Non-Computer Science students welcome!
Prerequisite: CS 4413, or permission of instructor

Fall 2020
Mondays, Wednesdays, and Fridays
9:45am – 10:35am

Devon Energy Hall (DEH)
Room 120 (Subject to change)

Dimitris Diochnos
School of Computer Science
244 Devon Energy Hall
diochnos@ou.edu
5-2319

Computer Science 5970
Computational Learning Theory

Computational learning theory views learning as a computational process and tries to understand the principles that govern learning using tools from computer science and mathematics. We will start the course with basic topics such as learning using membership and equivalence queries, version spaces, decision trees, and perceptrons. In sequence we will explore the probably approximately correct (PAC) model of learning, where we will study Occam algorithms, the VC-dimension, and the sample complexity of distribution-independent learning. We will discuss representation issues, proper learning, reductions, intractability, learning in the realizable case, and agnostic learning. Moreover, we will explore PAC learning under different noise models, the statistical query model as an approach of mitigating noisy oracles, and we will discuss robustness guarantees in adversarial settings (poisoning attacks, adversarial examples). Other topics include weak and strong learning (boosting), distribution-specific learning by drawing results from the framework of evolvability, online learning and learning with expert advice in the mistake bound model.

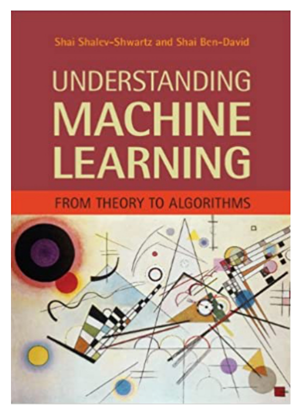

The seminar format will include reading, discussion and rigorous derivations of mathematical properties that characterize various learning problems. We will cover several seminal results spanning the last 40 years of research. Several handouts will be distributed and we will maintain our own notes for the course as we make progress. Nevertheless, two books can complement our discussion in class and either one of them can work well as a reference: (A) *Understanding Machine Learning: From Theory to Algorithms* by Shai Shalev-Schwartz and Shai Ben-David, and (B) *Foundations of Machine Learning* by Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. Both books are available for free online. The goals of this course are for students to: (1) develop a comprehensive understanding of the field; (2) understand and be able to apply formal approaches that provide guarantees on the behavior of machine learning algorithms, or argue about the inherent hardness of certain learning problems; (3) understand research papers in this broad field and be able to present the ideas to others.

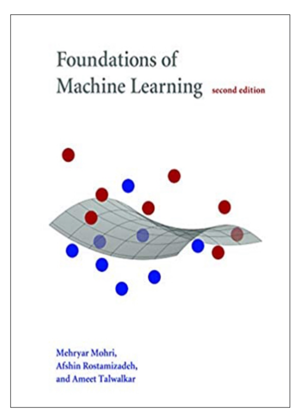

School of Computer Science
College of Engineering