Christopher A. Choquette-Choo

christopherchoquette.com
christopher-choquette-choo
cchoquette
CA, USA

Machine Learning Researcher *with 20+ papers as well as direct experience deploying my work into 6 products and indirectly into 30+ downstream products*

Research Experience

Google Brain & Google DeepMind — *Mountain View, CA, USA*
Machine Learning Researcher — *2022 – Present*

- Lead memorization analysis in large language models. Research how memorization manifests.
- Research into security vulnerabilities and auditing of machine learning and language models.
- Research and develop state-of-the-art differential privacy mechanisms for machine learning.
- Lead research into compression in federated learning.
- Deploy my techniques for compression, memorization analysis, and differential privacy into production.
- 3 spot bonuses for exceptional work, including for PaLM 2 tech awards and release as well as Gemini.
- 300+ CLs, 1 competition, 10+ papers released to date.

Google Research, Cerebra team — *New York, NY, USA*
Brain Resident — *2020 – 2022*

- Investigated concept interpretability of acoustic models. Presented at Google Research Conference.
- Led research into optimal privacy-communication-accuracy tradeoffs with sparsity in federated learning.
- Researched differentially private multi-winner voting mechanisms for machine learning.
- Guided and advise project into private semi-supervised learning for federated learning in dermatology.

Vector Institute, with Professor Nicolas Papernot — *Toronto, ON, Canada*
Research Assistant — *Sept 2019 – Oct. 2020*

- Led research into differentially private collaborative algorithms.
- Led Privacy-preserving machine learning.

Georgian Partners — *Toronto, ON, Canada*
Research Engineer — *Apr. 2019 – Aug. 2019*

- Owned development of a differentially private ML model, to guarantee user data privacy, in collaboration with Google's top machine learning library, TensorFlow/Privacy, which is used by 1000 people.
- Designed an AutoML package to intelligently tune an ML model on any dataset; used by 25+ people.

Vector Institute, with Professor Aspuru-Guzik — *Toronto, ON, Canada*
Undergraduate Researcher — *Apr. 2019 – Aug. 2019*

- Researched machine learning for molecular discovery via Gaussian processes and active learning.

Intel Corp. — *Toronto, ON, Canada*
Research Engineer — *May 2018 – May 2019*

- Spearheaded SOTA ML bug triager with 55% accuracy on 2000+ engineers and 76% on 500+ teams.
- Productionized triager with an engineering efficiency improvement of 25% and savings of >$10M annually.

Institute of Biomaterials and Biomedical Engineering with Professor Paul Santerre — *Toronto, ON, Canada*
Undergraduate Researcher — *Apr. 2016 – Sept. 2016*

- Studied mechanical properties of polyurethane scaffolds and dental resin composites. Used in patents.

Research and Papers

Peer-Reviewed Conference and Journal Proceedings

[20] *Multi-epoch matrix factorization mechanisms for private machine learning* Link — *2023*
Proceedings of the 40th International Conference on Machine Learning (ICML)
Christopher A. Choquette-Choo, H. Brendan McMahan, Keith Rush, Abhradeep Thakurta.

[19] *Private Federated Learning with Autotuned Compression* Link — *2023*
Proceedings of the 40th International Conference on Machine Learning (ICML)
Enayat Ullah*, **Christopher A. Choquette-Choo***, Peter Kairouz*, Sewoong Oh*.

[18] *Federated Learning of Gboard Language Models with Differential Privacy* Link — *2023*
The 61st Annual Meeting of the Association for Computational Linguistics
Zheng Xu, Yanxiang Zhang, Galen Andrew, **Christopher A. Choquette-Choo**, Peter Kairouz, H. Brendan McMahan, Jesse Rosenstock, Yuanbo Zhang.

[17] *Preventing verbatim memorization in language models gives a false sense of privacy* Link — *2023*
Proceedings of the 15th International Natural Language Generation Conference
Daphne Ippolito, Florian Tramèr*, Milad Nasr*, Chiyuan Zhang*, Matthew Jagielski*, Katherine Lee*, **Christopher A. Choquette-Choo***, Nicholas Carlini.

[16] *Proof-of-Learning is Currently More Broken Than You Think* Link — *2023*
2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P). IEEE Computer Society
Congyu Fang*, Hengrui Jia*, Anvith Thudi, Mohammad Yaghini, **Christopher A. Choquette-Choo**, Natalie Dullerud, Varun Chandrasekaran, Nicolas Papernot.

[15] *Private Multi-Winner Voting for Machine Learning* Link — *2023*
Proceedings on 23rd Privacy Enhancing Technologies Symposium (PETS)
Adam Dziedzic, **Christopher A. Choquette-Choo**, Natalie Dullerud, Vinith Menon Suriyakumar, Ali Shahin Shamsabadi, Muhammad Ahmad Kaleem, Somesh Jha.

[14] *The fundamental price of secure aggregation in differentially private federated learning* Link — *2022*
International Conference on Machine Learning. PMLR
Wei-ning Chen*, **Christopher A. Choquette-Choo***, Peter Kairouz*, Ananda Theertha Suresh*.

[13] *Label-Only Membership Inference Attacks* Link — *2021*
International Conference on Machine Learning (ICML)
Christopher A. Choquette-Choo, Florian Tramer, Nicholas Carlini, Nicolas Papernot.

[12] *Entangled Watermarks as a Defense against Model Extraction* Link — *2021*
USENIX Security Symposium (USENIX)
Hengrui Jia, **Christopher A. Choquette-Choo**, Varun Chandrasekaran, Nicolas Papernot.

[11] *Proof of Learning: Definitions and Practice* Link — *2021*
IEEE Symposium on Security and Privacy (IEEE S&P)
Hengrui Jia*, Mohammad Yaghini*, **Christopher A Choquette-Choo**,ˆ Natalie Dullerud,ˆ Anvith Thudi,ˆ Varun Chandrasekaran, Nicolas Papernot.

[10] *Machine Unlearning* Link — *2021*
IEEE Symposium on Security and Privacy (IEEE S&P)
Lucas Bourtoule*, Varun Chandrasekaran*, **Christopher A. Choquette-Choo***, Hengrui Jia*, Adelin Travers*, Baiwu Zhang*, David Lie, Nicolas Papernot.

[9] *CaPC Learning: Confidential and Private Collaborative Learning* Link — *2021*
International Conference on Learning Representations (ICLR)
Christopher A. Choquette-Choo*, Natalie Dullerud*, Adam Dziedzic*, Yunxiang Zhang*, Somesh Jha, Nicolas Papernot, Xiao Wang.

[8] *A Multi-label, Dual-Output Deep Neural Network for Automated Bug Triaging* Link *2019*
International Conference on Machine Learning and Applications (ICMLA)
Christopher A. Choquette-Choo, David Sheldon, Jonny Proppe, John Alphonso-Gibbs, Harsha Gupta.

Peer-Reviewed Workshop Proceedings

[7] *Communication Efficient Federated Learning with Secure Aggregation and Differential Privacy* Link *2021*
the Neural Information Processing Systems (NeurIPS) workshop on Privacy in Machine Learning
Wei-ning Chen*, Christopher A. Choquette-Choo*, Peter Kairouz*.

Pre-Prints (arXiv)

[6] *Palm 2 technical report* Link *2023*
arXiv
Anil, R., Dai, A. M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., ..., **Christopher A. Choquette-Choo**, ..., & Wu, Y.

[5] *Poisoning web-scale training datasets is practical* Link *2023*
arXiv
Nicholas Carlini, Matthew Jagielski, **Christopher A. Choquette-Choo**, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, Florian Tramèr.

[4] *Are aligned neural networks adversarially aligned?* Link *2023*
arXiv preprint arXiv:2306.15447
Nicholas Carlini, Milad Nasr, **Christopher A. Choquette-Choo**, Matthew Jagielski, Irena Gao, Anas Awadalla, Pang Wei Koh, Daphne Ippolito, Katherine Lee, Florian Tramèr, Ludwig Schmidt.

[3] *(Amplified) Banded Matrix Factorization: A unified approach to private training* Link *2023*
arXiv
Christopher A. Choquette-Choo, Arun Ganesh, Ryan McKenna, H. Brendan McMahan, Keith Rush, Abhradeep Guha Thakurta, Zheng Xu.

[2] *Students Parrot Their Teachers: Membership Inference on Model Distillation* Link *2023*
arXiv preprint arXiv:2303.03446
Matthew Jagielski, Milad Nasr, Katherine Lee, **Christopher A. Choquette-Choo**, Nicholas Carlini.

[1] *Fine-tuning with differential privacy necessitates an additional hyperparameter search* Link *2022*
arXiv
Yannis Cattan, **Christopher A Choquette-Choo**, Nicolas Papernot, Abhradeep Thakurta.

Under Review (and not yet released)

[0] *Doubly Robust Peer-To-Peer Learning Protocol* Link *Under Review*
Under Review
Nicholas Franzese, Adam Dziedzic, **Christopher A. Choquette-Choo**, Mark R. Thomas, Muhammad Ahmad Kaleem, Stephan Rabanser, Congyu Fang, Somesh Jha, Nicolas Papernot, Xiao Wang

[-1] *Privacy Side-Channels in Machine Learning* Link *Under Review*
Under Review
Edoardo Debenedetti, Giorgio Severi, Milad Nasr, **Christopher A. Choquette-Choo**, Matthew Jagielski, Eric Wallace, Nicholas Carlini, Florian Tramèr

Talks

Invited Talks

The Privacy Considerations of Production Machine Learning *2021*
MLOps New York Area Summit *Slides available upon request.*

Adversarial Machine Learning: Ensuring Security and Privacy of ML Models and Sensitive Data *2019*
REWORK Responsible AI Summit *Available as a part of the Privacy and Security in Machine Learning package*

Paper Presentations

Multi-Epoch Matrix Factorization Mechanisms for Private Machine Learning	*Oral presentation at ICML 2023 (Skip to 1:55:49)*
The Fundamental Price of Secure Aggregation in Differentially Private Machine Learning	*ICML 2022*
Label-Only Membership Inference Attacks	*Spotlight at ICML 2021*
Proof-of-Learning Definitions and Practice	*IEEE S&P 2021*
Machine Unlearning	*Oral presentation at IEEE S&P 2021*

Professional Activities

Program Committee

IEEE Security and Privacy (S&P) conference	*2024*
Generative AI + Law (GenLaw)'23 Workshop at ICML	*2023*

Session Chair

DL: Robustness at International Conference on Machine Learning (ICML)	*2022*

Reviewer

Nature Machine Intelligence Journal	*2023*
Neural Information Processing Systems (NeurIPS)	*2023*
International Conference on Machine Learning (ICML)	*2023*
Neural Information Processing Systems (NeurIPS)	*2022*
Nature Machine Intelligence Journal	*2022*
International Conference on Machine Learning (ICML) + Outstanding	*2022*
IEEE Transactions on Emerging Topics in Computing	*2022*
Machine Learning for the Developing World (ML4D) workshop at NeurIPS	*2021*
Journal of Machine Learning Research	*2021*
Machine Learning for the Developing World (ML4D) workshop at NeurIPS	*2020*

External Reviewer

USENIX Security Symposium	*2022*
IEEE Symposium on Security and Privacy	*2022*
International Conference on Machine Learning (ICML)	*2021*
USENIX Security Symposium	*2021*
IEEE Symposium on Security and Privacy	*2021*

Education

Bachelor of Applied Science in Engineering Science — *University of Toronto*
Major in Robotics Engineering — *2015-2020*
Thesis: Label-Only Membership Inference Attacks as Realistic Privacy Threats
Graduation with Honors (cGPA 3.73/4.00)

Honors and Awards

Schulich Leaders Full Scholarship — *University of Toronto*
$100,000 Value — *2015-2020*
Awarded on the basis of academic achievement and leadership to students pursuing a STEM degree.

Class of 9T7 Award — *University of Toronto*
$4000 Value — *2017*
Awarded on the basis of academic achievement and leadership.

Director's Summer Research Opportunities — *University of Toronto*
$5000 Value — *2016*
Awarded to fund a summer research opportunity in Canada at the Institute for Biomaterials and Biomedical Engineering.

Burger King Scholarship — *University of Toronto*
$1500 Value — *2015*
Awarded on the basis of academic achievement and leadership.

University of Toronto Scholarship — *University of Toronto*
$6000 Value — *2015*
Awarded on the basis of academic achievement.

Competitions

Undergraduate Science Case Competition (SCINAPSE) — *Western University*
(Finalist of 2) of 250+ teams. Upper Year Division. — *2017*

Microsoft Azure Machine Learning Case Competition — *University of Toronto*
(1st) of 20+ teams. — *2017*

UTEK Consulting Competition — *University of Toronto*
(Semi-Finalist) of 20+ teams. — *2016*

The Game, Engineering Design Competition — *University of Toronto*
(1st) of 10+ teams. $10,000 value. — *Sept. 2015 - Mar. 2016*

Community Outreach

Public Software

Google Research: Main Owner of Multi-Epoch Matrix Factorization package — *2023*

Google Research: Owner of Private Linear Compression — *2022*

TensorFlow Privacy: Sole Contributor of Bolt-On Method for Differentially Private Training — *2019*

CleverHans Blog

Arbitrating the integrity of stochastic gradient descent with proof-of-learning	*2021*
Beyond federation: collaborating in ML with confidentiality and privacy	*2021*
Teaching Machines to Unlearn	*2020*

Personal Blog

How to do Machine Unlearning	*2021*
Teaching Machines to Unlearn	*2020*

Community Service and Leadership

University of Toronto Consulting Association, University of Toronto *University of Toronto*
Director of Volunteer Consulting Group *2017-2018*

FoodSkrap Startup *Own Incorporation*
Co-Founder, CEO, and Software Developer *2016-2017*

You're Next Career Network *University of Toronto*
Director of Business Development, Startup *2016-2017*

Board of Directors *Plan Canada*
Youth Advisor *2015-2017*

Youth Advisory Council *Plan Canada*
Member *2014-2017*

Technical skills

Proficient in:	Python, C
Familiar with:	Java, MATLAB, Perl, SQL, Elasticsearch, JavaScript
Python libraries:	TensorFlow, Jax, Pax, SeqIO, T5X, PyTorch, NumPy, Pandas, Matplotlib, Scikit-learn, TensorFlow Federated, TensorFlow Privacy

Soft skills

Communication I focus on communicating complex ideas in a way anyone can understand.

Teamwork I care about being considerate and sharing responsibility in effective ways. Evidenced by 6 peer bonuses and 1 kudos at Google.

Leadership I believe that identifying strengths and clearing runways enables success.