Human Active Learning

Rui Castro[*], Charles Kalish, Robert Nowak, Ruichen Qian, Timothy Rogers, and Xiaojin Zhu

[*]Columbia University University of Wisconsin-Madison

Abstract

We investigate a topic at the interface of machine learning and cognitive science. Human active learning, where learners can actively query the world for information, is contrasted with passive learning from random examples. Furthermore, we compare human active learning performance with predictions from statistical learning theory.

We conduct a series of human category learning experiments inspired by a machine learning task for which active and passive learning error bounds are well understood, and dramatically distinct.

Our results indicate that humans are capable of actively selecting informative queries, and in doing so learn better and faster than if they are given random training data, as predicted by learning theory. However, the improvement over passive learning is not as dramatic as that achieved by machine active learning algorithms.

To the best of our knowledge, this is the first quantitative study comparing human category learning in active versus passive settings.

Acknowledgement:
This work is supported in part by the Wisconsin Alumni Research Foundation, and NSF Grant 0745423 from Developmental Learning Sciences.

The Task and Machine Learning Theory Bounds

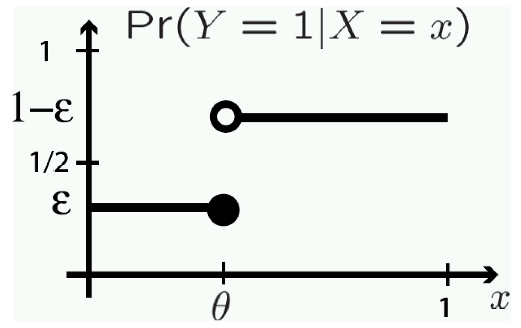

The two-category learning task with boundary θ and noise level ε.

For computers, active learning provably better than passive learning:

Passive learning

With n random training examples, the minimax lower bound for boundary estimate decreases polynomially as O(1/n):

$$\inf_{\hat{\theta}_n} \sup_{\theta\in[0,1]} \mathbb{E}[|\hat{\theta}_n-\theta|] \geq \frac{1}{4}\left(\frac{1+2\epsilon}{1-2\epsilon}\right)^{2\epsilon}\frac{1}{n+1}$$

Active learning

Probabilistic bisecting algorithm. Exponential decrease.

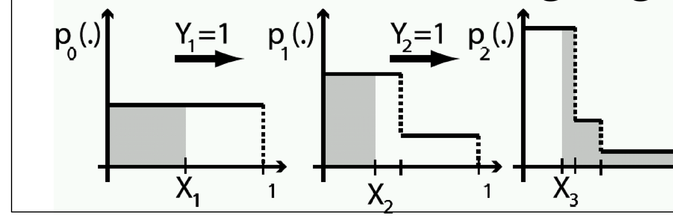

$$\sup_{\theta\in[0,1]} \mathbb{E}[|\hat{\theta}_n-\theta|] \leq 2\left(\sqrt{\frac{1}{2}+\sqrt{\epsilon(1-\epsilon)}}\right)^n$$

Human Passive and Active Learning Experiments

The same two-category learning task, with "alien egg" cover story:

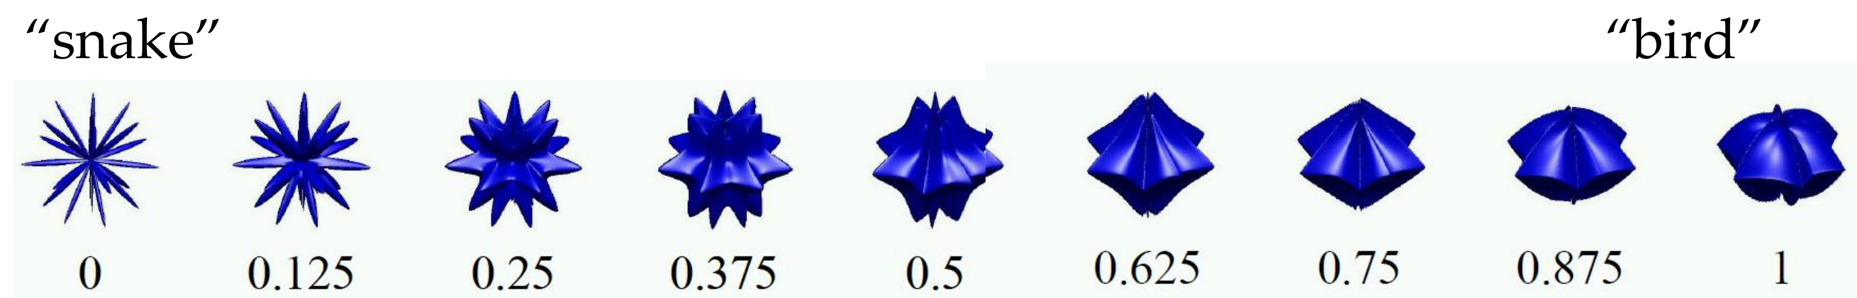

Three conditions

Condition "Random" (passive). Subjects cannot select the queries, and is instead presented sequentially with random examples and their noisy labels. The subject is regularly asked to guess the boundary from these observations (without feedback).

Condition "Human-Active". Subject at each iteration selects a query based on her previous queries and their noisy labels. She then receives a noisy label.

Condition "Machine-Yoked". This is a hybrid human-machine-learning condition in which the human subject passively observes samples selected by an active machine learning algorithm, observes the noisy label generated in response to each query, and is regularly asked to guess, without feedback, where the boundary is – as though the machine is teaching the human.

33 university students. After initial training, 5 sessions of 45 iterations each. True boundary θ randomly set in each session. Boundary guess every 3 iterations. Human-Active subjects use mouse scroll wheel to browse the shapes before determining a query.

Results

Human-Active better than passive. Noise makes human learning harder.

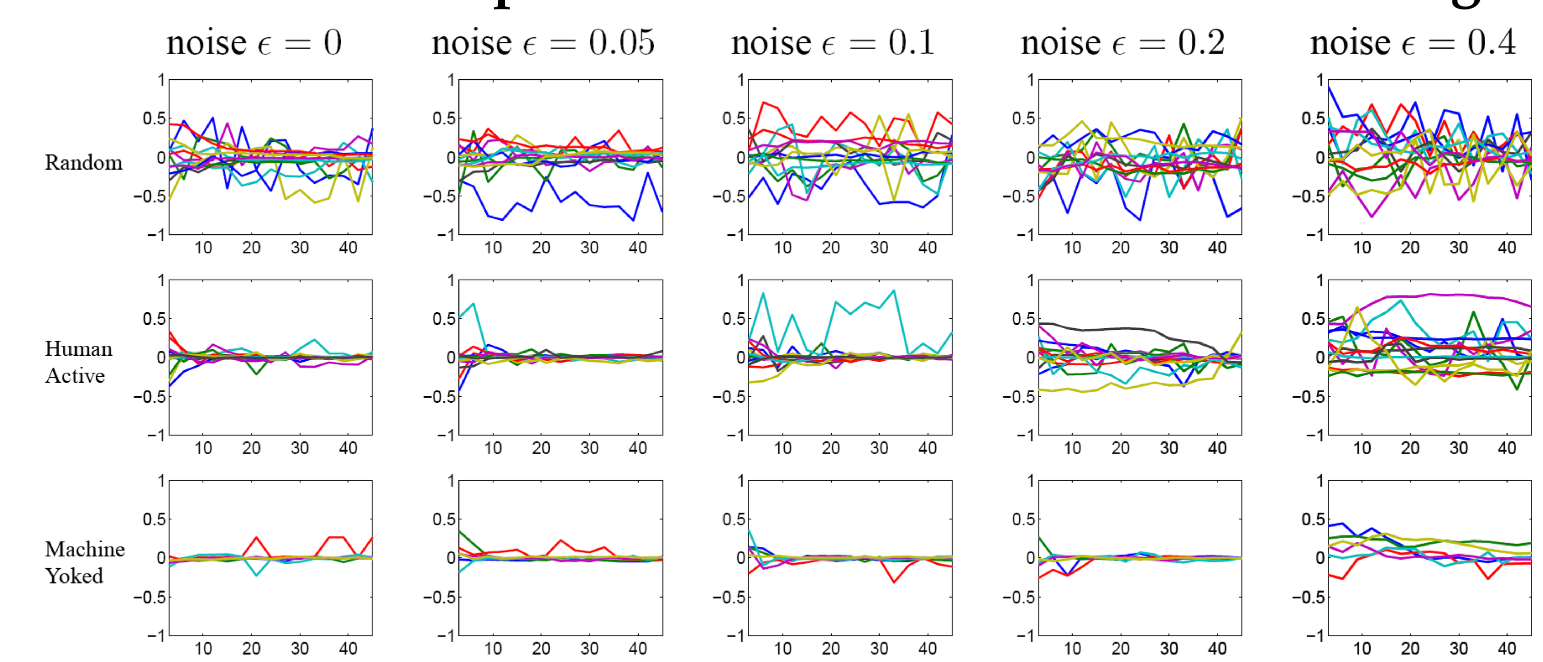

(x-axis: n, y-axis: guess-θ. Each curve is a subject)

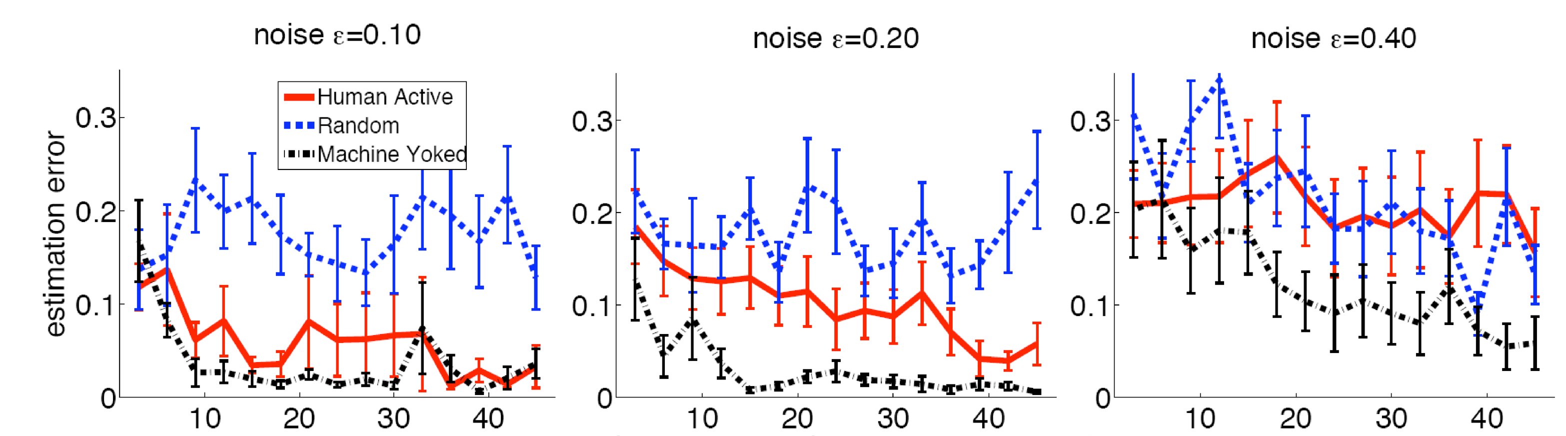

(x-axis: n, y-axis: |guess-θ|. Average of subjects)

Human active learning decreases error exponentially, decay constant smaller than learning theory predicts. Human passive slower than 1/n polynomial.

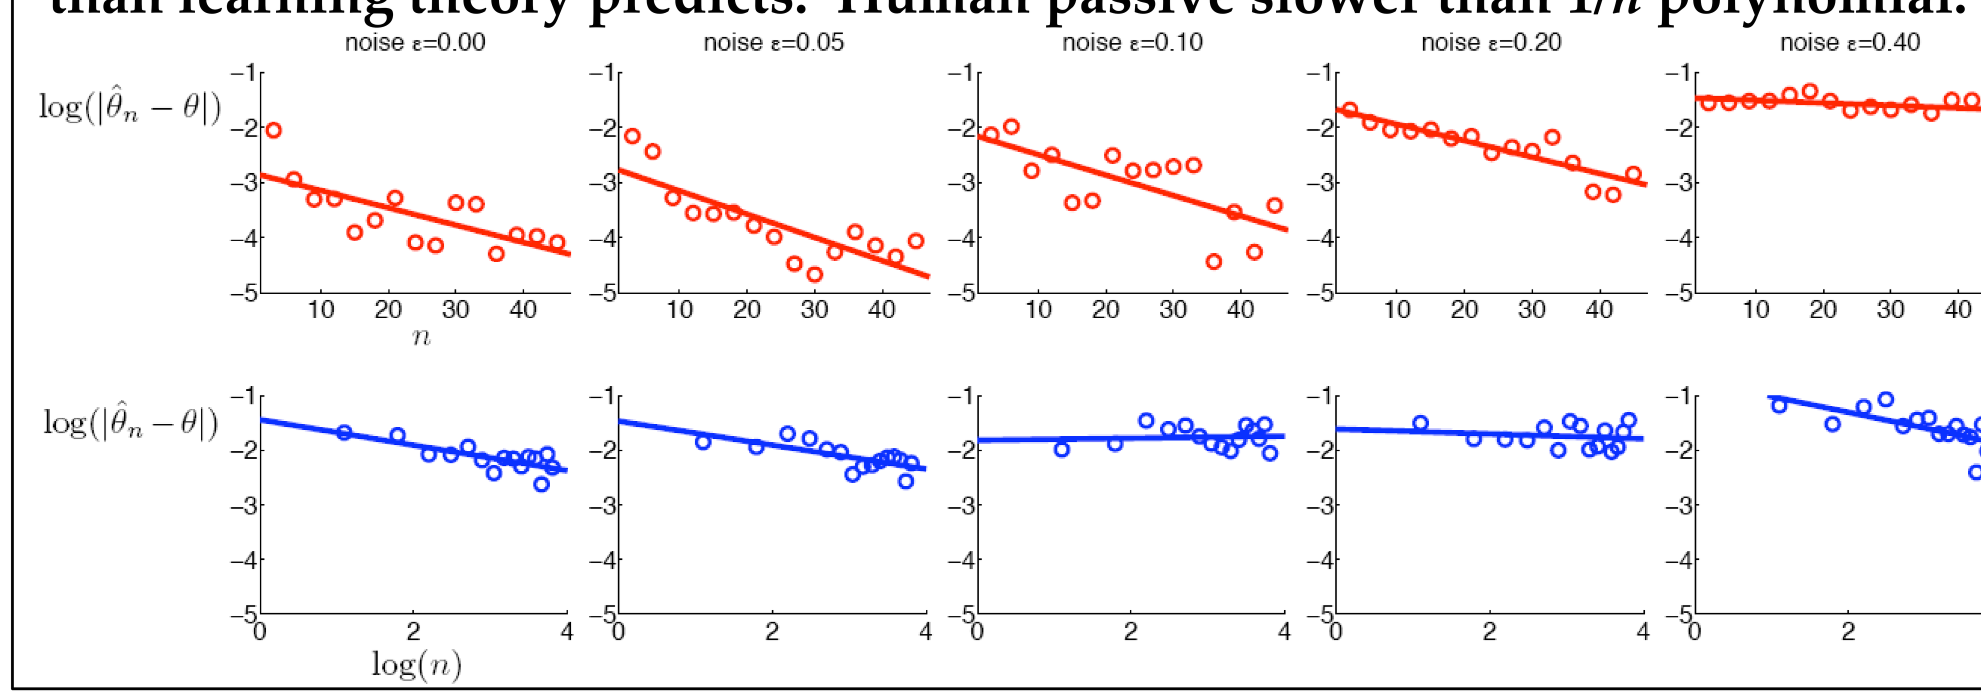

Contact: rmcastro@ee.columbia.edu, jerryzhu@cs.wisc.edu