Biologically-inspired adaptive learning in the Hopfield-network based self-optimization model

Aisha Belhadi
Embodied Cognitive Science Unit
Okinawa Institute of Science and Technology Graduate University
Okinawa, Japan
aisha.belhadi@gmail.com

Abstract

A significant portion of the recent growth of artificial intelligence can be attributed to the development of deep learning systems, going hand in hand with the accumulation of Big Data. It therefore makes sense that most often, these systems are based on supervised or reinforcement learning using massive datasets, and reward or error-based rules for training. Though these techniques have achieved impressive levels of accuracy and functionality, rivaling human cognition in some areas, they seem to work very differently from living systems that can learn, make associations and adapt with very sparse data, efficient use of energy and comparatively minimal training iterations. In the world of machine learning, Hopfield networks, with an architecture that allows for unsupervised learning, an associative memory, scaling, and modularity, offer an alternative way of looking at artificial intelligence, that has the potential to hew closer to biological forms of learning. This work distills some mechanisms of adaptation in biological systems, including metaplasticity, homeostasis, and inhibition, and proposes ways in which these features can be incorporated into Hopfield networks through adjustments to the learning rate, modularity, and activation rule. The overall aim is to develop deep learning tools that can recapitulate the advantages of biological systems, and to have a computational method that can plausibly model a wide range of living and adaptive systems of varying levels of complexity.

1 Introduction

Acting on information in an uncertain and volatile environment requires dealing with a number of trade-offs, for instance, between speed and precision [54], or optimization and generalizability. This is a common dilemma in machine learning systems, exemplified by the bias/variance trade-off [29]. Living systems seem to have developed multiple avenues to resolve these types of tensions, allowing them to learn about their changing environments and modify their learning strategies, in order to best adapt to their surroundings [54]. This adaptive learning is not strictly limited to cognitive or neural systems, and can be found at multiple timescales and levels of complexity, ranging from enzymatic cascades [39], gene regulatory networks [9], chemotaxis in prokaryotes [30], immune systems [44], and social systems [6, 21], to name a few.

To support adaptation, complex systems need to develop an ability to integrate new information while maintaining a level of stability [34]. Excessive destabilization would preclude the ability to learn consistent patterns, and to generalize skills, and could even lead to the disintegration of a system, whereas placing too much weight on past memory could maintain erroneous patterns and hinder learning when faced with a new situation. Mechanisms like the interplay between information integration and storage, between attention and memory [43], and attractor and integrator dynamics

Associative Memory & Hopfield Networks in 2023. NeurIPS 2023 workshop.

in the brain [31] underlie the robustness and flexibility required for adaptive learning. While in biological systems, these constraints are what seem to drive adaptation [22], achieving such robustness has been a challenge in the field of machine learning with deep neural networks, due to the tendency for these networks to forget previous training when learning new tasks sequentially or continuously, a phenomenon also known as catastrophic forgetting [33]. Furthermore, the resources they require in terms of energy and data to train supervised learning models may not be sustainable from an environmental and social perspective [52].

Hopfield networks [28], a type of artificial neural network with recurrent connections, have an architecture that may be adapted to solve some of these dilemma. Memory patterns are stored in the weights of connections between nodes in a Hopfield network, which is capable of one-shot and unsupervised learning. Interestingly, attractor dynamics are also at play in these networks. The coordination of the nodes' states, with respect to the constraints set by the weight matrix of all the connections in the network, defines an energy state for the system. After initiation with random states, the network's updating of the state of its nodes at every step gradually leads to convergence towards an attractor, a minimum in the energy landscape. In principle, the lower the energy state, the better the constraints are satisfied.

The strength of a Hopfield network, its encoding of memory patterns directly in the weights of connections between nodes, can also be one of its limitations in straightforward implementations of the network, as it restricts its memory storage capacity to approximately $0.14N$ patterns, with N being the number of nodes in the network [2]. Much work has since gone into extending the capabilities of Hopfield networks [35, 36, 50], though few are based on biologically-inspired mechanisms [3, 12, 22]. This is all the more surprising as the initial inspiration for the Hopfield network was from neurobiology, and the simplicity of its basic implementation makes it versatile, allowing it to be applied to model diverse biological systems, both neural and non-neural [14, 16, 42, 57, 58].

One specific implementation of Hopfield networks, the self-optimization (SO) model [59, 60] has a weight update rule based on Hebbian learning [26] and periodic resets of all network states, effectively creating a memory of visited attractors. Updating the weight matrix also transforms the energy landscape, such that the attraction basin for the global minimum is enlarged, increasing the likelihood of reaching it and finding stored memory patterns. These properties seem to combine and improve both optimization and memory stabilization, and are reminiscent of the higher stability of trajectory-coding neurons in the hippocampus [32], making the SO model an interesting candidate for solving the dilemma of adaptivity mentioned earlier. Other studies have extended the SO model [22, 62, 67] and shown its ability to solve combinatorial satisfiability problems [61]. As a Hopfield network with functionalities echoing living systems, could the SO model be further enhanced with other biological mechanisms, in order to widen its performance and applicability, and offer solutions to the dilemma of adaptation in artificial systems?

In Section 2, we explore features of adaptation and adaptive learning that are common in biological systems, with a special focus on metaplasticity and homeostatic regulation. Section 3 proposes ways in which these features could be translated into the SO model, and by extension, into Hopfield networks.

2 Adaptive mechanisms in biological systems

Hebbian learning already instantiates a form of activity-dependent neuronal adaptivity, as weight updates are meant to mimic long-term potentiation (LTP) and long-term depression (LTD) in synaptic plasticity, which modify the strength of synaptic connections ie. change how likely a post-synaptic neuron is going to activate (or be inhibited) after input from a specific synapse [10]. This plasticity itself is regulated, and its magnitude governed, by metaplasticity mechanisms [1]. Metaplasticity operates at different levels (chemical, molecular, architectural, network) and timescales (immediate, to several days). At the local synaptic level, metaplasticity is linked to changes to synaptic architecture, modifications to neurotransmitter-receptor sensitivity, ion-channel permeability, through transformation of their subunits, changes in receptor density via endocytosis, affecting the ratios of receptor types and hence, their downstream effects [1]. Plasticity effects induced by neuromodulators such as dopamine, acetylcholine, noradrenaline, and serotonin, allow for the integration of metaplasticity into functional networks [53]. Furthermore, the same neuromodulator can have opposite effects depending on the receptor types, whose distribution can vary in different brain regions. These properties scale up

metaplasticity dynamics over time and space, with differing local/global, and short/long-term effects. For example, a recent study describing the effects of dopamine on spontaneous behavior in mice found that dopamine seemed to enhance variability of behavior in the short timescale, and stabilize behavior patterns on longer timescales [40]. The synaptic-tagging-and-capture (STC) theory suggests another mechanism that spatially and temporally extends metaplasticity, through the consolidation of LTP in "tagged" synapses that are able to capture plasticity-related proteins generated by potentiation events in neighboring synapses [1]. Despite this diversity, many pathways seem to converge upon the regulation of calcium currents in neurons, directly and indirectly affecting neuronal excitability, with the net effect of dynamically supporting learning [53].

The dual role of some neuromodulators recalls the oscillatory nature of homeostasis. Indeed, metaplastic events can be thought of as synergistic, creating a positive feedback loop reinforcing the initial event. But a ubiquitous form of biological adaptation, homeostasis, seeks to dynamically regulate the system around a set point, counteracting runaway activation, depression or saturation of the system, which could threaten its viability [8]. It can also be pointed out that, though homeostasis is often described along the lines of control theory [8], due to the reliance on sensors monitoring the state of a system, and dynamic adjustments through negative feedback loops (inhibition, removal), its overall structure has echoes of attractor dynamics [15], a feature shared by Hopfield networks, as well as oscillatory dynamics [45, 56]. Homeostasis is the basis of vital physiological functions, such as the maintenance of body temperature, blood pressure, electrolyte balance, metabolic regulation. It has also been posited in synaptic regulation, as a means to restore balance in neural networks destabilized by activation and learning [13]. For instance, the Bienenstock-Cooper-Munro theory [7, 53], proposes a sliding threshold for LTP and LTD, whereby the stronger a synapse the higher the threshold to achieve LTP. Inhibitory processes are crucial to homeostasis, and are adaptive mechanisms in their own right. See Appendix 1 for a further discussion of inhibition and degeneracy in defining patterns in biological systems.

Just as complex systems have flexible ways of defining patterns, the ability to change patterns, by forgetting old ones, contributes to adaptation. This can occur through automatic processes such as the drift in spatio-temporal encoding in the hippocampus [17], protein turnover in the cell, or mRNA decay [51]. Or pattern changes can be elicited in a more specific and abrupt manner, in the form of resets. At the level of brain networks, the interplay between the neuromodulators acetylcholine and noradrenaline seems to be implicated in resetting functional brain networks in the context of surprise and uncertainty [5, 11]. Other forms of biological resets can be seen at the behavioral level in sleep [48], at the cellular level in dedifferentiation [64], and at the molecular level in chromatin reorganization during mitosis [49].

3 Implementing adaptive learning in the SO model

One way of increasing plasticity could be through scaling up the dimensions in the network, without necessarily scaling up the network itself. Rather than introducing new parameters, this could be achieved by providing the system a way to exploit the properties and behaviors of its existing components, as biological systems seem to do.

A form of local metaplasticity could be introduced, for instance, by making the learning rate α dependent on the weight of incoming inputs. Rather than a user-defined constant, α would thus be defined as a function of the weights w_{ij} between nodes *i* and *j*, transforming the Hebbian learning rule into:

$$w_{ij}\ (t+1) = w_{ij}\ (t) + dw_{ij}\ (t), \quad dw_{ij} = f(w_{ij})\ s_i\ s_j$$

where s_i and s_j are the states of nodes *i* and *j*, and $\alpha = f(w_{ij})$. Though the complexity of metaplasticity in neural systems precludes a linear approximation, simply defining f here as a linear function could still have differential effects on synapses depending on their activation state, with faster potentiation of co-activated synapses, and slower depression of synapses with opposing states, rather than increasing or decreasing weights by a fixed amount.

Linking α to changes in the energy state of the network is another way to implement a network-level metaplasticity, by computing the derivative of the energy function E with regards to time, or calculating ΔE between the current step and the previous step and defining α as a function of ΔE. This has similarities to a gradient descent, but rather than being used to minimize a loss function, or update weights, the slope of the descent in the energy landscape at a time point t (ie. the derivative

of the energy function E) serves to modify the learning rate $\alpha = f(\Delta E)$. In most cases, the function f will be chosen so as to ensure that α is low when ΔE is low, and high when the slope is higher, in order to avoid getting stuck in local minima and speed up learning when approaching a global energy minimum. But it would also be interesting to explore the behavior of the network with various definitions of $f(\Delta E)$, to see whether these variations can model different normal and pathological behavioral modes in biological systems.

Homeostasis' embeddedness in an environmental and evolutionary framework can make it challenging to implement in a meaningful way, and such work is beyond the scope of this paper (see [47]). However, rather than faithfully rendering homeostatic mechanisms, it might be interesting to mimic some of their overall dynamics, such as oscillation. One avenue would be through the network's activation function. Several types of activation functions have been trialed in the SO model [67], including sigmoid and Heaviside threshold functions. But it would be interesting to explore, for instance, the network's behavior with a trigonometric activation function, which could lead to an oscillation of activation states according to specific ranges of inputs. Though it may create node-level variability, would this contribute to global stabilization, similarly to homeostasis?

The use of a modular structure or clusters can allow for the implementation of other adaptive mechanisms. For instance, inter-cluster inhibitory connections were shown to improve energy levels in an SO model of the *C. elegans* connectome [42]. Similarly, we could imagine a structure with stronger intra-modular connections, and weaker inter-module connections, which could evolve over time, echoing functional networks in the brain, genes, enzymatic or cell-signaling pathways etc. Modules that have consistently increased their average weights over a certain number of steps could see an automatic boost to all intra-modular synaptic strengths, allowing for a form of temporal integration of a module's behavior. Or in a mechanism reminiscent of STC, if a certain number of nodes in a module increase the weights of their connections, this could trigger a weight increase for all nodes of the module. Concomitant weakening of extra-modular connections, or inhibitory inter-modular connections as described previously, could be further explored, as could the use of "silent" modules, which might introduce a form of degeneracy in the system, and may be of interest in continuous learning. Further implementations of resets and "forgetting" mechanisms, as well as examples of implementations in other network architectures can be found in Appendix 2.

4 Conclusion

Thanks to the multiplicity of features biological systems have access to, in the form of the chemical and physical properties of their components and environment, computational complexity can be scaled up in a myriad of ways by nesting, hierarchically ordering or intersecting such features, increasing the degrees of freedom within the system. Yet at the same time, these systems are highly constrained by mechanisms that maintain them within a range around a point of equilibrium. Rather than a contradiction between these two seemingly opposing dynamics, the constraints can be viewed as establishing a framework that then allows an opening up of the space of possibilities with the increase in degrees of freedom. Without the constraints giving them shape, these degrees of freedom would collapse into uninformative noise.

In Hopfield networks and the SO model, the constraining framework is given by the architecture of the network. This paper suggests ways in which higher degrees of freedom could be incorporated within this structure, thus increasing its flexibility and possibilities for adaptation. We propose that metaplasticity and homeostatic regulation could be modeled through a dynamical adjustment of the learning rate with regards to synaptic weights or the energy state of the system, through adaptively enhancing modularity in the system, and through modifications to the activation criteria. These proposals now need to be defined algorithmically and tested to observe how the network's behavior changes, and whether it will be able to solve new types of tasks. Some of these mechanism may already exist in deep learning models with more sophisticated architectures. It is also possible that once implemented, parallels may be found between these features, and features in supervised learning models. Nevertheless, it is far from certain whether the majority of adaptive systems, both biological and non-biological, have the capacity to implement supervised learning with sophisticated architectures and learning rules. Most learning in adaptive systems seems to proceed in an unsupervised manner [65, 66]. Therefore, achieving an associative memory model capable of performing complex adaptive learning tasks would be helpful not only to diversify the deep learning tools at our disposal,

but also to improve our understanding of how adaptive systems in the natural world solve the task of continually adjusting and thriving in their environments.

Acknowledgements

I thank Natalya Weber (OIST) for illuminating discussions on the SO model and Hopfield networks, as well as valuable feedback during the writing of this paper. I am grateful to Prof Tom Froese (OIST) for his guidance on the direction of this work.

References

[1] W. C. Abraham. Metaplasticity: Tuning synapses and networks for plasticity. *Nat Rev Neurosci*, 9(5):387–387, 2008-05.

[2] D. J. Amit, H. Gutfreund, and H. Sompolinsky. Storing Infinite Numbers of Patterns in a Spin-Glass Model of Neural Networks. *Phys. Rev. Lett.*, 55(14):1530–1533, 1985-09-30.

[3] R. C. Anafi and J. H. T. Bates. Balancing Robustness against the Dangers of Multiple Attractors in a Hopfield-Type Model of Biological Attractors. *PLOS ONE*, 5(12):e14413, 2010-12-22.

[4] P. Arkachar and M. D. Wagh. Criticality of lateral inhibition for edge enhancement in neural systems. *Neurocomputing*, 70(4):991–999, 2007-01-01.

[5] M. C. Avery and J. L. Krichmar. Neuromodulatory Systems and Their Interactions: A Review of Models, Theories, and Experiments. *Frontiers in Neural Circuits*, 11, 2017.

[6] M. L. Barnes, P. Wang, J. E. Cinner, N. A. J. Graham, A. M. Guerrero, L. Jasny, J. Lau, S. R. Sutcliffe, and J. Zamborain-Mason. Social determinants of adaptive and transformative responses to climate change. *Nat. Clim. Chang.*, 10(9):823–828, 2020-09.

[7] E. L. Bienenstock, L. N. Cooper, and P. W. Munro. Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, 2(1):32–48, 1982-01-01.

[8] G. E. Billman. Homeostasis: The Underappreciated and Far Too Often Ignored Central Organizing Principle of Physiology. *Frontiers in Physiology*, 11, 2020.

[9] S. Biswas, S. Manicka, E. Hoel, and M. Levin. Gene regulatory networks exhibit several kinds of memory: Quantification of memory in biological and random transcriptional networks. *iScience*, 24(3):102131, 2021-02-01.

[10] T. V. P. Bliss and G. L. Collingridge. A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361(6407):31–39, 1993-01.

[11] S. Bouret and S. J. Sara. Network reset: A simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neurosciences*, 28(11):574–582, 2005-11-01.

[12] T. F. Burns and T. Fukai. Simplicial Hopfield networks, 2023-05-09. Comment: 36 pages, 7 figures, published as a conference paper at ICLR 2023.

[13] J. Burrone and V. N. Murthy. Synaptic gain control and homeostasis. *Current Opinion in Neurobiology*, 13(5):560–567, 2003-10-01.

[14] L. Cantini and M. Caselle. Hope4Genes: A Hopfield-like class prediction algorithm for transcriptomic data. *Sci Rep*, 9(1):337, 2019-01-23.

[15] P. Curtin, C. Austin, A. Curtin, C. Gennings, C. Figueroa-Romero, K. A. Mikhail, T. M. Botero, S. A. Goutman, E. L. Feldman, and M. Arora. Dysregulated biodynamics in metabolic attractor systems precede the emergence of amyotrophic lateral sclerosis. *PLOS Computational Biology*, 16(4):e1007773, 2020-04-15.

[16] I. M. De la Fuente, L. Martínez, J. Carrasco-Pujante, M. Fedetz, J. I. López, and I. Malaina. Self-Organization and Information Processing: From Basic Enzymatic Activities to Complex Adaptive Cellular Behavior. *Frontiers in Genetics*, 12, 2021.

[17] S. DuBrow, N. Rouhani, Y. Niv, and K. A. Norman. Does mental context drift or shift? *Current Opinion in Behavioral Sciences*, 17:141–146, 2017-10-01.

[18] G. M. Edelman and J. A. Gally. Degeneracy and complexity in biological systems. *PNAS*, 98(24):13763–13768, 2001-11-20.

[19] S. Farashahi, C. H. Donahue, P. Khorsand, H. Seo, D. Lee, and A. Soltani. Metaplasticity as a Neural Substrate for Adaptive Learning and Choice under Uncertainty. *Neuron*, 94(2):401–414.e6, 2017-04-19.

[20] T. E. Faust, G. Gunner, and D. P. Schafer. Mechanisms governing activity-dependent synaptic pruning in the developing mammalian CNS. *Nat Rev Neurosci*, 22(11):657–673, 2021-11.

[21] T. Froese, C. Gershenson, and L. R. Manzanilla. Can Government Be Self-Organized? A Mathematical Model of the Collective Social Organization of Ancient Teotihuacan, Central Mexico. *PLOS ONE*, 9(10):e109966, 2014-10-10.

[22] T. Froese, N. Weber, I. Shpurov, and T. Ikegami. From autopoiesis to self-optimization: Toward an enactive model of biological regulation. *Biosystems*, 230:104959, 2023-08-01.

[23] M. Hallett. Neurophysiology of dystonia: The role of inhibition. *Neurobiology of Disease*, 42(2):177–184, 2011-05-01.

[24] N. Haroush and S. Marom. Inhibition increases response variability and reduces stimulus discrimination in random networks of cortical neurons. *Sci Rep*, 9(1):4969, 2019-03-21.

[25] J. L. He, I. Fuelscher, J. Coxon, P. Barhoun, D. Parmar, P. G. Enticott, and C. Hyde. Impaired motor inhibition in developmental coordination disorder. *Brain Cogn*, 127:23–33, 2018-11.

[26] D. O. Hebb. *The Organization of Behavior; a Neuropsychological Theory*. The Organization of Behavior; a Neuropsychological Theory. Wiley, 1949.

[27] L. J. Herstel and C. J. Wierenga. Network control through coordinated inhibition. *Current Opinion in Neurobiology*, 67:34–41, 2021-04-01.

[28] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982-04.

[29] G. James, D. Witten, T. Hastie, and R. Tibshirani. *An Introduction to Statistical Learning*, volume 103 of *Springer Texts in Statistics*. Springer New York, 2013.

[30] K. Kamino, J. M. Keegstra, J. Long, T. Emonet, and T. S. Shimizu. Adaptive tuning of cell sensory diversity without changes in gene expression. *Science Advances*, 6(46):eabc1087, 2020-11-13.

[31] M. Khona and I. R. Fiete. Attractor and integrator networks in the brain. *Nat Rev Neurosci*, 23(12):744–766, 2022-12.

[32] N. R. Kinsky, W. Mau, D. W. Sullivan, S. J. Levy, E. A. Ruesch, and M. E. Hasselmo. Trajectory-modulated hippocampal neurons persist throughout memory-guided navigation. *Nat Commun*, 11(1):2443, 2020-05-15.

[33] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017-03-28.

[34] H. Kitano. Biological robustness. *Nat Rev Genet*, 5(11):826–837, 2004-11.

[35] D. Krotov. A new frontier for Hopfield networks. *Nat Rev Phys*, 5(7):366–367, 2023-07.

[36] D. Krotov and J. J. Hopfield. Dense Associative Memory for Pattern Recognition. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[37] J. Luboeinski and C. Tetzlaff. Memory consolidation and improvement by synaptic tagging and capture in recurrent neural networks. *Commun Biol*, 4(1):1–17, 2021-03-03.

[38] C. M. MacLeod. The concept of inhibition in cognition. In D. S. Gorfein and C. M. MacLeod, editors, *Inhibition in Cognition.*, pages 3–23. American Psychological Association, 2007.

[39] K. MacVittie, J. Halámek, V. Privman, and E. Katz. A bioinspired associative memory system based on enzymatic cascades. *Chemical Communications*, 49(62):6962–6964, 2013.

[40] J. E. Markowitz, W. F. Gillis, M. Jay, J. Wood, R. W. Harris, R. Cieszkowski, R. Scott, D. Brann, D. Koveal, T. Kula, C. Weinreb, M. A. M. Osman, S. R. Pinto, N. Uchida, S. W. Linderman, B. L. Sabatini, and S. R. Datta. Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature*, 614(7946):108–117, 2023-02.

[41] P. Meier, A. Finch, and G. Evan. Apoptosis in development. *Nature*, 407(6805):796–801, 2000-10.

[42] A. Morales and T. Froese. Unsupervised Learning Facilitates Neural Coordination Across the Functional Clusters of the C. elegans Connectome. *Front Robot AI*, 7:40, 2020.

[43] E. J. Müller, B. R. Munn, M. J. Redinbaugh, J. Lizier, M. Breakspear, Y. B. Saalmann, and J. M. Shine. The non-specific matrix thalamus facilitates the cortical information processing modes relevant for conscious awareness. *Cell Reports*, 42(8):112844, 2023-08-29.

[44] G. Natoli and R. Ostuni. Adaptation and memory in immune responses. *Nat Immunol*, 20(7):783–792, 2019-07.

[45] L. K. Nguyen. Regulation of oscillation dynamics in biochemical systems with dual negative feedback loops. *Journal of The Royal Society Interface*, 9(73):1998–2010, 2012-03-14.

[46] G. F. Oster. Lateral inhibition models of developmental processes. *Mathematical Biosciences*, 90(1):265–286, 1988-07-01.

[47] E. A. D. Paolo. Organismically-inspired robotics: Homeostatic adaptation and teleology beyond the closed sensorimotor loop.

[48] G. R. Poe. Sleep Is for Forgetting. *J. Neurosci.*, 37(3):464–473, 2017-01-18.

[49] L. Ramos-Alonso, P. Holland, S. Le Gras, X. Zhao, B. Jost, M. Bjørås, Y. Barral, J. M. Enserink, and P. Chymkowitch. Mitotic chromosome condensation resets chromatin to safeguard transcriptional homeostasis during interphase. *Proceedings of the National Academy of Sciences*, 120(4):e2210593120, 2023-01-24.

[50] H. Ramsauer, B. Schäfl, J. Lehner, P. Seidl, M. Widrich, T. Adler, L. Gruber, M. Holzleitner, M. Pavlović, G. K. Sandve, V. Greiff, D. Kreil, M. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter. Hopfield Networks is All You Need, 2021-04-28. Comment: 10 pages (+ appendix); 12 figures; Blog: https://ml-jku.github.io/hopfield-layers/; GitHub: https://github.com/ml-jku/hopfield-layers.

[51] D. R. Schoenberg and L. E. Maquat. Regulation of cytoplasmic mRNA decay. *Nat Rev Genet*, 13(4):246–259, 2012-04.

[52] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni. Green AI, 2019-08-13. Comment: 12 pages.

[53] D. N. Scott and M. J. Frank. Adaptive control of synaptic plasticity integrates micro- and macroscopic network function. *Neuropsychopharmacol.*, 48(1):121–144, 2023-01.

[54] A. Soltani and A. Izquierdo. Adaptive learning under expected and unexpected uncertainty. *Nat Rev Neurosci*, 20(10):635–644, 2019-10.

[55] M. Stacho and D. Manahan-Vaughan. The Intriguing Contribution of Hippocampal Long-Term Depression to Spatial Learning and Long-Term Memory. *Frontiers in Behavioral Neuroscience*, 16, 2022.

[56] L. Susman, M. Kohram, H. Vashistha, J. T. Nechleba, H. Salman, and N. Brenner. Individuality and slow dynamics in bacterial growth homeostasis. *Proceedings of the National Academy of Sciences*, 115(25):E5679–E5687, 2018-06-19.

[57] A. Szedlak, S. Sims, N. Smith, G. Paternostro, and C. Piermarocchi. Cell cycle time series gene expression data encoded as cyclic attractors in Hopfield systems. *PLOS Computational Biology*, 13(11):e1005849, 2017-11-17.

[58] R. Watson, C. L. Buckley, R. Mills, and A. Davies. Associative memory in gene regulation networks. pages 659–666. MIT Press, 2010.

[59] R. A. Watson, C. L. Buckley, and R. Mills. The Effect of Hebbian Learning on Optimisation in Hopfield Networks, 2009-06-10.

[60] R. A. Watson, C. L. Buckley, and R. Mills. Optimization in "self-modeling" complex adaptive systems. *Complexity*, 16(5):17–26, 2011.

[61] N. Weber, W. Koch, O. Erdem, and T. Froese. On the use of associative memory in Hopfield networks designed to solve propositional satisfiability problems, 2023-08-14. Comment: 7 pages, 4 figures.

[62] N. Weber, W. Koch, and T. Froese. Scaling up the self-optimization model by means of on-the-fly computation of weights. In *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1276–1282, 2022-12-04. Comment: 7 pages, 7 figures.

[63] M. Wehr and A. M. Zador. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–446, 2003-11.

[64] Y. Yao and C. Wang. Dedifferentiation: Inspiration for devising engineering strategies for regenerative medicine. *npj Regen Med*, 5(1):1–11, 2020-07-31.

[65] L. Zaadnoordijk, T. R. Besold, and R. Cusack. Lessons from infant learning for unsupervised machine learning. *Nat Mach Intell*, 4(6):510–520, 2022-06.

[66] A. M. Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nat Commun*, 10(1):3770, 2019-08-21.

[67] M. Zarco and T. Froese. Self-Optimization in Continuous-Time Recurrent Neural Networks. *Frontiers in Robotics and AI*, 5, 2018.

Appendix 1: Inhibition and degeneracy as adaptive mechanisms in biology

Though homeostasis makes use of inhibition as a regulatory mechanism, inhibitory processes are not exclusive to it. They can have intrinsic functions beyond restoring balance to a system. Cognitive control, which relies on inhibition of certain behaviors, is central to executive functioning [38]. Similarly, inhibition is necessary for motor coordination [23, 25], for sharpening of sensory perception [24, 63], and for the proper functioning of brain networks, through maintenance of an excitatory/inhibitory (E/I) balance [27]. Considering LTD as a form of inhibition, its importance in memory formation in the hippocampus has also been highlighted by several studies [55].

Redundancy and degeneracy have been described as fundamental features of biological systems [18], contributing to their robustness and adaptivity [34]. Degeneracy is defined as "the ability of elements that are structurally different to perform the same function or yield the same output" [18], whereas redundancy involves structurally identical elements performing the same function. A system reliant on a single pathway to perform critical tasks is extremely vulnerable to the failure of said pathway, even if it is optimized for the task in question. On the other hand, a multiplicity of pathways, though they may be non-optimal or less efficient, allows for higher robustness at the system-level in case of failure. Examples of degeneracy can be seen in the genetic code, immune system and metabolic pathways.

Arguably, some developmental processes can be considered as a combination of degeneracy and inhibition (as negation or suppression). Apoptosis in morphogenesis [41, 46] and synaptic pruning in childhood [20] are two such examples. In both cases, the two-step process of generating a certain number of elements (cells or synapses/dendrites), followed by the removal of now-supernumerary elements, improves definition and functionality. This highlights the role of inhibition as a mechanism for pattern enhancement/segregation [46, 55].

Appendix 2: Further implementations of adaptive mechanisms

Metaplasticity has been introduced in sophisticated reward-based learning models, such as this study by Farashahi et al., which devised a model of adaptive learning under uncertainty, with metaplasticity governing the transition of synapses between hidden states of varying stability [19]. Interestingly, this model did not require explicit optimization or knowledge of the task structure to achieve results that were comparable to experimental data. This suggests that implementing metaplasticity in an unsupervised learning setting is not far-fetched. Another model, based on a recurrent neural network integrating calcium dynamics, and a leaky-integrate-and-fire mechanism, showed how incorporating STC improved memory consolidation and recall [37]. Lateral inhibition has been implemented in a number of feed-forward architectures [4]. These examples implemented in different network architectures demonstrate the usefulness of these adaptive mechanisms in improving learning.

A temporal dimension could be added to the synaptic weights by introducing a decay term in which the weights would be set to converge to zero in a stepwise or gradual manner, after a defined number of steps. This form of "forgetting" could be of interest in the context of continuous learning, as it could allow for new patterns to be encoded without a complete erasure of previous patterns, as presumably, the strongest weights would not have fully been reduced to 0.

The resetting of weights through decay is different from the resets defined in the SO model, which involve resetting network *states*. Incidentally, it would be interesting to compare the effects of these resets with the aforementioned suggestion to mimic homeostatic oscillations through a trigonometric activation function (see Section 3).

These partial state resets could be seen as mimicking cyclical biological processes such as sleep and cell division. A form of context-dependent reset, resembling brain network shifts triggered by noradrenaline, could be achieved by linking resets to network energy. For example, if the energy level remains above a certain threshold for a defined number of steps, a state reset could be triggered, with the idea that a high energy level signifies that the network is far from convergence, and that a reset might speed up constraint satisfaction. Another method might require specific modifications to the network structure, by defining a cluster of nodes that, when fully activated, trigger a network reset. But this might be more difficult to design as it is unclear what type of connections this cluster should have with the rest of the network.