IASD Anonymization, Privacy

Homework and Project

Debabrota Basu Pierre Senellart

January 22, 2021

1 Homework

For your homework, you have to read in detail one of the articles in the provided list or one you select yourself (which should be relevant for the class, and needs to be approved by the teachers) and write a critical report on it: summary of results (reformulated in your own words), discussion on the depth of the work, its impact, interestingness, limitations. This should essentially be what a reviewer of this article would have to do upon submission to a journal or a conference.

In addition, try to also discuss the context of the article: who were the authors and what is their expertise? where and when was it published? has it led to much further work? what other work does it rely on?

The entire report should be 2 to 5 pages long, depending on what there is to say about this particular article. Homework is individual, you cannot work in a group.

The article needs to be chosen on Moodle, by following the link "Selection of the homework article". Each student has to select a different article, the choice being made on a first-come, first-serve basis. Work on homework is individual.

The report on the article needs to be submitted by March 11 (23:59) on the Moodle Web site ("Homework").

2 Final project

The final project should consist on a practical or theoretical development based on the work of one or several articles. This can be the same article chosen for the homework, or a different one. Project can be worked on by individual students or by groups of two. The evaluation will expect more of a project with two students than with a single student.

The project will be evaluated on the basis of a project defense on March 25; along the defense, all outcomes of the project (code, datasets, experimental results, write-ups, etc.) need to be provided as well by March 25 as a ZIP file on Moodle ("Project").

3 List of articles

3.1 Anonymization

- t-closeness - Privacy Beyond k-Anonymity and l-Diversity
- Robust De-anonymization of Large Datasets
- From t-Closeness to Differential Privacy and Vice Versa in Data Anonymization

3.2 Differential privacy

- Broadening the Scope of Differential Privacy Using Metrics
- Differential Privacy for Functions and Functional Data
- Privacy: Theory meets Practice on the Map
- A Simple and Practical Algorithm for Differentially Private Data Release
- CoinPress: Practical Private Mean and Covariance Estimation
- Smooth Sensitivity and Sampling in Private Data Analysis

3.3 Privacy in ML

- Functional Mechanism: Regression Analysis under Differential Privacy
- Accuracy First: Selecting a Differential Privacy Level for Accuracy-Constrained ERM
- Differentially Private Distributed Convex Optimization via Functional Perturbation
- Differentially Private Coordinate Descent for Composite Empirical Risk Minimization.
- Tight Differential Privacy for Discrete-Valued Mechanisms and for the Subsampled Gaussian Mechanism Using FFT

3.4 Attacks on ML

- Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting
- Membership inference attacks against machine learning models
- Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures
- Hacking Smart Machines with Smarter Ones: How to Extract Meaningful Data from Machine Learning Classifiers