

Anvith Thudi

anvith.com

✉ anvith.thudi@mail.utoronto.ca

Education

University of Toronto

Ph.D. in Computer Science

Toronto, ON, Canada

Sep. 2022 - ongoing

- Advisors: Nicolas Papernot and Chris Maddison
- Thesis (tentative): “Data Problems in Modern Machine Learning”

University of Toronto

B.Sc in Mathematics, Spent Fall 2020 in Engineering Science

Toronto, ON, Canada

Sep. 2020 - May 2022

- GPA: 3.92/4.0

Simon Fraser University

Concurrent Studies Student (attended while in highschool)

Burnaby, BC, Canada

Sep. 2017 - May 2020

- GPA: 4.09/4.33

Awards and Honours

2026 Ontario Graduate Scholarship: Province of Ontario and University of Toronto

Gold Reviewer: ICML 2026

Notable Reviewer: ICLR 2025

2023 Canada Graduate Scholarship-Doctoral: Natural Sciences and Engineering Research Council of Canada

- *declined due to Vanier*

2023 Vanier Canada Graduate Scholarship: Natural Sciences and Engineering Research Council of Canada

- *Rank 1/173 of national round nominees (Ph.D. students in the Natural Sciences or Engineering)*

2022 Doctoral Entrance Scholarship: UofT Department of Computer Science

2022 Doctoral Recruitment Award: UofT Faculty of Arts and Science

2022 Galois Award: University College UofT

Dean’s List (2020-2022): UofT

2020 Loran Scholarship National Finalist: Loran Scholar’s Foundation

- *Top 88 highschool students in Canada*

Publications

Journal Proceedings

“k-Nearest Neighbour Adaptive Sampling (kNN-AS), a Simple Tool to Efficiently Explore Conformational Space”: Evianne M. Rovers, **Anvith Thudi**, Jérôme Hénin, Chris Maddison, Matthieu Schapira. *Journal of Chemical Theory and Computation* (2025)

“Selective Prediction via Neural Training Dynamics”: Stephan Rabanser, **Anvith Thudi**, Kimia Hamidieh, Adam Dziedzic, Nicolas Papernot. *Transactions on Machine Learning Research* (2025)

“From Differential Privacy to Bounds on Membership Inference: Less can be More”: **Anvith Thudi**, Ilya Shumailov, Franziska Boenisch, Nicolas Papernot. *Transactions on Machine Learning Research* (2024)

Conference Proceedings

“Efficient Public Verification of Private ML via Regularization”: Zoë Ruha Bell, **Anvith Thudi**, Olive Franzese-

McLaughlin, Nicolas Papernot, Shafi Goldwasser. *Proceedings of the 43rd International Conference on Machine Learning*. (2026)

"Gauss-Newton Unlearning for the LLM Era": Lev McKinney, **Anvith Thudi**, Juhan Bae, Tara Rezaei Kheirkhah, Nicolas Papernot, Sheila A. McIlraith, Roger Baker Grosse. *Proceedings of the 4th IEEE Conference on Secure and Trustworthy Machine Learning* (2026)

"Leveraging Per-Instance Privacy for Machine Unlearning": Nazanin Mohammadi Sepahvand, **Anvith Thudi**, Berivan Isik, Ashmita Bhattacharyya, Nicolas Papernot, Eleni Triantafillou, Daniel M. Roy, Gintare Karolina Dziugaite. *Proceedings of the 42nd International Conference on Machine Learning*. Oral at TPDP workshop 2025 (2025)

"Fast Exact Unlearning for In-context Learning Data for LLMs": Andrei Muresanu, **Anvith Thudi**, Michael R. Zhang, Nicolas Papernot. *Proceedings of the 42nd International Conference on Machine Learning* (2025)

"MixMin: Finding Data Mixtures via Convex Minimization": **Anvith Thudi**, Evianne Rovers, Yangjun Ruan, Tristan Thrush, Chris J. Maddison. *Proceedings of the 42nd International Conference on Machine Learning* (2025)

"MixMax: Distributional Robustness in Function Space via Optimal Data Mixtures": **Anvith Thudi**, Chris J. Maddison. *Proceedings of the 13th International Conference on Learning Representations* (2025)

"Gradients Look Alike: Sensitivity is Often Overestimated in DP-SGD": **Anvith Thudi**, Hengrui Jia, Casey Meehan, Ilia Shumailov, Nicolas Papernot. *Proceedings of the 33rd USENIX Security Symposium* (2024)

"Better Sparsifiers for Directed Eulerian Graphs": Sushant Sachdeva, **Anvith Thudi**, Yibin Zhao. *Proceedings of the 51st EATCS International Colloquium on Automata, Languages and Programming* (2024)

"Training Private Models That Know What They Don't Know": Stephan Rabanser, **Anvith Thudi**, Abhradeep Thakurta, Krishnamurthy Dvijotham, Nicolas Papernot. *Proceedings of the 37th Conference on Neural Information Processing Systems* (2023)

"Proof-of-Learning is Currently More Broken Than You Think": Congyu Fang, Hengrui Jia, **Anvith Thudi**, Mohammad Yaghini, Christopher A. Choquette-Choo, Natalie Dullerud, Varun Chandrasekaran, Nicolas Papernot. *Proceedings of the 8th IEEE European Symposium on Security and Privacy* (2023)

"On the Necessity of Auditable Algorithmic Definitions for Machine Unlearning": **Anvith Thudi**, Hengrui Jia, Ilia Shumailov, Nicolas Papernot. *Proceedings of the 31st USENIX Security Symposium* (2022)

"Unrolling SGD: Understanding Factors Influencing Machine Unlearning": **Anvith Thudi**, Gabriel Deza, Varun Chandrasekaran, Nicolas Papernot. *Proceedings of the 7th IEEE European Symposium on Security and Privacy* (2022)

"Proof of Learning: Definitions and Practice": Hengrui Jia, Mohammad Yaghini, Christopher A. Choquette-Choo, Natalie Dullerud, **Anvith Thudi**, Varun Chandrasekaran, Nicolas Papernot. *Proceedings of the 42nd IEEE Symposium on Security and Privacy* (2021)

Preprints.....

"SoK: Machine Learning Governance": Varun Chandrasekaran, Hengrui Jia, **Anvith Thudi**, Adelin Travers, Mohammad Yaghini, Nicolas Papernot

Experience

Microsoft Research Cambridge

Ph.D. Research Intern

Cambridge, UK

May. 2023 - July 2023

Radical Numerics

Consultant

Remote

May. 2026 - Ongoing

Talks

"Curating Training Data from Multiple Distributions": ETH Zurich 2026

"Unlearning Can Be Easy": OATML Group, Oxford 2026

"Making Datasets from Multiple Data Distributions": CIA Group, Cambridge 2026

"Unlearning Can Be Easy": Inria Montpellier 2026

"Leveraging Per-Instance Privacy for Machine Unlearning": Vector Institute ICML 2025 Conference Highlights 2025

"Leveraging Per-Instance Privacy for Machine Unlearning": Google DeepMind 2025

"Making Datasets from Multiple Data Distributions": Social Foundations of Computation Group, MPI Tübingen	2025
"Making Datasets from Multiple Data Distributions": University of British Columbia	2025
"Unlearning Can Be Easy": University of Wisconsin-Madison Security and Privacy Seminar	2025
"Datapoints that are Easy to Unlearn": Google DeepMind	2024
"Gradients Look Alike: Sensitivity is Often Overestimated in DP-SGD": USENIX Security 24'	2024
"Datapoints that are Easy to Unlearn": Harvard Efficient ML Seminar	2024
"The Unlearning Problem(s)": CS 562 at University of Illinois Urbana-Champaign	2024
"The Unlearning Problem(s)": The Alan Turing Institute	2023
"The Unlearning Problem(s)": Cambridge	2023
"The Unlearning Problem(s)": Google	2023
"The Unlearning Problem(s)": EPFL	2023
"The Unlearning Problem(s)": ETH Zurich	2023
"On the Necessity of Auditable Algorithmic Definitions for Machine Unlearning": USENIX Security 22' 2022	
"Unrolling SGD: Understanding Factors Influencing Machine Unlearning": Euro S&P 22'	2022
"The Unlearning Problem(s)": Meta	2022

Service

Reviewer

Euro S&P: IEEE European Symposium on Security and Privacy	2022
ICLR: International Conference on Learning Representations	2025,2026
ICML: International Conference on Machine Learning	2025,2026
L2M2 Workshop at ACL: Workshop on Large Language Model Memorization, Association for Computational Linguistics	2025
Neurips: Neural Information Processing Systems	2025,2026

Subreviewer

IEEE S&P: IEEE Symposium on Security and Privacy	2024
CCS: ACM Conference on Computer and Communications Security	2023
Neurips: Neural Information Processing Systems	2022
Asiacrypt: International Conference on the Theory and Application of Cryptology and Information Security 2026	

Panel

Neurips 2023 Unlearning Competition:	2026
--------------------------------------	------

Organizer

ML Lunch Talks at Vector: Vector Institute, Toronto	2024 - ongoing
Beyond Algorithms: Machine Unlearning from a Multidisciplinary Perspective Workshop: Toronto 2026	

Teaching

CSC2559 Trustworthy Machine Learning: Lecturer, University of Toronto	Winter 2027
---	-------------