CIS 625
Theory of Machine Learning
Prof. Michael Kearns

Outline/Agenda:

- MK background
- Course overview/topics
- Course mechanics
- A toy learning problem: model, algo, analysis

Nature of the Course

- Formal, mathematical models of ML problems
- Generally match implicit assumptions in practice
- Probabilistic/stat frameworks
- Course will be rigorous and proof-based
- Emphasis on general principles/methods, algo & statistical building blocks — i.e. theory.

Sample Questions/Topics

- What are some good general models for ML problems?
- What do learning algos look like in those models?
- What are the computation time & sample size required for learning? (other resources)
- What are the computational & probabilistic limitations to efficient ML?

- Relationships & reductions between models/problems
- Types/categories of ML problems (e.g. supervised, unsupervised, adversarial, reinforcement learning...)

Sample Tools/Methods

- Algorithms & complexity
- Probability & stats
- Optimization & approximation
- Game theory
- Touch briefly on many others: e.g. control, cryptography, ...

(Very) Rough Outline

- First half of course: deep-dive into "PAC" model of ML & variants
- Answer many basic questions & develop algo & prob. tools
- Based on K&V textbook (will provide chapters electronically)

- Second half: other models & topics:
  - online adversarial learning & game theory
  - ML & fairness
  - ML & differential privacy
  - reinforcement learning
  - query models, active & semi-supervised learning
  - theory of deep learning
  - ⋮

Course Mechanics

- Lectures Tue Thu 10:30 ET on Zoom
- Encourage "live" attendance with video on
- Course website for announcements/materials
- Readings: KV textbook chapters & papers from ML literature
- May create course chat forum
- TAs: Hua Wang & Sheng Gao (stats)

Course Requirements:

- Keep up with lectures & readings
- 3-4 problem sets
- final project/paper
- group work allowed, details TBA
- grading will be "doctoral style"
- (class is ~2/3 grad, 1/3 undergrad; CS, stats, math, eng., econ, ...)

Questions?

Comments?

Let's get started...