



**NYU**

**ROBERT F. WAGNER GRADUATE  
SCHOOL OF PUBLIC SERVICE**

## **PADM-GP 4700**

# Topics in Public Policy: Data, Evidence, Ethics, and Bias in an AI World Spring 2022

### **Instructor Information**

- Julia Lane
- Email: [jl4@nyu.edu](mailto:jl4@nyu.edu)
- Office Address: NYU-Wagner, Room 3038, 295 Lafayette Street, New York, NY 10012
- Office Hours: Virtual office hours can be made by appointment

### **Course Time and Location**

- Lecture: Tuesdays, 4:55-6:35pm

### **Course Description and Learning Objectives**

The importance of data and evidence in making decisions is critical to success in both the private and the public sector. The role of Artificial Intelligence (AI) in making sense of the vast amounts of data available to make evidence-based decisions has been highlighted by the establishment of a National Artificial Intelligence Initiative and by its pervasive use in virtually all parts of the economy and society. Yet there is potentially a dark side to the implementation of AI, particularly given its widespread use in criminal justice and immigration. This course will provide students with an overview of current approaches to examining potential biases and ethical challenges in that implementation, with a particular focus on data collection, model development and evidence building. It will focus in on particular use cases in social science, transportation, climate and industry.

The instructor serves on the AI Research Resources Task Force (NAIRRTF) as well as the Advisory Committee on Data for Evidence Building (ACDEB).

## Learning Assessment Table

Program Competencies	Corresponding Course Learning Objective	Corresponding Assignment Title
Understanding the current landscape	Characterizing the size, complexity, and variety of data The role of AI in data and evidence Defining bias and ethics	Summary of NAIRR charge Review of use case of interest
AI in practice	The role of record linkage Creating training and test data Validation	Analysis of specific use case of interest
Measuring bias	Using open source bias audit toolkit for machine learning developers, analysts, and policymakers to audit machine learning models	Audit toolkit exercise <a href="http://www.datasciencepublicpolicy.org/our-work/tools-guides/aequitas/">http://www.datasciencepublicpolicy.org/our-work/tools-guides/aequitas/</a>
Understanding privacy and confidentiality	Legal framework for access to and use of confidential data Current approaches (differential privacy, tiered access, homomorphic encryption, secure multiparty computation)	Summary of Advisory Committee for Data on Evidence Building charge Privacy literature review
Putting it all together	Synthesis of the current state of play and options for the future	Final project

## Housekeeping

- The NYU Classes site for this course will contain the lecture slides, additional reading materials, and assignments. In addition, all lectures are recorded and available on NYU classes after each session. Notifications and updates will be sent out through NYU classes on a regular basis.
- Punctuality is **very important**. We realize unforeseen circumstances arise, but please try to be on time. Disruptions affect not only us, but your fellow classmates as well.
- Active participation is a part of your overall grade. This means asking questions in chat, responding to questions when asked, and sharing information you come across that might be interesting for your classmates.
- We expect you to be prepared for class discussions and to keep up with what we have done in prior classes. The open exchange of ideas will be respected by all students. Respectful and inclusive discussion is required.
- Grades on assignments and class projects are non-negotiable.
- Late assignments are accepted. If you submit an assignment after the posted deadline it will be counted as late and will be penalized (see evaluation section). You can always turn an assignment in early to avoid penalties. There are no make-up assignments.

## Required Scholarly Activities

This is a graduate course, so I assume that you have the self-motivation and discipline to keep up with the readings on your own. This is a quickly moving area of interest, so you will be pointed to more readings during lectures as well as parts of the meetings of the NAIRRTF and the ACDEB.

## Course Structure

The course will be structured in weekly sessions. They will primarily be structured as lectures and active question and answer discussions. I am also likely to invite guest speakers from government, academia, and industry to provide brief overviews of both the challenges they face in their domains, and how they are currently addressed.

### Session 1: Introduction and overview

- Date: 03/22/2022
- Lecture:
  - Organizational details for class/housekeeping
  - Motivation and current activities
  - Definitions in theory and practice
- Readings
  - Review <https://www.ai.gov/nairrtf/>
  - <https://www.bea.gov/evidence> and <https://www.bea.gov/system/files/2021-10/acdeb-year-1-report.pdf>
  - Yarkoni, T., Eckles, D., Heathers, J. A., Levenstein, M. C., Smaldino, P. E., & Lane, J. (2021). Enhancing and accelerating social science via automation: Challenges and opportunities. *Harvard Data Science Review*.
  - [Implicit bias test](#)
  - Chapter 10 and 11 of <https://textbook.coleridgeinitiative.org/>
  - Responses to NAIRR TF Request for Information
    - [AI Now](#)
    - [American Civil Liberties Union](#)
    - [Center for Democracy and Technology](#)
    - [Consumer Reports](#)
    - [Electronic Privacy Information Center](#)
    - [Partnership on AI](#)
- Assignment 1: Identify use case of interest and explain why (due in four weeks)

### Session 2: Combining data sources to create evidence

- Date: 03/29/2022
- Lecture:
  - Understanding the Problem
  - The conceptual framework
  - AI approaches
  - Bias and ethics issues
- Readings:
  - Chapter 3 <https://textbook.coleridgeinitiative.org>

### Session 3: Data, Evidence, and AI

- Date: 04/05/2022
- Lecture:
  - Policy issues, data, and the need for AI
  - Generating training and test datasets
  - Challenges of curation, data management and governance
- Readings:
  - NAIRR TF meetings #1- #4
  - Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition (pp. 248-255). Ieee.
  - <https://www.kaggle.com/c/coleridgeinitiative-show-us-the-data>
- Assignment 3: Analysis of use case (due in four weeks)

### Session 4: Bias and ethics

- Date: 04/12/2022
- Lecture:
  - Conceptual framework
  - Empirical framework
  - Audit toolkit
- Readings
  - Hearing: Task Force on Artificial Intelligence: Equitable Algorithms: Examining Ways to Reduce AI Bias in Financial Services
    - <https://docs.house.gov/meetings/BA/BA00/20200212/110499/HHRG-116-BA00-Wstate-GhaniR-20200212-U1.pdf>
  - <http://www.datasciencepublicpolicy.org/our-work/tools-guides/aequitas/>
- Assignment 4
  - Generate an audit report using your own data or the sample data provided in Aequitas. Discuss the results (due in four weeks)

### Session 5: Bias and Ethics (continued)

- Date: 04/19/2022
- Lecture:
  - Definitional challenges
  - Evaluation challenges
  - The importance of access
- Readings:
  - Gibney, E. (2020). The battle for ethical AI at the world's biggest machine-learning conference. *Nature*, 577(7791), 609-610. <https://www.nature.com/articles/d41586-020-00160-y>
  - Raghu, M., Blumer, K., Corrado, G., Kleinberg, J., Obermeyer, Z., & Mullainathan, S. (2019). The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*. <https://arxiv.org/pdf/1903.12220.pdf>

### Session 6: Privacy and confidentiality

- Date: 04/26/2022

- Lecture:
  - Privacy/utility tradeoff
  - Different approaches
  - Current challenges
- Readings:
  - Chapter 12 of <https://textbook.coleridgeinitiative.org>
  - Jessica Hullman „Shots taken, shots returned regarding the Census’ motivation for using differential privacy (and btw, it’s not an algorithm)“ <https://statmodeling.stat.columbia.edu/2021/08/27/shots-taken-shots-returned-regarding-the-census-motivation-for-using-differential-privacy-and-btw-its-not-an-algorithm/>
  - Lane, J., Stodden, V., Bender, S., & Nissenbaum, H. (2014). Privacy, big data and the public good: Frameworks for engagement. Cambridge University Press.
  - Modernizing U.S. Data Infrastructure: Design Considerations for Implementing a National Secure Data Service to Improve Statistics and Evidence Building <https://www.datafoundation.org/modernizing-us-data-infrastructure-2020>

### Session 7: Final Project Presentations

- Date: 05/03/2022
  - Students present their final project

## Evaluation

### Project work

Students will work on their own research project during the entire semester. That project is supported by the individual assignments, and consists of 20% of the final grade.

### Assignments

You are required to complete 4 assignments throughout the class. The assignments constitute 40% of the grade: Please submit your assignments on time. Grades will be reduced 25% for each day that it is late – after four days, late assignments will receive no credit. Please turn in your assignment early if there is any uncertainty about your ability to turn it in on time.

### Class preparation and participation

Active participation in class and shared summaries will constitute 40% of the final grade (examples: participation in class and group discussions, posting on NYU Classes forum, responding to questions when asked, sharing information you come across that might be interesting for your classmates).

The breakdown of the evaluation activities:

Activity	Proportion of Grade
<b>Project work</b>	<b>20%</b>
Presentation	10%
Research memo (10 pages)	10%
<b>Assignments</b>	<b>40%</b>

Assignment 1: Review of use case	10%
Assignment 2: Analysis of use case	10%
Assignment 3: Audit toolkit	10%
Assignment 4: Privacy literature review	10%
<b>Class preparation and participation</b>	<b>40%</b>
Shared summaries of NAIRRTF and ACDEB activities	20%
Active class participation	20%

## General guidance

All project and individual assignments should be posted on NYU Classes before the deadline. Answers to the assignments should be well thought out and communicated precisely, as if reporting to your boss, client, or potential funding source. Avoid sloppy language, poor diagrams, irrelevant discussion, and irrelevant program output.

If you prepare and participate in the course you should be able to work on the assignments without major problems. But we all experience problems that we can't figure out right away. If you get stuck on something while preparing for class or working on the assignments, spend some time Googling to try to find the answer. If you seem to be moving forward, keep going. That search and discovery method will pay off, both in terms of the direct learning about how to do what you need to do, and also in terms of your learning how to find such things out.

However, in order to limit frustrations with class work we advise you to start your assignments early enough that if you experience problems without finding an answer, you still have enough time to ask about it.

If you are stuck after 30 minutes, just stop and ask your classmates or post on the forum on NYU classes. All class participants have access to this and can help you with your questions. You will most likely encounter the same problems as your peers. The forum is there for you to ask your peers for advice. If you don't find a solution, escalate it to the instructors.

## Plagiarism

All students must produce original work. Outside sources are to be properly referenced and/or quoted. Lifting copy from websites or other sources and trying to pass it off as your original words constitutes plagiarism. Such cases can lead to academic dismissal from the university.

## Academic Integrity

Academic integrity is a vital component of Wagner and NYU. All students enrolled in this class are required to read and abide by [Wagner's Academic Code](#). All Wagner students have already read and signed the [Wagner Academic Oath](#). Plagiarism of any form will not be tolerated and students in this class are expected to report violations to me. If any student in this class is unsure about what is expected of you and how to abide by the academic code, you should consult with me.

## **Henry and Lucy Moses Center for Students with Disabilities at NYU**

Academic accommodations are available for students with disabilities. Please visit the [Moses Center for Students with Disabilities \(CSD\) website](#) and click on the Reasonable Accommodations and How to Register tab or call or email CSD at (212-998-4980 or [mosescsd@nyu.edu](mailto:mosescsd@nyu.edu)) for information. Students who are requesting academic accommodations are strongly advised to reach out to the Moses Center as early as possible in the semester for assistance.

## **NYU's Calendar Policy on Religious Holidays**

[NYU's Calendar Policy on Religious Holidays](#) states that members of any religious group may, without penalty, absent themselves from classes when required in compliance with their religious obligations. Please notify me in advance of religious holidays that might coincide with exams to schedule mutually acceptable alternatives.