

Nicole Elyse Mitchell

B.S., M.S. Computer Science | Research Scientist

nicole.elyse.mitchell@gmail.com • San Francisco, CA • nicolemitchell.github.io • [Google Scholar](#) • [LinkedIn](#) • [Substack](#)

Education

Rice University, George R. Brown School of Engineering • Houston, TX

BS Dec 2018, MS May 2020

MS, Computer Science; GPA: 3.67/4.00

Research Focus: Machine Learning on Graph-Structured Data (advisor: Prof. Lydia Kavraki)

Awards: Adobe Research Women in Technology Scholarship (2019), Rice Computer Science Graduate Research Fellowship (2019)

BS, Computer Science; GPA: 3.79/4.00

Activities: NCAA Division I Track & Field

Awards: President's Honor Roll (2018), CRA-W Grace Hopper Celebration Research Scholar (2018)

Research Experience

Google Research, SF • Privacy, Unlearning, Federated Learning

Feb 2021 – Present

Research Scientist — May 2022 – Present

Advancing Formal Privacy Methods and Empirical Privacy Auditing

- Investigated how to best scale user-level differentially private training to LLMs under fixed compute budgets. [4]
- Introduced a new canary design leveraging high-temperature sampling to yield highly sensitive synthetic probes for detecting privacy leakage, plus an improved approach for auditing synthetically generated data. [1]

Models that Forget Innovation Pod

- Led a research effort leveraging modularity and unlearning to control data influence and address privacy and policy challenges. One of four proposals selected from 56 submissions to the Athena Innovation Pods program.
- Built a team of collaborators, drew connections across orgs, charted out a research roadmap, and secured buy-in from leadership.
- Proposed data compartmentalization as a strategy for AI governance, presenting a taxonomy mapping governance needs to technical requirements and strategies. [12]
- Engaged with product teams across Google to document internal use cases for unlearning and guide research.
- Hosted three student researchers investigating approximate unlearning for robustness and certifiable unlearning for privacy. [3]

Federated Learning

- Explored new algorithms for merging models in federated learning, drawing on ideas from meta-learning. [5]
- Developed infrastructure to scale federated learning simulation to datasets and models orders of magnitude larger than previously used. Released an open-source library, Dataset Grouper.
- Created an interactive AI Explorable on federated learning, conceiving the project and forging a cross-company collaboration. Worked with key stakeholders to balance the narrative with Google's privacy story (spot bonus); crafted a concise, informative, visually engaging piece highlighting research and responsibility issues (spot bonus). [10]

AI Resident — Feb 2021 – May 2022

Designing & Implementing a Compression Method for Federated Learning

- Leveraged entropy coding to design a method for compressing model updates in federated learning that optimizes the communication–accuracy trade-off. Deployed this custom compression method in Google's FL production system.
- Implemented compression-based aggregation methods for federated learning in TensorFlow Federated. This involved: custom TF ops hosted in TensorFlow Compression, TensorFlow Federated logic, and robust integration tests for system compatibility. [link]
- Presented at the Google Research Conference; published a TinyPaper at ICLR 2023 with an oral presentation. [8]
- Reduced communication 320x with <0.1% drop in model accuracy on GBoard next-word prediction (spot bonus).

Kavraki Computational Robotics, AI and Biomedicine Lab • Rice University Computer Science

Jan 2018 – Oct 2020

Graduate Research

- Built a deep graph convolutional network (GCN) using Keras to predict drug metabolism. Proposed the use of molecular representations learned through GCNs to identify metabolically labile atoms, and compared this to traditional feature extraction methods. [14]

Undergraduate Research

- Improved an incremental docking protocol (DINC) which computationally predicts how peptides bind to protein receptors. Experimented to identify unexpected behavior and strengthened DINC's robustness by handling these edge cases. [9]

Internships

Apple • iCloud Storage Analytics, Software Engineering Intern

May 2018 – Aug 2018

- Developed and implemented an anomaly detection system in Python using Pandas, SciPy, and Matplotlib to automatically detect regressions in quality of service and generate alert reports. Deployed this system, which surfaces one to two critical issues each day that otherwise went unnoticed.

Square • Appointments iOS, Software Engineering Intern

May 2017 – Aug 2017

- Optimized the calendar in the Square Appointments iOS app by identifying performance bottlenecks and improving the search algorithm, making a 16-fold improvement in CPU time and restoring calendar to 60 fps scrolling.
- Developed a customized market insights tool for merchants to compare their prices to nearby sellers, grouping similar transactions using the “word2vec” ML model. Built a Python Flask app with D3 visualization to display interactive reports.

Leadership & Community Contributions

- Co-lead the Brain Women community — conduct monthly socials, organize talks, and plan events.
- Help facilitate Research Retrospectives — host interviews with established researchers to shed light on their careers.

Publications

Under Review

- [1] [N. Mitchell](#), G. Andrew, A. Ganesh, H.B. McMahan, P. Kairouz, “Advancing the State-of-the-Art in Empirical Privacy Auditing,” under review at EMNLP (2026).
- [2] [N. Mitchell](#), D. Agarwal, M. Bohacek, R. Denton, R. Patel, “Long-term Measurements: Towards a Longitudinal Understanding of Human-AI Interactions,” under review at EMNLP (2026).

Journal Articles & Conference Papers

- [3] S. Schoepf, M. Mozer, [N. Mitchell](#), A. Brintrup, G. Kaissis, P. Kairouz, E. Triantafillou, “Redirection for Erasing Memory (REM): Towards a universal unlearning method for corrupted data,” ICLR (2026).
- [4] Z. Charles, A. Ganesh, R. McKenna, H.B. McMahan, [N. Mitchell](#), K. Pillutla, K. Rush “Learning with user-level differential privacy under fixed compute budgets,” SaTML (2025).
- [5] N. Dhawan, [N. Mitchell](#), Z.B. Charles, Z. Garrett, G.K. Dziugaite, “Leveraging Function Space Aggregation for Federated Learning at Scale,” TMLR (2024).
- [6] J.H. Lee, et al, “JaxPruner: A concise library for sparsity research,” Conference on Parsimony and Learning (2024).
- [7] Z.B. Charles, [N. Mitchell](#), K. Pillutla, M. Reneer, Z. Garrett, “Towards Federated Foundation Models: Scalable Dataset Pipelines for Group-Structured Learning,” NeurIPS Benchmarks & Datasets (2023).
- [8] [N. Mitchell](#), J. Ballé, Z.B. Charles, J. Konečný, “Optimizing the Communication–Accuracy Trade-off in Federated Learning with Rate–Distortion Theory,” ICLR (2023).
- [9] D. Devaurs, D.A. Antunes, S. Hall-Swan, [N. Mitchell](#), M. Moll, G. Lizeé, L.E. Kavraki, “Using parallelized incremental meta-docking can solve the conformational sampling issue when docking large ligands to proteins,” BMC Molecular and Cell Biology, vol. 20, no. 1, p. 42 (2019).

Public Writing

- [10] [N. Mitchell](#), A. Pearce, “How Federated Learning Protects Privacy,” PAIR AI Explorables (2022).

Workshop Papers

- [11] K. Lee, et al, “Machine Unlearning Doesn't Do What You Think: Lessons for Generative AI Policy, Research, and Practice,” GenLaw @ ICML (2024).
- [12] [N. Mitchell](#), E. Triantafillou, P. Kairouz, “Examining Data Compartmentalization for AI Governance,” GenLaw @ ICML (2024).
- [13] K. Rush, Z. Charles, Z. Garrett, S. Augenstein, [N. Mitchell](#), “DrJAX: Scalable and Differentiable MapReduce Primitives in JAX,” WANT @ ICML (2024).

Thesis

- [14] [N. Mitchell](#), “Machine Learning-Based Prediction of Sites of Metabolism in Drugs: Exploring Feature Extraction Methods on Molecular Graphs,” Master's Thesis, Rice University (2020).