
EDITORIAL FOR SPECIAL ISSUE ON THE THEORY AND PRACTICE OF DIFFERENTIAL PRIVACY 2018

ALEKSANDAR NIKOLOV

University of Toronto
e-mail address: anikolov@cs.toronto.edu

ABSTRACT. This special issue includes selected contributions from the 4th Workshop on Theory and Practice of Differential Privacy, which was held in Toronto, Canada on 15 October 2018 as part of the ACM Conference on Computer Security (CCS).

INTRODUCTION

The *4th Workshop on Theory and Practice of Differential Privacy* was held in Toronto, Canada on 15 October 2018 as part of the *ACM Conference on Computer Security (CCS)*. The workshop brought together researchers from computer science and statistics to discuss recent developments in both the theory and practice of differential privacy. This issue presents the following five contributions from that workshop:

- Aloni Cohen and Kobbi Nissim (2020) give an implementation of the reconstruction attack of Dwork, McSherry and Talwar (a variant of the Dinur and Nissim reconstruction attack) against the production system Diffix, run on real data. Diffix allows an unbounded number of queries to the data with a fixed amount of noise. The authors show that this allows exactly reconstructing the raw data, by implementing the Dwork, McSherry, and Talwar attack using pseudorandom queries.
- Naoise Holohan, Spiros Antonatos, Stefano Braghin and Pól Mac Aonghusa (2020) study a version of the Laplace noise mechanism in which the noise distribution is truncated to lie in a bounded range. This variant is convenient in implementations, and satisfies a slightly weaker (ϵ, δ) -differential privacy guarantee than the standard Laplace noise mechanism. The authors give tight bounds on the value of δ as a function of the range of truncation.
- Matthew Joseph, Aaron Roth, Jonathan Ullman and Bo Waggoner (2020) consider the problem of tracking population statistics in the local model of differential privacy when the data distribution may change over time. In the central curator model of differential privacy, the sparse vector technique allows for achieving error guarantees that only scale with the number of times the data distribution changes significantly. The main contribution of this work is an algorithm with similar properties in the local model.

Key words and phrases: differential privacy.

- Jordan Alexander Awan and Aleksandra Slavković (2020) offer a comprehensive set of tools for differentially private hypothesis testing with binary data. The authors propose the Tulap noise distribution as a tool to achieve (ϵ, δ) -differential privacy, and derive one- and two-sides hypothesis tests, and corresponding p -values, and confidence intervals, as post-processing of a test statistic perturbed with Tulap noise. The one-sides hypothesis test, in particular, is uniformly most powerful among all (ϵ, δ) -differentially private hypothesis tests at a given level.
- Finally, Borja Balle, Gilles Barthe and Marco Gaboardi (2020) take a functional view of (ϵ, δ) -differential privacy, and introduce the notion of a privacy profile of a mechanism, which describes, for any ϵ , the minimum δ for which the mechanism is (ϵ, δ) -differentially private. They study privacy profiles through the lens of divergences, and, as an application, give tighter analyses of privacy amplification under different ways to subsample the data.

Additional contributions will appear in the next issue of the Journal. The full program of the Workshop is available as part of this issue (Nikolov, 2020) and online at <https://tpdp.cse.buffalo.edu/2018/>.

REFERENCES

- Awan, Jordan Alexander, and Aleksandra Slavković.** 2020. “Differentially Private Inference for Binomial Data.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.725>.
- Balle, Borja, Gilles Barthe, and Marco Gaboardi.** 2020. “Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.726>.
- Cohen, Aloni, and Kobbi Nissim.** 2020. “Linear Program Reconstruction in Practice.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.711>.
- Dinur, Irit, and Kobbi Nissim.** 2003. “Revealing Information While Preserving Privacy.” *PODS '03*, 202–210. <https://doi.org/10.1145/773153.773173>.
- Dwork, Cynthia, Frank McSherry, and Kunal Talwar.** 2007. “The Price of Privacy and the Limits of LP Decoding.” *STOC '07*, 85–94. <https://doi.org/10.1145/1250790.1250804>.
- Holohan, Naoise, Spiros Antonatos, Stefano Braghin, and Pól Mac Aonghusa.** 2020. “The bounded Laplace mechanism in differential privacy.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.715>.
- Joseph, Matthew, Aaron Roth, Jonathan Ullman, and Bo Waggoner.** 2020. “Local Differential Privacy for Evolving Data.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.718>.
- Nikolov, Aleksandar.** 2020. “Program for TPDP 2018.” *Journal of Privacy and Confidentiality*, 10. <https://doi.org/10.29012/jpc.697>.