

# Thalamic Modulation of Memory in Recurrent Networks

P Stratton<sup>1,2</sup> ([pstratt@mit.edu](mailto:pstratt@mit.edu))

M Halassa<sup>2</sup> ([mhalassa@mit.edu](mailto:mhalassa@mit.edu))

<sup>1</sup> Queensland Brain Institute, The University of Queensland, St Lucia, QLD, Australia.

<sup>2</sup> Dept of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA.

## Abstract:

During delay tasks, some neurons in the murine thalamocortical system (Schmitt et al., 2017) and hippocampus ('time cells') (MacDonald et al., 2011) display transient spike responses with timing that is repeated reliably across trials. In higher mammals and primates, activity in some cells is consistently elevated. These transient responses during delays confer a short term memory of the stimulus. We wondered what neural network structures could facilitate the generation of such dynamic memory patterns. We show that in a simplified formalism of a dynamic recurrently-connected network (DRN), the number of unique dynamic patterns grows exponentially with network size. The DRN formalism emphasises the role in neural function of transient yet repeatable dynamics. Unlike reservoir networks, the connectivity matrix does not need to be finely tuned (random connectivity suffices), and the dynamics implement indefinite (not fading) memory. Gating of input patterns is assumed to be controlled by modulatory signals from the thalamus. In particular, recent experimental evidence suggests that inputs from the MD thalamus convey contextual information and can modulate cortical synaptic strengths. We show in a spiking neural network model that MD modulation of synaptic strength can indeed stabilize dynamic patterns of activity and hence short term memories.

**Keywords:** thalamus; short term memory; context; dynamics

## Methods

For initial investigations, we used a reduced network model that iterated activity through a random matrix (Fig 1). This allowed us to quickly test the effects of network size, activity level, and connection sparsity on:

1. The number of discrete reachable states.
2. The lengths of the generated sequences of activity.

Subsequently, we used a randomly-connected recurrent spiking network to show that the result also held for realistic networks of neurons. The simulated network contained 1000 integrate-and-fire (I&F) excitatory neurons (resting potential = 0, spiking threshold = 1, membrane time constant = 20 ms, connection probability = 0.5, connection weight = 0.01, synaptic time constant = 20 ms) and 250 I&F inhibitory

neurons (resting potential = 0, spiking threshold = 1, membrane time constant = 10 ms, connection probability to and from excitatory neurons = 0.05, connection weight from excitatory neurons = 0.2 with time constant = 5 ms, and to excitatory neurons = -0.2 with time constant = 3 ms). Network activity was initiated at  $t=0$  with spikes in a random selection of excitatory neurons, after which deterministic chaotic activity was self-sustaining.

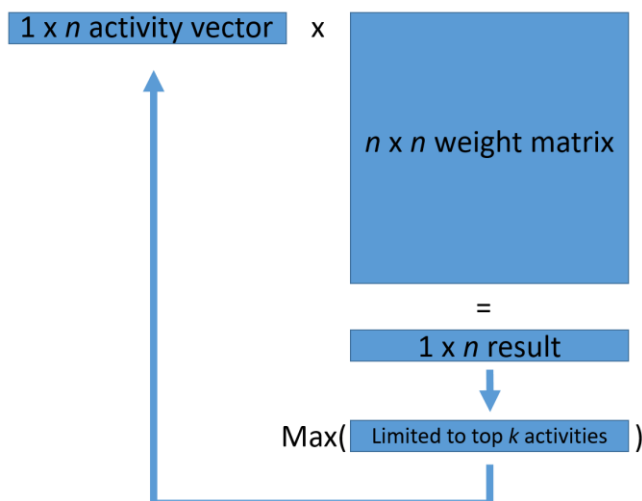


Figure 1. Reduced model. A network of  $n$  neurons is recurrently connected through an  $n \times n$  weight matrix. An iteration through the depicted update loop proceeds as follows: 1) A  $1 \times n$  activity vector represents all the neurons that just spiked as 1's and the remainder as 0's. 2) The activity vector is multiplied by the weight matrix to obtain the  $1 \times n$  result. 3) The top  $k$  neurons are deemed to emit a spike with the remainder assumed to be silenced by lateral inhibition (max() function). 4) The activity vector is updated with the neurons that just spiked and the next iteration follows.

## Results

### Reduced Model

Using the reduced model, we first investigated networks of between 5 and 500 neurons, with between 0.1% and

10% of neurons active at each step (1 step was equivalent to once around the activity update loop depicted in Fig 1). For these results, the network connectivity was set to 100% (all-to-all).

The theoretical maximum number of unique activity patterns for any given network is  $n \text{ choose } k$ , where  $n$  is the number of neurons and  $k$  is the number active at each step. For a network of  $n=500$  neurons with 10% active (i.e.  $k=50$ ) there are more than  $10^{60}$  possible patterns. 10% activity equates to a 5 Hz average firing rate if we assume there are 50 steps each second. At this rate, even for such a relatively small network, it would take approximately  $10^{40}$  universe lifetimes to discover all the patterns. However not all patterns are dynamically accessible; that is, they cannot all be reached by cycling through the weight matrix as in Fig 1. This is because the number of theoretically possible states grows combinatorially with the number of neurons, but the number of weights only grows quadratically, so the number of reachable states is significantly less than the naïve maximum, and the discrepancy is larger for larger networks. Additionally, states are not independent since all state transitions utilize the same weight matrix. To test these ideas, we generated all possible patterns for networks where the total number of possible patterns was  $10^7$  or less, and iterated them through a random weight matrix to see how many of the transformed patterns were dynamically accessible. Fig 2 (top) displays this as a percentage of all possible patterns for each network. We can see that the percentage of all possible patterns that were dynamically accessible decreased for larger networks, but that the decrease was significantly slower for sparser activity patterns. Sparse activity patterns therefore more efficiently use the range of the dynamically accessible state space of recurrent networks. To the best of our knowledge, this is a previously unrecognized advantage of sparse activity in recurrent networks.

To test the length of the sequences that could be generated, for each reduced-model network we created 1000 random starting patterns, then iterated each pattern through a random weight matrix 100 times (to simulate 2 seconds of activity at 50 steps/sec). If a pattern was repeated during these iterations, this was deemed a conflict, the iteration was stopped and the length of the cycle up to the conflict was noted. Fig 2 (bottom) shows the *minimum* cycle length for each network over the 1000 starting patterns. To reliably retrieve memories from dynamic patterns of activity, it is important that no cycles intersect within the required memory storage time. We can see that a network of 500 neurons with 1% activity (ie  $k=5$ ) had a minimum cycle length of only 100 steps (Fig 2 bottom, light orange markers). However, by extrapolation, simply doubling the number of neurons (ie to  $n=1000$ ,  $k=10$ ) increased

the minimum cycle length to approximately  $10^6$  (light orange fit line). Assuming 50 steps/sec, this equates to at least 5 hrs of continuous unique non-overlapping pattern generation using random synaptic weights. Extrapolating further, approximately only 1600 neurons would give a minimum pattern length of  $10^{11}$ , which is sufficient to produce 50 unique patterns every second for the average human lifespan of 75 years.

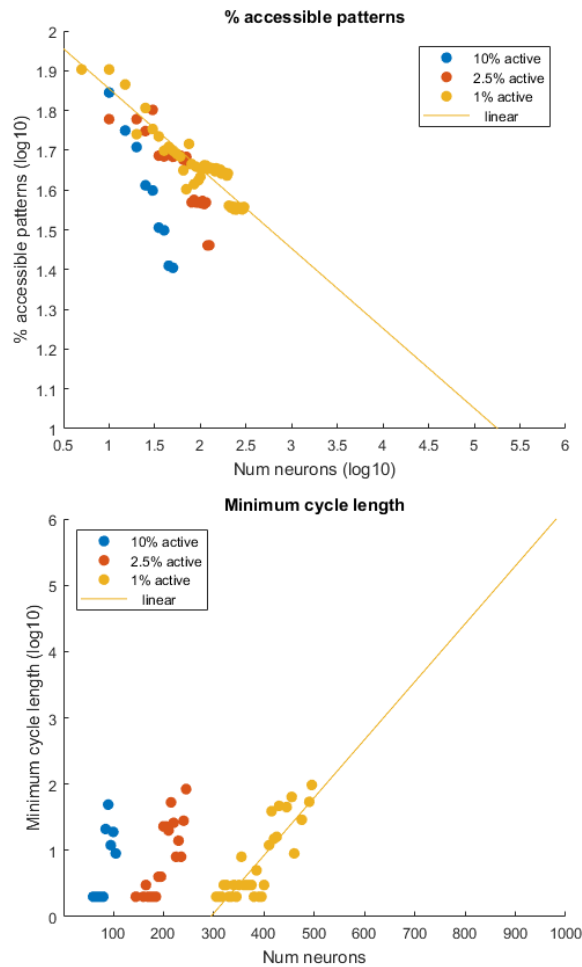


Figure 2. Effects of network size and sparse activity level on patterns of dynamic activity.

### Sparse Connectivity

Next we investigated the effects of sparse network connectivity. Sparse connectivity greatly reduced the minimum cycle length (Fig 3) and the percentage of accessible patterns (data not shown). However, the difference was smaller for networks that simultaneously used sparse activity patterns with sparse connectivity (1% active neurons per step – Fig 3 left) compared to those that used dense activity patterns (10% active neurons per step – Fig 3 right, where, for example the difference between 0.01 connected and 1.0 connected networks was much larger). This is potentially another

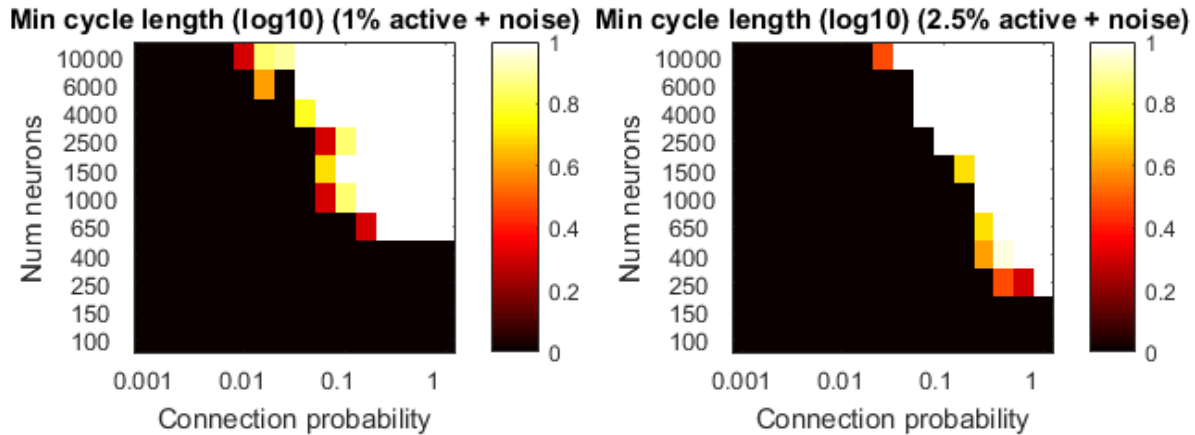


Figure 3. Minimum rescued cycle length with 10% noise (i.e. 1 noise spike for every 10 memory spikes) as a function of connection probability and number of neurons; dark regions show where patterns were not rescued (i.e. succumbed to noise) while light regions show where patterns were rescued when relevant connections were modulated up 2.0x.

advantage of sparse activity – relatively more of the dynamic state space is accessible when under a sparse connectivity constraint.

### Noise and Thalamic Modulation

The addition of noise spikes to the activity patterns immediately caused the patterns to deviate from their desired trajectories, which would clearly be catastrophic for the dynamically-maintained memories. However, as long as activity patterns were sparse and not all neurons were involved in any given trajectory, a desired trajectory could be rescued by increasing the weights of the connections between all the neurons involved in that trajectory (Fig 3). There is evidence that this transient weight increase is controlled by modulatory connections from the MD thalamus (Schmitt et al., 2017). The thalamus appears to provide a context signal which selectively increases the weights between neurons involved in maintaining the memory that has been learned to be relevant for that context. All synaptic connections that are involved at any time in a given trajectory are potentiated simultaneously. Sparse activity is therefore a critical component in rescuing noise-induced degradation of the memory patterns, since sparsity allows weight modulation to have a selective rather than a broad effect on network dynamics (i.e. if all neurons were involved in a given trajectory, then weight modulation would increase all synapses, and would not be selective for just the intended trajectory).

### Sustained Activity and Thalamic Modulation

Unlike in rodents, where memories are stored in transient activity patterns, memories in higher mammals and primates are often stored through the

consistent elevation of activity in selected cells. We tested the ability of the reduced model to maintain sustained elevated activity in a random selection of neurons; model parameters were unchanged except that the modulated weights were constrained to the diagonally symmetric recurrent connections between the selected neurons. In this case, sparse activity and sparse connectivity were *detrimental* to the ability of the network to sustain the activity (Fig 4), since randomly selected neurons are unlikely to be connected if connections are sparse, and hence cannot sustain each other's firing. We tested this result on a realistic network of spiking neurons. A network of 1000 excitatory neurons connected with probability=0.5 and with approximately 10% of neurons active at each step (in Fig 4, the bottom-left panel contains the relevant parameters) could sustain activity in a random selection of neurons (Fig 5).

### Conclusion

These results identify the minimum sufficient neural resources for dynamically storing memories through delay periods, and reveal that, when discounting the effects of noise, even modestly sized networks can store memories uniquely for practically arbitrary durations. However, when noise is introduced, patterns are destroyed unless there are neural mechanisms in place to counteract this effect. One possible mechanism is the modulation of synaptic weights, and it is possible that projections to the cortex from the MD thalamus play this role. Overall, our work posits conserved Thalamic involvement in cortical short term memory dynamics, but that in higher mammals, recurrent local cortical feedback may result in the emergence of dynamic changes in spike rate rather than temporally sparse sequences seen in rodents.

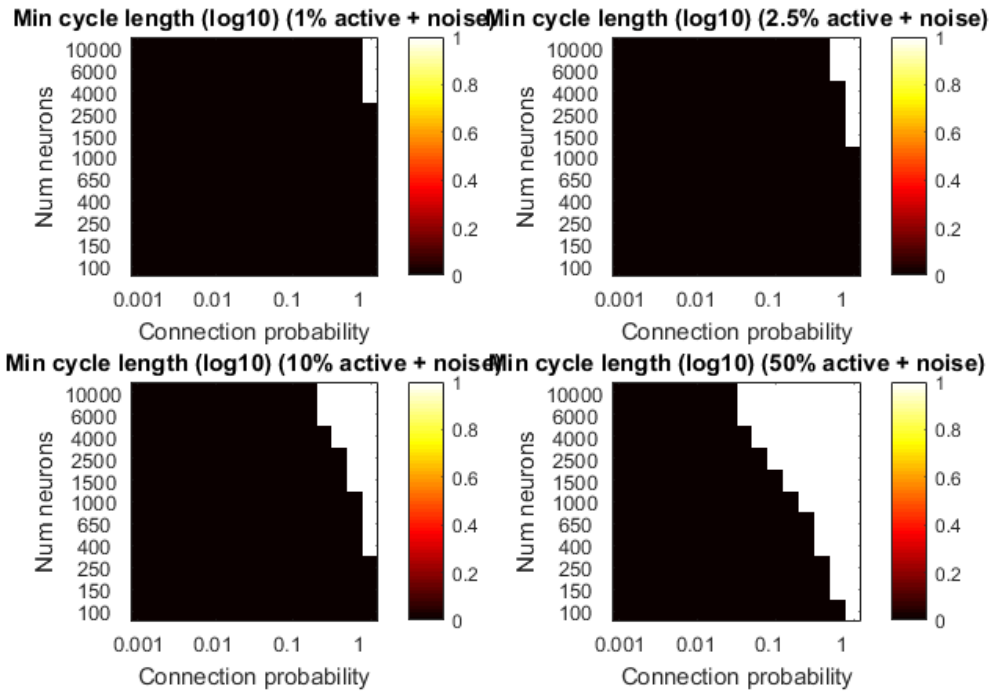


Figure 4. Ability of a network to sustain elevated activity in a random selection of neurons as a function of connection probability, number of neurons and percentage of active neurons at each step; dark regions show where patterns were not sustained while light regions show where patterns were sustained when relevant connections were modulated upwards 2.0x.

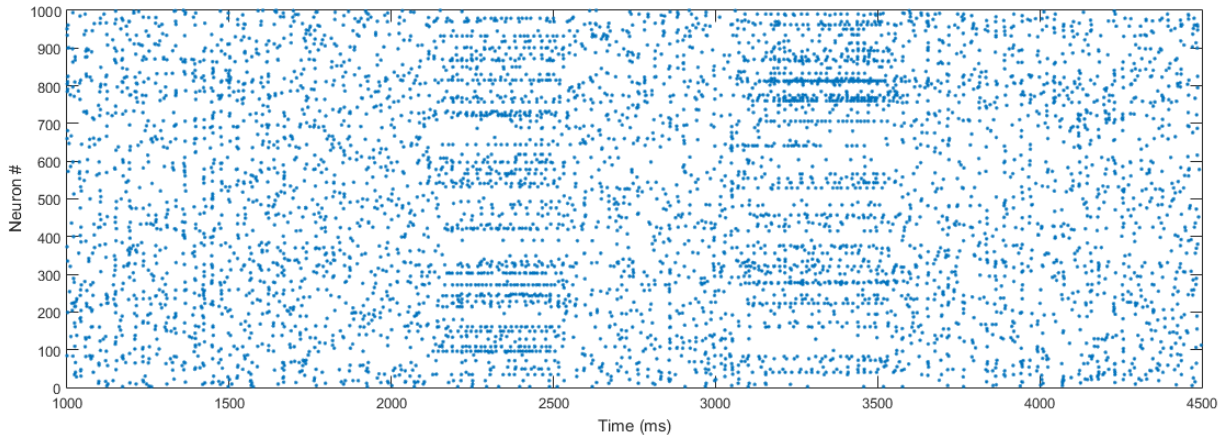


Figure 5. Thalamic modulation of connection weights causes sustained activity in a spiking network model. Connections between two randomly-selected groups of neurons were modulated up by 2.0x for 500 ms each.

## Acknowledgments

PS receives funding from the Queensland Brain Institute, Australia and the Massachusetts Institute of Technology, USA.

## References

- MacDonald C.J., Lepage K.Q., Eden U.T., & Eichenbaum H. (2011). Hippocampal “Time Cells” Bridge the Gap in Memory for Discontiguous Events. *Neuron*, 4, 737-749.
- Schmitt L.I., Wimmer R.D., Nakajima M., Happ M., Mofakham S., & Halassa M.M. (2017). Thalamic amplification of cortical connectivity sustains attentional control. *Nature*, 545, 219-223.