

abstract

COMPUTER SCIENCE/DISCRETE MATH SEMINAR, I
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

The goal of Statistical Learning Theory is to construct and understand algorithms that are able to generalize from a given training data set. Statistical learning algorithms are wildly popular now due to their excellent performance on many types of data; one of the most successful learning algorithms is AdaBoost, which is a classification algorithm designed to construct a "strong" classifier from a "weak" learning algorithm.

Just after the development of AdaBoost eight years ago, scientists derived margin-based generalization bounds to explain AdaBoost's good performance. Their result predicts that AdaBoost yields the best possible performance if it always achieves a "maximum margin" solution. Yet, does AdaBoost achieve a maximum margin solution? At least three large-scale empirical studies conducted within this eight year period conjecture the answer to be "yes".

In order to understand AdaBoost's convergence properties and answer this question, we look toward AdaBoost's dynamics. We simplify AdaBoost to reveal a nonlinear iterated map and analyze its behavior in specific cases. We find stable cycles for these cases, which can explicitly be used to solve for AdaBoost's output. Thus, we find the key to answering the question of whether AdaBoost always maximizes the margin - a set of examples in which AdaBoost's convergence can be completely understood. Using this unusual technique, we are able to answer the question of whether AdaBoost always converges to a maximal margin solution; the answer is the opposite of what was thought to be true!

In this talk, I will introduce AdaBoost, describe our analysis of AdaBoost when viewed as a dynamical system, and briefly mention a new boosting algorithm which always maximizes the

abstract

margin with a fast convergence rate.

This talk is designed for a general mathematical audience that likes to see pretty pictures of cyclic dynamics.