

IASD Anonymization, Privacy

Homework Instructions

Ashish Dandekar

Pierre Senellart

April 22, 2026

1 Description of the homework

For your homework, you have to read in detail one of the articles in the provided list 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.

2 Practical information

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.

The report on the article needs to be submitted by February 24 (23:59) on the Moodle Web site.

3 Link with final project

Reports submitted by all students will be made public and can serve as a basis for the project which needs to be completed by March 16, by groups of one or two students. The project should consist on a practical or theoretical development based on the work of a given research article, which can be one of the articles read for the homework, or another one. The project will be evaluated on the basis of a project defense on March 16; along the defense, all outcomes of the project (code, datasets, experimental results, write-ups, etc.) need to be provided as well by March 16.

4 List of articles

4.1 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

4.2 Differential privacy

- Smooth Sensitivity and Sampling in Private Data Analysis
- 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

4.3 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

4.4 Attacks on ML

- 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