Byrne Seminar—Data Science: The Good, the Bad, and the Ugly (Spring 2018)

Administrative

Instructors:
Prof. Waheed U. Bajwa (waheed.bajwa@rutgers.edu)
Prof. Anand D. Sarwate (anand.sarwate@rutgers.edu)
Department of Electrical and Computer Engineering, Rutgers University

Class time: Wednesdays, 5–6:20 PM in SERC 216 for 10 weeks (Jan. 17 – Mar. 28)

Credit requirements: This 1-credit seminar is pass/no credit. To receive credit for the seminar, you have to:

- Attend at least nine of the ten seminars
  - In the case of a missed seminar, submit a one-page report in lieu of the missed seminar
- Complete assigned readings and engage in class discussions
- Complete a final team project (details below)

Course Overview

Data science is a hot field, and the term "machine learning" has moved into the popular culture. "Artificial intelligence" is no longer the subject of sci-fi movies alone: we regularly interact with "smart systems" that are powered by sophisticated learning and inference algorithms. There is no question that these systems have made great improvements in the efficiency of services and quality of life. On the flip side, the decisions made by machines reflect the biases (implicit or explicit) of their designers. The last ten years have truly been a "Decade of Discovery" in terms of advances in data collection and processing. But how can we navigate the potentials and pitfalls in the decades to come?

In order to gain some critical perspective on these topics, people have to learn basic concepts of "data science," much like we understand basic concepts of biology, chemistry, and physics. In this course, students will be formally introduced to the terms of "data science," "machine learning," and "artificial intelligence," they will learn the principles behind feature learning as well as problem formulation and algorithms for data science, and they will understand what goes into common data science systems, how these systems work, and how decisions in system design can be reflected in the outputs. The students will also be exposed to topics of contemporary research, such as interpretability, fairness, bias, and privacy.

Final Projects

This course involves a final project where you work in teams of 3–5 with other students enrolled in the course. The project could be:

- A proposal for a data-driven solution to a real-world problem
- A detailed critical analysis of a data-driven system not covered in class