

Statistics 471/571/701: Modern Data Mining

Linda Zhao

Fall 2017

E-mail: lzhao@wharton.upenn.edu

Office: JMHH 470

Class Room: Vance B11

Web: canvas.upenn.edu

Office Hours: F 3:00 - 5:00 pm, or by appointment

Class Hours: 401: MW 10:30 - 11:50am; 402: MW 3:00 - 4:20pm

The Course

Course Description

Statistics has been evolving rapidly to keep up with the modern world, especially with computational methods for the explosion of data. As a significant part of data science we start the class with exploratory data analysis (EDA). We then show how to build, interpret, and adapt simple models; then go beyond with newer contemporary methods and techniques for handling large and complex data with applications in finance, marketing, medical fields, social science, entertainment, you name it. While this course makes extensive use of the statistical programming language R, no programming experience is required. By the end of the semester we hope that students have not only learned the modern statistical methods but have also become skilled in dealing with data of essentially any size.

This class is cross-listed as STAT 471 for undergraduates, STAT 571 as a graduate level course for students outside of the statistics department, and STAT 701 for MBA's.

Prerequisites

Two semesters of statistics courses, familiarity with multiple regressions is assumed. A quick review of Chapter 3 from the main textbook (ISLR) would be very helpful.

Methods covered

- Exploratory Data Analysis (EDA)
- Multiple Regression/Stepwise Regression
- Logistic Regression/Multi-Nomial regression
- K-nearest neighbors (KNN)
- Linear Discriminant Analysis (LDA), Quadratic Discriminant Analysis (QDA)
- Penalized regression: LASSO, Ridge Regression, Elastic Net
- Principal Components Analysis (PCA)
- Tree based methods such as Boosting, Random Forest

- Support Vector Machines.
- Text mining sentiment analysis
- Bootstrap and k-fold cross validation
- Training and Testing errors
- ROC/AUC and FDR
- Neural network/Deep learning
- Unsupervised learning

Case study/Datasets

Most of the following cases will be covered:

- Who tweets for Trump?
- Wharton Business Radio Audience Estimation via Amazon Mturk
- Diabetes/Healthcare (Predicting Readmission Probability for Diabetes Inpatients to Save Healthcare Cost)
- IQ=Success?
- Boost return by 80% in Lending Club?
- Handwriting recognition (image recognition)
- Can we do something to reduce crime rates?
- Framingham heart disease study
- Billion dollar Billy Beane
- What can we do to improve education – Texas third graders?
- Whose political bill is more likely to be approved in the sea of bills proposed by politicians?
- Can you predict housing prices?
- McGill Billboard – how long a song can sit on the board?
- Out of 502 stocks can we do better than S&P500?
- How to be successful at Kickstarter
- Hunting for important gene expression positions to help out with HIV positive patients
- Using Yelp reviews to predict the rating (text mining)
- Chinese Annual Industrial Survey

And more!

Required Materials

Software

The free and open source [statistical computing language R](#) is used. There are infinitely many new packages available for us to use; a [pretty interface to explore the publicly available R packages](#) is available via Microsoft.

It is *highly* recommended to use the free [RStudio](#), an interface for writing R documents and working with data.

R tutorial

The TA's will run an introduction to R during the first week of class. Bring a laptop to JMHH 375.

- Time: TBA
- Time: TBA

Textbook

Our required textbook is known as ISLR and is [freely available from the authors](#):

- Gareth James, Daniela Witten, Trevor Hastie and Robert Tibshirani, *An Introduction to Statistical Learning with Application in R*, First Edition, 2013, Springer New York.

Additional optional readings are recommended for getting familiar with and working with R. These include:

- Garrett Golemund & Hadley Wickham, *R for Data Science*, 2016, O'Reilly. Available [freely online](#).
- The R Core Team, *An Introduction to R*, available from [CRAN](#).
- Patrick Burns, *The R Inferno*, available [online](#).
- Peter Dalgaard, *Introductory Statistics with R*, Second Edition, 2008, Springer. Available on [Academia](#).
- Trevor Hastie, Robert Tibshirani, Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition, 2008, Springer. A PDF version of ESLR is available [from the authors](#).

Course Policies

Communication

Communication will be through [Canvas](#) and through [Piazza](#). Files will be uploaded to Canvas, including datasets, homeworks, and lecture notes. Piazza is a useful forum for students to ask questions.

Laptop Policy

A laptop is a must for the course. You are encouraged to bring the laptop to the classes so that you may run the lecture code simultaneously with the professor. However, it is not allowed to use the laptop for other purposes during the lectures. Cell phones must be turned off.

Assignments and Exams

Homework: We will give 4 or 5 homework assignments. These may be done in groups of up to 3 people; see the Group Policy for more details.

Mini project: This will be a take-home, *individual*, assignment. You will be given a dataset to analyze, step by step. There is a limit of **10 pages**. Further information will be given.

Quizzes: We will give 4 quizzes, which will be short, in-class, individual, quizzes. The first three will be 10 minutes each, and the lowest grade among these will be dropped. The final quiz will be 40 minutes long and cannot be dropped. *No makeup quizzes will be offered.*

1. **Quiz 1: TBA**
2. **Quiz 2: TBA**
3. **Quiz 3: TBA**
4. **Quiz 4: Mon 12/11 in class**

Midterm: This exam will be *individual*, open-book and done on the computer. You will be given an exam to work through, using R, to do under proctored conditions. You must bring a laptop

The midterm will be held on **Tuesday, November 7, from 6:00-8:00pm**. Place: TBA.

Final Project:: The ultimate goal of the class is to prepare/expose students to techniques that are suitable for modern data. The final project is designed so that each of you will bring a problem of personal interest to the class. You will need to identify a problem to tackle with a data set that either you collect/extract or find.

This project is done within a group of up to 3 members. A complete write up is required. This would be a good project to put in your CV if desired.

- Proposals can be submitted any time during the semester but no later than *TBA*
- A well-motivated, relevant topic is most desirable.
- Originality, complexity, and challenge will be another plus
- A complete write up is a must.
- **Maximum of 15 pages.**

Late Work Policy

It is imperative that you manage your workload properly for this course. We will allow late assignments up to 3 days late, with a 15% penalty per day. Note that lateness will be determined by timestamp on Canvas submissions, i.e. 12:01 AM is considered late.

Note also the above quiz policy: we will drop your lowest score of the first 3 quizzes, but there will not be any makeup quizzes.

Group Policy

The homework and the final project can be done by groups of up to three people. Sign up for groups on Canvas as soon as possible but no later than **Sunday, 09/10**. We will help out for those who need to find a group, with searches on Piazza.

Please note that at no time may a group have more than 3 members. In addition, while those within a group will submit a single homework file for the group, students must follow the code of academic integrity in regards to classmates outside their group. Finally, students do not have to complete the final project in the same group as for homework. They may form a new group though again no more than 3 people may be in a group. We prefer you keep the same groups through the semester but it is now required.

Grading Policy

- Homework: 20%
- Quizzes: 15%

- “Mini Project” (take-home): 20%
- Midterm Exam: 25%
- Final Project: 20%

Professor Zhao may make adjustments for those who actively contribute to the class throughout the semester.

Teaching Assistants

- Jeff Junhui Cai, junhui@wharton.upenn.edu
- John Brooke Waldt jbrooke@seas.upenn.edu
- Juan Marin jmarin@wharton.upenn.edu
- One more

R Markdown Advice

Make sure to include your Rmd file and a knitted file, in your submission.

1. Load your packages only once, and at the top of the document, in the setup chunk.
2. Use `library(pkg_name)` to load packages. Do not use `require()` - see [Hadley's advice](#).
3. You can use `command-option-i` to insert a new R chunk.
4. Most of the time, you will want to use `stringsAsFactors = FALSE` when reading csv files.
5. Do not include `install.packages("...")` in your code.
6. Any time R prints output, make sure to include text explaining what is being printed and why.

Class Schedule

Tentative and subject to change. Unless otherwise noted, readings refer to *Introduction to Statistical Learning*.