

Mānoa Seminar Series on Machine Learning
and Computational Neuroscience presents:

LEARNING TO LEARN TO ACT

Pieter Abbeel, EECS, UC Berkeley



Reinforcement learning and imitation learning have seen success in many domains, including autonomous helicopter flight, Atari, simulated locomotion, Go, robotic manipulation. However, sample complexity of these methods remains very high. In contrast, humans can pick up new skills far more quickly. To do so, humans might rely on a better learning algorithm or on a better prior (potentially learned from past experience), and likely on both. In this talk I will describe some recent work on meta-learning for action, where agents learn the imitation/reinforcement learning algorithms and learn the prior. This has enabled acquiring new skills from just a single demonstration or just a few trials. While designed for imitation and RL, our work is more generally applicable and also advanced the state of the art in standard few-shot classification benchmarks such as omniglot and mini-imagenet.

October 19, 2017, 12 noon, in ICS-space: POST 318B