Journal of Privacy and Confidentiality
Vol. 9 (1) 2019 TPDP 2016
Submitted March 2019
Published March 2019

SPECIAL ISSUE ON THE THEORY AND PRACTICE OF DIFFERENTIAL PRIVACY 2016

MARCO GABOARDI

University at Buffalo, SUNY
e-mail address: gaboardi@buffalo.edu

This special issue presents papers based on contributions to the second international workshop on the "Theory and Practice of Differential Privacy" (TPDP) held in New York, NY, 23 June 2016, as part of the International Conference on Machine Learning (ICML).

Differential privacy is a mathematically rigorous definition describing the privacy protection provided by some data release mechanisms. It offers a strong guaranteed bound on what can be learned about a user as a result of participating in a differentially private data analysis. Researchers in differential privacy come from several areas of computer science, statistics and data analysis. The workshop is intended to be an occasion for researchers from these different research areas to discuss the recent developments in the theory and practice of differential privacy.

The program of the workshop [1] included 6 contributed talks, 4 invited speakers and 19 poster presentations. Invited speakers and contributing speakers at the workshop were invited to submit papers to this special issue. Five papers were accepted, most of which directly reflect talks presented at the workshop.

In this special issue

Mark Bun, Thomas Steinke and Jonathan Ullman [2] study three different models for answering multiple queries in a differentially private way: offline, online, and adaptive. The main results of the paper show separations between these three models.

Uri Stemmer and Kobbi Nissim [3] present results about the generalization properties of differential privacy. They show that differential privacy can be used to prove concentration bounds for functions that are not low-sensitive.

Yue Wang, Daniel Kifer, and Jaewoo Lee [4] study algorithms for computing confidence intervals for the parameters of differentially private machine learning models based on objective perturbation and output perturbation techniques.

Or Sheffet [5] provides an analysis of statistical inference for the Ordinary Least Squares technique using differentially private statistical estimators. Under mild assumptions this analysis give t-values and confidence intervals.

Yu-Xiang Wang [6] studies a refinement of differential privacy based on the idea of protecting the privacy of a specific individual with respect to a fixed data set. The contribution shows mechanisms that provide this privacy guarantee and their use.

JPrivacy Confidentiality
www.journalprivacyconfidentiality.org
DOI:10.29012/jpc.717

References

[1] Gilles Barthe, Christos Dimitrakakis, Marco Gaboardi, Andreas Haeberlen, Aaron Roth, and Aleksandra B Slavković. "Program for TPDP 2016". In: *Journal of Privacy and Confidentiality* 9 (1 2019). DOI: 10.29012/jpc.699.

[2] Mark Bun, Thomas Steinke, and Jonathan Ullman. "Make Up Your Mind: The Price of Online Queries in Differential Privacy". In: *Journal of Privacy and Confidentiality* 9 (1 2019). DOI: 10.29012/jpc.655.

[3] Uri Stemmer and Kobbi Nissim. "Concentration Bounds for High Sensitivity Functions Through Differential Privacy". In: *Journal of Privacy and Confidentiality* 9.1 (2019). DOI: 10.29012/jpc.658.

[4] Yue Wang, Daniel Kifer, and Jaewoo Lee. "Differentially Private Confidence Intervals for Empirical Risk Minimization". In: *Journal of Privacy and Confidentiality* 9 (1 2019). DOI: 10.29012/jpc.660.

[5] Or Sheffet. "Differentially Private Ordinary Least Squares". In: *Journal of Privacy and Confidentiality* 9 (1 2019). DOI: 10.29012/jpc.654.

[6] Yu-Xiang Wang. "Per-instance Differential Privacy". In: *Journal of Privacy and Confidentiality* 9 (1 2019). DOI: 10.29012/jpc.662.