TPDP 2018 - Theory and Practice of Differential Privacy

Toronto, Canada - 15 October 2018 - part of CCS 2018

Location: Beanfield Centre, 203AB

Program

9:00 - 9:10	Welcome
9:10 - 10:00	Composition, Verification, and Differential Privacy Justin Hsu (Invited Speaker)
10:00 - 10:45	Coffee Break
10:45 - 11:35	Deploying Differential Privacy for Learning on Sensitive Data Úlfar Erlingsson (Invited Speaker)
11:35 - 11:55	Local Differential Privacy for Evolving Data Matthew Joseph, Aaron Roth, Jonathan Ullman and Bo Waggoner
11:55 - 2:00	Lunch Break
	Private PAC learning implies finite Littlestone dimension Noga Alon, Roi Livni, Maryanthe Malliaris and Shay Moran
2:00 - 3:00	Linear Program Reconstruction in Practice Aloni Cohen and Kobbi Nissim
	Optimizing error of high dimensional statistical queries under differential privacy Ryan McKenna, Gerome Miklau, Michael Hay and Ashwin Machanavajjhala
3:00 - 3:45	Coffee Break
3:45 - 4:35	PSI: A Private Data-Sharing Interface Salil Vadhan (Invited Speaker)
4:35 - 4:55	Towards Modeling Singling Out Aloni Cohen and Kobbi Nissim
4:55-6:00	Poster Session

Posters

- Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences. Borja Balle, Gilles Barthe and Marco Gaboardi
- Subsampled Renyi Differential Privacy and Analytical Moments Accountant. YuXiang Wang, Borja Balle and Shiva Kasiviswanathan
- Privately Learning HighDimensional Distributions. Gautam Kamath, Jerry Li, Vikrant Singhal and Jonathan Ullman
- Differentially Private ChangePoint Detection. Rachel Cummings, Sara Krehbiel, Yajun Mei, Rui Tuo and Wanrong Zhang
- Locally Private Mean Estimation: z-Tests and Confidence Intervals. Marco Gaboardi, Ryan Rogers and Or Sheffet
- Differentially Private Uniformly Most Powerful Tests for Binomial Data. Jordan Awan and Aleksandra Slavkovic
- Reasoning about Differential Privacy and its Relaxations. Tetsuya Sato, Gilles Barthe, Marco Gaboardi, Justin Hsu and ShinYa Katsumata
- Decentralized Differential Privacy via Mixnets. Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber and Maxim Zhilyaev
- Private Hypothesis Testing Using Rank-Based Methods. Andrew Bray, Simon Couch, Adam Groce, Zeki Kazan and Kaiyan Shi
- Test without Trust: Optimal Locally Private Distribution Testing. Jayadev Acharya, Clement Canonne, Cody Freitag and Himanshu Tyagi
- Efficient Distributed Differential Privacy for Noisy Sums. Mark Bun and Noah Stephens-Davidowitz
- INSPECTRE: Privately Estimating the Unseen. Jayadev Acharya, Gautam Kamath, Ziteng Sun and Huanyu Zhang
- Towards a probability measure approach for Differential Privacy in Bayesian inference. Gian Pietro Farina, Jiawen Liu, Marco Gaboardi and Mark Bun
- A Software Framework for Testing and Evaluation of Differentially Private Active Learning Schemes. Daniel Bittner, Shantanu Rane, Alejandro Brito and Rebecca Wright
- Shrinkwrap: Efficient SQL Query Processing in Differentially Private Data Federations. Johes Bater, Xi He, William Ehrich, Ashwin Machanavajjhala and Jennie Rogers
- The bounded Laplace mechanism in differential privacy. Naoise Holohan, Spiros Antonatos, Stefano Braghin and Pol Mac Aonghusa
- Differentially-Private Draw and Discard Machine Learning. Vasyl Pihur, Aleksandra Korolova, Frederick Liu, Subhash Sankuratripati, Moti Yung, Dachuan Huang and Ruogu Zeng
- Individual Sensitivity Preprocessing for Data Privacy. Rachel Cummings and David Durfee
- A Hybrid Trust Model for Distributed Differential Privacy. Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig and Rui Zhang
- Differentially Private Synthetic Data Generation via GANs. Digvijay Boob, Rachel Cummings, Dhamma Kimpara, Uthaipon Tantipongpipat, Chris Waites and Kyle Zimmerman
- Boosting Model Performance through Differentially Private Model Aggregation. Sophia Collet, Robert Dadashi, Zahi N. Karam, Chang Liu, Parinaz Sobhani, Yevgeniy Vahlis and Ji Chao Zhang

Context

Differential privacy is a promising approach to privacy-preserving data analysis. Differential privacy provides strong worst-case guarantees about the harm that a user could suffer from participating in a differentially private data analysis, but is also flexible enough to allow for a wide variety of data analyses to be performed with a high degree of utility. Having already been the subject of a decade of intense scientific study, it has also now been deployed in products at government agencies such as the U.S. Census Bureau and companies like Apple and Google.

Researchers in differential privacy span many distinct research communities,

Invited Speakers

Úlfar Erlingsson
Google
Justin Hsu
University of Wisconsin, Madison
Salil Vadhan
Harvard University

Important Dates

Abstract Submission

including algorithms, computer security, cryptography, databases, data mining, machine learning, statistics, programming languages, social sciences, and law. This workshop will bring researchers from these communities together to discuss recent developments in both the theory and practice of differential privacy.

Submission

The overall goal of TPDP is to stimulate the discussion on the relevance of differentially private data analyses in practice. For this reason, we seek contributions from different research areas of computer science and statistics.

Authors are invited to submit a short abstract (2-4 pages maximum) of their work. Abstracts must be written in English. You can find the deadline and submission server link on the right.

Submissions will undergo a lightweight review process and will be judged on originality, relevance, interest and clarity. Submission should describe novel works or works that have already appeared elsewhere but that can stimulate the discussion between different communities at the workshop. Accepted abstracts will be presented at the workshop either in technical sessions or as posters.

The workshop will not have formal proceedings and is not intended to preclude later publication at another venue.

Specific topics of interest for the workshop include (but are not limited to):

- theory of differential privacy,
- privacy preserving machine learning,
- differential privacy and statistics,
- differential privacy and security,
- differential privacy and data analysis,
- trade-offs between privacy protection and analytic utility,
- differential privacy and surveys,
- programming languages for differential privacy,
- relaxations of the differential privacy definition,
- differential privacy vs other privacy notions and methods,
- experimental studies using differential privacy,
- differential privacy implementations,
- differential privacy and policy making,
- applications of differential privacy.

Call for Papers: pdf

~~July 20 (anywhere on earth)~~
July 27 (anywhere on earth)

Notification
~~August 13~~
August 18

Workshop
October 15

Submission website

Easychair TPDP 2018

Organizing and Program Committee

Aleksandar Nikolov
University of Toronto
Raef Bassily
Ohio State University
Mark Bun
Boston University
Michael Hay
Colgate University
Vishesh Karwa
Temple University
Katrina Ligett
Hebrew University
Anand Sarwate
Rutgers University
Thomas Steinke
IBM
Reza Shokri
National University of Singapore
Kunal Talwar
Google