Privacy for Machine Learning

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

- Motivations, traditional approaches, randomized response
- Definition and properties of differential privacy
- Mechanisms for discrete/categorical data
- Mechanisms for continuous data
- Alternative notions of differential privacy
- Differential privacy for statistical learning
- Attacks and connections with robustness
- Local differential privacy and federated learning

Compétence à acquérir :

This course covers the basics of Differential Privacy (DP), a framework that has become, in the last ten years, a de facto standard for enforcing user privacy in data processing pipelines. DP methods seek to reach a proper trade-off between protecting the characteristics of individuals and guaranteeing that the outcomes of the data analysis stays meaningful.

The first part of the course is devoted the basic notion of epsilon-DP and understanding the trade-off between privacy and accuracy, both from the empirical and statistical points of view. The second half of the course will cover more advanced aspects, including the different variants of DP and the their use to allow for privacy-preserving training of large and/or distributed machine learning models.

Mode de contrôle des connaissances :

- Individual homework (Python)
- Group project on a research paper (with report and defense)

Bibliographie, lectures recommandées :

- The Algorithmic Foundations of Differential Privacy, C. Dwork & A. Roth, Foundations and Trends in Theoretical Computer Science (2014)
- Programming Differential Privacy, J. P. Near & C. Abuah, online book (2021)