

[PADP 9200] Big Data and Artificial Intelligence in Public Administration and Policy

University of Georgia

FALL 2021

Professor L. Jason Anastasopoulos

Email: ljanastas@uga.edu

Course Information

Office Hours: TuTh 3:00-4:00

Course Time and Day: Mondays, 4:10-6:55 PM

Location: Baldwin 301.

Course Objectives

The big data revolution is transforming public policy and governance. The goal of this course is to provide a non-technical overview of some of the methods driving big data methodologies and to explore how these technologies are shaping and will shape the future of public policy and government with an eye towards the ethical dilemmas that these technologies raise. We begin the course with a discussion of some of the fundamental theories and applications of machine learning methods, which form the basis of big data and artificial intelligence technologies. We then move on to an in-depth discussion of the ethical and societal promises and perils that these technologies pose for decision making in government more broadly, focusing on the potential of these techniques to shape government and policy.

Attendance and Participation

The most important content from this class will come from the lectures and group assignments during lecture time. Because of this and the rather technical nature of this class, attendance and participation in class is extremely important. If you cannot attend a lecture you must present me with a valid excuse at least 24 hours prior to the start of class unless the situation you encountered was an emergency. Either way, absence requires explanation and documentation if you do not want points taken off your final participation grade.

Grading and Requirements

- *Participation/discussion leader: 40%.*
- *Three (3) problem sets: 20% (lowest of 3 is dropped)*
- *Final group report and presentation: 35%.*

Discussion Leaders

Groups of students will be assigned the role of “discussion leaders.” Discussion leaders will lead the class discussion for that week by reading the assigned content, preparing a 15-30 minute presentation summarizing the readings and will propose a series of 3-5 discussion points to start off our discussion about the content. Every student **MUST** participate in a group as a discussion leader at least once. Discussion leader groups can have a maximum of 2 people and if you do not sign up for one week as a discussion leader you will be assigned to a week.

PLEASE SIGN UP TO BE A DISCUSSION LEADER AT LEAST TWICE WITH YOUR NAME AND THE READING FOR THAT WEEK THAT YOU WOULD LIKE TO DISCUSS. Here is the sign-up sheet:

<https://docs.google.com/spreadsheets/d/13z5636GkmDaOlcUBMHxrbhfWyqx7JGRsLMn4XVhszwQ/edit?usp=sharing>

Problem Sets

During the first three weeks of the course, there will be three problem sets which you can work on in groups and which are designed to give you some hands on experience with machine learning algorithms in a policy context. These problem sets will involve some rudimentary programming in the statistical language **R** and will teach you about some of the basic machine learning algorithms used in practice today.

Group Project

A major portion of your grade will involve writing an essay in response to a question that I will assign two weeks prior to the memorandum due date. For guidelines on how to write a policy memorandum, if you choose to write in this style, please see this excellent guide by Iris Malone: [policy memorandum guidelines](#). The essay will be graded on the basis of the quality and clarity of your writing and the quality and clarity of the ideas that you present.

Required and Recommended Texts*

Kearns, Michael, and Aaron Roth. [The Ethical Algorithm: The Science of Socially Aware Algorithm Design](#). Oxford University Press, 2019.

[Machine Learning: The New AI](#) Ethem Alpaydin (2016) MIT Press. Referred to in the schedule as **EA**.

James, Witten, Hastie and Tibshirani. 2015. *An Introduction to Statistical Learning with Applications in R*. Springer Science. Available for free here: <http://www-bcf.usc.edu/~gareth/ISL/>. Referred to in the schedule as **JWHT**.

* Most texts will be available online on the course site.

Course Outline

1. Introduction to machine learning and artificial intelligence.
2. Introductions to algorithms in government.
3. Algorithms, ethics and public policy.
 - Introduction.
 - Normative approaches.
 - Causal inference approaches.
4. Algorithms, behavior and decision-making in government.
 - Decision making by humans and machines.
 - Strength and weaknesses of machine learning and AI.
 - Machine learning and fairness.
5. Special topics: Blockchain
 - Introduction to blockchain and smart contracts.
 - Applications of blockchain in government: social credit, monetary policy, elections.
6. Special topics (time permitting): Biotechnology and Transhumanism.

COURSE SCHEDULE

Class 1 08/23: *Introduction to Machine Learning and Artificial Intelligence: Fundamentals*

Programming fundamentals.

- Introduction to programming in **R**.

Overview of machine learning

- Machine learning in public organizations.
- What is machine learning?
- Supervised & unsupervised learning.
- Inference versus prediction.

Readings:

- ❖ Larson et al “How We Analyzed the COMPAS Recidivism Algorithm”
<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>
- ❖ **EA** Chapter 1.
- ❖ Kleinberg, J., Ludwig, J., Mullainathan, S. and Obermeyer, Z., 2015. Prediction policy problems. American Economic Review, 105(5), pp.491-95.
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4869349/pdf/nihms776714.pdf> .
- ❖ **JWHT** – Introduction, pp 1-15.

Class 2 08/30: *Introduction to statistical learning theory*

- Training, testing and cross-validation.
- Assessing model accuracy.
- Overfitting.
- Regression vs. classification problems.
- The Bias-Variance tradeoff.
- **Application:** H1-B Visa Certification. [H1-B Application Data](#).

Required Readings

- ❖ **JWHT** – Statistical Learning, pp 15-37, 176–184.
- ❖ **EA** - Chapter 2.

NO CLASS: LABOR DAY 09/06

Class 3 09/13: *Understanding supervised machine learning through examples: decision trees and regression.*

- “Pure ML”: Decision tree algorithms and CART.
- “Statistical ML”: Linear regression as a machine learning algorithm.
- Application 1: Preventative policing: pre-crime targeting and detection. [NYC Stop and Frisk Data: 2003–2016](#)

Required Readings

- ❖ **JWHT** – Chapter 3.
- ❖ **EA** - Chapter 3.

Class 4 09/20: Algorithms and government - overview.

- Overview of administrative decision making and machine learning.
- Bureaucracy and technology.

Required Readings

- ❖ Coglianese and Lehr. 2017. "[Regulating by Robot Administrative Decision Making in the Machine-Learning Era](#)". Georgetown Law Journal.
- ❖ [A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions](#). Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, Rhema Vaithianathan ; PMLR 81:134-148
- ❖ Movie: *The Minority Report*. You can watch this for free on Internet Archive: <https://archive.org/details/MinorityReport>

Optional Readings

- ❖ Lee, Ronald M. "[Bureaucracies, bureaucrats and information technology](#)." *European Journal of Operational Research* 18, no. 3 (1984): 293-303.

Class 5 09/27: Algorithms, ethics and public policy I: Introduction and approaches.

- Overview of algorithmic fairness.

Required Readings

- Kearns, Michael, and Aaron Roth. [The Ethical Algorithm: The Science of Socially Aware Algorithm Design](#). Oxford University Press, 2019. **pp 1-94**

Optional Readings

- ❖ Cath, C., 2018. Governing artificial intelligence: ethical, legal and technical opportunities and challenges. <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2018.0080>

- ❖ Piano, S.L., 2020. Ethical principles in machine learning and artificial intelligence: cases from the field and possible ways forward. *Humanities and Social Sciences Communications*, 7(1), pp.1-7. <https://www.nature.com/articles/s41599-020-0501-9>

Class 6 10/04: *Algorithms, ethics and public policy II: Normative approaches.*

- Normative approaches to creating fair algorithms.

Required Readings

- ❖ Binns, Reuben. "[What Can Political Philosophy Teach Us about Algorithmic Fairness?](#)" *IEEE Security & Privacy* 16, no. 3 (2018): 73-80.
- ❖ Rawls, John. *A theory of justice*. Harvard university press, 2009. Pgs 1-15 (of the book not the PDF)
<https://philosophyintrocourse.files.wordpress.com/2016/05/rawls-theory-of-justice.pdf>
- ❖ Leben, Derek. "A Rawlsian algorithm for autonomous vehicles." *Ethics and Information Technology* 19, no. 2 (2017): 107-115.
<https://link.springer.com/content/pdf/10.1007/s10676-017-9419-3.pdf>

Optional Readings

- ❖ Keeling, Geoff. "Against Leben's Rawlsian collision algorithm for autonomous vehicles." In 3rd Conference on Philosophy and Theory of Artificial Intelligence, pp. 259-272. Springer, Cham, 2017.
https://www.researchgate.net/profile/Geoff_Keeling/publication/327274610_Against_Leben%27s_Rawlsian_Collision_Algorithm_for_Autonomous_Vehicles/links/5c9fb00e45851506d736183b/Against-Lebens-Rawlsian-Collision-Algorithm-for-Autonomous-Vehicles.pdf

Class 7 10/11: *Algorithms and government - decision making in theory.*

- Decision making by humans v. machines.

Required Readings

- ❖ Keiser, Lael. 2010. "[Understanding Street Level Bureaucrats Decision Making](#)," *Public Administration Review*. 70 (02) pp.247-57. **JSTOR**
- ❖ Dawes et al. [Clinical versus actuarial judgement](#)

- ❖ Dietvorst. [Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err](#)

Class 8 10/18: *Algorithms and government - decision making in practice - police and judges*

- Technology and street level bureaucracy.
- Applied examples: judges and police.

Required Readings

- ❖ Harcourt, Bernard E. "[Against prediction: Sentencing, policing, and punishing in an actuarial age.](#)" (2005).
- ❖ Green, Ben, and Yiling Chen. "Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 90-99. ACM, 2019.
<https://dl.acm.org/doi/pdf/10.1145/3287560.3287563?download=true>

Optional Readings

- ❖ Kleinberg, Jon, et al. "[Human decisions and machine predictions.](#)" *The quarterly journal of economics* 133.1 (2017): 237-293.

Class 9 10/25: *Strengths and Weaknesses of Machine Learning Systems for Public Policy*

- Strengths and weaknesses of the machine learning approach and how it might apply to public policy.

Required Readings

- ❖ Breiman. [Statistical Modeling: The Two Cultures](#)
- ❖ Lazer et al. [The parable of Google Flu.](#)
- ❖ Olteanu et al. [Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries](#)

Optional Readings

- ❖ Norvig. [On Chomsky and the Two Cultures of Statistical Learning](#)

Class 10 11/01: Machine Learning and Bias I: Overview

- Defining bias in the machine learning context.

Required Readings

- ❖ Angwin et al. [Machine Bias](#)
- ❖ Angwin & Larson. [Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say](#)
- ❖ Chouldechova. [Fair Prediction with Disparate Impact: A study of bias in recidivism prediction instruments.](#)
- ❖ Kleinberg et al. [Inherent Trade-Offs in the Fair Determination of Risk Scores](#)

Optional Readings

- ❖ Corbett-Davies et al. [Algorithmic Decision Making and the Cost of Fairness.](#)

Class 11 11/08: Machine Learning and Bias II: Sources and Pathways in Practice

- Machine learning bias in practice.

Required Readings

- ❖ Pierson et al. [A large-scale analysis of racial disparities in police stops across the United States](#)
- ❖ Caliskan et al. [Semantics Derived Automatically from Language Corpora Contain Human-like Biases](#)
- ❖ Torralba & Efros. [Unbiased Look at Dataset Bias](#)

Class 12 11/15: Machine learning, big data and ethics: the modern panopticon?

- Ethical considerations of big data and machine learning.

Required Readings

- ❖ Foucault. ["Panopticism" in Discipline and Punishment. Pp 195-228](#)
- ❖ Ohm & Peppet. [What if Everything Reveals Everything?](#)

- ❖ Kosinski, Stillwell, and Graepel. [Private Traits and Attributes Are Predictable From Digital Records of Human Behavior](#)
- ❖ Wu & Zhang. [Automated Inference on Criminality using Face Images](#)
- ❖ Wang & Kosinsky. [Deep neural networks are more accurate than humans at detecting sexual orientation from facial images](#)

Class 13 11/22: *Special topics: Introduction to blockchain in government*

- ❑ Blockchain Technology Explained (YouTube Video):
<https://www.youtube.com/watch?v=qOVAbKkSH10>

Blockchain in Government: Social Credit

- ❑ Social Credit and E-Government
 - ❑ Hou, H., 2017, July. "The application of blockchain technology in E-government in China." In 2017 26th International Conference on Computer Communication and Networks (ICCCN) (pp. 1-4). IEEE.
 - ❑ Xu et al. 2018. "Constructing trustworthy and safe communities on a blockchain-enabled social credits system". In Proceedings of the 15th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (pp. 449-453).
https://dl.acm.org/doi/pdf/10.1145/3286978.3287022?casa_token=jc7wi2ZSOmkAAAAA:Rnlm74_jX6PQ4IU4DrrgZKI4DojCnLjPYQvpi8ldyJ96ScR20k6eEj0gxZLI_L4X9jUFUBEscvgVv

OPTIONAL READINGS

Blockchain Monetary Policy: Central Bank Digital Currency (CBDC) and Cryptocurrency.

- ❑ Bordo, M.D. and Levin, A.T., 2017. Central bank digital currency and the future of monetary policy (No. w23711). National Bureau of Economic Research.

Blockchain and E-Voting

- ❑ [Culnane, C., Essex, A., Lewis, S.J., Pereira, O. and Teague, V., 2019. Knights and knaves run elections: Internet voting and undetectable electoral fraud. *IEEE Security & Privacy*, 17\(4\), pp.62-70.](#)
- ❑ [Park, S., Specter, M., Narula, N. and Rivest, R.L., 2020. Going from bad to worse: from internet voting to blockchain voting.](#)

Class 14 11/29: *Internet of Things (IoT) and the Internet of Bodies (IoB)*

READINGS TBD.

Class 16 12/06: *Governing the Internet of Things (IoT) and the Internet of Bodies (IoB)*

READINGS TBD.

Statement about Students with Disabilities

Students with special needs that require accommodation should notify me and the Office for Disability Services in the first two weeks of the course so appropriate arrangements can be made. All information and documentation of special needs is confidential.

Statement about Plagiarism and Academic Dishonesty

Students are responsible for maintaining the highest standards of honesty and integrity in every phase of their academic careers. The penalties for academic dishonesty are severe and ignorance of the policy is not an acceptable defense. See also <https://ovpi.uga.edu/academic-honesty>.

CDC Recommendations Regarding COVID-19

- Wear a mask to protect yourself and others to reduce the spread of COVID-19.
- Stay at least 6 feet from others who don't live with you and avoid crowded areas.
- Wash your hand with soap and water for 20 seconds or use hand sanitizer with at least 60% alcohol.