

EECS 598

Statistical Learning Theory

University of Michigan, Winter 2014

In this course we will prove performance guarantees that quantify the ability of a machine learning algorithm to generalize from training data to unseen test data. Potential topics to be covered include concentration of measure, uniform deviation bounds, empirical and structural risk minimization, Rademacher complexity, Vapnik-Chervonenkis theory, consistency and rates of convergence, margin-based bounds, stability bounds, and application of these theories to learning algorithms such as decision trees, boosting, support vector machines, and kernel density estimators. Additional topics as time allows may include online learning, active learning, transfer learning, and ranking.

Prerequisites: Graduate level probability such as EECS 501, and some experience with rigorous mathematical proofs.

Instructor: Clayton Scott

Meeting time: MW 9:00 – 10:30

