CS 594: Responsible Data Science and Algorithmic Fairness

Fall 2020: T-Th 11:00 am - 12:15pm

- Instructor: Abolfazl Asudeh
- Contact: asudeh@uic.edu; www.cs.uic.edu/ asudeh
- Course page: www.cs.uic.edu/ asudeh/teaching/cs594fall20/
- Office hours: Th: 12:30am, 2:30pm (Skype: a.asudeh)

(Optional) Textbook and references

- Barocas, Solon, Moritz Hardt, and Arvind Narayanan. "Fairness in machine learning"
- Barocas & Selbst (2016) "Big Data's Disparate Impact"
- White House Report on Big Data (2014) "Big Data: Seizing Opportunities, Preserving Values"
- Kroll et al. (2017) "Accountable Algorithms"
- Cathy O'Neil "Weapons of Math Destruction"
- (Tutorial) A. Asudeh, H. V. Jagadish. Fairly Evaluating and Scoring Items in a Data Set. PVLDB, 13(12): 3445-3448, 2020

Rationale for the course

This course views data-driven and algorithmic decision making through the lens of data ethics and societal impacts. It shall cover the important aspects of the timely research area of responsible data science. The course will empower the graduate students with tools to start exploring/conducting research in this area.

Course Objective/description

The course aims at exploring different aspects of responsible data science and societal impacts of data-driven and algorithmic decision making. While a major focus of the course will be on Algorithmic Fairness, other aspects including transparency, accountability, equity, and stability. The course will cover the breadth from different aspects:

- Introduction and background: Course outline, aspects of responsibility in data science through recent examples.
- Data Ethics Terms: Fairness, Transparency, Accountability, Stability, Equity, and Diversity
- A Taxonomy of Fairness Definitions: Individual v.s. Group fairness, Statistical v.s. Societal Norms of Fairness, Fairness based on Model Independence v.s. Separation, Intersectional Fairness, Diversity as Fairness
- Impossibility results
- Fairness in Machine Learning Models; Fairness in Classification
- Fairness Beyond Classification: Fairness in Assignment, Human-designed models, Ranking, and Non-predictive Models
- Bias in Social Data
- Post-process techniques for achieving Fairness
- Data Preparation for achieving Fairness by Preprocess techniques
- In-process techniques for achieving Fairness
- Fairness and causality
- Stable Decision Making
- Dataset Investigation for Fairness: Coverage, Data Profiling, and Nutritional Labels

Class Organization

Method of instruction

- The first half of the class includes lectures given by the instructor. The instructor will provide the background and different aspects of responsible data science. This includes data ethics terms such as fairness, accountability, transparency, and stability.
- In second half, depending upon the size and background of the class students, students will be asked to read new research papers, summarize the papers and make a brief presentation in the class. The students are also required to select a research project on one of the related topics. Depending on students' interest the projects may range from implementing existing techniques to targeting open research problems. Should a publishable work develop, the instructor will provide guidance towards its eventual publication in a suitable journal or conference.

Evaluation/Grading

The students will be evaluated based on three criteria:

- Active Class Participation: 30%
- Presentation: 35%
- Final Project: 35+10%

Pre-requisites

The course will be accessible to students with a wide range of backgrounds, and we expect CS students working in a variety of applied areas to enroll for the course. The required pre-requisites are:

- CS 401-Algorithms I or equivalent: an undergraduate background in algorithms analysis
- CS 412 or CS 418 or equivalent: an undergraduate background in data science/machine learning