Shaping Attractor Landscapes in Boolean Liquid State Machines via STDP and Global Plasticity

Jérémie Cabessa[1] and Alessandro E.P. Villa[2]

[1] DAVID Lab, University of Versailles (UVSQ) – Paris-Saclay
78035 Versailles, France

[2] NeuroHeuristic Research Group, DESI – HEC, University of Lausanne (UNIL)
CH-1015 Lausanne, Switzerland

jeremie.cabessa@uvsq.fr alessandro.villa@unil.ch

Abstract. Small Boolean Liquid State Machines (B-LSMs) offer a simplified yet expressive biologically inspired model of recurrent computation, in which network attractor dynamics can be systematically analyzed. In their untrained form, B-LSMs exhibit complex, often chaotic dynamics with short-lived memory traces. This study investigates how local synaptic plasticity (STDP) and a global plasticity (GP) mechanism jointly shape the attractor landscapes of these networks. Specifically, we show that synaptic modifications can drive B-LSMs to exhibit exponentially many attractors, each corresponding to a potential memory. Such high attractor regimes are attainable through global synaptic crafting. Under noisy background conditions, STDP tends to drive the networks back to low attractor regimes; however, when receiving carefully designed inputs, STDP maintains the networks' rich attractor dynamics. Overall, our findings highlight the theoretical potential for storing an impressive number of memories in recurrent neural networks, with significant implications for theoretical neuroscience and neuromorphic computing.

Keywords: Reservoir computing · Liquid state machines · Boolean networks · Synaptic plasticity · STDP · Global plasticity · Attractors · finite state automata.

1 Introduction

Attractor dynamics constitute a core framework in computational neuroscience and machine learning. In neuroscience, attractors refer to stable firing patterns in the cerebral cortex and have been critically linked to long-term and short-term memory, attention, decision-making, and mental disorders [18]. The mechanisms underlying attractor formation, their roles in representation and memory, and the empirical evidence supporting their existence in the brain have been thoroughly studied [10] (and references therein).

In machine learning, attractor states are used to model associative memory, wherein a system evolves toward stable configurations that represent stored

patterns—typically corresponding to local or global minima of an energy function [6,19]. Hopfield networks provide a classical model of fixed-point attractors, employing symmetric connectivity to ensure convergence to stable states [6]. Subsequent developments have extended this framework to accommodate asymmetric interactions [19], higher-order synaptic terms [12], and dense memory representations [3]. More recent innovations introduce continuous energy landscapes and update rules inspired by attention mechanisms, significantly enhancing both memory capacity and convergence speed [17]. Boltzmann machines—and their restricted variants—generalize this paradigm by incorporating stochastic dynamics, enabling the modeling of richer probability distributions and laying the groundwork for modern generative models [4, 5].

Reservoir Computing (RC) frameworks [13], such as Echo State Networks (ESNs) [8] and Liquid State Machines (LSMs) [14], leverage fixed recurrent network topologies with rich, dynamic internal states to project input streams into high-dimensional temporal representations. In particular, LSMs were introduced as a biologically inspired model of computation based on spiking neural networks with fading memory, offering a principled alternative to traditional rate-based neural networks [14]. Here, we consider Boolean Liquid State Machines (B-LSMs), a biologically inspired instantiation in which units operate in binary states, enabling a systematic analysis of the network's attractor dynamics [1].

The mechanisms and computational implications of spike-timing-dependent plasticity (STDP) have been extensively studied [2]. Notably, the emergence of cell assemblies – underlying attractor dynamics – can be promoted by stimulus-driven synaptic pruning in conjunction with STDP in large neural networks operating amid background noise [7]. Furthermore, the interaction between excitatory and inhibitory couplings has been shown to give rise to new attractor states, characterized by stable, synchronous activity patterns [21]. The combined action of STDP and homeostatic plasticity mechanisms also enhances both the storage and maintenance of time-varying attractor dynamics [20].

In this work, we investigate the attractor landscapes of small Boolean Liquid State Machines (B-LSMs) subject to the combined effects of spike-timing-dependent plasticity (STDP) and global synaptic plasticity. This simplified yet expressive framework enables a systematic analysis of the networks' attractor dynamics. We demonstrate that global plasticity can drive the network into rich attractor regimes, which can be sustained over time through carefully designed input patterns. Our findings highlight the theoretical potential for storing an exponential number of memories in recurrent neural networks, with implications for both theoretical neuroscience and the development of neuromorphic hardware.

2 Model

2.1 Boolean Liquid State Machines (B-LSMs)

Boolean Liquid State Machines (B-LSMs) are recurrent neural networks composed of binary-valued inputs and reservoir units [15]. Formally, a B-LSM is

defined as a tuple

$$\mathcal{N} = \left(U, X, \left\{\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})\right\}_{t\geq 0}, \left\{\mathbf{W}^{[\mathbf{in}]}(\mathbf{t})\right\}_{t\geq 0}, \left\{\mathbf{b}(\mathbf{t})\right\}_{t\geq 0}\right)$$

where $U = \{1, \ldots, M\}$ denotes the indices of the M input units, and $X = \{1, \ldots, N\}$ the indices of the N reservoir units. At any time t, the input-to-reservoir weight matrix is $\mathbf{W}^{[\mathbf{in}]}(\mathbf{t}) \in \mathbb{R}^{N\times M}$, the recurrent reservoir weight matrix is $\mathbf{W}^{[\mathbf{res}]}(\mathbf{t}) \in \mathbb{R}^{N\times N}$, and $\mathbf{b}(\mathbf{t}) \in \mathbb{R}^N$ is the bias vector of the reservoir cells. The network $\mathcal{N}$ is said to be *static* if the weights and biases $\mathbf{W}^{[\mathbf{in}]}(\mathbf{t})$, $\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})$ and $\mathbf{b}(\mathbf{t})$ remain constant over time t; it is called *evolving* otherwise. A B-LSM is illustrated in Figure 1 (left). The dynamics of a B-LSM is governed by the following equation

$$\mathbf{x}(\mathbf{t}+\mathbf{1}) = \theta\left(\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})\cdot\mathbf{x}(\mathbf{t}) + \mathbf{W}^{[\mathbf{in}]}(\mathbf{t})\cdot\mathbf{u}(\mathbf{t}) + \mathbf{b}(\mathbf{t})\right) \tag{1}$$

where $\mathbf{u}(\mathbf{t}) \in \mathbb{B}^M$ is the binary input vector, $\mathbf{x}(\mathbf{t}) \in \mathbb{B}^N$ and $\mathbf{x}(\mathbf{t}+\mathbf{1}) \in \mathbb{B}^N$ the reservoir state at time t and $t+1$, respectively, and θ the element-wise hard threshold function:

$$\theta(x) = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{otherwise} \end{cases}$$

Given an infinite input stream

$$\bar{\mathbf{u}} = \left(\mathbf{u}(\mathbf{0}), \mathbf{u}(\mathbf{1}), \mathbf{u}(\mathbf{2}), \ldots\right) \in \left(\mathbb{B}^M\right)^\infty ,$$

the system evolves to generate a corresponding state trajectory

$$\mathcal{N}(\bar{\mathbf{u}}) = \left(\mathbf{x}(\mathbf{0}), \mathbf{x}(\mathbf{1}), \mathbf{x}(\mathbf{2}), \ldots\right) \in \left(\mathbb{B}^N\right)^\infty ,$$

where each state $\mathbf{x}(\mathbf{t})$ is determined by Equation (1), for all $t > 0$. The sequence $\mathcal{N}(\bar{\mathbf{u}})$ is called the *dynamics* of $\mathcal{N}$ for the input stream $\bar{\mathbf{u}}$.,

2.2 Attractors

Attractors in a B-LSM are state configurations toward which the network dynamics evolve and stabilize into fixed-point, cyclic, or chaotic behavior. Formally, a finite set of states $A = \{\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k\} \subseteq \mathbb{B}^N$ is said to be an *attractor* of a B-LSM if there exist some input stream $\bar{\mathbf{u}} = (\mathbf{u}(0), \mathbf{u}(1), \mathbf{u}(2), \ldots)$ and time step t_0 such that the induced network dynamics $\mathcal{N}(\bar{\mathbf{u}}) = (\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \ldots)$ satisfies $\mathbf{x}(t) \in A$ for all $t \geq t_0$.

Boolean recurrent neural networks, and consequently B-LSMs, are computationally equivalent to finite-state automata (FSA) [11, 15, 16]. More precisely, for any B-LSM $\mathcal{N}$ of size N (i.e., with N reservoir cells), one can construct a corresponding FSA $\mathcal{A}$ of size $O(2^N)$. The nodes of $\mathcal{A}$ are the Boolean states of $\mathcal{N}$, and there is a transition from node $\mathbf{x}$ to node $\mathbf{x}'$ labeled by input $\mathbf{u}$ in $\mathcal{A}$ if and only if $\mathcal{N}$ transitions from state $\mathbf{x}$ to $\mathbf{x}'$ upon receiving input $\mathbf{u}$. According to this construction, the *dynamics* of $\mathcal{N}$ correspond to *paths* in the automaton

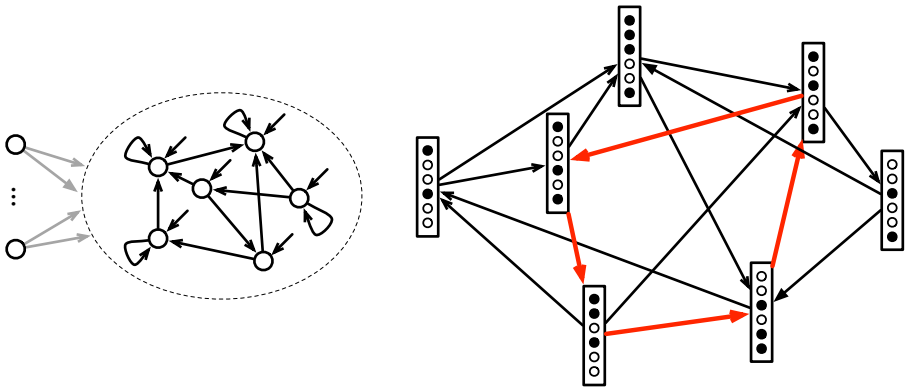

Fig. 1: A Boolean liquid state machine (left) and its corresponding automaton (right). The different dynamics of the network correspond to the different paths in the automaton. The attractors of the network correspond to the cycles in the automaton.

$\mathcal{A}$, and, in particular, the *attractors* of the network $\mathcal{N}$ correspond to *(simple) cycles* in $\mathcal{A}$ [1]. The reverse translation from a given automaton $\mathcal{A}$ to its corresponding network $\mathcal{N}$ is of no interest for our purposes [16]. This correspondance from B-LSM to FSA is illustrated in Figure 1 (right).

Based on these considerations, the number of attractors of a B-LSM corresponds to the number of elementary circuits in its corresponding FSA:

Fact 1 *The number of attractors of a B-LSM with N reservoir cells is in $\Theta\big(2^N!\big)$.*

The attractors of $\mathcal{N}$ can therefore be systematically enumerated using the following procedure:

1. Construct the FSA $\mathcal{A}$ associated to $\mathcal{N}$;
2. Enumerate all simple cycles (or elementary circuits) of $\mathcal{A}$ using Johnson's algorithm [9].

Note that the number of attractors in a B-LSM fluctuates in response to synaptic modifications. As the network adjusts its synaptic configuration, the structure of the associated automaton changes accordingly, revealing a new set of cycles that define the updated attractor regime of the network. Overall, the combinatorial explosion in the number of attractors (Fact 1) underscores the substantial memory capacity of B-LSMs and motivates a tractable analysis framework focused on small networks ($N \leq 8$).

2.3 Synaptic Plasticity

We consider two forms of synaptic plasticity that govern learning in the B-LSM framework:

(i) Spike-Timing-Dependent Plasticity (STDP): While receiving random background inputs, the network undergoes synaptic modifications governed by an STDP rule, which updates the reservoir weights $w_{ij}^{\text{res}}(t)$ at each time step based on local pre- and post-synaptic activity, as follows:

$$\begin{aligned} \Delta w_{ij}^{\text{res}}(t+1) &= \eta_{ij} \cdot (x_j(t) \cdot x_i(t+1) - x_j(t+1) \cdot x_i(t)) \\ w_{ij}^{\text{res}}(t+1) &= \texttt{clip}\left[w_{ij}^{\text{res}}(t) + \Delta w_{ij}^{\text{res}}(t+1)\right] . \end{aligned}$$

Here, $\eta_{ij} = \eta \cdot \epsilon$ is a noisy learning rate, with $\epsilon \sim \mathcal{U}([0.95, 1.05])$ and $\eta \in \mathbb{R}$, and weights are bounded via the `clip` function to the interval $[w_-^{\text{res}}, w_+^{\text{res}}]$, with $w_-^{\text{res}}, w_+^{\text{res}} \in \mathbb{R}$.

(i) global plasticity (GP): Upon receiving a predefined trigger input pattern, the entire reservoir undergoes a *global plasticity (GP)* rule modeled as a simulated annealing process that adjusts the reservoir synapses for the duration of the pattern as follows:

$$\begin{cases} \mathbf{W}^{[\mathbf{res}]}(\mathbf{t}+\mathbf{1}) = \mathbf{W}^{[\mathbf{res}]}(\mathbf{t}) + \boldsymbol{\epsilon}\,, \text{ with } \boldsymbol{\epsilon} \sim \mathcal{U}\left([-\epsilon, +\epsilon]^{M \times N}\right) \\ \Delta E(t+1) = \text{attr}\left(\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})\right) - \text{attr}\left(\mathbf{W}^{[\mathbf{res}]}(\mathbf{t}+\mathbf{1})\right) \\ P_{\text{accept}}(t+1) = \exp\left(-\frac{\Delta E(t+1)}{T_t}\right) \\ T_{t+1} = T_t \cdot \alpha\,, \text{ with } T_0 \in \mathbb{R} \text{ and } 0 < \alpha < 1 \end{cases}$$

where $\mathbf{W}^{[\mathbf{res}]}(\mathbf{t}+\mathbf{1})$ is a noisy update of $\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})$, which is accepted if it results in more attractors than $\mathbf{W}^{[\mathbf{res}]}(\mathbf{t})$ (computed using the attr function) or with a probability $P_{\text{accept}}(t+1)$, which depends on the decreasing temperature T_{t+1}. The objective is to explore synaptic configurations in order to maximize the number of attractors, balancing stochastic exploration (through noise addition) with convergence (through a cooling mechanism).

3 Results

3.1 Experimental setup

We consider Boolean Liquid State Machines (B-LSMs) with one input cell and $N \in \{5, 6, 7, 8\}$ reservoir cells. The reservoir weights are drawn from a normal distribution, i.e., $\mathbf{W}^{[\mathbf{res}]} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$, and the input weights and biases are set to $\mathbf{W}^{[\mathbf{in}]} = \mathbf{1}$ and $\mathbf{b} = \mathbf{0}$, respectively.

These networks are submitted to background input streams and trigger patterns. For the background activity, we generate a random binary input stream $\bar{u} = \big(u(0), u(1), u(2), \dots\big)$ of length 1000, where each inter-spike interval satisfies $\text{ISI}_i \sim \text{Pois}(\lambda = 2)$. For the trigger patterns, we generate a random binary sequence $\bar{p} = \big(p(0), p(1), p(2), \dots\big) \sim \text{Binomial}(n = 50, p = 0.5)$, and then insert it at P non-overlapping random positions inside the input stream $\bar{u}$, where $P \in \{0, 1, 3, 5, 7, 9, 11\}$ (blue regions in Figure 3).

A simulation involves the random generation of a Boolean network $\mathcal{N}$ (reservoir weights) and a background binary input stream $\bar{u}$ interspersed with P trigger patterns. The network processes this input and undergoes STDP or GP, depending on whether it receives background activity or a trigger pattern (white and blue regions in Figure 3). At each time step, these plasticity mechanisms modify the network's synaptic weights, thereby modifying its corresponding automaton, and thus altering its number of attractors, referred to as the *attractor regime* (cf. Section 2.2). The successive attractor regimes are computed at every time step, and their evolution over time is analyzed (see Figure 3). Each simulation was repeated 10 times with different random seeds.

3.2 Attractor regimes

Figure 2 presents the distribution of the number of attractors observed in B-LSMs composed of N reservoir nodes ($N \in \{5, 6, 7, 8\}$) and driven by 11 trigger patterns (a configuration giving rise to the largest number of attractors). As the network size increases from 5 to 8 internal nodes, the maximum number of attractors rises sharply from 32 to 23042. The average number of attractors also grows with the network size, indicating that larger networks can settle into higher attractor regimes. These results underscore the capacity of B-LSMs to support a large number of attractors, each corresponding to a potential memory state, as formally stated in Fact 1.

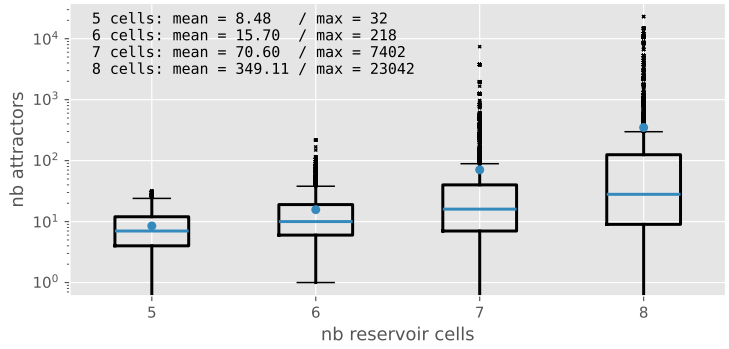

Fig. 2: Distribution of the number of attractors achieved by B-LSMs composed of $N \in \{5, 6, 7, 8\}$ reservoir nodes and driven by input streams of length 1000 containing 11 trigger patterns. Each box plot summarizes results from 10 simulations, with the median indicated by a blue line and the mean by a blue circle. The y-axis is shown on a logarithmic scale.

3.3 Effect of global plasticity (GP)

Every time the network encounters a trigger pattern, it undergoes the global plasticity (GP) process for the duration of the pattern (see Section 2.3). GP is a network-level mechanism designed to amplify the number of attractors, implemented as a simulated annealing process. The network adjusts all its reservoir synaptic weights and, with high probability, transitions into successive configurations that increase the number of attractors.

As expected, the GP mechanism significantly enhances the attractor regime of the network. This effect is clearly illustrated in Figure 3, where the trigger periods (indicated by blue bands) correspond to pronounced increases in the attractor regime. These substantial rises in attractor numbers are generally accompanied by large synaptic changes, quantified as the sum of absolute differences between synaptic weights at successive time steps ($t-1$ and t). The GP mechanism is essential for achieving large attractor regimes, as disabling this process leads to stagnation within very low regimes. In this context, the GP mechanism can be viewed as a form of *learning*.

To quantify this phenomenon on a larger scale, we analyzed all simulations involving B-LSMs subjected to 11 trigger patterns. For each pattern, we computed the difference between the number of attractors two time steps prior to the pattern's onset and at the pattern's termination. The average of these differences, referred to as the *rise in the attractor regime*, are reported in Table 1. We see that the expansion of the attractor regime induced by trigger patterns grows significantly with network size, indicating greater responsiveness and dynamical complexity in larger reservoirs.

Nb nodes	**5 nodes**	**6 nodes**	**7 nodes**	**8 nodes**
drop / rise	-5.5 / 5.4	-15.6 / 16.9	-142.0 / 141.6	-1015.8 / 1012.1

Table 1: Rises and drops in attractor regimes induced by the onsets and terminations of trigger patterns, respectively. Each pair in the form x/y represents the average decrease ("drop") and increase ("rise") in attractor count associated with the terminations and onsets of triggers pattern, respectively.

3.4 Effect of spike-timing-dependent plasticity (STDP)

Between trigger patterns, the network undergoes spike-timing-dependent plasticity (STDP) driven by random background inputs (see Section 2.3). STDP is a local mechanism that adjusts synaptic strengths by reinforcing or weakening connections based on the timing of pre- and post-synaptic spikes. This mechanism is commonly associated with learning, as it refines neural circuits to encode salient input patterns.

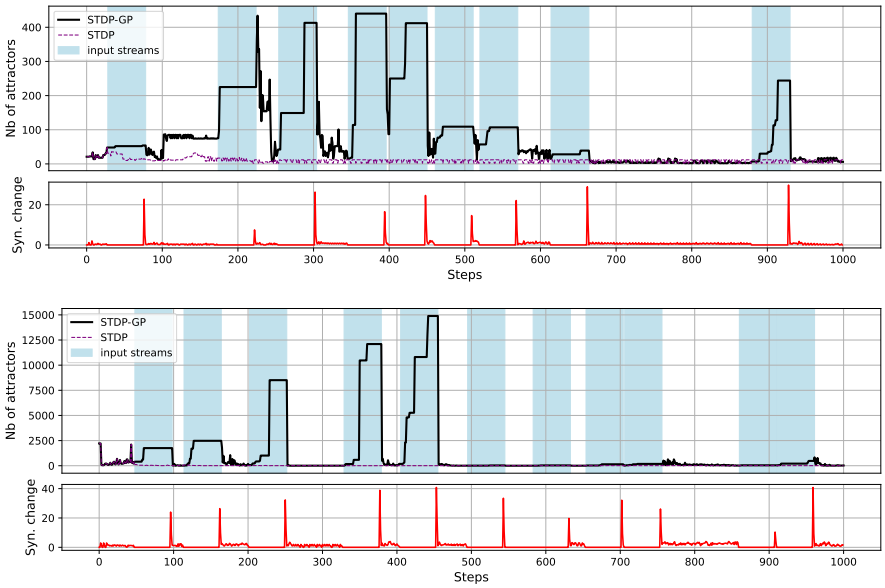

Fig. 3: Simulations of Boolean networks subject to random background activity and trigger patterns (indicated by blue bands). Top: a B-LSMs with 7 reservoir nodes. Bottom: a B-LSM with 8 reservoir nodes. The networks undergo STDP during background activity and GP during trigger patterns. At each time step, the number of attractors is computed, and the evolution of the attractor regime is displayed (black trace). For comparison, the attractor regime evolution without GP is shown (dotted trace). The synaptic change induced by the STDP and GP plasticity mechanisms are also shown (red trace).

The STDP mechanism drives the network abruptly into low attractor regimes. This effect is illustrated in Figure 3, where the terminations of trigger periods – marked by the reactivation of STDP – coincide with sudden drops in the attractor regime. STDP typically induces small but targeted synaptic changes that can disrupt the attractor structure of the network. We hypothesize that maintaining a high-dimensional attractor regime incurs an energetic cost that STDP alone, when driven by random background input, is unable to support. In this context, STDP can be interpreted as a form of *memory recalibration*, enabling the network to adapt and support new learning. It is important to emphasize that the behavior of STDP differs markedly when applied to structured input streams, as opposed to random ones, as detailed in Section 3.5.

To better understand this phenomenon, we analyzed simulations involving networks with 11 trigger patterns. For each pattern, we measured the difference between the number of attractors attained during the pattern and two time steps after its termination. The averages of these differences, referred to as the *drop in the attractor regime*, is reported in Table 1. The results reveal that the magnitude of the drops mirrors that of the rises, both showing a significant correlation with network size.

3.5 Stability of attractor regimes

Consider a network operating in a high attractor regime. If its synaptic weights were frozen, the network would necessarily preserve this regime over time, as its associated automaton (and the cycles it contains) would remain unchanged (see Section 2.2). However, in biological neural networks, synaptic weights are not abruptly frozen after learning; rather, they are continuously shaped by ongoing network activity. Furthermore, STDP induced by background activity tends to significantly reduce the attractor regime, as shown in Section 3.4. This raises a natural question: Can neural networks maintain high attractor regimes over time despite the continuous influence of activity-dependent plasticity?

We demonstrate that a network exposed to specific input streams – rather than random background activity – can effectively maintain high attractor regimes over extended periods, even in the presence of STDP. More precisely, we consider a synaptic configuration of a B-LSM associated with a regime R of 878 attractors. We then compute a specific input stream that drives the network through attractors of R, as long as no STDP is enabled. Afterward, we re-enable the STDP mechanism and track the evolution of the attractor regime in response to this specific input stream.

The results of this simulation, presented in Figure 4, demonstrate that the network is capable of maintaining its attractor regime over time, depending on the learning rate η. For a very small learning rate ($\eta = 0.0001$), the network retains its high attractor regime indefinitely. In this case, the synaptic changes induced by this minimal learning rate are too small to alter the dynamical automaton associated with the network, thereby preserving its attractor regime. At an intermediate learning rate ($\eta = 0.001$), the network initially sustains its high attractor regime but gradually transitions to lower regimes, displaying plateauing effects. In contrast, for larger learning rates ($\eta = 0.01$ and $\eta = 0.1$), the network rapidly collapses into low attractor regimes.

These observations suggest that the magnitude of synaptic changes directly influences the stability of the dynamical automaton underlying the network. Small synaptic changes are associated with stability and preservation of the attractor regime, while large changes lead to its disruption. Overall, this mechanism may be interpreted as a form of *progressive memory loss* in the absence of memory reactivation.

4 Conclusion

This study investigates the attractor dynamics of small Boolean Liquid State Machines (B-LSMs) subjected to both local and global synaptic plasticity mechanisms: spike-timing-dependent plasticity (STDP) and global plasticity (GP). This bio-inspired model offers the key advantage of enabling precise enumeration of attractors. Our theoretical and experimental results demonstrate that B-LSMs can exhibit remarkably rich attractor landscapes.

Our findings demonstrate that STDP alone is insufficient to drive a substantial expansion of a network's attractor regime. In most cases, local synaptic

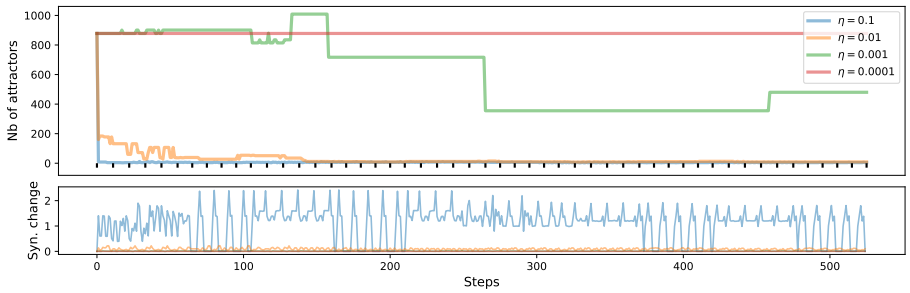

Fig. 4: Evolution of the attractor regime in the a B-LSM composed of 7 nodes, under a specific input stream and different STDP learning rates. (Top) The network begins in a high attractor regime of 878 attractors and receives a crafted inputs (black ticks) that would confine it to this regime in the absence of STDP. The evolution of the attractor regime is shown for three STDP learning rates ($\eta = 0.0001$, $\eta = 0.001$, $\eta = 0.01$). (Bottom) Magnitude of synaptic changes induced by STDP over time.

adjustments induced by unstructured inputs result in a collapse of the attractor dynamics. In contrast, global synaptic modifications enable a marked expansion of the attractor landscape of the networks. The association between global synaptic changes and specific trigger patterns can be interpreted as a form of learning or memory retrieval, whereby additional attractor states are (re)incorporated into the network's dynamics. Our results also indicate that networks can sustain high-dimensional attractor regimes, provided their inputs consistently guide them into specific attractor patterns rather than subjecting them to random, unstructured input.

This study is limited to small networks, as the sheer number of attractors in larger systems renders direct computation impractical. A promising direction for future research involves identifying computable networks' metrics capable of predicting their attractor regime, thereby enabling the systematic study of attractor dynamics in larger networks.

References

1. Cabessa, J., Villa, A.E.P.: An attractor-based complexity measurement for boolean recurrent neural networks. PLoS ONE **9**(4), e94204+ (2014)
2. Caporale, N., Dan, Y.: Spike timing-dependent plasticity: a Hebbian learning rule. Annu. Rev. Neurosci. **31**, 25–46 (2008)
3. Demircigil, M., Heusel, J., Löwe, M., Upgang, S., Vermet, F.: On a model of associative memory with huge storage capacity. Journal of Statistical Physics **168**(2), 288–299 (Jul 2017)
4. Fischer, A., Igel, C.: Training restricted boltzmann machines: An introduction. Pattern Recognition **47**(1), 25–39 (2014). `https://doi.org/https://doi.org/10.1016/j.patcog.2013.05.025`, `https://www.sciencedirect.com/science/article/pii/S0031320313002495`
5. Hinton, G.E., Sejnowski, T.J.: Optimal perceptual inference. Proceedings of the IEEE conference on Computer Vision (1983)

6. Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. Proceedings of the national academy of sciences **79**(8), 2554–2558 (1982)
7. Iglesias, J., Eriksson, J., Pardo, B., Tomassini, M., Villa, A.: Stimulus-driven unsupervised synaptic pruning in large neural networks. Lect Notes Comput Sci **3704**, 59–68 (2005). https://doi.org/10.1007/11565123_6
8. Jaeger, H.: The 'echo state' approach to analysing and training recurrent neural networks. Tech. rep., German National Research Center for Information Technology GMD Technical Report (2001)
9. Johnson, D.B.: Finding all the elementary circuits of a directed graph. SIAM J. Comput. **4**(1), 77–84 (1975). https://doi.org/10.1137/0204007, https://doi.org/10.1137/0204007
10. Khona, M., Fiete, I.R.: Attractor and integrator networks in the brain. Nature Reviews Neuroscience **23**, 744–766 (2022). https://doi.org/10.1038/s41583-022-00642-0, https://doi.org/10.1038/s41583-022-00642-0
11. Kleene, S.C.: Representation of events in nerve nets and finite automata. In: Shannon, C., McCarthy, J. (eds.) Automata Studies, pp. 3–41. Princeton University Press, Princeton, NJ (1956)
12. Krotov, D., Sompolinsky, H.: Dense associative memory for pattern recognition. In: Advances in Neural Information Processing Systems (NeurIPS 2016). pp. 1172–1180 (2016)
13. Lukosevicius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. Comput. Sci. Rev. **3**(3), 127–149 (2009). https://doi.org/10.1016/J.COSREV.2009.03.005, https://doi.org/10.1016/j.cosrev.2009.03.005
14. Maass, W., Natschläger, T., Markram, H.: Real-time computing without stable states: A new framework for neural computation based on perturbations. Neural computation **14**(11), 2531–2560 (2002)
15. McCulloch, W.S., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysic **5**, 115–133 (1943)
16. Minsky, M.L.: Computation: finite and infinite machines. Prentice-Hall, Inc., Englewood Cliffs, N. J. (1967)
17. Ramsauer, H., Schäfl, B., Lehner, J., Seidl, P., Widrich, M., Adler, T., Gruber, L., Holzleitner, M., Pavlović, M., Sandve, G.K., Greiff, V., Kreil, D., Kopp, M., Klambauer, G., Brandstetter, J., Hochreiter, S.: Hopfield networks is all you need. arXiv preprint arXiv:2008.02217 (2020), https://arxiv.org/abs/2008.02217
18. Rolls, E.T.: Attractor networks. Wiley Interdisciplinary Reviews: Cognitive Science **1**(1), 119–134 (2010). https://doi.org/10.1002/wcs.3, https://doi.org/10.1002/wcs.3
19. Sompolinsky, H., Kanter, I.: Temporal association in asymmetric neural networks. Phys. Rev. Lett. **57**, 2861–2864 (Dec 1986). https://doi.org/10.1103/PhysRevLett.57.2861, https://link.aps.org/doi/10.1103/PhysRevLett.57.2861
20. Susman, L., Brenner, N., Barak, O.: Stable memory with unstable synapses. Nature Communications **10**(1), 4441 (Sep 2019). https://doi.org/10.1038/s41467-019-12306-2, https://doi.org/10.1038/s41467-019-12306-2
21. Yazdanbakhsh, A., Babadi, B., Rouhani, S., Arabzadeh, E., Abbassian, A.: New attractor states for synchronous activity in synfire chains with excitatory and inhibitory coupling. Biol Cybern **86**(5), 367–78 (May 2002). https://doi.org/10.1007/s00422-001-0293-y