1147-92-753 **Benjamin Lansdell***, ben.lansdell@gmail.com, and **Sofia Triantafillou** and **Konrad Kording**.
Optimizing policies with thresholds in neuroscience.

Many natural processes and human-designed policies make decisions via thresholding. This design in fact provides a powerful tool for separating the effect of the policy from other variables that may be correlated with the thresholded variable, which can then be used to adjust the policy's threshold to maximize utility. Here I present an application of this idea to learning in neural networks. Neurons are active when their synaptic inputs place their membrane potential above a threshold. Inputs that place a neuron marginally above or below threshold are near identical, so any average difference in outcome can then only be attributed to the fact that the neuron spiked. I show how the spiking threshold of a neuron can be used to efficiently learn the effect of the neuron's activity on a utility function, even in the presence of confounding. This can be used to update the neuron's parameters to increase overall utility. I then describe an analysis of the general method as a contextual bandit that shows how to design conservative policies, that only change their threshold when it will increase utility with high probability. This model has applications not just to learning in neural networks but also the many settings in medicine and economics where decisions are made with thresholds. (Received January 28, 2019)