Gov 391K: Machine Learning in Political Science

Fall 2022 ∗ Mondays, 3:00-6:00 pm ∗ BAT 1.104

Connor T. Jerzak
Assistant Professor of Government
University of Texas at Austin

Substantive inquiry in social science increasingly involves massive datasets or non-traditional data types containing image, audio, text, network, temporal, or geospatial information. In this context, traditional quantitative methods, many of which are based on the generalized linear modeling framework, can be difficult to apply due to the problems that they face in the situation where the number of covariates is greater than the number of data points or where the data is best represented not as a matrix but as a more general kind of data object. Since the 1950s and with explosive growth since the 1990s, a class of modeling strategies has been developed to address these problems while flexibly learning associations present in high-dimensional or non-traditional data types. These modeling strategies fall under the umbrella terms "machine learning" (ML) or "artificial intelligence" (AI).

In this course, we will explore this universe of ML modeling strategies—their motivation, properties, and manner of implementation. Models examined will include tree-based methods, convolutional network networks, machine learning approaches for social networks, time series network models, ensembles, and probabilistic machine learning. We will interrogate how the use of these strategies can be used to support social science inquiry. We will also examine topics explicitly at the intersection of social science and ML itself—topics related to fairness, privacy, and the implications of ML for politics and society. By the end of the course, students should be able to employ state-of-art machine learning methods in their research and policy applications, understand the strengths and limitations of these methods, and know how to modify the proposed methods to maintain legal compliance, privacy, and ethics.

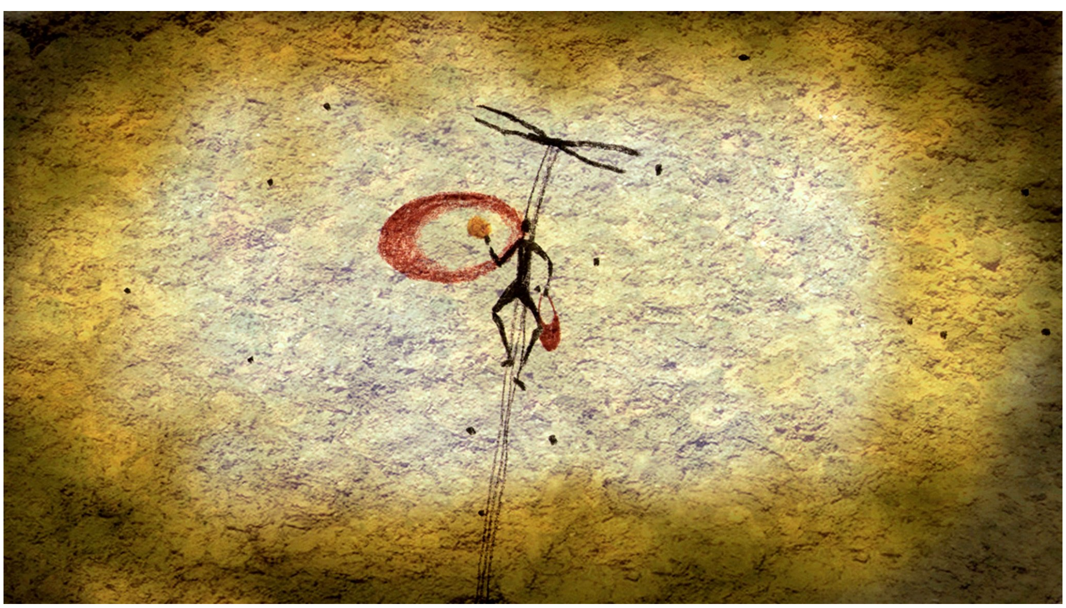

A Stone Age cave painting in Spain depicting the gathering of honey with fire.
Credit: Thierry Berrod/Mona Lisa Production/Science Source

1 Contact Information

Instructor
NAME: Connor T. Jerzak
EMAIL: connor.jerzak@austin.utexas.edu
URL: ConnorJerzak.com
OFFICE: BAT 3.108
VIRTUAL OFFICE: utexas.zoom.us/my/jerzak
OFFICE HOURS: Wednesdays, 11:40 am - 1:40 pm tinyurl.com/2w5dmucr
or send me an email to make an appointment.
Feel free to drop by my office anytime as well:
As long as I'm there and the door is open, I'm happy to talk anytime.

Teaching Assistant
NAME: Siyun Jiang
EMAIL: syjiang@utexas.edu
URL: SiyunJiang.weebly.com
OFFICE HOURS: Thursdays, 9:30-12:30 pm, BAT 1.118

2 Prerequisites

By far the most important prerequisite for the course is a desire to learn the material and willingness to embrace the feeling of discomfort that can arise when learning new things.

This course also assumes some knowledge of

- probability theory
- statistical theory
- linear algebra
- calculus
- data analysis using R

If you are unsure of whether your background will be sufficient for the course, reach out to the instructor.

Note: This course is approved to serve as a major elective for any student concurrently pursuing the University of Texas at Austin Master's in Statistics.

3 Instructional Tools

We will use a variety of tools to run this course.

- **Canvas** (utexas.instructure.com/courses/1345570): Our course website will be hosted by Canvas. The website will contain all course information related to the course. We will use the "Discussions" feature in Canvas to encourage opportunities for students to help each other learn by answering questions other students pose about the material and about problem sets.
- **Gradescope**: Gradescope is the online platform you will use to submit and receive feedback on your assignments. You can access our Gradescope page via Canvas.

4 Course Framework

The course has four main components: in-class lectures/discussions, section, bi-weekly exercises done in pairs, an individual mid-term exam, and a collaborative end-of-semester project. We believe it is important to have multiple ways in which students can be evaluated since some students may learn best with each modality (e.g., in group work, in solo work, in project-based work).

Some guidelines.

- **Class Meetings**: In class meetings, key material will be introduced. We will typically begin class with a lecture. A break will occur at the class half-way point. After the break, we will usually have a more practical tutorial. We ask that students observe the following guidelines during class:

  1. *Be on time.* Late entry of students can be disruptive to the focus of students who arrived on time. There may be cases where, due to external circumstances, students may need to regularly arrive late; in that case, let the instructors know and we will think of a plan to minimize the degree of class disruption.
  2. *No computer or cell phone use during lecture components.* Computers can be a useful aide to learning in some situations, but evidence has accumulated that they can be a source of distraction for those around you in others. Therefore, use of laptops and cell phone during the lecture component of the course is not allowed.

- **Exercises**: Exercises will be done in partnership with a randomly selected peer pair. These pairs will be rotated every exercise. You and your partner will work together to answer the questions. Submissions are to be made to Gradescope by 5 pm Sunday. See next section for additional details.

5 Course Requirements

The course grade is based on the following components:

- **Class participation** (10% of the course grade): The class participation portion of one's grade is important but can sometimes feel ill-defined. By "participation," we mean engagement broadly defined: Does a student attend class? Is the student engaged in class or section? Is the student helpful to others on the Canvas discussion page? Does the student treat other students with courtesy and respect in collaborations or discussions? The later aspect of the grade is not meant to be punitive; in your career, your success will be affected by your ability to work well with others, so this aspect is also a part of the course evaluation.

- **Exercises** (30% of the course grade): Students will be required to complete exercises on a bi-weekly basis. You will be randomly assigned pairs. In these pairs, you will expected to support each other. The motivation for this pairing is to encourage students to (a) get to know a broad set of other students while (b) helping each other in your understanding of the material. The use of bi-weekly exercises also encourages regular engagement with the material, something critical to learning. The following guidelines will apply to these exercises:

  – *Collaboration policy.* Students are allowed to discuss the exercises with the instructional staff and any other students in class in addition to your assigned pair. After all, teaching one's peers can be helpful in learning the material better oneself. However, you and your assigned partner are required to write up your own answers. You are not allowed to copy the work of other teams. Copying others' work deprives you of the opportunity to learn the material; if you feel like you are struggling in the course, this information is very important for the instructional staff to know so we can support you. If you feel like you are struggling in the course, it is better to let us know sooner rather than later so we can identify impactful solutions.

  - *Online help and office hours.* You are strongly encouraged to reach out to the instructional staff through Canvas and office hours about any questions you might have about the course materials. Students should also feel free to ask questions and answer the questions posed by others at Canvas, which will be considered in the overall evaluation of class participation.
  - *Submission guidelines.* Answers to all questions should be incorporated into a single Rmarkdown (or LaTeX) file. The instructional staff will provide an Rmarkdown template. For the exercises, each student should submit the PDF file electronically to Gradescope. Once you upload the PDF file, you will see a list of the questions in the assignment and thumbnails of your file. For each assigned question, click the PDF page(s) that contains your answer. No late submission will be accepted unless you obtain a prior approval from the instructor.

- **Mid-term take-home exam** (30% of the course grade): There will be a mid-term take-home exam including some analytical and some empirical questions. The motivation for having an exam is that the existence of an individual evaluation event can be a good stimulus for students to hone in on learning the material in a more focused way than in the bi-weekly exercises. You may not discuss the exam with anyone else. You are allowed to ask the instructional staff clarifying questions. You will have 24 hours to complete the exam once you've opened it. We will give more detailed instructions before the exam is released.

- **Final project** (30% of the course grade): The final project can be a collaborative project with another student in this class or can be done on an individual basis. There are pro's and con's to both choices. In collaborative projects, you can often implement more ambitious analyses, as you will have the support of a peer. In an individual project, you can tailor it specifically to a part of your dissertation. The final project can be methodological (e.g., development of a method) or empirical (e.g., application of an existing method). We highly recommend the use of LaTeX or Rmarkdown. The quality of project will be judged based on the originality of ideas presented and the clarity with which they are presented. Students are required to meet the following three checkpoints and are encouraged to discuss their projects with the teaching staff throughout the semester:

  - October 5: The proposal should include a brief statement of the problem to be solved (or the question to be answered) and propose a feasible plan for conducting research. The proposal should be one page long (single spaced).
  - November 9: A memo that outlines the core results of the project. For an empirical project, you should include main tables and figures with brief explanations. For a methodological project, you should include the description of methods and any analytical results. A project may combine both aspects. The memo should not exceed 5 single-spaced pages. You will be asked to give feedback on another student/group's project, which will be graded based on its usefulness (5% of the course grade).
  - December 12: Final report (no more than 15 single-space pages including the title, tables, and figures).

6 Course Materials

There is no single textbook for this course. The following texts are excellent references for some of the material:

PML Murphy, Kevin P. (2022). *Probabilistic Machine Learning*. MIT Press. https://probml.github.io/pml-book/book1.html

DL Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2019). *Deep Learning*. MIT Press. https://www.deeplearningbook.org/

ESL Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning*. Springer. https://link.springer.com/content/pdf/10.1007/978-0-387-84858-7.pdf

7 Course Plan

This course intends to provide an introduction to the following topics related to machine learning and its application to social science:

Module 1 Introduction & Key Theory

Introduction to Machine Learning

Class 1 (Aug. 22) Overview of the course; ML vs. statistical inference vs. causal inference; a brief history of ML

Reading
- Daoud, Adel and Devdatt Dubhashi (2020).“Statistical modeling: the three cultures.” https://arxiv.org/abs/2012.04570
- Athey, S. (2018). The impact of machine learning on economics. In *The economics of artificial intelligence: An agenda* (pp. 507-547). University of Chicago Press.

Class 2 (Aug. 29) Ten things you need to know about the theory of ML; gradient-based vs. convexity-based ML; autodifferentiation

Reading
- ESL 7
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018). Automatic differentiation in machine learning: a survey. *Journal of Marchine Learning Research*, 18, 1-43.

ML with Tabular Data

Module 2 ML with High-Dimensional Tabular Data (Application: Treatment Effect Heterogeneity)

Class 1 (Sept. 12) Regularized regression models (Lasso, Ridge, `hiernet`); cross-validation

Reading
- ESL 3 (excluding 3.3)
- Bien, J., Taylor, J., & Tibshirani, R. (2013). A lasso for hierarchical interactions. *Annals of statistics*, 41(3), 1111.

Exercise 1 Due Sept. 18

Class 2 (Sept. 19) Tree-based methods; random forests; BART; causal forests

Reading
- Strobl, C., Malley, J., & Tutz, G. (2009). An introduction to recursive partitioning: rationale, application, and characteristics of classification and regression trees, bagging, and random forests. *Psychological methods*, 14(4), 323.
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27), 7353-7360.

Module 3 Model Validation: Visualization, Metrics, Simulation

Class 1 (Sept. 26) Validation via visualization, via quantitative metrics, and via simulation; diagnosing ML modeling problems

Reading
- Hossin, M., & Sulaiman, M. N. (2015). A review on evaluation metrics for data classification evaluations. *International journal of data mining & knowledge management process*, 5(2), 1.
- Wood-Doughty, Z., Shpitser, I., & Dredze, M. (2021). Generating synthetic text data to evaluate causal inference methods. *arXiv preprint arXiv:2102.05638.*
- Tufte, Edward (2020). *Seeing with Fresh Eyes: Meaning, Space, Data, Truth.* Graphics Press. *(Optional)*

ML with Non-Tabular Data

Exercise 2 Due Oct. 2

Module 4 An Introduction to Neural Networks (Application: Social Science Uses of Satellite Images)

Class 1 (Oct. 3) What are neural networks? What are the common tools to improve training? What are CNNs?; When use this model class in the social science context?

Reading
- DL 1.2 *(if you didn't have a chance to read this in Module 1)*
- DL 6
- DL 9

Final Project Checkpoint 1 due Oct. 5

Module 5 Probabilistic Machine Learning (Application: Network Analysis)

Class 1 (Oct. 10) Some more theory: autodifferentiation of stochastic computation nodes, re-parameterization sampling, "dropout" connection, variational inference

Reading
- Jospin, L. V., Laga, H., Boussaid, F., Buntine, W., & Bennamoun, M. (2022). Hands-on Bayesian neural networks—A tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2), 29-48.
- Ranganath, R., Gerrish, S., & Blei, D. (2014, April). Black box variational inference. In *Artificial intelligence and statistics* (pp. 814-822). PMLR.
- DL 16 *(optional)*

Module 6 Representation Learning (Application: Politics in Text Data (1))

Exercise 3 Due Oct. 16

Class 1 (Oct. 17) What is representation learning? Word vectors; causal embeddings; transfer learning

Reading
- Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532-1543).
- Rodman, E. (2020). A Timely Intervention: Tracking the Changing Meanings of Political Concepts with Word Vectors. *Political Analysis*, 28(1), 87-111. doi:10.1017/pan.2019.23
- Veitch, V., Sridhar, D., & Blei, D. M. (2019). Using text embeddings for causal inference.
- DL 15.2

Mid-term Released on Oct. 18

Mid-term Due by Oct. 23

Module 7 Time Series Models (Application: Politics in Text Data (2))

Class 1 (Oct. 24 LSTM; transformers; language models

Reading
- PML 15.1-15.2 (LSTM)
- PML 15.4-15.5 (Transformers)
- PML 15.7 (Language model applications)

Module 8 Ensembles (Application: Time Series Prediction of Covid Death Rates)

Class 1 (Oct. 31) Super-learners; Bayesian model averaging

Reading
- Van der Laan, M. J., Polley, E. C., & Hubbard, A. E. (2007). Super learner. *Statistical applications in genetics and molecular biology*, 6(1).
- Montgomery, J. M., Hollenbach, F. M., & Ward, M. D. (2012). *Improving predictions using ensemble Bayesian model averaging.* Political Analysis, 20(3), 271-291.

Module 9 Double Machine Learning (Application: Confounder Adjustment)

Exercise 4 Due Nov. 6

Class 1 (Nov. 7) Double machine learning

Reading
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters.
- Felderer, B., Kueck, J., & Spindler, M. (2022). Using Double Machine Learning to Understand Nonresponse in the Recruitment of a Mixed-Mode Online Panel. *Social Science Computer Review*, 08944393221095194.

Final Project Checkpoint 2 due Nov. 9

ML in Society

Module 10 Privacy, Fairness & Interpretability in Machine Learning (Application: Bail Recommendation Systems)

Class 1 (Nov. 14) Quantitative fairness in ML; social utility and ML models

Reading
- Kleinberg, J., & Mullainathan, S. (2019, June). Simplicity creates inequity: implications for fairness, stereotypes, and interpretability. In *Proceedings of the 2019 ACM Conference on Economics and Computation* (pp. 807-808).

- Rudin, C., Chen, C., Chen, Z., Huang, H., Semenova, L., & Zhong, C. (2022). Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys*, 16, 1-85.
- Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudik, M., & Wallach, H. (2019, May). Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2019 CHI conference on human factors in computing systems* (pp. 1-16).
- Bibal, A., Lognoul, M., De Streel, A., & Frénay, B. (2021). Legal requirements on explainability in machine learning. *Artificial Intelligence and Law*, 29(2), 149-169.
- Evans, G., King, G., Schwenzfeier, M., & Thakurta, A. (2020). Statistically valid inferences from privacy protected data. URL: GaryKing. org/dp.

Module 11 A Crash Course into Reinforcement Learning (Application: Anti-Poverty Experiments)

CLASS 1 (NOV. 28) What is reinforcement learning? What are the core algorithms? How is it relevant to the social sciences?

READING
- Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., & Wen, Z. (2018). A Tutorial on Thompson Sampling. *Foundations and Trends in Machine Learning*, 11(1), 1-96.
- Charpentier, A., Elie, R., & Remlinger, C. (2021). Reinforcement learning in economics and finance. *Computational Economics*, 1-38.
- Shortreed, S. M., Laber, E., Lizotte, D. J., Stroup, T. S., Pineau, J., & Murphy, S. A. (2011). Informing sequential clinical decision-making through reinforcement learning: an empirical study. *Machine learning*, 84(1), 109-136.

CONCLUSION

Module 12 Wrapping Up

CLASS 1 (DEC. 5) Reviewing key points; discussing the future of ML in policy and social science research; improving the credibility of ML-based social science; wrapping up.

READING
- Kapoor, Sayash and Narayanan Arvind (2022). "Leakage and the Reproducibility Crisis in ML-based Science." https://arxiv.org/abs/2207.07048

FINAL PROJECT **Due by 11:59 pm Austin time, December 12 (tentative)**

8 A Note on Programming Language Choice

Students should be aware that there are multiple computer programming languages that can be used for implementing machine learning models. In this course, we will at times make use of `tensorflow`, which has well-supported `R` interface through the package `Rtensorflow`. `tensorflow` functions can also be converted to `JAX` functions (`JAX` is another language which focuses, among other things, on high-performance automatic differentiation for advanced machine learning models). Students should be aware that `PyTorch` is a competing machine learning language which is also popular in academia and industry. Syntax in `tensorflow` and `PyTorch` are similar, so students who know one should be able to learn the other with relative ease.

9 A Note on the Readings

When doing the readings, there are several strategies that you might pursue.

1. First, you could start by skimming over the article, looking at the equations (or figures) and trying to understand what they're telling you. After thinking about that, then you may have a better framework for where you should focus your attention in the article, and what material is relatively less important to linger on.

2. A second approach is to read the introduction and conclusion, getting a sense for the big picture, and then diving into the details.

Don't ever feel discouraged if you feel like you aren't understanding parts of the readings!

Get what you can out of it in your first read-through—then talk to friends, consult YouTube videos, or come to Office Hours to work through your understanding. A great way to understand the intuition for something better is to implement it yourself (as Richard Feynman famously said, "What I cannot create, I do not understand.").

No one's understanding of a topic is total, we are all in a learning journey that sometimes involves valleys but sometimes also involves peaks!

10 Other Resources

There are other resources that could be useful to students in different aspects of the course, such as when giving presentations or when undertaking the creative process of formulating a research project.

On giving presentations:

- Winston, Patrick. "How to Speak." youtube.com/watch?v=Unzc731iCUY

On the creativity process:

- De Bono, Edward. *Laternal Thinking: Creativity Step by Step.* Harper Colophon.

On writing:

- Pinker, Stephen. *The Sense of Style: The Thinking Person's Guide to Writing in the 21st Century.* Penguin Books.

- Williams, Joseph and Joseph Bizup. *Style: Lessons in Clarity and Grace.* Pearson.

11 Arrangements

Students with arrangements through the UT Disability & Access Services Office should ensure that those arrangements are communicated to the instructor so we can set up a constructive environment for you.

12 Integrity

"The core values of The University of Texas at Austin are learning, discovery, freedom, leadership, individual opportunity, and responsibility. Each member of the University is expected to uphold these values through integrity, honesty, trust, fairness, and respect toward peers and community."

— UT Austin Code of Conduct

"As a student of The University of Texas at Austin, I shall abide by the core values of the University and uphold academic integrity."

— UT Austin Student Honor Code

13 Acknowledgements

I am grateful to Matthew Blackwell, Adel Daoud, Frances Hagopian, Kosuke Imai, Gary King, Michael Peress, Xiaolong Yang, and Xiang Zhou; the planning of the course builds on conversations or prior teaching experience with them.

14 Final Reflections

The instructors reserve the right to modify aspects of the course as it proceeds, both in response to student feedback and our own assessment of how to improve the course.