



Course Information

CS 59300 – Machine Learning Theory

Instructional modality: In-person lectures

3 credit hours

Prerequisites: grad standing or permission of the instructor. MA/STAT 41600 Probability (or equivalent), CS 38100

Introduction to the Analysis of Algorithms (or equivalent)

Instructor(s) Contact Information

Instructor: Steve Hanneke

Office Location: LWSN 2116K

Homepage: <http://www.stevehanneke.com>

Email Address: hanneke@purdue.edu

Weekly office hour day/time TBD

Course Description

This course will focus on the theory of machine learning. We will start with the basics: discussing what kinds of accuracy guarantees are possible for learning algorithms, what classes of concepts are learnable, and how much training data is necessary and sufficient to achieve a given accuracy guarantee. Further topics may include the theories of online learning, boosting, sample compression schemes, multiclass classification, semi-supervised learning, active learning, universal consistency, contextual bandits, reinforcement learning, and others. The course will also introduce students to recent trends in machine learning theory, including many open problems currently ripe for progress. The material covered will make connections to related subjects, such as statistics and probability theory, combinatorics, algorithm design and analysis, game theory, and information theory.

Learning Outcomes

By the end of the semester, you should be able to:

1. Understand the state of the art of machine learning theory, including the major results and their proofs.
2. Read and comprehend learning theory results and their proofs, such as those that may appear in papers published at the Conference on Learning Theory (COLT) or other ML Theory conferences.
3. Produce mathematical proofs of claims relevant to machine learning theory.

How to Succeed in this Course

- Attend the lectures
- Do the readings
- Do the assignments
- Ask questions if you need clarification



Learning Resources, Technology & Texts

- Textbook: Shalev-Shwartz & Ben-David
 - There will be required readings listed on the Brightspace page.
- Additional Readings: A list of additional readings (mostly papers or notes) will be given in Brightspace. Assignments may sometimes be linked to those readings. Some additional readings will be marked “optional”, indicating they are purely for students wishing to dig deeper into the subject.
- We will use Brightspace for releasing assignments, making announcements, etc.

Course Logistics

The instruction is mostly lecture-based.

The class schedule will be posted and updated in Brightspace.

Homework assignments will be released approximately every 2-3 weeks.

There will be one midterm exam, and one final exam.

Assignments

There are approximately 5 homework assignments that require solving mathematical problems, including producing proofs.

Grading will be as follows:

Assignments	60%
Midterm	20%
Final exam	20%

More details are given in Brightspace associated with each assignment.

Homework groups of 2-3 students are acceptable, but each student must write up his/her own solution. Solutions that appear to have been copied will not be accepted.

Students may also complete an optional course project for extra credit if needed. These projects should represent original research contributions to machine learning theory relevant to the material presented in class, and demonstrating mastery of the material covered in class. These optional projects must be completed by an individual student, without any collaboration with others. The total point value of a student project is potentially unbounded, though earning enough extra credit to significantly affect a student’s letter grade may require a truly significant research contribution to learning theory.