network was trained with eight attractors and placed in one of them. The average overlap between the network state and the original pattern in the period 800–1000 ms afterwards was plotted. Noise was measured relative to the average value of the unit potential h, dilution by the probability of removing a synapse. Time constants were uniformly distributed around 7.2 s. Average over 100 trials in each point.

(a)   (b)
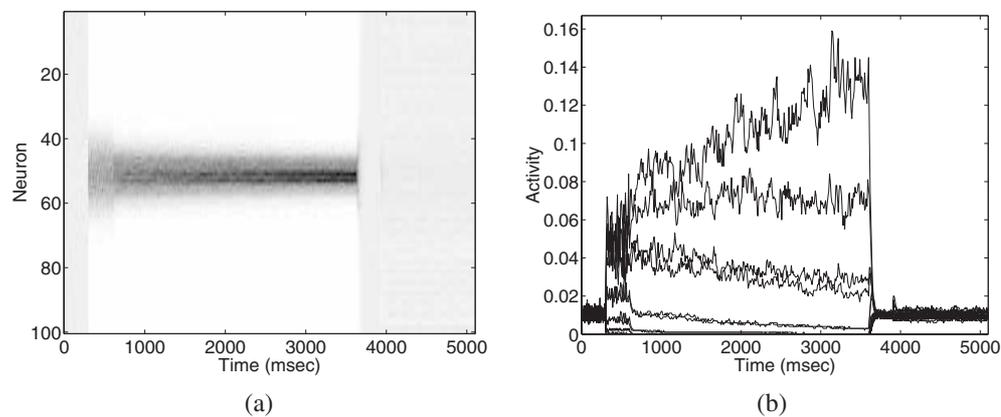
Figure 6. (a) Single stored target, with slow learning ($\tau_L = 72$ s) during delay stage. (b) Plot of selected unit activities over time.

(figure 7). Hence it was possible to both have a single stable state and time multiplexed multiple attractors without having to change the time constant or gain of adaptation.

The upper limit of the number of bump states that could be stored when adaptation occurred appeared to be (for these parameters) approximately 10. For more targets increasing overlap between the attractors created a band weight matrix with a single continuous attractor state. Adaptation then caused the delay state to turn into a continuous moving wave instead of discrete shifts between targets. This is similar to the adaptation-induced instability of bump states in other models [28]. For a smaller number of stored targets the regions of strong local excitation in the weight matrix were separated by mutual inhibition, producing a barrier for the translation of the bump state from one location to another.

(a) One target   (b) Two targets
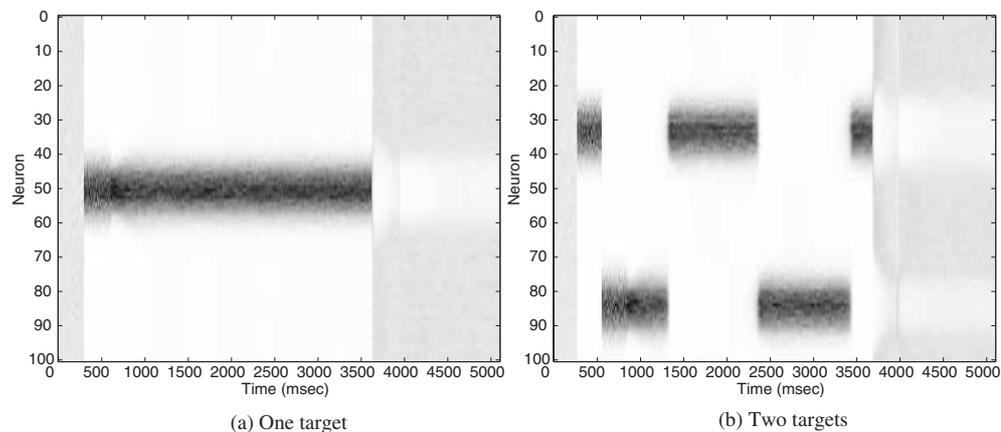
Figure 7. Network activity with adaptation, for $g_A = 0.35$. The bump state was stable for a single learned target, while it flipped between different locations if two or more targets were stored.

4.2. Non-Hebbian plasticity

Non-Hebbian fast adaptation like e.g. augmentation is from a biological point of view a less radical hypothesis than fast Hebbian plasticity for the stabilization of attractor memory states [22]. However, in the current setting of an attractor network with spatially located input, synaptic plasticity dependent only on presynaptic activity (such as augmentation) will produce identical weight changes from each activated unit to all other units, precluding the formation of a localized bump.

A learning rule dependent only on postsynaptic activity can produce a weight matrix that is capable of sustaining a bump state: input-activated neurons will receive input equally weighed from all other neurons, but since they have greater weights than the rest they will be more strongly activated. This is identical to a learning rule changing only cellular excitability. The postsynaptic rule is however not able to learn two or more different bump states, since these will activate each other. It also appears to be insufficient to adequately maintain bump attractor states against destabilizing influences such as adaptation, at least with the current network (simulations were done with plasticity of only β_i in a pre-set weight matrix; results not shown).

We conclude that if fast plasticity is involved in the creation of bump states it needs to be dependent on both pre-and postsynaptic activity if multiple states are to be learned, and at least postsynaptic activity if only one state at the time is stored.

5. Discussion

We have described an attractor neural network with fast Hebbian plasticity that performs as a working memory in the oculomotor delayed response task. It exhibits bump states similar to previous models while being based on synaptic plasticity rather than a hard-wired weight matrix. The modified connectivity maintains the persistent activity and stabilizes it against noise, distractors and network inhomogeneity. This network can readily store multiple memory states, with external cues activating one memory at a time.

In most models to date the maintenance of line attractors requires fine tuning of network parameters [46, 51]. In our network the main parameter issue instead becomes setting learning parameters to a range suitable for the task at hand, which is far less sensitive. Learning also

stabilizes the network against inhomogeneity and noise, making it possible to maintain bump states even when the individual neurons have different parameters and connectivity.

This is similar to the models of eye position control and invariant object recognition described by Seung [44, 45] where line attractors are formed through learning. However, in these models the learning is assumed to occur on a far slower timescale than the network dynamics. Also, our model is not based on the assumption that the learned attractors form or approximate a line attractor; while multiple overlapping stimuli can produce it, it is no requirement in our model for a functional working memory even in the oculomotor task. It may further be possible that slower forms of plasticity create an underlying line attractor weight matrix, which is transiently modified by faster plasticity [50].

Persistent activity is an integral part of this model. There is no fundamental contradiction between the hypothesis that persistent activity within cell assemblies is necessary for working memory and the hypothesis that working memory is stored by changes in synaptic efficacy. Both possibilities can be combined; the activity is necessary for readout to activate the proper behavioural responses and synaptic plasticity is necessary or helpful in maintaining the activity.

It has been argued that a good reason to assume that working memory is stored as persistent activity is that experimental disruption of such activity causes the animal to make an error [12]. However, this would also be the case if the memory was stored synaptically and read out as reverberatory activity. Then instead of interfering with the storage as such the read out from memory would be disturbed. Thus, these two hypotheses cannot be distinguished based on such experimental results.

The existence of sufficiently fast Hebbian plasticity is a basic assumption of this model. While non-Hebbian plasticity of suitable speed is known [22], the existence of fast Hebbian synaptic change in the PFC remains conjectural. Long lasting (one day) LTP at hippocampo-prefrontal synapses can be induced by trains of five short bursts of 10 pulses at 200 Hz given at the frequency of the theta rhythm [10], but the time before the LTP is expressed appears to be of the order of seconds to minutes [19, 23]. On the other hand it is possible that there can exist Hebbian effects within post-tetanic potentiation [4] and that short-lived Hebbian effects are just not easily demonstrated in the current experimental set-ups.

The extension of the model to include adaptation [41] enables attractor dynamics that moves between the different stored states. Due to the existence of a single attractor state in the one cue case, the network avoids the destabilizing effect of adaptation reported by Laing and Longtin [28]. For the same parameter values but with several cues it can also sustain time-multiplexed activity where the network state jumps between the memory states during the delay period, similar to the switching in the binocular rivalry model of Laing and Chow [27].

The addition of an adaptation term destabilizes traditional ring-attractor models [28]. Adaptation acts as an inhibitory current which is preferentially enhanced on the trailing side of a moving bump state, causing it to drive the motion forward. In this model this does not occur due to the non-band structure of the weight matrix, unless enough overlapping memories have been stored to create such a band matrix.

In addition to assigning a functional role in working memory for cellular and synaptic phenomena like plasticity and adaptation, the model presented here makes several specific predictions:

- There could be inter-trial interference due to incompletely erased synaptic changes. Mistakes should tend to confuse the current and a similar previous target, as has been observed [6]. The presentation of several targets at the same place followed by a target nearby should lead to a probability of mistake increasing with the number of times the first target was shown and decreasing with the distance to the second target.

- There should exist a print-now signal resetting the synaptic state of the network concurrent with the reset of neuron activity at the end of the task, regardless of success or failure.
- The shape of the input to the working memory network determines the memory state shape. Hence a modification of the form of the input should result in a corresponding change in the delay activity. If, for example, the previously sharply spatially tuned input was replaced with a broadly tuned input more cells would become less active in the delay period. This would not happen in a fixed synapse model, where the bump shape is set by the pre-existing synaptic structure.
- A prediction and a possible problem with this kind of model is the lack of drift. Although it remains uncertain whether drift of bump states plays a major role in the task [51], the difference between the saccade target and the achieved target in oculomotor tasks appears to increase monotonically with delay time [15, 38, 52]. In the current model it could be due to the presence of other processes, such as a high intrinsic noise level combined with weak learning.

Since the network is not based on pre-set synaptic strengths, it can be generalized to arbitrary attractor states. For example, it can be used to learn from a 2D retinotopic map, enabling the persistence of 2D bumps of activity at locations learned during the cue period. The attractor states could also be distributed representations as in many regular attractor network memories. Instead of using a specific kind of network for working memory and another for long-term memory, our model suggests the possibility that the cerebral cortex may be using a canonical network architecture with a spectrum of learning time constants for different functions. Such general networks would become modality specific due to their afferent and efferent connections rather than any particular architectural features, and shift their function due to task demands. This could explain the apparent discrepancy between different studies on the domain specificity of PFC (as discussed in [49]).

Though we investigated within the context of a specific computational model (BCPNN) we expect the qualitative aspects of our results to be largely independent of the exact simulation model. A possible future extension would be to implement the model using spiking neurons to further examine its generality and properties, such as the effect of spike synchrony in memory reset and the effects of AHP modulation on network dynamics.

In conclusion, we have shown that fast Hebbian learning is sufficient to reproduce many properties of a working memory task which has previously been modelled in terms of persistent activity states in a fixed connectivity attractor network. It remains to be examined to what extent plasticity is necessary to maintain stable activity and how quickly it can be modulated to fit the task.

Appendix

The BCPNN update equations for unit activity are

$$\tau_{\mathrm{m}} \frac{\mathrm{d}h_i(t)}{\mathrm{d}t} = \beta_i(t) + \ln\left(\sum_j^N w_{ij}(t)x_j(t)\right) + g_{\mathrm{I}}I_i + g_{\mathrm{N}}\eta_i(t) - h_i(t) \tag{A.1}$$

$$x_i(t) = \frac{\mathrm{e}^{h_i}}{\sum_j^N \mathrm{e}^{h_i}} \tag{A.2}$$

where h_i is the membrane potential of unit i, x_i the rate of fire. τ_{m} is the membrane time constant, g_{I} the gain of the external input I_i, and $g_{\mathrm{N}}\eta_i(t)$ a noise term.

Table A.1. Default parameters of the model.

Symbol	Parameter	Value
	Cue length (ms)	300
	Delay period length (ms)	3000
	Reset period length (ms)	300
N	Number of neurons	100
σ	Input tuning width	10
g_I	Input gain	1
g_N	Noise gain	0.1
g_A	Adaptation gain	-0.35
τ_L	Learning time constant (ms)	7200
τ_A	Adaptation time constant (ms)	160
τ_m	Membrane time constant (ms)	10
$\kappa(t)$	Print-now signal	0 or 1, 90 for reset, 15 in figure 4

The weights w_{ij} and their corresponding biases β_i are set by the following learning rule:

$$\kappa(t)\tau_L \frac{d\hat{P}_i(t)}{dt} = x_i(t) - \hat{P}_i(t) \tag{A.3}$$

$$\kappa(t)\tau_L \frac{d\hat{P}_{ij}(t)}{dt} = x_i(t)x_j(t) - \hat{P}_{ij}(t) \tag{A.4}$$

$$\beta_i(t) = \log(\hat{P}_i(t)) \tag{A.5}$$

$$w_{ij}(t) = \frac{\hat{P}_{ij}(t)}{\hat{P}_i(t)\hat{P}_j(t)} \tag{A.6}$$

where $\hat{P}_i$ and $\hat{P}_{ij}$ are moving averages of unit activity and co-activity.

τ_L is the learning time constant of the associative projection. The learning time constant is assumed to be modulated by an external input $\kappa(t)$ such as a dopamine signal, acting as a 'print-now' signal [8, 42, 53]. During memory encoding it is decreased, enabling the network to learn quickly, while during delay periods the time constant is long, making the synaptic strengths stable.

The adaptation was modelled using an extra associative projection with learning time constant τ_A and (negative) gain g_A added to the update equation (A.1).

References

[1] Amari S-I 1977 Dynamics of pattern formation in lateral-inhibition type neural fields *Biol. Cybern.* **27** 77–87
[2] Amit D and Brunel N 1997 Dynamics of a recurrent network of spiking neurons before and following learning *Network: Comput. Neural Syst.* **8** 373–404
[3] Amit D J 1994 The Hebbian paradigm reintegrated: local reverberations as internal representations *Behav. Brain Sci.* **18** 617–26
[4] Bao J-X, Kandel E R and Hawkins R D 1997 Involvement of pre- and postsynaptic mechanisms in posttetaic potentiation at aplysia synapses *Science* **275** 969–70
[5] Ben-Yishai R, Bar-Or R L and Sompolinsky H 1995 Theory of orientation tuning in visual cortex *Proc. Natl Acad. Sci. USA* **92** 3844–8
[6] Bichot N P and Schall J D 1999 Effects of similarity and history on neural mechanisms of visual selection *Nat. Neurosci.* **2** 549–54
[7] Cartling B 1997 Control of computational dynamics of coupled integrate-and-fire neurons *Biol. Cybern.* **76** 383–95
[8] Cohen J, Braver T and Brown J 2002 Computational perspectives on dopamine function in prefrontal cortex *Curr. Opin. Neurobiol.* **12** 223–9

[9] Compte A, Brunel N, Goldman-Rakic P S and Wang X-J 2000 Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model *Cereb. Cortex* **10** 910–23

[10] Doyere V, Burette F, Negro C R-D and Laroche S 1993 Long-term potentiation of hippocampal afferents and efferents to prefrontal cortex: implications for associative learning *Neuropsychologia* **31** 1031–53

[11] Durstewitz D, Kelc M and Güntürkün O 1999 A neurocomputational theory of the dopaminergic modulation of working memory functions *J. Neurosci.* **19** 2807–22

[12] Durstewitz D, Seamans J K and Sejnowski T J 2000 Neurocomputational models of working memory *Nat. Neurosci.* **3** 1184–91

[13] Fransén E and Lansner A 1995 Low spiking rates in a population of mutually exciting pyramidal cells *Network: Comput. Neural Syst.* **6** 271–88

[14] Fransén E and Lansner A 1998 A model of cortical associative memory based on a horizontal network of connected columns *Network: Comput. Neural Syst.* **9** 235–64

[15] Funahashi S, Bruce C J and Goldman-Rakic P 1989 Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex *J. Neurophysiol.* **61** 331–49

[16] Fuster J 1989 *The Prefrontal Cortex* 2nd edn (New York: Raven)

[17] Gaiarsa J-L, Caillard O and Ben-Ari Y 2002 Long-term plasticity at GABAergic and glycinergic synapses: mechanisms and functional significance *Trends Neurosci.* **25** 564–70

[18] Goldman-Rakic P 1995 Cellular basis of working memory *Neuron* **14** 477–85

[19] Gustafsson B, Asztely F, Hanse E and Wigström H 1989 Onset characteristics of long-term potentiation in the guinea-pig hippocampal ca1 region *in vitro Eur. J. Neurosci.* **1** 382–94

[20] Gutkin B S, Laing C R, Colby C L, Chow C C and Ermentrout G B 2001 Turning on and off with excitation: the role of spike-timing asynchrony and synchrony in sustained neural activity *J. Comput. Neurosci.* **11** 121–34

[21] Hansel D and Sompolinsky H 1996 Chaos and synchrony in a model of a hypercolumn in visual cortex *J. Comput. Neurosci.* **3** 7–34

[22] Hempel C M, Hartman K H, Wang X-J, Turrigiano G G and Nelson S B 2000 Multiple forms of short-term plasticity at excitatory synapses in rat medial prefrontal cortex *J. Neurophysiol.* **83** 3031–41

[23] Hirsch J and Crepel F 1990 Use-dependent changes in synaptic efficacy in rat prefrontal neurons *in vitro J. Physiol.* **427** 31–49

[24] Holst A 1997 The use of a Bayesian neural network model for classification tasks *PhD Thesis* Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, Sweden, Sept. TRITA-NA-P9708

[25] Hopfield J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8

[26] Laing C R and Chow C C 2001 Stationary bumps in networks of spiking neurons *Neural Comput.* **13** 1473–94

[27] Laing C R and Chow C C 2002 A spiking neuron model for binocular rivalry *J. Comput. Neurosci.* **12** 39–53

[28] Laing C R and Longtin A 2001 Noise-induced stabilization of bumps in systems with long-range spatial coupling *Physica* D **160** 149–72

[29] Lansner A and Ekeberg Ö 1989 A one-layer feedback artificial neural network with a Bayesian learning rule *Int. J. Neural Syst.* **1** 77–87

[30] Lansner A and Fransén E 1992 Modeling Hebbian cell assemblies comprised of cortical neurons *Network: Comput. Neural Syst.* **3** 105–19

[31] Lansner A and Holst A 1996 A higher order Bayesian neural network with spiking units *Int. J. Neural Syst.* **7** 115–28

[32] Lashley K 1951 The problem of serial order in behavior *Cerebral Mechanisms in Behavior* ed L Jeffress (New York: Wiley)

[33] McCormick D, Connors B, Lighthall J and Prince D 1985 Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex *J. Neurophysiol.* **54** 782–805

[34] Orlov T, Yakovlev V, Amit D, Hochstein S and Zohary E 2002 Serial memory strategies in macaque monkeys: behavioral and theoretical aspects *Cereb. Cortex* **12** 306–17

[35] Otani S, Blond O, Desche J-M and Crépel F 1998 Dopamine facilitates long-term depression of glutamatergic transmission in rat prefrontal cortex *Neuroscience* **85** 669–76

[36] Pantic L, Torres J, Kappen H and Gielen S C 2002 Associative memory with dynamic synapses *Neural Comput.* **14** 2903–23

[37] Perez Y, Morin F and Lacaille J-C 2001 A Hebbian form of long-term potentiation dependent on mGluR1a in hippocampal inhibitory interneurons *Proc. Natl Acad. Sci. USA* **96** 9401–6

[38] Ploner C J, Gaymard B, Rivaud S, Agid Y and Pierrot-Deseilligny C 1998 Temporal limits of spatial working memory in humans *Eur. J. Neurosci.* **10** 794–7

[39] Romo R, Brody C D, Hernandez A and Lemus L 1999 Neuronal correlates of parametric working memory in the prefrontal cortex *Nature* **399** 470–3

[40] Sandberg A 2003 Bayesian attractor neural network models of memory *PhD Thesis* Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, Sweden TRITA-NA-0310, ISBN 91-7265-684-0

[41] Sandberg A and Lansner A 2002 Synaptic depression as an intrinsic driver of reinstatement dynamics in an attractor network *Neurocomputing* **44–46** 615–22

[42] Sandberg A, Lansner A and Petersson K 2001 Selective enhancement of recall through plasticity modulation in an autoassociative memory *Neurocomputing* **38–40** 867–73

[43] Sandberg A, Lansner A, Petersson K-M and Ekeberg Ö 2002 A Bayesian attractor network with incremental learning *Network: Comput. Neural Syst.* **13** 179–94

[44] Seung H 1996 How the brain keeps the eyes still *Proc. Natl Acad. Sci. USA* **93** 13339–44

[45] Seung H 1998 Learning continuous attractors in recurrent networks *Adv. Neural Inf. Process. Syst.* **10** 654–60

[46] Seung H, Lee D, Reis B and Tank D 2000 Stability of the memory of eye position in a recurrent network of conductance-based model neurons *Neuron* **26** 259–71

[47] Tegnér J, Compte A and Wang X 2002 The dynamical stability of reverberatory dynamics *Biol. Cybern.* **87** 471–81

[48] Torres J J, Pantic L and Kappen H J 2002 Storage capacity of attractor neural networks with depressing synapses *Phys. Rev.* E **66** 061910

[49] Undergleider L G, Courtney S M and Haxby J V 1998 A neural system for human visual working memory *Proc. Natl Acad. Sci. USA* **95** 883–90

[50] von der Malsburg C 1981 The correlation theory of brain function *Technical Report* Department of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Göttingen
Reprinted in: Domany E, van Hemmen J L and Schulten K (ed) 1994 *Models of Neural Networks* vol 2 (Berlin: Springer) chapter 2 pp 95–119

[51] Wang X-J 2001 Synaptic reverberation underlying mnemonic persistent memory *Trends Neurosci.* **24** 455–63

[52] White J M, Sparks D L and Stanford T R 1994 Saccades to remembered target locations: an analysis of systematic and variable errors *Vis. Res.* **34** 79–92

[53] Wickens J and Kötter R 1995 Cellular models of reinforcement *Models of Information Processing in the Basal Ganglia* ed J C Houk, J L Davis and D G Beiser (Cambridge, MA: MIT Press) pp 187–214

[54] Willshaw D, Buneman O and Longuet-Higgins H 1969 Non-holographic associative memory *Nature* **222** 960–2

[55] Wu X and Liljenström H 1994 Regulating the nonlinear dynamics of olfactory cortex *Network: Comput. Neural Syst.* **5** 47–60