

# CS281: Advanced Machine Learning (Fall 2013)

Instructor: Prof. Ryan P. Adams  
Teaching Fellows: Eyal Dechter, Scott Linderman, Dougal Maclaurin  
Lectures: Monday and Wednesday, 1-2:30pm  
Location: Maxwell-Dworkin G125  
Section: TBD  
Office Hours: Monday 2:30-4:00pm in MD 233 (Adams)  
URL: <http://www.seas.harvard.edu/courses/cs281>  
Contact: [cs281-f13-staff@seas.harvard.edu](mailto:cs281-f13-staff@seas.harvard.edu)

## Course Description

This course is about learning to extract statistical structure from data, for making decisions and predictions, as well as for visualization. The course will cover many of the most important mathematical and computational tools for probabilistic modeling, as well as examine specific models from the literature and examine how they can be used for particular types of data. There will be a heavy emphasis on implementation. You may use Matlab, Python or R. Each of the five assignments will involve some amount of coding, and the final project will almost certainly require the running of computer experiments.

## Prerequisites

You should feel comfortable with the basics of probability: joint densities, conditional distributions, etc. An undergraduate level is fine — there will not be much measure theory and such background is not required. There will be some basic linear algebra, e.g., solving linear systems and thinking about eigenvalues/eigenvectors. We will regularly use multivariable calculus.

## Readings and Lecture Preparation

Each week will have some required reading, typically from books and papers. There will also be optional reading and videos to watch. Machine learning is unique in that there are many videos that are freely available on the web of world-leading researchers discussing their work. It is highly recommended to watch these videos. Everyone learns in a different way and so it may be helpful to hear the same material presented from a slightly different point of view. Moreover, the availability of these videos means that class time can be used for dynamic interaction rather than just lecturing.

To help encourage you to do the reading, there will be an online quiz before each lecture. This will be a very simple, multiple choice quiz that is not intended to be difficult if you have read the required material. These quizzes will be on iSites at <http://isites.harvard.edu/icb/icb.do?keyword=k98807> and will comprise 5% of your final grade. They are due at 13:00 on the day of each lecture.

There will be a strongly-recommended weekly section led by the teaching staff. **You will be evaluated on material that may be presented only in section.**

## Assignments

There will be five assignments due over the course of the term. Each of these will be worth 10% of your final grade, after dropping the lowest one. These will be assigned and due according to the following schedule, in which you have two weeks for each one:

	<i>Out</i>	<i>Due</i>
Assignment 1	6 September 2013	20 September 2013
Assignment 2	20 September 2013	4 October 2013
Assignment 3	4 October 2013	18 October 2013
Assignment 4	25 October 2013	8 November 2013
Assignment 5	8 November 2013	22 November 2013

The assignments should be completed using  $\text{\LaTeX}$ ; the template file will be available with the problem set. All assignments should be submitted electronically by 23:59 on the due date, via the iSites dropbox at <http://isites.harvard.edu/icb/icb.do?keyword=k98807>. It is your responsibility to ensure that the PDF files are readable. We will make every effort to grade the assignments promptly and return them to you in the next section.

**Collaboration Policy** You are encouraged to discuss high-level solutions and problem-solving approaches with other students in the class, but ultimately you must write your own code and produce your own results. If you have collaborated with one or two other students in the planning and design of solutions to the assignments, provide their names on your writeup in the space provided.

**Regrading Policy** If you think there is a mistake with the grading of your midterm or one of your assignments, you may submit it to your section TF for a regrade. Note that this regrade is a new draw from the grading distribution and may result in your grade moving up or down. There is a time limit on these requests: you may only ask for a regrade within two weeks of the section at which the assignment/midterm was returned.

**Late Policy** There are no late days or extensions under any circumstances. Instead:

- Your lowest assignment grade will be dropped. Keep this in reserve in case something unexpected arises.
- If you fail to turn in an assignment on time, you can turn it in up to a week late with a penalty of 50%.

## Midterm Exam

There will be an in-class midterm exam, scheduled for Wednesday 23 October 2013. This exam comprises 15% of your grade and will cover material from the beginning of the class, through the Friday 18 October section on loopy belief propagation.

## Final Project

In the second half of the course, you will complete a project. The ideal outcome of this project would be a paper that could be submitted to a top-tier machine learning conference such as NIPS, ICML or AISTATS. There are different ways to approach this project, which are discussed in a more comprehensive document that is available from the course website at <http://www.seas.harvard.edu/courses/cs281/files/project.pdf>. There are four separate components of the project:

**Proposal, Due 25 October 2013** This is a two-page document that describes the problem you intend to solve, your approach to solving it and the experiments that you intend to run to evaluate your solution. The proposal represents 10% of your overall course grade.

**Abstract and Status Report, Due 22 November 2013** This is a one page document that contains a draft of your final abstract, as well as a brief status report on the progress of your project. This represents 5% of your overall course grade.

**Poster Session, 5 December 2013** You will make a conference-style poster about your project, with a class poster session during reading week. SEAS will pay for the cost of producing the poster. You will also submit a PDF file of your poster. This will represent 5% of your course grade.

**Final Report, Due 11 December 2013** You will write a report of up to ten pages, in the style of a mainstream CS conference paper. This will represent 20% of your overall grade.

You will upload these materials via the iSites dropbox. Please see the document referenced above for a more thorough description of the final project and policies related to collaboration, etc.

### Grading

Each Assignment (lowest dropped)	10%
Online Quizzes	5%
Midterm	15%
Final Project Proposal	10%
Final Project Abstract and Status Report	5%
Final Project Poster	5%
Final Project Report	20%

### Textbooks and References

The following book is required for the course:

**Machine Learning: A Probabilistic Perspective** Kevin P. Murphy, MIT Press, 2012. This is a very new book that covers a wide set of important topics. As the book is fresh and comprehensive, there are still quite a few errors. We will try to maintain lists of errata as they are discovered.

The following book is strongly recommended, but not required:

**Pattern Recognition and Machine Learning** Christopher M. Bishop, Springer, 2006. An excellent and affordable book on machine learning, with a Bayesian focus. It covers fewer topics than the Murphy book, but goes into greater depth on many of them and you may find that you prefer Bishop's exposition.

These are other (free online!) books on machine learning and related topics that you may find helpful, but that are completely optional.:

**Information Theory, Inference, and Learning Algorithms** David J.C. MacKay, Cambridge University Press, 2003. Freely available online at <http://www.inference.phy.cam.ac.uk/mackay/itila/>. A very well-written book with excellent explanations of many machine learning topics.

**Bayesian Reasoning and Machine Learning** David Barber, Cambridge University Press, 2012. Freely available online at <http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.Online>.

**The Elements of Statistical Learning** Trevor Hastie, Robert Tibshirani, and Jerome Friedman, Springer, 2009. Freely available online at <http://www-stat.stanford.edu/~tibs/ElemStatLearn/>.

These are books on some specialized topics that you may find useful:

**Gaussian Processes for Machine Learning** Carl Edward Rasmussen and Christopher K.I. Williams, MIT Press, 2006. Freely available online at <http://www.gaussianprocess.org/gpml/>.

**Non-Uniform Random Variate Generation** Luc Devroye, Springer-Verlag, 1986. Freely available online at <http://luc.devroye.org/rnbookindex.html>.

**Probabilistic Graphical Models: Principles and Techniques** Daphne Koller and Nir Friedman, MIT Press, 2009.

**Numerical Optimization** Jorge Nocedal and Stephen J. Wright, Springer, 2006.

**Bayesian Data Analysis** Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. CRC, 2013.

**Elements of Information Theory** Thomas M. Cover and Joy A. Thomas, Wiley, 1991.

**Monte Carlo Statistical Methods** Christian P. Robert and George Casella, Springer, 2005.

## Discussion Board

Most questions about the course, lecture or section material, or the assignments should be addressed via Piazza at <https://piazza.com/harvard/fall2013/cs281>. The course instructors will check this discussion board and make an effort to post responses within a day. Students taking the class are also encouraged to post responses. Code examples can be posted, but don't post anything you wouldn't be expected to share with other students in the class as per the collaboration policy. Long, detailed questions are probably best answered during office hours. Questions that are not appropriate for the discussion board may be sent to the staff via email. Use your judgement.

## Anticipated Schedule

**4 Sep 2013** Lecture 1: Introduction to Inference and Learning

**9 Sep 2013** Lecture 2: Simple Discrete Models

**11 Sep 2013** Lecture 3: Simple Gaussian Models

**12/13 Sep 2013** Section: Math Review

**16 Sep 2013** Lecture 4: Bayesian Statistics

**18 Sep 2013** Lecture 5: Linear Regression

**19/20 Sep 2013** Section: Practical Optimization; Assignment 1 Due

**23 Sep 2013** Lecture 6: Linear Classifiers

**25 Sep 2013** Lecture 7: Generalized Linear Models

**26/27 Sep 2013** Section: Undirected Graphical Models and Factor Graphs  
**30 Sep 2013** Lecture 8: Directed Graphical Models  
**2 Oct 2013** Lecture 9: Mixture Models  
**3/4 Oct 2013** Section: Factor Analysis and PCA; Assignment 2 Due  
**7 Oct 2013** Lecture 10: Sparse Linear Models  
**9 Oct 2013** Lecture 11: Exact Inference  
**10/11 Oct 2013** Section: The Junction Tree Algorithm  
**16 Oct 2013** Lecture 12: Variational Inference  
**17/18 Oct 2013** Section: Loopy Belief Propagation; Assignment 3 Due  
**21 Oct 2013** Lecture 13: Monte Carlo Basics  
**23 Oct 2013** In-Class Midterm  
**24/25 Oct 2013** Section: Particle Filtering; Final Project Proposals Due  
**28 Oct 2013** Lecture 14: Markov Chain Monte Carlo  
**30 Oct 2013** Lecture 15: Advanced MCMC  
**31 Oct / 1 Nov 2013** Section: MCMC Practicalities  
**4 Nov 2013** Lecture 16: Latent Dirichlet Allocation  
**6 Nov 2013** Lecture 17: State Space Models  
**7/8 Nov 2013** Section: Graph Models; Assignment 4 Due  
**11 Nov 2013** Lecture 18: Kernels  
**13 Nov 2013** Lecture 19: Gaussian Processes  
**14/15 Nov 2013** Section: Practical Gaussian Processes  
**18 Nov 2013** Lecture 20: Dirichlet Processes I  
**20 Nov 2013** Lecture 21: Dirichlet Processes II  
**21/22 Nov 2013** Section: Practical Dirichlet Processes; Assignment 5 Due; Final Project Abstract and Status Report Due  
**25 Nov 2013** Lecture 22: Boltzmann Machines  
**2 Dec 2013** Lecture 23: Neural Networks  
**4 Dec 2013** Lecture 24: Advanced Neural Networks  
**5 Dec 2013** Final Project Poster Session (Tentative)  
**11 Dec 2013** Final Project Reports Due