Associative memory and deep learning with Hebbian synaptic and structural plasticity

Naresh Ravichandran[1], Anders Lansner[1,2], Pawel Herman[1,3]

1. CST, EECS, KTH Royal Institute of Technology, Stockholm,
2. Department of Mathematics, Stockholm University, Stockholm
3. Digital Futures, KTH Royal Institute of Technology, Stockholm

{nbrav,ala,paherman}@kth.se

Abstract

The brain achieves complex information processing and cognitive functions leveraging synaptic learning mechanisms that are local, asynchronous, online and Hebbian in nature. Our work here investigates a neural network model with localized Hebbian plasticity that can perform associative memory and multilayer representation learning. This functionality is achieved with a brain-like modular hybrid architecture combining feedforward and recurrent processing pathways. We evaluate the model on the MNIST and F-MNIST datasets and propose that several aspects of the model are attractive for machine learning and brain-like neuromorphic hardware design.

1. Introduction

Learning in the brain is largely due to the aggregate result of changes in the efficacies of trillions of synaptic connections spread across the brain. This change in synaptic efficacies is governed by asynchronous activity- and experience-dependent mechanisms that are spatiotemporally local to the synapse and, at least in part, due to Hebbian plasticity, i.e., based on the coactivation of pre- and post-synaptic neuronal activities. Hence, large-scale neural networks with localized learning might hold the key to developing powerful brain-like algorithms performing complex information processing and cognitive functions paralleling the brain. Chief among them is the ability to build internal representations of sensory data and use these representations for computations such as memory processing and decision making.

The feedforward pathways along the sensory cortex extracts hierarchical internal representations from sensory data in a predominantly unsupervised manner. This property has been adopted in deep neural networks with the use of backprop based stochastic gradient descent algorithms and has met enormous success on pattern recognition tasks with complex real-world data [1–4]. Since straightforward implementation of backprop is not biologically plausible, there has been growing interest in recent years in more brain-plausible localized plasticity rules for learning representations [5–13].

While feedforward connections with their long-range projections from one cortical area to the next in the hierarchy are prominent, it is the recurrent connections local to each area that is ubiquitous in the neocortex [14–16]. It is estimated that 80% of the synapses made on cortical excitatory neurons are recurrent and less than 10% are feedforward, even in early sensory cortical areas [15, 17]. It thus seems apparent that such extensive recurrent circuitry in the cortex has a significant computational role that has not yet been captured by deep feedforward neural networks. It is unclear what role such numerous recurrent connections lend in terms of cortical information processing. One prominent theory is that they perform associative memory, where assemblies of coactive neurons (called a cell assembly) act as internal mental representations of memory objects [18–23], allowing several crucial brain computations such as pattern completion, prototype extraction, figure-ground segmentation, Gestalt perception, etc. Several theoretical studies have shown that recurrently connected neuron-like binary units with symmetric connectivity can implement attractor dynamics: the network is guaranteed to converge to attractor states corresponding to local energy minima [18]. Learning memories in such networks typically follows some form of Hebbian plasticity. Early artificial neural network models, such as Boltzmann machines built on such work but have since fallen out of favor with more recent deep learning architectures [24, 25]. In computational neuroscience, however, attractor networks have continued to offer mechanistic explanations for many complex cognitive functions and provide a source of rich network dynamics [26–30].

The main contribution of this work is to demonstrate that a neural network model with modular architecture and fully online localized Hebbian form of synaptic and structural plasticity is sufficient for learning sparse distributed multi-layer representations in an unsupervised manner and use the learned representations for associative memory processing. We approached this by building the network in two connection motifs: multi-layer feedforward architecture, and shallow feedforward network coupled with recurrent associative memory. We evaluated our model on the MNIST and Fashion MNIST (F-MNIST) benchmarks.

2. Model description

2.1 Modular architecture

The network is based on the BCPNN (Bayesian Confidence Propagation Neural Network) model of neocortical information processing (Lansner & Ekeberg, 1989; Lansner & Holst, 1996; Lundqvist et al., 2006; Ravichandran et al., 2020; Sandberg et al., 2002; Tully et al., 2014). BCPNN converts probabilistic inference into localized computations in terms of neural and synaptic operations [31, 32, 36]. The architecture derives from the discrete columnar organization of the neocortex of large mammalian brains, which posits that the neocortex constitutes many identical hypercolumns modules (Fig. 1A), each of which in turn comprises many minicolumns [37, 38]. The hypercolumn module is defined by local competition [38, 39] and shared receptive field within its constituted minicolumns (implemented by local inhibitory neurons in the neocortex). The minicolumn comprises around 80-100 tightly interconnected neurons having functionally similar response properties [37, 38, 40].

Our multilayer network consists of an input layer and three hidden layers connected in a feedforward architecture (Fig. 1C). The associative memory network consists of one input layer and one hidden layer, with feedforward connections from input to hidden layer, and recurrent connections within the hidden layer (Fig. 1B). We model this hidden layer with two laminae L4 (granular layer) and L2/3 (supragranular). The L4 units receive the feedforward connections from the input layer and the L2/3 units receive recurrent connections from other L2/3 units in the hidden layer and the L4 units within the same minicolumn.

2.2 Soft-winner-takes-all activation

When source units (indexed by i) send connections to a target unit (index by j), the activity propagation rule is:

$$s_j = b_j + \sum_{i \in src} \pi_i \, w_{ij} \, c_{ij},$$

$$\pi_j = \frac{\exp(s_j)}{\sum_{k \in hypercol} \exp(s_k)}$$

where s_j is the total input received by the j-th target unit, b_j, w_{ij}, c_{ij}, are the bias, weight and, connectivity parameters respectively. The activation π_j is calculated by a softmax non-linear activation function that implements a soft-winner-takes-all competition between the minicolumns within each hypercolumn and creates sparse activity [39].

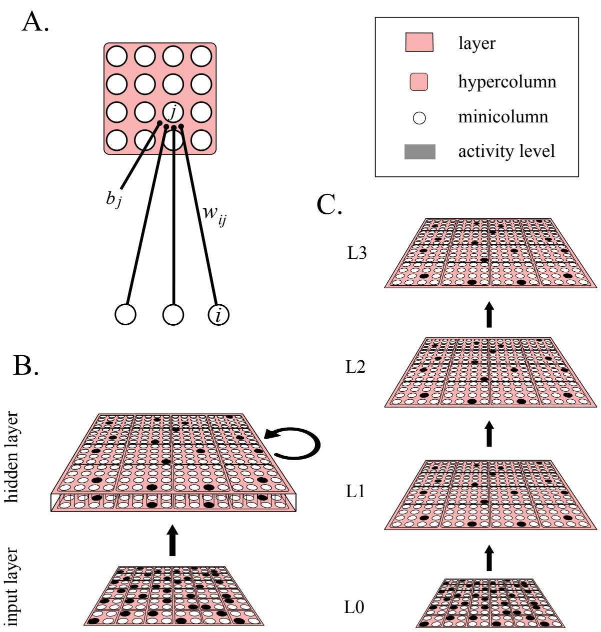

***Figure 1: A.** Schematic of one hypercolumn module showing one constituent minicolumn under local competition within the hypercolumn and receiving connections from source minicolumns. Network schematic for **B.** hybrid associative memory network, and **C.** multi-layer feedforward network with each layer constituting many hypercolumn modules.*

2.3 Localized learning

The synaptic plasticity in our model uses local traces of co-activation of pre- and post- synaptic units to compute Bayesian weights. The learning step involves incrementally updating three p-traces: probability of pre-synaptic activity, p_i, probability of post-synaptic activity, p_j, and joint probability of pre-synaptic and post-synaptic activities, p_{ij}, as follows:

$$p_i := (1-\alpha)\, p_i + \alpha\, \pi_i,$$

$$p_j := (1-\alpha)\, p_j + \alpha\, \pi_j,$$

$$p_{ij} := (1-\alpha)\, p_{ij} + \alpha\, \pi_i\, \pi_j,$$

where α is the learning rate. The bias and weight parameters are computed from the p-traces as follows:

$$b_j = \log\, p_j,$$

$$w_{ij} = \log \frac{p_{ij}}{p_i p_j}.$$

As a crucial departure from traditional backprop-based DNNs, the learning rule above is online, localized, correlative, and Hebbian. This learning rule along with the softmax activation function was shown to be implementing a discrete mixture model with the learning update fitting the parameters of the model and the activation function computing the posterior probability of the mixture components conditioned on the input [9].

2.4 Localized rewiring

Structural plasticity in the brain is the rewiring mechanism that changes the connectivity structure of the network by removing existing synaptic connections and creating new ones [41–45]. This continuous synaptic rewiring is, at least in part, an experience-dependent and activity-dependent process [42, 43]. The rewiring mechanism adaptively finds efficient sparse connectivity between the layers modeling structural plasticity in the brain. This mechanism uses the p-traces locally available at each synapse to maximize a mutual information score and updates the sparse binary connectivity matrix c_{ij} accordingly [9]. The activation, learning, and rewiring mechanisms are applicable to both recurrent and feedforward components.

2.5 Implementation and experimental setup

The input layer had 784 hypercolumns (number of pixels in MNIST and F-MNIST) with 2 minicolumns each (binary coding). All the hidden layers had 100 hypercolumns ($H = 100$) and 100 minicolumns each ($M = 100$). The connectivity was set by fixing the number of incoming hypercolumn connections per target layer (fan-in). This was set to 25 for the input to hidden layer connections and to 10 for hidden-to-hidden layer connections. All the training was through online incremental (not batch) updates and the learning rate α was set to 1e-4. For the multilayer network, learning was carried out across all the layers simultaneously through 10 epochs of training data. For the associative memory network, feedforward connections were trained first for 5 epochs, followed by 5 epochs of recurrent connections learning. The recurrent attractor was run for $T = 20$ timesteps.

We implemented a custom C++ implementation of the network with CUDA acceleration to leverage the asynchronous and localized nature of the computations in terms of model parallelism. To this end, we ran each layer's computations (activation, learning, and rewiring) on one GPU while transferring only the layer activities to the subsequent layer (MPI asynchronous data transfers), which greatly sped up our computations. We ran the simulation on NVIDIA A100 GPUs.

3. Results

3.1 Associative memory extracts prototypes and renders representations robust to distortions

To visualize the attractor representation, we set the input layer to the test samples, drove the feedforward connections to hidden layer, ran the recurrent attractor network, and back projected (top down) the hidden representations to the input layer using the same weights as feedforward connectivity. We observed that the recurrent attractors converged in a few steps, and their reconstruction corresponded to prototypical digits (not shown).

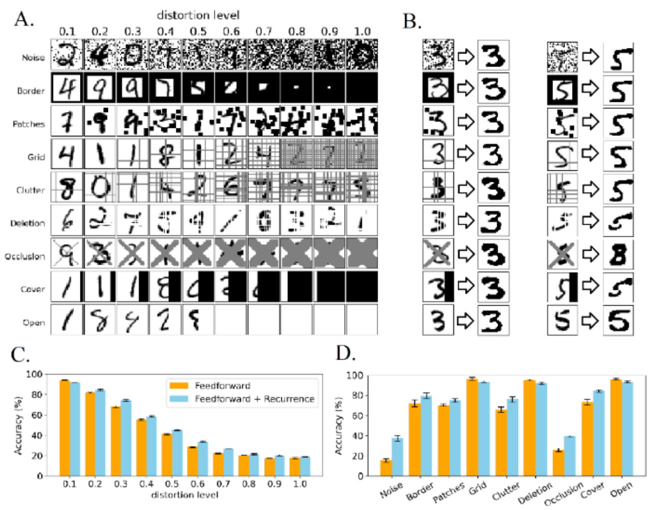

Figure 2: Robustness to distortion. A. *Samples from the distorted MNIST dataset, with 9 distortion types (rows) and varying distortion levels (columns).* ***B.*** *Examples of digits "3" (left) and "5" (right) under all distortion types (distortion level = 0.3), and the input reconstruction after convergence of the recurrent layer. For most cases, the attractor reaches the prototype pattern while removing all distortions.* ***C.*** *Classification performance for different distortion levels comparing feedforward-driven and recurrent attractor representations.* ***D.*** *Classification performance for all distortion types with distortion level at 0.3*

We examined if the recurrent attractor network added value to the feedforward-driven representations when tested on severely distorted samples. For this, we first created a distorted version of the MNIST dataset (examples shown in Fig. 2A) following the work of George et al., (2017). In particular, we introduced nine different distortions (Fig. 2A rows) and controlled the level of distortion with a distortion level parameter ranging from 0.1 (minimum distortion) to 1.0 (maximum distortion) in steps of 0.1 (Fig. 2A columns). Then we tested the network (trained beforehand on the undistorted MNIST data) with these distorted samples to comparatively study the recurrent network attractor activations and their corresponding input reconstructions. Except for high distortion levels (0.5 or higher), the attractor network was very robust to the distortions, and the reconstructed images showed that they converged to the prototypical digits and most of the distortions were removed upon attractor convergence. Fig 2B shows examples of two digits under all distortion types (distortion level of 0.3) and the reconstructions of the corresponding attractor activations. To quantify the network's robustness to input distortions in the pattern recognition context, we used a linear classifier, which was trained on feedforward-driven activations. We then compared the classification accuracy from five random trials obtained for recurrent attractor representations and the feedforward-driven hidden representations on the distorted MNIST dataset (Fig. 2C and 2D). We found that the recurrent attractor representations performed better compared to feedforward-driven representations on all distortion levels greater than 0.1 (Fig. 2C). We also examined the performance across the different distortion types and the recurrent attractor representations turned out more resilient in most cases (Fig. 2D). Further investigation is needed to understand why feedforward-driven representations scored higher on three distortion types: "Grid", "Deletion", and "Open". Both feedforward-driven and recurrent attractor representations are the result of unsupervised learning, so the distortion resistance provided by the recurrent component is achieved without any access to data labels.

3.2 Feedforward multilayer network extracts hierarchical representations

For evaluating the multilayer network, we trained the network on MNIST and F-MNIST data and used a linear classifier to classify the labels from representations from each layer in the network (Fig. 3). We saw that, in both datasets, the accuracy improved significantly from layer 0 (raw data) to layer 1, and from layer 1 to layer 2, with a small drop in performance from layer 2 to layer 3 (still higher than layer 0).

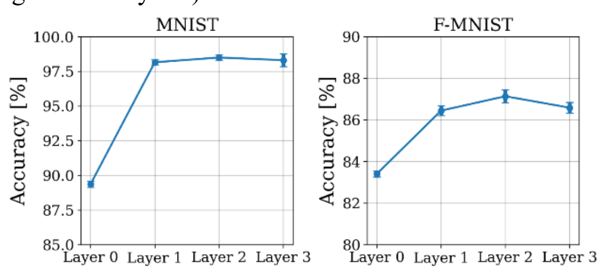

Figure 3: *Performance of a linear classifier trained on the internal representations from each layer in the multilayer network for MNIST (left) and F-MNIST (right) datasets.*

We then plotted the receptive fields learned by the structural plasticity mechanism across layers 1, 2 and, 3 (Fig. 4C, 4B and, 4A respectively) for 24 randomly chosen hypercolumns when trained on MNIST data. Given that the connections were randomly initialized, structural plasticity formed contiguous local patches in the image space across all layers, with layers higher in the hierarchy showing progressively larger receptive fields. Next, we visualized the response properties of individual minicolumns (24 randomly chosen within the same hypercolumn with shared receptive fields) by plotting the input image that evoked the most activity (Fig. 4F, 4E and, 4D respectively). We observed that layer 1 minicolumns are activated by small edges of different orientations, while layers 2 and layer 3 progressively are activated by larger regions of the image resembling whole digits.

4. Conclusions

The direction we pursued here integrated various architectural and functional details from neurobiology, especially from the mammalian neocortex: unsupervised learning, Hebbian synaptic plasticity, Hebbian structural plasticity, sparse distributed activities, sparse connectivity, columnar and laminar cortical architecture, etc., into an abstract neural network model. The hybrid architecture showed how the integration of recurrent attractors networks with feedforward networks can avoid interference between memory pattern that is commonly observed in flat attractor networks (without any hidden layers) and render the representations robust to severe distortions. Our multilayer network demonstrated a functioning prototype of hierarchical learning without requiring any label information or global backprop signal.

Our modular architecture accompanied with the localized learning and rewiring mechanisms derive from the powerful computational primitives of the canonical neocortical circuits [37, 38] that, we believe, need to be understood for building brain-like algorithms as well as for exploring the perceptual and cognitive functions of the

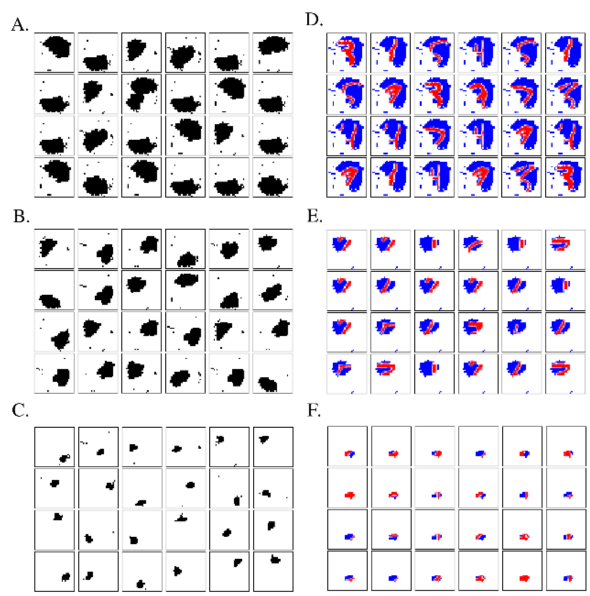

Figure 4: *Receptive fields of layers 1, 2 and, 3 (C, B and A respectively) formed by the structural plasticity mechanism shows localized contiguous patches with progressively larger sizes. The input data sample that evoked the highest activity of various minicolumns within the same hypercolumn in layer 1, 2 and, 3 (F, E and D respectively) shows diverse response profiles suitable for hierarchical feature detection.*

cortex. An additional advantage of such networks equipped with brain-like algorithms is that they allow for straightforward conversion into spiking neural networks that enable efficient event-driven communication and low-power neuromorphic hardware [47–49]

References

1. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature 2015 521:7553. 521, 436–444 (2015). https://doi.org/10.1038/nature14539.
2. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems (2012). https://doi.org/10.1145/3065386.
3. Schmidhuber, J.: Deep learning in neural networks: An overview. Neural Networks. 61, 85–117 (2015). https://doi.org/10.1016/J.NEUNET.2014.09.003.
4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. pp. 5999–6009 (2017).
5. Krotov, D., Hopfield, J.J.: Unsupervised learning by competing hidden units. Proc Natl Acad Sci U S A. 116, 7723–7731 (2019). https://doi.org/10.1073/pnas.1820458116.
6. Bartunov, S., Santoro, A., Hinton, G.E., Richards, B.A., Marris, L., Lillicrap, T.P.: Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In: Advances in Neural Information Processing Systems. pp. 9368–9378 (2018).
7. Illing, B., Gerstner, W., Brea, J.: Biologically plausible deep learning — But how far can we go with shallow networks? Neural Networks. 118, 90–101 (2019). https://doi.org/10.1016/j.neunet.2019.06.001.
8. Ravichandran, N.B., Lansner, A., Herman, P.: Learning representations in Bayesian Confidence Propagation neural networks. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp. 1–7. IEEE (2020). https://doi.org/10.1109/IJCNN48605.2020.9207061.
9. Ravichandran, N.B., Lansner, A., Herman, P.: Brain-Like Approaches to Unsupervised Learning of Hidden Representations - A Comparative Study. In: Artificial Neural Networks and Machine Learning – ICANN 2021. pp. 162–173. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86383-8_13.
10. Journé, A., Rodriguez, H.G., Guo, Q., Moraitis, T.: Hebbian Deep Learning Without Feedback. (2022). https://doi.org/10.48550/arxiv.2209.11883.
11. Sa-Couto, L., Wichert, A.: Using brain inspired principles to unsupervisedly learn good representations for visual pattern recognition. Neurocomputing. 495, 97–104 (2022). https://doi.org/10.1016/J.NEUCOM.2022.04.130.
12. Bengio, Y.: How Auto-Encoders Could Provide Credit Assignment in Deep Networks via Target Propagation. (2014).
13. Payeur, A., Guerguiev, J., Zenke, F., Richards, B.A., Naud, R.: Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. Nature Neuroscience 2021 24:7. 24, 1010–1019 (2021). https://doi.org/10.1038/s41593-021-00857-x.
14. van Bergen, R.S., Kriegeskorte, N.: Going in circles is the way forward: the role of recurrence in visual inference. Curr Opin Neurobiol. 65, 176–193 (2020). https://doi.org/10.1016/J.CONB.2020.11.009.
15. Douglas, R.J., Martin, K.A.C.: Recurrent neuronal circuits in the neocortex. Current Biology. 17, R496–R500 (2007). https://doi.org/10.1016/J.CUB.2007.04.024.
16. Lamme, V.A.F., Roelfsema, P.R.: The distinct modes of vision offered by feedforward and recurrent processing. Trends Neurosci. 23, 571–579 (2000). https://doi.org/10.1016/S0166-2236(00)01657-X.
17. Stepanyants, A., Martinez, L.M., Ferecskó, A.S., Kisvárday, Z.F.: The fractions of short- and long-range connections in the visual cortex. Proc Natl Acad Sci U S A. 106, 3555–3560 (2009). https://doi.org/10.1073/pnas.0810390106.
18. Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities (associative memory/parallel processing/categorization/content-addressable memory/fail-soft devices). Proc Natl Acad Sci U S A. 79, 2554–2558 (1982).
19. Lansner, A.: Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. Trends Neurosci. 32, 178–186 (2009). https://doi.org/10.1016/j.tins.2008.12.002.
20. Hebb, D.O.: The Organization of Behavior. Psychology Press (1949). https://doi.org/10.4324/9781410612403.
21. Harris, K.D.: Neural signatures of cell assembly organization,

https://www.nature.com/articles/nrn1669, (2005). https://doi.org/10.1038/nrn1669.

22. Willshaw, D.J., Buneman, O.P., Longuet-Higgins, H.C.: Non-holographic associative memory. Nature. 222, 960–962 (1969). https://doi.org/10.1038/222960a0.

23. Fransén, E., Lansner, A.: Low Spiking Rates in a Network with Overlapping Assemblies. The Neurobiology of Computation. 203–208 (1995). https://doi.org/10.1007/978-1-4615-2235-5_33.

24. Ackley, D.H., Hinton, G.E., Sejnowski, T.J.: A learning algorithm for Boltzmann machines. Cogn Sci. 9, 147–169 (1985). https://doi.org/10.1016/S0364-0213(85)80012-4.

25. Hinton, G.E., Osindero, S., Teh, Y.W.: A fast learning algorithm for deep belief nets. Neural Comput. 18, 1527–1554 (2006). https://doi.org/10.1162/neco.2006.18.7.1527.

26. Lundqvist, M., Herman, P., Lansner, A.: Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. J Cogn Neurosci. 23, 3008–3020 (2011). https://doi.org/10.1162/jocn_a_00029.

27. Silverstein, D.N., Lansner, A.: Is attentional blink a byproduct of neocortical attractors? Front Comput Neurosci. 5, 13 (2011). https://doi.org/10.3389/FNCOM.2011.00013/BIBTEX.

28. Fiebig, F., Lansner, A.: A spiking working memory model based on hebbian short-term potentiation. Journal of Neuroscience. 37, 83–96 (2017). https://doi.org/10.1523/JNEUROSCI.1989-16.2016.

29. Sussillo, D.: Neural circuits as computational dynamical systems, (2014). https://doi.org/10.1016/j.conb.2014.01.008.

30. Vogels, T.P., Rajan, K., Abbott, L.F.: Neural network dynamics, www.annualreviews.org, (2005). https://doi.org/10.1146/annurev.neuro.28.061604.135637.

31. Sandberg, A., Lansner, A., Petersson, K.M., Ekeberg, Ö.: A Bayesian attractor network with incremental learning. Network: Computation in Neural Systems. 13, 179–194 (2002). https://doi.org/10.1080/net.13.2.179.194.

32. Lansner, A., Ekeberg, Ö.: A one-layer feedback artificial neural network with a Bayesian learning rule. Int J Neural Syst. 01, 77–87 (1989). https://doi.org/https://doi.org/10.1142/S0129065789000499.

33. Lundqvist, M., Rehn, M., Lansner, A.: Attractor dynamics in a modular network model of the cerebral cortex. Neurocomputing. 69, 1155–1159 (2006). https://doi.org/10.1016/J.NEUCOM.2005.12.065.

34. Lansner, A., Holst, A.: A higher order Bayesian neural network with spiking units., (1996). https://doi.org/10.1142/S0129065796000816.

35. Lundqvist, M., Herman, P., Lansner, A.: Attractor Hypothesis of Associative Cortex: Insights from a Biophysically Detailed Network Model. In: Functional Brain Mapping and the Endeavor to Understand the Working Brain. IntechOpen (2013). https://doi.org/10.5772/56229.

36. Tully, P.J., Hennig, M.H., Lansner, A.: Synaptic and nonsynaptic plasticity approximating probabilistic inference. Front Synaptic Neurosci. 6, 8 (2014). https://doi.org/10.3389/FNSYN.2014.00008/ABSTRACT.

37. Mountcastle, V.B.: The columnar organization of the neocortex, (1997). https://doi.org/10.1093/brain/120.4.701.

38. Douglas, R.J., Martin, K.A.C.: Neuronal circuits of the neocortex, (2004). https://doi.org/10.1146/annurev.neuro.27.070203.144152.

39. Carandini, M., Heeger, D.J.: Normalization as a canonical neural computation. Nature Reviews Neuroscience 2011 13:1. 13, 51–62 (2011). https://doi.org/10.1038/nrn3136.

40. Buxhoeveden, D.P., Casanova, M.F.: The minicolumn hypothesis in neuroscience. Brain. 125, 935–951 (2002). https://doi.org/10.1093/BRAIN/AWF110.

41. Bailey, C.H., Kandel, E.R.: Structural changes accompanying memory formation. Ann. Rev. Physiol. 55, 397–426 (1993). https://doi.org/10.1146/annurev.ph.55.030193.002145.

42. Holtmaat, A., Svoboda, K.: Experience-dependent structural synaptic plasticity in the mammalian brain. Nature Reviews Neuroscience 2009 10:9. 10, 647–658 (2009). https://doi.org/10.1038/nrn2699.

43. Butz, M., Wörgötter, F., van Ooyen, A.: Activity-dependent structural plasticity. Brain Res Rev. 60, 287–305 (2009). https://doi.org/10.1016/J.BRAINRESREV.2008.12.023.

44. Lamprecht, R., LeDoux, J.: Structural plasticity and memory. Nature Reviews Neuroscience 2004

5:1. 5, 45–54 (2004). https://doi.org/10.1038/nrn1301.

45. Stettler, D.D., Yamahachi, H., Li, W., Denk, W., Gilbert, C.D.: Axons and Synaptic Boutons Are Highly Dynamic in Adult Visual Cortex. Neuron. 49, 877–887 (2006). https://doi.org/10.1016/J.NEURON.2006.02.018.

46. George, D., Lehrach, W., Kansky, K., Lázaro-Gredilla, M., Laan, C., Marthi, B., Lou, X., Meng, Z., Liu, Y., Wang, H., Lavin, A., Phoenix, D.S.: A generative vision model that trains with high data efficiency and breaks text-based CAPTCHAs. Science (1979). 358, (2017). https://doi.org/10.1126/SCIENCE.AAG2612.

47. Ravichandran, N., Lansner, A., Herman, P.: Spiking neural networks with Hebbian plasticity for unsupervised representation learning. (2023).

48. Xu, J., Wang, D., Li, F., Zhang, L., Stathis, D., Yang, Y., Jin, Y., Lansner, A., Hemani, A., Zou, Z., Zheng, L.R.: A Memristor Model with Concise Window Function for Spiking Brain-Inspired Computation. 2021 IEEE 3rd International Conference on Artificial Intelligence Circuits and Systems, AICAS 2021. (2021). https://doi.org/10.1109/AICAS51828.2021.9458424.

49. Podobas, A., Svedin, M., Chien, S.W.D., Peng, I.B., Ravichandran, N.B., Herman, P., Lansner, A., Markidis, S.: StreamBrain: An HPC Framework for Brain-like Neural Networks on CPUs, GPUs and FPGAs. Proceedings of the 11th International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies. 1–6 (2021). https://doi.org/10.1145/3468044.3468052.