

CS 559: Machine Learning: Fundamentals and Applications

Instructor: Philippos Mordohai (mordohai@cs.stevens.edu)

Textbook

Pattern Classification (2nd Edition) (Hardcover)
by Richard O. Duda, Peter E. Hart and David G. Stork
Publisher: Wiley-Interscience; 2 edition (October 2000)
ISBN-10: 0471056693
ISBN-13: 978-0471056690

Office Hours: Tuesday 5-6

Evaluation

Project 25%

4 homework sets 20%

Pop-up quizzes and participation 10%

Midterm 20%

Final 25%

Policies

Collaboration Policy. Homework assignments will be done individually: each student must hand in their own answers. It is acceptable for students to collaborate in understanding the material but not in solving the problems. Use of the Internet is allowed, but should not include searching for previous solutions or answers to the specific questions of the assignment. We will assume that, as participants in a graduate course, you will be taking the responsibility of making sure that you personally understand the solution to any work arising from collaboration.

Late Submission Policy. The penalty for late submission is 20% of the grade per day, enforced at 6:15 each day after the due date. If urgent or unusual circumstances prohibit you from submitting a homework assignment in time, please e-mail the instructor explaining the situation.

Week 1: Introduction, including probability theory background (DHS Ch. 1 + Appendix A.4)

Week 2: Bayesian decision theory (DHS Ch. 2)

Week 3: Maximum likelihood; Bayesian parameter estimation (DHS Ch.3)

Week 4: Expectation Maximization; Hidden Markov Models (DHS Ch.3)

Week 5: Principal Component Analysis (notes)

Week 6: Eigenfaces; Fisher Linear Discriminant; Nonparametric techniques (DHS Ch. 3 and 4 + notes)

Week 7: Linear Discriminant Functions and Perceptron (DHS Ch. 5)

Week 8: Midterm; Perceptron (part II) (DHS Ch. 5 + notes)

Week 9: Project proposals; MSE procedures for linear discriminant functions (DHS Ch. 5)

Week 10: Support Vector Machines (DHS Ch. 5 + notes)

Week 11: Support Vector Machines (part II); Boosting (DHS Ch. 5 + notes)

Week 12: Graphical Models and Markov Random Fields (notes)

Week 13: Unsupervised methods and clustering (DHS Ch.10)

Week 14: Project presentations