Gavin R. Brown **Curriculum Vitae**

grbrown@cs.washington.edu (419) 560-6660

Education

PhD in Computer Science, January 2024
Boston University, Boston, Massachusetts
Advisor: Adam Smith
Dissertation: *Memorization and Privacy in Learning: Fundamental Limits and New Algorithms*

Bachelor of Science in Mathematics, May 2015
Case Western Reserve University, Cleveland, Ohio
Cumulative GPA: 3.36 / 4.0

Research Experience

Postdoctoral Scholar, advised by Sewoong Oh September 2023 - Present
Paul G. Allen School of Computer Science and Engineering, University of Washington

Graduate Research Assistant, advised by Adam Smith September 2019 - September 2023
Computer Science Department, Boston University

Graduate Research Assistant, advised by Peter Chin September 2017 – December 2019
Computer Science Department, Boston University

Undergraduate Researcher, advised by David Gurarie December 2014 – May 2015
Mathematics Department, CWRU

Undergraduate Researcher, advised by Philip Taylor May 2012 – May 2013
Physics Department, CWRU

Professional Experience

Instructor October 2016 – August 2017
Kaplan Test Prep; Boston, Massachusetts

- Taught GRE, SAT, and ACT preparatory classes.

Junior Associate June 2015 – May 2016
Mu Sigma; Bangalore, India, and Wallingford, Connecticut

- Used statistics and machine learning techniques on large-scale marketing problems.

Awards

- COLT Best Student Paper Award, 2023
- Research Excellence Award 2022-2023, Boston University Department of Computer Science.
- Teaching Fellow Excellence Award, 2019-2020, Boston University Department of Computer Science.

Publications

- **Gavin Brown** and Lydia Zakynthinou. "Tukey depth mechanisms for practical private mean estimation." https://arxiv.org/abs/2502.18698
- Maryam Aliakbarpour, Konstantina Bairaktari, **Gavin Brown**, Adam Smith, Nathan Srebro, and Jonathan Ullman. "Metalearning with Very Few Samples Per Task." *COLT 2024*. https://arxiv.org/abs/2312.13978

- **Gavin Brown**, Jonathan Hayase, Samuel B. Hopkins, Weihao Kong, Xiyang Liu, Sewoong Oh, Juan C. Perdomo, and Adam Smith. "Insufficient Statistics Perturbation: Stable Estimators for Private Least Squares." *COLT 2024*. https://arxiv.org/abs/2404.15409
- **Gavin Brown**, Krishnamurthy (Dj) Dvijotham, Georgina Evans, Daogao Liu, Adam Smith, and Abhradeep Thakurta. "Private Gradient Descent for Linear Regression: Tighter Error Bounds and Instance-Specific Uncertainty Estimation." *ICML 2024*. https://arxiv.org/abs/2402.13531
- **Gavin Brown**, Samuel B. Hopkins, and Adam Smith. "Fast, Sample-Efficient, Affine-Invariant Private Mean and Covariance Estimation for Subgaussian Distributions." *COLT 2023*. https://arxiv.org/abs/2301.12250
  Best Student Paper.
- **Gavin Brown**, Mark Bun, and Adam Smith. "Strong Memory Lower Bounds for Learning Natural Models." *COLT 2022.* https://arxiv.org/abs/2206.04743
- **Gavin Brown**, Shlomi Hod, and Iden Kalemaj. "Performative Prediction in a Stateful World." *AISTATS 2022.* https://arxiv.org/abs/2011.03885
- **Gavin Brown**, Marco Gaboardi, Adam Smith, Jonathan Ullman, and Lydia Zakynthinou. "Covariance-Aware Private Mean Estimation Without Private Covariance Estimation." *NeurIPS 2021 Spotlight*. https://arxiv.org/abs/2106.13329
- **Gavin Brown**, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. "When is Memorization of Irrelevant Training Data Necessary for High-Accuracy Learning?" *STOC 2021.* https://arxiv.org/abs/2012.06421
- Louis Jensen, **Gavin Brown**, Xiao Wang, Jacob Harer, and Sang Chin. "Deep Learning for Minimal-context Block Tracking through Side-channel Analysis." *ICASSP 2019.* IEEE, 2019.
- Xiao Wang, Quan Zhou, Jacob Harer, **Gavin Brown**, Shangran Qiu, Zhi Dou, John Wang, Alan Hinton, Carlos A. Gonzalez, and Peter Chin. "Deep learning-based classification and anomaly detection of side-channel signals." *Cyber Sensing 2018*. Vol. 10630. International Society for Optics and Photonics, 2018.
- Sang Chin, Jonathan Cohen, Alison Albin, Mykola Hayvanovych, Elizabeth Reilly, **Gavin Brown,** and Jacob Harer. "A Mathematical Analysis of Network Controllability Through Driver Nodes." *IEEE Transactions on Computational Social Systems* 4, no. 2 (2017): 40-51.
- Jiayuan Miao, **Gavin Brown**, and Philip Taylor. "Theoretically guided design of efficient polymer dielectrics." *Journal of Applied Physics* 115.9 (2014): 094104.

Talks

- Workshop on Privacy-Preserving Machine Learning. "Tukey Depth Mechanisms for Practical Private Mean Estimation." March 21, 2025.
- Google Privacy Seminar. "Stable Estimators for Fast Private Statistics." November 13, 2024.
- COLT 2024. "Insufficient Statistics Perturbation: Stable Estimators for Private Least Squares." July 1, 2024.
- TTIC Workshop on New Frontiers in Algorithmic Robust Statistics. "Stable Estimators for Fast Private Statistics." June 13, 2024.
- UW Theory Seminar. "Stable Estimators for Fast Private Statistics." May 28, 2024.
- MIT A&C Seminar. "Stable Estimators for Fast Private Statistics." May 15, 2024.
- COLT 2023. "Fast, Sample-Efficient, Affine-Invariant Private Mean and Covariance Estimation for Subgaussian Distributions." July 12, 2023.
- FORC 2023. "Fast, Sample-Efficient, Affine-Invariant Private Mean and Covariance Estimation for Subgaussian Distributions." June 9, 2023.
- Harvard Privacy Tools DP Meeting. "Fast and Sample-Efficient Private Mean Estimation." May 19, 2023.

- Google Privacy Seminar. "Example Memorization in Learning: Batch and Streaming." August 17, 2022.
- Fields Workshop on Differential Privacy and Statistical Data Analysis. "Covariance-Aware Private Mean Estimation Without Private Covariance Estimation." July 26, 2022.
- COLT 22. "Strong Lower Bounds for Learning Natural Models." July 2, 2022.
- NeurIPS 21. "Covariance-Aware Private Mean Estimation Without Private Covariance Estimation." January 9, 2021. Prerecorded.
- University of Washington Theory Lunch. "Covariance-Aware Private Mean Estimation Without Private Covariance Estimation." October 29, 2021.
- Boston University AI Research Initiative Seminar. "When Is Memorization of Irrelevant Training Data Necessary for High-Accuracy Learning?" October 25, 2021.
- STOC 21. "When Is Memorization of Irrelevant Training Data Necessary for High-Accuracy Learning?" June 21, 2021.
- Penn State University Statistical Data Privacy Seminar. "When Is Memorization of Entire Examples Necessary for High-Accuracy Learning?" May 6, 2021.
- Hebrew University Theory Seminar. "When Is Memorization of Entire Examples Necessary for High-Accuracy Learning?" May 5, 2021.
- Workshop on the Theory of Overparameterized Machine Learning. "When Is Memorization of Irrelevant Training Data Necessary for High-Accuracy Learning?" April 20, 2021.
- Boston University Probability and Statistics Seminar. "When Is Memorization of Entire Examples Necessary for High-Accuracy Learning?" April 11, 2021.

Program Committees Service and Reviewing

TMLR; ICLR 2025; AAAI 2025; COLT 2024; ICML 2024; ICLR 2024; NeurIPS 2023; AISTATS 2023; FAccT 2022, 2023; TPDP 2022, 2024.

Teaching Experience

Boston University Teaching Fellow

- CS 537 – Randomness in Computing, Spring 2020, with Sofya Raskhodnikova
- CS 330 – Introduction to Algorithms, Fall 2019, with Adam Smith
- CS 542 – Machine Learning, Summer I 2019, with Peter Chin
- CS 112 – Introduction to Computer Science II, Fall 2018, with Christine Papadakis-Kanaris
- CS 542 – Machine Learning, Spring 2018, with Peter Chin

GRE, SAT, and ACT Instructor, Kaplan Test Prep