Syllabus: Governing the Algorithmic Society

Spring 2019
Professor Andrew D. Selbst
Tuesday 4–5:50 p.m., Room 4-04

Course Description:

Society is increasingly controlled by algorithms. New technologies based on finding patterns in data are used in choosing who gets jobs, credit, and housing, where to send police, how long people stay in prison, how speech is regulated and opinions influenced. How should society respond? What role does law play in responding to a world increasingly dominated by algorithmic decision-making? What must we understand about the relationship between technology and society in order to answer that question?

In this course, we will explore how the ubiquity of algorithmic decisionmaking challenges some of our most fundamental stated and unstated assumptions about what law is and how it works, while asking some more concrete questions related to existing technology and corresponding law and policy. We will begin with a brief introduction to theories of technology, society, and law, which will serve as the analytical frame for the semester. In the following weeks, we will move through different social and legal domains, such as employment, policing, trials, and injuries from autonomous vehicles, to understand how and why algorithmic decision-making is challenging to govern. The readings will draw on cutting edge research in law, computer science, social science, and social theory, as well as contemporary news articles and opinion pieces. By the end of the semester, we will better understand not just the relationship between algorithms and society, but how to think about the governance of technology generally, which will become ever more important as new technologies develop in the future.

Assessment will primarily be based on class participation and a final paper. While a willingness to scrutinize the details of technology is required, no math is required or expected.

About Me:

I am an attorney and postdoctoral scholar at Data & Society Research Institute. For the last several years, I have been working with computer scientists, sociologists, anthropologists and other legal scholars to study the legal and social issues surrounding various aspects of algorithmic decisionmaking. The issues discussed in this seminar are therefore the central focus of my work. But while I have been thinking about these issues for a while, there are a great many open questions that I hope we can collaboratively think through over the course of the semester, and which you will all write about by the end.

In addition, before I went to law school, I was an electrical engineer. This means that I have enough technical knowledge not to be intimidated by technology, but I cannot write a machine

learning algorithm. I fully believe, then, that you do not need to be able to code well—or at all—to intimately understand this subject matter; you must only be willing to deeply engage with the technical concepts, and be able to examine technology with a critical eye.

The best way to reach me is by email: aselbst@fordham.edu. Office hours are by appointment only. My office is located at 36 W. 20th St., on the 11th Floor, and I will happily meet you there, or if I am meeting with many of you in a row (e.g. after your outline is due), then I may sit in a coffee shop near the school. I can also usually stay a few minutes after class to answer any brief or immediate questions.

Electronic Courseware

This course will use TWEN, where I will post the readings, you will post your reaction pieces, and we can all post interesting news articles and discussion that crop up from time to time. To me, one of the most fun things about this subject is that it's constantly in the news, so please do post things you see and we can discuss them in class.

Required Texts

There is no textbook for this course. Each week, there will be a reading assignment listed on the syllabus. The reading assignment may include journal articles, news articles, cases, or book chapters. I hope to post the readings on TWEN, but almost everything will be available online should you wish to retrieve it yourself; the law library website has access to non-legal journals.

Optional background reading:

- VIRGINIA EUBANKS, AUTOMATING INEQUALITY (2018)
- CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION (2016)
- FRANK PASQUALE, THE BLACK BOX SOCIETY (2015)
- WHITE HOUSE OFFICE OF SCIENCE AND TECHNOLOGY POLICY, BIG DATA: A REPORT ON ALGORITHMIC SYSTEMS, OPPORTUNITY, AND CIVIL RIGHTS (2016)

These texts are not required for the course, but they are good reads for those interested in this subject, and they are often referenced in popular media.

In addition, for assistance with the research paper, I recommend Eugene Volokh's book *Academic Legal Writing.* (It doesn't really matter which edition, though later is probably at least a little better.)

Learning Outcomes

I have two goals for you by the end of this semester. The first is that you intimately understand the social and legal consequences and challenges of the transformation society is going through with respect to algorithms and AI. The second, and perhaps even more important, is

that you understand how to critically interrogate new technologies, and understand at a high level the complicated relationship between law, technology, and society. This understanding is crucial so that you can apply the skills you learn here to future technologies that have not yet been invented or popularized.

Attendance Policy

We only meet thirteen times this semester, and class participation is essential to any seminar. Therefore, attendance is required, and as discussed below, class participation is part of your grade. If you miss a class, it must be for a good reason. As we are all professionals, I will take you at your word that you have such a reason; you need not tell it to me. You must, however, alert me at least twenty-four hours in advance of your absence, except in the case of true emergency. You are entitled to miss exactly one class this way; each following missed class will result in automatic one step deductions from your final grade (e.g. from an A to an A-, or an A- to a B+). Emergencies and other truly exceptional circumstances will be evaluated on an individual basis, at my discretion.

Laptop Policy

Laptops may be used in class, as long as everyone remains engaged in the discussion.

Grading Policy:

Assessment will be based on three components:

20% Class Participation
20% Weekly response papers
60% Final research paper

Weekly Response Papers:

Before each class, you must write a short (~1 page) response piece reflecting on the readings for that week, due by 10am on the day of class. These response pieces are meant to spur discussion, and will hopefully serve as the basis for the class discussion each week. They are not meant to be difficult; I just want to see that you are reading and thinking about the material. The readings will cover much more material than a one-page response can, so do not feel the need to cover all the different ideas from the readings, and do not do independent research. Full credit will be given if you appear to be making a good faith effort to read and reflect.

Life happens, and you may not be able to write a response paper each week. Each of you may skip up to two response papers, no questions asked, as long as you notify me by the deadline. Each missed response paper beyond that will result in a loss of one quarter of the response paper portion of your grade (5% of your overall grade).

Final Research Paper

The majority of your grade will be determined by an original research paper, due at the end of finals period. In it, you should do new factual and legal research to explore a topic related to algorithms and the law that we do not cover in class, or a new aspect of a topic we do cover. The paper should be about 10,000-12,000 words in length. This word count is not itself important, and I will not be counting; you should write exactly as many words as needed to make your arguments well. I am giving you an approximate word count solely to explain that I anticipate that something in that range will accomplish the task. If you greatly exceed that range, your claim should be large enough to justify the extra length used to defend it.

In the process of writing the research paper, you must complete two assignments earlier in the semester. Each of these must be submitted at the same time as the week's response paper.

Paper Proposal

For Class 5, you will submit a paper proposal. This should be about one page and no more than two, and will consist of a proposed abstract for your paper, identifying the topic you wish to explore, some open questions, and a rough sketch of the paper's expected thesis. For this assignment, you will need to do some initial research to understand what is out there already. I also expect this thesis to change as you research more, so do not worry about it being perfect; I just want to make sure you're on the right track.

In the following week, I will return the proposals with comments and may ask to meet with you if there is something I feel we should discuss.

Annotated Paper Outline and Initial Bibliography

For Class 10, you will submit an annotated paper outline and initial bibliography. By this time, you should have done substantial research and understand your argument decently well. I expect you to outline the paper at a level of detail such that I can understand your argument by reading the outline. Annotate the outline at the key points needed to make this happen. For the bibliography, you should provide ten sources and for each, a few sentences explaining what it is and how it fits in the argument of your paper.

In the following week or two, I will meet with each of you for half an hour to discuss the paper outlines and offer suggestions. It is imperative that you lay out the outline and annotation in a tight, organized manner. I'll have twenty of these to read; if you do not give me information in an organized enough way, I will not be able to help you in that time.

Class Schedule (subject to change):

Unfortunately, I must cancel class January 29 due to an unavoidable conflict. We will make up that class on a mutually agreed upon date. I initially propose that we meet at our normal time

on February 19, ignoring the Monday schedule, but I will ask the class when we first meet what will work for everyone.

- Jan. 15: Class 1
- Jan. 22: Class 2
- Jan. 29: CLASS CANCELLED
- Feb. 5: Class 3
- Feb. 12: Class 4
- Feb 19: MONDAY SCHEDULE; NO CLASS SCHEDULED.
    - Proposed: Make-up Class 5
- Feb. 26: Class 6
- Mar. 5: Class 7
- Mar. 12: Class 8
- Mar. 19: SPRING BREAK; NO CLASS
- Mar. 26: Class 9
- Apr. 2: Class 10
- Apr. 9: Class 11
- Apr. 16: Class 12
- Apr. 23: Class 13

Topics and Assignments:

These are the weekly reading and occasional paper-related assignments. For the reading, I recommend reading the materials in the order listed.

Class 1: The Law and Politics of Technology

Required Readings:

- Langdon Winner, *Do Artifacts Have Politics?* 109 DAEDALUS 121 (1980)
- Lawrence Lessig, *The Law of the Horse: What Cyberlaw Might Teach*, 113 HARV. L. REV. 501 (1999)
- Joy Buolamwini, *Algorithms Aren't Racist. Your Skin Is just too Dark*, HACKERNOON (May 29, 2017)

Recommended Additional Readings:

- GEOFFREY C. BOWKER & SUSAN LEIGH STAR, SORTING THINGS OUT: CLASSIFICATION AND ITS CONSEQUENCES 1-32 (1999) (*Introduction*)

Class 2: Introduction to Algorithms, Machine Learning, and Artificial Intelligence

Required Readings:

- David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 669–702 (2017) (Part II)
- M.C. Elish & danah boyd, *Situation Methods in the Magic of Big Data and AI*, COMMUNICATION MONOGRAPHS 57 (2018)
- Oscar Schwartz, *"The Discourse Is Unhinged": How the Media Gets AI Alarmingly Wrong* THE GUARDIAN (July 25, 2018)

Recommended Additional Readings:

- Pedro Domingos, *A Few Useful Things to Know About Machine Learning* (skim)
    - Parts of this may be challenging from a technical perspective, but it's short and contains useful nuggets.
- David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 702–17 (2017) (Part III).

Class 3: Introduction to Algorithmic Discrimination

Required Readings (and one Viewing):

- Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & M. L. REV. 857, 869–74 (2017) (Part I.A)
- Solon Barocas and Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV 671, 671-713 (2016) (read Introduction; skim Part I, as a more expansive taxonomy was covered in Lehr & Ohm; read Part II)
- Deborah Hellman, What Is Discrimination, When Is It Wrong and Why?, Keynote at 2018 FAT* Conference (watch until 51:20—presentation only)

Recommended Additional Readings:

- Tal Z. Zarsky, *Understanding Discrimination in the Scored Society*, 89 Wash. L. Rev. 1375 (2014)
- ROBINSON + YU, KNOWING THE SCORE: NEW DATA, UNDERWRITING, AND MARKETING IN THE CONSUMER CREDIT MARKETPLACE (2014)
    - This report discusses alternative credit models. While we'll mostly talk about employment, U.S. law's approach to credit discrimination is similar, and credit is one of the main topics in which algorithmic discrimination is discussed, so it is useful to know what's out there.

Class 4: Criminal Justice Risk Assessments

Required Readings:

- Julia Angwin, Jeff Larson, Surya Mattu & Lauren Kirchner, *Machine Bias*, PROPUBLICA (May 23, 2016)
- *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016)
- Jessica Eaglin, *Constructing Recidivism Risk*, 67 Emory L.J. 59, 88-104 (2018) (Part II)
- Christopher Slobogin, *Principles of Risk Assessment: Sentencing and Policing*, 15 OHIO ST. J. CRIM. L. 583, 589–93 (2018) (PART II.C)
- Rebecca Wexler, *When a Computer Program Keeps You in Jail*, N.Y. TIMES (June 13, 2017)

Recommended Additional Readings:

- Sonja B. Starr, *Evidence-Based Sentencing and the Scientific Rationalization of Discrimination*, 66 STAN. L. REV. 803 (2014)
- Sandra G. Mayson, *Bias In, Bias Out*, 128 Yale L.J. __ (forthcoming 2019)

Class 5: Predictive Policing

Assignment: Paper Proposal Due

Required Readings:

- UPTURN, STUCK IN A PATTERN: EARLY EVIDENCE ON "PREDICTIVE POLICING" AND CIVIL RIGHTS (2016)
- Andrew E. Taslitz, *What Is Probable Cause, and Why Should We Care?: The Costs, Benefits, and Meaning of Individualized Suspicion*, 73 LAW & CONTEMP. PROBS., Summer 2010, at 145, 145–172 (Parts I & II)
- Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. Pa. L. Rev. 327, 376–88 (2016) (Part III)
- Whren v. United States, 517 U.S. 806 (1996)

Recommended Additional Readings:

- Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. Pa. L. Rev. 327 (2016) (the rest)
- CLAIRE GARVIE ET AL., THE PERPETUAL LINE-UP (2016)
    - We may not discuss facial recognition specifically, but facial recognition is going to play a big role in policing—and predictive policing—soon enough, so it might be of interest.

Class 6: Transparency and Explanations in Private Sector Algorithms

Required Readings:

- Jenna Burrell, *How the Machine 'Thinks'*, BIG DATA & SOC'Y (2016)
- Zack Lipton, *The Mythos of Model Interpretability*, PROCEEDINGS OF THE 2016 ICML WORKSHOP ON HUMAN INTERPRETABILITY IN MACHINE LEARNING (WHI 2016)
- Ed Felten, *What Does It Mean to Ask for an "Explainable" Algorithm?*, FREEDOM TO TINKER (May 31, 2017)
- SARAH AMMERMANN, ADVERSE ACTION NOTICE REQUIREMENTS UNDER THE ECOA AND THE FCRA
- Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 FORDHAM L. REV. 1085, 1099–1106 (2018) (Part II.A)
- Margot Kaminski, *The Right to Explanation, Explained*, 34 BERKELEY TECH. L.J. __ (forthcoming 2019)

Recommended Additional Readings:

- Mike Ananny & Kate Crawford, *Seeing Without Knowing*, 20 NEW MEDIA & SOC'Y 973 (2018)
- Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014)

Class 7: Transparency and Explanations in Government Algorithms

Required Readings:

- AI NOW, ALGORITHMIC ACCOUNTABILITY POLICY TOOLKIT 7-9 (examples of algorithmic systems used in government)
- Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward A Framework to Redress Predictive Privacy Harms*, 55 B.C. L. Rev. 93, 111–21, 124–28 (2014) (Part II.B-II.D; Part III.B)
- Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1278–1300 (2008) (Part II)
- Virginia Eubanks, *A Child Abuse Prediction Model Fails Poor Families*, WIRED (Jan. 15, 2018)

Recommended Additional Readings:

- AI NOW, LITIGATING ALGORITHMS: CHALLENGING GOVERNMENT USE OF ALGORITHMIC DECISION SYSTEMS (2018)
- Cary Coglianese And David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147 (2017)

Class 8: Speech, Algorithms, and the Digital Public Sphere

Required Readings:

- Zeran v. America Online, Inc., 129 F.3d 327, 328 (4th Cir. 1997)
- Philip Oltermann, *Tough New German Law Puts Tech Firms and Free Speech in Spotlight*, THE GUARDIAN (Jan. 5, 2018)
- An Act to Improve Enforcement of the Law in Social Networks (Network Enforcement Act)
    - This is a translation of NetzDG, the German law that Oltermann discusses.
- Drew Harwell, *AI Will Solve Facebook's Most Vexing Problems, Mark Zuckerberg Says. Just Don't Ask When or How.*, WASH. POST (Apr. 11, 2018)
- Kate Klonick, *Facebook Released Its Content Moderation Rules. Now What?*, N.Y. TIMES (Apr. 20, 2018)
- James Grimmelmann, *The Platform Is the Message*, 2 GEO. TECH. L. REV. 217 (2018)
- Miami Herald Pub. Co. v. Tornillo, 418 U.S. 241, 2837 (1974)

Recommended Additional Readings:

- CENTER FOR DEMOCRACY & TECHNOLOGY, MIXED MESSAGES? THE LIMITS OF AUTOMATED SOCIAL MEDIA CONTENT ANALYSIS (2017)
- Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 HARV. L. REV. 1598 (2018)
- ROBYN CAPLAN, CONTENT OR CONTEXT MODERATION? ARTISANAL, COMMUNITY-RELIANT AND INDUSTRIAL APPROACHES (2018)
- Jack M. Balkin, *Free Speech in the Algorithmic Society: Big Data, Private Governance, and New School Speech Regulation* 51 U.C. DAVIS L. REV. 1149 (2018)

Class 9: Persuasion, Manipulation, and Consumer Protection

Required Readings:

- Zeynep Tufekci, *Engineering the Public*, FIRST MONDAY, July 2014
- Woodrow Hartzog, *Unfair and Deceptive Robots*, 74 MD. L. REV. 785, 789–807, 811–23 (2015) (Parts I.A–I.C & II)
- Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS. L. REV. 1183, 1205–09, 1221–25 (Parts III & V)
- Faye Flam, *Cambridge Analytica Knew How to Sell … Cambridge Analytica*, BLOOMBERG (Mar. 28, 2018)

Recommended Additional Readings:

- Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1007–24, 1031–34 (2014) (Parts II.A-C, III.A, III.B.3)
- James Grimmelmann, *Speech Engines*, 98 MINN. L. REV. 868 (2014)
- Daniel Susser, Beate Roessler, Helen Nissenbaum, *Online Manipulation* (forthcoming)

Class 10: Robots, Algorithms, and Tort Law

Assignment: Annotated Paper Outline and Initial Bibliography Due

Required Readings:

- Curtis Karnow, *The Application of Traditional Tort Theory to Embodied Machine Intelligence, in* ROBOT LAW 51, 59–74 (Ryan Calo, A. Michael Froomkin & Ian Kerr, eds. 2016) (Parts 2.2–3)
- Ryan Calo, *Is Law Ready for Driverless Cars?*, 61 COMM. ACM 34 (2018)
- Jonathan Dyble, *Understanding SAE Automated Driving – Levels 0 To 5 Explained*, GIGABIT (Apr. 23, 2018)
- David G. Owen, *Figuring Foreseeability*, 44 WAKE FOREST L. REV. 1277, 1281–82, 1291–1306 (2009) (Part II.A & Part III)
    - Read this carefully. It's subtle, but important to the arguments about algorithmic harms and foreseeability.
- Johana Bhuiyan, *Uber's Self-Driving Software Detected the Pedestrian in the Fatal Arizona Crash but Did not React in Time*, RECODE (May 7, 2018)
- Angela Chen, *IBM's Watson Gave Unsafe Recommendations for Treating Cancer*, THE VERGE (Jul. 26, 2018)
- W. Nicholson Price II, *Medical Malpractice and Black-Box Medicine*, *in* Big Data, Health Law, and Bioethics 295 (I. Glenn Cohen, Holly Fernandez Lynch, Effy Vayena & Urs Gasser eds. 2018).

Recommended Additional Readings:

- Kyle Graham, *Of Frightened Horses and Autonomous Vehicles: Tort Law and Its Assimilation of Innovations*, 52 SANTA CLARA L. REV. 1241 (2012)
- Andrew D. Selbst, *Negligence and AI's Human Users* (draft to be posted on SSRN)

Class 11: Algorithmic Management of Workers (Guest Speaker: Alex Rosenblat)

Required Readings:

- ALEX ROSENBLAT, UBERLAND: HOW ALGORITHMS ARE REWRITING THE RULES OF WORK 1-20, 138–166 (2018) (skim the Introduction; read Chapter 5).

- Solon Barocas & Karen Levy, *What Customer Data Collection Could Mean for Workers*, HARV. BUS. REV. (Aug. 31, 2016)
- Benjamin Means & Joseph A. Seiner, *Navigating the Uber Economy*, 49 U.C. DAVIS L. REV. 1511, 1524–35, 1539–45 (2016) (Parts II & III.B)
- Complaint, *FTC v. Uber*, No. 3:17-cv-00261 (N.D. Cal. 2017)

Recommended Additional Readings:

- ALEX ROSENBLAT, UBERLAND: HOW ALGORITHMS ARE REWRITING THE RULES OF WORK (2018) (the rest)
    - I did not want to assign the whole book, but it's a fun, engaging, and relatively quick read if you're interested, and Alex will be joining us for the first half of class.
- Ryan Calo & Alex Rosenblat, *The Taking Economy: Uber, Information, and Power*, 117 COLUM. L. REV. 1623 (2017)

Class 12: Proposals for Oversight of the Algorithmic Society

Required Readings:

- Christian Sandvig, Kevin Hamilton, Karrie Karahalios & Cedric Langbort, *Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms*, DATA AND DISCRIMINATION: CONVERTING CRITICAL CONCERNS INTO PRODUCTIVE INQUIRY (2014)
- AI NOW, ALGORITHMIC IMPACT ASSESSMENTS: TOWARD ACCOUNTABLE AUTOMATION IN PUBLIC AGENCIES (2018)
- AI in Government Act of 2018, S. 3502, 115th Cong.
- Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83, 105–11, 119–22 (2017) (Parts II and IV)

Recommended Additional Readings:

- RYAN CALO, THE CASE FOR A FEDERAL ROBOTICS COMMISSION (2014)
- ACLU Press Release, *First Amendment Lawsuit Brought on Behalf of Academic Researchers Who Fear Prosecution Under the Computer Fraud and Abuse Act* (Apr. 2 2018)

Class 13: Using Technology to Solve Sociotechnical Problems

Required Readings:

- Julia Angwin, et al., *Bias in Criminal Risk Scores Is Mathematically Inevitable*, PROPUBLICA (Dec. 30, 2016)

- Alexandra Chouldechova, et al, *A Case Study Of Algorithm-Assisted Decision Making In Child Maltreatment Hotline Screening Decisions*, 2018 CONFERENCE ON FAIRNESS, ACCOUNTABILITY AND TRANSPARENCY (FAT*) (skim the technical parts)
    - This is about the same system that Eubanks discussed in the Week 7 reading.
- Andrew D. Selbst et al., *Fairness and Abstraction in Sociotechnical Systems*, 2019 ACM Conference on Fairness, Accountability and Transparency (FAT*)
- Deven R. Desai & Joshua A. Kroll, *Trust but Verify: A Guide to Algorithms and the Law*, 31 HARV. J. L. & TECH. 1, 35–55 (Parts V, VI.A–B).

Recommended Additional Readings/Viewings:

- Arvind Narayanan, Tutorial, *21 Fairness Definitions and Their Politics*, 2018 CONFERENCE ON FAIRNESS, ACCOUNTABILITY AND TRANSPARENCY (FAT*)
    - Video of a tutorial at the FAT* conference
- Sharad Goel, Maya Perelman, Ravi Shroff & David Alan Sklansky, *Combatting Police Discrimination in the Age of Big Data,* 20 NEW CRIM. L. REV. 181 (2017)