

Syllabus for 46-926: Statistical Machine Learning I

Spring 2018

Description: This is the first in a two-part sequence on statistical machine learning. This first course covers tools and approaches for prediction, including both regression and classification. The focus is on understanding the foundations of the methods so that they can be both applied and modified. Topics include foundations of supervised learning, regularized and nonparametric regression, bias-variance tradeoffs, model validation and assessment, classification, and tree-based methods. The second course will expand into advanced topics, including boosting and ensemble methods, support vector machines, mixture models and topic modeling, natural language processing, markov decision processes and reinforcement learning, and neural networks. There may be some movement of topics between the two courses as time and interest permits.

Professor: Max G'Sell (sounds like "gazelle")
Email: mgsell@cmu.edu. Please post questions on Piazza instead of email if possible.
Office: Baker Hall 132B
Office hours: Sundays at 3pm on Canvas. Additional office hours on Friday mornings are being considered.

Calendar: Mondays 5:30–8:30 PM.

Lecture 1: January 15 - Pittsburgh
Lecture 2: January 22 - Pittsburgh
Lecture 3: January 29 - Pittsburgh
Lecture 4: February 5 - NYC
Lecture 5: February 12 - Pittsburgh
Lecture 6: February 19 - Pittsburgh
Lecture 7: February 26 - NYC

Final Exam: Wednesday, March 7th at 3:00 PM.

Text: The following three texts are good references for the course. Note that the first two are available for free online from the authors, and the third is available for free download from SpringerLink when you are on the CMU network.

Elements of Statistical Learning, by Trevor Hastie, Robert Tibshirani, and Jerome Friedman. Springer. ISBN: 978-0-387-84858-7

Introduction to Statistical Learning, by Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. Springer. ISBN: 978-1-4614-7138-7

Applied Predictive Modeling by Max KuhnKjell Johnson. Springer. ISBN: 978-1-4614-6848-6

Other references for special topics will be given in class.

TAs: Ben LeRoy
Kwhangho Kim
Minshi Pen

TA office hours will be held on Canvas at 3:00pm on Saturdays.

If ten minutes pass without any questions, then the TA session may end.

Course Grade: You are responsible for checking your scores regularly on Canvas, and reporting any errors or discrepancies to us.

Homework and a final exam will contribute to your final course grade. Your overall homework average will count for 50%, and the final exam will count for the other 50%.

Computing: This course will rely primarily on python for computation. We will fall back on R for topics which are not well-developed in current python packages. The lectures and homework will assume the general proficiency at programming, simulation, and R that you have developed in your previous courses, but will assume no familiarity with python. As in learning any new programming language, some exploration and internet searching on your part is assumed.

I would highly recommend using jupyter notebooks for workingg with python in this course, including for writing up homework assignments.

Homework: Homework will generally be due at 5:20 PM on canvas the days of lecture. Homework will usually be available the Tuesday evening after lecture. Because we often discuss homework in detail at the beginning of lecture, late homeworks will not be accepted without prior arrangement.

The homeworks are a key component of this course, and a large amount of learning will occur during the completion of the exercises. There will be instances where the homework is used to extend material that was taught during the lecture. There will be cases where important results will be presented in the homework. These new results should be considered as being equally important to results presented during lecture.

Further, the difficulty level of the homework exercises is, on average, greater than that of the examples done during lecture. It is not a wise use of our limited lecture time to work through many extended examples, and watching me work an exercise is not a replacement for doing the work yourself.

Email: Please do not email TAs with concerns over grading, or any other issues with the course. Instead, email the instructor.

I assume that you check, at least daily, the email account which is linked up to Canvas (in other words, your @andrew.cmu.edu address). Email can be used for important announcements, so make sure that you are regularly reading mail sent this account.

Please post any questions regarding homework or the lecture to Piazza. We will do our best to be helpful and respond promptly. *However*, you should be reasonable about expectations of responses to questions on the due dates of homeworks. The time leading up a lecture tends to be the busiest time of the week for me. This is especially true if I am traveling to New York that day.

Final Exam: The final exam will be on Wednesday, March 7th at 3:00 PM

The final exam will be three hours in length. You will be allowed to bring one sheet of **handwritten** notes, which you have written yourself. (Both sides of an 8.5 by 11 inch sheet of paper.) **You will not be allowed to use a calculator or computer.**

If you understand everything that was covered in the lectures, and everything that appeared on the homeworks, then you are completely ready for the final exam. There is material in the texts that will not be covered in lecture or the homework; you are not responsible for this material on the final. There will be material that is covered in lecture or in the homework, but does not appear in any of the texts; you are responsible for this material on the final.

Lectures:

I expect you to attend all lectures. Also, I assume you hear everything I say during lecture. In other words, if I have to make an important announcement, and I do so during normal lecture time, I am not obligated to deliver the announcement via any other means. **In particular, I often make important announcements at the beginning of lecture, and you are responsible for being aware of these. These announcements are often on a slide titled “Announcements.”**

The lecture notes will be delivered via the document camera or tablet. These will be scanned into a .pdf file and posted on the Canvas site the day following the lecture. If needed, I will clarify, expand, and/or correct the notes prior to posting them.

Prior to lecture, I will post a version of the lecture slides with blanks to be filled in. You can print a copy of this prior to lecture to follow along. **These may not be the entirety of what we will cover during lecture.**

You are encouraged to ask questions at any time during the lectures. If I feel that your question ventures into technical details which are not pertinent to the discussion, I may ask you to either talk after class, come to office hours, or send me an email. You should not interpret this as me discouraging these (or any other) questions.

Integrity:

Please read CMU’s policy on cheating and plagiarism:
<http://www.cmu.edu/policies/documents/Cheating.html>

You are not allowed to utilize any materials from past versions of this course. Any such use will be considered cheating, and treated as such.

Collaboration is a tricky issue. The support and assistance of your classmates can go a long way towards helping you to understand the material. Ultimately, however, you are responsible for preparing yourself for the final exam, later courses, and your future. I encourage you to collaborate on understanding the material and your homework. However, everything you submit must be your own work. Your final writeup should be done independently and you must write your own code for computational problems. You are responsible for understanding everything that you submit. **Please ask me if you have any confusion.** Instances where students copy the work of another student will be treated as cheating.