Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Neurobiology

Perspective

Energetically efficient learning in neuronal networks

Aaron Pache[1] and Mark C.W. van Rossum[1,2]

Abstract

Human and animal experiments have shown that acquiring and storing information can require substantial amounts of metabolic energy. However, computational models of neural plasticity only seldom take this cost into account, and might thereby miss an important constraint on biological learning. This review explores various ways to reduce energy requirements for learning in neural networks. By comparing the resulting learning rules to cognitive and neurophysiological observations, we discuss how energy efficiency might have shaped biological learning.

Addresses

[1] School of Mathematical Sciences, University of Nottingham, Nottingham NG7 2RD, United Kingdom
[2] School of Psychology, University of Nottingham, Nottingham, United Kingdom

Corresponding author: van Rossum, Mark CW (lpzmv@nottingham.ac.uk)

Current Opinion in Neurobiology 2023, **83**:102779

This review comes from a themed issue on **Computational Neuroscience 2023**

Edited by **Jeanette Hellgreny Kotaleski** and **Tatjana Tchumatchenko**

For a complete overview see the Issue and the Editorial

Available online xxx

https://doi.org/10.1016/j.conb.2023.102779

Introduction

Metabolism is central to biology. Survival requires collecting sufficient amounts of energy, and once obtained, spending it carefully. In humans, the brain is the most energy-consuming organ per weight, after the kidneys. It consumes about 20W, the same as a dim light bulb. While locally brain energy consumption can double when neural activity is high [1], globally consumption appears fairly constant, even during sleep [2]. In comparison, cyclists can produce a peak power of about 1 kW— ten times the resting whole-body metabolism. So while the brain consumes a moderate amount, it does so incessantly. It therefore makes sense that animals would seek to reduce their brain's energetic needs. On short time scales neural processing adapts depending on current energy reserves, e.g. Refs. [3,4], while on long timescales evolutionary pressure can improve the brain's energy efficiency, for instance by minimizing wiring length [5].

Energy efficiency is now widely seen as a neural design principle [6,7]. A number of physiological phenomena, that would be counter-intuitive otherwise, can be understood from an energy perspective. For instance, synaptic release is unreliable, thereby introducing noise in neural processing. However, it has been argued that unreliable release is energetically beneficial [8,9]. As another example, neural codes are typically sparse, so that a given stimulus only activates a small fraction of neurons. While computationally a dense code would have a higher information capacity, sparse codes maximize the information capacity per energy [10]. Thus energy efficiency considerations provide a unifying explanation for disparate findings.

The metabolic cost of learning

More recently the metabolic cost of learning has been added to the brain's energy budget, see Ref. [11] for a recent review. Experiments in Drosophila indicate that already for simple tasks these costs are substantial. In Mery and Kawecki [12] flies were exposed to a conditioning protocol and learned to associate an odor to a mechanical shock, so that the odor would subsequently be avoided. After the protocol, all feeding was stopped and the time to die from starvation was measured. It was found that the conditioning substantially reduced the lifespan. After controlling for increased energy consumption after exposure to just unconditioned or conditioned stimuli, which for instance could increase motor activity or stress level, the decrease in lifespan was still some 20%. Likewise, if instead the flies were not starved, they binged on food intake right after learning [13,14]. The energy associated with learning in these experiments was estimated at some 10 mJ per bit of stored information, 6 orders of magnitude more than a modern computer solid-state drive [15].

In mammals, starvation experiments are less common for ethical reasons, but there are hints that high glucose availability improves memory [16,17], and experiments are starting to reveal an intimate link between metabolism and memory, e.g. Ref. [18]. In humans, there is evidence, mainly in children and elderly, for a link between energy availability and memory, for a review see Ref. [19].

The high energy costs of learning suggest that animals try to learn while expending as little energy as possible and that it would be beneficial to study how efficiency has shaped biological learning. Ideally, performance measures such as learning speed, capacity, or retention would be unaffected, but typically one would expect a trade-off between efficiency and performance. Currently, however, most computational studies pose no limit on plasticity or its energy costs.

A synaptic plasticity energy model

It is unclear why learning is metabolically so demanding. Learning could be demanding on the synaptic level, cellular, or organism level. The biophysical cost of synaptic plasticity in mammals was first estimated by Karbowski [20], who found protein phosphorylation to be the largest consumer, far exceeding cost estimates for protein synthesis, transport, and costs such as actin thread milling. It was estimated at 3×10^6ATP/synapse/min. However this still amounts to a relatively small amount of energy, some 10% of the total energy in a rat brain [20]. Moreover, phosphorylation is more characteristic of inexpensive, early-phase long-term potentiation (LTP) than of the expensive late-phase LTP (see below). So a large amount of energy use remains unaccounted for [15].

To make progress we assume that the energy requirements are dominated by synaptic plasticity. A parsimonious model of energy M accumulates the amount of synaptic changes over time t and over synapses i.

$$M = \sum_{t,i} |w_i(t+1) - w_i(t)|.$$

A more general variant includes a power α, such that $M = \sum_{t,i} | w_i(t+1) - w_i(t) |^{\alpha}$. For instance, in the limit $\alpha \to 0$, it counts the number of synaptic updates, irrespective of size [21,22]. Note that the model ignores any synergistic or competitive interactions across synapses or across time. These likely exist, but have not been quantified.

While energy efficiency likely pervades many forms of learning, to examine efficiency we focus here on an artificial task, the MNIST digit classification task, trained with the well-known ‘back propagation of error‘ algorithm. While this task is somewhat artificial, it is a staple machine-learning benchmark and one can draw on an extensive literature that has explored this task, albeit in the absence of energy considerations.

Briefly, in the MNIST task, one is provided with a set of 60,000 images of handwritten digits (0.9), as well as the digit that each image is supposed to represent, i.e. the class label. The network has ten output units. The desired goal of the network is that each input image activates its corresponding output unit, while the other output units remain silent (one-hot encoding).

When training the network, samples are provided one by one (for biological realism we do not consider batching), and the output of the network is calculated. This output is compared to the desired output using a (surrogate) loss function such as a mean squared error loss or a cross-entropy loss, Figure 1 top. The loss function and its derivatives with respect to each synaptic weight are calculated. Each weight is then updated with an amount proportional to the derivative, a procedure known as gradient descent. Performance is periodically evaluated during training by testing the accuracy of a subset of the data that was kept aside. The presentation of samples and updating of the weights is repeated until a desired accuracy is reached.

Using this setup we explored various strategies to save energy. That is, can we reach the same accuracy on the task using less energy? We describe three methods to save energy that we recently developed. The first two sparsify plasticity [23,21], while the third one relies on switching between different forms of plasticity [22].

Saving method I: skip over correct samples

As soon as a training sample activates the target output unit more than the other output units, the sample is classified correctly. However, because the standard neural network training scheme minimizes the loss function described above and not the number of errors, correctly classified samples will still cause synaptic updates in the network. While updating the weights for correctly classified samples leads to robust decision boundaries, it is energetically inefficient. From the machine learning literature, it is known that is not necessary to modify the synaptic weights on every sample, although knowing which samples to skip is not trivial [24,25].

In line with this, people tend to pay particular attention to errors when learning. Performance improves when feedback is provided following an error. Error feedback elicits a large negative deflection in the voltage in the electroencephalogram signal (EEG). The larger the EEG signal, the more that memory item improves [26,27]. Meanwhile, feedback on correct responses has little effect on learning [28]. We hypothesized that blocking plasticity on correct examples could save energy and maintain performance.

The first saving method [23] only learns when the sample is misclassified and skips over samples that are correctly classified, Figure 1bottom row, left. This by itself leads to good training performance, but generalization performance is poor and fragile, because it finds decision boundaries that run close to the data points

Figure 1

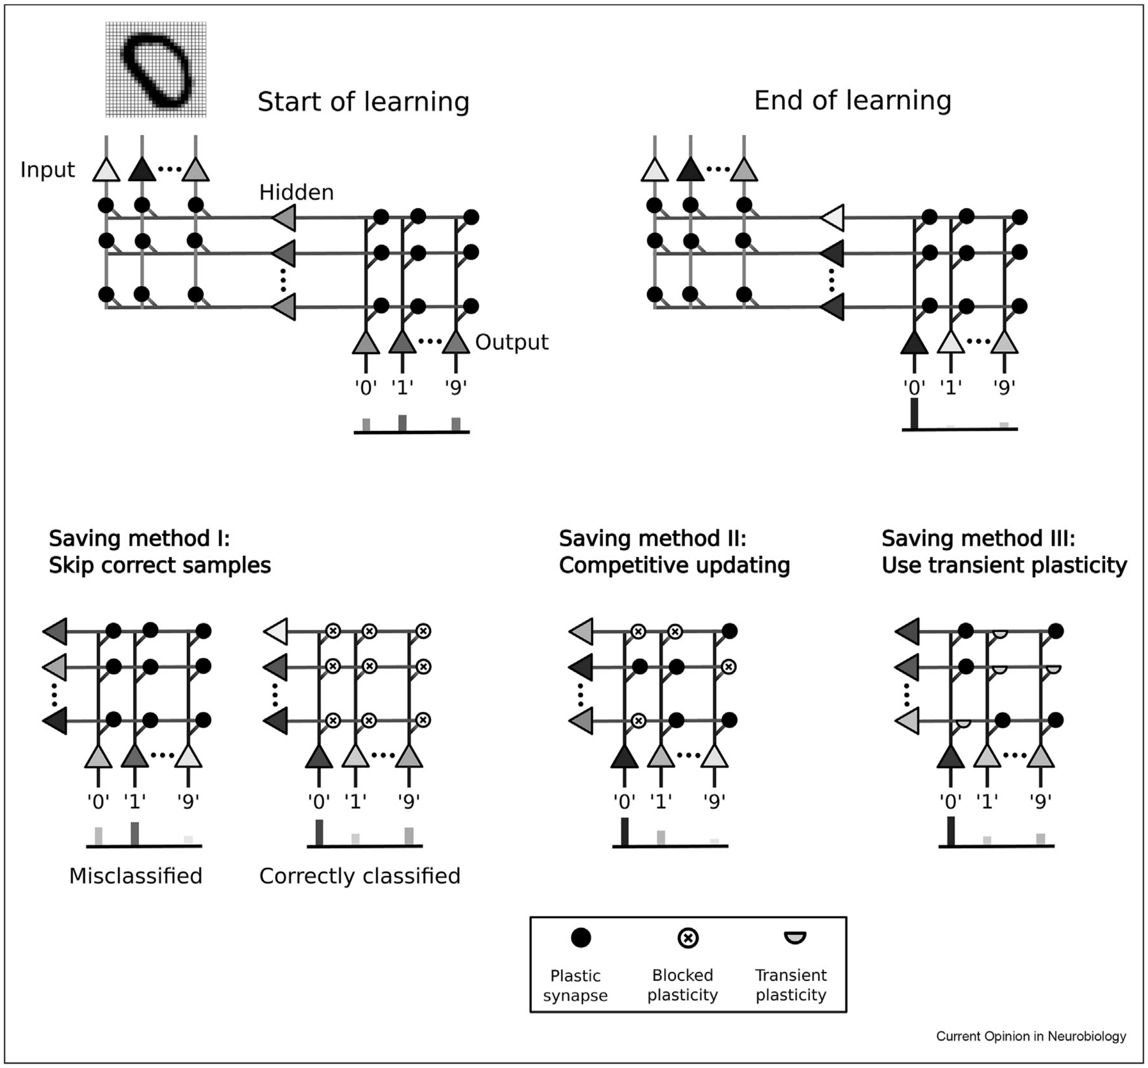

Ways to reduce energy requirements of synaptic plasticity. Top: Traditional neural network training on a classification task. Input images are provided and all synapses (circles) are modified on every sample until the output closely matches the target output.
Bottom row: Three methods to save plasticity energy. Left: by blocking plasticity for correctly classified inputs. Middle: by restricting plasticity to a subset of synapses. Right: by combining metabolically inexpensive transient plasticity with expensive persistent plasticity. For conciseness, the input layer neurons and synapses are not shown on the bottom row.

from the training set. With a tweak, better generalization is achieved at only little extra energetic cost, namely by not only training on the samples that are currently incorrect but also those that were ever incorrectly classified in the past. In line with this, it was shown that currently correct but previously incorrect memory items evoke a larger EEG response than always-correct items, and similar to incorrect items [27].

How much energy can be saved with this method, depends on the size of the data set relative to the task complexity. For standard MNIST, this method used 2.2 × less energy (as defined above) than an equivalent backprop network, irrespective of the network size [23]. But when the data set is large, for instance, because data augmentation was used, one can skip over many more samples. For instance, for the EMNIST dataset (equivalent to MNIST but with 4 times more data [29]) savings increased to 8.4 × . The reason for using large data sets is that they lead to better performance. Interestingly, because the algorithm automatically cherry picks the data it learns from, the benefit of using larger data sets is preserved under this method, despite ignoring many samples.

How might such an algorithm be implemented in the brain? We suggest that plasticity is gated by error signals, so that little or no learning takes place when the response of the animal is correct. We suggest that the hippocampus could play a key role. The hippocampus is a medium-term storage buffer that allows for off-line learning, for instance during sleep replay. We propose that it preferentially stores data that is in conflict with the current knowledge framework, guided by conflict monitoring in the prefrontal cortex [30].

As an aside, the method is also of use for machine learning practitioners. While skipping over correct samples increases the number of cycles over the data set required to reach criterion performance, feed-forward passes through the network are very quick. As a result, the increase in cycles is outweighed by the reduction in required synaptic updates. The net effect is a speed-up of network training in computers: 1.2 × faster for MNIST, and 7.6 × faster for EMNIST.

Saving method II: restrict plasticity to some synapses

The second method sparsifies plasticity on a single trial [21]. The backprop algorithm in principle prescribes updates to all synapses in the network. (It can happen that the update is actually zero, for instance when the gradient is zero, but it is not systematically controlled). A more energy-efficient method is to limit plasticity to a few synapses, Figure 1 bottom row, middle. As above, such sparsification of plasticity reduces the number of updates and can save energy. While it is possible to designate some synapses as static and others as plastic from the outset, a more efficient method is to competitively update only the synapses with the largest changes. Specifically, on each input sample, the proposed weight changes are calculated, but only the ones with largest magnitude are actually implemented. The larger the network, the more synapses can remain the same. On different samples, different synapses will be updated. Hence plasticity remains distributed across the synapses and the benefit of using large networks remains [21]. In very large networks, further savings were found when plasticity is restricted to sub-nets so that the input and output plasticity of neurons are coordinated.

The saving achieved with this method depends on the number of synapses in the network. If the network has just enough synapses for the task complexity, not much efficiency can be gained, but the more synapses, the larger the saving. This makes this strategy particularly relevant for mammalian brains with their large number of neurons and synapses. For instance, macaque V1 has some 150 million neurons and some 300×10^9 synapses [31]. Suppose one learns to recognize, say, a new car model. If one were to model this with a standard neural network with standard backprop rule [32], *all* synapses in the visual cortex are a little bit updated on *every* exposure. Such excessive updating is wasteful and unbiological. For MNIST, by restricting plasticity, energy requirements can be reduced by 3 orders of magnitude for a V1-sized network [21]. As an added benefit, restricting synaptic updates can help to prevent catastrophic forgetting, e.g. Ref. [33].

Experimentally, synaptic plasticity is indeed constrained in space and time. For instance, during motor learning plasticity appears to be restricted to certain dendritic branches [34], but see Ref. [35]. Moreover, it is not possible to induce late-phase LTP twice in rapid succession [36,37]. A prediction of a competitive mechanism is that if plasticity were blocked at one synapse, another one would take over. Some indications of such competition exist, but the mechanism and general rules are currently unknown [38–40].

Saving method III: synaptic caching

The above saving methods sparsify plasticity. The third method [22] exploits that both in insects and mammals there are multiple forms of plasticity with distinct molecular pathways and different temporal persistence [41–43]. In the fly, persistent LongTerm Memory (LTM) is metabolically more expensive than transient Anaesthesia Resistant Memory (ARM) memory [12,14]. In mammals, one distinguishes between early-phase LTP and late-phase LTP and similar evidence for a difference in metabolic cost is arising [44,43]. This cost difference opens up the possibility to save energy by combining different forms of plasticity, Figure 1 bottom row, right. In particular, one can initially ‘cache’ information in transient forms of plasticity and only intermittently consolidate into expensive persistent forms, so-called synaptic caching [22]. There is a trade-off between consolidating too often, which is costly, and too rarely so that information will have been lost. The maximum saving for the MNIST task is achieved when all updates can be accumulated in the transient memory at no cost and only are consolidated at the end of learning. In that case, the energy required to train the network to the same accuracy is 1–2 orders of magnitude smaller than a network that uses only persistent plasticity [22].

Outlook

In summary, we have described a number of ways to increase the energy efficiency of network learning. Each saving algorithm addresses a different inefficiency: Skipping saves most when training on large data sets; competitive updating saves most when networks are over-dimensioned relative to the task; and synaptic caching works best when updates can be temporarily stored. This complementarity also opens the possibility to combine the different saving algorithms.

Incorporating metabolic demands of learning leads to modifications of learning rules that are in line with experimental data. As with the metabolic cost of activity, alternative explanations could be raised. For instance, we argued above that hippocampus should preferentially store wrong outcomes to save energy. But alternatively, the capacity of the hippocampus might be limited, or, plasticity might be sparse to prevent overwriting. Such explanations need not be mutually exclusive; however, the metabolic argument is attractive because of its ability to unify explanations.

While we have concentrated on supervised learning of classification tasks, we expect that efficiency concerns run like a red thread across learning paradigms. For instance, the insect experiments discussed use a reinforcement learning setup, while the synaptic caching algorithm has also been applied to perceptron and STDP learning [22]. Similarly, saving strategies likely run from the biophysical level, e.g. intracellular receptor transport, to the synaptic plasticity rules and all the way up to the algorithmic level, e.g. deciding if novel information is worth replay. Energy-efficient learning could also suggest specific network architectures that allow for efficient learning.

When it comes to the energy needed for synaptic transmission and spike generation, precise models exist [7]. In contrast, there is no accurate model of the energy consumed by learning. We have used a cost model that is proportional to the number and size of synaptic updates, but this is likely an oversimplification. We also ignored other memory-related cost, such as might come from replay. As more detailed insight is gained into the energetics of learning, the saving strategies can be re-evaluated.

Memory storage in traditional and neuromorphic computer hardware is expensive as well [45], so that our insight gained from biological learning might be adopted there. Modern hardware is highly optimized for traditional backprop learning; however, biological inspiration can still improve efficiency. For instance, the first strategy (skipping over correct samples) is easy to implement and, as we saw, can substantially reduce network training time.

We have only considered fixed efficiency strategies. But learning and plasticity are adaptive, so that they are reduced in periods of energy scarcity. The difference in energy cost of different plasticity pathways means that any efficiency mechanism mainly acts on expensive long-lasting persistent plasticity. For example, in flies, aversive LTM is halted under low energy conditions [14]. Such adaptive regulation is also found in mammals where late-phase LTP is halted under low energy conditions, while early-phase LTP is not [44]. Dopamine seems to play a key role to promote long-lasting storage in both systems. Optimal adaptive strategies and their biological correlates are an obvious candidate for future investigation. Such adaptive strategies could again extend to the network level [46], for example, changing metabolic demands by modifying coding sparsity, or adjusting replay processes, are some of the intriguing possibilities.

Declaration of competing interest

The authors declare no competing interests.

Data availability

No data was used for the research described in the article.

Acknowledgments

We would like to thank Simon Laughlin, HoLing Li, Jiamu Jiang, Emilie Foyard, and Maxime Girard for discussions.

References

Papers of particular interest, published within the period of review, have been highlighted as:

* of special interest

1. Hyder F, Rothman DL, Shulman RG: **Total neuroenergetics support localized brain activity: implications for the interpretation of fmri**. *Proc Natl Acad Sci USA* 2002, **99**: 10771–10776, https://doi.org/10.1073/pnas.132272299.
2. DiNuzzo M, Nedergaard M: **Brain energetics during the sleep–wake cycle**. *Curr Opin Neurobiol* 2017, **47**:65–72.
3. Longden KD, Muzzu T, Cook DJ, Schultz SR, Krapp HG: **Nutritional state modulates the neural processing of visual motion**. *Curr Biol* 2014, **24**:890–895.
4. Padamsey Z, Katsanevaki D, Dupuy N, Rochefort NL: **Neocortex saves energy by reducing coding precision during food scarcity**. *Neuron* 2022, **110**:280–296.
5. Bullmore E, Sporns O: **The economy of brain network organization**. *Nat Rev Neurosci* 2012, **13**:336–349, https://doi.org/10.1038/nrn3214.
6. Lennie P: **The cost of cortical computation**. *Curr Biol* 2003, **13**: 493–497, https://doi.org/10.1016/S0960-9822(03)00135-0.
7. Attwell D, Laughlin SB: **An energy budget for signaling in the grey matter of the brain**. *J Cerebr Blood Flow Metabol* 2001, **21**: 1133–1145, https://doi.org/10.1097/00004647-200110000-00001.
8. Levy WB, Baxter RA: **Energy-efficient neuronal computation via quantal synaptic failures**. *J Neurosci* 2002, **22**:4746–4755.
9. Harris JJ, Jolivet R, Attwell D: **Synaptic energy use and supply**. *Neuron* 2012, **75**:762–777.
10. Levy WB, Baxter RA: **Energy efficient neural codes**. *Neural Comput* 1996, **8**:531–543, https://doi.org/10.1162/neco.1996.8.3.531.
11. Padamsey Z, Rochefort NL: **Paying the brain's energy bill**. *Curr Opin Neurobiol* 2023, **78**, 102668, https://doi.org/10.1016/j.conb.2022.102668.
12. * Mery F, Kawecki TJ: **A cost of long-term memory in drosophila**. *Science* 2005, **308**:1148, https://doi.org/10.1126/science.1111331.

Shows that persistent learning is metabolically expensive in fruitflies

13. Plaçais P-Y, de Tredern È, Scheunemann L, Trannoy S, Goguel V, Han K-A, Isabel G, Preat T: **Upregulated energy metabolism in the drosophila mushroom body is the trigger for long-term memory**. *Nat Commun* Jun. 2017, **8**, https://doi.org/10.1038/ncomms15510.

14. * Plaçais P-Y, Preat T: **To favor survival under food shortage, the brain disables costly memory**. *Science* 2013, **339**: 440–442, https://doi.org/10.1126/science.1226018.

Show that learning adapts depending on starvation.

15. Girard M, Jiang J, van Rossum MC: *Estimating the energy requirements for long term memory formation*. 2023, 2301.09565. arxiv.

16. Gold PE: **Glucose modulation of memory storage processing**. *Behav Neural Biol* 1986, **45**:342–349.

17. Hall JL, Gonder-Frederick L, Chewning W, Silveira J, Gold P: **Glucose enhancement of performance of memory tests in young and aged humans**. *Neuropsychologia* 1989, **27**: 1129–1138.

18. Sgammeglia N, Sprecher SG: **Interplay between metabolic energy regulation and memory pathways in drosophila**. *Trends Neurosci* 2022, **45**:539–549, https://doi.org/10.1016/j.tins.2022.04.007.

19. Smith MA, Riby LM, van Eekelen JAM, Foster JK: **Glucose enhancement of human memory: a comprehensive research review of the glucose memory facilitation effect**. *Neurosci Biobehav Rev* 2011, **35**:770–783.

20. * Karbowski J: **Metabolic constraints on synaptic learning and memory**. *J. of neurophysiology* 2019, **122**:1473–1490.

First bottom-up estimate of plasticity cost.

21. * van Rossum MCW: **Competitive plasticity to reduce the energetic costs of learning**. *bioRxiv* 2023:2023. 04.

Energy saving model that competitively selects which synapses to update

22. * Li H, van Rossum MCW: **Energy efficient synaptic plasticity**. *Elife* 2020, **9**, e50804.

Energy saving model that trades off transient and persistent plasticity.

23. * Pache A, van Rossum MCW: *Lazy learning: a biologically-inspired plasticity rule for fast and energy efficient synaptic plasticity*. eprint arXiv; 2023. 2303.16067.

Energy saving model that skips over correct samples.

24. Agarwal PK, Har-Peled S, Varadarajan KR, *et al.*: **Geometric approximation via coresets**. *Combinatorial and computational geometry* 2005, **52**:1–30.

25. Paul M, Ganguli S, Dziugaite GK: *Deep learning on a data diet: finding important examples early in training*. 2021, 07075. arXiv: 2107.

26. van der Helden J, Boksem MAS, Blom JHG: **The importance of failure: feedback-related negativity predicts motor learning efficiency**. *Cerebr Cortex* 2009, **20**:1596–1603.

27. de Bruijn ER, Mars RB, Hester R: **Processing of performance errors predicts memory formation: enhanced feedback-related negativities for corrected versus repeated errors in an associative learning paradigm**. *Eur J Neurosci* 2020, **51**:881–890.

28. Guthrie JT: **Feedback and sentence learning**. *J Verb Learn Verb Behav* 1971, **10**:23–28. doi:10.1016/S0022-5371(71)80088-9.URL.

29. Cohen G, Afshar S, Tapson J, van Schaik A: *Emnist: an extension of mnist to handwritten letters*. 2017, https://doi.org/10.48550/ARXIV.1702.05373. URL, https://arxiv.org/abs/1702.05373.

30. Anderson MC, Hulbert JC: **Active forgetting: adaptation of memory by prefrontal control**. *Annu Rev Psychol* 2021, **72**: 1–36, https://doi.org/10.1146/annurev-psych-072720-094140.

31. Vanni S, Hokkanen H, Werner F, Angelucci A: **Anatomy and physiology of macaque visual cortical areas V1, V2, and V5/MT: bases for biologically realistic models**. *Cerebral cortex (New York, N.Y. : 1991* 2020, **30**:3483–3517, https://doi.org/10.1093/cercor/bhz322.

32. Yamins DL, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ: **Performance-optimized hierarchical models predict neural responses in higher visual cortex**. *Proc Natl Acad Sci USA* 2014, **111**:8619–8624.

33. Leimer P, Herzog M, Senn W: *Synaptic weight decay with selective consolidation enables fast learning without catastrophic forgetting*. BioRxiv; 2019, 613265.

34. Cichon J, Gan W-B: **Branch-specific dendritic ca-2+ spikes cause persistent synaptic plasticity**. *Nature* 2015, **520**: 180–185, https://doi.org/10.1038/nature14251.

35. Otor Y, Achvat S, Cermak N, Benisty H, Abboud M, Barak O, Schiller Y, Poleg-Polsky A, Schiller J: **Dynamic compartmental computations in tuft dendrites of layer 5 neurons during motor behavior**. *Science (New York, NY)* 2022, **376**:267–275, https://doi.org/10.1126/science.abn1421.

36. Rioult-Pedotti MS, Friedman D, Donoghue JP: **Learning-induced LTP in neocortex**. *Science* 2000, **290**:533–536.

37. Kramár EA, Babayan AH, Gavin CF, Cox CD, Jafari M, Gall CM, Rumbaugh G, Lynch G: **Synaptic evidence for the efficacy of spaced learning**. *Proc Natl Acad Sci USA* 2012, **109**: 5121–5126.

38. Sajikumar S, Morris RG, Korte M: **Competition between recently potentiated synaptic inputs reveals a winner-take-all phase of synaptic tagging and capture**. *Proc Natl Acad Sci USA* 2014, **111**:12217–12221.

39. Triesch J, Vo AD, Hafner A-S: **Competition for synaptic building blocks shapes synaptic plasticity**. *Elife* 2018, **7**, e37836, https://doi.org/10.7554/eLife.37836.

40. Jeong Y, Cho H-Y, Kim M, Oh J-P, Kang MS, Yoo M, Lee H-S, Han J-H: **Synaptic plasticity-dependent competition rule influences memory formation**. *Nat Commun* 2021, **12**:3915, https://doi.org/10.1038/s41467-021-24269-4.

41. Isabel G, Pascual A, Preat T: **Exclusive consolidated memory phases in drosophila**. *Science* 2004, **304**:1024–1027, https://doi.org/10.1126/science.1094932.

42. Abraham WC, Robins A: **Memory retention–the synaptic stability versus plasticity dilemma**. *Trends Neurosci* 2005, **28**: 73–78, https://doi.org/10.1016/j.tins.2004.12.003.

43. Dembitskaya Y, Piette C, Perez S, Berry H, Magistretti PJ, Venance L: **Lactate supply overtakes glucose when neural computational and cognitive loads scale up**. *Proc Natl Acad Sci USA* 2022, **119**, e2212004119, https://doi.org/10.1073/pnas.2212004119.

44. Potter WB, O'Riordan KJ, Barnett D, Osting SM, Wagoner M, Burger C, Roopra A: **Metabolic regulation of neuronal plasticity by the energy sensor AMPK**. *PLoS One* 2010, **5**, e8996.

45. Han S, Liu X, Mao H, Pu J, Pedram A, Horowitz MA, Dally WJ: **Eie: efficient inference engine on compressed deep neural network**. *Comput Architect News* 2016, **44**:243–254.

46. Tingley D, McClain K, Kaya E, Carpenter J, Buzsáki G: **A metabolic function of the hippocampal sharp wave-ripple**. *Nature* 2021, **597**:82–86, https://doi.org/10.1038/s41586-021-03811-w.