THE UNIVERSITY OF BRITISH COLUMBIA

Curriculum Vitae — Mathias Lécuyer

Surname: Lécuyer
First name: Mathias
Department: Computer Science
Faculty: Science
Present rank: Assistant Professor
Since: July 1, 2021

`https://mathias.lecuyer.me`
mathias.lecuyer@ubc.ca

Academic degrees

Columbia University	PhD	Computer Science	09/2013–08/2019
Columbia University	MSc	Computer Science	09/2011–05/2013
École Polytechnique	MSc	Computer Science	08/2008–08/2011

Employment record

Assistant Professor	University of British Columbia	2021/07 – present
Microsoft Research	Post-doctoral Researcher	09/2019 – 06/2021
Columbia University	Research Assistant	2013 – 2019
Microsoft Research	Research Intern	05/2017 – 07/2017
Columbia University	Teaching Assistant	2012 – 2014

Research interests

I work on trustworthy Artificial Intelligence (AI) systems, with a focus on enforcing provable guarantees in models and their data ecosystems. Despite all their promises, AI systems also introduce new safety challenges: AI models rely on large scale data collection and processing, introducing new risks to privacy and intellectual property; adversaries can manipulate model predictions without explicitly breaking into host machines, opening new attack surfaces in systems processing sensitive data; and AI models are often too opaque for end-users to understand and trust, and too brittle for system designers' security requirements.

A key challenge towards safe AI is the community's focus on average-case performance. This tradition is visible in theoretical foundations focusing on expected error, and in empirical benchmarks that assess models on average performance metrics. Conversely, AI safety challenges arise from pervasive adversarial interactions in AI deployments, and require worst-case analysis. My research aims to better understand the failure cases of AI systems, and develop rigorous data and AI systems with worst-case safety guarantees. My recent contributions tackle these challenges in four broad directions.

Privacy preserving data systems. AI systems rely on large scale data collection and aggregation, which exposes sensitive information. Indeed, even the output of a computation, such as a trained ML model, enables reconstruction attacks that expose individual inputs from the training data. Differential Privacy (DP) is the leading approach to preventing such data leakage in data-driven applications. However, DP suffers from practical challenges, including reduced utility, a mismatch between theoretical assumptions and empirical requirements in continuously operating systems, and a lack of practical mechanisms for resource allocation.

I develop new theory, algorithms, and system mechanisms for DP, to make end-to-end privacy preserving systems more practical. This effort follows three complementary directions. First, I design DP theory [C6, O1] and systems abstractions [C6, C8, C12, C18, C15] to support whole workloads on continuously growing data streams. My recent work [C15] serves as the DP blueprint for the Privacy-Preserving Attribution (advertising measurement) API under standardization efforts with the W3C. Second, I develop new algorithms to optimize AI models with DP [C13], to improve the utility of privacy preserving models. Third, I develop techniques to audit privacy leakage from trained models and AI systems [C17], to enable third parties to quantify privacy leakage from non DP AI models, even without access to, or control of, the training pipeline.

Adversarial robustness. AI models show a worrying susceptibility to adversarial attacks, in which an attacker applies imperceptible changes to the input to arbitrarily influence a target model. Such attacks can bypass fraud detection, trick self-driving cars, or jailbreak aligned foundation models. My work laid the foundations for Randomized Smoothing (RS) [C7], the only technique to provably ensure robustness that scales to the largest AI models. This work has been cited more than a thousand times, and is still actively developed. RS also serves as a building block for other AI safety tools, such as to enforce fairness guarantees, or create robust watermarks and unlearnable examples. I recently used Differential Privacy composition to extend RS to test-time adaptive models [C16]. This technique opens a new design space to address the shortcomings of RS.

Explainability. Another important tool to assess safety properties in AI models is to understand the impact of data on test-time behavior. I design explainability and transparency tools for AI models, that build on techniques from causal inference to understand how data influences model predictions, with explanations that are faithful to model behavior in manipulated settings [C1, C3, C10].

Causal Machine Learning. Deep learning time-series models often inform downstream decisions, using repeated forecasts on different values of a controllable feature (e.g., the price of a good) to select the best outcome (e.g., the demand yielding the highest revenue). Since optimizing decisions can lead to different actions than those present in the training set, there is an implicit requirement that time-series models will generalize to actions outside of the training distribution. Despite this core requirement, time-series models are typically only trained and evaluated on in-distribution predictive tasks. I design deep learning architectures, optimization procedures, and evaluation methods to learn causal models, that generalize better when forecasting the effect of actions outside of the training distribution. My first project [O2] is deployed at `wiremind.io`, to optimize prices for passenger rail and airlines.

Publications

Student authors whom I was **formally supervising** at the time of the work are in bold. Shared first authorship is denoted with *. Corresponding authors, usually a sign of involvements and contribution, are denoted with †.

Citation metrics. According to Google Scholar as of June 2025, my work has been cited 1,892 times, with an h-index of 12, and 15 papers with 10 or more citations. `https://scholar.google.com/citations?user=WeIvMTUAAAAJ&hl=en`.

Refereed publications

Journals

[J1] Mathias Lécuyer, *Riley Spahn*, *Kiran Vodrahalli*, Roxana Geambasu, and Daniel Hsu. "Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform". *ACM SIGOPS Operating Systems Review (OSR)* (2019).

Conference Proceedings

[C20] ***Bingshan Hu***, *Zhiming Huang*, *Tianyue H. Zhang*, Mathias Lécuyer, and Nidhi Hegde. "Connecting Thomson Sampling and UCB: Towards More Efficient Trade-offs Between Privacy and Regret". *International Conference on Machine Learning (ICML)*. 2025. AR 3260/12107=26.9%

[C19] *Qifan Yan*, *Andrew Liu*, *Shiqi He*, Mathias Lécuyer, and Ivan Beschastnikh. "FedFetch: Faster Federated Learning with Adaptive Downstream Prefetching". *IEEE International Conference on Computer Communications (INFOCOM)*. 2025. AR 272/1458=18.6%

[C18] *Pierre Tholoniat**, *Kelly Kostopoulou**, Mosharaf Chowdhury, Asaf Cidon, Roxana Geambasu, Mathias Lécuyer, and Junfeng Yang. "DPack: Efficiency-Oriented Privacy Budget Scheduling". *European Conference on Computer Systems (EuroSys)*. 2025. AR 241/1308=18.4%

[C17] 10+ cites ***Mishaal Kazmi****, *Hadrien Lautraite**, *Alireza Akbari**, ***Qiaoyue Tang****, ***Mauricio Soroco***, Tao Wang, Sébastien Gambs, and Mathias Lécuyer. "PANORAMIA: Privacy Auditing of Machine Learning Models without Retraining". *Conference on Neural Information Processing Systems (NeurIPS)*. 2024. AR 4043/15671=25.8%

[C16] ***Saiyue Lyu****, ***Shadab Shaikh****, ***Frederick Shpilevskiy****, Evan Shelhamer, and Mathias Lécuyer. "Adaptive Randomized Smoothing: Certified Adversarial Robustness for Multi-Step Defences". *Conference on Neural Information Processing Systems (NeurIPS)*. **(Spotlight)**. 2024. AR 4043/15671=25.8%; spotlight: 387/15671=2.5%

[C15] *Pierre Tholoniat*, *Kelly Kostopoulou*, *Peter McNeely*, *Prabhpreet Singh Sodhi*, *Anirudh Varanasi*, Benjamin Case, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. "Cookie Monster: Efficient On-Device Budgeting for Differentially-Private Ad-Measurement Systems". *ACM Symposium on Operating Systems Principles (SOSP)*. **(Distinguished Artifact Honorable Mention)**. 2024. AR 43/248=17.3%

[C14] ***Amir Sabzi***, *Rut Vora*, *Swati Goswami*, Margo Seltzer, Mathias Lécuyer, and Aastha Mehta. "NetShaper: A Differentially Private Network Side-Channel Mitigation System". *USENIX Security Symposium*. 2024. AR 98/515=19%

[C13] 10+ cites ***Qiaoyue Tang***, ***Frederick Shpilevskiy***, and Mathias Lécuyer. "DP-AdamBC: your DP-Adam is actually DP-SGD (unless you apply Bias Correction)". *AAAI Conference on Artificial Intelligence*. **(Oral)**. 2024. AR 2342/9862=24%; oral: 224/9862=2.3%

[C12] *Kelly Kostopoulou**, *Pierre Tholoniat**, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. "Turbo: Effective Caching in Differentially-Private Databases". *ACM Symposium on Operating Systems Principles (SOSP)*. 2023. AR 43/229=19%

[C11] 10+ cites ***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". *Conference on Machine Learning and Systems (MLSys)*. 2023. AR 46/207=22%

[C10] 50+ cites *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". *International Conference on Machine Learning (ICML)*. 2022. AR 1235/5630=22%

[C9] Mathias Lécuyer, *Sang Hoon Kim*, Mihir Nanavati, Junchen Jiang, Siddhartha Sen, Amit Sharma, and Aleksandrs Slivkins. "Sayer: Using Implicit Feedback to Optimize System Policies". *ACM Symposium on Cloud Computing (SoCC)*. 2021. AR 45/145=32%

[C8] 10+ cites *Tao Luo**, *Mingen Pan**, *Pierre Tholoniat**, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. "Privacy Budget Scheduling". *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 2021. AR 31/165=19%

[C7] **1000+ cites** *Mathias Lécuyer*, *Vaggelis Atlidakis*, Roxana Geambasu, Daniel Hsu, and Suman Jana. "Certified Robustness to Adversarial Examples with Differential Privacy". *IEEE Symposium on Security and Privacy (S&P Oakland)*. 2019. AR 60/457=13%

[C6] 50+ cites *Mathias Lécuyer*, *Riley Spahn*, *Kiran Vodrahalli*, Roxana Geambasu, and Daniel Hsu. "Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform". *ACM Symposium on Operating Systems Principles (SOSP)*. 2019. AR 38/276=14%

[C5] 10+ cites *Mathias Lécuyer**, *Riley B. Spahn**, Roxana Geambasu, Tzu-Kuo Huang, and Siddhartha Sen. "Pyramid: Enhancing Selectivity in Big Data Protection with Count Featurization". *IEEE Symposium on Security and Privacy (S&P Oakland)*. 2017. AR 60/457=13%

[C4] 10+ cites *Mathias Lécuyer**, *Max Tucker**, and Augustin Chaintreau. "Improving the Transparency of the Sharing Economy". *International World Wide Web Conference (WWW)*. 2017. AR 164/966=17%

[C3] 50+ cites *Mathias Lécuyer*, *Riley B. Spahn*, *Giannis Spiliopoulos*, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. "Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidences". *ACM Conference on Computer and Communications Security (CCS)*. 2015. AR 128/646=20%

[C2] 50+ cites *Nicolas Viennot*, *Mathias Lécuyer*, *Jonathan Bell*, Roxana Geambasu, and Jason Nieh. "Synapse: New Data Integration Abstractions for Agile Web Application Development." *European Conference on Computer Systems (EuroSys)*. 2015. AR 110/657=17%

[C1] **100+ cites** *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." *USENIX Security Symposium*. 2014. AR 67/350=19%

Workshop Proceedings

[W8] ***Frederick Shpilevskiy***, ***Saiyue Lyu***, Krishnamurthy Dj Dvijotham, Mathias Lécuyer, and Pierre-André Noël. "Adaptive Diffusion Denoised Smoothing : Certified Robustness via Randomized Smoothing with Differentially Private Guided Denoising Diffusion". *Workshop on Test-Time Adaptation: Putting Updates to the Test! at ICML*. **(Oral)**. 2025.

[W7] ***Qiaoyue Tang***, ***Alain Zhiyanov***, and Mathias Lécuyer. "On the Performance of Differentially Private Optimization with Heavy-Tail Class Imbalance". *High-dimensional Learning Dynamics (workshop at ICML)*. 2025.

[W6] ***Mishaal Kazmi****, *Hadrien Lautraite**, *Alireza Akbari**, ***Mauricio Soroco***, ***Qiaoyue Tang***, Tao Wang, Sébastien Gambs, and Mathias Lécuyer. "PANORAMIA: Privacy Auditing of Machine Learning Models without Retraining". *Theory and Practice of Differential Privacy (TPDP) workshop*. 2024.

[W5] ***Shadab Shaikh***, ***Saiyue Lyu***, ***Frederick Shpilevskiy***, Evan Shelhamer, and Mathias Lécuyer. "Adaptive Randomized Smoothing for Certified Multi-Step Defence". *Workshop on Test-Time Adaptation: Model, Adapt Thyself! (MAT), at CVPR*. 2024.

[W4] ***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". *Tackling Climate Change with Machine Learning: workshop at NeurIPS*. 2023.

[W3] ***Qiaoyue Tang*** and Mathias Lécuyer. "DP-Adam: Correcting DP Bias in Adam's Second Moment Estimation". *Trustworthy and Reliable Large-Scale Machine Learning Models (RTML) Workshop at ICLR*. 2023.

[W2] *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". *Advances in Interpretable Machine Learning and Artificial Intelligence (AIMLAI) Workshop*. 2022.

[W1] 10+ cites *Mathias Lécuyer*, *Joshua Lockerman*, *Lamont Nelson*, Siddhartha Sen, Amit Sharma, and Aleksandrs Slivkins. "Harvesting Randomness to Optimize Distributed Systems". *The Seventeenth ACM Workshop on Hot Topics in Networks (HotNets)*. 2017. AR 28/124=23%

Non-refereed publications

Invited Papers

[I1] 10+ cites *Mathias Lécuyer*, *Riley B. Spahn*, Roxana Geambasu, Tzu-Kuo Huang, and Siddhartha Sen. "Enhancing Selectivity in Big Data". *IEEE Security and Privacy Symposium Magazine* (2017).

Selected pre-prints

[O2] ***Thomas Crasson***, Yacine Nabet, and Mathias Lécuyer. "Training and Evaluating Causal Forecasting Models for Time-Series". *arXiv preprint arXiv:2411.00126*. Presented at the 2025 AGIFORS Revenue Management SG meeting. 2024.

[O1] 10+ cites Mathias Lécuyer. "Practical Privacy Filters and Odometers with Rényi Differential Privacy and Applications to Differentially Private Deep Learning". *arXiv preprint arXiv:2103.01379*. 2021.

Patents

[P1] Mathias Lécuyer. "Privacy Filters and Odometers For Deep Learning". US 2022/0327227 A1. United States of America. **Published** (June 11, 2024). 2021.

Consultant

2024-now Scientific Advisor, Wiremind

Awards and distinctions

Awards for Teaching

2021 21WT2 Positive Teaching Letter from the Dean of Science, for receiving some of the highest student evaluations in the Faculty of Science in 2021 Winter Term 2, for my class on Differential Privacy (CPSC 538L).

Awards for Scholarship

2024 ACM Symposium on Operating Systems Principles Distinguished Artifact Honorable Mention

2022 Google Research Scholar award

2011 Bourse Carnot Fellowship (an excellency fellowship for French students focused on Research and Entrepreneurship)

Awards for Service

2025 Distinguished Reviewer Award, 2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)

Other Awards

At Machine Learning conferences, the area chairs often select a small number of top papers for spotlights or oral presentations. These distinctions are considered prestigious.

2024 Spotlight (2.5% acceptance rate), Conference on Neural Information Processing Systems (NeurIPS).

2024 Oral (2.3% acceptance rate), AAAI Conference on Artificial Intelligence.

Invited Presentations

- *Training Causal Time-Series Models for Generalizable Forecasting*, International Conference on Statistics and Data Science, Vancouver, Canada, June 24, 2025.
- *Adversarial Robustness and Privacy Measurements using Hypothesis-tests*, PrivSec Lab, LATECE, UQAM, Montréal, Canada, May 5, 2025.
- *Adversarial Robustness and Privacy Measurements using Hypothesis-tests*, International Laboratory and Learning Systems (IRL–ILLS), Montréal, Canada, May 2, 2025.
- *Adversarial Robustness and Privacy Measurements using Hypothesis-tests*, University of Waterloo Cryptography, Security, and Privacy (CrySP), Canada, Apr 29, 2025.
- *Adversarial Robustness and Privacy Measurements using Hypothesis-tests*, CleverHans Lab for security and privacy of machine learning, Vector Institute and University of Toronto, Canada, Apr 28, 2025.
- *Adversarial Robustness and Privacy Measurements using Hypothesis-tests*, joint SRI and Vector AI Safety Reading Group, Toronto, Canada, Apr 23, 2025.
- *Adaptive Randomized Smoothing: Certified Adversarial Robustness for Multi-Step Defences*, Mathematics of Machine Learning, Canadian Mathematical Society (CMS) Winter meeting, Vancouver, Canada, Nov 30-Dec 01, 2024.
- *Security and Privacy in the age of Foundation Models*, MSRA Vancouver Talk Series, Vancouver, Canada, Aug 14, 2024.
- Panel: *Privacy and Connected Cars*, Intelligent Transportation Systems Society of Canada (ITS Canada) Annual Conference & Expo, Vancouver, Canada, Jun 20, 2024.
- *PANORAMIA: Efficient Privacy Auditing of Machine Learning Models without Retraining*, Statistical Aspects of Trustworthy Machine Learning, Banff International Research Station (BIRS) workshop, Banff, Canada, Feb 11-16, 2024.
- *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 on Privacy. Edmonton, Canada, May 23, 2023.
- *Towards a Practical Differentially Private Machine Learning Platform.* Northwest Data Science Seminar Series, Jul 22, 2020. Online.
- *Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform.* University of California Berkeley, USA, Security Seminar, Oct 15, 2019.
- *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, University of British Columbia, Canada, 2019.
- *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, University of California Los Angeles, USA, 2019.
- *Security, Privacy, and Transparency Guarantees for Machine Learning Systems*, Microsoft Research New York, USA, 2019.
- *Certified Robustness to Adversarial Examples with Differential Privacy*, University of Washington, Security Seminar, Seattle, USA, Oct 31, 2018.
- *Certified Robustness to Adversarial Examples with Differential Privacy*, University of California Berkeley, Security Seminar, Berkeley, USA, Oct 5, 2018.
- *Certified Robustness to Adversarial Examples with Differential Privacy*, Google Brain, June 8, USA, 2018.
- *Harvesting Randomness for Counterfactual Evaluation of Systems*, Stanford, NetSeminar, Jun 7, USA, 2018.
- *Certified Robustness to Adversarial Examples with Differential Privacy*, Stanford, Security Lunch, USA, Jun 6, 2018.
- *XRay: Enhancing the Web's Transparency with Differential Correlation*, University of Washington and Microsoft Research Summer Institute, USA, Jul 29, 2014.
- *Dispatch: weaving a safe web of news*, Columbia Journalism & Technology Breakfast, Nov 29, USA, 2012.

Other Presentations

- *Privacy & AI*: Technical Executive Summary, UBC CS Vancouver Senior Women in Tech event, Jun 13, 2024.
- *Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform*: The 27^{th} ACM Symposium on Operating Systems Principles, Huntsville, Ontario, Canada, Oct 28, 2019.
- *Certified robustness to adversarial examples with differential privacy*: The 40^{th} IEEE Symposium on Security and Privacy, San Francisco, California, USA, May 21, 2019.
- *Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence*: The 22^{nd} ACM Conference on Computer and Communications Security, Denver, Colorado, USA, Oct 14, 2015.
- *XRay: Enhancing the Web's Transparency with Differential Correlation*: The 23^{rd} USENIX Security Sumposium, San Diego, California, USA, Aug 20, 2014.

Tutorials and guest lectures

- *Privacy as hypothesis testing: linking Differential Privacy, membership attacks, and privacy audits*, Responsible AI, Canadian AI, Montreal, Jun 09, 2023.
- *Guest lecture on privacy infrastructure*, Decentralized Security: Theory and Systems (CS 294-163). Instructor: Raluca Ada Popa. University of California Berkeley, 2019.

Selected media and publicity

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

Student Supervision

In the supervisory role column, *Supervisor** denotes a co-supervision arrangement in which I acted as the main supervisor. *Co-supervisor** denotes equal supervisory responsibility.

Postdoc Supervision

Student Name	Program Type	Year		Supervisory Role (supervisor, co-supervisor)
		Start	Finish	
Bingshan Hu	Postdoc	2023/05	ongoing	Co-supervisor

PhD Supervision

Student Name	Program Type	Year		Supervisory Role (supervisor, co-supervisor)
		Start	Finish	
Saiyue Lyu	PhD	2022/09	ongoing	Supervisor
Qiaoyue Tang	PhD	2021/09	ongoing	Supervisor
Frederick Shpilevskiy	PhD	2023/09	ongoing	Supervisor

Distinctions: · Frederick Shpilevskiy received UBC's Four Year Doctoral Fellowship (4YF).

MSc supervision

Student Name	Program Type	Year		Supervisory Role (supervisor, co-supervisor)
		Start	Finish	
Shadab Shaikh	MSc	2021/09	2024/12	Supervisor
Mishaal Kazmi	MSc	2021/09	2024/12	Supervisor*
Haley Li	MSc (essay)	2021/09	2024/08	Supervisor
Amir Sabzi	MSc	2021/09	2024/08	Co-supervisor*
Shiqi He	MSc	2019/09	2023/08	Co-supervisor

Notable next affiliations: · Mishaal Kazmi: Ph.D. student at Northeastern University · Amir Sabzi: Ph.D. student at Princeton · Shiqi He: Ph.D. student at the University of Michigan.

Undergraduate Supervision

Student Name	Program Type	Year		Supervisory Role (supervisor, co-supervisor)
		Start	Finish	
Alain Zhiyanov	Directed Studies	2025/01	ongoing	Supervisor
Jessica Bator	Directed Studies (2 terms)	2024/06	2024/12	Supervisor
Helen Chen	Directed Studies (2 terms)	2023/09	2024/04	Supervisor
Mauricio Matias Soroco	Directed Studies (2 terms)	2023/09	2024/04	Supervisor
	Summer internship (USRA)	2023/05	2023/08	Supervisor
Ryan Shar	Honors Thesis	2023/09	2024/04	Supervisor
	Summer internship (USRA)	2023/05	2023/08	Supervisor
Eric Xiong	Honors Thesis	2023/09	2024/04	Supervisor
	Summer internship (SURE)	2023/05	2023/08	Supervisor
Joel Hempel	Directed Studies (2 terms)	2023/05	2023/08	Supervisor
Frederick Shpilevskiy	Summer internship (SURE)	2023/05	2023/08	Supervisor
	Honors Thesis	2022/09	2023/04	Supervisor*

Notable next affiliations: · Mauricio Matias Soroco: Ph.D. student at Simon Fraser University · Eric Xiong: M.Sc. student (research thesis) at the University of Alberta · Helen Chen: Amazon · Ryan Shar: MS student (course based) at Carnegie Mellon University · Frederick Shpilevskiy: Ph.D. student at the University of British Columbia

Other supervision

Student Name	Program Type	Year		Supervisory Role (supervisor, co-supervisor)
		Start	Finish	
Erell Boutin Jeanniard du Dot	MSc École des Mines de Saint-Étienne	2025/05	2025/08	Internship supervisor
Perrine Porcher	MSc École Polytechnique	2025/03	2025/07	Internship supervisor
Bastien Gless	MSc ENSTA	2024/04	2024/08	Internship supervisor
Matthieu Marquis-Lorber	MSc École Polytechnique	2024/04	2024/08	Internship supervisor
Nicolas Welti	MSc École Polytechnique	2024/04	2024/08	Internship supervisor
Romain Guth	MSc École Polytechnique	2024/04	2024/08	Internship supervisor
Thomas Crasson	MSc École Polytechnique	2023/04	2023/08	Internship supervisor
Ivan Bettannier	MSc École Polytechnique	2023/04	2023/08	Internship supervisor

Service to the community

Conference program committee: In Computer Science, conferences publish proceedings which are the most prestigious publication venues in the field. Program committees (PCs) are responsible for reviewing submitted papers, and making decisions on which papers are published. PCs is systems and security & privacy have a heavy load, with 15-20 papers to review. System PCs (e.g., OSDI, SOSP) are particularly selective and prestigious. Security & Privacy Associate Chair are selected to oversee reviewers on a large number of papers, and make final decisions: this is also a prestigious position in the community.

- The 47^{th} IEEE Symposium on Security and Privacy (S&P 2026), **Associate Chair**
- The 3^{rd} IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2025)
- The 42^{nd} International Conference on Machine Learning (ICML 2025)
- The 2^{nd} IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2024)
- The 38^{th} Conference on Neural Information Processing Systems (NeurIPS 2024)
- The 17^{th} USENIX Symposium on Operating Systems Design and Implementation (OSDI 2023)
- The 43^{rd} IEEE Symposium on Security and Privacy (S&P 2023)
- The 42^{nd} IEEE Symposium on Security and Privacy (S&P 2022)
- The 39^{th} International Conference on Machine Learning (ICML 2022)
- The 35^{th} Conference on Neural Information Processing Systems (NeurIPS 2021)
- The 30^{th} USENIX Security Symposium (2021)
- The 4^{th} Conference on Machine Learning and Systems (MLSys 2021)
- The 38^{th} International Conference on Machine Learning (ICML 2020)
- The 29^{th} USENIX Security Symposium (2020)
- The EuroSys 2020 conference (2020)
- The ACM Symposium on Cloud Computing 2019 (SoCC 2019)
- The Workshop on Systems for ML and Open Source Software at NeurIPS 2018

Scholarly committees

- Research Council Member, Canadian AI Safety Institute (CAISI) Research Program at CIFAR.

Workshop organization

- Organized a one-day workshop between the Systopia Lab at UBC and researchers from Microsoft Research Asia, Vancouver lab. The event was successful, with about 60 attendees from CS and ECE, 3 MSR talks, 7 UBC faculty lightning talks, and 10+ posters (October 4, 2024).
- Co-organized the TrustML @ UBC workshop, including finding and inviting speakers, working on the schedule, advertising, and chairing one session. The event was successful, with about 100 attendees, 4 long talks, 8 short talks, and 15+ posters (February 28, 2024).
- Co-organized the TrustML @ UBC workshop, including finding and inviting speakers, working on the schedule, advertising, and chairing two sessions. The event was successful, with about 100 attendees, 7 long talks, 8 short talks, and 15+ posters (June 23, 2023).

Outreach

- Lead an outreach event at Southlands Elementary School, the catchment school of the Musqueam people (Feb 26, 2025). I got a small grant to purchase 5 Sphero bots; organized 10 student volunteers to run an hour long programming activity to teach block programming, with the end-goal of guiding a Sphero bot through a printed maze; ran the activity with two 7^{th} grade classes, for a total of 51 students over two hours.
- Mentored two high school students from First Nations for an initiation to ML research (second edition), for the Verna J. Kirkness Education Foundation Program (July 8 to 12, 2024).
- Designed a one week initiation to ML research, and mentored two high school students from First Nations for the Verna J. Kirkness Education Foundation Program (May 29 to June 3, 2023).
- Tutorial on web application architecture through "person-in-the-middle" attacks, at the Microsoft Research NY Data Science summer school, 2017.
- Co-designed and delivered a programming workshop for the Annual Engineering Exploration, organized by the Society of Women Engineers for New York City high-school female students, 2016.
- Co-designed and delivered a programming workshop for the Annual Engineering Exploration, organized by the Society of Women Engineers for New York City high-school female students, 2014.