

TPDP 2016 - Theory and Practice of Differential Privacy

New York, USA - 23 June 2016 - part of ICML 2016

Program - room O'Neil, Marriott Hotel

- 8:20 - 8:30 Welcome
- 8:30 - 9:10 [Machine Learning and Privacy: Friends or Foes?](#)
[Vitaly Shmatikov](#) - invited speaker.
- 9:10 - 10:00 [Make Up Your Mind: The Price of Online Queries in Differential Privacy](#),
Mark Bun, Thomas Steinke and Jonathan Ullman.
[Principled Evaluation of Differentially Private Algorithms](#),
Michael Hay, Ashwin Machanavajhala, Gerome Miklau, Yan Chen and Dan Zhang.
[The Cost of Provable Privacy: A Case Study on Linked Employer-Employee Data](#),
Samuel Haney, Ashwin Machanavajhala, John Abowd, Matthew Graham, Mark Kutzbach and Lars Vilhuber.
- 10:00 - 10:30 Coffee Break
- 10:30 - 11:10 [New Directions in Privacy-Preserving Data Analysis](#)
[Kamalika Chaudhuri](#) - invited speaker.
- 11:10 - 12:00 [Differentially Private Integer Partitions and their Applications](#),
Jeremiah Blocki.
[Proving Differential Privacy via Probabilistic Couplings](#),
Gilles Barthe, Marco Gaboardi, Benjamin Gregoire, Justin Hsu and Pierre-Yves Strub.
[Challenges of Visualizing Differentially Private Data](#),
Dan Zhang, Michael Hay, Gerome Miklau and Brendan O'Connor.
- 12:00 - 1:30 Lunch Break and Poster Setup
- 1:30 - 2:10 [Bridging the Gap between Computer Science and Legal Approaches to Privacy](#)
[Kobbi Nissim](#) - invited speaker.
- 2:10 - 3:00 [Generalization and Learnability under Differential Privacy and its Variants](#)
[Yu-Xiang Wang](#) - invited speaker.
- 3:00 - 3:20 Coffee Break
- 3:20 - 4:40 Poster Session

Posters

- Mark Bun, Thomas Steinke and Jonathan Ullman. Make Up Your Mind: The Price of Online Queries in Differential Privacy
- Jeremiah Blocki. Differentially Private Integer Partitions and their Applications
- Rachel Cummings, Jennifer Wortman Vaughan and David Pennock. The Possibilities and Limitations of Private Prediction Markets
- Borja Balle, Maziar Gomrokchi and Doina Precup. Differentially Private Policy Evaluation
- Mijung Park and Max Welling. Differentially Private Iteratively Reweighted Least Squares
- Xi He, Ashwin Machanavajjhala, Cheryl Flynn and Divesh Srivastava. Composing Differential Privacy and Secure Multiparty Computation for Efficient Private Record linkage
- Ryan Rogers, Aaron Roth, Adam Smith and Om Thakkar. Max-Information, Differential Privacy, and Post-Selection Hypothesis Testing
- Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen and Dan Zhang. Principled Evaluation of Differentially Private Algorithms
- Jaewoo Lee and Daniel Kifer. Postprocessing for Iterative Differentially Private Algorithms
- Or Sheffet. Differentially Private Ordinary Least Squares: t-Values, Confidence Intervals and Rejecting Null-Hypotheses
- Antti Honkela, Mrinal Das, Onur Dikmen and Samuel Kaski. Efficient differentially private regression
- Gilles Barthe, Gian Pietro Farina, Marco Gaboardi, Emilio Jesus Gallego Arias, Andy Gordon, Justin Hsu and Pierre-Yves Strub. PrivInfer: A Framework for Differentially Private Bayesian Machine Learning
- Samuel Haney, Ashwin Machanavajjhala, John Abowd, Matthew Graham, Mark Kutzbach and Lars Vilhuber. The Cost of Provable Privacy: A Case Study on Linked Employer-Employee Data
- Gilles Barthe, Marco Gaboardi, Benjamin Gregoire, Justin Hsu and Pierre-Yves Strub. Proving Differential Privacy via Probabilistic Couplings
- Daniel Winograd-Cort, Andreas Haeberlen, Benjamin Pierce and Aaron Roth. A Framework for Adaptive Differential Privacy
- Dan Zhang, Michael Hay, Gerome Miklau and Brendan O'Connor. Challenges of Visualizing Differentially Private Data
- The Privacytools Group at Harvard and James Honaker. PSI: A Private data Sharing Interface
- Aristide Charles Yedia Tossou and Christos Dimitrakakis. Differential Privacy and Multi-Armed Bandits
- Jalaj Upadhyay. Fast and Space-optimal Low-rank Factorization in the Streaming Model With Application in Differential Privacy

Context

Differential privacy is a promising approach to the privacy-preserving release of data: it offers a strong guaranteed bound on the increase in harm that a user incurs as a result of participating in a differentially private data analysis. Several mechanisms and software tools have been developed to ensure differential privacy for a wide range of data analysis tasks.

Researchers in differential privacy come from several area of computer science as algorithms, programming languages, security, databases, machine learning, as well as from several areas of 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.

Submission

Invited Speakers

[Kamalika Chaudhuri](#)

University of California, San Diego,

[Kobbi Nissim](#)

Ben-Gurion University & Harvard University,

[Vitaly Shmatikov](#)

Cornell Tech,

[Yu-Xiang Wang](#)

Carnegie Mellon University.

Important Dates

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 by May 1, 2016. Abstracts must be written in English and be submitted as a single PDF file at [EasyChair page for TPDP](#).

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: [txt](#)

Abstract Submission

May 1, 2016

Notification

May 10, 2016

Deadline registration ICML

May 15, 2015

Workshop

June 23, 2016

Submission website

[EasyChair page for TPDP](#)

Organizing and Program Committee

[Gilles Barthe](#)

IMDEA Software

[Christos Dimitrakakis](#)

University of Lille

[Marco Gaboardi](#)

University at Buffalo, SUNY

[Andreas Haeberlen](#)

University of Pennsylvania

[Aaron Roth](#)

University of Pennsylvania

[Aleksandra B. Slavkovic](#)

Penn State University