

COS424/SML302: Fundamentals of Machine Learning

Spring 2019

Course description

Problems about data abound. Here are some examples:

- Netflix collects ratings about movies from millions of its users. From these ratings, how can they predict which movies a user will like?
- JSTOR scans and runs OCR software on millions of scholarly articles. Scholars want to search and explore their collection. How should JSTOR organize it?
- A biologist has collected hundreds of thousands of measurements about the genotypes and traits of a large population. Can she make a hypothesis about which genotypes regulate which traits?
- Google sends and receives hundreds of millions of email messages each day. Are some of them spam? Which advertisements should they show next to each user?

Data analysis is central to many modern problems in science, industry, and culture. In science and engineering, it is essential to be fluent in solving modern data analysis problems. This class puts you on the path towards that fluency.

In this course, we will learn about a suite of tools in modern data analysis: when to use them, the assumptions they make about data, their capabilities, and their limitations. More importantly, we will learn about the language for and process of solving data analysis problems. On completing the course, you will be able to approach the analysis of large, complex data sets. In particular, you will be able to, given a data set, define the data analysis problem, learn about new methods, apply these methods to data, and understand the meaning of the results.

Administration

Lectures: Mondays and Wednesdays, 8:30AM-10:20AM
Friend Center 101

Precepts: Wednesdays, 2:30PM-3:20PM; Thursdays, 10:00AM-10:50AM; Fridays, 3:30PM-4:20PM
Computer Science 105

Instructor: Prof. Barbara Engelhardt

Office hours: Mondays 10:30AM-11:30AM in COS 302

Email: bee@princeton.edu

Teaching assistants:

Diana Cai, Jonathan Lu, Guillaume Martinet, Matthew Meyers, Archit Verma, Tianju Xue. See Piazza for contact information and office hours.

Piazza

We will use Piazza to host all communication.

1. Sign up for the Piazza site at piazza.com/princeton/spring2019/cos424sml302.
2. Use it to ask and answer questions about the course.
3. Use it to communicate with the instructors privately.
4. Use it to receive important announcements from the instructors.
5. Use it to download course notes, programming assignments, and reading materials.
6. Use it to access weekly quizzes.

Prerequisites

The prerequisite knowledge is calculus, linear algebra, computer programming (Python), and some exposure to probability and statistics. Contact Prof. Engelhardt if you have concerns about your prerequisite coursework.

Course programming

We will require the code for the data analysis homework assignments to be done in Python. Python has emerged as an easy and fast platform to develop many machine learning methods. In particular, the library SciKit-Learn has a large number of ML methods and approaches for use (including regression, classification, cross validation, etc.).

For visualization and downstream analysis of the results, R is a powerful open-source platform for statistical computing and visualization. You can download R for many platforms at <http://www.r-project.org/>.

To get started with R, see *Introductory Statistics with R* by Peter Daalgard. It is available as a PDF from the Princeton Library. There are a number of excellent packages for data visualization in R, such as `ggplot2`.

Writing with L^AT_EX

We will use L^AT_EX to write the homework assignments and the final project. We will post templates for the homework assignments and the final project on the website. To jointly edit a single L^AT_EX file among collaborators, consider using *Overleaf* or *Git* (all free and Princeton includes paid accounts for free too).

Course requirements

There are three kinds of work required for the course.

- Homework assignments. (60%) There are three homework assignments due throughout the semester at 11:59PM on the due date. These will all be the analysis of a specific data set, disseminated with the homework description, using methods discussed in class; the deliverables will be a five page write up of the data, analyses, one page of methods, and results (see Piazza page for the write-up template and an example write-up), and the Python code used to analyze the data. All homework assignments may be done alone or in pairs. If you choose to pair with another classmate, you may not pair with that same classmate for more than one of the homework assignments. Because of the nature of the team structure, any additional late days are given at the discretion of the professor, and these are very rare.

- Reading quizzes online. (10%) There will be weekly multiple choice question quizzes online about the reading that week. Each quiz will close at the start of class on Wednesdays. These will consist of a few short questions about the assigned reading material for the week. There are 12 weeks of class, and you are expected to complete 10 of these reading quizzes (i.e., you are excused from two of them with no penalty). There are no extensions and no late days. There is no extra credit for completing more than ten.
- Final project. (30%) The class project will be either a dramatic extension of one of the three homework projects in the course, or your own work on the development or application of machine learning methods to a large data set. You will turn in an eight-page write-up of your project on Dean's date on May 14th by 5pm; on May 13th, you will present your work at a poster session for the Princeton community. You may work alone on your project, but we encourage you to work in groups of up to four; you may pair with a classmate that you worked with on a previous assignment for the project.

Failure to complete any significant component of the course may result in a D or F.

Important Dates

- 6-Feb HW 1 out
- 27-Feb HW 1 due; HW 2 out
- 27-Mar HW 2 due; HW 3 out
- 17-Apr HW 3 due
- 22-Apr Project proposals due
- 13-May 9am-4pm Final project poster session
- 14-May Final project due (Dean's Date)

Grading and Collaboration Policies

As mentioned above, each homework assignment may be done alone or in pairs, although no more than two people may work on a homework assignment together. If individuals or groups discuss important details of their approach outside of these groups, please acknowledge this discussion in the Acknowledgements section of the homework. No pair of students may re-pair to work on a later homework assignment without the instructor's permission beforehand.

The final project may be done with zero to three partners (i.e., the largest teams may be of size four); we scale our expectations in the deliverables with the size of the groups. You are allowed to collaborate with people with whom you worked on a previous homework assignment.

The late day policy is that, across the three assignments, you may take one week of additional time to hand in one of the assignments. The week may not be split among assignments, but must be used on only one. Notice that the additional week will be taken on a specific assignment must be requested before the due date of that assignment. For every day that an assignment is late, 5% of the grade (out of 100%) will be taken off. When one person in a pair has used their late days, and another has not, we cannot extend the late days to both partners, and the partners may receive different grades on the assignment – this is something to consider when finding assignment partners. No late days will be given for any of the 12 reading quizzes – due every Wednesday during the semester before class – nor for the final project – due Dean's Date at 5pm. However, only 10 of the 12 quizzes are required – you may skip any two you choose with no penalty.

Note that we have software for examining the similarity of code and homework/project write ups, and we will be running these on everything that is handed in. If we find that there is substantial replication of

code or of text in the write up, we will give that assignment a 0 grade and send the assignment and evidence of the violation to the appropriate honor code committee.

If you feel your work has been graded unfairly, please bring the full set of graded work for the semester — not just the graded work in question — to Prof Engelhardt. I will work with you to make sure the grades are fair and appropriate, but, to be clear, this does not always mean that the grades will improve. Often problems in work that the graders missed will be identified as well as the opposite.

Syllabus and Readings

Most readings come from: Murphy, K. *Machine Learning: A Probabilistic Approach*. MIT, in press. (MLAPA); the e-book is posted online on Blackboard. Additional resources for material will be posted on Slack, as well as the slides from lecture.

- Hastie, T., Tibshirani, R. and Freedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd Edition, Springer, 2009. (ESL)
- Bishop, C. *Pattern Recognition and Machine Learning*. Springer-Verlag, 2006. (PRML)

Readings below are tentative. Please see Piazza for any changes to the reading.

Lecture	Date	Subject	Reading
01	M 04 Feb	Introduction	MLAPA Ch 1
02	W 06 Feb	Probability and statistics review	MLAPA Ch 2; [Opt] MLAPA Ch 3.1-3.4
03	M 11 Feb	Graphical models	MLAPA Ch 10.1-10.2, 10.4
04	W 13 Feb	Probabilistic classification	MLAPA Ch 3.5
05	M 18 Feb	Features and kernels	MLAPA 14.1-14.2
06	W 20 Feb	Kernel classifiers	MLAPA 14.3-14.5
07	M 25 Feb	Linear regression	MLAPA 7.1-7.3; [Opt] ESL Ch 3.1-3.2
08	W 27 Feb	Regularized linear regression	MLAPA 7.5.1,7.6.1,7.6.2; [Opt] ESL Ch 3.4
09	M 04 Mar	Logistic regression	MLAPA 8.1-8.2
10	W 06 Mar	Generalized linear models	MLAPA 9.1-9.3.2; [Opt] McCullagh and Nelder, Ch 2
11	M 11 Mar	Snow day	
12	W 13 Mar	K-Means	MLAPA 11.1-11.3
	M 18 Mar	Spring break	
	W 20 Mar	Spring break	
13	M 25 Mar	Mixture models	
14	W 27 Mar	Expectation-maximization	MLAPA 11.4-11.6
15	M 01 Apr	Hidden Markov models	MLAPA 17.1-17.2
16	W 03 Apr	Dimension reduction and PCA	MLAPA Ch 12.1-12.2
17	M 08 Apr	Factor analysis	
18	W 10 Apr	Probabilistic topic models	Blei (2011)
19	M 15 Apr	Dirichlet processes	
20	W 17 Apr	Communities in networks	Airoldi et al. (2008)
21	M 22 Apr	Gaussian process regression	Roberts et al. (2013)
22	W 24 Apr	Markov chain Monte Carlo	MLAPA 23 (optional), 24.1-24.3.5
23	M 29 Apr	Scalable machine learning	MLAPA 21.1-21.5 (not 21.4)
24	W 01 May	Summary and discussion	
	M 13 May	Poster session	
	T 14 May	Dean's Date	Projects due (5pm)