Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Neurobiology

Dendritic solutions to the credit assignment problem

Blake A Richards[1,2,3] and Timothy P Lillicrap[4]

Guaranteeing that synaptic plasticity leads to effective learning requires a means for assigning credit to each neuron for its contribution to behavior. The 'credit assignment problem' refers to the fact that credit assignment is non-trivial in hierarchical networks with multiple stages of processing. One difficulty is that if credit signals are integrated with other inputs, then it is hard for synaptic plasticity rules to distinguish credit-related activity from non-credit-related activity. A potential solution is to use the spatial layout and non-linear properties of dendrites to distinguish credit signals from other inputs. In cortical pyramidal neurons, evidence hints that top-down feedback signals are integrated in the distal apical dendrites and have a distinct impact on spike-firing and synaptic plasticity. This suggests that the distal apical dendrites of pyramidal neurons help the brain to solve the credit assignment problem.

Addresses
[1] Department of Biological Sciences, University of Toronto Scarborough, Toronto, ON, Canada
[2] Department of Cell and Systems Biology, University of Toronto, Toronto, ON, Canada
[3] Learning in Machines and Brains Program, Canadian Institute for Advanced Research, Toronto, ON, Canada
[4] DeepMind, London, United Kingdom

Current Opinion in Neurobiology 2018, **54**:xx–yy

This review comes from a themed issue on **Neurobiology of learning and plasticity**

Edited by **Scott Waddell** and **Jesper Sjotrom**

https://doi.org/10.1016/j.conb.2018.08.003

Introduction: the credit assignment problem

The flexibility of learning in animals indicates that the brain possesses general purpose *learning algorithms*. A learning algorithm is a set of rules for translating the experiences an animal has into changes in their neural circuits (e.g. synaptic changes). The ultimate goal of a learning algorithm is to alter the behavioral phenotype of the animal, helping it to adapt to the environment. Understanding the brain's learning algorithms is key to understanding the biological basis of animal intelligence.

The formal study of learning algorithms often utilizes the concept of a *loss function* (also known as a cost function) [1,2]. Within neuroscience, a loss function provides a metric for the failure of the current phenotype in achieving an animal's goals (Figure 1a) [3]. For example, a loss function could measure motor slips or sensory prediction errors. Ideally, the brain would have some way of ensuring that changes in a neural circuit reduce a given loss function [3], at least within the environments that the animal is likely to encounter [4]. To do this, it is useful to assign 'credit' (or 'blame') to each neuron or synapse for its contribution to the loss function [5,6]. However, outside of very simple neural circuits, credit assignment calculations are difficult. In a hierarchical sensorimotor circuit with multiple stages of processing, such as the mammalian neocortex, the credit that a neuron in a sensory area deserves for any motor errors depends on that neuron's downstream connections to motor circuits (Figure 1b) [7]. The difficulty of assigning credit in the context of hierarchical circuits is known as the *credit assignment problem* [8].

Typically, solutions to the credit assignment problem have been explored in neural network models that treat each neuron as a single voltage compartment with a single type of output (e.g. a scalar firing-rate or spike train) [7,9••,10••,11–14,15•]. This strategy is reasonable at face value: it fits with the basic properties of neural computation and helps to reduce mathematical complexity. However, there are two reasons that this strategy may have inadvertently made it more difficult to identify the brain's solution to the credit assignment problem. First, if each neuron is calculating *everything* using a single voltage value, then any incoming signals about credit (e.g. feedback from another cortical area) must be integrated with other signals about sensory data, or they must arrive at a separate time. The result is that any credit related signals must be carefully timed or they risk becoming entangled with other ongoing calculations (Figure 1c,d). There is some evidence of clock-like phasic activity in various parts of the brain [16], but none of these seem to exhibit the clear segregation between feedforward and feedback activity required for credit assignment. Second, if a neuron only has one type of output, for example, a firing rate, then it is not immediately obvious how neural circuits can disambiguate credit related activity from basic information transmission (Figure 1f).

Of course, real neurons are not single compartments — they possess complex dendritic trees that integrate different signals in different locations [17–27], often in non-linear manners that have important functional implications [28–44]. Moreover, active channels in dendrites can drive spiking behavior that is different from regular

Figure 1

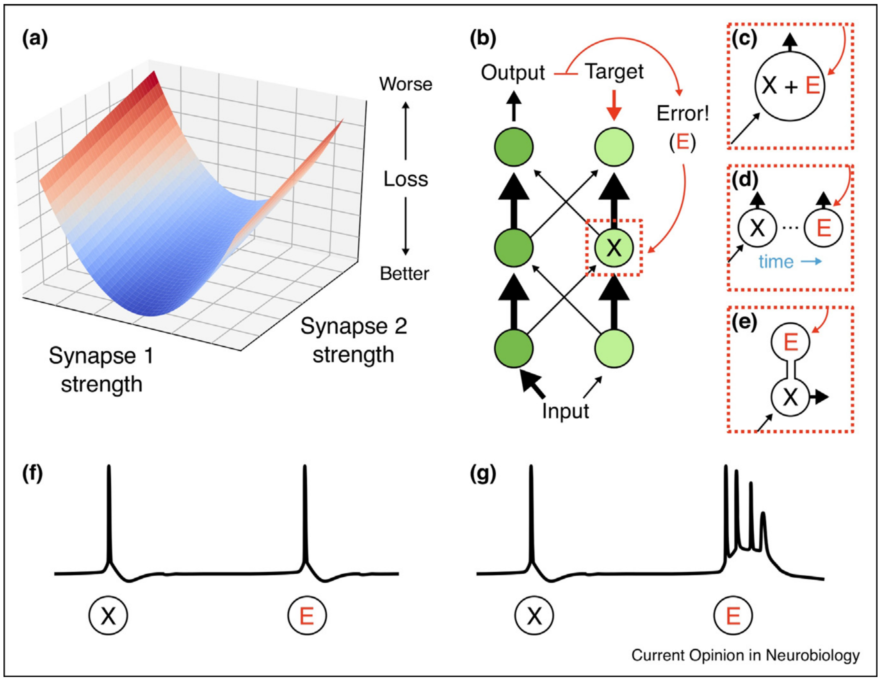

Loss functions and credit assignment. **(a)** Illustration of a loss function. A loss function provides a metric for the performance of an agent on some learning task. In a neural circuit, loss functions are functions of synaptic strength. The goal of learning is to find synaptic strengths that minimize the loss function. Here, an arbitrary loss function is plotted for a network with only two synapses. **(b)** Illustration of the credit assignment problem. A multilayer neural network with two neurons per layer is shown. Circles indicate neurons, with green circles indicating highly active neurons. Arrows indicate synaptic connections and the width of the arrows indicates synaptic strength. If an input arrives at the left-hand neuron, its activity causes strong activation in the downstream left-hand neurons, due to strong synaptic connections. However, if the loss function specifies that the target was to give an output at the right-hand, then an error is generated. To make it more likely that the right-hand output neuron would be activated, it would help to increase the feedforward activity of the right-hand middle neuron, X. In other words, this neuron deserves some 'credit' for the incorrect output. Credit assignment can be achieved if the error signal at the top-level is sent back to the middle-layer. **(c)** However, if the middle-layer neuron is a single compartment, this error signal, E, would be integrated with the ongoing activity, X, thereby altering the 'forward' computation being performed by this neuron. **(d)** A possible solution is to have carefully timed phases where feedforward and feedback signals are received at distinct times. **(e)** An alternative is to integrate the credit assignment signal in a separate dendritic compartment. **(f)** and **(g)** Illustration of the use of specialized spike-waveforms for credit assignment. **(f)** If incoming inputs and credit signals both produce the same type of spiking output in a neuron (indicated by 'X' and 'E', respectively), it is difficult to differentiate credit assignment from ongoing processing. **(g)** In contrast, if credit signals drive dendritic non-linearities that produce unique spike-waveforms (e.g. a complex spike or high-frequency burst), then it is easy to differentiate credit assignment from other processes.

spiking [45,46]. One possibility, then, is to segregate credit signals into dendritic compartments, where (i) they can be kept separate from other ongoing calculations (Figure 1e), and (ii) they can drive unique spike-waveforms that signal credit information (Figure 1g). Thus, there has been a growing interest in understanding whether one of the solutions to the credit assignment problem lies in dendritic computation [47,48••,49••,50•] (and see also IMN Sacramento *et al.* arXiv: 1801.00062).

What counts as evidence for credit assignment?

The ideal experiment for understanding credit assignment in the brain would be to measure a loss function explicitly, then demonstrate that a given synaptic plasticity mechanism was responsible for ensuring reductions in that loss function during learning. Such experiments are currently outside of our technical reach, though, because it is often unclear how we can identify a loss function in the brain and track its progress over time [3]. Furthermore, there is no reason to assume that the brain explicitly represents any of the loss functions it may be reducing. Indeed, at the neural level, it is possible to reduce a loss function without there being any direct neural correlate of said loss function to find [51,52].

Given these realities, the best strategy for scientists to study credit assignment depends on the level of analysis. For example, if the desire is to examine whether credit assignment actually shapes activity in the brain based on the extent to which different neurons contribute to a task [53], then it is possible to use tetrode recordings and similar approaches [54]. In contrast, if the desire is to understand the cellular mechanisms by which credit is