EECS 496-10 – Computational Learning Theory

Syllabus

Lev Reyzin

Winter 2019

Time and location: T-R, 9:30-10:50am, Tech MG28

Instructor: Lev Reyzin, Mudd 3005, `lreyzin@northwestern.edu`

Prerequisite background: Familiarity with the design and analysis of algorithms, basic computational complexity theory, and mathematical maturity.

Office hours: TBD

Website: `http://www.levreyzin.com/teaching/w19_eecs496-10/`

Required textbook: Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning, 2nd edition*

Optional textbook: Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*

Topics: This course will introduce some of the central topics in computational learning theory, a field which approaches the question "whether machines can learn" from the perspective of theoretical computer science. We will study well defined and rigorous mathematical models of learning where it will be possible to give precise and rigorous analysis of learning problems and algorithms. A big focus of the course will be the computational efficiency of learning in these models. We will develop some provably

efficient algorithms and explain why such provable algorithms are unlikely for other models.

We will only cover topics which permit a mathematically rigorous analysis and hence mathematical maturity is absolutely necessary. In particular, some familiarity with basic probability (such as linearity of expectation) and basic linear algebra will be necessary. Also, the emphasis of the course will be on proofs, so if you are in this course, you should enjoy proofs and math.

Example topics include inductive inference, query learning, PAC learning and VC theory, Occam's razor, online learning, boosting, support vector machines, bandit algorithms, statistical queries, Rademiacher complexity, and neural networks.

Grading:

- 60% take-home problem sets
- 40% final exam

Grades may also be adjusted upward or downward depending on class participation.

Problem set collaboration policy: Unless otherwise specified on an assignment, students may discuss problem sets with one another, but they *must write the solutions on their own.* Collaborators (people you speak to about an assignment) must be named at the top of the assignment. No collaboration will be allowed on exams.

Late work policy: In general, late work will not be accepted.