17.835: Machine Learning and Data Science in Politics

Spring 2018

Instructor: Professor In Song Kim
TAs: Tugba Bozcaga, Nicolas Dumas, Andreas Wiedemann

Department of Political Science
MIT

1 Contact Information

	In Song Kim	Nicolas Dumas	Andreas Wiedemann	Tugba Bozcaga
Office:	E53–407	E53–458	E53–448	E53–408
Email:	insong@mit.edu	ndumas@mit.edu	awiedem@mit.edu	bozcaga@mit.edu
Office Hours:	Thurs 4pm-6pm	Mon 11:30am-12:30pm	Wed 11:30am-12:30pm	Fri 11.30am-12.30pm

2 Logistics

- Lectures: Mondays and Wednesdays 9:30–11:00am, 3-133
- Recitations: Thursdays 1-2 pm, 1-242; Thursdays 4-5 pm, 1-379; Fridays 10-11 am, 1-379
  - There will be three weekly recitation sessions. We will cover a review of the theoretical material and also provide help with computing issues. The teaching assistants will run the sessions and can give more details. Attendance is strongly encouraged.

3 Course Description

Empirical studies in political science is entering a new era of "Big Data" where a diverse range of data sources have become available to researchers. Examples include network data from political campaigns, data from social media generated by individuals, campaign contribution and lobbying expenditure made by firms and individuals, and massive amount of international trade flows data. How can we take advantage of these new data sources and improve our understanding of politics? This course introduces various machine learning methods and their applications in political science research. Students will

1. Be introduced to various quantitative political science research topics in its four subfields: American Politics, International Relations, Comparative Politics, and Political Methodology.
2. Learn basic machine learning algorithms and data science tools that are applied in political science research

3. Apply data analysis tools using R programming language through problem sets.

4. Collect and analyze data to learn substantive topics of own interest.

5. Learn how to communicate data-driven findings and insights.

Note: the topics covered in this class represent only a very small subset of political science research. If you enjoy this class, please consider a HASS concentration in Political Science. We also offer a major and a minor in Political Science, as well as a minor in Public Policy and a minor in Applied International Studies. Internships and research opportunities too. Check out these programs and more at: `https://polisci.mit.edu/undergraduate`.

4 Prerequisites

This class will assume that you do not have any prior exposure to political science and machine learning. One prerequisite for this course is basic programming skills in at least one language (e.g., Python). Students who have taken **6.0001: Introduction to Computer Science and Programming in Python** or the equivalent are ready to take this course. If you have any questions about whether you are prepared for this course, please talk to the instructor.

5 Notes on Computing

In this course we use R, an open-source statistical computing environment that is very widely used in statistics and data science. (If you are already well versed in another statistical software, you are free to use it, but you will be on your own). We will begin the course with an introduction to R and no prior exposure to the programming language is required. Each problem set will contain computing and/or data analysis exercises which can be solved with R but often require going beyond canned functions to write your own program.

6 Course Requirements

The final grades are based on the following items:

- **Problem sets** (35%): **Five** problem sets will be given throughout the semester. Problem sets will contain analytical, computational, and data analysis questions. Each problem set will contribute equally toward the calculation of the final grade. The following instructions will apply to all problem sets unless otherwise noted.

  - Late submission will not be accepted unless you ask for special permission from the instructor in advance. (Permission may be granted or not granted, with or without penalty, depending on the specific circumstances.)
  - Working in groups is encouraged, but each student must submit their own writeup of the solutions. In particular, you should not copy someone else's answers or computer code. We also ask you to write down the names of the other students with whom you solved the problems together on the first sheet of your solutions.
  - For analytical questions, you should include your intermediate steps, as well as comments on those steps when appropriate. For data analysis questions, include annotated code as part of your answers. All results should be presented so that they can be easily understood.

  - All answers should be typed. Students are strongly encouraged to use a typesetting system such as LaTeX.

- **Final problem set** (15%): The last problems set will be a special problem set, which will be weighted more heavily toward the calculation of the final grade. You will not be allowed to collaborate with anybody on the final problem set. This is to check if you have developed sufficient experience to work through problems on your own.

- **Final group project** (40%): Students are expected to form a group. Each group will apply methods they learned in this course to an empirical problem of own substantive interest. The group project will be evaluated based on the performance across the following three tasks.

  - **Three Tasks**
    1. Data Collection and Cleaning (10%): Collecting political science data. Final output: (1) dataset and (2) replication codes.
       * Your group will either engage in your own data collection (e.g., using web-scraping) or utilize multiple existing datasets in political science. The dataset should be submitted in a standard data format (e.g., `.csv`, `sql`, `.json`) as an output of this task.
         · Collecting new data: Your team will choose a topic of interest related to various subfields in political science such as American Politics, International Relations, Comparative Politics, and Political Economy. The instructor will provide guidance to identify potential sources for novel data collection.
         · Utilizing existing datasets: You should merge various datasets available in political science research. The instructor will also make two of his own "big" datasets available: (1) Lobbying Database (see `www.LobbyView.org`), and (2) International Trade data.
    2. Descriptive Analysis (15%): Your group will then conduct descriptive data analysis. You should submit tables and figures that effectively illustrate key patterns in your data. A five-page report should be submitted as an output of this task by April 9.
    3. Data Analysis (15%): You will utilize various tools that you learn from the course to conduct an in-depth data analysis. Each group will give 10 minutes in-class presentation of their main findings in the last week of the semester. A poster should be submitted as a final output of the project on the last day of the class.
  - **Deadlines:** Please be aware of the following deadlines. Late submission will *not* be accepted. You are welcome to arrange a meeting (during the office hours) with the instructor and the TAs as you make a progress over the semester.
    * **February 21**: By this date, please form your team. A group should have at least three students and consist of no more than 5 students. Your team should arrange a meeting with the instructor within a week after this date.
    * **March 2**: By this date, your team should identify the dataset to analyze. Please submit one-page description of your project that explains (a) the specific dataset that your team is going to collect/analyze, (b) the main puzzle/problem that your team plans to study.
    * **April 9**: By this date, your team should submit a five-page long report summarizing the results from your descriptive data analysis. Please submit a document with at most 5 figures or tables that summarize your data with informative caption for each.

    - ∗ **May 16**: By this date, your team should submit the final poster of your project. There will be in-class poster session on this date. Poster session will be at E51-095 (9:30AM – 11:00AM) on May 16 (Wed)

- **Participation** (10%): Students are expected to attend all classes and actively participate in their group project. Students are strongly encouraged to ask questions and participate in discussions during lectures and recitation sessions. In addition, there will be required readings for each section of the course which students are expected to complete *prior to* the lectures.

7 Course Website

You can find the Stellar website for this course at:

http://stellar.mit.edu/S/course/17/sp18/17.835/

We will distribute course materials, including readings, lecture slides and problem sets, on this website.

8 Questions about Course Materials

In this course, we will utilize an online discussion board called *Piazza*. In addition to recitation sessions and office hours, please use the Piazza Q&A board when asking questions about lectures, problem sets, and other course materials. You can access the Piazza course page either directly from the below address or the link posted on the Stellar course website:

https://piazza.com/mit/spring2018/17835

Using Piazza will allow students to see other students' questions and learn from them. Both the TA and the instructor will regularly check the board and answer questions posted, although everyone else is also encouraged to contribute to the discussion. A student's respectful and constructive participation on the forum will count toward his/her class participation grade. *Do not email your questions directly to the instructor or TA* (unless they are of a personal nature)— we will not answer them!

9 Notes on Poster

Poster presentation is an efficient way to get valuable feedback from a large number of people. A poster should effectively summarize your key findings. Here are some notes.

1. **Use keywords and bullet points:** You should not use full sentences—your audience will never read them. Try to use keywords (or half sentences when needed), and make sure that you use only one line to deliver each point.

2. **Use LaTeX:** There are many online templates to help you make posters easily, e.g., http://www-i6.informatik.rwth-aachen.de/ dreuw/latexbeamerposter.php

10 Books

- Recommended books: We will read chapters from these books throughout the course. We strongly recommend that you at least purchase "Quantitative Social Science An Introduction." (QSS) These books will be available for purchase at COOP and online bookstores (e.g. Amazon) and on reserve in the library.
  - Imai, Kosuke. 2017 *Quantitative Social Science An Introduction*. Princeton University Press.
  - Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning*. Springer.
  - Christopher M. Bishop. 2007. *Pattern Recognition and Machine Learning*, Springer (A great introduction to machine learning).

11 Tentative Course Outline

11.1 Introduction

- Machine Learning and Data Science in Political Science

  Optional Reading:
  - Justin Grimmer. "We Are All Social Scientists Now: How Big Data, Machine Learning, and Causal Inference Work Together." Available at `https://stanford.edu/~jgrimmer/bd_2.pdf`

11.2 Data Analysis using R

- Dataframe
- Descriptive data analysis
- Visualization

  Optional Reading:
  - QSS: Chapter 1.3

11.3 Causality

- Causal Inference
- Average Treatment Effect (ATE) and Average Treatment Effect for the Treated (ATT)

  Required Reading:
  - QSS: Chapter 2
  - Hersh , Eitan D. 2013. "Long-Term Effect of September 11 on the Political Behavior of Victims Families and Neighbors." *Proceedings of the National Academy of Sciences* 110 (52): 20959 -63.

  Recommended Reading:
  - Fowler, James. 2008. "The Colbert Bump in Campaign Donations: More Truthful than Truthy." *PS: Political Science & Politics* 41(3): 533–539

11.4 Linear Regression

- OLS (Ordinary Least Squares)
- Difference in means estimator
- Regression and Causation

  Required Reading:

  - Chapter 4.2 (First Week)
  - Chapter 4.3 (Second Week)

  Optional Reading:

  - Stephens-Davidowitz , Seth I . 2014 a. "The Cost of Racial Animus on a Black Presidential Candidate: Evidence Using Google Search Data." Journal of Public Economics. 118 : 26–40

11.5 Supervised Learning

- Introduction to Supervised Learning
- K-Nearest-Neighbor (KNN) Classifier
- Support Vector Machine (SVM)
- Over fitting
- Ridge Regression
- Least Absolute Shrinkage and Selection Operator (LASSO)

  Optional Reading:

  - Francisco Cantú and Sebastián M. Saiegh. 2011 "Fraudulent Democracy? An Analysis of Argentinas Infamous Decade Using Supervised Machine Learning." *Political Analysis*. 19: 409–433
  - Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning*. Ch 3.1–3.4, Ch 7.

11.6 Unsupervised Learning Methods

- Principal Component Analysis (PCA)
- Clustering Algorithm

  Reading:
  - QSS: Chapter 3.7

- Mixture Models and EM Algorithm

  Reading:
  - Bishop Ch.9
  - Kosuke Imai, In Song Kim, and Steven Liao. "Measuring Trade Profile with Two Billion Observations of Product Trade." Working paper available at `http://web.mit.edu/insong/www/research/research.html`

11.7 Text Analysis

- Introduction to Text Analysis

  Required Reading:

  – QSS: Chapter 5.1

  Recommended Reading:

  – Grimmer, Justin, and Brandon M. Stewart. "Text as data: The promise and pitfalls of automatic content analysis methods for political texts." *Political Analysis* (2013): 28.

- Latent Dirichlet Analysis (LDA)

  Reading:

  – Blei, David M., Andrew Y. Ng, and Michael I. Jordan. "Latent Dirichlet allocation." *Journal of Machine Learning Research* 3 (2003): 993-1022.

  – Roberts, Margaret E., et al. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* (2014).

11.8 Network Analysis

- Network Analysis

  Recommended Reading:

  – QSS: Chapter 5.2

  – Kim, In Song and Dmitriy Kunisky. "Mapping Political Communities: A Statistical Analysis of Lobbying Networks in Legislative Politics." Working paper available at `http://web.mit.edu/insong/www/pdf/network.pdf`

11.9 Applications in Political Science

- International Trade with Big Data

  Reading:

  – C. A. Hidalgo, B. Klinger, A.-L. Barabási, R. Hausmann. "The Product Space Conditions the Development of Nations." *Science* 317.5837 (2007): 482-487

- Lobbying and Campaign Contribution

  Reading:

  – In Song Kim. "Political Cleavages within Industry: Firm-level Lobbying for Trade Liberalization." *American Political Science Review*, 111.1: 1-20.

  – Stephen Ansolabehere, John M. de Figueiredo, and James M. Snyder. "Why is There so Little Money in U.S. Politics?" *Journal of Economic Perspectives*, 17.1 (2003): 105-130

- Identifying Behavioral Patterns using Massive Data

  Reading:

- Gary King, Jennifer Pan, and Margaret E Roberts. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review*, 107.2: 326-343.
- Pierson, E., Simoiu, C., Overgoor, J., Corbett-Davies, S., Ramachandran, V., Phillips, C., and Goel, S. (2017). "A large-scale Analysis of Racial Disparities in Police Stops across the United States." arXiv preprint arXiv:1706.05678.

- Measuring Ideological and Political Preferences using Social Network Data

  Reading:

  - Robert Bond and Solomon Messing. "Quantifying Social Media's Political Space: Estimating Ideology from Publicly Revealed Preferences on Facebook." *American Political Science Review* 109.1 (2015): 62-78.
  - Pablo Barberá "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23.1 (2014): 76-91

- What do Politicians Do?

  Reading:

  - Justin Grimmer, Solomon Messing, and Sean Westwood. "How Words and Money Cultivate a Personal Vote: The Effect of Legislator Credit Claiming on Constituent Credit Allocation." *American Political Science Review*, 106.4 (2012), 703-719

- Big Administrative Data: Promises and Pitfalls

  Reading:

  - Connelly, R., Playford, C.J., Gayle, V., Dibben, C., 2016. "The Role of Administrative Data in the Big Data Revolution in Social Science Research." *Social Science Research*, Special issue on Big Data in the Social Sciences 59, 112
  - Kopczuk, W., Saez, E., Song, J., 2010. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937." *The Quarterly Journal of Economics* 125, 91128.

- Machine Learning Algorithms in Society

  Reading:

  - Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. "Human Decisions and Machine Predictions." *The Quarterly Journal of Economics* 133 (1):23793