Overview

Daniel Hsu

COMS 6998-4

About COMS 6998-4

- "Topics in learning theory"
  - This semester: "interactive learning"
- Website:
  `http://www.cs.columbia.edu/~djhsu/coms6998-f17`
  - Course information, policies, academic rules of conduct, etc.
- Courseworks, Piazza: links on website
- Homework 0
  - Required; due 9/13 at 11:59 PM, even if you are still on wait list

About you

- "Mathematical maturity"
  - Ready to read and write lots of proofs.
  - See Homework 0.
- Some useful tools for learning theory
  - Chernoff and union bounds, Taylor's theorem, convexity
  - Reductions
  - Etc.
  - We'll review some basics

About the course staff

- Instructor: Prof. Daniel Hsu
  - Website: `http://www.cs.columbia.edu/~djhsu/`
  - Research in algorithmic statistics, machine learning
  - Office hours: Mon 3-5 PM in 426 Mudd
  - Also have office hours today (Wed Sept 6) 3-5 PM
- Teaching assistant: Ji (Mark) Xu
  - Office hours: TBD

About COMS 6998-4 (again)

1. Online and statistical learning
2. Active learning
3. Teaching and curricula
4. Explanations and interpretations
5. Other forms of feedback

- Major omission: *reinforcement learning*.
  - Bandits and reinforcement learning (COMS 6998-1), Alekh Agarwal and Alex Slivkins.
  - Dynamic programming and reinforcement learning (B9140-1), Daniel Russo.

Online and statistical learning

In online learning, a learner repeatedly interacts with its environment and is judged based on the quality of these interactions.

- Example: content personalization

Active learning

Active learning refers to settings where a learner may adaptively query a teacher.

- Example: restaurant review classifier

Teaching and curricula

Specialized mechanisms for teaching can, in some cases, provide advantages over (say) teaching with random examples.

- Example: the hope of teachers everywhere?
- Example: data set selection to achieve specific machine learning result.
- Question: How can this benefit from interaction?

Explanations and interpretations

As machine learning becomes increasingly prevalent in our everyday lives, it is also becoming important to make its effects as transparent as possible. One way this is achieved is by providing "human-understandable explanations" for the outcomes of machine learning, whatever that means.

- Example: data-driven loan decision

Other forms of feedback

- Example: crowdsourcing
- Example: scientific method
- Etc.

About you (again)

- Read papers/notes (posted on website)
  - Reading responses due before every lecture
  - Instructions on website
- Attend lectures and participate in discussion
- Solve one or two problem sets early on in semester
  - Write-up should pass through a LaTeX compiler
- **Present a paper and/or scribe a lecture**
  - Twice (most likely)
  - Instructions on website; sign-up by 9/20
- **Work on research project**
  - E.g., new, interesting theoretical result
  - E.g., simplify an existing, complex result in a non-trivial way
  - E.g., *high quality* empirical study for real application
- Abide by course policies, academic rules of conduct
  - See website
  - Violators reported to the Dean's office, get failing grade for assignment and/or course

Schedule

- First eight lectures given by me or the TA.
  - Online and statistical learning, some active learning.
- Student paper presentations after that.
- Last lecture (and maybe final exam time): project presentations