GAUTAM KAMATH

ADDRESS
Cheriton School of Computer Science
Davis Centre
Room 3124
200 University Ave W
Waterloo, ON N2L 3G1

CONTACT
Personal Website: `www.gautamkamath.com`
Group Website: `thesalon.github.io`
Cell: (657) 206-7724
Email: gckamath@uwaterloo.ca

RESEARCH INTERESTS

Reliable and trustworthy algorithms, statistics, and machine learning, particularly privacy and robustness.

PROFESSIONAL APPOINTMENTS

University of Waterloo
David R. Cheriton School of Computer Science
Assistant Professor
July 2019 - Present

Vector Institute
Faculty Member, Canada CIFAR AI Chair
March 2023 - Present
Faculty Affiliate
December 2020 - March 2023

Simons Institute for the Theory of Computing
Microsoft Research Fellow
August 2018 - May 2019

EDUCATION

Massachusetts Institute of Technology
Ph.D., September 2018
S.M., September 2014
Electrical Engineering and Computer Science

Cornell University
B.S., summa cum laude, May 2012
Computer Science, Electrical and Computer Engineering

SELECTED HONOURS AND AWARDS

Best Paper Award, International Conference on Machine Learning (ICML 2024) — July 2024
For "Position: Considerations for Differentially Private Learning with Large-Scale Public Pretraining"
Awarded to 10 best papers, out of 2610 accepted and 9473 submitted papers

2024 Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies — July 2024
For "The Discrete Gaussian for Differential Privacy"
Awarded annually to a single paper for an outstanding contribution to privacy enhancing technologies

2023 University of Waterloo Faculty of Math Golden Jubilee Research Excellence Award — July 2023
For outstanding research contributions, early-career recipient

Canada CIFAR AI Chair — March 2023
To recruit and retain world leading AI researchers to Canada

NSERC Discovery Accelerator Supplement — April 2020
To accelerate an established, superior research program

SELECTED PROFESSIONAL ACTIVITIES

Association for Computational Learning, Director — July 2024 - Present
36th International Conference on Algorithmic Learning Theory (ALT 2025), PC Co-Chair — June 2024 - Present
Association for Algorithmic Learning Theory, Steering Committee (ex-officio) — June 2024 - Present
Learning Theory Alliance, Executive Committee — November 2023 - Present
Transactions on Machine Learning Research (TMLR), Editor in Chief — July 2023 - Present
52nd Annual ACM Symposium on Theory of Computing (STOC 2020), General Chair — March 2020 - June 2020

PUBLICATIONS

Metrics: 4180 citations, h-index **33** (according to Google Scholar, July 25, 2024)
Primary publication venues: NeurIPS, ICML, COLT, STOC, FOCS, SODA
In computer science, conference proceedings are the primary venue for publishing complete research works.

Most authorships are in alphabetical order. Papers with contribution-order authorship are indicated, and equal contributions are marked with * or ˆ. Generally, these will put the students as first-author, with equal contribution amongst the senior authors. Ⓡ is used for randomized author order.

Differentially Private Post-Processing for Fair Regression
Ruicheng Xian, Qiaobo Li, Gautam Kamath, Han Zhao (Contribution order)
Proceedings of the 41st International Conference on Machine Learning (ICML 2024)

Disguised Copyright Infringement of Latent Diffusion Models
Yiwei Lu*, Matthew Y.R. Yang*, Zuoqiu Liu*, Gautam Kamathˆ, Yaoliang Yuˆ (Contribution order)
Proceedings of the 41st International Conference on Machine Learning (ICML 2024)

Position: Considerations for Differentially Private Learning with Large-Scale Public Pretraining
Florian Tramèr*, Gautam Kamath*, Nicholas Carlini* (Reverse alphabetical order)
Proceedings of the 41st International Conference on Machine Learning (ICML 2024)
ICML 2024 Best Paper
Oral Presentation

Indiscriminate Data Poisoning Attacks on Pre-trained Feature Extractors
Yiwei Lu, Matthew Y.R. Yang, Gautam Kamath*, Yaoliang Yu* (Contribution order)
Proceedings of the 2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2024)

Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment
Rachel Cummings, Damien Desfontaines, David Evans, Roxana Geambasu, Yangsibo Huang, Matthew Jagielski, Peter Kairouz, Gautam Kamath, Sewoong Oh, Olga Ohrimenko, Nicolas Papernot, Ryan Rogers, Milan Shen, Shuang Song, Weijie Su, Andreas Terzis, Abhradeep Thakurta, Sergei Vassilvitskii, Yu-Xiang Wang, Li Xiong, Sergey Yekhanin, Da Yu, Huanyu Zhang, Wanrong Zhang
Harvard Data Science Review, 6(1), 2024

Unbiased Statistical Estimation and Valid Confidence Intervals Under Differential Privacy
Christian Covington, Xi He*, James Honaker*, Gautam Kamath* (Contribution order)
Statistica Sinica special issue on Data Privacy, to appear, 2024

Not All Learnable Distribution Classes are Privately Learnable
Mark Bun, Gautam Kamath, Argyris Mouzakis, Vikrant Singhal
Proceedings of the 35th International Conference on Algorithmic Learning Theory (ALT 2024)

Private Distribution Learning with Public Data: The View from Sample Compression
Shai Ben-David, Alex Bie, Clément L. Canonne, Gautam Kamath, Vikrant Singhal
Advances in Neural Information Processing Systems 36 (NeurIPS 2023)
Spotlight Presentation

Distribution Learnability and Robustness
Shai Ben-David, Alex Bie, Gautam Kamath, Tosca Lechner
Advances in Neural Information Processing Systems 36 (NeurIPS 2023)

Hidden Poison: Machine Unlearning Enables Camouflaged Poisoning Attacks
Jimmy Z. Di, Jack Douglas, Jayadev Acharya*, Gautam Kamath*, Ayush Sekhari* (Contribution order)
Advances in Neural Information Processing Systems 36 (NeurIPS 2023)

Private GANs, Revisited
Alex Bie, Gautam Kamath*, Guojun Zhang* (Contribution order)
Transactions on Machine Learning Research (TMLR), 2023
Survey Certification

Individual Privacy Accounting for Differentially Private Stochastic Gradient Descent
Da Yu, Gautam Kamath*, Janardhan Kulkarni*, Tie-Yan Liu*, Jian Yin*, Huishuai Zhang* (Contribution order)
Transactions on Machine Learning Research (TMLR), 2023

Exploring the Limits of Model-Targeted Indiscriminate Data Poisoning Attacks
Yiwei Lu, Gautam Kamath*, Yaoliang Yu*. (Contribution order)
Proceedings of the 40th International Conference on Machine Learning (ICML 2023)

Robustness Implies Privacy in Statistical Estimation
Samuel B. Hopkins, Gautam Kamath, Mahbod Majid, Shyam Narayanan
Proceedings of the 55th ACM Symposium on Theory of Computing (STOC 2023)

Indiscriminate Data Poisoning Attacks on Neural Networks
Yiwei Lu, Gautam Kamath*, Yaoliang Yu* (Contribution order)
Transactions on Machine Learning Research (TMLR), 2022

New Lower Bounds for Private Estimation and a Generalized Fingerprinting Lemma
Gautam Kamath, Argyris Mouzakis, Vikrant Singhal
Advances in Neural Information Processing Systems 35 (NeurIPS 2022)

Private Estimation with Public Data
Alex Bie, Gautam Kamath, Vikrant Singhal
Advances in Neural Information Processing Systems 35 (NeurIPS 2022)

Improved Rates for Differentially Private Stochastic Convex Optimization with Heavy-Tailed Data
Gautam Kamath, Xingtu Liu, Huanyu Zhang
Proceedings of the 39th International Conference on Machine Learning (ICML 2022)
Long Talk

Robust Estimation for Random Graphs
Jayadev Acharya, Ayush Jain, Gautam Kamath, Ananda Theertha Suresh, Huanyu Zhang
Proceedings of the 35th Annual Conference on Learning Theory (COLT 2022)

A Private and Computationally-Efficient Estimator for Unbounded Gaussians
Gautam Kamath, Argyris Mouzakis, Vikrant Singhal, Thomas Steinke, Jonathan Ullman
Proceedings of the 35th Annual Conference on Learning Theory (COLT 2022)

The Price of Tolerance in Distribution Testing
Clément L. Canonne, Ayush Jain, Gautam Kamath, Jerry Li
Proceedings of the 35th Annual Conference on Learning Theory (COLT 2022)

Calibration with Privacy in Peer Review
Wenxin Ding, Gautam Kamath ® Weina Wang ® Nihar B. Shah (Contribution order, with randomization)
Proceedings of the 2022 IEEE International Symposium on Information Theory (ISIT 2022)

Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism
Samuel B. Hopkins, Gautam Kamath, Mahbod Majid
Proceedings of the 54th ACM Symposium on Theory of Computing (STOC 2022)
Presented at the 3rd Symposium on Foundations of Responsible Computing (FORC 2022, non-archival track)

Differentially Private Fine-tuning of Language Models
Da Yu, Saurabh Naik, Arturs Backurs*, Sivakanth Gopi*, Huseyin A. Inan*, Gautam Kamath*, Janardhan Kulkarni*, Yin Tat Lee*, Andre Manoel*, Lukas Wutschitz*, Sergey Yekhanin*, Huishuai Zhang* (Contribution order)
Journal of Privacy and Confidentiality, to appear, 2024
Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)

The Role of Adaptive Optimizers for Honest Private Hyperparameter Selection
Shubhankar Mohapatra*, Sajin Sasy*, Xi He^, Gautam Kamath^, Om Thakkar^ (Contribution order)
Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022)
Oral Presentation

Robustness Meets Algorithms
Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, Alistair Stewart
Communications of the ACM, 64(5), 2021
Invited Research Highlight

Remember What You Want to Forget: Algorithms for Machine Unlearning
Ayush Sekhari, Jayadev Acharya*, Gautam Kamath*, Ananda Theertha Suresh* (Contribution order)
Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

Enabling Fast Differentially Private SGD via Just-in-Time Compilation and Vectorization
Pranav Subramani*, Nicholas Vadivelu*, Gautam Kamath (Contribution order)
Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

PAPRIKA: Private Online False Discovery Rate Control
Wanrong Zhang, Gautam Kamath*, Rachel Cummings* (Contribution order)
Proceedings of the 38th International Conference on Machine Learning (ICML 2021)
Presented at the 2nd Symposium on Foundations of Responsible Computing (FORC 2021, non-archival track)

On the Sample Complexity of Privately Learning Unbounded High-Dimensional Gaussians
Ishaq Aden-Ali, Hassan Ashtiani, Gautam Kamath
Proceedings of the 32nd International Conference on Algorithmic Learning Theory (ALT 2021)

Random Restrictions of High-Dimensional Distributions and Uniformity Testing with Subcube Conditioning
Clément L. Canonne, Xi Chen, Gautam Kamath, Amit Levi, Erik Waingarten
Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)

CoinPress: Practical Private Mean and Covariance Estimation
Sourav Biswas, Yihe Dong, Gautam Kamath, Jonathan Ullman
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)

The Discrete Gaussian for Differential Privacy
Clément L. Canonne, Gautam Kamath, Thomas Steinke
Journal of Privacy and Confidentiality, 12(1), 2022
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)
2024 Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies Deployed in the 2020 US Census

Private Identity Testing for High-Dimensional Distributions
Clément L. Canonne, Gautam Kamath, Audra McMillan, Jonathan Ullman, Lydia Zakynthinou
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)
Spotlight Presentation

Privately Learning Markov Random Fields

Huanyu Zhang, Gautam Kamath*, Janardhan Kulkarni*, Zhiwei Steven Wu* (Contribution order)
Proceedings of the 37th International Conference on Machine Learning (ICML 2020)

Private Mean Estimation of Heavy-Tailed Distributions
Gautam Kamath, Vikrant Singhal, Jonathan Ullman
Proceedings of the 33rd Annual Conference on Learning Theory (COLT 2020)

Locally Private Hypothesis Selection
Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, Huanyu Zhang
Proceedings of the 33rd Annual Conference on Learning Theory (COLT 2020)

Differentially Private Algorithms for Learning Mixtures of Separated Gaussians
Gautam Kamath, Or Sheffet, Vikrant Singhal, Jonathan Ullman
Advances in Neural Information Processing Systems 32 (NeurIPS 2019)

Private Hypothesis Selection
Mark Bun, Gautam Kamath, Thomas Steinke, Zhiwei Steven Wu
IEEE Transactions on Information Theory, 67(3), 2021
Advances in Neural Information Processing Systems 32 (NeurIPS 2019)

Sever: A Robust Meta-Algorithm for Stochastic Optimization
Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Jacob Steinhardt, Alistair Stewart
Proceedings of the 36th International Conference on Machine Learning (ICML 2019)

Privately Learning High-Dimensional Distributions
Gautam Kamath, Jerry Li, Vikrant Singhal, Jonathan Ullman
Proceedings of the 32nd Annual Conference on Learning Theory (COLT 2019)

The Structure of Optimal Private Tests for Simple Hypotheses
Clément Canonne, Gautam Kamath, Audra McMillan, Adam Smith, Jonathan Ullman
Proceedings of the 51st ACM Symposium on Theory of Computing (STOC 2019)

Anaconda: A Non-Adaptive Conditional Sampling Algorithm for Distribution Testing
Gautam Kamath, Christos Tzamos
Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)

INSPECTRE: Privately Estimating the Unseen
Jayadev Acharya, Gautam Kamath, Ziteng Sun, Huanyu Zhang
Journal of Privacy and Confidentiality, 10(2), 2020
Proceedings of the 35th International Conference on Machine Learning (ICML 2018)

Actively Avoiding Nonsense in Generative Models
Steve Hanneke, Adam Kalai, Gautam Kamath, Christos Tzamos
Proceedings of the 31st Annual Conference on Learning Theory (COLT 2018)

Which Distribution Distances are Sublinearly Testable?
Constantinos Daskalakis, Gautam Kamath, John Wright
Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)

Testing Ising Models
Constantinos Daskalakis, Nishanth Dikkala, Gautam Kamath
IEEE Transactions on Information Theory, 65(11), 2019
Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)

Robustly Learning a Gaussian: Getting Optimal Error, Efficiently
Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, Alistair Stewart

Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)

Concentration of Multilinear Functions of the Ising Model with Applications to Network Data
Constantinos Daskalakis, Nishanth Dikkala, Gautam Kamath
Advances in Neural Information Processing Systems 30 (NIPS 2017)

Being Robust (in High Dimensions) Can Be Practical
Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, Alistair Stewart
Proceedings of the 34th International Conference on Machine Learning (ICML 2017)

Priv'IT: Private and Sample Efficient Identity Testing
Bryan Cai, Constantinos Daskalakis, Gautam Kamath
Proceedings of the 34th International Conference on Machine Learning (ICML 2017)

Robust Estimators in High Dimensions without the Computational Intractability
Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, Alistair Stewart
Invited to SIAM Journal on Computing Special Issue for FOCS 2016, 48(2), 2019 (SICOMP)
Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2016)
Invited to Highlights of Algorithms 2017 (HALG 2017)
Invited to Communications of the ACM, Research Highlights (CACM)

A Size-Free CLT for Poisson Multinomials and its Applications
Constantinos Daskalakis, Anindya De, Gautam Kamath, Christos Tzamos
Proceedings of the 48th ACM Symposium on Theory of Computing (STOC 2016)

Optimal Testing for Properties of Distributions
Jayadev Acharya, Constantinos Daskalakis, Gautam Kamath
Advances in Neural Information Processing Systems 28 (NIPS 2015)
Spotlight Presentation

On the Structure, Covering, and Learning of Poisson Multinomial Distributions
Constantinos Daskalakis, Gautam Kamath, Christos Tzamos
Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2015)

A Chasm Between Identity and Equivalence Testing with Conditional Queries
Jayadev Acharya, Clément Canonne, Gautam Kamath
Theory of Computing, 14(19), 2018
Proceedings of the 19th International Workshop on Randomization and Computation (RANDOM 2015)

Adaptive Estimation in Weighted Group Testing
Jayadev Acharya, Clément Canonne, Gautam Kamath
Proceedings of the 2015 IEEE International Symposium on Information Theory (ISIT 2015)

Faster and Sample Near-Optimal Algorithms for Proper Learning Mixtures of Gaussians
Constantinos Daskalakis, Gautam Kamath
Proceedings of the 27th Annual Conference on Learning Theory (COLT 2014)

An Analysis of One-Dimensional Schelling Segregation
Christina Brandt, Nicole Immorlica, Gautam Kamath, Robert Kleinberg
Proceedings of the 44th ACM Symposium on Theory of Computing (STOC 2012)

PREPRINTS AND OTHER WRITINGS

Machine Unlearning Fails to Remove Data Poisoning Attacks
Martin Pawelczyk*, Jimmy Z. Di*, Yiwei Lu, Gautam Kamathˆ, Ayush Sekhariˆ, Seth Neelˆ (Contribution order)

Avoiding Pitfalls for Privacy Accounting of Subsampled Mechanisms under Composition

Christian Janos Lebeda*, Matthew Regehr*, Gautam Kamathˆ, Thomas Steinkeˆ (Contribution order)
Manuscript

Private Mean Estimation with Person-Level Differential Privacy
Sushant Agarwal, Gautam Kamath, Mahbod Majid, Argyris Mouzakis, Rose Silver, Jonathan Ullman
Manuscript

Choosing Public Datasets for Private Machine Learning via Gradient Subspace Distance
Xin Gu, Gautam Kamath*, Zhiwei Steven Wu* (Contribution order)
Manuscript

A Bias-Variance-Privacy Trilemma for Statistical Estimation
Gautam Kamath, Argyris Mouzakis, Matthew Regehr, Vikrant Singhal, Thomas Steinke, Jonathan Ullman
Manuscript (Revise and Resumit at Journal of the American Statistical Association)

Report of the 1st Workshop on Generative AI and Law
A. Feder Cooper, Katherine Lee, James Grimmelmann, Daphne Ippolito, Christopher Callison-Burch, Christopher A. Choquette-Choo, Niloofar Mireshghallah, Miles Brundage, David Mimno, Madiha Zahrah Choksi, Jack M. Balkin, Nicholas Carlini, Christopher De Sa, Jonathan Frankle, Deep Ganguli, Bryant Gipson, Andres Guadamuz, Swee Leng Harris, Abigail Z. Jacobs, Elizabeth Joh, Gautam Kamath, Mark Lemley, Cass Matthews, Christine McLeavey, Corynne McSherry, Milad Nasr, Paul Ohm, Adam Roberts, Tom Rubin, Pamela Samuelson, Ludwig Schubert, Kristen Vaccaro, Luis Villa, Felix Wu, Elana Zeide
Manuscript

A Primer on Private Statistics
Gautam Kamath, Jonathan Ullman
Manuscript

Bounds on the Expectation of the Maximum of Samples from a Gaussian
Gautam Kamath
Manuscript

TALKS

8th Annual Toronto Machine Learning Summit 2024 — July 2024
Privacy Risks and Protections in Machine Learning Systems (Opening Keynote)

2024 Canadian Computing Olympiad — June 2024
Selection, with Rounds or Adversarial Comparisons

Yale CSPC 471/571: Trustworthy Deep Learning — April 2024
Differentially Private Machine Learning (Guest Lecture)

New York University Courant CS Colloquium — March 2024
Principled Approaches for Trustworthy Algorithms, Statistics, and Machine Learning

Cornell University ORIE Colloquium — March 2024
Principled Approaches for Trustworthy Algorithms, Statistics, and Machine Learning

Yale University CS Talk — February 2024
Principled Approaches for Trustworthy Algorithms, Statistics, and Machine Learning

UC San Diego HDSI Seminar — February 2024
Principled Approaches for Trustworthy Algorithms, Statistics, and Machine Learning

Vector Institute Distinguished Talk Series — November 2023
The Promise and Pitfalls of Public Data in Private ML

Vector Institute Machine Learning Theory Workshop — November 2023
Statistical Estimation with Privacy Constraints

University of Guelph CARE-AI Seminar — November 2023
Protecting Individual Privacy in Machine Learning

Columbia University Robust Statistics and Privacy Workshop — October 2023
Robust Estimators for Private Estimation

6th Eastern Great Lakes (EaGL) Theory of Computation Workshop — October 2023
Differentially Private Mean Estimation

JSM 2023 Invited Session on Robust Statistics and Differential Privacy August 2023
Private Estimators from Robust Statistics

BIRS Workshop on Contextual Integrity for Differential Privacy July 2023
Considerations for Differentially Private Learning with Large-Scale Public Pretraining

ICML 2023 Workshop on Generative AI + Law July 2023
What does Differential Privacy have to do with Copyright?

ICML 2023 Black in AI Social July 2023
AI and Society (Panel member)

Vector Machine Learning Security and Privacy Workshop July 2023
The Promise and Pitfalls of Public Data in Private ML

IISC Bangalore Joint Telematics Group Summer School on Information Theory June 2023
Introduction to Differential Privacy (Invited Tutorial, 8 hours of lectures)

Upper Bound Workshop Bridge the Gap: Differential Privacy and Statistical Analysis May 2023
Recent Connections between Differential Privacy and Robustness

Cybersecurity, Privacy, and AI in Health Data: Advancements and Challenges May 2023
An Introduction to Differential Privacy for Analysis of Sensitive Data

Vector Faculty Research Meeting April 2023
Public Data for Private Machine Learning

Waterloo CPI Conference The Weaponization of Disinformation in Canada April 2023
Problematic Disinformation (Panel member)

Waterloo Computer Science Club Prof Talk March 2023
Differential Privacy in Machine Learning

KAUST Rising Stars in AI Symposium 2023 February 2023
Differentially Private Fine-tuning of Language Models

IEEE Conference on Secure and Trustworthy Machine Learning February 2023
An Introduction to Differential Privacy (Invited Tutorial)

Georgia Tech AI4OPT Seminar February 2023
Efficient Private Mean Estimation

IMS International Conference on Statistics and Data Science December 2022
The Role of (Statistical) Bias in Private Estimation

Cornell Computer Science Theory Seminar December 2022
Efficient Private Mean Estimation

UC Berkeley BLISS Seminar November 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

US Census ML/AI Discussion Group November 2022
CoinPress: Practical Private Point Estimation and Confidence Intervals

LinkedIn Data Tech Talk Series October 2022
Differentially Private Fine-tuning of Language Models

Canadian AI Federated Learning Workshop October 2022
Differentially Private Fine-tuning of Language Models

Rutgers Business School MSIS Department Seminar Series October 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

Columbia Statistics Seminar Series October 2022
CoinPress: Practical Private Point Estimation and Confidence Intervals

Huawei Strategy and Technology Conference AI 2022 September 2022
Differentially Private Fine-tuning of Language Models

University of British Columbia CAIDA Seminar Series August 2022
Differentially Private Fine-tuning of Language Models

PIMS Mathematics of Ethical Decision-making Systems Seminar August 2022
Statistical Estimation with Differential Privacy

Meta Lunch and Learn August 2022
Statistical Estimation with Differential Privacy

Fields Workshop on Differential Privacy and Statistical Data Analysis July 2022
Premonitions of Public Data for Private ML

International Conference on Robust Statistics July 2022
Robust Estimation for Random Graphs

Sino-EU Doctoral School for Logistics, Information, Management, and Service Science July 2022

Differentially Private Machine Learning (Invited Tutorial)

Google Privacy Seminar June 2022
Differentially Private Fine-tuning of Language Models

10th Iran Workshop on Communication and Information Theory May 2022
An Introduction to Differential Privacy (Invited Tutorial)

University of Washington Theory Seminar April 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

University of Waterloo Probability Seminar Series April 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

Apple Workshop on Privacy Preserving ML April 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

Apple Workshop on Privacy Preserving ML April 2022
Differentially Private Fine-tuning of Language Models

Simons Institute Data Privacy: Foundations and Applications Reunion March 2022
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

Huawei Noah's Ark Lab Federated Learning Group Seminar March 2022
Differentially Private Fine-tuning of Language Models

UMass Amherst Machine Learning & Friends Lunch February 2022
Differentially Private Fine-tuning of Language Models

BIRS Workshop on Mathematical Statistics and Learning November 2021
Efficient Mean Estimation with Pure Differential Privacy via a Sum-of-Squares Exponential Mechanism

Google Workshop on Federated Learning and Analytics November 2021
Differentially Private Fine-tuning of Language Models

ML Collective Deep Learning: Classics and Trends November 2021
Differentially Private Fine-tuning of Language Models

IDEAL Mini-workshop on New Direction on Robustness in ML November 2021
Statistical Estimation with Differential Privacy

London Symposium on Information Theory May 2021
Private Hypothesis Selection

TrustML Seminar April 2021
CoinPress: Practical Private Estimation

Boston-Area DP Seminar April 2021
Hypothesis Selection with Privacy

Virtual Conference on Robustness and Privacy March 2021
Differentially Private Mean and Covariance Estimation

Google Privacy and Machine Learning Seminar March 2021
CoinPress: Practical Private Mean and Covariance Estimation

McGill Statistics Seminar February 2021
CoinPress: Practical Private Point Estimation and Confidence Intervals

University of Waterloo ML + Logic Seminar January 2021
Robustness in Unsupervised and Supervised Machine Learning

Simons Institute Reading Group November 2020
Differentially Private Statistical Estimation

University of Toronto Theory Seminar October 2020
Hypothesis Selection with Privacy Constraints

University of Pennsylvania Wharton Statistics Seminar September 2020
CoinPress: Practical Private Point Estimation and Confidence Intervals

Northwestern University IDEAL Seminar August 2020
Theory and Practice for Private Statistical Estimation - Gaussians and Beyond

Harvard Privacy Tools Group Meeting August 2020
CoinPress: Practical Private Mean and Covariance Estimation

Joint Statistical Meetings August 2020
Differentially Private Mean and Covariance Estimation

Carnegie Mellon University Theory Lunch August 2020
Hypothesis Selection with Privacy Constraints

University of Waterloo Algorithms and Complexity Seminar May 2020
Robustness in Unsupervised and Supervised Machine Learning

Google Mountain View Algorithms Group Meeting February 2020
Privately Learning High-Dimensional Distributions
Information Theory and Applications Workshop February 2020
Private Hypothesis Selection
Stanford University Management Science & Engineering Seminar February 2020
Principled Tools for Modern Statistical Challenges
National Technical University of Athens Corelab Seminar July 2019
Privately Learning High-Dimensional Distributions
Workshop on Algorithms for Learning and Economics July 2019
Efficient Multivariate Robust Statistics
Google Seattle Cerebra Journal Club April 2019
Estimating a Gaussian: Robustly or Privately
Simons Institute Workshop on Data Privacy: From Foundations to Applications March 2019
Privately Learning High-Dimensional Distributions
Berkeley Theory Lunch March 2019
Privately Learning High-Dimensional Distributions
MIT Algorithms and Complexity Seminar February 2019
Privately Learning High-Dimensional Distributions
Berkeley BLISS Seminar February 2019
Privately Learning High-Dimensional Distributions
Information Theory and Applications Workshop February 2019
Privately Learning High-Dimensional Distributions
Caltech Mathematics of Information Seminar January 2019
Privately Learning High-Dimensional Distributions
Simons Institute Data Privacy: Foundations and Applications Boot Camp January 2019
Statistical Inference and Privacy
Symposium on Discrete Algorithms January 2019
Anaconda: A Non-Adaptive Conditional Sampling Algorithm for Distribution Testing
Microsoft Research Machine Learning and Optimization Lunch November 2018
Privately Learning High-Dimensional Distributions
Simons Institute Workshop on Robust and High-Dimensional Statistics November 2018
Realizing Robustness
TTIC Workshop on Computational Efficiency and High-Dimensional Robust Statistics August 2018
Beyond Theory: Realizing Robustness
BIRS Mathematical Foundations of Data Privacy Workshop May 2018
Differentially Private Hypothesis Testing and Property Estimation
CRM Modern Challenges of Learning Theory Workshop April 2018
Robustness in Unsupervised and Supervised Machine Learning
MIT LIDS and Stats Tea April 2018
INSPECTRE: Privately Estimating the Unseen
Conference on Information Sciences and Systems March 2018
Hypothesis Testing with Alternative Distances
Boston University Computer Science Seminar February 2018
Principled Tools for Modern Statistical Data Science
McGill University Computer Science Seminar February 2018
Principled Tools for Modern Statistical Data Science
University of Waterloo Computer Science Seminar February 2018
Principled Tools for Modern Statistical Data Science
Symposium on Discrete Algorithms January 2018
Which Distribution Distances are Sublinearly Testable?
University of Pennsylvania Theory Seminar December 2017
Statistical Hypothesis Testing in the Modern Age
Boston University Algorithms and Theory Seminar November 2017
Statistical Hypothesis Testing in the Modern Age
University of Massachusetts Amherst Theory Seminar October 2017
Statistical Hypothesis Testing in the Modern Age
McMaster Seminar in Computing and Software October 2017

Statistical Hypothesis Testing in the Modern Age

FOCS Workshop on Frontiers in Distribution Testing October 2017
Testing with Alternative Distances

Cornell Theory Lunch September 2017
Robust Estimators in High Dimensions without the Computational Intractability

Cornell Theory Seminar September 2017
Statistical Hypothesis Testing in the Modern Age

International Conference on Machine Learning August 2017
Priv'IT: Private and Sample Efficient Identity Testing

ICML Workshop on Private and Secure Machine Learning August 2017
Priv'IT: Private and Sample Efficient Identity Testing

Northeastern Theory Seminar March 2017
Some Frontiers in Distribution Testing

University of Pennsylvania Theory Lunch September 2016
Optimal Testing for Properties of Distributions

China Theory Week August 2016
Robust Estimators in High Dimensions without the Computational Intractability

Symposium on Theory of Computing June 2016
A Size-Free CLT for Poisson Multinomials and its Applications

MIT Signals, Information, and Algorithms Laboratory Group Meeting March 2016
Optimal Testing for Properties of Distributions

University of Massachusetts Boston Computer Science Seminar February 2016
Optimal Testing for Properties of Distributions

Berkeley Theory Lunch September 2015
Optimal Testing for Properties of Distributions

Conference on Learning Theory June 2014
Faster and Sample Near-Optimal Algorithms for Proper Learning Mixtures of Gaussians

Interdisciplinary Workshop on Information and Decision in Social Networks November 2012
An Analysis of One-Dimensional Schelling Segregation

Symposium on Theory of Computing May 2012
An Analysis of One-Dimensional Schelling Segregation
Winner of Best Student Presentation Award

GRADUATE STUDENTS

Jimmy Z. Di (Fall 2023 - Present)
MMath, Computer Science, University of Waterloo
Awarded **Vector Scholarship in Artificial Intelligence**

Matthew Regehr (Fall 2023 - Present)
PhD, Computer Science, University of Waterloo
Awarded **Ontario Graduate Scholarship**
Awarded **NSERC Canada Graduate Scholarship - Doctorate**

Sabrina Mokhtari (Fall 2022 - Present)
MMath, Computer Science, University of Waterloo
Awarded **Vector Scholarship in Artificial Intelligence**
Awarded **NSERC Canada Graduate Scholarship - Master's**

Sara Kodeiri (Fall 2022 - Present)
MMath, Computer Science, University of Waterloo

Argyris Mouzakis (Fall 2020 - Present)
PhD, Computer Science, University of Waterloo
Awarded **Onassis Foundation Scholarship**

Matthew Regehr (Fall 2021 - Summer 2023, co-advised with Shai Ben-David)
MMath, Computer Science, University of Waterloo
Thesis: A Bias-Variance-Privacy Trilemma for Statistical Estimation
Next Position: PhD Student in Computer Science at University of Waterloo
Awarded **Vector Scholarship in Artificial Intelligence**
Awarded **NSERC Canada Graduate Scholarship - Master's**

Yaxian Alex Bie (Fall 2021 - Summer 2023, co-advised with Shai Ben-David)
MMath, Computer Science, University of Waterloo
Thesis: Private Distribution Learning with Public Data
Next Position: Research Engineer at Huawei
Current Position: Research Scientist at Google Research
Awarded **Vector Scholarship in Artificial Intelligence**
Awarded **Ontario Graduate Scholarship**

Mahbod Majid (Fall 2020 - Fall 2022)
MMath, Computer Science, University of Waterloo
Thesis: Efficient and Differentially Private Statistical Estimation via a Sum-of-Squares Exponential Mechanism
Next Position: PhD Student in Machine Learning at Carnegie Mellon University
Awarded **Waterloo CPI Cybersecurity and Privacy Excellence Graduate Scholarship**
Awarded **Faculty of Mathematics Graduate Research Excellence Award**
Awarded **University Finalist for the Governor General's Gold Medal**
Awarded **University Finalist for the Alumni Gold Medal**

Christian Covington (Fall 2020 - Summer 2022, co-advised with Xi He)
MMath, Computer Science, University of Waterloo
Thesis: Unbiased Statistical Estimation and Valid Confidence Intervals Under Differential Privacy
Next Position: PhD Student in Biostatistics at Harvard University

POSTDOCS

Vikrant Singhal (Fall 2021 - Fall 2023)
Next Position: Research Associate at OpenDP

UNDERGRADUATE RESEARCH ADVISING

Chris Trevisan (Spring 2022 - Present)
Published "Sorting and Selection in Rounds with Adversarial Comparisons" in SODA 2024

Matthew Y.R. Yang (Fall 2022 - Summer 2024)
Awarded **CRA Outstanding Undergraduate Researcher, Finalist**
Published "Indiscriminate Data Poisoning Attacks on Pre-trained Feature Extractors" in SaTML 2024
Published "Disguised Copyright Infringement of Latent Diffusion Models" in ICML 2024
Next position: MS Student in Machine Learning at Carnegie Mellon University

Jimmy Z. Di (Fall 2021 - Summer 2023)
Published "Hidden Poison: Machine Unlearning Enables Camouflaged Poisoning Attacks" in NeurIPS 2023
Next position: MMath Student in Computer Science at University of Waterloo

Ruiyun Chao (Fall 2022)

Olivia Ma (Fall 2022)
Next position: Master of Science in Computing (AI & ML) at Imperial College London

Valerie Liu (Fall 2022)
Next position: Master's student in Computing Science at University of Alberta

Jack Douglas (Summer 2022)
Published "Hidden Poison: Machine Unlearning Enables Camouflaged Poisoning Attacks" in NeurIPS 2023

Andrew Guo (Summer 2022)

Chirag Jindal (Summer 2022)

Landy Xu (Spring 2021 - Fall 2021)
Next position: Master of Science in Applied Computing Student at University of Toronto

Xingtu Liu (Fall 2020 - Summer 2021)
Published "Improved Rates for Differentially Private Stochastic Convex Optimization with Heavy-Tailed Data" in ICML 2022
Next position: Master of Science in Computer Science at Simon Fraser University

Nicholas Vadivelu (Fall 2020 - Summer 2021)
Awarded **CRA Outstanding Undergraduate Researcher, Runner-Up**
Awarded **Jessie W.H. Zou Memorial Award**
Published "Enabling Fast Differentially Private SGD via Just-in-Time Compilation and Vectorization" in NeurIPS 2021
Next position: Quantitative Research and Data Scientist at Citadel

Pranav Subramani (Fall 2019 - Summer 2021)
Published "Enabling Fast Differentially Private SGD via Just-in-Time Compilation and Vectorization" in NeurIPS 2021
Next position: Quantitative Researcher at Cubist Systematic Strategies

Sourav Biswas (Fall 2019 - Summer 2021)
Awarded **CRA Outstanding Undergraduate Researcher, Honorable Mention**
Published "CoinPress: Practical Private Mean and Covariance Estimation" in NeurIPS 2020
Next position: PhD Student in Computer Science at University of Toronto

OTHER ADVISING

Ishaq Aden-Ali (Summer 2020 - Summer 2021)
Published "On the Sample Complexity of Privately Learning Unbounded High-Dimensional Gaussians" in ALT 2020
Next position: PhD Student in Computer Science at UC Berkeley

Sushant Agarwal (Spring 2022 - Summer 2022)
Next position: PhD Student in Computer Science at Northeastern University

SELECTED PRESS COVERAGE AND QUOTES

Artists Are Slipping Anti-AI 'Poison' into Their Art. Here's How It Works
Scientific American, March 2024

The modern Luddites' fight for the human premium in AI
Canvas8 Report, February 2024

1 big thing: It can be hard for AI to forget
Axios Science, January 2024

New tools help artists fight AI by directly disrupting the systems
NPR All Things Considered, November 2023

This new data poisoning tool lets artists fight back against generative AI
MIT Technology Review, October 2023

Police in Essex County have started using licence plate scanners. Here's how they work
CBC, April 2023

Who Is Working to End the Threat of AI-Generated Deepfakes, and Why Is It So Difficult?
Gizmodo, November 2022

Can AI Learn to Forget?
Communications of the ACM, April 2022

Now That Machines Can Learn, Can They Unlearn?
Wired Magazine, August 2021

Canadian educator gains following in China after posting online course to Chinese video sharing site Bilibili
Global Times, January 2021

Foreign Professor Becomes an Uper at Bilibili Wowing Chinese Audience. University of Waterloo Differential Privacy Class is Available Online (Translated from Chinese)
Heart of the Machine, January 2021

GRANTS

Canada CIFAR AI Chair
Sole PI
03/2023 - 03/2028

Ontario Research Fund: Research Infrastructure
Co-PI, with Xi He
11/2022 - 11/2024

Apple Unrestricted Gift
Sole PI
8/2022

Compute Canada Resources for Research Groups
Co-PI, with Xi He
4/2022 - 3/2023

Google Unrestricted Gift
Sole PI
2/2022

Canada Foundation for Innovation John R. Evans Leaders Fund
Co-PI, with Xi He
8/2021 - 1/2024

NSERC Discovery Grant
Sole PI
4/2020 - 3/2025

NSERC Discovery Grant - Accelerator Supplement
Sole PI
4/2020 - 3/2023

NSERC Discovery Grant - Launch Supplement
Sole PI
4/2020 - 3/2021

Compute Canada Resources for Research Groups
Co-PI, with Xi He
4/2021 - 3/2022

Compute Canada Resources for Research Groups
Co-PI, with Xi He
4/2020 - 3/2021

University of Waterloo Startup Grant
Sole PI
7/2019 - 6/2024

ADDITIONAL HONOURS AND AWARDS

Senior Member, IEEE	December 2023
Notable Reviewer, SaTML 2023	February 2023
Best Reviewer Award, CCS 2021	November 2021
Top Graduate Instructor for CS 761 in Fall 2019	March 2020
Top 5% Highest-Scoring Reviewer for ICML 2019	June 2019
Top 30% Highest-Scoring Reviewer for NeurIPS 2018	December 2018
MIT Akamai Presidential Graduate Fellowship	September 2012 - May 2013
Best Student Presentation Award, STOC 2012	May 2012
Cornell Computer Science Prize for Academic Excellence	May 2012
Eight time Dean's list at Cornell University	Fall 2008 - Spring 2012
Recognized by Cornell CS for outstanding work as TA for CS 3110 and CS 4820	Spring 2012
John G. Pertsch Jr. Prize for second highest GPA in ECE	Spring 2011
Recognized by Cornell CS for outstanding work as TA for CS 1114	Spring 2010
Canadian Open Mathematics Challenge Gold Medalist in Central Ontario Region	Spring 2007

TEACHING

Instructor	University of Waterloo	Fall 2019 - Present
CS 240: Data Structures and Data Management		(1 term)
CS 480: Introduction to Machine Learning		(4 terms, 8 sections)
CS 761: Randomized Algorithms		(1 term)
CS 860: Algorithms for Private Data Analysis		(2 terms)
Teaching Assistant	Massachusetts Institute of Technology	Spring 2015, 2017
6.853: Algorithmic Game Theory and Data Science		(1 semester)
6.856: Randomized Algorithms		(1 semester)
Teaching Assistant	Cornell University	Spring 2010 - Spring 2012
CS 1114: Intro to Computing with Matlab and Robotics		(2 semesters)
CS 2850: Networks		(1 semester)
CS 3110: Data Structures and Functional Programming		(5 semesters)
CS 4820: Introduction to Algorithms		(3 semesters)

PROFESSIONAL ACTIVITIES

Journal Editor-in-Chief: TMLR
Conference Program Committee Chair: ALT 2025
Conference General Chair: STOC 2020
Workshop Program Chair: TPDP 2021, TPDP 2022, UpML 2022
Conference Core Program Committee or Area Chair: SODA 2020, ICML 2020, ICALP 2020, RANDOM 2020, ALT 2021, ICLR 2021, FORC 2021, COLT 2021, CCS 2021, NeurIPS 2021, ESA 2021, ALT 2022, SODA 2022, ICLR 2022, AAAI 2022, COLT 2022, NeurIPS 2022, SaTML 2023, FOCS 2023, ICLR 2023, ALT 2023, USENIX Security 2023, COLT 2023, FAccT 2023, ICML 2023, NeurIPS 2023, COLT 2024, NeurIPS 2024
Machine Learning Conference Program Committee Member (i.e., reviewer): NIPS 2016, ICML 2018, NeurIPS

2018, AISTATS 2019, ICML 2019, NeurIPS 2019, AAAI 2020, AISTATS 2020, FAccT 2021, UAI 2022, CANAI 2024
Workshop Program Committee Member: TPDP 2019, PriML 2019, TPDP 2020, PPML 2020, PriML 2021, ICBINB 2021, ICBINB 2022, ICBINB 2023, Regulatable ML 2023, PPAI 2024, TPDP 2024
Journal Guest Editor: TALG Special Issue for SODA 2020
Journal Action Editor: Transactions on Machine Learning Research (March 2022 - July 2023)
Other: ICML 2021 Workshop Reviewer, NeurIPS 2021 Social Chair, NeurIPS 2022 Ethics Reviewer, ICML 2023 Workshop Reviewer, SaTML 2024 Social Media Chair, NeurIPS 2023 Ethics Reviewer, TMLR 2022 Outstanding Paper Committee, ICML 2024 Workshop Reviewer, 2024 Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies Committee, SaTML 2025 Social Media Chair

Conference external reviewer: AAAI, AISTATS, ALT, COLT, FAccT, FOCS, ICALP, ICML, ICML Workshops, ISAAC, ISIT, ITCS, NeurIPS, RANDOM, SODA, STACS, STOC
Journal reviewer: Algorithmica, Annals of Statistics, Behavior Research Methods, Foundations and Trends in Theoretical Computer Science, Foundations of Data Science, Journal of Machine Learning Research, Journal of Privacy and Confidentiality, Statistica Sinica, Theory of Computing Systems
Grant reviewer: Natural Sciences and Engineering Research Council, Blavatnik Interdisciplinary Cyber Research Centre, National Science Foundation, Singapore Ministry of Education

Organizer of Vector Institute Workshop "Vector Machine Learning Security and Privacy Workshop" (July 2024)
Co-organizer of TTIC Workshop "New Frontiers in Robust Statistics" (June 2024)
Organizer of Vector Institute Workshop "Vector Machine Learning Security and Privacy Workshop" (July 2023)
Lead organizer of Fields Institute Workshop "Workshop on Differential Privacy and Statistical Data Analysis" (June 2022)
Co-organizer of ICML 2022 Workshop "Updatable Machine Learning" (July 2022)
Co-organizer of ICML 2022 Workshop "Theory and Practice of Differential Privacy" (July 2022)
Co-organizer of ICML 2021 Workshop "Theory and Practice of Differential Privacy" (July 2021)
Co-organizer of ICLR 2021 Workshop "Distributed and Private Machine Learning" (May 2021)
Co-organizer of NeurIPS 2020 Social "Data Privacy: Academia, Industry, Policy, and Society" (December 2020)
Co-organizer of FOCS 2019 Workshop "A TCS Quiver" (November 2019)
Co-organizer of FOCS 2017 Workshop "Frontiers in Distribution Testing" (October 2017)
Co-organizer of FOCS 2016 Workshop "Orthogonal Polynomials and Applications" (October 2016)
Organizer of the Second Annual Sublinear Algorithms and Big Data Day (April 2015)
Cofounder and organizer of MIT Theory Lunch (Fall 2012 - Summer 2013)
Advisor for Danny Lewin MIT Theory Student Retreat (Fall 2014, 2016, 2017)
Organizer of Second Annual Danny Lewin MIT Theory Student Retreat (October 2013)

Executive Committee of Learning Theory Alliance (November 2023 - Present)
Steering Committee of Theory and Practice of Differential Privacy (October 2023 - Present)
Founder and co-organizer of DifferentialPrivacy.org (July 2020 - Present)
Editor of ALT Highlights (April 2021 - July 2021)
Maintainer of CS Theory Blog Aggregator (January 2019 - Present)
Co-organizer for the TCS+ online seminar series in Theoretical Computer Science (August 2014 - Present)
Editor of Property Testing Review (March 2016 - June 2020)
Editor of MIT Theory of Computation Student Blog (November 2013 - October 2016)

Member of University of Waterloo CS Women in Computer Science Committee (August 2023 - Present)
Member of University of Waterloo CS Awards Committee (August 2023 - Present)
Member of University of Waterloo CS School Advisory Committee on Appointments (August 2022 - August 2023)
Member of University of Waterloo CS Equity, Diversity, and Inclusion Committee (August 2020 - August 2022)
Member of University of Waterloo CS Graduate Recruitment Committee (August 2019 - August 2020)

Reviewer for Vector Scholarship in AI (Spring 2023)
Organizer for CIFAR Deep Learning + Reinforcement Learning Summer School (Fall 2023 - Summer 2024)
Reviewer for Vector Institute Visiting Researcher Program (Winter 2024, Spring 2024)

THESIS COMMITTEES

PhD Thesis: Amit Levi (Waterloo), Vikrant Singhal (Northeastern University), Guojun Zhang (Waterloo), Jimit Majmudar (Waterloo), Tosca Lechner (Waterloo), Nathan Harms (Waterloo), Kelly Ramsay (Waterloo), Bailey Kacsmar (Waterloo), Tim Dockhorn (Waterloo), Yiwei Lu (Waterloo), Aseem Baranwal (Waterloo)

Master's Thesis: Sachin Vernekar (Waterloo), Sushant Agarwal (Waterloo), Shubhankar Mohapatra (Waterloo), Amur Ghose (Waterloo), Beracira Chen (Waterloo), Kaiwen Wu (Waterloo), Lingyi Zhang (Waterloo), Nivasini Ananthakrishnan (Waterloo), Thomas Humphries (Waterloo), Harry Sivasubramaniam (Waterloo), Xinda Li (Waterloo), Haolin Yu (Waterloo), Emily Lepert (Waterloo), Niki Hasrati (Waterloo), Abdulrahman Diaa (Waterloo)