CS 229
Machine Learning
Handout #1: Course Information

Meeting Times and Locations

Lectures	Mondays and Wednesdays, 9:30 AM - 10:50 AM Bishop Auditorium
Discussion Sections	Friday 1:30-2:50 PM Section Bishop Auditorium (optional attendance)

Teaching Staff

Professor
Andrew Ng
Office: Gates 112

Professor
Ron Dror
Office: Gates 204

Course Coordinator
Swati Dube Batra
Office: Gates 108

Head TAs
Anand Avati
Office: Gates 108

Raphael Townshend
Office: Clark S221

Teaching Assistants

Shervine Amidi	Tony Duan	Yu Wang
Christopher Chute	Xuhua Gao	Xiaoyan Wu
Gael Colas	Shreyansh Pandey	Jian Zhang
Ashwini Ramamoorthy	Suvadip Paul	Cindy Wang
Jaspreet Kaur	Mario Srouji	Akhila Yerukola
Hoon Ji Kim (Andy)	Ethan Steinberg	Ashkon Farhangi
Zahra Koochak	Atharva Parulekar	Vikranth Dwaracheria
Christopher Yeh	Fantine Huot	

Remote Teaching Assistants

Varsha Shankar	Kushaagra Goyal
Sagar Honnungar	Hojat Ghorbani
Ishan Patil	

Contact Information

If you and have a homework, technical or general administrative question about CS229, for you to get the fastest possible response, please post it on our Piazza forum. For information about office hours (where we can help you in person), see refer to the Office Hours calendar on the website.

Course Description

This course provides a broad introduction to machine learning and statistical pattern recognition. Topics include: supervised learning (generative/discriminative learning, parametric/non-parametric learning, neural networks, support vector machines); unsupervised learning (clustering, dimensionality reduction, kernel methods); learning theory (bias/variance tradeoffs; VC theory; large margins); reinforcement learning and adaptive control. The course will also discuss recent applications of machine learning, such as to robotic control, data mining, autonomous navigation, bioinformatics, speech recognition, and text and web data processing.

Prerequisites

Students are expected to have the following background:

- Knowledge of basic computer science principles and skills, at a level sufficient to write a reasonably non-trivial computer program.
- Familiarity with the basic probability theory. (CS 109 or STATS 116)
- Familiarity with the basic linear algebra (MATH 104, MATH 113, or CS205)

Course Materials

There is no required text for this course. Notes will be posted periodically on the course web site. The following books are recommended as optional reading:

- Richard Duda, Peter Hart and David Stork, *Pattern Classification*, 2nd ed. John Wiley & Sons, 2001.
- Tom Mitchell, *Machine Learning*. McGraw-Hill, 1997.
- Richard Sutton and Andrew Barto, *Reinforcement Learning: An introduction*. MIT Press, 1998
- Trevor Hastie, Robert Tibshirani and Jerome Friedman, *The Elements of Statistical Learning*. Springer, 2009

Course handouts and other materials can be downloaded from
http://www.stanford.edu/class/cs229/syllabus.html

Online Resources

- Home page: http://cs229.stanford.edu/
- Assignments: https://www.gradescope.com
- Current quarter's class videos: https://mvideox.stanford.edu/Course/1184
- Piazza forum: https://piazza.com/class/jkbylqx4kcp1h3

Homeworks and Grading: There will be four written **homeworks** (4 x 10 = 40%), one **midterm** (20%), and one major open-ended **term project** (40%). The homeworks will contain written questions and questions that require some programming. In the term project, you will investigate some interesting aspect of machine learning or apply machine learning to a problem that interests you. The term project may be done in teams of up to three persons. We try very hard to make questions unambiguous, but some ambiguities may remain. Ask if confused or state your assumptions explicitly. Reasonable assumptions will be accepted in case of ambiguous questions.

Honor code: We strongly encourage students to form study groups. Students may discuss and work on homework problems in groups. However, each student must write down the solutions independently, and without referring to written notes from the joint session. In other words, each student must understand the solution well enough in order to reconstruct it by him/herself. In addition, each student should write on the problem set the set of people with whom s/he collaborated. Further, since we occasionally reuse problem set questions from previous years, we expect students not to copy, refer to, or look at the solutions in preparing their answers. **It is an honor code violation to intentionally refer to a previous year's solutions.** This applies both to the official solutions and to solutions that you or someone else may have written up in a previous year.

Late Assignments: Each student will have a total of **seven (7) free late (calendar) day**s to use for homeworks, project proposals and project milestones. Once these late days are exhausted, any assignments turned in late will be penalized 20% per late day. However, **no assignment will be accepted more than three days after its due date**, and late days cannot be used for the final project writeup. Each 24 hours or part thereof that a homework is late uses up one full late day.

Assignment Submission: All assignments (both on-campus and SCPD students) will be submitted online in Gradescope. You should have received an email invite to the Gradescope course at the email registered in Axess.

Regrade requests: After grades are released on Gradescope, there will be a regrade period of **3 days** during which you can request regrades for specific questions. No regrade requests will be accepted after the 3 day period.

Midterm exam: The midterm will be a take-home exam for both on-campus and SCPD students, in the week of Nov 5th. Details of exact duration and format will be posted soon.

Sections: To review material from the prerequisites or to supplement the lecture material, there will occasionally be extra discussion sections held on Friday. An announcement will be made whenever one of these sections is held. Attendance at these sections is optional.

Communication with the Teaching Staff: If you have a question that is not confidential or personal, we encourage you to post it on our forum on Piazza. To contact the teaching staff directly, we strongly encourage you to come to office hours. By having questions sent to all of us, you will get answers much more quickly. Of course, confidential or personal questions can still be sent directly to Professor Ng, Professor Dror, or the TAs via a private post on Piazza. For grading questions, please talk to us after class or during office hours.

Answers to commonly asked questions and clarifications to the homeworks will be posted on the FAQ.

Schedule

- **Introduction** (1 class)
Basic concepts.

- **Supervised learning.** (6 classes)
Supervised learning setup. LMS.
Logistic regression. Perceptron. Exponential family.
Generative learning algorithms. Gaussian discriminant analysis. Naive Bayes.
Support vector machines.
Model selection and feature selection.
Evaluating and debugging learning algorithms.

- **Learning theory.** (2 classes)
Bias/variance tradeoff.
Practical advice on how to use learning algorithms.

- **Deep Learning.** (2 classes)
Neural Networks. Backpropagation.
Vectorization.

- **Unsupervised learning.** (5 classes)
Clustering. K-means.
EM. Mixture of Gaussians.
Factor analysis.
PCA (Principal components analysis).
ICA (Independent components analysis).

- **Reinforcement learning and control.** (4 classes)
MDPs. Bellman equations.
Value iteration and policy iteration. Linear quadratic regulation (LQR). LQG. Q-learning. Value function approximation. Policy search. Reinforce. POMDPs.

Dates for assignments

- Assignment 1: Out 10/03. Due 10/17.
- Assignment 2: Out 10/17. Due 10/31.
- Assignment 3: Out 10/31. Due 11/14.
- Assignment 4: Out 11/14. Due 12/05.
- Take-Home Midterm: Out 11/07 (6:30pm). Due 11/09 (6:30pm).

Term project

- Proposals due 10/19 (11:59pm).
- Milestone due 11/16 (11:59pm).
- Poster presentations on 12/11 (8:30am-11:30am).
- Final report due on 12/13 (11:59pm, no late days).