[PADP 9200] Big Data in Public Administration and Policy or Machine Learning, Artificial Intelligence and the Administrative State

Professor Jason Anastasopoulos
Class: 6:30-9:15 pm, 311 Baldwin Hall
Email: ljanastas@uga.edu
Office: Baldwin Hall 280E
Office Hours: 11-12pm, T,Th or by Appointment

Course Objectives

The big data revolution is transforming public organizations and governance from the micro (bureaucrats) to the macro (organizational) level. 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 organizations, policy and governance. 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 promise and perils of these techniques for decision making and design of public organizations more broadly, focusing on the potential of these techniques for shaping public organizations and government more broadly. Finally, we discuss some of the challenges surrounding regulation of artificial intelligence technologies and the implications of these technologies for the future of the administrative state.

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 want points taken off your final participation grade.

Grading and Requirements

- *Participation/discussion leader: 20%.*
- *Two (2) problem sets: 20% (10% each) (lowest will be dropped).*
- *Three (3) policy memoranda: 60%.*

Discussion Leaders

After Week 2, 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 and 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 4 people and if you do not sign up for one week as a discussion leader you will be assigned to a week. Each student can serve as a discussion leader a **maximum** of 2 times. If you serve as a discussion leader twice, you can replace one of your policy memorandum grades with your grade as a discussion leader.

Problem Sets

During the first month of the course, there will be two brief 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 very rudimentary programming in the statistical language **R** and will teach you about some of the simplest supervised and unsupervised machine learning algorithms used in practice today.

Policy Memoranda

A major portion of your grade will involve writing three (3) policy memoranda in response to a question or questions that I will assign two weeks prior to the memorandum due date. For guidelines on how to write a policy memorandum, please see this excellent guide by Iris Malone: http://web.stanford.edu/~imalone/Teaching/ps1winter17/PS1-PolicyMemo.pdf. Policy memoranda 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.

Texts

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

Scott, James C. Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed Yale University Press. Referred to in the schedule as **JS.**

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**.

Monogan III, James E. 2015. Political Analysis Using R , Springer. http://link.springer.com/book/10.1007%2F978-3-319-23446-5 .

Referred to in the schedule as **M3** .

Course Outline

1. Introduction to machine learning and artificial intelligence.
2. Overview of algorithms and the administrative state.
3. Algorithms, administrative behavior and decision-making.
    - Decision making by humans and machines.
    - Strength and weaknesses of machine learning and AI.
    - Machine learning and fairness.
4. Algorithms and political institutions.
    - Bureaucratic discretion.
    - Accountability.
    - Centralization/decentralization.

COURSE SCHEDULE

Week 1: *Introduction to Machine Learning and Artificial Intelligence: Fundamentals*

Programming fundamentals.

- Introduction to programming in **R**.
- **R** markdown and notebooks.
- APIs and webscraping 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.
- **M3** Chapters 1, 2, 10, 11.1-11,4.
- 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.

Week 2: *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.

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

Week 3: *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

Readings

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

Week 4: *Understanding unsupervised machine learning through examples: k-means clustering.*

- Understanding unsupervised learning through k-means clustering.

Readings

- ❖ **JWHT** -- Chapter 10.1, 10.3.1.
- ❖ **EA** - Chapter 5.

Week 5: *Algorithms and the administrative state - overview.*

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

Readings

- ❖ **JS** - Introduction, Chapters 1 & 2.
- ❖ Coglianese and Lehr. 2017. “Regulating by Robot Administrative Decision Making in the Machine-Learning Era”. Georgetown Law Journal.

Week 6: *Algorithms, Administrative Behavior and Decision Making in Theory*

- Decision making by street level bureaucrats.
- Decision making by humans v. machines.

Readings

- ❖ *Lipsky, Toward a theory of street-level bureaucracy.*
- ❖ *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*

Week 7: *Algorithms, Administrative Behavior and Decision Making in Practice - Police and Judges*

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

Readings

- ❖ Buffat, Aurélien. "Street-level bureaucracy and e-government." *Public Management Review* 17.1 (2015): 149-161.
- ❖ Kleinberg, Jon, et al. "Human decisions and machine predictions." *The quarterly journal of economics* 133.1 (2017): 237-293.
- ❖ Harcourt, Bernard E. "Against prediction: Sentencing, policing, and punishing in an actuarial age." (2005).

Week 8: *Strengths and Weaknesses of Machine Learning Systems for Public Administration*

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

- ❖ *Breiman. Statistical Modeling: The Two Cultures*
- ❖ *Norvig. On Chomsky and the Two Cultures of Statistical Learning*
- ❖ *Lazer et al. The parable of Google Flu.*
- ❖ *Olteanu et al. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries*

Week 9: *Machine Learning and Bias I: Overview*

- Defining bias in the machine learning context.

- ❖ *Harrits, Gitte Sommer. "Stereotypes in Context: How and When Do Street-Level Bureaucrats Use Class Stereotypes?." Public Administration Review (2018).*
- ❖ *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
- ❖ *Corbett-Davies et al. Algorithmic Decision Making and the Cost of Fairness.*

Week 10: *Machine Learning and Bias II: Sources and Pathways in Practice*

- Machine learning bias in practice.

❖ *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]*

Week 11: *Machine learning, big data and ethics: the modern panopticon?*

- Ethical considerations of big data and machine learning.

❖ *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]*

Week 12: *Bureaucratic discretion and technological change.*

Readings TBD

Week 13: *Machine learning and political institutions: accountability and transparency*

❖ Coglianese and Lehr, Transparency and Algorithmic Governance
❖ *Sandvig, Hamilton, Karahalios, and Langbort, [Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms]*
❖ *Ananny & Crawford, [Seeing without knowing Limitations of the transparency ideal and its application to algorithmic accountability]*

Week 14: *Machine learning and political institutions: centralization, machine learning and AI systems*

❖ Waldo, The Administrative State Chapter 8: Centralization versus Decentralization
❖ Bloomfield, Brian P., and Rod Coombs. "Information technology, control and power: The centralization and decentralization debate revisited." *Journal of management studies* 29.4 (1992): 459-459.

- ❖ Markus, M. Lynne, and Daniel Robey. "Information technology and organizational change: causal structure in theory and research." *Management science* 34.5 (1988): 583-598.
- ❖ Bretschneider, Stuart. "Information technology, e-government, and institutional change." *Public Administration Review* 63.6 (2003): 738-741.

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.