

Lectures

Stat 154/254: Statistical Machine Learning

Lectures marked with ★ are still in progress.



Introduction

- [Class introduction](#)

Unit 1: Regression in a machine learning setting

- [Linear regression in matrix notation \(review\)](#)
- [Linear regression in an ML setting](#)
- [Basis expansions and splines for linear regression](#)
- [L1 and L2 penalization](#)

Unit 2: Classification

- [Empirical risk minimization for classification](#)
- [Trading off false positives and negatives](#)
- [Perceptron algorithm \(bonus content\)](#)

Unit 3: Risk and complexity

- [Asymptotics and bias–variance tradeoff in linear regression](#)
- [Regret decomposition](#)
- [Cross validation](#)
 - [Held-out data and cross-validation](#)
 - [Simulated bias–variance tradeoff example](#)
 - [Cross-validation cautionary note](#)
- [Generalization bounds](#)
 - [Generalization error and uniform laws](#)
 - [Uniform laws for smooth function classes](#)
 - [VC dimension and zero–one loss](#)

Unit 4: Trees and weak learners

- [Classification and regression trees](#)
- [Bagging weak learners](#)
- [Boosting weak learners](#)

Unit 5: Neural networks and optimization

- [Optimization notes](#)

Prerequisites

Mathematical statistics, probability, linear algebra, multivariate calculus, and Python are all prerequisites.

For assistance with mathematics prerequisites, the [Student Learning Center](#) provides additional assistance.

A summary of the required linear algebra skills can be found [here](#).