
A self-organizing memory network

Callie Federer¹ and Joel Zylberberg^{1,2}

¹Department of Physiology and Biophysics, University of Colorado Anschutz Medical Campus

²Learning in Machines and Brains Program, Canadian Institute For Advanced Research

Abstract

Working memory requires information about external stimuli to be represented in the brain even after those stimuli go away. This information is encoded in the activities of neurons, and neural activities change over timescales of tens of milliseconds. Information in working memory, however, is retained for tens of *seconds*, suggesting the question of how time-varying neural activities maintain stable representations. Prior work shows that, if the neural dynamics are in the ‘null space’ of the representation - so that changes to neural activity do not affect the downstream read-out of stimulus information - then information can be retained for periods much longer than the time-scale of individual-neuronal activities. The prior work, however, requires precisely constructed synaptic connectivity matrices, without explaining how this would arise in a biological neural network. To identify mechanisms through which biological networks can self-organize to support memory function, we derived biologically plausible synaptic plasticity rules that dynamically modify the connectivity matrix to enable information storing. Networks implementing this plasticity rule can successfully learn to store information even if only 10% of the synapses are plastic, they are robust to synaptic noise, and they can store information about multiple stimuli.

1 Introduction

Working memory is a key cognitive function, and it relies on us retaining information about external stimuli even after they go away. Stimulus-specific elevated firing rates have been observed in the prefrontal cortex during the delay period of working memory tasks, and are the main neural correlates of working memory [10, 11]. Perturbations to the delay period neural activities cause changes in the animal’s subsequent report of the remembered stimulus [16, 23]. These elevated delay-period firing rates are not static but have time-varying dynamics with activities changing over timescales of tens of *milliseconds* [3, 22, 1], yet information can be retained for tens of *seconds* (Fig. 1A). This suggests the question of how time-varying neural activities keeps representing the same information.

Prior work from Druckmann and Chklovskii shows that, if the neural dynamics are in the “null space” of the representation – so that changes to neural activity do not affect the downstream read-out of stimulus information – then information can be retained for periods much longer than the time-scale of individual neuronal activities (called the FEVER model; Fig. 1B) [9]. That model has a severe fine-tuning problem, discussed below. We identified a synaptic plasticity mechanism that overcomes this fine-tuning problem, enabling neural networks to learn to form stable representations.

While the dynamics of neurons in the FEVER model closely match that which is observed in the monkey prefrontal cortex during a working memory task [20], the model itself requires that the network connectivity matrix have one or more eigenvalues very near to unity. According to the Gershgorin Circle Theorem, this will almost surely not happen in randomly-connected networks: finely tuned connectivity is needed. Druckmann and Chklovskii suggest a mechanism of Hebbian

learning by which this fine-tuned connectivity can be learned. That mechanism requires the read-out weights to form a ‘tight frame’ [9], which will not necessarily be true in biological circuits. Thus, the prior work leaves it unknown how synaptic plasticity can form and/or maintain functional working memory networks.

In this paper, we identify biologically plausible synaptic plasticity rules that can solve this fine-tuning problem without making strong assumptions like ‘tight frame’ representations. Our plasticity rules dynamically re-tune the connectivity matrix to enable persistent representations of stimulus information. We perform experiments to demonstrate that networks using these plasticity rules are able to store information about multiple stimuli, work even if only a fraction of the synapses are tuned, and are robust to synaptic noise. We also show that these networks improve over time with multiple presented stimuli, and that the learning rules work within densely or sparsely connected networks.




Figure 1: Stimulus retention in working memory. **(A)** When presented with an external stimulus, s , neural activity patterns initially encode an internal representation of that stimulus, $\hat{s}(t=0)$. These neural activity patterns change over time on timescales of tens of milliseconds, and yet somehow the same information is stored for up to tens of seconds. **(B)** While the firing rates, $r_i(t)$, change over time, information about stimulus, $\hat{s}(t)$, can remain unchanged as long as the projection of the firing rates onto the ‘‘read-out’’ vector \vec{d} , remains constant [9].

2 Model

2.1 The Rate-Based Network Model

We use a rate-based network like that of the FEVER model [9] but with positive rectifying activation functions (ReLU - rectified linear units) [5]. We use standard linear dynamics in the network model:

$$\tau \frac{da_i}{dt} = -a_i(t) + \sum_j L_{ij} r_j(t), \quad (1)$$

where $a_i(t)$ is the internal state (membrane potential) of the i^{th} neuron at time t , $r_j(t)$ is the output (firing rate) of the j^{th} neuron, τ the time constant, and L_{ij} represents the strength of synaptic connection from neuron j to neuron i . The firing rates, $r_j(t)$, are given by a positive rectifying function of the internal states: $r_j(t) = [a_j(t)]_+$. After an external stimulus, s , is presented to the network, the network’s representation of the stimulus, $\hat{s}(t)$, is obtained from a weighted combination of firing rates:

$$\hat{s}(t) = \sum_i d_i r_i(t), \quad (2)$$

where d_i is the weight of the contribution of neuron i to the stimulus (the ‘‘read-out weight’’). For now we consider a single scalar stimulus value: we study the encoding of multiple stimuli in section 3.2.

2.2 Plasticity Rules

In order to organize the network to store information, we update the synaptic weights, L_{ij} , so as to minimize changes in the stimulus representation. To do this, we differentiated Equation 2 with respect to time, to calculate $\frac{d\hat{s}}{dt}$. We then used gradient descent with respect to L_{ij} on loss function $(\frac{d\hat{s}}{dt})^2$ to calculate the update rule:

$$\Delta L_{ij} = -\eta \frac{d\hat{s}}{dt} d_i r'_i(t) r_j(t), \quad (3)$$

where η is the learning rate of the network and r' is the slope of the (ReLU) activation function. In section 3.6, we discuss the biological plausibility of this plasticity rule. We chose the elements of \vec{d} to be positive based on candidate sources of the global error signal discussed in section 3.6. The source code and scripts for reproducing our experiments are available at <https://github.com/cfederer/SOMN>.

3 Results

3.1 Stimulus retention in the self-organizing memory network

To evaluate the performance of our working memory networks, we determined how well the networks could store information about a scalar stimulus value. We quantified the fraction of stimulus retained, $\frac{\hat{s}(t)}{\hat{s}(t=0)}$, over 3 seconds. This is the duration of heightened activity observed during working memory tasks [10]. The networks were all-to-all connected (partial connectivity discussed in section 3.5) and contained 100 neurons. We first determined how well random networks without plastic synapses store information, by initializing each network with random neural activities, $\vec{a}(t=0)$, random read-out weights, \vec{d} , and random connection weight matrices, L , and simulating the dynamical evolution of the representation. We then compared these “constant random synapse” networks to ones with identical random initial conditions, but in which our plasticity rule (Eq. 3) dynamically updated the synapses (Figs. 2B, C). Finally, we compared both of these randomly-initialized networks to ones that had the fine-tuned connectivity specified by the FEVER model.

The stimulus value in the randomly-initialized plastic network initially decreased slightly, but the plasticity rules quickly reorganized the connectivity, and the representation remained constant after the first ~ 50 ms. In the networks with fixed random synaptic weights, the representation quickly decayed to 0 (Figs. 2A,B). Thus, our plasticity rule enables initially random networks to quickly become effective working memory systems.

To ensure that the success of our plasticity rule at forming an effective memory network was not limited to a fortuitous random initialization, we quantified the fraction of stimulus retained over 100 different networks, each with a different initial connectivity matrix, read-out vector, and initial activity vector $\vec{a}(t=0)$. In the FEVER networks, stimulus retention is perfect across all networks. The models with plastic random synapses perform almost as well as the FEVER models, but require some time to self-organize before the representations remain constant (Figs. 2B, C). In the random constant synapse networks, information is quickly lost (Fig. 2B).

3.2 Remembering multiple stimuli

The previous section shows how networks can self-organize to store information about one stimulus, but working memory capacity in adult humans is typically 3–5 items [7]. To incorporate this working memory capacity into our models, we adapted the representation such that there were multiple read-out vectors, one for each stimulus value. We then derived plasticity rules via gradient descent on the squared and summed time derivatives of these representations: the loss was $\sum_k (\frac{d\hat{s}_k}{dt})^2$. This led to the plasticity rule

$$\Delta L_{ij} = -\eta \sum_{k=1}^n \frac{d\hat{s}_k}{dt} d_{ik} r'_i(t) r_j(t), \quad (4)$$

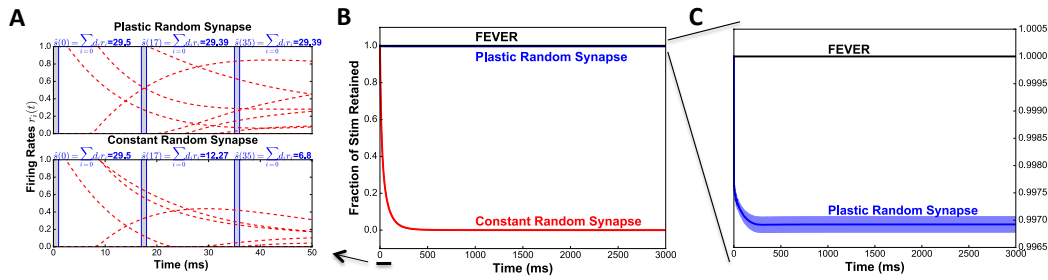


Figure 2: Stimulus retention in self-organizing memory networks. **(A)** Neural dynamics of the first 50 ms of a simulation of two different networks with the same random initial conditions. One network had constant random synapses (bottom panel), while the other one had plastic synapses that were updated via Eq. 3 (top panel). Red dashed lines show the firing rates of 10 of the neurons. The remembered stimulus values at times 0, 17 ms and 35 ms (indicated by the shaded vertical bars) are shown. In the plastic random network, the remembered stimulus value does not change much even though the neural activities keep evolving. **(B)** The average fraction of stimulus retained over 100 random initializations for the FEVER model (black), the plastic synapse model (blue) and the constant random synapse model (red). **(C)** A zoomed in look at the average fraction of stimulus retained for the FEVER model and the plastic random synapse model from (B). Shaded areas in B and C represent \pm standard error of the mean.

where n is the number of stimuli to be remembered. We chose $n = 4$ for our experiments, and we quantified how well the networks remember these stimuli (S1-S4 in Fig. 3). (The networks can store up to 100 stimuli; data not shown). In our experiments, each neuron in the network contributed to the representation of every stimulus value: *in vivo*, most neurons are sensitive to multiple aspects of stimuli [19]. This is not a requirement: the models successfully represent multiple stimuli even when subsets of the neurons participate in each representation.

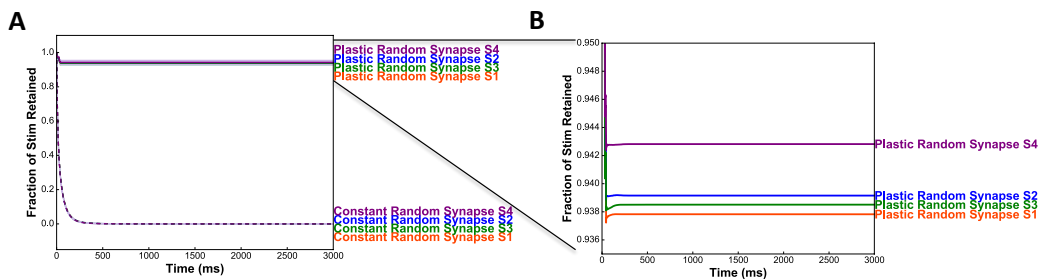


Figure 3: Remembering four things. **(A)** The average fraction of four stimuli retained for 100 randomly initialized networks with constant random synapses (dashed lines) and plastic random synapses (solid lines). Different colors are for different stimulus values. Shaded areas represent \pm standard error of the mean. **(B)** A zoomed in look at the average fraction of four stimuli retained in the plastic model in panel A. Shading omitted from lines in panel B for clarity.

3.3 Network Robustness

3.3.1 Partial Tuning

While synapses are plastic, it is not known if *all* synapses change. To determine how well the network performs if only some synapses are updated, we simulated networks in which different fractions of the synapses were updated using Eq. 3: the other synapses were held constant. We quantified the fraction of stimulus retained by these networks (Fig. 4A). Even with just 10% of the synapses being tuned, the networks learn to store information about the stimulus. In hindsight, this makes

sense. To store n stimulus values, n constraints must be satisfied by the connectivity matrix: it must have n eigenvalues near 1 (we chose $n = 1$ for Fig. 4A). Because the connectivity matrix has many more than n elements, many configurations can satisfy the constraint, so it is possible to satisfy the constraint without updating every synapse.

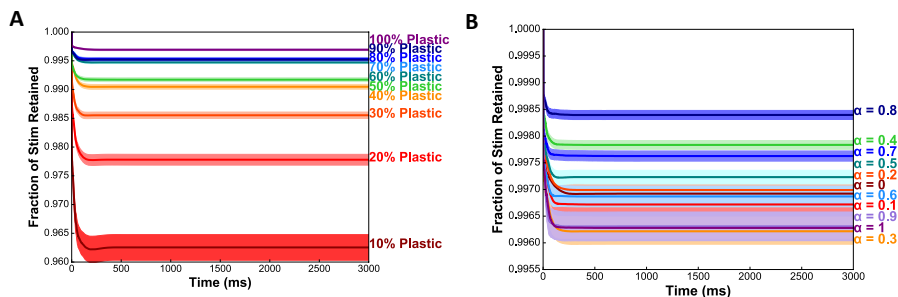


Figure 4: Network robustness to partial plasticity and noise. **(A)** The fraction of stimulus retained for 100 random initializations. Different lines are for different fractions of plastic synapses. **(B)** The fraction of stimulus retained for 100 random initializations of networks with differing levels of synaptic update noise (α). Shaded areas in A and B represent \pm standard error of the mean.

3.3.2 Noisy Synapses

Working memory must be robust to noise and imprecise components [4]. To determine how robust our networks are to synaptic noise, we simulated networks with various levels of random noise added to the synaptic updates. We added Gaussian noise with mean 0 and standard deviation α times the update to the synapse, ΔL_{ij} , where α was varied from 0 to 1. We quantified the fraction of stimulus retained for networks with various noise levels and found that the noise did not have a noticeable effect on the network performance (Fig. 4B). The fraction of stimulus retained is almost equivalent for all values of α , with minor differences due to random initial conditions. This should not be surprising considering that multiplying error signals by random synaptic weights does not hinder learning, so long as the network is still being pushed down the loss gradient, even if not directly [17].

3.4 Pre-training the Network

In the previous examples, each network is initialized with random connection weights. In reality, the working memory network will be continuously learning and will not start over with random connection weights when each new stimulus is presented. Consequently, we speculated that, once the network had learned to store one stimulus, it should be able to remember subsequently presented stimuli, even with minimal re-training. Relatedly, experimental work shows that performance in working memory tasks in children and young adults can be increased not only for trained tasks but for new tasks not part of the training: this coincides with strengthening of connectivity in the prefrontal cortex [6].

To determine if our synaptic update rule enables the network to store new stimuli without further training, we first trained the networks (Eq. 3) to remember 1, 5 or 10 individual stimuli, one at a time: each new stimulus corresponded to another random initialization of the activity patterns $\vec{a}(t = 0)$. We quantified the networks' abilities to represent these training stimuli, and found that the networks performed better on each subsequent stimulus: training improved performance (Fig. 5A). Next, we asked if after training on 1, 5, or 10 prior stimuli, the network could store information about a new stimulus without any more synaptic updates. We found that a network was able to store information about a new stimulus after being trained on at least 1 previous stimulus (Fig. 5B). Once the connectivity weight matrix (L) has obtained one or more eigenvalues near unity, it is able to stably store novel stimuli without additional training.

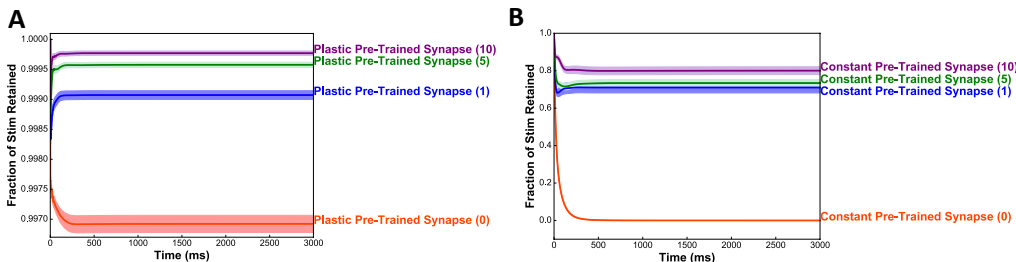


Figure 5: Training improves performance. **(A)** The fraction of stimulus retained over 100 random initializations for a plastic random synapse network that has seen 0 previous stimuli (orange), 1 previous stimulus (blue), 5 previous stimuli (green) or 10 previous stimuli (purple). **(B)** The fraction of stimulus retained over 100 random initializations for a constant synapse network that has previously been trained on 0, 1, 5, or 10 previous stimuli, but with no training during the simulation period shown. Color coding as in panel A. Shaded areas in A and B represent \pm standard error of the mean.

3.5 Partially Connected Networks

In the previous results, the connectivity was all-to-all (100% connectivity). Real neural circuits are not 100% connected. In visual cortex, for example, connection probabilities range from 50% to 80% for adjacent neurons [14]. To ask if our synaptic update rule could self-organize partially connected networks, we simulated networks with different connection probabilities, and in which the synapses were updated using Eq. 3. We quantified the fraction of stimulus retained by these networks. Performance declines somewhat as connectivity decreases, but even networks with 10% connection probabilities can learn to store stimulus information (Fig. 6). Experiments show that working memory performance declines with age, which correlates with a reduction in the number of synapses [21, 2].




Figure 6: Network robustness to partial connectivity. The fraction of stimulus retained over 100 random initializations. Different lines are for different connection probabilities. Shaded areas represent \pm standard error of the mean.

3.6 Biological Plausibility

Synaptic updates are thought to rely on synaptically local information, like the activities of the pre- and post-synaptic neurons. Our plasticity rules (Eqs. 3,4) involve this information, in addition to “global” error value(s) $\{\frac{ds_k}{dt}\}$. Thus, we obtained three-factor rules: synaptic changes depend on the pre- and post-synaptic neurons’ activities, and a global error signal [17]. We propose two possible sources of the global error signal, one involving segregated dendrites, and the other using neuromodulators.

3.6.1 Calculating the Global Error Signal Locally Using Segregated Dendrites

The global error signal could be calculated locally by each neuron, by exploiting the fact that, in pyramidal cells, feedback arrives at the apical dendrites, and modulates synaptic plasticity at the

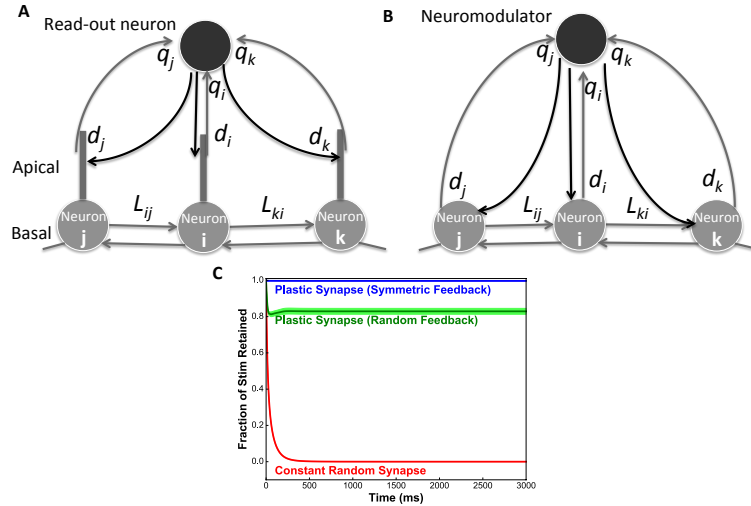


Figure 7: Potential biological implementations. **(A)** Cartoon depicting the error signals being calculated locally by the neurons' apical dendrites [12]. The read-out neuron calculates the remembered stimulus value via Eq. 7, where q_i is the weight of the synapses from cell i in the memory circuit to the read-out neuron. The apical dendrite calculates the error signal, weighted by d_i , by subtracting the feedback values at adjacent time points. The basal dendrite receives inputs from neurons in the working memory circuit. The soma transmits the error signal to the basal dendrites, where the synaptic updates, L_{ij} are calculated by correlating the pre- and post-synaptic activities with the error signal. **(B)** Cartoon depicting the error signals being communicated by neuromodulators. Neurons in the modulatory system calculate the remembered stimulus values via Eq. 7, where q_i is the weight of the synapse from cell i in the memory circuit to the modulatory read-out neuron. That cell releases an amount of neuromodulator that tracks the changes in the represented stimulus value. The modulatory chemical affects synaptic plasticity by an amount that depends on the receptor density, d_i , at the synapses on neuron i . Both implementations require asymmetric feedback: the weights to the read-out neuron from cell i (q_i) will not necessarily match the weights, d_i , with which cell i receives the feedback signals. **(C)** The fraction of stimulus retained over 100 random initializations of a network with a plastic random synapse model with symmetric feedback ($q = d$) (blue), plastic random synapses with random feedback and readout weights ($q \neq d$) (green), and constant synapses (red). The shaded areas represent \pm standard error of the mean.

basal dendrites (where information comes in from other cells in the memory network) [17, 12] (Fig. 7A). Here, a readout layer provides feedback to the apical dendrites that specifies the represented stimulus value $\hat{s}(t)$: the weight of the feedback synapse to neuron i from read-out neuron k is $d_{i,k}$, and so the apical dendrite receives a signal $\sum_k d_{i,k} \hat{s}_k(t)$. The apical dendrites track the changes in these feedback signals, sending that information ($\sum_k d_{i,k} d\hat{s}_k/dt$) to basal dendrites via the soma. Correlating those signals with the pre- and post-synaptic activity at each of the synapses on the basal dendrites, the synaptic updates specified by Eq. 4 are obtained. Thus, the neurons locally compute the synaptic updates.

3.6.2 Signalling the Global Error Signal with Neuromodulators

Alternatively, the global error signal(s) could be communicated throughout the network by neuromodulators, like dopamine, acetylcholine, serotonin, norepinephrine or nitric oxide. These have all been shown to be important in synaptic plasticity in the prefrontal cortex and in working memory [18, 8]. This is reward learning, with the reward values coming from the neuromodulator concentrations. Experimental work shows that synapses have activity-dependent "eligibility traces" that are converted into changes in synaptic strength by reward-linked neuromodulators [13]. In this scenario the concentration of different modulators tracks the error signals, $\frac{d\hat{s}_k}{dt}$, and the densities of the receptors to the modulators at each synapse are $d_{i,k}$. Thus, at each synapse, the modulators bring information

$\sum_k d_{ik} d\hat{s}_k/dt$ that, when correlated with the pre- and post-synaptic activities, yields the updates from Eq. 4 (Fig. 7B).

3.6.3 Calculating the stimulus with random feedback weights

In either of the two scenarios discussed above, d_i must be accessed in two separate places: at the synapse, L_{ij} , to calculate the update, and at the read-out layer to calculate the remembered stimulus value (Eq. 2). Weight transport is not biologically plausible [17]. To remove this source of implausibility, we let the top-down feedback impinge on the neurons at synapses with weights d_i , and the read-out layer calculate the remembered stimulus as:

$$\hat{s}(t) = \sum_i q_i r_i(t), \quad (5)$$

where $d \neq q$ [17]. To test whether networks with this asymmetric feedback could learn to store stimulus information, we simulated such asymmetric networks and quantified the fraction of stimulus remaining. The results (Fig. 7C) show that even with asymmetric feedback, networks can learn to store stimulus information. This can be understood by noting that, so long as both q and d contain only positive elements, the feedback update signal to each synapse has the same sign as the update calculated from gradient descent. Thus, the synaptic updates with asymmetric feedback are in the same direction as those obtained from gradient descent, which suffices for learning [17].

4 Discussion

We derived biologically plausible synaptic plasticity rules through which networks self-organize to store information in working memory. Networks implementing these plasticity rules are robust to synaptic noise, to having only some of the synapses updated, and to partial connectivity. These networks can store multiple stimuli and have increased performance after previous training. We suggest two candidate sources for the global error signal necessary for the plasticity rule, and demonstrate that our networks can learn to store stimulus information while satisfying the added requirements imposed by these biological mechanisms. This flexibility suggests that other types of synaptic plasticity updates may also be able to organize working memory circuits.

The results presented here were obtained for networks of 100 neurons – as opposed to larger networks – to speed up the simulations. Tests on networks with 10,000 neurons show that the update rule works in larger networks. The optimal learning rate, η , decreases as the network size increases. Aside from network size, a potential caveat in using a rate-based network model is losing information about spike-timing dependency. A future direction would be to create a spike-based model and determine what, if anything, must be adjusted to account for spike timing, and for the discretization that spiking neurons entail for the information shared between cells.

Along with understanding how information is stored in working memory, this work may have implications in training recurrent neural networks (RNNs). Machine learning algorithms are generally unrealistic from a biological perspective: most algorithms rely on non-local synaptic updates or symmetric synapses. We show that a recurrent network can learn to store information using biologically plausible synaptic plasticity rules which require local information plus a global error signal (or signals), that can be calculated on the apical dendrite or via neuromodulators. This same synaptic update setup could be utilized in RNNs to make them more biologically realistic. This would let us better understand how the brain learns, and could lead to novel biomimetic technologies: prior work on biologically realistic machine learning algorithms has led to hardware devices that use on-chip learning [15, 24]. Synaptically local updates do not have to be coordinated over all parts of the chip, enabling simpler and more efficient hardware implementations.

Acknowledgments

The authors would like to acknowledge Z. Kilpatrick, A. Person, G. Felson, J. Costello and J. Gilmer for their help. This work was supported by a Canadian Institute for Advanced Research (CIFAR), Azrieli Global Scholar Award, and a Google Faculty Research Award.

References

- [1] E. H. Baeg, Y. B. Kim, K. Huh, I. Mook-Jung, H. T. Kim, and M. W. Jung. Dynamics of population code for working memory in the prefrontal cortex. *Neuron*, 40(1):177–188, 2003.
- [2] K.L. Bopp and P. Verhaeghen. Working Memory and Aging: Separating the Effects of Content and Context. *Psychol Aging*, 24(4):968–980, 2009.
- [3] CD Brody, A Hernández, A Zainos, and R. Romo. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex*, (11), 2003.
- [4] CD Brody, R Romo, and A. Kepecs. Basic mechanisms for graded persistent activity: Discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol*, 13(2):204–211, 2003.
- [5] M. Carandini. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biology*, 2(9), 2004.
- [6] C Constantinidis and T Klingberg. The neuroscience of working memory capacity and training. *Nat Rev Neurosci.*, 17(7):438–449, 2016.
- [7] N. Cowan. NIH Public Access. *Curr Dir Psychol Sci*, 19(1):51–57, 2010.
- [8] M D ’esposito and BR. Postle. The Cognitive Neuroscience of Working Memory. *Annu. Rev. Psychol*, 66:115–42, 2015.
- [9] S Druckmann and DB. Chklovskii. Neuronal circuits underlying persistent representations despite time varying activity. *Curr Biol*, 22(22):2095–2103, 2012.
- [10] S Funahashi, M.V. Chafee, and P.S. Goldman-Rakic. Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, (365):753–756, 1993.
- [11] J. M. Fuster and G. E. Alexander. Neuron Activity Related to Short-Term Memory. *Science*, 173(3997):652–654, 1971.
- [12] J Guerguiev, T.P. Lillicrap, and B.A. Richards. Deep learning with segregated dendrites. *arXiv preprint*, 1610.00161:1–29, 2016.
- [13] K He, M Huertas, S.Z. Hong, X Tie, J.W. Hell, H Shouval, and A. Kirkwood. Distinct Eligibility Traces for LTP and LTD in Cortical Synapses. *Neuron*, 88(3):528–538, 2015.
- [14] B Hellwig. A quantitative analysis of the local connectivity between pyramidal neuron in layers 2/3 of the visual cortex. *Biol. Cybern.*, 82:111–121, 2000.
- [15] P Knag, J.K. Kim, T Chen, and Z Zhang. A Sparse Coding Neural Network ASIC With On-Chip Learning for Feature Extraction and Encoding. *IEEE J. Solid-State Circuits*, 50(4):1070–1079, 2015.
- [16] N Li, K Daie, K Svoboda, and S Druckmann. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 532(7600):459–464, 2016.
- [17] T.P. Lillicrap, D Cownden, D.B.. Tweed, and C.J. Akerman. Random feedback weights support learning in deep neural networks. *Nat. Commun*, pages 1–27, 2016.
- [18] C.N.J. Meunier, P Chameau, and P. M. Fossier. Modulation of Synaptic Plasticity in the Cortex Needs to Understand All the Players. *Front. Synaptic Neurosci.*, 9(February):2, 2017.
- [19] E.K Miller and Fusi S. Limber neurons for a nimble mind. *Neuron*, 78(2):211–213, 2013.
- [20] J.D. Murray, A Bernacchia, N.A Roy, C Constantinidis, R Romo, X Wang, S Fusi, and J.C. Martinez-Trujillo. Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl. Acad. Sci.*, 114(2):1–6, 2017.
- [21] A. Peters, C. Sethares, and J. I. Luebke. Synapses are lost during aging in the primate prefrontal cortex. *Neuroscience*, 152(4):970–981, 2008.
- [22] R Romo, C D Brody, a Hernández, and L Lemus. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, 399(6735):470–473, 1999.
- [23] K Wimmer, D Nykamp, C Constantinidis, and A Compte. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.*, 17(3):431–439, 2014.
- [24] J Zylberberg, J.T. Murphy, and M.R. DeWeese. A sparse coding model with synaptically local plasticity and spiking neurons can account for the diverse shapes of V1 simple cell receptive fields. *PLoS Comput. Biol.*, 7(10), 2011.