

---

# Interactive machine learning via reductions to supervised learning\*

---

Daniel Hsu

## Abstract

I'll describe new approaches for two interactive machine learning problems that operate via reductions to supervised learning. The first problem is contextual bandit learning, where the learner repeatedly selects an action based on the current context, observing and collecting the reward only for the selected action. The second problem is agnostic active learning, where data for a classification task are initially unlabeled, and each individual label must be explicitly requested. In each of these problems, the decision-making of the learner is guided by policies constructed using repeated calls to an oracle for solving cost-sensitive classification problems. [Joint work with Alekh Agarwal, Tzu-Kuo Huang, Satyen Kale, John Langford, Lihong Li, and Rob Schapire.]