DS-UA 202: Responsible Data Science (Spring 2024)

Lecture: Tuesday/Thursday 6:20-7:35pm at 6 Washington Pl (Meyer Hall) Room 121

Lab Section 005: Fridays at 7 East 12th St Room 121, 9:00-9.50am
Lab Section 004: Fridays at 7 East 12th St Room 121, 10:15-11.05am
Lab Section 003: Fridays at 7 East 12th St Room 121, 11:15-12.05pm
Lab Section 002: Fridays at 7 East 12th St Room 121, 12:30-1.20pm
Lab Section 006: Fridays at 40 W 4th St (Tisch Hall) Room LC9, 2:45-3:35pm

Instructor: Umang Bhatt (umangbhatt@nyu.edu)

Teaching assistants: See BrightSpace for office hours
Andrew Bell (alb9742@nyu.edu), section leader
Raphael Meyer (ram900@nyu.edu), section leader
Aradhita Bhandari (aradhita.bhandari@nyu.edu)
Marcia Ma (xm618@nyu.edu), grader
Venetia Pliatsika (venetia@nyu.edu), grader

The instructor and the teaching assistants are also available by appointment.

Prerequisites: Successful completion of these courses, or approval of the instructor:

- DS-UA 112: Introduction to Data Science

Course description and impact

The first wave of data science focused on accuracy and efficiency: on what we *can* do with data. The second wave is about responsibility: what we *should* and *should not* do. Accordingly, this technical course tackles the issues of ethics and responsibility in data science, including legal compliance, data quality, algorithmic fairness and diversity, transparency of data and algorithms, privacy, and data protection.

Data science promises to improve people's lives, accelerate scientific discovery and innovation, and bring about positive societal change. Yet, if not used responsibly—in accordance with ethical and moral norms, and legal and policy considerations—this same technology can cause harm on an unprecedented scale. Algorithmic changes in search engines can sway elections and incite violence; irreproducible results can influence global economic policy; models based on biased data can legitimize and amplify racist policies in the criminal justice system; algorithmic hiring practices can silently and scalably violate equal opportunity laws, exposing companies to lawsuits and reinforcing the feedback loops that lead to lack of diversity, which is both socially undesirable and can negatively impact performance of organizations. These strategic issues become more important as the economy globalizes. Therefore, as we develop and deploy data science methods, we are compelled to think about the effects these methods have on individuals, population groups, and on society at large.

The European Union recently enacted the General Data Protection Regulation (GDPR) that mandates legal protections of data subjects on the part of government entities and companies that employ algorithms and data to make decisions. The US is following suit with a plethora of local efforts, including a recently passed algorithmic transparency law in New York City that applies to City agencies. These legal frameworks, and many others that will soon follow, compel us to develop skills and acquire methodologies for operationalizing responsibility.

The course has four modules: (1) Fairness; (2) Data Science Lifecycle; (3) Transparency and Interpretability; and (4) Data Protection. An important feature of this course is its holistic treatment of the data science lifecycle, beginning with data discovery and acquisition, through data cleaning, integration, querying, analysis, and result interpretation. For example, when considering fairness and diversity in hiring and college admissions, we will be looking at how the data was collected, cleaned and otherwise pre-processed, before analyzing the fairness of a particular classification or ranking method.

Learning objectives

After successfully completing the course, students are able to:

- Construct an end-to-end case study that illustrates the role of data science in society.
- Explain the ethical and/or legal constraints in the collection and sharing of data according to a framework of the student's choice.
- Implement a computer program that applies anonymization and privacy techniques to a dataset, and explain the trade-offs with utility.
- Articulate the differences between various interpretations of algorithmic fairness, and relate these interpretations to the points of view of different stakeholders.
- Implement a computer program that audits a black-box classifier.

Topics covered

The course is structured into a sequence of lectures and accompanying assignments. The assignments include of labs, quizzes, homeworks, and a course project:

- Labs are short exercises done in class and submitted in class.
- Homework assignments are longer exercises designed to take one to two weeks to complete.
- The course project is designed to take three weeks to complete.

This course uses the python programming language. All labs and homework assignments are expected to be formulated as notebooks in [Google Colaboratory].

While none of the homeworks are stand-alone writing assignments, writing is embedded in the homeworks. During the homeworks, students write computer programs and provide short essays on the interpretation of the data and the implications for the data around some decision or problem that the data informs. The writing length would typically be one to two pages for these essays.

This course does not have a required textbook. Each topic will be accompanied by required reading, as listed in the weekly schedule. In some cases, expert-level technical research papers are listed as assigned reading.

Weekly topics and reading are posted on the course website.

Course assessment

All assignments (labs and homeworks) must be entirely the student's own submissions. Any sharing or copying of assignments is considered cheating and will result in an F in the course. A second cheating incident will result in a one-semester suspension from the program.

Students accumulate up to 100 points during the course:
- **Labs: 10 points.** 11 labs x 1 points per lab = 11 points. 10 points sufficient for full credit.
- **Quizzes: 10 points.** 6 quizzes x 2 points per quiz = 12 points. 10 points sufficient for full credit.
- **Homeworks: 3 homeworks x 10 points per homework = 30 points.** Homeworks are assigned on Monday in class, and due before class, at 10am on Monday. Homeworks must be submitted on time. If a homework is submitted late, the student will receive no credit. Each student can make use of 2 additional days for late homework submission. If a late day is used, it is used in full. That is, if a student submits homework 2 hours late, this counts as a full day.
- **Project: 30 points.**
- **Final exam: 20 points.**

Grades will be determined using this scale:

Grade in Course	Points Earned
A	94 - 100
A-	90 - 93
B+	87 - 89
B	84 - 86
B-	80 - 83

C+	76 - 79
C	70 - 75
C-	65 - 69
F	Less than 65

Moses statement

Disability Disclosure Statement: Academic accommodations are available for students with disabilities. The Moses Center website is www.nyu.edu/csd. Please contact the Moses Center for Students with Disabilities (212-998-4980 or mosescsd@nyu.edu) for further information. Students who are requesting academic accommodations are advised to reach out to the Moses Center as early as possible in the semester for assistance.

Academic integrity and honesty

All students are expected to do their own work. Collaboration on homework assignments, quizzes and the exam is disallowed. For the course project, students are expected to collaborate with their team partner only. Collaboration with students other than those on their team is *disallowed*. Questions regarding acceptable collaboration should be directed to the class instructor prior to the collaboration. It is a violation of the honor code to copy or derive solutions from other students (or anyone at all), textbooks, previous instances of this course, or other courses covering the same topics. Copying solutions from other students, or from students who previously took a similar course, is also clearly a violation of the honor code. Finally, a good point to keep in mind is that you must be able to explain and/or re-derive anything that you submit. This is particularly important if you should adapt solutions from online sources. Please also refer to the general NYU academic integrity statement.

AI policy. We live in the age of viable generative AI. Banning these tools is neither realistic, nor desirable. In fact, learning to use these tools is an emerging skill. Note that AI tools do not always produce correct or accurate results. In addition, it is unwise to rely on them too much. There are situations where you won't have access to these tools, for instance during technical interviews. In addition, there are also skills someone with an advanced degree in Data Science is just expected to have on tap - without AI assistance or looking anything up. To integrate both considerations, you can use generative AI tools to do the assignments in this class, **except** the final exam (which will need to be done entirely without any electronic devices of any kind). If you use an AI to guide you in completing an assignment, you have to disclose which parts were generated by the AI. When in doubt about whether to disclose the use of AI - speak to the class instructor.