Title: Privacy Regulation and Protection in Machine Learning

Summary

Recent advances in artificial intelligence greatly benefit from data-driven machine learning methods that train deep neural networks with large scale data. The usage of data should be responsible, transparent, and comply with privacy regulations. This workshop aims to bring together industry and academic researchers, privacy regulators and legal, policy people to have a conversation on privacy research. We hope to (re)visit major privacy considerations from both technical and nontechnical perspectives through discussions with interdisciplinary discussions.

Topics of interest include, but are not limited to

- Relationship of privacy regulation (such as GDPR, DMA) to machine learning
- Interpolation and explanation of data privacy
- Efficient methods for privacy preserving machine learning
- Federated learning for data minimization
- Differential privacy theory and practice
- Threat model and privacy attacks
- Encryption methods for machine learning
- Privacy in machine learning systems
- Privacy for large language models
- Relationship between privacy, transparency, auditability, verifiability
- Relationship between privacy, robustness, fairness etc.

Modality

The workshop prioritizes in-person discussions, with a hybrid component of streaming invited talks and panel discussions to virtual attendees in order to ensure accessibility to those who cannot attend in person. We will also make the recordings of the presentations available to the virtual attendees.

Tentative schedule

We have planned for invited talks, spotlight talks from contributed papers, poster sessions, and a panel discussion. We highlight the exchange of ideas through question and answer for talks, and active discussion in panel and poster sessions. We have listed a tentative schedule.

9-9:05 Introduction and opening remarks
9:05-9:40 Invited speaker 1
9:40-10:00 Spotlight paper talks
10-10:15 Break
10:15-10:50 Invited speaker 2
10:50-11:25 Invited speaker 3
11:25-12:30 Poster

12:30-1:30 Lunch break
1:30-2:25 Panel discussion
2:25-3:00 Invited speaker 4
3:00-3:15 Break
3:15-3:50 Invited speaker 5
3:50-4:20 Spotlight talks
4:20-4:35 Break
4:35-5:10 Invited speaker 6
5:10-5:45 Invited speaker 7
5:45-5:50 Concluding remarks

Invited speakers and panelists

We have invited the following speakers and panelists from diverse backgrounds to ensure success of the workshop's above mentioned goal, confluence of ideas from regulators, practitioners, and researchers, with six confirmed speakers and one speaker pending reply. Among the speakers, we have invited Massimo Attoresi who has rich experience as privacy regulator; Dan Kifer who is deeply engaged in the usage of privacy techniques in the US Census Bureau; Janel Thamkul who has rich experience for policy counseling at companies. In addition, We also invited (co-)inventors of important privacy techniques such as differential privacy and federated learning, and industry leaders who have experience of privacy in production.

1. Kobbi Nissim (confirmed): Professor at the Department of Computer Science, Georgetown University and an Affiliate Professor at Georgetown Law; one of the inventors of differential privacy
2. Dan Kifer (confirmed): Professor at the Department of Computer Science & Engineering, Penn State University; deeply involved in U.S. Census 2020
3. Massimo Attoresi (pending reply): Acting Head of the Technology & Privacy unit of the EDPS
4. Janel Thamkul (confirmed): Deputy General Counsel at Anthropic; Anthropic serves one of the best generative language models.
5. Daniel Ramage (confirmed) Director at Google leading a team that invents federated learning, and works on privacy technologies that have been successfully shipped to production
6. Rachel Cummings (confirmed): Associate Professor of Industrial Engineering and Operations Research at Columbia University; Affiliate in the Department of Computer Science (by courtesy) and a Co-chair of the Cybersecurity Research Center at the Data Science Institute.
7. Will Bullock (confirmed): Director at Meta leading a team working on privacy for ads and commerce business.

We may additionally invite experts in the field to serve as panelists, and may reach out to the following candidates: Andrea Jelinek (Austria DPA), Gauri Joshi (Associate Professor at CMU), Golnoosh Farnadi (Assistant Professor at McGill University & Mila), Aureline Bellet (tenured researcher at Inria), Nic Lane (University of Cambridge & Flower Labs), Sergei Vassilvitskii (Director at Google).

Organizers and biographies

Contact: Zheng Xu (xuzheng@google.com) and Sewoong Oh (sewoongo@google.com)

Salman Avestimehr

- avestime@usc.edu
- https://www.avestimehr.com/
- https://scholar.google.com/citations?user=Qhe5ua0AAAAJ&hl=en
- Bio: Dr. Avestimehr is a Dean's Professor, the inaugural director of the USC-Amazon Center on Trustworthy AI, and the director of the Information Theory and Machine Learning (vITAL) research lab at the Electrical and Computer Engineering Department of University of Southern California. He is also the co-founder and CEO of FEDML. He received his Ph.D. in 2008 and M.S. degree in 2005 in Electrical Engineering and Computer Science, both from the University of California, Berkeley. Prior to that, he obtained his B.S. in Electrical Engineering from Sharif University of Technology in 2003. His research interests include information theory, decentralized and federated machine learning, secure and privacy-preserving machine learning and distributed computing. Dr. Avestimehr has received a number of awards for his research, including the James L. Massey Research & Teaching Award from IEEE Information Theory Society, an Information Theory Society and Communication Society Joint Paper Award, a Presidential Early Career Award for Scientists and Engineers (PECASE) from the White House (President Obama), a Young Investigator Program (YIP) award from the U. S. Air Force Office of Scientific Research, a National Science Foundation CAREER award, the David J. Sakrison Memorial Prize, and several Best Paper Awards at Conferences and Workshops. He has been an Associate Editor for IEEE Transactions on Information Theory, a general Co-Chair of the 2020 International Symposium on Information Theory (ISIT), and co-organizer of many workshops. He was an Amazon Scholar in Alexa AI in 2021. He is a fellow of IEEE.

Tian Li

- litian@uchicago.edu
- https://www.cs.cmu.edu/~litian/
- https://scholar.google.com/citations?hl=en&user=8JWoJrAAAAAJ
- **Bio:** Tian Li is currently a postdoctoral researcher at FAIR Labs at Meta, and will join the University of Chicago as an Assistant Professor in 2024. Her research interests are in distributed optimization, federated learning, and trustworthy ML, specifically on principled and scalable approaches to enabling learning across diverse data sources. She earned her Ph.D. in Computer Science at Carnegie Mellon University in August 2023. Before CMU, she received her undergraduate degrees in Computer Science and Economics from Peking University. She received the Best Paper Award at the ICLR Workshop on

Security and Safety in Machine Learning Systems, was invited to participate in the EECS Rising Stars Workshop, won the first place in UK-US Privacy Enhancing Technologies Challenge, and was recognized as a Rising Star in Machine Learning/Data Science by multiple institutions.

Niloofar (Fatemeh) Mireshghallah

- niloofar@cs.washington.edu
- https://cseweb.ucsd.edu/~fmireshg/
- https://scholar.google.com/citations?user=WUCu45YAAAAJ&hl=en&authuser=2
- **Bio**: Niloofar Mireshghallah is a post-doctoral scholar at the Paul G. Allen Center for Computer Science & Engineering at University of Washington. She received her Ph.D. from the CSE department of UC San Diego in 2023. Her research interests are Trustworthy Machine Learning and Natural Language Processing. She is a recipient of the National Center for Women & IT (NCWIT) Collegiate award in 2020 for her work on privacy-preserving inference, a finalist of the Qualcomm Innovation Fellowship in 2021 and a recipient of the 2022 Rising star in Adversarial ML award.

Sewoong Oh

- sewoongo@google.com
- https://homes.cs.washington.edu/~sewoong/
- https://scholar.google.com/citations?user=N6AWeX0AAAAJ
- **Bio:** Sewoong Oh is a Professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington and a Staff Research Scientist at Google's Federated Learning team. Previous to joining University of Washington in 2019, he was an Assistant Professor in the department of Industrial and Enterprise Systems Engineering at University of Illinois at Urbana-Champaign since 2012. Sewoong's research focuses on foundations of private machine learning in topics including differential privacy, secure and robust machine learning, and federated learning. He was co-awarded the ACM SIGMETRICS best paper award in 2015, NSF CAREER award in 2016, ACM SIGMETRICS rising star award in 2017, and GOOGLE Faculty Research Awards in 2017 and 2020.

Florian Tramer

- florian.tramer@inf.ethz.ch
- www.floriantramer.com
- https://scholar.google.com/citations?user=ijH0-a8AAAAJ
- **Bio:** Florian Tramèr is an assistant professor of computer science at ETH Zurich. Before that he was a PhD student at Stanford University, and a visiting researcher at Google Brain. His research interests lie in Computer Security, Cryptography and Machine Learning security. In his current work, he studies the worst-case behavior of Deep Learning systems from an adversarial perspective, to understand and mitigate long-term threats to the safety and privacy of users. He is a recipient of the 2021 Rising star in Adversarial ML award. Together with collaborators, his research was selected as

runner-up for the 2022 Caspar Bowden Award, and awarded a USENIX Security 2023 distinguished paper award.

Zheng Xu

- xuzheng@google.com
- https://research.google/people/106689/
- https://scholar.google.com/citations?user=TfWIMTYAAAAJ
- **Bio**: Zheng Xu is a research scientist working on federated learning and privacy at Google. He earned his Ph.D. in optimization and machine learning from University of Maryland, College Park, in 2019. Before that, he got his master's and bachelor's degree from the University of Science and Technology of China. He has published 30+ papers at top research conferences and journals with 9000+ citations, and received two best student paper awards. He is a co-author of Advances and Open Problems in Federated Learning, and a lead author of A Field Guide to Federated Optimization, both of which resulted from 20+ collaborators in workshop discussions. He is a co-organizer of the Google Federated Learning and Analytics Workshop 2020, TTIC workshop on New Frontiers in Federated Learning 2023, and the lead organizer of Federated Learning and Analytics in Practice Workshop at ICML 2023.

Anticipated audience size

We expect more than 100 participants. The estimation is based on FL and GenLaw workshops at ICML’23.

Plan to get an audience

We will take the following actions:

- create a workshop webpage with all the information;
- advocate call for papers via sharing it with relevant mailing lists in both academia and industry;
- send email announcements to organizer institutions and the broader privacy research community;
- reach out to traditionally underrepresented institutions;
- and share information on social media such as twitter and linkedin, and advertisement on relevant slack workspaces.

Diversity commitment

We encourage diversity in both organizers and speakers, which considers gender, affiliation, location, career stages, knowledge and cultural background. Specifically, the current tentative schedule considers the balance of gender, industry and academia, background of policy, legal and techniques, and the location. We are particularly excited about the possibility to discuss European leading privacy regulations such as GDPR and DMA at Vienna Austria.

The organizers are at different career stages ranging from early career postdoc to full professor. We intentionally invite relatively senior speakers, and will balance it with highlighting the novel contributed research from students. We will seek sponsorship from Google, Meta, FedML, Flower and other companies to potentially provide financial support for students, and consider the diversity when providing such support.

Access

All information will be updated on a webpage on a regular basis. Though the workshop is non-archival, we will use OpenReview for double-blind review process, and host the accepted papers. If we can secure funding, we can provide more virtual access, and provide financial support for attending the workshop.
Already published work at main machine learning venues (ICLR/ICML/NeurIPS) including papers accepted to the ICLR main conference will be explicitly discouraged in call for papers.

Previous related workshops

This inaugural workshop is built upon the success of closely related privacy and/or federated learning workshops at ICLR/ICML/NeurIPS in recent years (2019-2022), in which some subset of the organizers have been involved in varying capacities. For example, GenLaw-ICML’23, FL-ICML’23, FL-NeurIPS'22, TPDP-ICML’22, PriML-NeurIPS’21, TPDP-ICML’21, DPML-ICLR’21. Notably, Zheng Xu and Tian Li are (co-)organizers of FL-ICML’23, Niloofar (Fatemeh) Mireshghallah is the (co-)organizer of GenLaw-ICML’23 and DPML-ICLR’21.

Privacy is a multifaceted, technological and social issue of today’s information age. Foundational advances will inevitably involve a combination of legislation, economics view points, and technical solutions. This workshop emphasizes bringing together regulators, industry, and academics to initiate discussions and synergistic cross pollination of ideas. Additionally, this workshop will discuss privacy research in major machine conferences, for the first time in recent years that will be located in Europe, to highlight the interdisciplinary nature of privacy .

Reviewing

In addition to the double-blind reviewing as part of our diversity commitment, we will attract a diverse program committee. Each submission will be reviewed by 3 reviewers. Decisions will be made in a transparent way by the organizers. We will encourage presentation of work with novel perspectives and ambitious goals. We will avoid conflict of interests by explicitly asking the reviewers to indicate any. The final list of accepted papers will be published on the workshop website.