Fair decision making

Keywords: machine learning, statistics, fairness, reinforcement learning

Supervisors: Christos Dimitrakakis, christos.dimitrakakis@gmail.com

1 Background

Fairness is a desirable property of decision rules applied to a population of individuals. For example, college admissions should be decided on variables that inform about merit, but fairness may also require taking into account the fact that certain communities are inherently disadvantaged. At the same time, a person should not feel that another in a similar situation obtained an unfair advantage. All this must be taken into account while still caring about optimizing for a decision maker's utility function.

As another example, consider mortgage decisions: while lenders should take into account the creditworthiness of individuals in order to make a profit, society must ensure that they do not unduly discriminate against socially vulnerable groups. The problem becomes even more challenging when we take into account potential uncertainties in decision making models, which can make some notions of fairness impossible to satisfy.

This project will examine fairness in decision making for a topic of the student's choice, but the focus should always be in data driven problems. The thesis will be synergistic with our research with the ECON-CS group at Harvard. Our previous work on this topic includes *subjective fairness* [Dimitrakakis et al., 2017b] and *fairness in bandits* [Dimitrakakis et al., 2017a].

2 Goals

The project can be either theoretical or practical. For the former, a solid background in mathematics is necessary, as the student will attempt to develop and prove fairness properties of algorithms. For the latter, the student should be able to envisage an appropriate case study, such as mortgage decisions or university admissions, collect data for it. Mathematical maturity is still necessary, but the emphasis is more on research methodology and programming skills.

1. Formalize the decision problem.
2. Specify the fairness criteria used.
3. Develop an algorithm solving the problem while satisfying the fairness criteria.

3 Skills

The student should have good working knowledge of calculus and probability.

4 Benefits

The student will obtain background in statistics, machine learning and constrained optimisation, as well the links between these areas, as well as practical experience in C++ and/or optimality proofs. Good theoretical or experimental results will lead to writing and submitting a paper to a suitable peer-reviewed venue.

References

Christos Dimitrakakis, Yang Liu, , Debmalya Mandal, David Parkes, and Goran Radanovic. Calibrated fairness in bandits. In *Fairness, Accountability and Transparency in Machine Learning, at KDD*, 2017a.

Christos Dimitrakakis, Yang Liu, David Parkes, and Goran Radanovic. Subjective fairness: Fairness is in the eye of the beholder. Technical Report 1706.00119, arXiv, 2017b. URL https://arxiv.org/abs/1706.00119.