

IsyE 8813: Algorithmic Foundations of Ethical ML - Final Project

Proposal: Tuesday, Nov 9th. Class presentations: Dec 7th. Write-ups: December 14th.

Your project proposal should be emailed to jziani3@gatech.edu.

The course project will involve engagement with a current topic of research in data privacy. You can choose to do the project by yourself or with a partner. You should strive to produce something novel, such as new experiments, algorithms, models, or theorems. The project involves several deliverables:

1. A proposal describing related work and the intended topic and scope of the project.
2. A brief in-class presentation (10 minute per person, 20 if you are in a group of 2).
3. The final write-up: a short technical report describing the problem you considered, related work, your approach, and your results.

Your project proposal should be emailed to me by the end of the day on Tuesday, November 9th, and should include:

- Names of group members (1-2 people)
- A description of what you plan to do.
- Brief discussion of the main challenges you expect to face.
- Brief work plan, including dates for various intermediate tasks.

I will aim to get back to you about the project proposals on November 16th, so that you have a bit of time to work on the project. Note that I am not expecting the project to be a full research paper, but maybe a “toy” version of one with a few first-cut experimental or theoretical results that show a novel or interesting insight in the space of fairness or privacy.

Below are a few (non-exclusive) broad categories of project types. You are welcome to use these suggestions or design your own project related to the content of this course.

1 Implementation of an existing algorithm

Your project could produce a clean implementation of an existing DP or fairness algorithm, and test it out on (simulated or real) data.

For implementation projects, your description should include the specification of the input/output of your algorithm, the language you plan to use, the data you plan to test it on, what will be involved in preparing the data, what the testing will consist of, and some thoughts on what aspects of the algorithm might benefit from (heuristic) optimization.

2 Theoretical research project

Your project could also be a theoretical research project that solves (or begins to solve) open problems in the privacy and fairness literature. Your final write-up should include a summary of existing work, your new theoretical results with proofs, along with a discussion of challenges faced, other approaches attempted, and suggestions for future work.

3 Experimental research project

Your project could also be an experimental research project, in which you aim to use the techniques learned in class to solve a research question of practical interest, and to use simulations and experiments to validate the approaches you are using.

4 Where can I find ideas on what to work on?

A few possibilities below:

- Start with your research interests. Is there a natural question in your research area where privacy and fairness may be important considerations?
- Look at recent research. Some workshop and conferences that specialize in responsible data analysis, and have a lot of relevant papers, are the following:
 - EAAMO: <https://eaamo.org/>
 - FAccT: <https://facctconference.org/2021/index.html>
 - FORC: <https://responsiblecomputing.org/forc-2021/>
 - TPDP: <https://tpdp.journalprivacyconfidentiality.org/2021/>
- Take some inspiration from the real-life issues we have seen in class, and see if you can formulate a problem to solve them, or find real data to work with and apply existing techniques on.
- Follow-up on some of the reading you have been doing for class presentations.