Local Methods for Privacy Protection and Impact on Fairness

Catuscia Palamidessi

Local Methods for Privacy Protection and Impact on Fairness

Catuscia Palamidessi
Inria Saclay & Institut Polytechnique de Paris
Palaiseau, France
catuscia@lix.polytechnique.fr

ABSTRACT

The increasingly pervasive use of big data and machine learning is raising various ethical issues, in particular privacy and fairness. In this talk, I will discuss some frameworks to understand and mitigate the issues, focusing on iterative methods coming from information theory and statistics. In the area of privacy protection, *differential privacy* (DP) and its variants are the most successful approaches to date. One of the fundamental issues of DP is how to reconcile the loss of information that it implies with the need to preserve the utility of the data. In this regard, a useful tool to recover utility is the *iterative Bayesian update* (IBU), an instance of the *expectation-maximization* method from statistics. I will show that the IBU, combined with a version of DP called *d-privacy* (also known as *metric differential privacy*), outperforms the state-of-the-art, which is based on algebraic methods combined with the *randomized response* mechanism, widely adopted by the Big Tech industry (Google, Apple, Amazon, ...). Then, I will discuss the issue of biased predictions in machine learning, and how DP can affect the level of *fairness* and accuracy of the trained model. Finally, I will show that the IBU can be applied also in this domain to ensure fairer treatment of disadvantaged groups and reconcile fairness and accuracy.

CCS CONCEPTS

• **Security and privacy → Privacy-preserving protocols**; • **Mathematics of computing → Expectation maximization**; **Information theory**.

KEYWORDS

Differential privacy, metric differential privacy, iterative Bayesian update, fairness.

ACM Reference Format:
Catuscia Palamidessi. 2023. Local Methods for Privacy Protection and Impact on Fairness . In *Proceedings of the Thirteenth ACM Conference on Data and Application Security and Privacy (CODASPY '23), April 24–26, 2023, Charlotte, NC, USA.* ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3577923.3587263

CODASPY '23, April 24–26, 2023, Charlotte, NC, USA

ACM ISBN 979-8-4007-0067-5/23/04.
https://doi.org/10.1145/3577923.3587263

BIOGRAPHY

Catuscia Palamidessi is a director of research at Inria Saclay (since 2002), where she leads the team COMETE. She has been a professor at the University of Genova, Italy (1994-1997) and Penn State University, USA (1998-2002). Palamidessi's research interests include Privacy, Machine Learning, Fairness, Secure Information Flow, Formal Methods, and Concurrency. In 2019 she has obtained an ERC advanced grant to conduct research on Privacy and Machine Learning. In 2022 she has been awarded the Gran Prix of the French Academy of Science. She has been PC chair of various conferences including LICS and ICALP, and PC member of more than 120 international conferences. She is on the Editorial board of several journals, including the IEEE Transactions in Dependable and Secure Computing, the ACM Transactions on Privacy and Security, Mathematical Structures in Computer Science, Theoretics, the Journal of Logical and Algebraic Methods in Programming, and Acta Informatica. She is serving on the Executive Committee of ACM SIGLOG, CONCUR, and CSL.

ACKNOWLEDGEMENT

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme. Grant agreement № 835294.

REFERENCES

[1] Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: differential privacy for location-based systems. In *Proceedings of the 20th ACM Conference on Computer and Communications Security (CCS 2013).* ACM, 901–914. https://doi.org/10.1145/2508859.2516735

[2] Cynthia Dwork, Frank Mcsherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *In Proceedings of the Third Theory of Cryptography Conference (TCC) (Lecture Notes in Computer Science, Vol. 3876)*, Shai Halevi and Tal Rabin (Eds.). Springer, 265–284.

[3] Ehab ElSalamouny and Catuscia Palamidessi. 2020. Full Convergence of the Iterative Bayesian Update and Applications to Mechanisms for Privacy Protection. arXiv:1909.02961 [cs.CR] To appear in the proceedings of EuroS&P.

[4] Karima Makhlouf, Sami Zhioua, and Catuscia Palamidessi. 2021. On the Applicability of Machine Learning Fairness Notions, In Proceedings of BIAS 2020. *SIGKDD Explor* 23, 1, 14–23.

[5] Simon Oya, Carmela Troncoso, and Fernando Pérez-González. 2017. Back to the Drawing Board: Revisiting the Design of Optimal Location Privacy-preserving Mechanisms. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas, Texas, USA). ACM, 1959–1972. https://doi.org/10.1145/3133956.3134004