TPDP 2017 - Theory and Practice of Differential Privacy

Dallas, TX, USA - October 30, 2017 - part of CCS 2017

Program (Dallas Ballroom D3)

8:45 - 9:00	Welcome!
9:00 - 9:45	Challenges of Differential Privacy in Practice Dan Kifer - Pennsylvania State University (Invited Speaker)
9:45 - 10:00	Accuracy First: Selecting a Differential Privacy Level for Accuracy-Constrained ERM Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Zhiwei Steven Wu
10:00 - 10:45	Coffee Break
10:45 - 11:30	Privacy with Constraints: Challenges & Opportunities Xi He - Duke University (Invited Speaker)
11:30 - 12:00	Ektelo: A Framework for Defining Differentially-Private Computations Dan Zhang, Ryan McKenna, Ios Kotsogiannis, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala Finite Sample Differentially Private Confidence Intervals Vishesh Karwa and Salil Vadhan
12:00 - 2:00	Lunch Break
2:00 - 2:45	Modular Approach to Differential Privacy Ilya Mironov - Google (Invited Speaker)
2:45 - 3:00	Practical Locally Private Heavy Hitters Raef Bassily, Kobbi Nissim, Uri Stemmer and Abhradeep Thakurta
3:00 - 3:45	Coffee Break
3:45 - 4:30	Concentrating on the Foundations of Differential Privacy Thomas Steinke - IBM Research (Invited Speaker)
4:30 - 5:00	BLENDER: Enabling Local Search with a Hybrid Differential Privacy Model Brendan Avent, Aleksandra Korolova, David Zeber, Torgeir Hovden and Benjamin Livshits Differential Privacy on Finite Computers Victor Balcer and Salil Vadhan
5:00 - 6:00	Poster Session

Posters

- Accuracy First: Selecting a Differential Privacy Level for Accuracy-Constrained ERM. Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner and Zhiwei Steven Wu.

- Practical Locally Private Heavy Hitters. Raef Bassily, Kobbi Nissim, Uri Stemmer and Abhradeep Thakurta.
- BLENDER: Enabling Local Search with a Hybrid Differential Privacy Model. Brendan Avent, Aleksandra Korolova, David Zeber, Torgeir Hovden and Benjamin Livshits.
- Differential Privacy on Finite Computers. Victor Balcer and Salil Vadhan.
- Ektelo: A Framework for Defining Differentially-Private Computations. Dan Zhang, Ryan McKenna, Ios Kotsogiannis, Gerome Miklau, Michael Hay and Ashwin Machanavajjhala.
- Finite Sample Differentially Private Confidence Intervals. Vishesh Karwa and Salil Vadhan.
- Towards an Optimal Algorithm for Concentrated Differential Privacy. Jaroslaw Blasiok, Mark Bun, Aleksandar Nikolov and Thomas Steinke.
- One-sided Privacy Stylianos Doudalis, Ios Kotsogiannis, Ashwin Machanavajjhala and Sharad Mehrotra.
- Differential Privacy for Growing Databases. Rachel Cummings, Gi Heung Robin Kim, Sara Krehbiel and Uthaipon Tao Tantipongpipat.
- Composable and Versatile Privacy via Truncated CDP. Mark Bun, Cynthia Dwork, Guy Rothblum and Thomas Steinke.
- Privacy Loss in Apple's Implementation of Differential Privacy on MacOS 10.12. Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang and Xiaofeng Wang.
- Revisiting the Gaussian Mechanism. Borja Balle.
- Bootstrap Inference and Differential Privacy. Thomas Brawner and James Honaker.
- Hiding Data Deception Attacks in Differential-Privacy Noise. Jairo Giraldo, Alvaro Cardenas, Murat Kantarcioglu and Jonathan Katz.
- Vulnerability in Floating Point Implementation of Exponential Mechanism. Matthew Burke.
- On the Correct Use of the Gaussian Mechanism for Differential Privacy. Jun Zhao.
- Locally Differentially Private Heavy Hitter Identification. Tianhao Wang, Ninghui Li and Somesh Jha.
- Competitive Differentially Private Algorithms for Interactive Queries. Aleksandar Nikolov and Jonathan Ullman.

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, 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.

Invited Speakers

Dan Kifer
Pennsylvania State University
Ilya Mironov
Google
Thomas Steinke
IBM Research
Xi He
Duke University

Important Dates

Abstract Submission
~~August 4, 2017~~
August 11, 2017 **extended**
(Anywhere on Earth)

Notification
September 4, 2017

Workshop
October 30, 2017

Submission website

Abstracts must be written in English and do not need to be anonymous.

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.

Easychair TPDP 2017

Organizing and Program Committee

Rachel Cummings
Caltech → Georgia Tech
Marco Gaboardi
University at Buffalo, SUNY
Justin Hsu
University of Pennsylvania
Aleksandra Korolova
University of Southern California
Ashwin Machanavajjhala
Duke University
Gerome Miklau
UMass Amherst
Abhradeep Guha Thakurta
UC Santa Cruz
Jonathan Ullman (chair)
Northeastern University