

MS&E 226: “Small” Data

Time and place

Lectures: Mondays, Wednesday, and Fridays, 10:30-11:50 AM, Skilling Auditorium

Discussion sections: Fridays, 3:30-4:20 PM, Skilling Auditorium

Instructor

Ramesh Johari, Associate Professor, MS&E, CS (by courtesy), EE (by courtesy)

Huang Engineering Center, Room 311

E-mail: rjohari@stanford.edu

Admin: Jim Fabry, jim.fabry@stanford.edu

Office hours: Tuesdays 3:45-5:15 PM, Spilker 143 (except on November 6th)

Additional office hours by appointment only

Teaching assistants

Je-ok Choi (co-Head TA)

E-mail: jchoi89@stanford.edu

Office hours: Fri 1-2pm, Huang B020

Lin Fan (co-Head TA)

E-mail: linfan@stanford.edu

Office hours: Wed 5:30-7:30pm, GESB 131

Nikhil Garg

E-mail: nkgarg@stanford.edu

Office hours: Mon 4-6pm, GESB 131

Mel Tutar

E-mail: mtutar@stanford.edu

Office hours: Wed 2:30-4:30pm, McCullough 126

Allison Park

E-mail: ajypark@stanford.edu

Office hours: Tue 6-8pm, Huang B016

Course websites

Public (lecture notes, resources): <http://web.stanford.edu/class/msande226>

Calendar (lectures, discussions, office hours): <https://tinyurl.com/2018-msande226-cal>

Piazza (communication, announcements, materials): piazza.com/stanford/fall2018/mse226

Gradescope (grading, problem set submission): <https://gradescope.com/> Entry code: **97G3N8**

Lectures

There will be 23 lectures for the course this year. **Note: We will not be meeting every Monday, Wednesday, and Friday.** In particular, we will have lecture on the following dates:

9/24, 9/26, 9/28

10/1, 10/3, 10/8, 10/10, 10/15, 10/17, 10/19, 10/22, 10/24, 10/29, 10/31 (exam)

11/2, 11/9, 11/12, 11/14, 11/26, 11/28, 11/30

12/3, 12/5, 12/7

For your convenience, the exact schedule of lectures, discussion sections, and office hours is available at the following Google calendar:

<https://tinyurl.com/2018-msande226-cal>

Lectures are recorded and available through SCPD at: <https://mvideox.stanford.edu/>

Catalog course description

This course is about understanding "small data": these are datasets that allow interaction, visualization, exploration, and analysis on a local machine. The material provides an introduction to applied data analysis, with an emphasis on providing a conceptual framework for thinking about data from both statistical and machine learning perspectives. Topics will be drawn from the following list, depending on time constraints and class interest: approaches to data analysis: statistics (frequentist, Bayesian) and machine learning; binary classification; regression; bootstrapping; causal inference and experimental design; multiple hypothesis testing. Class lectures will be supplemented by data-driven problem sets and a project. Prerequisites: CME 100 or MATH 51; 120, 220 or STATS 116; experience with R at the level of CME/ STATS 195 or equivalent.

Detailed course description

Before figuring out what a "small" data class is about, it's important to understand what "big data" means.

It's hard to walk down the street these days (at least at Stanford) without hearing the phrase "big data". A casual Google search for headlines including "big data" turns up the following:

- Big oil turns to Big Data as oil prices plummet
- Can Big Data help us fight rising suicide rates?
- When accounting meets Big Data
- Big Data: Searching for drug side effects
- Using Big Data, IoT to track & control building energy efficiency

Informally, the phrase "big data" is often used to refer to the information explosion that is providing data across every walk of human endeavor at finer granularity than ever possible

before. More formally, “big data” refers to data that *cannot be stored or analyzed on a single machine*, and instead requires a cluster of machines (or more) to be processed and studied.

Big data has prompted a revolution in algorithms for managing, storing, and operating on data: MapReduce, Hadoop, Spark, etc. These are all household names (at least in the technology industry) to make working with massive datasets easier.

But there are two points that should give us pause in the rush to big data. First, even as datasets grow ever larger, computational technology is also advancing -- so the size of a dataset that can be processed on a single machine is becoming much larger than before (for example, it is now common to find 64GB of RAM on a typical household PC, and much more on enterprise machines). This is formally what we mean when we say “small data”: in contrast to big data, it is any dataset that can be stored, processed, analyzed, and visualized on a single machine, typically interactively.

Second, and more importantly, often the most important part of the analytical pipeline in understanding a large dataset is actually understanding small datasets first. It is in the realm of small data where we first understand techniques for summarizing and exploring data, for drawing inferences from data, for making predictions, and ultimately for making decisions. This course is about building that foundation.

Our approach is built on the old cliché: “Give a man a fish and you feed him for a day; teach a man to fish and you feed him for a lifetime.” This is not a vocational course: you will not learn a large number of different tools for analyzing data, and we will not be spending a lot of time on the details of good data processing practice. Among other things, this means (if you intend to be a good data scientist) you should also find ways to learn SQL, R, Python (or your favorite language), and a range of different algorithmic techniques for analyzing data. There are plenty of ways to do this now, both through other courses at Stanford and through a large number of resources online. Data processing tools change fast, and more than anything else you will need to *learn how to learn*.

With that in mind the primary goal of this course is to give you a footing in which to ask critical questions about different methods you will encounter over a lifetime of working with data. We will teach you: how to be careful in defining your objective; how to compare and contrast different approaches to the same problem; the differences and similarities between frequentist statistics, Bayesian statistics, and machine learning approaches; checking model assumptions; and how to use data to draw causal inferences. We use linear regression as a playground for our study, but a successful student in this class should be able to use the insights in a class like this to become a more educated consumer of an entire “toolbox” of data analysis methods.

Course outline

1. *Summarization* (2 weeks). Given a single data set, how do we summarize it? Basic sample statistics. Using models to succinctly summarize data. The algebra of linear regression and logistic regression. In-sample measures of fit: R^2 and residuals.
2. *Prediction* (2-3 weeks). How do we generalize our understanding of a data set to *new* samples? Formalizing the prediction problem. Binary classification. Linear regression and logistic regression as approaches to prediction. Model complexity and the bias-variance tradeoff. Training vs. test sets and cross validation.
3. *Inference* (2-3 weeks). How do we generalize our understanding of a data set to draw inferences about the population or system from which the data came? The basics of frequentist estimation and hypothesis testing. Application to linear regression. The bootstrap. The multiple hypothesis testing problem. Comparison to Bayesian estimation and hypothesis testing.
4. *Causality* (2 weeks). How do we determine the effect that *changing* a system will have? The Rubin causal model, potential outcomes, and counterfactuals. The “gold standard”: randomized experiments. The basics of causal inference from observational data. From causal inference to data-driven decisions.

Prerequisites

This is a quantitative, mathematical, and computational course. Accordingly, these are the prerequisites.

- (1) *Math 51*. A hard constraint; do not enroll if you have not had multivariable calculus and linear algebra. I will **not** be recapping basic concepts in class, so this should be technology you have at your fingertips.
- (2) *Probability at the level of MS&E 120/220 or Stats 116*. Also a hard constraint, just like Math 51. We will not be reviewing probability in this class.
- (3) *Computational methodology*. You will need to be able to work with datasets. While you don't need to be a master, some familiarity (going into the class) with one of R, Matlab, or Python is important. The most significant computational work in the class will happen in the second half. Therefore, concurrent enrollment in Stats/CME 195 is an acceptable way to fulfill this requirement.
- (4) *Quantitative curiosity*. This class is about learning how to think critically about quantitative methods. Therefore getting the most out of the class requires “opening the hood”, and being curious about why things work the way they do. If you are uncomfortable thinking quantitatively, this is not the class for you.

Evaluation

Your evaluation will be based on a combination of five problem sets, a midterm exam (including both an in-class and a take-home component), an in-class final exam, and a guided mini-project that will run through the final three problem sets. The grade will be determined as follows:

- Problem sets: 30%

- In-class component of the midterm: 15%
- Take-home component of the midterm: 15%
- In-class final exam: 15%
- Mini-project: 25%

Note: I award A+'s in the course to those students who receive A's and also demonstrate exceptional performance on the project.

Problem sets

There will be a total of 5 problem sets. Problem sets must be submitted online through Gradescope (<https://gradescope.com/>; entry code **97G3N8**). *Note that Gradescope will only be used for homework submission and grading; all announcements and discussions will be handled through Piazza (above).* Problem set sheets will also be posted on Piazza (in the resources section). Problem sets will generally be due at 5 PM (except Problem Set 1, which is due at 10 AM on October 1).

Except for medical necessity, no late problem sets will be accepted. All assignments will be posted to the course website. Problem sets are assigned and due as follows:

- Problem set 1: Handed out on 9/24, due on 10/1
- Problem set 2: Handed out on 9/27, due on 10/11
- Problem set 3: Handed out on 10/11, due on 10/25
- Problem set 4: Handed out on 11/2, due on 11/15
- Problem set 5: Handed out on 11/15, due on 11/29

Depending on their length, the total number of points in each set might vary. Each part of each problem will be graded as follows:

- You will receive zero points if you do not attempt it.
- If you attempt it, but there are either substantial methodological errors or major conceptual misunderstandings of the material, you will receive 2 points.
- If you attempt it and there are no substantial methodological errors or major conceptual misunderstandings of the material, you will receive 3 points.

We expect students who make a reasonable effort, even if not perfect, will receive 3 points.

Receiving 2 points is intended to be a sign of significant comprehension issues and should only happen in cases where you haven't been able to keep up with the material. Since you can get full credit without having been perfect, you should make sure to read the solutions to aid your comprehension of the material.

You can discuss the assignments among yourselves, but everybody must turn in his/her own written solutions in his/her own words. If you do a substantial subset of the work on your problem set with others, document on each assignment the other students that you worked with.

Mini-project

Starting with Problem Set 3, we will guide you through an applied mini-project, learning to apply some of the methods in this class in greater depth. You will select a dataset to work with, and use the methods from the course to understand the data. We will provide more details of the structure of this project early in the quarter. You will also peer review projects of two other teams.

For the mini-project, you can work in pairs. Only one copy of your project needs to be turned in per pair.

Because we assign teaching assistants to projects, it is important to register your project team (solo or teams of 2) on time. The deadline to register is **October 18 at 5 PM**.

You can register here: <https://tinyurl.com/2018-msande226-project>

Project deadlines are as follows:

- Team registration: Due on 10/18
- Project part 1: Handed out on 10/11, due on 10/25
- Peer review 1: Handed out on 11/2, due on 11/9
- Project part 2: Handed out on 11/2, due on 11/15
- Peer review 2: Handed out on 11/15, due on 11/27
- Project part 3: Handed out on 11/15, due on 12/6

Submitting on Gradescope

As noted above, problem sets and project submissions are submitted and graded through Gradescope. To ensure this process is smooth, there are a few things to keep in mind:

- You are required to tag your answers correctly. The graders will ignore any part of your solution that is not tagged. Note that this means you also have to correctly tag your code. Allow enough time prior to submission to ensure you are able to tag correctly.
- In order to grade code you submitted, we need to be able to copy your code. Make sure this is possible (e.g., do not upload screenshot images of your code). We will deduct points if we cannot check your code.
- If you believe we have made a mistake grading your work, you should submit a regrade request through Gradescope. This sends your request directly to the grader on that particular question. *You must submit your regrade request within 14 days of the grades of that particular problem set or project part being published.*
- If you decide to do the mini-project in a team of two (see below), only one team member should submit the project to Gradescope and then add the other team member to the submission.

Course communications: Piazza

As noted above, we will use Piazza to manage course announcements and a discussion forum. You can sign up for the course on Piazza here:

piazza.com/stanford/fall2018/mse226

Please use Piazza for all course-related communication with us. We will aim to respond to questions in a 24-48 hour period, except of course for those of an urgent nature (e.g., typos on problem sets or lecture notes, clarifying course logistics, etc.). Among other things this means you should not wait until the last day before a problem set is due to message us; we will likely not respond in time.

You are encouraged to attend office hours to ask questions of a technical nature; these are best discussed in a face-to-face setting, given the quantitative nature of the course.

If you are having difficulty, find help right away— *do not wait until you fall even further behind!* There is an obvious temptation to wait until the day before the due date to do all the work on the problem sets, and I can assure you this approach will almost certainly lead to very poor performance in the class.

Student participation on Piazza

We would love to have students help answer Piazza questions -- particularly in the “waiting period” before course staff answers. Thus we are providing an incentive for you to participate in Q&A! Every time you answer a question that is marked “good answer” by an instructor, you will receive one lottery ticket. At the end of the quarter these tickets will be entered into a lottery for 4 \$50 Amazon gift cards. At most one gift card can be won by a single student.

Exam policy

The midterm exam will be held in two parts. On October 31, there will be an in-class multiple choice exam. In addition, on October 31 at the end of class we will post a take-home component, that will be due back by the beginning of class on November 2 (submission via Gradescope). The midterm will cover all topics from the first part of the class.

The final exam will be held at the time set by the registrar, from 8:30 AM - 11:30 AM on December 13. It will be a multiple choice exam, of similar length to the midterm. The final exam will cover all topics from the second half of the class.

SCPD students who do not come to campus will take the exams remotely through exam monitors. For all other students, except for medical necessity, *there will be no alternate exam dates or times. You should only register for the class if you are certain you can take the exams on these dates.*

Discussion sections

We are planning to schedule weekly discussion sections; location and time TBA. This will be an opportunity for students to get additional help on the material, in a more guided session. We will use these sessions to go over material from the lectures, and to walk through problems that will help you solve the problem sets.

Computation

Many of the problem sets, as well as the guided mini-project, will require you to be comfortable carrying out computations on data. For this purpose, the “official” language of the course is R; if you want to develop facility with R, it is sufficient to enroll concurrently in Stats/CME 195.

We will provide links to R on the website, as well as any datasets that are needed through the course of the quarter.

As noted above, *you are welcome to work in pairs on any computational component in the class, and on the mini-project.*

Suggested textbooks

I will distribute notes as the course goes on. In addition, you may find it helpful to have the following textbooks on hand. I’m not requiring them, because they are available online, and we will not be linearly working our way through any of them. At the same time, they are not particularly expensive (by textbook standards). Some problems will also be drawn from these books.

1. *All of Statistics*, by Larry Wasserman. This book is as ambitious in scope as the title suggests. The obvious tradeoff for this breadth is a lack of depth. In particular, this book provides less of the nuances of how data analysis can easily go wrong; on the other hand, it is one of the few places where you can find statistics and machine learning treated together (in a somewhat accessible way).

The book can be accessed online (for Stanford students) here:

<http://link.springer.com/book/10.1007/978-0-387-21736-9>

It is available through Amazon here:

<http://www.amazon.com/All-Statistics-Statistical-Inference-Springer/dp/0387402721>

(Though note that you can buy the “MyCopy” softcover edition for less through Springer at the first link above.)

2. *Statistical Models*, by David Freedman. This book is a great complement to the Wasserman book. Although targeted nominally at statistical methods for the social sciences, it is unique in its healthy skepticism of statistical methods applied without regard to modeling assumptions. Freedman was well known for his advocacy of “shoe leather” statistics: in his view good data analysis required a substantial investment of on-the-ground effort to understand the context of the data you were analyzing. As a result his book spends a lot of time asking the reader to critically evaluate the methods

that are being taught.

This book can be accessed online (for Stanford students) here:

<http://ebooks.cambridge.org/ebook.jsf?bid=CBO9780511815867>

It is available through Amazon here:

<http://www.amazon.com/Statistical-Models-Practice-David-Freedman/dp/0521743850>

3. *Data Analysis Using Regression and Multilevel/Hierarchical Models*, by Andrew Gelman and Jennifer Hill. Though primarily about methods for dealing with multilevel models (e.g., models with random effects, fixed effects, and mixed effects), this book is also one that provides great insight into how to think about data more generally. I will be drawing on the first two parts of this book to help describe various modeling and interpretation issues in applying linear regression. Unfortunately this book is not available online; here is the book site:

<http://www.stat.columbia.edu/~gelman/arm/>

It is available through Amazon here:

<http://www.amazon.com/Analysis-Regression-Multilevel-Hierarchical-Models/dp/052168689X>

Other references

Here is an assortment of other books that you may find useful to consult.

1. [Computer Age Statistical Inference](#), by Efron and Hastie. This brand new text covers a range of topics relevant to statistical inference in the age of big data. It is notable for being written by two of the most well-known living statisticians, and for the breadth of scope. Highly recommended.
2. [Elements of Statistical Learning](#), by Hastie, Tibshirani, and Friedman. This is a comprehensive reference on prediction models for classification and regression.
3. [An Introduction to Statistical Learning with Applications in R](#), by James, Witten, Hastie, and Tibshirani. An introductory version of the preceding book, with a greater emphasis on hands-on practice; the text for Stats 202.
4. [Bayesian Data Analysis](#), by Gelman, Carlin, Stern, Dunson, Vehtari, Rubin. A great reference on applied Bayesian statistics.
5. [Statistics for Experimenters](#), by Box, Hunter, and Hunter. A classic reference on experimental design.
6. [Causal Inference for Statistics, Social, and Biomedical Sciences](#), by Imbens and Rubin. A recent and comprehensive book on methods for causal inference from data, especially with deep coverage of methods for observational data analysis.
7. [A Guide to Econometrics](#), by Kennedy. A good companion to the Imbens and Rubin text, Kennedy's book helps you understand common ways in which causal inference and econometric analysis can go awry.
8. [Mostly Harmless Econometrics](#), by Angrist and Pischke. A succinct and accessible treatment of the basics of causal inference.

9. Theory of Point Estimation, Lehmann and Casella; and Testing Statistical Hypotheses, by Lehmann and Romano. These are advanced theoretical texts for the foundations of estimation and hypothesis testing (often used for Stats 300A).