A general error-modulated STDP learning rule applied to reinforcement learning in the basal ganglia

Trevor Bekolay, Chris Eliasmith

We present a novel error-modulated spike-timing-dependent learning rule that utilizes a global error signal and the tuning properties of neurons in a population to learn arbitrary transformations on n-dimensional signals. This rule addresses the gap between low-level spike-timing learning rules modifying individual synaptic weights and higher-level learning schemes that characterize behavioural changes in an animal.

The learning rule is first analyzed in a small spiking neural network. Using the encoding/decoding framework described by Eliasmith and Anderson (2003), we show that the rule can learn linear and non-linear transformations on n-dimensional signals. The learning rule arrives at a connection weight matrix that differs significantly from the connection weight matrix found analytically by Eliasmith and Anderson's method, but performs similarly well.

We then use the learning rule to augment Stewart et al.'s biologically plausible implementation of action selection in the basal ganglia (2009). Their implementation forms the "actor" module in the actor-critic reinforcement learning architecture described by Barto (1995). We add a "critic" module, inspired by the physiology of the ventral striatum, that can modulate the model's likelihood of selecting actions based on the current state and the history of rewards obtained as a result of taking certain actions in that state.

Despite being a complicated model with several interconnected populations, we are able to use our learning rule without any modifications. As a result, we suggest that this rule provides a unique and biologically plausible characterization of supervised and semi-supervised learning in the brain.

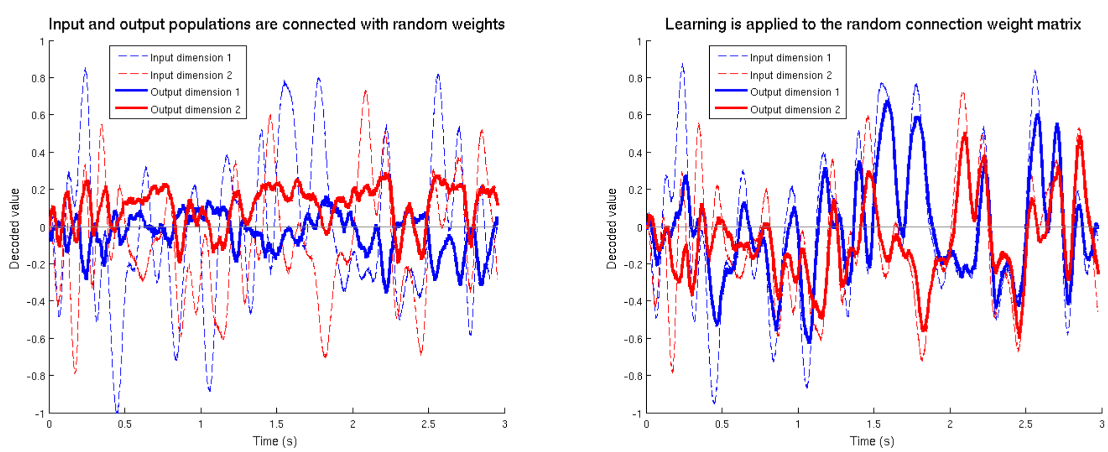

Figure 1: (Left) Without learning, the output population does not perform a useful transformation on the two-dimensional input signal. (Right) With an appropriate error signal, the learning rule quickly learns a linear transformation: a communication channel.