Peter Lathan â€” Synaptic plasticity as probabilistic inference

Organisms face a hard problem: based on noisy sensory input, they must set a large number of synaptic weights. However, they do not receive enough information in their lifetime to learn the correct, or optimal weights (i.e., the weights that ensure the circuit, system, and ultimately organism functions as effectively as possible). Instead, the best they could possibly do is compute a probability distribution over the optimal weights. Based on this observation, we hypothesize that synapses represent probability distributions over weights -- in contrast to the widely held belief that they represent point estimates. From this hypothesis, we derive learning rules for both supervised and unsupervised learning. This introduces a new feature: the more uncertain the brain is about the optimal weight of a synapse, the more plastic it is. Consequently, the learning rate of each synapse is adjusted on the fly. This framework makes several testable predictions and, combined with the ansatz that more uncertain synapses are more variable, it is consistent with current data.