# Synaptic state matching: a dynamical architecture for predictive internal representation and feature perception

Saeed Tavazoie

Department of Biochemistry and Molecular Biophysics, Columbia University, 1150 St. Nicholas Ave., Room 406. New York, NY 10032.

Correspondence: Email: st2744@columbia.edu

Classification: Biological Sciences/Neuroscience

**Abstract**

Here we consider the possibility that a fundamental function of sensory cortex is the generation of an internal simulation of sensory environment in real-time. A logical elaboration of this idea leads to a dynamical neural architecture that oscillates between two fundamental network states, one driven by external input, and the other by recurrent synaptic drive in the absence of sensory input. Synaptic strength is modified by a proposed synaptic state matching (SSM) process that ensures equivalence of spike statistics between the two network states. Remarkably, SSM, operating locally at individual synapses, generates accurate and stable network-level predictive internal representations, enabling pattern completion and unsupervised feature detection from noisy sensory input. SSM is a biologically plausible substrate for learning and memory because it brings together sequence learning, feature detection, synaptic homeostasis, and network oscillations under a single parsimonious computational framework. Beyond its utility as a potential model of cortical computation, artificial networks based on this principle have remarkable capacity for internalizing dynamical systems, making them useful in a variety of application domains including time-series prediction and machine intelligence.

## Introduction

The search for function of cortical circuits has been a central focus of neuroscience since the pioneering works of Mountcastle (1), Hubel, and Wiesel (2). A guiding hypothesis behind much of this work is the existence of a canonical microcircuit whose basic computation is critical to sensory processing throughout the cortex. Detailed micro-architectural maps (3) show great promise in advancing our understanding of these circuits. However these efforts will greatly benefit from constraints on the nature of the computational task itself. At the highest level, cognitive systems generate internal representations of the outside world that facilitate adaptive behaviors. To what extent can this high-level truism inform us about neural architecture and computation operating at much lower levels? Here I show that a logical extension of this principle, down to the scale of individual synapses, naturally leads to a dynamical neural architecture that generates predictive internal representations, enabling feature detection from noisy sensory input.

To bridge the gap between the abstract notion of internal representations and their physical neural substrate, let us consider *dynamic* internal representations which, in their perfect form, are able to internally generate the spatiotemporal dynamics of sensory input-in essence, an internal simulation of the sensory environment. Here, sensory input is considered from the perspective of cortex (e.g. activity of Lateral Geniculate Nucleus (LGN) input neurons into V1). Such inputs typically preserve relevant topographic information in two dimensional space (e.g. retinotopy).

How do populations of neurons form a faithful simulation of spatiotemporal dynamics of sensory input? The most obvious substrate is Hebbian synaptic plasticity (4) and its modern variant, spike-timing dependent plasticity (STDP) (5, 6), where a pre-synaptic action-potential (spike) that immediately precedes a post-synaptic spike strengthens the synapse, while one with the opposite temporal order weakens it. In principle, this form of synaptic strength modification would lead to a causal chain of activity reflecting the input pattern. However, learning in model neural networks based on this form of plasticity has clear stability issues since the inherent positive feedback will increase synaptic strength without bounds. Various heuristics have been proposed to address the stability problem by imposing additional constraints on the scale and rate of synaptic potentiation, often utilizing non synapse-local information (5, 7). The biological plausibility of specific solutions may be in doubt. However, it is clear from experimental evidence that synaptic plasticity strength is somehow homeostatically modulated, presumably in order to maintain neural networks within stable operational bounds (8).



**Results**

Arguing from first principles, I will motivate a parsimonious neural architecture design capable of simulating dynamical systems through an inherently stable synaptic modification process operating on strictly local information. Our discrete implementation is composed of a population of neurons with potential all-to-all synaptic connectivity mediated by axons with conduction delay. For the sake of simplicity and symmetry, neurons can form both positive (activating) and negative (inhibitory) synapses onto each other. The neurons are generic integrate-and-fire variety whose composite synaptic input determines spiking based on a

thresholded sigmoidal response. Sensory input is applied to every neuron in the form of a sequence of spike trains. Synaptic weights are initially set to zero. The goal is to arrive at a distribution of synaptic weights that enables the network to generate sensory input patterns autonomously. Some form of Hebbian synaptic plasticity could serve as the driving force here. However, we have to ensure that: 1) modifications will lead to robust convergence of synaptic weights to a distribution (*a priori* unknown) that generates accurate predictive internal representations and 2) potentiation will remain within bounds (*a priori* unknown) that ensures long-term network stability.



**Synaptic state matching**

Let us focus down on an individual synapse which has access only to local information on pre and post-synaptic spiking events and their temporal correlation-structure over time. From the perspective of this synapse, perfect internal representation of sensory input is achieved when these observable spike statistics are identical between two proposed states: 1) network activity is generated by sensory input (open), and 2) network activity is generated by recurrent synaptic connections in the absence of sensory input (closed). I propose that the imposition of this synaptic state-matching (SSM) constraint at individual synapses can give rise to accurate and stable internally generated dynamics at the level of the entire network.

To ensure smooth convergence and long-term stability, synaptic state matching occurs in real-time (**Fig. 1*A***), with the entire network switching back-and-forth between activity imposed by sensory input (*open state*) and that generated by the internal simulation of the most likely sensory input (*closed state*). At each synapse, a presynaptic spike followed closely by a

postsynaptic spike will lead to specific biochemical changes with the potential of strengthening the synapse. Let us refer to the strength of this effect as 'potentiation strength'. Let us further assume that each synapse has the capacity to store a running average estimate of its potentiation strength in each of these two distinct states. During the early phase of learning, the closed state is quiescent because weak synapses are unable to generate activity, and potentiation strength computed during this phase is low. During the open state, the strong drive of sensory input generates spiking activity with high potentiation strength among synapses with tightly coupled pre-post spiking. Upon the occurrence of a candidate potentiation event, the synapse makes a decision: the stored mean potentiation strength is compared between the two states; if it is higher in the open state, the synapse is strengthened proportional to the instantaneous value of potentiation strength; otherwise, if the mean potentiation strength is higher in the closed state, then the synapse has clearly overshot and therefore will correct itself by undergoing a depression penalty with a scale proportional to its current mean potentiation strength during the closed state. This simple local algorithm ensures that potentiation occurs until pre-post spiking statistics are matched between open and closed states. From the perspective of an individual synapse, this is the best local indication that recurrent synaptic activity is recapitulating the spatiotemporal dynamics of sensory input.

Potentiation strength is computed by a temporally asymmetric Hebbian process that is modulated by the rates of pre- and post-synaptic activity (**Fig. 1*B*, Fig.S1** and **Methods**). We call this plasticity rule spike-timing dependent covariance plasticity (STCP) because it is a temporally asymmetric, event-driven, version of covariance learning (9). In this way, the

relative strength of potentiation is proportional to the statistical support for the causal chain of activity from pre-synaptic to post-synaptic neuron. As we show below, this simple local algorithm ensures smooth convergence to synaptic weight distributions that generate accurate and stable dynamic internal representations of sensory input.

**Predictive internal representations**

Simulation of a small SSM network demonstrates how predictive internal representations are formed (**Fig. 2**). Details of implementation are available in **Methods**. For the sake of demonstration, the input spike trains form a triangular wave pattern with a period of 26 time-steps. Also see **Fig S2** for an example of sensory input in the form of repeating circles. The thirty neurons of the network in **Fig. 2** are driven by sensory input during the open state (Gaussian with mean 15 and s.d. 5 time-steps), and free to generate autonomous activity during the closed state (mean 15 and s.d. 5 time-steps). Over repeated presentations of the stimulus, plasticity drives synaptic weights to a distribution that leads to generation of internal dynamics (**Fig. 2*A***). As can be seen, the internally generated pattern is a perfect simulation of the missing input, in essence a short-term prediction of the sensory environment (**Fig. 2*B***). In this simulation, every neuron has the potential capacity to interact with every other neuron through connections with time-delay (latencies). The range of latencies (1-5 time-steps) reflects both axonal conduction and synaptic transmission delays that can vary substantially due to axonal diameter/myelination and synapto-dendritic time-constants.

**Robustness to parameter choice, sensory noise, and synaptic perturbations**

The learning performance is remarkably robust to large variations in almost every parameter, including state-switching frequency, neuron firing threshold, range of latencies, and global potentiation scale (**Table S1**). Strikingly, potentiation scale can vary over many orders of magnitude and still give rise to highly accurate and stable dynamics. The key to this remarkable robustness is the synaptic state matching constraint that, although operating locally at individual synapses, ends up imposing absolute constraints on the fidelity of network-level internal representation (**Fig. S3**). One reflection of stability is the near balance of positive and negative synaptic drive into each neuron across three orders of magnitude in potentiation scale (**Fig. 2*C***). Experimental evidence suggests that this matching of inhibitory and excitatory drive is a common operating characteristic of neural networks *in vivo* (10, 11).

Robustness to input noise and structural perturbations are desirable features of biological neural networks. SSM networks show remarkable noise-resilience, forming accurate and stable internal representations in environments with signal-to-noise ratios (SNR) of ~ 1 (**Fig. 3*A***). As can be seen, this allows SSM networks to perform perfect noise-filtering for stored patterns. Structural perturbations in the form of random synaptic ablation are also well tolerated. In a SSM network trained with the triangular-wave input, thirty percent of synapses can be removed without substantial degradation in performance (**Fig. 3*B***).

**Feature detection and spatiotemporal maps**

During the closed state, the network internally generates the most likely input activity pattern, conditioned on the imposed sensory input during the preceding open state. As such, neural

activity during the closed state is a short-term prediction of sensory environment. However, the network will not predict the trajectory of sensory environment for all stimuli. Indeed, only recurrent patterns with high statistical regularity have the potential of being internalized within the network's synaptic matrix. As such, SSM networks can carry out feature-detection because any activity during the closed state must correspond to pattern completion in response to a *recognizable* feature presented during the preceding open state. In this way, SSM networks filter sensory environment for recurrent features on the timescale of state-switching. An illustrative example of this can be seen for spatiotemporal patterns that represent edges (lines) moving in different directions within a simulated visual field (**Fig. 4*A*, Fig. S4**). In this example, the input neurons are arranged in a 20 by 20 plane preserving 'retinotopy'. The single 400 neuron SSM network efficiently learns these four patterns and performs robust pattern completion in the presence of noise (**Fig 4*A***).

One can think of this small SSM as corresponding to a single cortical column in area V1. The ability of this column to identify a feature relies on a long-enough temporal sequence of spikes observed during the open state in order to recall the appropriate stored pattern. Naturally, this means that closed-state activity is biased towards the end of the temporal sequence of spikes that constitute a feature. When there is a topographic projection of sensory input into cortex, this temporal asymmetry in the recall of internal representation translates to a spatial bias in the response of neurons to detectable features. This phenomenon is clearly seen in our simple conception of a V1 cortical column (**Fig. 4*A***). Mapping the response specificity of this patch of 'cortex' to bars moving in different directions, produces receptive field maps with clear spatial preference, reminiscent of those observed

experimentally (12) (**Fig. 4*B***). It should be emphasize that feature learning and detection happen by an entirely unsupervised process with intrinsic long-term accuracy and stability.

**Multi-layer networks**

The single-layer networks, presented so far, are directly driven by (clamped to) the external input. In this setting, synaptic state matching leads to a generative model that captures first and second order correlations in the spatiotemporal dynamics of sensory input. In the realm of machine learning, additional internal hidden layers, with neurons whose activities are not directly clamped to external input, are free to capture higher-order statistical regularities in the distribution of static input patterns (13). The representational power of SSM networks can be similarly expanded by addition of hidden layers. For example, large hidden layers can capture complexity in a long sequence of spike trains (**Fig S5**). In turn, small hidden layers can carry out dimensionality reduction by forming compressed predictive internal representations of high-dimensional sensory input (**Fig. S6**).

**Discussion**

I propose that a fundamental function of sensory cortex is the generation of an internal simulation of sensory environment in real time. Here, we see that a parsimonious neural architecture can achieve this through synaptic state matching driven by a spike-timing dependent covariance plasticity rule. Synaptic state switching has superficial similarities to a phase-switching algorithm used to train Boltzmann machines (13). As a stochastic extension of Hopfield's associative networks (14), Boltzmann machines have shown great success in modeling the probability distribution of static input patterns. During learning, the stochastic

units are clamped to temporally uncorrelated input patterns during one phase (learning), and allowed to run freely in the second phase (unlearning). The goal of each phase is to accurately sample pair-wise correlation distributions which, because of the necessity for reaching thermal equilibrium, can take extremely long timescales to accomplish (13). As such, these training phases have little correspondence to real-time state switching and matching in the sense described here.

Real time state matching is crucial for stability and accuracy (**Fig. S3**), but it also provides a convenient substrate for rapid and noise-resilient feature detection. In this way, the SSM network efficiently filters high-dimensional sensory input and generates lower-dimensional feature perceptions that, can in turn, serve as input for another SSM network higher in the hierarchy. Low-level SSM networks, with high-frequency state-switching, perceive features at small spatiotemporal scales; lower-frequency state-switching at higher levels of the hierarchy can support feature perception at larger spatiotemporal scales. In turn, top level SSM networks can project down and bias feature detection at lower levels of the hierarchy. This view exposes a system-intrinsic definition of an 'object' as an embedded pattern of sensory input that gives rise to internally generated activity (i.e. recognition). This is a behaviorally salient perspective since any internally generated activity is limited to spatiotemporal patterns with high statistical regularity which, in turn, are the most likely to convey information about future reward.

Do experimental observations support the existence of state-switching? An obvious substrate is network oscillations that span a wide dynamic range in frequency and are ubiquitous in

brains (15) including those of insects (16). The required gating of sensory input could be regulated locally by these oscillations, or more centrally, perhaps through thalamic gating (17). Cortical oscillations are thought to underlie coordination and synchrony at large spatiotemporal scales. Synaptic state matching would extend the role of oscillations down to the smallest neurobiological scales, at individual synapses and single potentiation events. If the phase of the oscillation defines the open-closed state boundary, SSM would predict that candidate potentiation events are interpreted in a highly phase dependent manner, with opposite effects on synaptic strength. Indeed, this is exactly what was observed in an experimental imposition of activity during an LTP/LTD protocol (18). Bursts of spikes delivered during the peak of the theta oscillation induced LTP, and those delivered during the trough induced LTD (18).

It has been elegantly argued that prediction is an operating principle throughout the brain (19). Indeed, recent work has shown that even microbial regulatory networks are capable of predictive behavior (20). Here we see that an atomic prediction principle, locally operating at individual synapses, generates stable and accurate network-level dynamical models that extract recurrent features of sensory environments. SSM is an attractive biological mechanism because, in addition to its simplicity, it is remarkably tolerant to noise, structural perturbations, and choice of parameters. Beyond their relevance to cortical information processing, networks based on SSM can serve as general purpose devices in a variety of application domains including sequence learning, sequence generation, feature/object detection, time-series prediction, noise-filtering, and signal decomposition. From the

experimental perspective, the search for molecular and cellular correlates of synaptic state matching represents an important area for future research.

**Acknowledgements**

The author wishes to thank John Hopfield for insightful comments on an earlier version of the manuscript. The author is also thankful to Sohail Tavazoie and Masoud Tavazoie for providing critical feedback during the early phase of the project.

## Methods

### SSM network implementation

SSM networks were simulated in discrete-time using Matlab (Mathworks). Each neuron receives strongly activating synapses from a single input axon. Therefore, during the open state, the activity of each neuron is clamped to the input state and follows it precisely (**Fig. 1*A***). During the closed state, input is gated off and neuronal activity is determined by recurrent synaptic drive. In the simplest implementation, every neuron can synapse onto every other neuron through connections with a range of time-delays (e.g. 3-5 time-steps). In biological neural networks, inhibitory synapses are formed by specialized inter-neurons. In our simulations, neurons can directly form both activating and inhibitory synapses. These simpler symmetric networks are essentially equivalent to ones where each excitatory neuron imposes a dominant drive on an inhibitory neuron with broad connections in the network.

### Neuron model

For each neuron, contributions from all synapses are added linearly proportional to their weights and a sigmoidal function with a threshold determines spiking:

$$V_i = \frac{\left[tanh\left(S \cdot \sum_{j=1}^{N} w_{ij} X_j - \frac{1}{2}\right) - 1\right]}{2}$$

$$X_i = 1 \quad if \quad V_i \geq threshold$$

$$X_i = 0 \quad otherwise$$

Here, $V_i$ is the membrane voltage of neuron $i$ . $w_{ij}$ is the synaptic weight (strength) of connection from neuron $j$ to neuron $i$. $X_j$ is the spiking state of neuron $j$. $N$ is the total number of neurons and parameter $S$ determines the sharpness of the hyperbolic tangent sigmoid function. If the membrane voltage reaches threshold, the neuron fires.

**Synaptic plasticity**

Synaptic potentiation occurs during the open state according to a spike-timing dependent covariance plasticity rule (**Fig. S1**). Each synapse maintains a local accounting of pre- and post-synaptic spike rates, and of the time-averages of potentiation strength computed under both open and closed states. The 'aim' of each synapse is to efficiently drive the strength of the synapse to a value that leads to a close matching of the mean potentiation strength during open and closed states. This is each synapse's best attempt at producing an internal representation given that it only has access to local information. Parameters spike-rate memory and potentiation memory determine the time-scale of the running average for spike-rate and potentiation strength respectively.

If the scale of the step-wise change in synaptic strength ($\alpha$) is large, the synapse may overshoot its optimal weight. This may be reflected in a mean potentiation strength that is higher in the closed state as compared to the open state. In such a scenario, the synapse will undergo a weight depression penalty proportional to the mean potentiation strength during the closed state. This local homeostatic mechanism ensures long-term synaptic state matching which, in turn, guarantees long-term accuracy and stability (**Fig. S3**).

## References


1. Mountcastle VB (1957) Modality and topographic properties of single neurons of cat's somatic sensory cortex. *J Neurophysiol* 20(4):408-434 (in eng).
2. Hubel DH (1959) Single unit activity in striate cortex of unrestrained cats. *J Physiol* 147:226-238 (in eng).
3. Bock DD*, et al.* (2011) Network anatomy and in vivo physiology of visual cortical neurons. *Nature* 471(7337):177-182 (in eng).
4. Hebb DO (1949) *The organization of behavior; a neuropsychological theory* (Wiley, New York,) pp xix, 335 p.
5. Abbott LF & Nelson SB (2000) Synaptic plasticity: taming the beast. *Nat Neurosci* 3 Suppl:1178-1183 (in eng).
6. Dan Y & Poo MM (2006) Spike timing-dependent plasticity: from synapse to perception. *Physiol Rev* 86(3):1033-1048 (in eng).
7. Bienenstock EL, Cooper LN, & Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2(1):32-48 (in eng).
8. Pozo K & Goda Y (2010) Unraveling mechanisms of homeostatic synaptic plasticity. *Neuron* 66(3):337-351 (in eng).
9. Sejnowski TJ (1977) Storing covariance with nonlinearly interacting neurons. *J Math Biol* 4(4):303-321 (in eng).
10. Haider B, Duque A, Hasenstaub AR, & McCormick DA (2006) Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *J Neurosci* 26(17):4535-4545 (in eng).
11. Shu Y, Hasenstaub A, & McCormick DA (2003) Turning on and off recurrent balanced cortical activity. *Nature* 423(6937):288-293 (in eng).
12. Weliky M, Bosking WH, & Fitzpatrick D (1996) A systematic map of direction preference in primary visual cortex. *Nature* 379(6567):725-728 (in eng).
13. Hinton GE, Sejnowski, T. J. (1986) Learning and relearning in Boltzmann machines. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.*, ed Rumelhart DE, McClelland, J. L (MIT Press, Cambridge, MA), Vol 1, pp 282-317.
14. Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A* 79(8):2554-2558 (in eng).
15. Buzsaki G & Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304(5679):1926-1929 (in eng).
16. Tanaka NK, Ito K, & Stopfer M (2009) Odor-evoked neural oscillations in Drosophila are mediated by widely branching interneurons. *J Neurosci* 29(26):8595-8603 (in eng).
17. Lorincz ML, Kekesi KA, Juhasz G, Crunelli V, & Hughes SW (2009) Temporal framing of thalamic relay-mode firing by phasic inhibition during the alpha rhythm. *Neuron* 63(5):683-696 (in eng).
18. Huerta PT & Lisman JE (1995) Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. *Neuron* 15(5):1053-1063 (in eng).
19. Hawkins J & Blakeslee S (2004) *On intelligence* (Times Books, New York) 1st Ed p 261 p.

20. Tagkopoulos I, Liu YC, & Tavazoie S (2008) Predictive behavior within microbial genetic networks. *Science* 320(5881):1313-1317 (in eng).

## Figure legends

**Fig. 1. Synaptic state matching architecture and spike-timing dependent covariance plasticity.** (***A***) Dynamical architecture of synaptic state switching and state matching. Neurons continually oscillate between two global network states: activity imposed by sensory input (open) and activity generated by recurrent synaptic drive, in the absence of sensory input (closed). (***B***) synaptic plasticity rule (potentiation strength) in discrete time. X and Y are binary events corresponding to pre-synaptic and post-synaptic spikes, respectively and <X> and <Y> are their continuous valued running averages. For an activating synapse (+), a pre-synaptic spike, followed immediately by a post-synaptic spike gives rise to a potential weight increase ($\Delta w>0$). In the case of an inhibitory synapse (-), a pre-synaptic spike preceding a quiescent post-synaptic neuron gives rise to a potential weight increase ($\Delta w<0$). Potentiation can only occur following a pre-synaptic spike. Candidate potentiation events lead to synaptic strength modulation such that the mean potentiation strength is matched between the open and closed states. As mean potentiation strength is a compact measure of the locally observed spike statistics, Synaptic state matching assures that recurrent synaptic drive is recapitulating the spatiotemporal dynamics of sensory input.

**Fig. 2. A thirty neuron SSM network trained on input spikes in the form a triangular wave.** (***A***) Synaptic weight matrices for different latency conduction delays (1-5 time steps). Positive values (red) correspond to activating connections and negative values (blue) to inhibitory ones. (***B***) full pattern of activity in the sensory environment (full pattern), input activity during the open state (open), internal activity during the closed state (closed), and the

combined (complete). Internal activity is a perfect match to the missing input. SSM parameter choices were as follows: potentiation strength ($\alpha$): $4x10^{-5}$, spike-rate memory ($m_s$): 100, potentiation memory ($m_p$): 100, state switching period: (Gaussian, $\tau_\mu$=15,$\tau_\sigma$=5), neuron firing threshold ($V_t$): 0.5, sigmoid sharpness (S): 10, latency range (L): 5. (***C***) Input synaptic drive into a single neuron for activating input (red), and inhibitory input (blue). Spikes are represented in green. The two simulations are identical except for a thousand-fold higher potentiation scale ($\alpha$). Each neuron has 290 synaptic inputs (29 neurons x 5 latencies x 2 polarities).



**Fig. 3. Robustness to sensory noise and synaptic perturbations.** (***A***) Patten of sensory environment with noise (full pattern), input activity during the open state (open), internal activity during the closed state (closed), and the combined (complete). (***B***) Match (accuracy) of internally generated activity to the missing input during the same interval versus fraction of random synapses removed. Accuracy is a conservative measure of how well internally generated pattern matches the missing input during the same time interval. It is defined as one minus the fraction of discordant spikes between the missing input spike trains and the internally generated spike trains. The value here is the average for all thirty neurons. Accuracy can be negative in cases where internally generated activity is noisy and/or unstable.

**Fig. 4. Feature detection and spatiotemporal maps in the context of moving bars.** A 400 neuron network arranged in a 20x20 sheet with 'retinotopic' input from the visual field is trained on bars moving left, right, up and down. (***A***) Pattern completion in response to bars

moving right (top) and left (bottom). The clean input is presented with equal power random noise during the open state. The network performs perfect pattern completion during the second closed state (dashed boxes) following the presentation of a long-enough sequence of moving bars during the preceding open state. Open/closed durations are sampled from a Gaussian with a mean of 7 and sd of 2 time steps. For the sake of space, bars moving up and down are omitted and only every other time-point is shown. (***B***) Spatial receptive field map of internally generated activity in response to bars moving in four different directions. Color-code reflects each neuron's most preferred direction. Brighter colors reflect stronger preference.

**Figure 1**

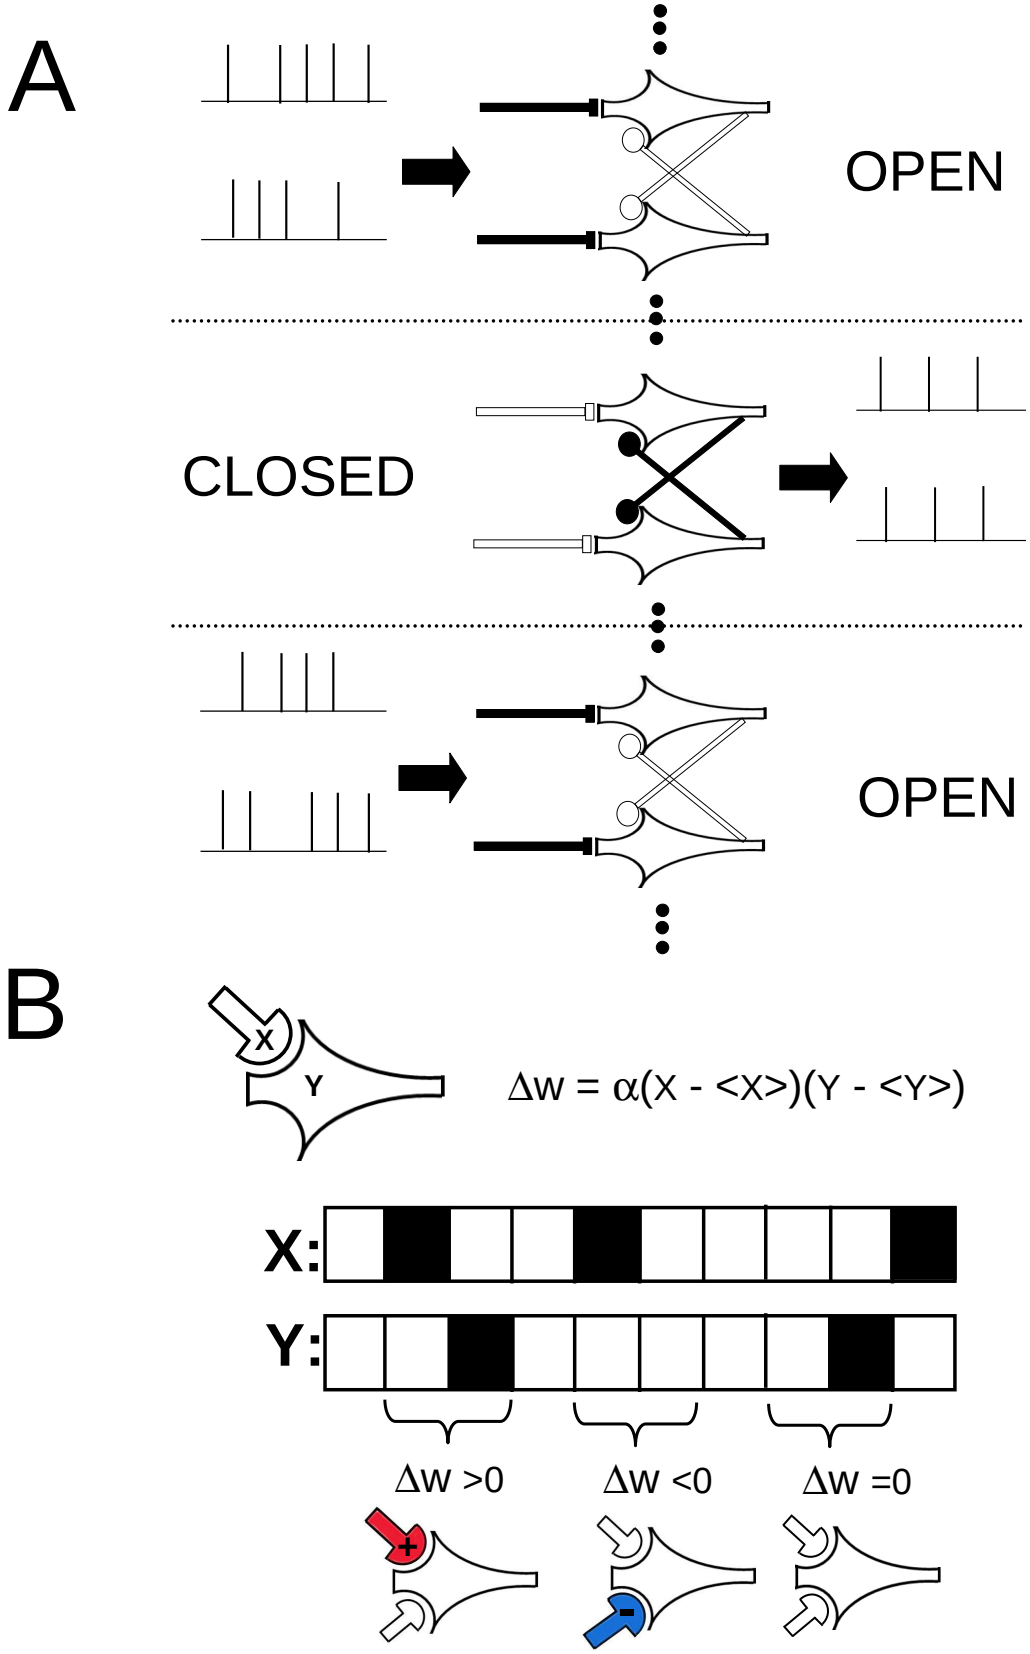

**Figure 2**

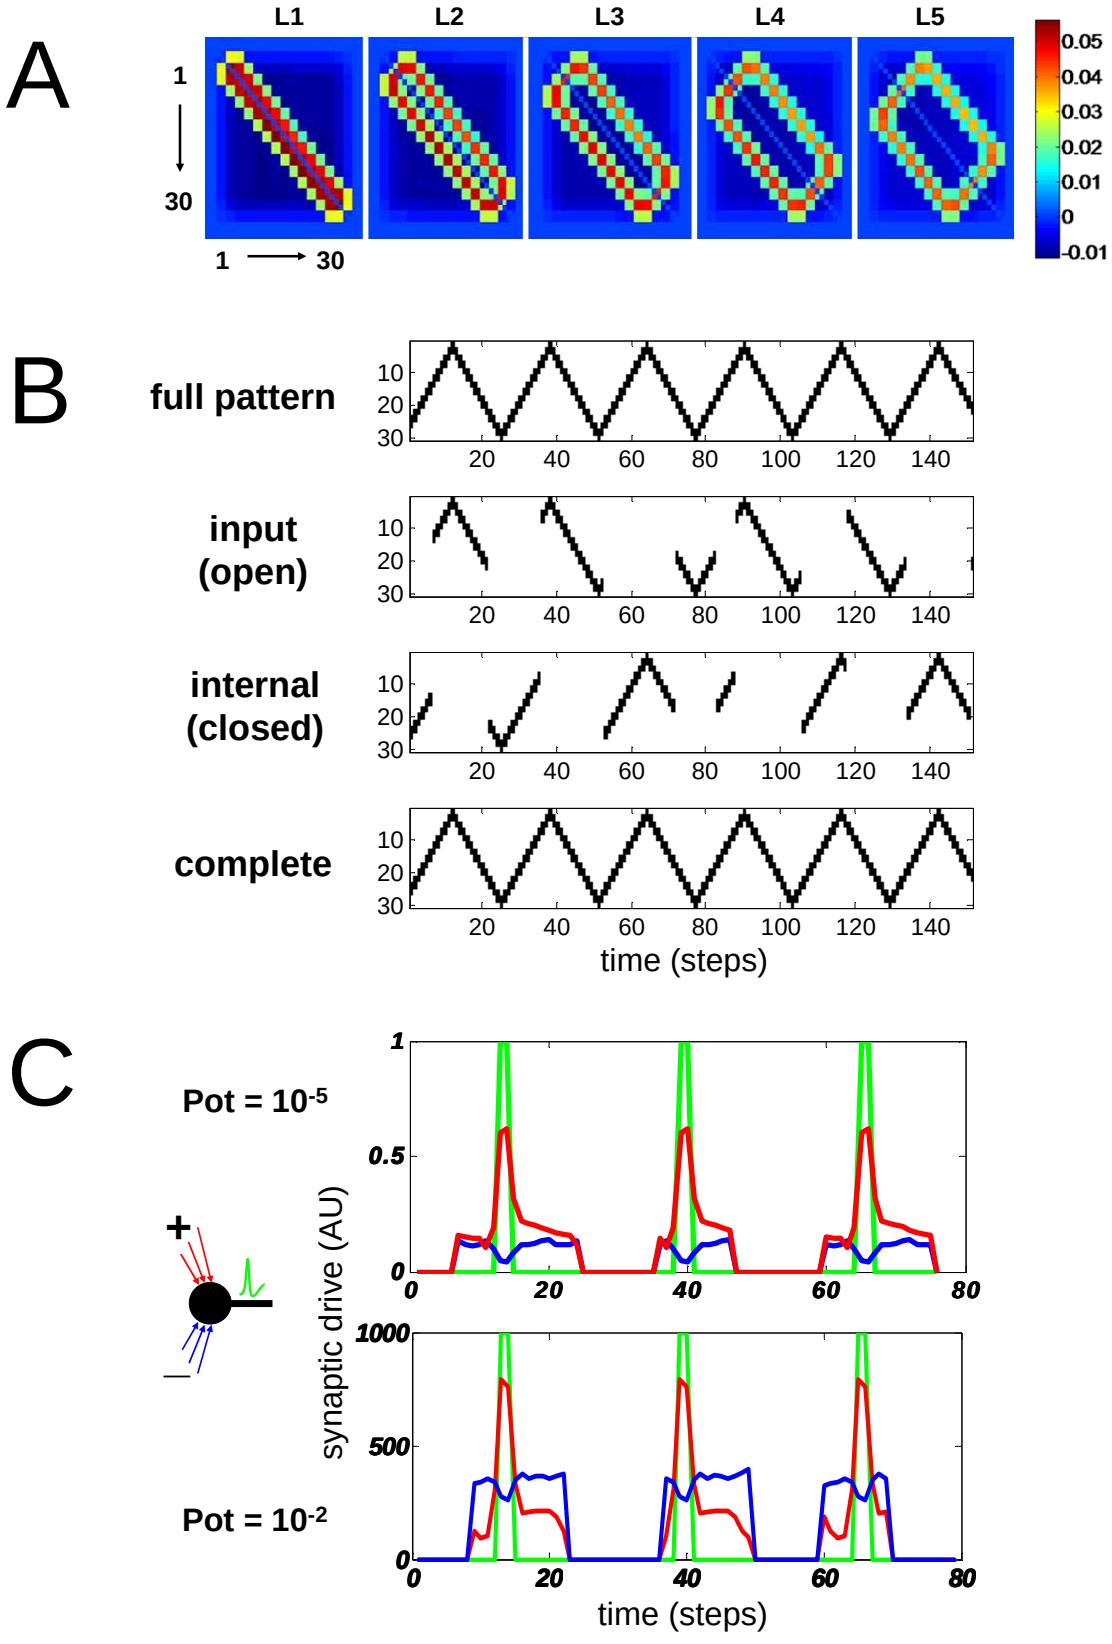

**Figure 3**

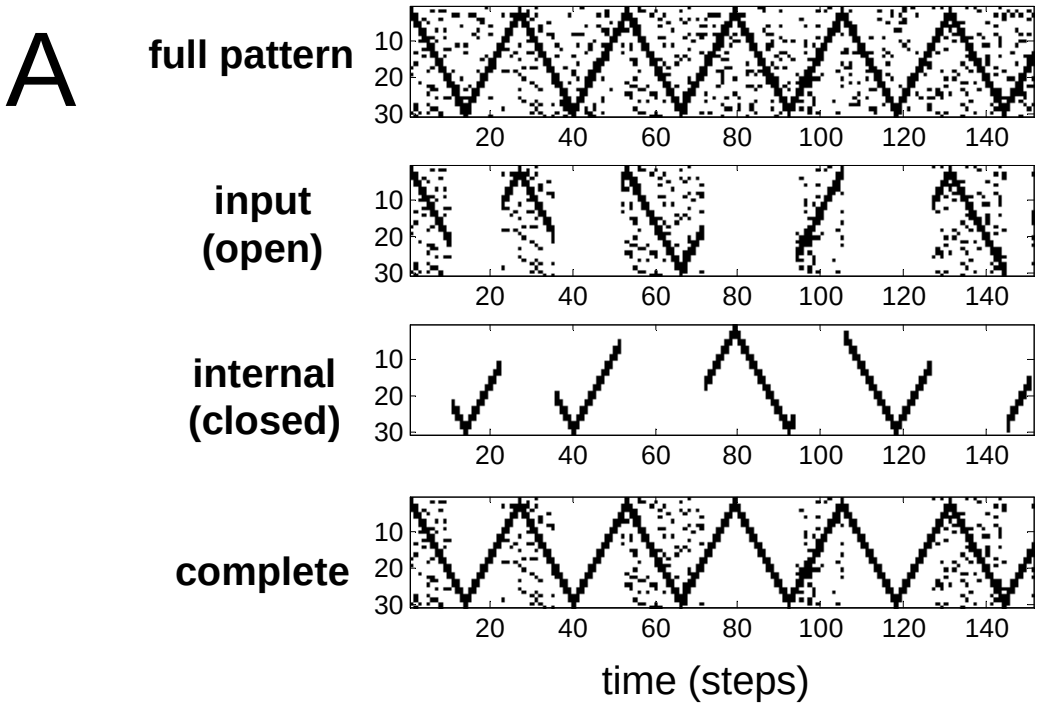


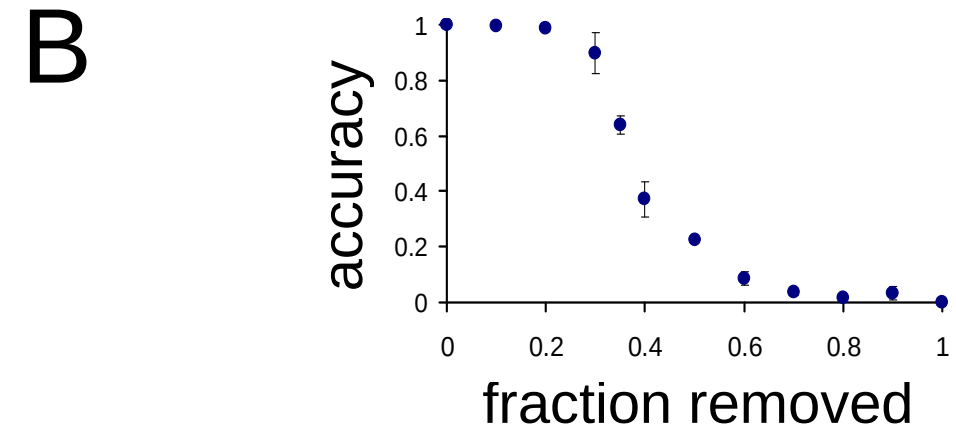

**Figure 4**

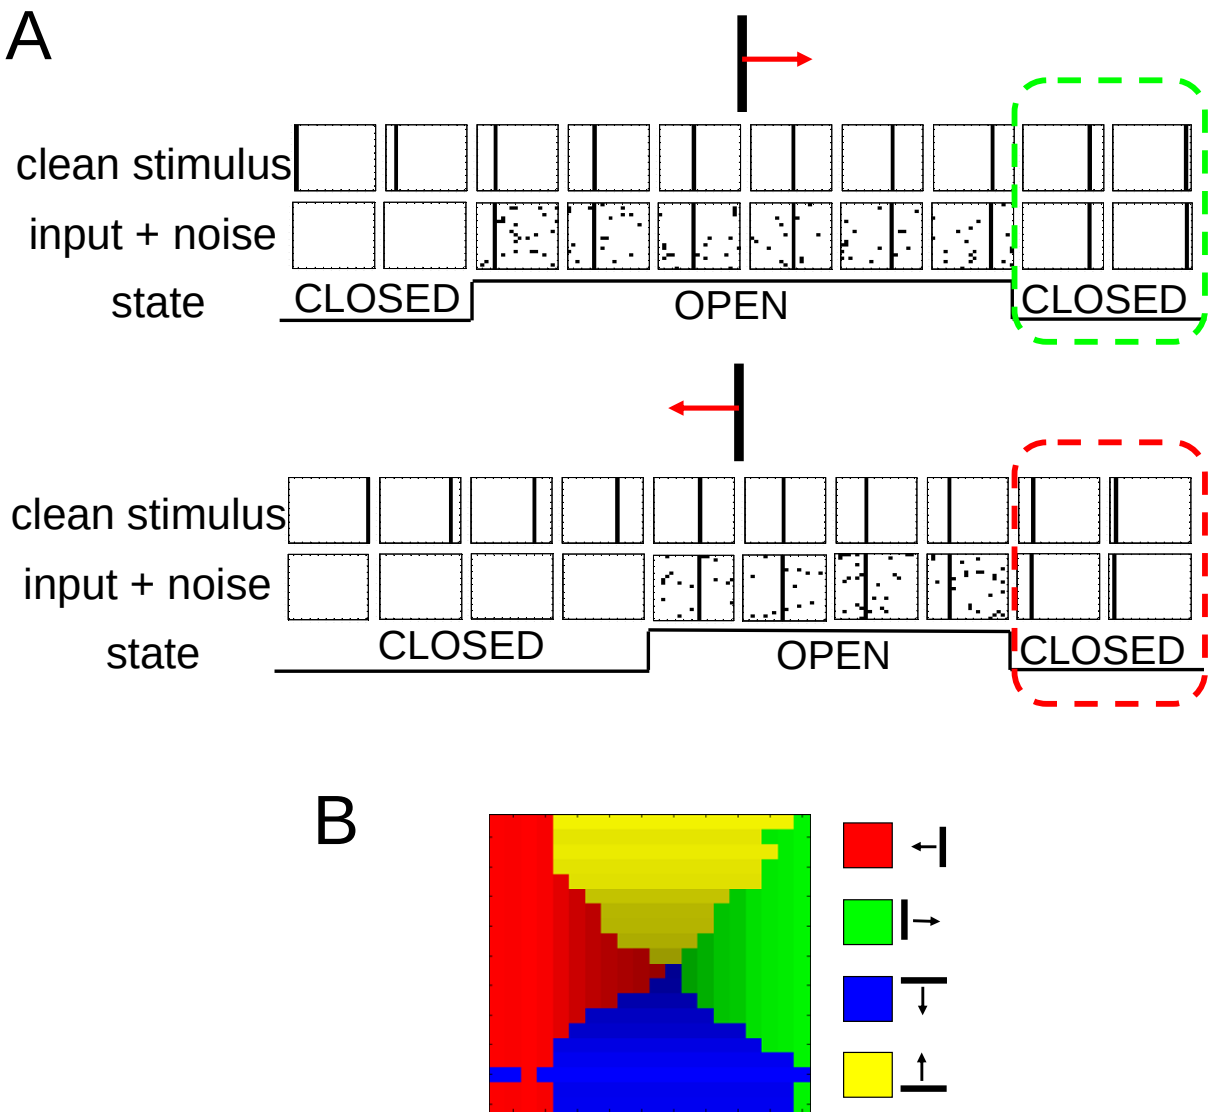

**Supplementary Table**

| Parameter | Operational Range |
|---|---|
| mean state switching period ($\tau_\mu$) | 7 - 240* |
| neuron firing threshold ($V_t$) | 0.1 - 0.7 |
| sigmoid sharpness (S) | 1 - 100* |
| latency (L) | > 3 |
| potentiation scale ($\alpha$) | $5x10^{-5}$ - 10* |
| spike-rate memory ($m_s$) | 5-5000* |
| potentiation memory ($m_p$) | 5-5000* |

**Table S1. Parameter insensitivity.** Operational ranges that maintain accuracy above 0.80 on the single triangular wave input pattern (**Fig. 2**). Asterisks denote conservative bounds because values beyond were not tested.

**Supplementary Figures**

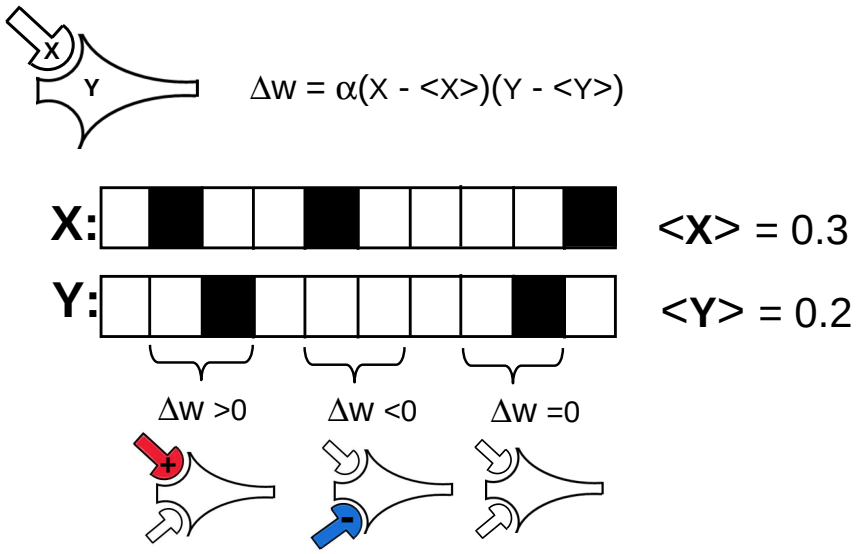


$\Delta w^{+} = \alpha(1 - 0.3)(1 - 0.2) = 0.56\alpha$

$\Delta w^{-} = \alpha(1 - 0.3)(0 - 0.2) = -0.14\alpha$

$\Delta w^{0} = 0$

**Fig. S1. Spike-timing dependent covariance plasticity (STCP)**. Following every presynaptic event (X=1), potentiation strength is computed by a temporally asymmetric event-driven Hebbian process that takes into account the average activities of presynaptic (<X>) and postsynaptic (<Y>) neurons. In this cartoon example, three potential scenarios are depicted: (1) a presynaptic spike, immediately followed by a post-synaptic spike, gives rise to a positive potentiation strength ($\Delta w>0$) with the potential of strengthening an activating synapse. (2) a presynaptic spike immediately followed by a quiescent postsynaptic neuron, gives rise to a negative potentiation strength ($\Delta w<0$) with the potential of strengthening an inhibitory synapse. (3) In the absence of a presynaptic event, potentiation strength is zero ($\Delta w=0$).

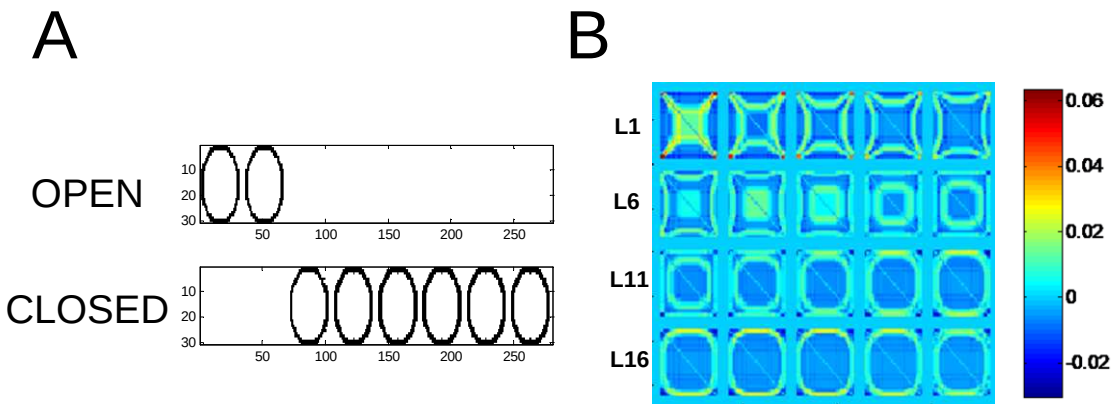


**Fig. S2. A 30 neuron SSM network trained on input spikes in the form of repeating circles.** (***A***) Input presented to the network during the open state (open). The network generates internal activity during the closed state (closed). This SSM network was trained with the potential of long latency synaptic connections (1-20). (***B***) Synaptic weight matrices for conduction-delay (latencies) of 1-20 time steps. SSM network parameter choices were as follows: potentiation strength ($\alpha$): $5x10^{-5}$, spike-rate memory ($m_s$): 50, potentiation memory ($m_p$): 50, state switching period: (Gaussian, $\tau_\mu$=15,$\tau_\sigma$=5), neuron firing threshold ($V_t$): 0.5, sigmoid sharpness (S): 10, latency range (L): 20.

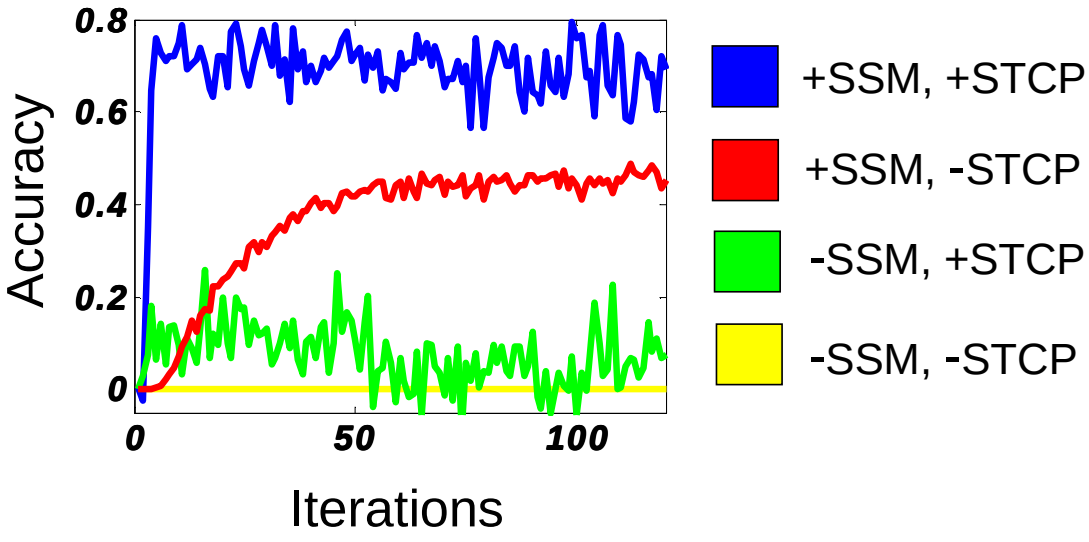


**Fig. S3. The role of SSM and STCP in the accuracy-trajectory of learning.** Synaptic state matching (+SSM) is crucial for accuracy and long term stability. Spike-timing dependent covariance plasticity (+STCP) substantially improves convergence rate and accuracy as compared to a plasticity rule that is not modulated by spike-rate history (-STCP). The 40 neuron SSM is trained on two alternating random spike patterns (features), each 40 time-steps long, with an intervening 20 time step quiescent period. The accuracy does not reach maximum (1.0) because the network is unable to generate the earliest part of each random pattern due to the absence of any input during the quiescent period preceding the closed state.

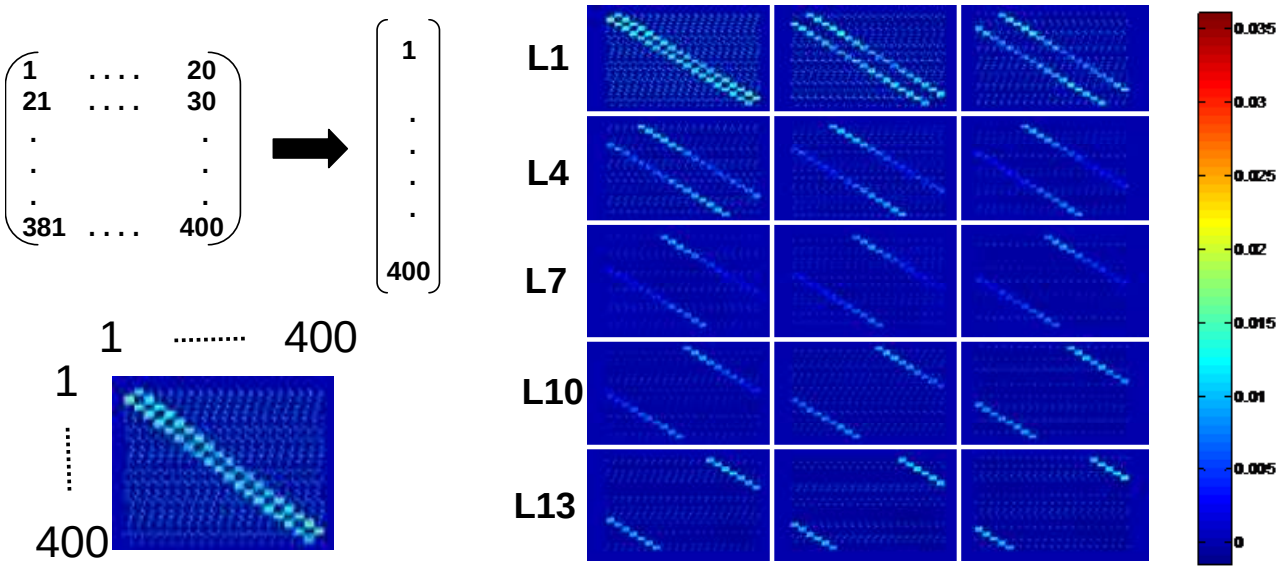


**Fig. S4. An SSM network for moving bars across a model visual field.** Top left: sensory input into a two-dimensional field (20x20) generates activities in 400 neurons arranged as shown. Right: synaptic weight matrices (for latencies 1-15). Each of these is a 400 by 400 matrix of synaptic weights (bottom left for latency=1). Positive values correspond to activating synapses, and negative values to inhibitory synapses. SSM parameter choices were as follows: potentiation strength ($\alpha$): $5x10^{-5}$, spike-rate memory ($m_s$): 50, potentiation memory ($m_p$): 50, state switching period: (Gaussian, $\tau_\mu=7,\tau_\sigma=2$), neuron firing threshold ($V_t$): 0.5, sigmoid sharpness (S): 10, latency range (L): 15. Each neuron has 11970 synaptic inputs (399 neurons x 15 latencies x 2 polarities).

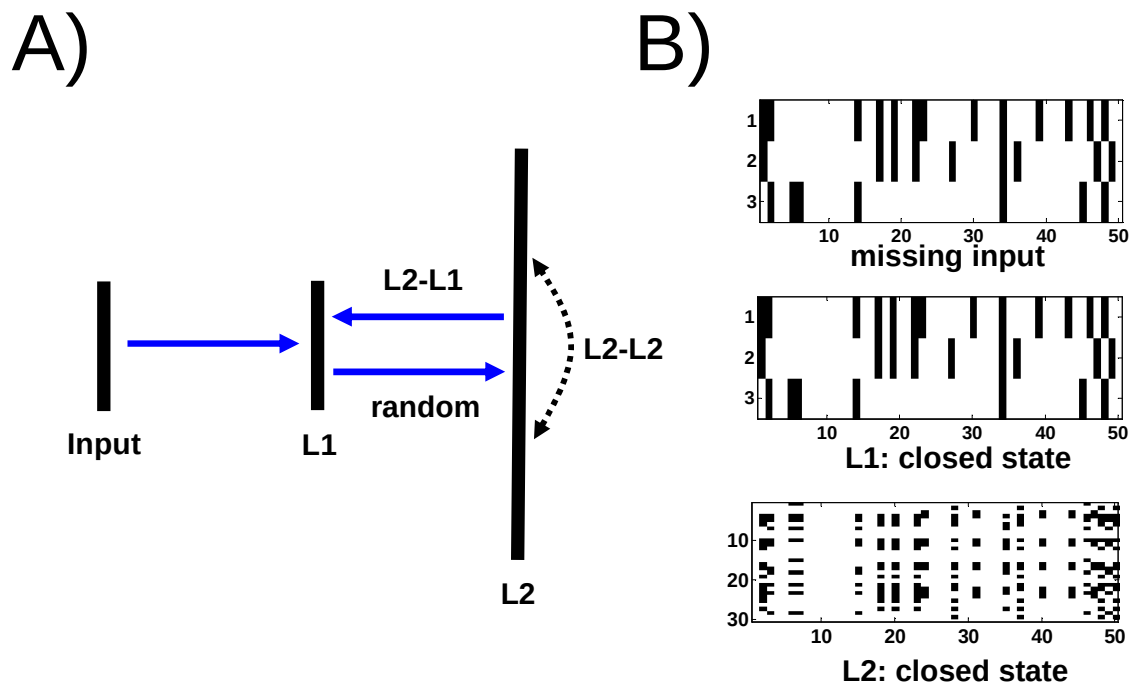


**Fig. S5. Hidden layer learns a generative model of a long sequence of spike trains.** (***A***) The three neurons in layer 1 (L1) are clamped to input spike trains (in this example, a random pattern with spike frequency of 0.2). During the open state, these neurons project (through synapses with random weights) to a much larger hidden layer (L2) which in this example contains thirty neurons. Synaptic state matching in the intra-layer 2 synaptic connections (L2-L2) gives rise to an accurate predictive model within this hidden layer. Synaptic state matching of synapse projections from layer 2 to layer 1 (L2-L1) generates an accurate predictive model of the missing input at layer 1. Layer 1 neurons do not have any internal recurrent connections. (***B***) top: missing input into layer 1 during the closed state; middle: predicted pattern of missing input created by the projection of hidden layer axons to layer 1 during the closed state. The pattern is a perfect match to the missing input; bottom: the coincident activity of the 30 neurons in the hidden layer during the closed state. SSM parameter choices for both L2-L2 and L2-L1 connections were as follows: potentiation strength ($\alpha$): $5\text{x}10^{-4}$, spike-rate memory ($m_s$): 200, potentiation memory ($m_p$): 50, state switching period: (Gaussian, $\tau_\mu$=30,$\tau_\sigma$=10), neuron firing threshold ($V_t$): 0.5, sigmoid sharpness (S): 10, latency range (L): 20.

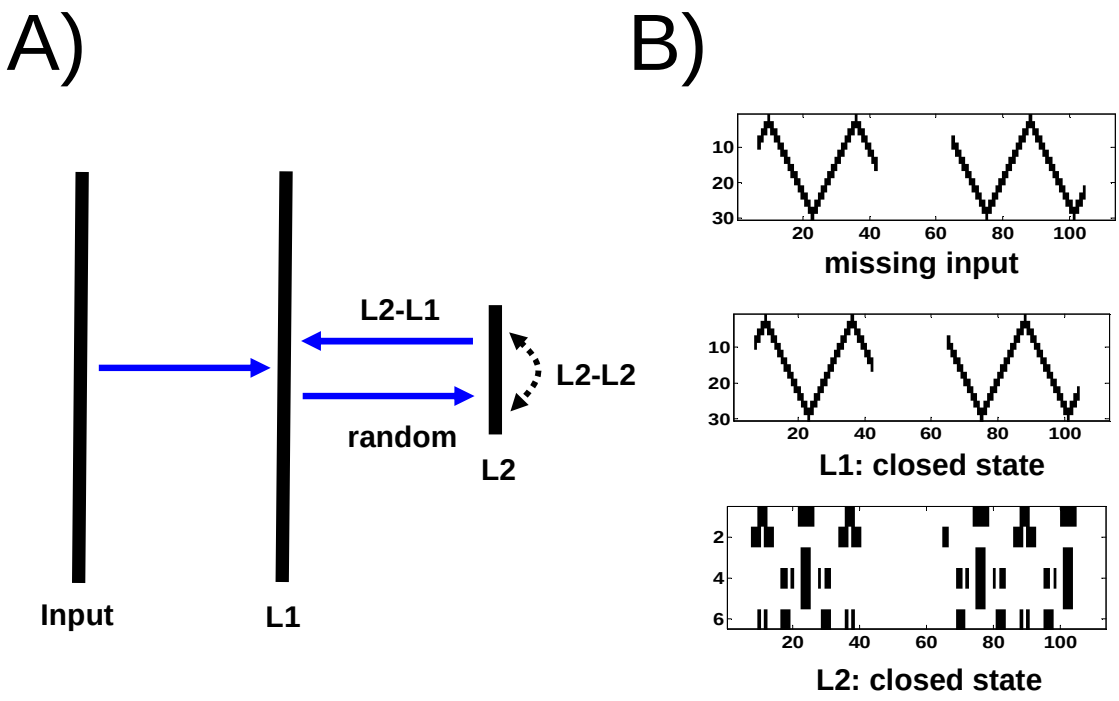


**Fig. S6. Hidden layer learns a low-dimensional representation of spike trains in the form of a repeating triangular wave.** (***A***) The thirty neurons in layer 1 (L1) are clamped to input strains in the form a repeating triangular wave. During the open state, these neurons project (through synapses with random weights) to a much smaller hidden layer (L2) which in this example contains six neurons. Synaptic state matching in the intra-layer 2 synaptic connections (L2-L2) gives rise to an accurate predictive model within this hidden layer. Synaptic state matching of synapse projections from layer 2 to layer 1 (L2-L1) generates an accurate predictive model of the missing input at layer 1. Layer 1 neurons do not have any internal recurrent connections. (***B***) top: missing input into layer 1 during the closed state; middle: predicted pattern of missing input created by the projection of hidden layer axons to layer 1 during the closed state. The pattern is a perfect match to the missing input; bottom: the coincident activity of the six neurons in the hidden layer during the closed state. SSM parameter choices for both L2-L2 and L2-L1 connections were as follows: potentiation strength ($\alpha$): $5x10^{-4}$, spike-rate memory ($m_s$): 200, potentiation memory ($m_p$): 50, state switching period: (Gaussian, $\tau_\mu$=30,$\tau_\sigma$=10), neuron firing threshold ($V_t$): 0.5, sigmoid sharpness (S): 10, latency range (L): 20.