

Sean Commins*

Efficiency: an underlying principle of learning?

<https://doi.org/10.1515/revneuro-2017-0050>

Received July 6, 2017; accepted July 18, 2017; previously published online September 11, 2017

Abstract: Learning is essential. It allows animals to change circumstances, deal with new situations and adapt to environments. Here, we argue that learning, at behavioral and neural levels, involves efficiency, reflected in metabolic cost reductions. Behaviourally, although multiple solutions to a novel problem may be available, all solutions are not learnt – it is too costly. Furthermore, once a strategy has been selected, it is reinforced producing an efficiency that leads to a maximisation of performance and metabolic cost reductions. Learning can be represented in the brain through many mechanisms; however, if learning is truly efficient, then, all such mechanisms should also be accompanied by a reduction in measurable metabolic costs. By thinking about learning in terms of efficiency, not simply as a descriptive term but rather in terms of metabolic costs, it allows learning to be examined more carefully and provides predictions that can be easily tested (and indeed refuted).

Keywords: costs; efficiency; energy; hippocampus; learning; metabolism.

Introduction

‘Nature is thrifty in all its actions’

–Pierre Louis Maupertuis (1698–1759)

Across many disciplines, the idea that bodies try and reduce the effort made to achieve a goal has been well studied. In the 1600s, the famous French mathematician Pierre de Fermat postulated that light travels between two points along the shortest path leading to the principle of least time. In physics, of all the paths that a thrown ball can take, it seems to use the one that minimises the total energy required; this idea is known as the principle of least action. Similar ideas have been proposed in the behavioural sciences. For example, an interesting effect is observed in the area of linguistics whereby the most

common word in a book or that is spoken will occur twice as much as the second most common word, and three times more frequent than the third word, etc. In what has become known as Zipf’s, law the frequency of a word is inversely proportional to its rank (Zipf, 1949). Zipf suggested that word distribution is due to an attempt to communicate efficiently using the least amount of effort, that is, communicators do not want to work harder than is necessary for understanding to occur. Decision-making processes may follow the same rule. Allport (1954/1979) proposed that people do not have the time to form a separate opinion about all things suggesting that ‘Life is just too short to have differentiated concepts about everything. A few pathways are enough for us to walk in.’ (p. 173). As such, people evaluate information or form decisions that are thought to rely on simple strategies that are time efficient (Fiske and Taylor, 1984). Such decision-making process can lead to categorisation, generalisation, stereotypes and heuristics (Tversky and Kahneman, 1973; Macrae et al., 1994). The question arises whether a similar principle can be applied to the behaviour of learning and by extension the neural underpinnings of learning. This article proposes the idea of ‘learnt efficiency’.

Defining learnt efficiency, efficiency and cost

Despite being a major research topic across multiple disciplines (e.g. psychology, neuroscience, computer science and behavioural ecology), learning is difficult to define with many authors either avoiding the topic or providing definitions specific to the discipline (De Houwer et al., 2013; Barron et al., 2015). Notwithstanding the various difficulties, learning can be defined as a dynamic process leading to a change in behaviour and its underlying biology, brought about through interaction with the environment. The contention proposed here is that ‘if there are multiple ways to solve a problem, animals will not learn all possibilities but will select one that requires the least effort for the animal’. At the start of the learning process, multiple solutions may be available and tested, as often observed by high variability in individual responses during an early learning phase (e.g. see Meade et al., 2005). However, it is not cost effective to learn and

*Corresponding author: Sean Commins, Department of Psychology, Maynooth University, Co. Kildare, Ireland, e-mail: Sean.Commins@mu.ie

maintain all such solutions. Thus, as training progresses, an optimum strategy for that individual is selected and/or ineffective strategies are eliminated (Benson-Amram and Holekamp, 2012; Thornton and Samson, 2012; Chow et al., 2016), thereby, reducing the variability in response. The selected strategy is then repeated and reinforced, allowing behaviours to be performed quicker and easier, which, in turn, allows for an increased likelihood that such strategies are repeatedly chosen. This, it is argued here, leads to an optimisation in performance and the minimisation of effort or cost. Any changes to this strategy, particularly during the late phase of learning, will lead to additional effort and cost.

The term ‘efficiency’ is purposefully used; the term does not necessarily relate to the best choice, the most direct route, or even the best performance. For example, Meade et al. (2005) trained pigeons to find their way home. The initial flights of each bird are highly variable, but they become increasingly stereotyped by the end of training. Such flights are very individualised and often might not be the most direct or best route. Efficiency has been used as an explanatory term across many domains; however, most research, including our own (e.g. Farina et al., 2015), use the term merely as a description of an effect that has shown some sort of change of activity in one group compared to another that brings about a better performance. Time spent performing the behaviour and the number of errors are usually reduced, and this is *assumed* to translate into a reduced effort and, therefore, efficiency. On one level, this is fine, but as outlined by Poldrack (2015), it reveals nothing about the scientific mechanism underlying the effect or, indeed, provides testable predictions; direct measurements of efficiency or effort are seldom taken. Traditionally, efficiency is described by what we can get out of a system given a particular input, maximising performance or output while minimising the effort or input. However, there is not just a single way to describe efficiency. For example, in an efficient system, we might be expected to see an increase in output or performance while the level of input or effort remains constant. Alternatively, we might see the level of performance remaining the same or even increasing while the level of input simultaneously decreases (red arrows in Figure 1A). An inefficient system is one in which the output remains the same or decreases while the level of input increases. Further, we might see a decrease in output with a constant level of input (blue arrows, Figure 1A). An increase in output with increased effort or a decreased performance with a decreased input (black arrows, Figure 1A) is a system that shows neither efficiency nor inefficiency.

Expanding this idea further to reflect learning, Figure 1B–G shows a number of different scenarios whereby performance and cost, often associated with effort (see Christie and Schrater, 2015), could putatively change with learning across training. Figure 1B shows a situation in which the performance increases (black arrow) and the cost simultaneously decreases (red arrow). An alternative scenario depicted in Figure 1E shows that while performance increases, the cost remains the same. Note here that although cost levels remain constant throughout training, individual cost levels may differ [this is represented by the red line in Figure 1D that may be high (dashed), medium (solid), or low (dashed)]. Another situation represented in Figure 1F shows the cost decreasing but the performance remains the same. Again, the performance level may differ between individuals and represented by the location of the black line [high (dashed), medium (solid) or low (dashed)]. These three figures represent a system that is increasing in efficiency and are represented by an increase in the difference between the performance and cost lines (increasing black hashed area or decreasing red hashed area). Figure 1C, E, and G show the opposite – an increasing inefficient system (increasing hashed red area or decreasing hashed black region).

So far, cost has been mentioned, but one of the major difficulties is how to define and measure cost. For example, time may have a real cost for animals in a natural context; spending extra time trying to solve something may result in the increased risk in predation and death (Chow et al., 2016). Cost, therefore, can refer to a number of variables but should be somewhat related to energy (Poldrack, 2015). However, energetic cost is also a very general term and may refer to mechanical energy, neural effort, cognitive effort, muscle activity and metabolic cost (Huang et al., 2012). Metabolic cost is considered the best measure of energy consumption and can be done by measuring heat loss, referred to as direct calorimetry, or through oxygen consumption or carbon dioxide production (Leonard, 2010). While other measurements such as increased blood pressure and heart rate may also be used, these measures are indirect and often do not correlate well with the more direct means (Westbrook and Braver, 2015). Using this concept of efficiency, rather than as purely descriptive term, and taking cost as to represent metabolic costs (that can be measured) allows learning to be examined more carefully, perhaps provide a better explanation and allow better predictions to be made. Table 1 provides an overview of some predictions that might be made using these concepts, with sample research evidence. The following sections will discuss these ideas in further detail.

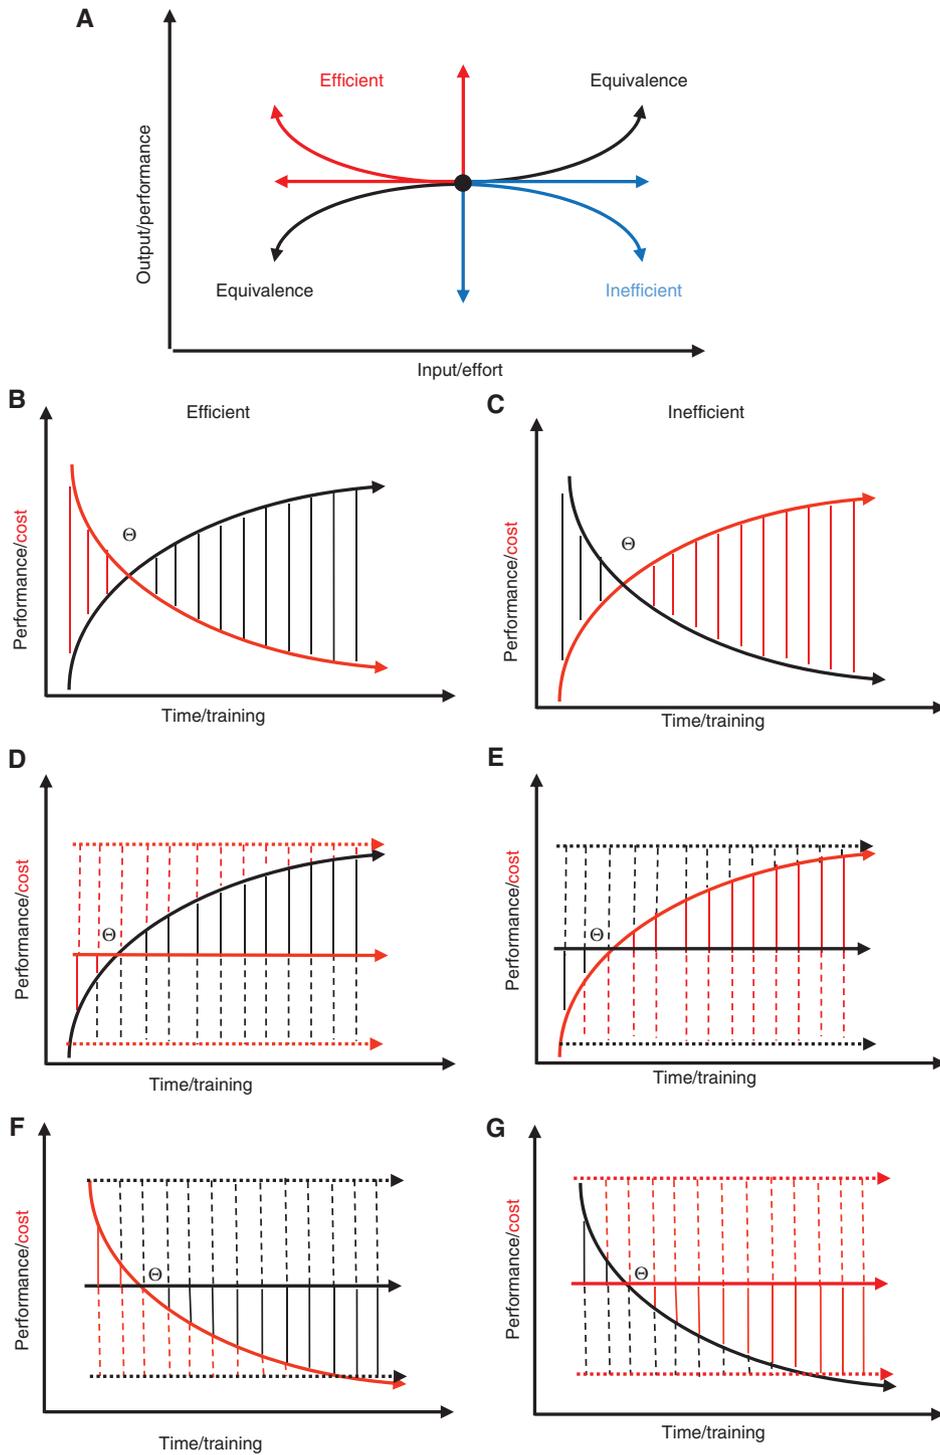


Figure 1: Changing dynamics of efficiency that is dependent on the interaction between performance (output) and effort (input). (A) A plot demonstrating how efficiency (red) or inefficiency (blue) can emerge depending on the changing dynamics between the inputs and outputs to a particular system. (B–G) A variety of putative situations whereby efficiency can emerge with learning. The simplest situation is where the performance increases as the costs decrease (B); alternatively, either performance may increase despite a constant level of costs (D), or costs may decrease with a constant level of performance (F). Both situations show an increase in efficiency (represented by increasing solid straight black lines), but this depends on the level of cost and performance. A greater efficiency (greater black hashed area) is demonstrated with high performance (upper dashed black line) compared to a medium performance (solid black line). Greater efficiency is also observed with low cost levels (lower red dashed line) compared to medium levels (solid red line). Note in this model the point of interaction, Θ , where the performance and cost lines cross is not fixed and may scale upwards or downwards and is very different across different individuals. Plots C, E, G represent the opposite effect, whereby there is an increase in inefficiency (represented by increasing solid straight red lines) with time.

Table 1: General and specific predictions related to efficiency during learning.

General predictions	Specific predictions	Evidence?
As performance on a task maximises, there is a minimisation of metabolic costs	As one becomes proficient at a task, less energy, in the form of oxygen/glucose consumption, heat loss or ATP, CO ₂ production is required to complete it	Studies such as those of Huang et al. (2012) and Lay et al. (2002) have shown decreases in metabolic power with motor learning. See also the 'Expressing behavioral efficiency during learning' and 'Expressing neural efficiency during learning' sections
Different stages of learning have different metabolic costs	Early-stage learning would incur higher metabolic costs compared to late-stage learning (see, e.g. Figure 1B) Additional learning, particularly during late-stage learning, would incur extra metabolic costs	Studies such as that McNay et al. (2000) have found large metabolic changes during early learning, but these were attenuated late in training
Efficiency should be observed at different levels of the organism	Organisms operate as a whole. Animals use behavioural mechanisms that are efficient and should show reduced metabolic costs. Such learning mechanisms are reflected at systems, neural and molecular levels, which should also show reduced metabolic costs	See the sections on 'Expressing behavioral efficiency during learning' and 'Expressing neural efficiency during learning' for details and predictions regarding behavioural and neural efficiency, respectively Studies like that of Picard et al. (2013) hint at multiple level efficiency
Efficiency should be reflected across learning irrespective of task type or time span	Animals are capable of various types of learning (spatial, verbal, motor, etc.) and can learn across multiple trials or with a single trial, with instruction or without. For efficiency to hold true, it should be reflected across all	Reduced metabolic costs have been shown across different types of learning, but there may be differences depending on trial type, task familiarity and whether changes are distributed or reorganised in the brain (see Kelly and Garavan, 2005)
Efficiency may depend on an individual's basal cost and performance level at the time of learning	IQ, age, sex, physiology, health status, etc., may all affect efficiency to different degrees	Individuals that are more intelligent have a greater efficiency in brain use after training compared to those with less intelligence (Neubauer et al., 2004; Neubauer and Fink, 2009) Older adults that show good memory performance often have extra neural activation compared to younger adults (Morcom et al., 2007) Although these studies refer to 'efficiency/inefficiency', metabolic markers were not examined
Efficiency levels are not fixed and may be modifiable	Metabolic costs/performance may be modified by physical exercise, diet, mental activity, sleep, etc	Erickson et al. (2011) found that exercise and hippocampal volumes were correlated with strong performance in memory tasks of older adults Duzel et al. (2016) suggest that running-induced hippocampal neurogenesis is associated with increased synaptic plasticity and improved spatial memory in animals Metabolic markers were not examined in these studies

Expressing behavioural efficiency during learning

Despite the difficulties associated with defining and measuring cost, there is evidence to suggest that repetition of a behaviour, including learning, does minimise energetic costs. Take locomotion as an example: humans walk with a step width that corresponds to the minimum metabolic cost (Arellano and Kram, 2011); likewise, swinging arms while walking seems to reduce metabolic cost compared to walking without arms swinging (Ortega et al., 2008). Repeated practice also reduces metabolic costs. Lay et al. (2002), for example, showed a decline in

metabolic energy cost with practice on a rowing task. As summarised by Sparrow and Newell (1994) if mechanical power output is held constant, then, individuals who are well-practised on a task will have a reduced heart rate and metabolic energy expenditure. Similarly, in an arm-reaching task, Huang et al. (2012) demonstrated that net metabolic power, measured directly through expired gas analysis, decreased with motor learning. Interestingly, such metabolic costs continued to decrease late in learning even when movements were fine tuned and had stabilised. Indeed, the authors observed greater cost reductions when muscle activity had reached asymptotic levels. This effect has also been observed by others (e.g. Galna and Sparrow, 2006).

Detailed examination of animal learning also reveals many examples of efficiency, at least at a behavioural level. If, as suggested, animals do not learn all possible solutions but select one to minimise effort, then, a number of behaviours should be observed including (1) the use of multiple strategies to solve a problem will be generally limited, (2) as animals want to avoid unnecessary effort, some associations during learning are easier made than others, (3) once learning has been achieved and particular associations are made, it will be difficult to make further changes, as this would require more effort, and (4) the chosen strategy or solution may not necessarily be the best one. Some of these features are well recognised in the learning literature. For example, blocking (Kamin, 1969) is a well-described phenomenon from the classical conditioning literature. In the first phase of a typical blocking paradigm, an animal is exposed to the pairing of a tone and shock. In the second phase of the experiment, light and tone are paired with a shock. Upon testing, it has been shown that animals have learnt little about the light. Therefore, learning the association between the tone and shock blocks any subsequent learning about the light. This phenomenon suggests first that animals do not learn multiple strategies readily, although in this case, both the light and the tone are predictors of a shock, just the tone-shock association is learnt; second, once a strategy has been adopted or an association has been made, i.e. the tone predicts a shock, it is difficult to change it and even update that particular association. Further examples of this can be illustrated by another behavioural effect termed latent learning (Reiss and Wagner, 1972). If a tone is initially not paired with a shock, but is then subsequently paired with one, then, this association is slower to learn than if an animal had been initially exposed to the tone and shock pairing. Therefore, once something has been learnt, in this case the tone does not predict a shock, it is more difficult to form a new association. It is predicted that any new associations are not only more difficult to form but would also be accompanied by an increase in metabolic costs. Such experiments have not been done to date.

Another relevant concept in behavioural psychology is overshadowing (Pavlov, 1927; Mackintosh, 1974). When two or more stimuli are presented together, all will not produce an equal level of response. One stimulus, depending on its salience, novelty, ecological relevancy, or other factors, may provide a stronger or more reliable response compared to the other stimuli. This more salient stimulus is said to overshadow the others. Therefore, although all stimuli may provide equal predictability, not all stimuli are learnt; stimuli are selectively associated

– some stimuli are more easily learnt than others. Animals in the natural environment use this very effectively. For example, nonvenomous scarlet kingsnakes (*Lampropeltis elapsoides*) look very similar to the poisonous coral snake (*Micrurus fulvius*): in terms of colour, both have alternating red, yellow and black rings; however, the sequence of coloured rings differs between the animals. Despite this, the kingsnakes are avoided by predatory birds as much as coral snakes (Kikuchi and Pfennig, 2013) suggesting that the predators have not learnt all stimuli equally but have selected one (in this case colour) and have learnt to avoid it. The predators have learnt efficiently, and the mimicry observed in some potential prey exploits this fully. A recent study conducted by Kazemi et al. (2014) examined this overshadowing effect more formally in a laboratory setting. The authors trained blue tits to find worms hidden under a paper. One particular paper that had a combination of a particular colour, shape and pattern did not contain a reward and so had to be avoided. Birds were then tested with prey that contained just one of these stimuli and the combined stimulus. In theory, birds could have learnt each of the stimuli and avoided all test conditions, or they could have avoided just the combined stimulus. The authors found that the birds avoided the combined stimulus and the ones that had the particular colour. Therefore, colour overshadowed the other attributes of pattern and shape. Importantly, all attributes were not learnt (see also Farina et al., 2015). While colour may be more relevant to birds, this may not be the most important/salient attribute to other animals. Learning efficiency is specific to each species as well as the individual. To emphasise this, biologically and ecologically relevant stimuli are easier to associate and learn, whereas other stimuli may incur more costs. For example, pigeons receiving food that is paired with light tend to respond more than when food is paired to a tone. In contrast, rats respond better when food is paired to a tone. In addition, taste and poison are associated easily than taste and shock (Krane and Wagner, 1975). Physical similarity in stimuli can also lead to quicker associations; for example, different colours paired together or different line orientations when paired are easier to learn than colour-line orientation associations (see Rescorla and Gillan, 1980 for similar examples). Such overshadowing examples would be predicted to involve a reduction in metabolic costs for the animal.

When an animal is rewarded in the presence of a particular stimulus, it is likely to respond in the presence of a similar but recognisably different stimulus. This effect is termed stimulus generalisation (Ghirlanda and Enquist, 2003). Animals essentially ‘bet’ on the results being the

same as before. This ability to generalise has been shown across a range of stimuli (including light wavelength, Hanson, 1959; tone frequency, Brennan and Riccio, 1972; object orientation, Bloomfield, 1967; object location, Cheng et al., 1997) and across different species (pigeons, Cheng et al., 1997; honeybees, Cheng, 2002; rats, Brennan and Riccio, 1972; fish, Fay, 1970 and humans, Nicholson and Gray, 1971). The effect is so widespread that it has been suggested to constitute a universal law for psychology (see Shepard, 1987). It is argued here that rather than being a law, itself; stimulus generalisation is a product of learnt efficiency. What is learnt in one situation can be transferred to another; therefore, the need to learn about every possible situation that may arise is avoided. Time, effort and metabolic costs should be reduced. While the ability to generalise may be a useful way for dealing with novelty, it may come at a long-term price, for example, it is only by testing different responses and being flexible that animals can adapt to a new situation (Ghirlanda and Enquist, 2003). Further difficulties arise in the case of fear generalisation, which may have debilitating consequences (see Dymond et al., 2015).

For learning efficiency, to hold all such behaviours described above should be associated with a minimisation of cost in terms of energy. Measuring the amount of oxygen consumed to reach asymptotic performance may be applied across these various conditions. For example, if animals do not maintain multiple solutions as it is not very cost effective, then, this should be observed and measurable. Where animals must use and maintain multiple solutions, we should see evidence of increased costs compared to those that do not require to do so. Likewise, we would predict that training with a salient, ecologically-, biologically-, or other relevant cues would not only aid strategy selection and quicken learning but would be associated with lower metabolic costs compared to those trained with non-salient stimuli. If relearning or adding extra information requires more effort, this should also come with a cost that is measurable. Similarly, an animal that generalises should show reduced costs compared to an animal that displays greater flexibility. Very few studies have examined oxygen or glucose changes that accompany such behavioural phenomena.

Expressing neural efficiency during learning

While selection of a strategy and fine tuning of behaviour may underlie metabolic reduction, decreasing costs

should also be reflected at the neural level (Huang et al., 2012). The brain is one of the most energy-expensive organs in the body, ranked third behind skeletal muscle and the liver. The energy cost of the vertebrate brain is between 2% and 10% of the total body metabolic cost (Mink et al., 1981), with the human brain costing even more in energy terms coming in at 20%; this is despite being only 2% of the body mass (Clarke and Sokoloff, 1999). Although the energy requirements of a neuron vary across different parts of a neuron and the type of neuron (Sengupta et al., 2010; Hallermann et al., 2012), it is estimated that action potentials cost between 21% and 47%, synaptic processes 34–59% and resting potential 13–20% of the total energy requirements (Attwell and Laughlin, 2001; Poldrack, 2015). Neural activity is particularly expensive with Hasenstaub et al. (2010) suggesting that energy minimisation may be a unifying principle underlying general neural dynamics. Furthermore, Herculano-Houzel (2011) suggests that the energy budget per neuron, in contrast to other cells, has been stretched in evolution to remain constant and that it may be close to its limit. As such, neural processes, including those involving learning, must live within this budget, be able to adjust accordingly so that excessive synaptic activity is avoided and, as is argued here, energy costs are minimised. Energy minimisation is suggested, here, to be also a principle underlying neural learning dynamics.

According to our definition, learning is also associated with changes in underlying biological, including neural, mechanisms (see also West-Eberhard, 2003). The brain has developed multiple mechanisms that have been associated with learning. The most prominent is the idea that learning, across a range of domains (associative, spatial, motor, etc.) adjusts the synaptic strength of connections, which, in turn, can lead to longer-lasting structural changes. Hebb (1949) suggested that ‘the persistence or repetition of a reverberatory activity (or ‘trace’) tends to induce lasting cellular changes that add to its stability... When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased’. One phenomenon that fulfils this idea is long-term potentiation (LTP). First discovered by Bliss and Lomo in 1973, LTP has long been considered as a putative mechanism that underlies learning and memory (Takeuchi et al., 2013). LTP is the stable and enduring enhancement of synaptic strength following high-frequency neural stimulation (HFS). Blocking LTP in specific regions leads to impaired learning. For example, blocking hippocampal glutamate NMDA receptors impair LTP

and produce learning deficits in hippocampal-dependent tasks (such as spatial learning) (Morris et al., 1986). Conversely, enhancement of the function of the NMDA receptor leads to enhanced synaptic potentiation and a concurrent improvement on memory tests (Tang et al., 1999). However, direct evidence showing synaptic enhancement, in the form of LTP, being actually induced by learning has only been demonstrated recently (see Whitlock et al., 2006). The key feature of LTP is that it enhances the efficiency of the neural connections. Neurons post-HFS have an increased response to the same stimulation compared the response pre-HFS; neurons now show an increased output for the same level of input (see Figure 1A). Other forms of synaptic enhancement may bring about similar efficiency in the system such as paired-pulse facilitation (Commins et al., 1998; Craig and Commins, 2005; Fernández-Fernández et al., 2015), metaplasticity (Li et al., 2014), or other mechanisms underlying fast learning (Yger et al., 2015). Other forms of synaptic plasticity such as long-term depression (LTD) or paired-pulse depression have been also associated with learning and memory (Brigman et al., 2010). However, these phenomena result in a decreased neural response compared to baseline using the same stimulation. These processes suggest that there is less output for a similar level of input. Although this form of plasticity may seem inefficient (and come at a cost), it also seems to play a crucial role in learning, particularly, in the ability to readjust the range in which synapses can operate. LTD, for example, has been shown to inactivate memories (Nabavi et al., 2014), weaken previously potentiated synapses allowing for greater flexibility (Nicholls et al., 2008) and suppress interference from competing information (Connor and Wang, 2015). Such mechanisms may play an important role during the early phase of learning when specific strategies are selected and refined. Therefore, irrespective of the type of plasticity, it is argued here that all forms of synaptic plasticity associated with learning should also be accompanied by a measurable change in metabolic cost, echoing Hebb's sentiment. Very few studies have examined oxygen/glucose changes that are associated with LTP or other plastic effects.

Many reports have shown that the adult brain undergoes constant remodelling, and this is necessary for learning and memory (Grillo et al., 2015). For example, motor learning correlates with the increased number and stabilisation of postsynaptic spines on excitatory neurons (Yang et al., 2009). Furthermore, a recent study by Chen et al. (2015) shows that motor learning is associated with structural rearrangement of presynaptic boutons in different classes of inhibitory interneurons. In this study, there is a counterbalancing of excitation with increased

pruning and inhibition, pointing to a drive toward synaptic homeostasis with learning (Turrigiano, 2008). As well as restructuring at the neural level, learning is also associated with re-organisation at the systems level. Maguire et al. (2000) have shown that taxi drivers that are required to learn thousands of streets in London have a larger hippocampus compared to non-taxi drivers. Furthermore, the longer the time spent as a taxi driver the larger the hippocampus becomes (Maguire et al., 2000). Such changes in organisation are seen with other learning tasks including motor learning (Pascual-Leone et al., 1995, 2005) as well as in other species (e.g. see Krebs et al., 1989). Some reports have suggested that such reorganisation is associated with neural efficiency. For example, Dehaene et al. (2015) suggest that learning to read (literacy) affects neural organisation, and this is done by both 'creating and automating an efficient interface' between language and vision (p. 241). This may be achieved by the physical thickening of connections of regions such as the splenium and corpus callosum (Carreiras et al., 2009) and other regions including the visual word form area (VWFA) (Thiebaut de Schotten et al., 2014). Although such studies use the term efficient as a descriptive term and suggest a mechanism by which it may be achieved, they do not show directly whether such re-organisational changes (at cellular or system levels) are also accompanied by metabolic changes. It is predicted here that they should.

One method to examine metabolic changes that may accompany learning is via analysis of imaging studies. Typically, brain imaging using fMRI and PET, measuring blood oxygen and glucose, respectively, has examined how blood flow increases to different areas of the brain during a task. However, imaging is not that straightforward. Many researchers have shown task-related decreases in regional blood flow (e.g. Amedi et al., 2005). While others have shown a reorganisation or redistribution of functional networks after learning (see Kelly and Garavan, 2005, for details). Furthermore, neural activity can elicit changes in both metabolic and vascular dynamics, whereby initial oxygen is consumed and later replenished via cerebral blood flow (Li et al., 2011). Regional blood flow seems to increase more than oxygen consumption suggesting that the supply of oxygen exceeds its demand (Fox et al., 1988). More importantly, it has been shown that metabolic demands associated with a particular task seem to be very low (Kurzban, 2010), with Raichle and Mintun (2006) suggesting that glucose consumption related to a task could be as low as 1% and rarely exceeding 10% of resting blood flow. Others have reported a mere 8–12% increase in adenosine triphosphate (ATP) production during visual stimulation (Lin et al., 2010).

Such figures have prompted the suggestion that although the brain is indeed overall costly, most of the cost is associated with intrinsic resting state activities (Raichle and Mintun 2006) rather than external demands. However, using more direct *in vivo* microdialysis McNay et al. (2000) found a 32% decrease in hippocampal glucose concentrations as animals learnt a complex maze task (spontaneous alternation test). With a less demanding task, the authors still found an 11% decrease in glucose concentration. Furthermore, systemic injections of glucose reversed this decreased concentration and increased performance on the more complex task. The decreases in glucose concentrations were interpreted as being related to the demands of the task and the increased take-up of glucose into neurons or glia. In a more recent study, Francois et al. (2014) used *in vivo* amperometric techniques to monitor oxygen concentrations in the nucleus accumbens as animals learnt a task requiring the discrimination between a rewarded and a nonrewarded cue. Oxygen signals in the nucleus accumbens increased during acquisition, specifically to the rewarded cue and also the first days of extinction. In this study and also in McNay et al. (2000), it is interesting to note that any metabolic changes observed were attenuated as training progressed. This perhaps gives credence to the suggestion that different learning stages have different metabolic costs whereby initial engagement with a task may be costly in metabolic terms but that this decreases as learning progresses, hinting at efficiency in the system. More direct evidence comes from Picard et al. (2013) who trained monkeys in a motor task. The authors monitored single unit recordings and 2-deoxyglucose (2DG) simultaneously in the primary motor area. The authors found no change in neural firing but a decrease in metabolic (2DG) activity with extended training. The authors speculate that such efficiency, less metabolism required for a given amount of neural activity, may be due to a number of factors including changes in synaptic plasticity, restructuring of synapses, or increased synchronisation of synaptic activity.

Another mechanism linked to the learning process is neurogenesis. Adult neurogenesis occurs throughout the lifespan but is mainly confined to two regions of the brain: the subventricular zone of the lateral ventricles and the subgranular zone of the dentate gyrus (Zhao et al., 2008). The dentate gyrus is an area of particular importance as it is part of the hippocampus, a structure critically involved in learning and memory (Morris et al., 1986). There is much evidence demonstrating that factors that enhance or impair hippocampal neurogenesis can affect cognition and learning. Mutant mice (e.g. *mbdt^{-/-}*) with decreased dentate neurogenesis show impaired learning

and memory using the water maze task and impaired LTP (Zhao et al., 2003). Similarly, using protocols such as exercise or enriched environment that increase cell proliferation also improves learning and memory (Kempermann et al., 1997; van Praag et al., 1999). However, whether neurogenesis is directly involved in learning is controversial, particularly given that new neurons are confined to a small region of the hippocampus and the limitation with some of the ablation techniques used (Zhao et al., 2008). There are a limited number of studies that have examined correlations between hippocampal neurogenesis and metabolic markers. One study, conducted in non-human primates, has shown a positive correlation between hippocampal neurogenesis and blood glucose levels, suggesting that increased neurogenesis is associated with increased metabolism (Perera et al., 2011). Given this and the idea that each additional neuron has its own energy budget that would incur extra costs (Herculano-Houzel, 2011), neurogenesis may not be the primary or most suitable mechanism underlying learning; rather, changing the synaptic weightings and/or structural rearrangement of existing neurons may be a more cost-effective method. This may explain why neurogenesis is not observed more widely throughout the brain.

Finally, although there are many computational models of learning and memory, most of these focus on information capacity and not on energy efficiency (Sacramento et al., 2015). Many models would suggest that an effective method for reducing the high cost of neural activity is through sparse distributed networks (Hu et al., 2012). This can be achieved by keeping the number of active synapses low, whereby only a few neurons are active at a time. For example, Attwell and Laughlin (2001) suggest that only 15% or less of a group of neurons (with action potentials of 4 Hz) need to be active to encode a particular condition. Interestingly, many specific cells that have been discovered throughout the brain only represent a small percentage of the total neurons of the particular region. For example, in area CA1 of the hippocampus, the majority of cells are silent with approximately 20–30% of cells being cells that represent a specific location (place cells, see Nakazawa et al., 2004); similarly approximately 25% of neurons in area F5 of the monkey are considered ‘mirror neurons’ (Rizzolatti et al., 2014), and there are 20% of the cells in areas TE of the inferotemporal cortex that are considered ‘face cells’ (Gauthier and Logothetis, 2000). Perhaps, such sparse representations for certain specific cells may form the neural equivalent of the Pareto principle and follow a power law distribution, whereby 80% of the work is performed by 20% of neurons (see Reed, 2001). An alternative method to achieve sparseness

and minimise costs during learning is to remove synapses through pruning (Mimura et al., 2003) or maximising silent synapses (see also Sacramento et al., 2015). Despite such models, direct analysis of metabolic costs associated with sparse neural networks is limited and difficult to achieve. In sum, although the brain has developed multiple ways for learning including synaptic plasticity, neurogenesis, neural remodelling and system reorganisation, it is argued here that in order for the idea of learnt efficiency to hold, all of these ways should be accompanied by a reduction in metabolic costs.

Understanding how neural efficiency, in whatever form it takes, and behavioural efficiency are directly linked is difficult, especially given that many tasks are, themselves, complicated and require multiple regions, which may, after learning, become reorganised. How can metabolic costs be measured across such dynamic networks? One possibility for human research is through connectivity studies (see Kelly and Garavan, 2005); another is by conducting studies along the lines of Picard et al. (2013) that have targeted a particular behaviour and area, directly linking neural activity, learning and metabolic costs. One structure that might provide a good starting point for this type of analysis is the hippocampus. This structure has been anatomically very well characterised; much is known regarding the various cell types, cell distribution, pattern of connectivity and neurochemistry make-up (Szi­lágyi et al., 2011). As well as its established role in learning and memory (Scoville and Milner, 1957; Morris et al., 1986; Harvey et al., 2008; Diviney et al., 2013), the hippocampus is known to undergo various forms of synaptic plasticity (Bliss and Lomo, 1973; Commins et al., 1998), show neurogenesis (Kempermann et al., 1997), and is also very susceptible to stress (Lupien and Lepage, 2001), age, and many disorders including Alzheimer's disease (Braak and Del Tredici, 2015). Furthermore, neurons in the hippocampus are thought to have the highest energy requirements compared to neurons elsewhere (Stranahan and Mattson, 2008). Examination of the hippocampus may very well reveal the structure as a possible 'mediator of efficiency'.

A possible mechanism for learning and learnt efficiency

Figure 2 attempts to summarise how learning may occur and incorporate some of the ideas that have been discussed so far. At the start of the learning process, (a) there may be many solutions to a particular problem, and animals over the initial few trials may use a variety of

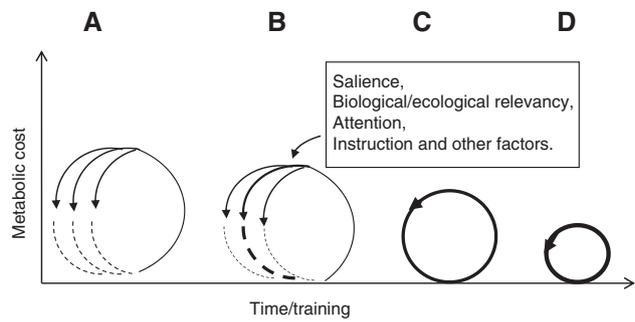


Figure 2: Proposed simple model of learning and efficiency.

(A) Initial learning is often associated with testing of different solutions (multiple arrows) which has a high cost. All solutions can not be maintained. (B) A particular strategy is chosen (dark circular line) aided by a variety of factors. (C) The selected strategy is reinforced and repeated (darker circular lines); metabolic costs are lowered. (D) Selected strategy continues to be reinforced leading to a maximisation of performance and a lowering of metabolic costs.

different strategies, as represented by the multiple arrows. The costs associated with having and trying different solutions are high; all solutions cannot be maintained over the long term. With time, multiple trials and repetition (represented by circular arrows), a particular strategy is selected (b). The behavioural selection may be quickened and aided by a variety of factors relevant to the environment including the saliency, relevancy, stability of cues, as well as the circumstances of the animal, and other factors. At this point, (c) ineffective strategies are eliminated, and the selected strategy is repeated and reinforced, as represented by the darkened single line. Strategy selection and reinforcement, as well as strategy elimination may result from a form of operant conditioning (see also Chow et al., 2016). For example, if a particular strategy is associated with reward or success, then the selected behaviour may be repeated and reinforced, while unrewarded strategies are eliminated, a form of behavioural competition. Synaptic plasticity (LTP and LTD, see above) and neural competition may also be critical at this stage. Competition between neurons is necessary during development for refining neural circuits (Changeux and Danchin, 1976) and shaping behaviour. This is observed particularly in the early development of the primary visual pathway; axons become segregated into eye-specific regions in the visual cortex as a result of activity-dependent synaptic competition between axons representing the two eyes (Cabelli et al., 1995). Higher-order cognition is also subject to competition during development. There is evidence for competition between reading and face recognition in the ventral visual cortex with regions devoted to face responses becoming gradually displaced with increasing literacy (Dehaene et al., 2010; Cantlon et al., 2011).

In the adult, brain competition, too, has recently been proposed to underlie the plasticity associated with learning and memory. For example, although 70% of neurons in the lateral amygdala receive sensory input, only about 25% of them exhibit plasticity related to fear conditioning (Rumpel et al., 2005; Han et al., 2007; see also Rashid et al., 2016). This suggests that certain neurons are preferentially selected over others, using the minimum number to maximise performance. Furthermore, Han et al. (2007) have provided evidence that CREB is the key molecular component that underlies this competition and the selection of certain neurons to participate in the memory trace. Synaptic plasticity and re-organisation of synapses and regions, as mentioned previously, may also play important roles in both the selection and reinforcement of a strategy, as well as the elimination of other ineffective ones, all of which leads to a maximisation of performance levels and importantly the reduction in metabolic costs (d).

The downside of efficiency

Learning is generally seen as being beneficial. It allows an animal to modify its behaviour so that it can better adapt to its environment. However, learning may come at a price (see Dukas, 1999; Burns et al., 2011 for a detailed discussion). Although efficiency has the benefit of reducing metabolic costs, it, too, may have a downside. For example, we have recently shown (Farina et al., 2015) that when animals are trained on a spatial task in the presence of two cues (of unequal salience in terms of brightness) and are presented upon re-test with just the brighter one, they can find the goal easily. However, if presented with the dimmer cue, animals are impaired. This selectivity, while efficient at least in behavioural terms – having to learn one cue rather than two, reduces the animals' ability to be flexible should the situation change. Training squirrels on a problem-solving task, Chow et al. (2016) also found that there is a trade-off between behavioural selectivity and flexibility. Those animals that displayed high levels of flexibility also engaged more regularly in ineffective behaviours. Another issue is that efficiency in learning may also lead to inaccuracies and misattribution errors. In a recent study, Hunt and Chittka (2015) found that bumblebees can create 'false memories'. Bees were initially trained to find a reward with a yellow flower stimulus followed by training to find a reward with a black and white-patterned stimulus. When bees were tested on the following day, they opted to go for the most recent stimulus, i.e. the black and white one. However, if tested 3 days afterward, the bees went for a stimulus-containing yellow and white stripes,

a novel stimulus, which was a hybrid of two previous ones. Interestingly, this effect only emerged after multiple testing trials. This finding suggests that bees found it difficult to keep track of multiple memories and were unable to store an accurate memory of each flower visited, thereby, leading to a mixed-up false memory. Bees did not learn all the features of the flower but extracted the general features. Recall that all of these general features may have been then applied to other flowers inappropriately (Hunt and Chittka, 2015). Humans, too, create 'false memories', with eyewitness testimony being a classic example of this (Loftus and Palmer, 1974). We have argued that this inability to learn about every situation encountered may lead to a generalisation effect, an arguably efficient behavioural strategy, but a strategy that could also prove debilitating. Fear generalisation, for example, occurs when a fear response to certain stimuli spreads to other related stimuli. Such responses may result in fear, avoidance and anxiety in many individuals, often leading to hospitalisation (Dymond et al., 2015). As a consequence of inaccuracies in learning and recall, poor decisions are often made; see Kahneman (2011) for a more detailed discussion.

Comparisons with other models

The ideas presented here have parallels with other theories, particularly those dealing with cognitive effort. Many cognitive processes such as attention, decision making and working memory are often thought to be costly (Kool et al., 2010), as such humans try to avoid such mental expense and fatigue by seeking the least cognitively expensive method (Solomon, 1948; Balle, 2002). Many 'cost/benefit models' (Krebs et al., 2010) suggest that cognitive demand is generally avoided in favour of more rewarding approaches, and this switch is facilitated through the subjective feeling of mental fatigue. The difficulty with many of these models is that they do not speak to any underlying biological processes (Christie and Schrater, 2015). Other models that fall under the category of 'limited resources models' (see Gailliot et al., 2007) suggest that mental effort is related to the available levels of blood glucose. When glucose levels are diminished, cognitive tasks are more difficult to perform, resulting in mental fatigue. More recently, Christie and Schrater (2015) provided a very nice detailed account of cognitive effort using an optimum control model that captures elements from both. The authors suggest that costly cognitive processes are avoided due to the allocation of energy resources. Depending on the cost/benefit analysis, only limited recourses may be committed to the current task

in order to conserve energy, in the form of glycogen, and to be further allocated at some future date. The brain, therefore, makes a strategic decision; resources are only committed when the payoff is worthwhile. While in this article the focus has mainly been on cost changes during a dynamic learning process and neural expression, the underlying metabolic mechanisms as outlined by Christie and Schrater (2015) may be very applicable. Further examination along these lines is required.

Other very influential models that speak of efficiency include the neural efficiency model of intelligence by Haier et al. (1992). In this model, it is suggested that smarter individuals display lower, and as a result, more efficient neural activation when engaged in various cognitive tasks. While there is generally good evidence for this idea, recent research has moderated the universality of the model; sex, task demand, task type and the brain region being taxed, may all play a role (Neubauer and Fink, 2009). On the other hand, there have been recent proposals of aging inefficiency (Rypma and D'Esposito, 2000; Morcom et al., 2007; Nyberg et al., 2014). Memory decline and impaired task performance is observed in many older adults, particularly with regard to episodic memory and recall of details of a specific event. Functional imaging analysis has shown a different pattern of neural activity in older adults that have good memory compared to young adults (Davis et al., 2008). Despite having similar retention scores and engaging a similar brain network during a particular task, older adults display extra neural activation of this network. Furthermore, older adults show fewer retrieval-related activity reductions compared to younger adults. Such age-related differences, as suggested by Morcom et al. (2007), are related to inefficiency – extra neural activation for a similar performance level. Both aging and intelligence efficiency models may be united under the general learning efficiency concept proposed here. As suggested in Figure 1B–G and Table 1, performance and cost levels are not fixed and may vary from individual to individual across a lifespan.

The emphasis throughout this manuscript has been on efficiency, but are we and our brains simply lazy learners? We can, of course, think deeply, learn multiple strategies, learn extra information that can be incorporated into our existing knowledge, but the argument here is that these actions come with a greater metabolic cost and, as such, is not our immediate response. Such purposeful actions are similar to those included in many of the dual-processing models, particularly those that concern decision-making processes, where comparisons are made between a fast, intuitive and affect-dependent system, and another that is more analytical, slow and logical (Evans, 2007; Mukherjee,

2010; Kahneman, 2011). Evidence from the animal literature has shown that many animals can acquire extra information and learn multiple strategies. For example, Kamil and Cheng (2001) proposed a multiple heading hypothesis that suggests that the more landmarks that are available, the more accurate Clark's nutcrackers tend to be in finding a goal. They suggest that all landmarks available are encoded. Although the authors themselves questioned this in terms of redundancy, they counter argued that having accuracy is vital for the birds' survival and also prevents change in landmark availability. But is this strategy typical? Indeed, a recent study by the same authors (Kelly et al., 2010) concede that perhaps all cues are not treated equally with some being relied on more than others, depending on the salience of the landmark. Furthermore, Kealy et al. (2008) attempted to examine whether animals hold a hierarchy of strategies, an idea that suggests that once one strategy has been compromised, another can take its place. The authors found no evidence for this and reported that animals were simply impaired once the learnt strategy was removed. These and other studies show that under certain circumstances, multiple strategies and extra information can be acquired, but this may not be the default position, and it is argued here that it would incur extra metabolic costs.

Conclusion

The article attempts to present the idea that learning tries to be efficient. From behavioural strategies such as selectivity and use of certain strategies, the general ignoring of multiple stimuli in favour of one or two prominent ones, to the generalisation of information to different situations, all try to minimise effort for the animal. However, for these behaviours and others to be truly efficient, all must be accompanied by a reduction in metabolic costs that can be measured. Furthermore, while learning at the neural level may also be represented by a variety of mechanisms, such as cellular and system reorganisation, changes in synaptic strength, neurogenesis and the promotion of sparseness, as well as bringing about an increased accuracy or quicker performance, these mechanisms should also be accompanied by reduced metabolic costs for learning to be truly efficient.

Acknowledgments: The author would like to thank Dr. Daragh McLoughlin for comments on an earlier draft of the manuscript.

References

- Allport, G.W. (1954/1979). *The Nature of Prejudice* (Cambridge, MA, USA: Perseus Books).
- Amedi, A., Malach, R., and Pascual-Leone, A. (2005). Negative BOLD differentiates visual imagery and perception. *Neuron* 48, 859–872.
- Arellano, C.J. and Kram, R. (2011). The effects of step width and arm swing on energetic cost and lateral balance during running. *J. Biomech.* 44, 1291–1295.
- Attwell, D. and Laughlin, S.B. (2001). An energy budget for signalling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* 21, 1133–1145.
- Balle, M. (2002). La loi du moindre effort mental: les représentations mentales. *Sciences Humaines* 128, 36–39.
- Barron, A.B., Hebets, E.A., Cleland, T.A., Fitzpatrick, C.L., Hauber, M.E., and Stevens, J.R. (2015). Embracing multiple definitions of learning. *Trends Neurosci.* 38, 405–407.
- Benson-Amram, S. and Holekamp, K.E. (2012). Innovative problem solving by wild spotted hyenas. *Proc. R. Soc. B Biol. Sci.* 279, 4087–4095.
- Bliss, T.V. and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356.
- Bloomfield, T.M. (1967). A peak shift on a line-tilt continuum. *J. Exp. Anal. Behav.* 10, 361–366.
- Braak, H. and Del Tredici, K. (2015). The preclinical phase of the pathological process underlying sporadic Alzheimer's disease. *Brain* 138, 2814–2833.
- Brennan, J.F. and Riccio, D.C. (1972). Stimulus control of shuttle avoidance in young and adult rats. *Can. J. Psychol.* 26, 361–373.
- Brigman, J.L., Wright, T., Talani, G., Prasad-Mulcare, S., Jinde, S., Seabold, G.K., Mathur, P., Davis, M.I., Bock, R., Gustin, R.M., et al. (2010). Loss of GluN2B-containing NMDA receptors in CA1 hippocampus and cortex impairs long-term depression, reduces dendritic spine density, and disrupts learning. *J. Neurosci.* 30, 4590–600.
- Burns, J.G., Foucaud, J., and Mery, F. (2011). Costs of memory: lessons from 'mini' brains. *Proc. R. Soc. B Biol. Sci.* 278, 923–929.
- Cabelli, R.J., Hohn, A., and Shatz, C.J. (1995). Inhibition of ocular dominance column formation by infusion of NT-4/5 or BDNF. *Science* 267, 1662–1666.
- Cantlon, J.F., Pinel, P., Dehaene, S., and Pelphrey, K.A. (2011). Cortical representations of symbols, objects, and faces are pruned back during early childhood. *Cereb. Cortex* 21, 191–199.
- Carreiras, M., Seghier, M.L., Baquero, S., Estévez, A., Lozano, A., Devlin, J.T., and Price, C.J. (2009). An anatomical signature for literacy. *Nature* 461, 983–986.
- Changeux, J.P., and Danchin, A. (1976). Selective stabilisation of developing synapses as a mechanism for the specification of neuronal networks. *Nature* 264, 705–712.
- Chen, J.L., Margolis, D.J., Stankov, A., Sumanovski, L.T., Schneider, B.L., and Helmchen, F. (2015). Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. *Nat. Neurosci.* 18, 1101–1108.
- Cheng, K. (2002). Generalisation: mechanistic and functional explanations. *Anim. Cogn.* 5, 33–40.
- Cheng, K., Spetch, M.L., and Johnson, M. 1997. Spatial peak shift and generalization in pigeons. *J. Exp. Psychol. Anim. Behav. Process.* 23, 469–481.
- Chow, P.K.Y., Lea, S.E.G., and Leaver, L.A. (2016). How practice makes perfect: the role of persistence, flexibility and learning in problem solving efficiency. *Anim. Behav.* 112, 273–283.
- Christie, S.T. and Schrater, P. (2015). Cognitive cost as dynamic allocation of energetic resources. *Front. Neurosci.* 9, 289.
- Commins, S., Gigg, J., Anderson, M., and O'Mara, S.M. (1998). Interaction between paired-pulse facilitation and long-term potentiation in the projection from hippocampal area CA1 to the subiculum. *Neuroreport* 9, 4109–4113.
- Connor, S.A. and Wang, Y.T. (2015). A place at the table: LTD as a mediator of memory genesis. *Neuroscientist* 22, 359–71.
- Clarke, D.D. and Sokoloff, L. (1999). *Circulation and Energy Metabolism of the Brain. Basic Neurochemistry: Molecular, Cellular and Medical Aspects.* G.J. Siegel, B.W. Agranoff, R.W. Albers, S.K. Fisher and M.D. Uhler, eds. (Philadelphia: Lippincott-Raven), pp. 637–669.
- Craig, S. and Commins, S. (2005). Interaction between paired-pulse facilitation and long-term potentiation in the projection from hippocampal area CA1 to the entorhinal cortex. *Neurosci. Res.* 53, 140–146.
- Davis, S.W., Dennis, N.A., Daselaar, S.M., Fleck, M.S., and Cabeza, R. (2008). Que PASA? The posterior–anterior shift in aging. *Cereb. Cortex* 18, 1201–1209.
- Dehaene, S., Pegado, F., Braga, L.W., Ventura, P., Nunes Filho, G., Jobert, A., Dehaene-Lambertz, G., Kolinsky, R., Morais, J., and Cohen L. (2010). How learning to read changes the cortical networks for vision and language. *Science* 330, 1359–1364.
- Dehaene, S., Cohen, L., Morais, J., and Kolinsky, R. (2015). Illiterate to literate: behavioural and cerebral changes induced by reading acquisition. *Nat. Rev. Neurosci.* 16, 234–244.
- De Houwer, J., Barnes-Holmes, D., and Moors, A. (2013). What is learning? On the nature and merits of a functional definition of learning. *Psychon. Bull. Rev.* 20, 631–642.
- Diviney, M., Fey, D., and Commins, S. (2013). Hippocampal contribution to vector model hypothesis during cue-dependent navigation. *Learn. Mem.* 20, 367–378.
- Dukas, R. (1999). Costs of memory: ideas and predictions. *J. Theor. Biol.* 197, 41–50.
- Duzel, E., van Praag, H., and Sendtner, M. (2016). Can physical exercise in old age improve memory and hippocampal function? *Brain* 139(Pt 3), 662–673.
- Dymond, S., Dunsmoor, J.E., Vervliet, B., Roche, B., and Hermans, D. (2015). Fear generalization in humans: systematic review and implications for anxiety disorder research. *Behav. Ther.* 46, 561–582.
- Erickson, K.I., Voss, M.W., Prakash, R.S., Basak, C., Szabo, A., Chaddock, L., Kim, J.S., Heo, S., Alves, H., White, S.W., et al. (2011). Exercise training increases size of hippocampus and improves memory. *Proc. Natl. Acad. Sci. USA* 108, 3017–3022.
- Evans, J.S.T.B.T. (2007). *Hypothetical Thinking. Dual Processes in Reasoning and Judgement* (New York, USA: Psychology Press: Taylor and Francis Group).
- Farina, F.R., Burke, T., Coyle, D., Jeter, K., McGee, M., O'Connell, J., Taheny, D., and Commins S. (2015). Learning efficiency: the influence of cue salience during spatial navigation. *Behav. Processes* 116, 17–27.

- Fay, R. (1970). Auditory frequency generalization in the goldfish (*Carassius auratus*). *J. Exp. Anal. Behav.* 14, 353–360.
- Fernández-Fernández, D., Rosenbrock, H., and Kroker, K.S. (2015). Inhibition of PDE2A, but not PDE9A, modulates presynaptic short-term plasticity measured by paired-pulse facilitation in the CA1 region of the hippocampus. *Synapse* 69, 484–496.
- Fiske, S.T. and Taylor, S.E. (1984). *Social Cognition* (Reading, MA, USA: Addison-Wesley Publishing Co).
- Fox, P.T., Raichle, M.E., Mintun, M.A., and Dence, C. (1988). Nonoxidative glucose consumption during focal physiologic neural activity. *Science* 241, 462–464.
- Francois, J., Huxter, J., Conway, M.W., Lowry, J.P., Tricklebank, M.D., and Gilmour, G. (2014). Differential contributions of infralimbic prefrontal cortex and nucleus accumbens during reward-based learning and extinction. *J. Neurosci.* 34, 596–607.
- Gailliot, M.T., Baumeister, R.F., DeWall, C.N., Maner, J.K., Plant, E.A., Tice, D.M., Brewer, L.E., and Schmeichel, B.J. (2007). Self-control relies on glucose as a limited energy source: willpower is more than a metaphor. *J. Pers. Soc. Psychol.* 92, 325–336.
- Galna, B. and Sparrow, W.A. (2006). Learning to minimize energy costs and maximize mechanical work in a bimanual coordination task. *J. Motor. Behav.* 38, 411–422.
- Gauthier, I. and Logothetis, N.K. (2000). Is face recognition not so unique after all? *Cogn. Neuropsychol.* 17, 125–142.
- Ghirlanda, S. and Enquist, M. (2003). A century of generalization. *Anim. Behav.* 66, 15–36.
- Grillo, F.W., West, L., and De Paola, V. (2015). Removing synaptic brakes on learning. *Nat. Neurosci.* 18, 1062–1064.
- Haier, R.J., Siegel, B., Tang, C., Abel, L., and Buchsbaum, M.S. (1992). Intelligence and changes in regional cerebral glucose metabolic rate following learning. *Intelligence* 16, 415–426.
- Hallermann, S., de Kock, C.P., Stuart, G.J., and Kole, M.H. (2012). State and location dependence of action potential metabolic cost in cortical pyramidal neurons. *Nat. Neurosci.* 15, 1007–1014.
- Han, J.H., Kushner, S.A., Yiu, A.P., Cole, C.J., Matynia, A., Brown, R.A., Neve, R.L., Guzowski, J.F., Silva, A.J., and Josselyn, S.A. (2007). Neuronal competition and selection during memory formation. *Science* 316, 457–460.
- Hanson, H. (1959). Effects of discrimination training on stimulus generalization. *J. Exp. Psychol.* 58, 321–333.
- Harvey, D.R., McGaughan, A.M., Murphy, J., Burns, L., McMonagle, E., and Commins, S. (2008). Emergence of an egocentric cue guiding and allocentric inferring strategy that mirrors hippocampal brain-derived neurotrophic factor (BDNF) expression in the Morris water maze. *Neurobiol. Learn. Mem.* 89, 462–479.
- Hasenstaub, A., Otte, S., Callaway, E., and Sejnowski, T.J. (2010). Metabolic cost as a unifying principle governing neuronal biophysics. *Proc. Natl. Acad. Sci. USA* 107, 12329–12334.
- Hebb, D.O. (1949). *The Organization of Behavior* (New York, USA: Wiley and Sons).
- Herculano-Houzel, S. (2011). Scaling of brain metabolism with a fixed energy budget per neuron: implications for neural activity, plasticity and evolution. *PLoS One* 6, e17514.
- Hu, T., Genkin, A., and Chklovskii, D.B. (2012). A network of spiking neurons for computing sparse representations in an energy-efficient way. *Neural. Comput.* 24, 2852–2872.
- Huang, H.J., Kram, R., and Ahmed, A.A. (2012). Reduction of metabolic cost during motor learning of arm reaching dynamics. *J. Neurosci.* 32, 2182–2190.
- Hunt, K.L. and Chittka, L. (2015). Merging of long-term memories in an insect. *Curr. Biol.* 25, 741–745.
- Kahneman, D. (2011). *Thinking, Fast and Slow* (New York, USA: Farrar, Straus and Giroux).
- Kamil, A.C. and Cheng, K. (2001). Way-finding and landmarks: the multiple-bearings hypothesis. *J. Exp. Biol.* 204, 103–113.
- Kamin, L.J. (1969). Predictability, surprise, attention and conditioning. *Punishment and Aversive Behavior*. B.A. Campbell and R.M. Church, eds. (New York, USA: Appleton-Century-Crofts).
- Kazemi, B., Gamberale-Stille, G., Tullberg, B.S., and Leimar, O. (2014). Stimulus salience as an explanation for imperfect mimicry. *Curr. Biol.* 24, 965–969.
- Kealy, J., Diviney, M., Kehoe, E., McGonagle, V., O’Shea, A., Harvey, D., and Commins S. (2008). The effects of overtraining in the Morris water maze on allocentric and egocentric learning strategies in rats. *Behav. Brain Res.* 192, 259–263.
- Kelly, A.M. and Garavan, H. (2005). Human functional neuroimaging of brain changes associated with practice. *Cereb. Cortex* 15, 1089–1102.
- Kelly, D.M., Kamil, A.C., and Cheng, K. (2010). Landmark use by Clark’s nutcrackers (*Nucifraga columbiana*): influence of disorientation and cue rotation on distance and direction estimates. *Anim. Cogn.* 13, 175–188.
- Kempermann, G., Kuhn, H.G., and Gage, F.H. (1997). More hippocampal neurons in adult mice living in an enriched environment. *Nature* 386, 493–495.
- Kikuchi, D.W. and Pfennig, D.W. (2013). Imperfect mimicry and the limits of natural selection. *Quart. Rev. Biol.* 88, 297–315.
- Kool, W., McGuire, J.T., Rosen, Z.B., and Botvinick, M.M. (2010). Decision making and the avoidance of cognitive demand. *J. Exp. Psychol. Gen.* 139, 665–682.
- Krane, R.V. and Wagner, A.R. (1975). Taste aversion learning with a delayed shock US: implications for the generality of the laws of learning. *J. Comp. Physiol. Psychol.* 88, 882–889.
- Krebs, J.R., Sherry, D.F., Healy, S.D., Perry, V.H., and Vaccarino, A.L. (1989). Hippocampal specialization of food-storing birds. *Proc. Natl. Acad. Sci. USA* 86, 1388–1392.
- Krebs, R.M., Boehler, C.N., and Woldorff, M.G. (2010). The influence of reward associations on conflict processing in the Stroop task. *Cognition* 117, 341–347.
- Kurzban, R. (2010). Does the brain consume additional glucose during self-control tasks? *Evol. Psychol.* 8, 245–260.
- Lay, B.S., Sparrow, W.A., Hughes, K.M., and O’Dwyer, N.J. (2002). Practice effects on coordination and control, metabolic energy expenditure, and muscle activation. *Hum. Mov. Sci.* 21, 807–830.
- Leonard, W.R. (2010). Measuring human energy expenditure and metabolic function: basic principles and methods. *J. Anthropol. Sci.* 88, 221–230.
- Loftus, E.F. and Palmer, J.C. (1974). Reconstruction of automobile destruction: an example of the interaction between language and memory. *J. Verbal. Learn. Verbal. Behav.* 13, 585–589.
- Li, J., Bravo, D.S., Upton, A.L., Gilmour, G., Tricklebank, M.D., Fillenz, M., Martin, C., Lowry, J.P., Bannerman, D.M., and McHugh, S.B. (2011). Close temporal coupling of neuronal activity and tissue oxygen responses in rodent whisker barrel cortex. *Eur. J. Neurosci.* 34, 1983–1996.
- Li, Q., Rothkegel, M., Xiao, Z.C., Abraham, W.C., Korte, M., and Sajikumar, S. (2014). Making synapses strong: metaplasticity prolongs associativity of long-term memory by switching synaptic tag mechanisms. *Cereb. Cortex* 24, 353–363.

- Lin, A.-L., Fox, P.T., Hardies, J., Duong, T.Q., and Gao, J.-H. (2010). Nonlinear coupling between cerebral blood flow, oxygen consumption and ATP production in the human visual cortex. *Proc. Natl. Acad. Sci. USA* 107, 8446–8451.
- Lupien, S.J. and Lepage, M. (2001). Stress, memory, and the hippocampus: can't live with it, can't live without it. *Behav. Brain Res.* 127, 137–158.
- Mackintosh, N. (1974). *The Psychology of Animal Learning* (London, UK: Academic Press).
- Macrae, N.C., Milne, A.B., and Bodenhausen, G.V. (1994). Stereotypes as energy-saving devices: a peek inside the cognitive toolbox. *J. Pers. Soc. Psychol.* 66, 37–47.
- Maguire, E.A., Gadian, D.G., Johnsrude, I.S., Good, C.D., Ashburner, J., Frackowiak, R.S., and Frith, C.D. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proc. Natl. Acad. Sci. USA* 97, 4398–4403.
- Meade, J., Biro, D., and Guilford, T. (2005). Homing pigeons develop local route stereotypy. *Proc. Biol. Sci.* 272, 17–23.
- McNay, E.C., Fries, T.M., and Gold, P.E. (2000). Decreases in rat extracellular hippocampal glucose concentration associated with cognitive demand during a spatial task. *Proc. Natl. Acad. Sci. USA* 97, 2881–2885.
- Mimura, K., Kimoto, T., and Okada, M. (2003). Synapse efficiency diverges due to synaptic pruning following overgrowth. *Phys. Rev. E*, 68, 031910.
- Mink, J.W., Blumenshine, R.J., and Adams, D.B. (1981). Ratio of central nervous system to body metabolism in vertebrates: its constancy and functional basis. *Am. J. Physiol.* 241, 203–212.
- Morcom, A.M., Li, J., and Rugg, M.D. (2007). Age effects on the neural correlates of episodic retrieval: increased cortical recruitment with matched performance. *Cereb. Cortex* 17, 2491–2506.
- Morris, R.G., Anderson, E., Lynch, G.S., and Baudry, M. (1986). Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature* 319, 774–776.
- Mukherjee, K. (2010). A dual system model of preferences under risk. *Psychol. Rev.* 177, 243–255.
- Nabavi, S., Fox, R., Proulx, C.D., Lin, J.Y., Tsien, R.Y., and Malinow, R. (2014). Engineering a memory with LTD and LTP. *Nature* 511, 348–352.
- Nakazawa, K., McHugh, T.J., Wilson, M.A., and Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nat. Rev. Neurosci.* 5, 361–372.
- Neubauer, A.C. and Fink, A. (2009). Intelligence and neural efficiency. *Neurosci. Biobehav. Rev.* 33, 1004–1023.
- Neubauer, A.C., Grabner, R.H., Freudenthaler, H.H., Beckmann, J.F., and Guthke, J. (2004). Intelligence and individual differences in becoming neurally efficient. *Acta Psychol. (Amst)* 116, 55–74.
- Nicholls, R.E., Alarcon, J.M., Malleret, G., Carroll, R.C., Grody, M., Vronskaya, S., and Kandel, E.R. (2008). Transgenic mice lacking NMDAR-dependent LTD exhibit deficits in behavioral flexibility. *Neuron* 58, 104–117.
- Nicholson, J.N. and Gray, J.A. (1971). Behavioural contrast and peak shift in children. *Br. J. Psychol.* 62, 367–373.
- Nyberg, L., Andersson, M., Kauppi, K., Lundquist, A., Persson, J., Pudas, S., and Nilsson, L.G. (2014). Age-related and genetic modulation of frontal cortex efficiency. *J. Cognit. Neurosci.* 26, 746–754.
- Ortega, J.D., Fehlmann L.A., and Farley C.T. (2008). Effects of aging and arm swing on the metabolic cost of stability in human walking. *J. Biomech.* 41, 3303–3308.
- Pascual-Leone, A., Nguyet, D., Cohen, L.G., Brasil-Neto, J.P., Cammarota, A., and Hallett, M. (1995). Modulation of muscle responses evoked by transcranial magnetic stimulation during the acquisition of new fine motor skills. *J. Neurophysiol.* 74, 1037–1045.
- Pascual-Leone, A., Amedi, A., Fregni, F., and Merabet, L.B. (2005). The plastic human brain cortex. *Annu. Rev. Neurosci.* 28, 377–401.
- Pavlov, I.P. (1927). *Conditioned Reflexes* (Oxford, UK: Oxford University Press).
- Perera, T.D., Lu, D., Thirumangalakudi, L., Smith, E.L., Yaretskiy, A., Rosenblum, L.A., Kral, J.G., and Coplan, J.D. (2011). Correlations between hippocampal neurogenesis and metabolic indices in adult nonhuman primates. *Neural Plast.* 2011, 1–6.
- Picard, N., Matsuzaka, Y., and Strick, P.L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in M1. *Nat. Neurosci.* 16, 1340–1347.
- Poldrack, R.A. (2015). Is 'efficiency' a useful concept in cognitive neuroscience? *Dev. Cognit. Neurosci.* 11, 12–17.
- Raichle, M.E. and Mintun, M.A. (2006). Brain work and brain imaging. *Annu. Rev. Neurosci.* 29, 449–476.
- Rashid, A.J., Yan, C., Mercaldo, V., Hsiang, H.L., Park, S., Cole, C.J., De Cristofaro, A., Yu, J., Ramakrishnan, C., Lee, S.Y., et al. (2016). Competition between engrams influences fear memory formation and recall. *Science* 353, 383–387.
- Reed, W.J. (2001). The Pareto, Zipf and other power laws. *Econ. Lett.* 74, 15–19.
- Reiss, S. and Wagner A.R. (1972). CS habituation produces a 'latent inhibition effect' but no active 'conditioned inhibition'. *Learn. Motiv.* 3, 237–245.
- Rescorla, R.A. and Gillan, D.J. (1980). An analysis of the facilitative effect of similarity on second-order conditioning. *J. Exp. Psychol. Anim. Behav. Process.* 6, 339–351.
- Rizzolatti, G., Cattaneo, L., Fabbri-Destro, M., and Rozzi, S. (2014). Cortical mechanisms underlying the organization of goal-directed actions and mirror neuron-based action understanding. *Physiol. Rev.* 94, 655–706.
- Rumpel, S., LeDoux, J., Zador, A., and Malinow, R. (2005). Postsynaptic receptor trafficking underlying a form of associative learning. *Science* 308, 83–88.
- Rypma, B. and D'Esposito, M. (2000). Isolating the neural mechanisms of age-related changes in human working memory. *Nat. Neurosci.* 3, 509–515.
- Sacramento, J., Wichert, A., and van Rossum, M.C. (2015). Energy efficient sparse connectivity from imbalanced synaptic plasticity rules. *PLoS Comput. Biol.* 11, e1004265.
- Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science* 237, 1317–1323.
- Scoville, W.B. and Milner, B.J. (1957). Loss of recent memory after bilateral hippocampal lesions. *Neurol. Neurosurg. Psychiatry* 20, 11–21.
- Sengupta, B., Stemmler, M., Laughlin, S.B., and Niven, J.E. (2010). Action potential energy efficiency varies among neuron types in vertebrates and invertebrates. *PLoS Comput. Biol.* 6, e1000840.
- Solomon R.L. (1948). The influence of work on behavior. *Psychol. Bull.* 45, 1.

- Sparrow, W.A. and Newell, K.M. (1994). Energy expenditure and motor performance relationships in humans learning a motor task. *Psychophysiology* 31, 338–346.
- Stranahan, A.M. and Mattson, M.P. (2008). Impact of energy intake and expenditure on neuronal plasticity. *Neuromol. Med.* 10, 209–218.
- Szilágyi, T., Orbán-Kis, K., Horváth, E., Metz, J., Pap, Z., and Pávai, Z. (2011). Morphological identification of neuron types in the rat hippocampus. *Rom. J. Morphol. Embryol.* 52, 15–20.
- Tang, Y.P., Shimizu, E., Dube, G.R., Rampon, C., Kerchner, G.A., Zhuo, M., Liu, G., and Tsien, J.Z. (1999). Genetic enhancement of learning and memory in mice. *Nature* 401, 63–69.
- Takeuchi, T., Duzskiewicz, A.J., and Morris, R.G. (2013). The synaptic plasticity and memory hypothesis: encoding, storage and persistence. *Phil. Trans. R. Soc. B* 369, 20130288.
- Thiebaut de Schotten, M., Cohen, L., Amemiya, E., Braga, L.W., and Dehaene S. (2014). Learning to read improves the structure of the arcuate fasciculus. *Cereb. Cortex* 24, 989–995.
- Thornton, A. and Samson, J. (2012). Innovative problem solving in wild meerkats. *Anim. Behav.* 83, 1459–1468.
- Turrigiano, G.G. (2008). The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* 135, 422–435.
- Tversky, A. and Kahneman, D. (1973). Availability: a heuristic for judging frequency and probability. *Cogn. Psychol.* 5, 207–232.
- van Praag, H., Christie, B.R., Sejnowski, T.J., and Gage, F.H. (1999). Running enhances neurogenesis, learning, and long-term potentiation in mice. *Proc. Natl. Acad. Sci. USA* 96, 13427–13431.
- Westbrook, A. and Braver, T.S. (2015). Cognitive effort: a neuroeconomic approach. *Cognit. Affect. Behav. Neurosci.* 15, 395–415.
- West-Eberhard, M.J. (2003). *Developmental Plasticity and Evolution* (Oxford, UK: Oxford University Press).
- Whitlock, J.R., Heynen, A.J., Shuler, M.G., and Bear, M.F. (2006). Learning induces long-term potentiation in the hippocampus. *Science* 313, 1093–1097.
- Yang, G., Pan, F., and Gan, W.B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature* 462, 920–924.
- Yger, P., Stimberg, M., and Brette, R. (2015). Fast learning with weak synaptic plasticity. *J. Neurosci.* 35, 13351–13362.
- Zhao, X., Ueba, T., Christie, B.R., Barkho, B., McConnell, M.J., Nakashima, K., Lein, E.S., Eadie, B.D., Willhoite, A.R., Muotri, A.R., et al. (2003). Mice lacking methyl-CpG binding protein 1 have deficits in adult neurogenesis and hippocampal function. *Proc. Natl. Acad. Sci. USA* 100, 6777–6782.
- Zhao, C., Deng, W., and Gage, F.H. (2008). Mechanisms and functional implications of adult neurogenesis. *Cell* 132, 645–660.
- Zipf, G.K. (1949). *Human Behavior and the Principle of Least Effort* (Cambridge, MA, USA: Addison-Wesley).

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.