

Areas

Machine Learning, Causal Inference, Privacy, Security, Software Systems.

Education

- 2013–2019 **PhD in Computer Science**, *Columbia University*, New York, *GPA 4.0*.
◦ Thesis: *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*.
◦ Advisors: Roxana Geambasu, Augustin Chaintreau, and Daniel Hsu.
- 2011–2012 **Master of Science in Computer Science**, *Columbia University*, New York, *GPA: 3.90*.
- 2008–2011 **Ingénieur Diplômé**, *Ecole Polytechnique*, Paris, France, *GPA: 3.90*.
Major in Computer Science, multidisciplinary training in Physics, Economics, and Biology.
- 2006–2008 **Classes Préparatoires in Science**, *Lycée Louis-le-Grand*, Paris, France, *GPA: 4.0*.
Intensive preparation in Mathematics and Physics for competitive admission into top French universities.

Work Experience

- 2021–Current **Assistant Professor**, *The University of British Columbia*, Vancouver.
- 2019–2021 **Post-Doctoral Researcher**, *Microsoft Research*, New York.
◦ Developed a new counterfactual estimators for system policies [9].
◦ Developed a new methodology to explain ML predictions based on training data [7].
- 2013–2019 **Research Assistant**, *Columbia University*, New York.
- Summer 2017 **Research Intern**, *Microsoft Research*, New York.
◦ Applied reinforcement learning to systems, with a focus on off-policy evaluation [14].
- 2012–2014 **Teaching Assistant**, *Columbia University*, New York.
◦ Distributed Systems. Instructor: Roxana Geambasu.
◦ Data Journalism. Instructor: Mark Hansen.
◦ Computer Networking. Instructor: Augustin Chaintreau.
- Summer 2014 **Freelance**, *Floatingapps*, New York.
◦ Developed a Rails application for rental data visualization, using Cassandra, Elasticsearch, d3, and React.
◦ Built an iOS application for email visualization as timelines in Objective-C.
- Summer 2013 **Freelance**, *Milky*, Paris, France.
◦ Built an extended iOS application for *Le Grand Journal*, a show of major French TV Channel Canal+.
- 2009–2012 **Cofounder**, *Pionid*, Paris, France.
Developed the main product, an iOS group messaging application.

Research

I am broadly interested in machine learning systems, with a specific focus on applications that provide rigorous guarantees of robustness, privacy, and security. In my research, I leverage, adapt, and improve theoretical tools (differential privacy, causal inference, reinforcement learning) to enable specific applications (ML attacks/defenses, privacy preserving data management, system decisions optimization) with sound guarantees. I currently focus on two broad directions:

ML Security, Privacy & Explainability While data-driven systems can yield social and economic benefits, they also open new security and privacy threats, and their opacity can undermine users' trust. I develop ML models with guarantees of security, privacy, and explainability, and build the infrastructure necessary to soundly deploy and maintain them.

Certified Defenses Against Adversarial Examples.

I developed PixelDP (now called randomized smoothing), the first certified defense that both offers a guaranteed level of robustness against these attacks and scales to large models and datasets, such as Google's Inception on the ImageNet dataset [11].

Infrastructure Support for Differential Privacy in Machine Learning Workloads.

I am developing infrastructure support for differential privacy at the level of entire ML and DB query workloads. This includes designing the DP theory and systems abstractions required to keep track of privacy loss on continuously growing data streams [12, 8], adding privacy as a resource in cluster management systems, and designing scheduling algorithms for the privacy resource [10].

Increasing Explainability in Machine Learning.

I leverage causal inference theory to design black-box explainability techniques, to understand data usage in web services [20, 18], or how training data can influence a model's predictions [7]. I also work on explainable models [6].

Counterfactual Evaluation & Optimization for Systems Machine learning and statistics present new opportunities to solve traditional systems challenges in improved ways. I leverage causal inference and reinforcement learning to evaluate and optimize systems policies with sound statistical approaches [14, 9].

Papers

- [1] Amir Sabzi, Rut Vora, Swati Goswami, Margo Seltzer, Mathias Lécuyer, and Aastha Mehta. Netshaper: A differentially private network side-channel mitigation system. In *Tackling Climate Change with Machine Learning: workshop at NeurIPS*, 2024.
- [2] Shiqi He, Qifan Yan, Feijie Wu, Lanjun Wang, Mathias Lécuyer, and Ivan Beschastnikh. Gluefl: Reconciling client sampling and model masking for bandwidth efficient federated learning. In *Conference on Machine Learning and Systems (MLSys)*, 2023.
- [3] Kelly Kostopoulou, Pierre Tholoniati, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. Turbo: Effective caching in differentially-private databases. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 2023.
- [4] Mauricio Soroco, Joel Hempel, Xinze Xiong, Mathias Lécuyer, and Joséphine Gantois. Flowering onset detection: Traditional learning vs. deep learning performance in a sparse label context. In *Tackling Climate Change with Machine Learning: workshop at NeurIPS*, 2023.
- [5] Qiaoyue Tang and Mathias Lécuyer. Dp-adam: Correcting dp bias in adam's second moment estimation. In *Trustworthy and Reliable Large-Scale Machine Learning Models (RTML) Workshop at ICLR*, 2023.
- [6] Ali Behrouz, Mathias Lécuyer, Cynthia Rudin, and Margo Seltzer. Fast optimization of weighted sparse decision trees for use in optimal treatment regimes and optimal policy design. In *Advances in Interpretable Machine Learning and Artificial Intelligence (AIMLAI) Workshop*, 2022.
- [7] Jinkun Lin, Anqi Zhang, Mathias Lécuyer, Jinyang Li, Aurojit Panda, and Siddhartha Sen. Measuring the effect of training data on deep learning predictions via randomized experiments. In *International Conference on Machine Learning (ICML)*, 2022.
- [8] Mathias Lécuyer. Practical privacy filters and odometers with rényi differential privacy and applications to differentially private deep learning. *Preprint (arXiv)*, 2021.
- [9] Mathias Lécuyer, Sang Hoon Kim, Mihir Nanavati, Junchen Jiang, Siddhartha Sen, Amit

- Sharma, and Aleksandrs Slivkins. Sayer: Using implicit feedback to optimize system policies. In *ACM Symposium on Cloud Computing (SoCC)*, 2021.
- [10] Tao Luo, Mingen Pan, Pierre Tholoniati, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. Privacy budget scheduling. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2021.
 - [11] Mathias Lécuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)*, 2019.
 - [12] Mathias Lécuyer, Riley Spahn, Kiran Vodrahalli, Roxana Geambasu, and Daniel Hsu. Privacy accounting and quality control in the sage differentially private ML platform. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 2019.
 - [13] Mathias Lécuyer, Riley Spahn, Kiran Vodrahalli, Roxana Geambasu, and Daniel Hsu. Privacy accounting and quality control in the sage differentially private ML platform. In *ACM SIGOPS Operating Systems Review (OSR)*, 2019.
 - [14] Mathias Lécuyer, Joshua Lockerman, Lamont Nelson, Siddhartha Sen, Amit Sharma, and Aleksandrs Slivkins. Harvesting randomness to optimize distributed systems. In *The Seventeenth ACM Workshop on Hot Topics in Networks (HotNets)*, 2017.
 - [15] Mathias Lécuyer, Riley B. Spahn, Roxana Geambasu, Tzu-Kuo Huang, and Siddhartha Sen. Enhancing selectivity in big data. *Invited paper in the IEEE Security and Privacy Symposium Magazine*, 2017.
 - [16] Mathias Lécuyer, Riley B. Spahn, Roxana Geambasu, Tzu-Kuo Huang, and Siddhartha Sen. Pyramid: Enhancing selectivity in big data protection with count featurization. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)*, 2017.
 - [17] Mathias Lécuyer, Max Tucker, and Augustin Chaintreau. Improving the transparency of the sharing economy. In *Proceedings of the International World Wide Web Conference (WWW)*, 2017.
 - [18] Mathias Lécuyer, Riley B. Spahn, Giannis Spiliopoulos, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. Sunlight: Fine-grained targeting detection at scale with statistical confidence. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2015.
 - [19] Nicolas Viennot, Mathias Lécuyer, Jonathan Bell, Roxana Geambasu, and Jason Nieh. Synapse: New data integration abstractions for agile web application development. In *Proceedings of the European Conference on Computer Systems (EuroSys)*, 2015.
 - [20] Mathias Lécuyer, Guillaume Ducoffe, Francis Lan, Andrei Papancea, Theofilos Petsios, Riley Spahn, Augustin Chaintreau, and Roxana Geambasu. XRay: Increasing the web's transparency with differential correlation. In *Proceedings of the USENIX Security Symposium*, 2014.

Awards & Distinctions

- 2022 Google Research Scholar award (<https://research.google/outreach/research-scholar-program/>).
- 2017 Finalist for the Facebook PhD Fellowship (<https://research.fb.com/announcing-the-2017-facebook-phd-fellows/>).
- 2015 Finalist for the Microsoft Research Fellowship.
- 2011 “Bourse Carnot” Fellowship: an excellency fellowship for French students in the Unites-States, focused on Research and Entrepreneurship.

Invited Talks

- 2023 [Tutorial] *Privacy as hypothesis testing: linking Differential Privacy, membership attacks, and privacy audits*, Canadian AI, Responsible AI Track, Montréal.
- 2023 *DP-AdamBC: your DP-Adam is actually DP-SGD (unless you apply Bias Correction)*, Bridge the gap: Differential Privacy and Statistical Analysis, Amii Upper Bound Workshop, Edmonton.
- 2020 *Towards a Practical Differentially Private Machine Learning Platform*, Northwest Data Science Seminar Series.
- 2019 *Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform*, University of California Berkeley, Security Seminar.
- 2019 *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, University of British Columbia.
- 2019 *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, University of California Los Angeles.
- 2019 *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, Microsoft Research New York.
- 2018 *Certified Robustness to Adversarial Examples with Differential Privacy*, University of Washington, Security Seminar.
- 2018 *Certified Robustness to Adversarial Examples with Differential Privacy*, University of California Berkeley, Security Seminar.
- 2018 *Certified Robustness to Adversarial Examples with Differential Privacy*, Google Brain.
- 2018 *Harvesting Randomness for Counterfactual Evaluation of Systems*, Stanford, NetSeminar.
- 2018 *Certified Robustness to Adversarial Examples with Differential Privacy*, Stanford, Security Lunch.
- 2014 *XRay: Enhancing the Web's Transparency with Differential Correlation*, University of Washington and Microsoft Research Summer Institute.
- 2012 *Dispatch: weaving a safe web of news*, Columbia Journalism & Technology Breakfast.

Media Coverage

- 2021 Geneviève Lasalle, "Ottawa finance la création d'un outil pour déchiffrer les mots de passe," Radio Canada, 09 novembre 2021.
- 2020 Alexandra Pihen, "La vie privée en voie d'extinction," Science et Vie, Hors Série 290, 04 mars 2020.
- 2016 *The data republic*, "To safeguard democracy, the use of data should be made as transparent as possible," The Economist.
- 2016 Priya Kumar, *When Was the Last Time You Read a Privacy Policy?*, Slate.com.
- 2015 Tom Simonite, *Probing the Dark Side of Google's Ad-Targeting System*, MIT Technology Review.
- 2014 Steve Lohr, *XRay: A New Tool for Tracking the Use of Personal Data on the Web*, The New York Times.

Service

- 2024 Program Committee, Conference on Secure and Trustworthy Machine Learning (SaTML) 2024.
- 2023 Program Committee, Symposium on Operating Systems Design and Implementation (OSDI) 2023.

- 2023 Workshop Reviewer, International Conference on Learning Representations (ICLR) 2023.
- 2022/2023 Program Committee, Security & Privacy (Oakland) 2023.
- 2022 Reviewer, International Conference on Machine Learning (ICML) 2022.
- 2021/2022 Program Committee, Security & Privacy (Oakland) 2022.
- 2021 External Reviewer, USENIX Security 2022.
- 2021 Reviewer, NeurIPS 2021.
- 2020/2021 Program Committee, USENIX Security 2021.
- 2021 Program Committee, Conference on Machine Learning and Systems (MLSys) 2021.
- 2020 Reviewer, International Conference on Machine Learning 2020.
- 2019 Guest Lecture, graduate class “Decentralized Security: Theory and Systems” (CS 294-163).
Instructor: Raluca Ada Popa University of California Berkeley
- 2020 External Reviewer, Security & Privacy (Oakland) 2020.
- 2020 Program Committee, Eurosys 2020.
- 2019/2020 Program Committee, USENIX Security 2020.
- 2019 Reviewer, Journal of Machine Learning Research.
- 2019 Program Committee, ACM Symposium on Cloud Computing (SoCC) 2019.
- 2018 Program Committee, Systems for Machine Learning Workshop at the Conference on Neural Information Processing Systems, <http://learningsys.org/nips18/>.
- 2017 Reviewer, AISys Workshop at the ACM Symposium on Operating Systems Principles, <http://learningsys.org/sosp17/>.
- 2016 Reviewer, ACM Transactions on Internet Technology.
- 2016 Co-designed and delivered a programming workshop for the Annual Engineering Exploration,
& 2014 organized by the Society of Women Engineers for New York City high-school female students.