

TAHSEEN W. RABBANI

trabbani(AT)uchicago(DOT)edu <https://twr7bm.github.io>

EDUCATION

University of Maryland College Park, MD • Doctor of Philosophy, Computer Science, 2018-2024

Advisor: Furong Huang

Advisor: Dr. Furong Huang

New York University New York City, NY • Master's Program in Mathematics, 2017-2018

Cumulative GPA: 3.80/4.00

GPA: 3.81/4.00

University of Virginia Charlottesville, VA • Bachelor of Arts in Mathematics, 2011-2015

Cumulative GPA: 3.64/4.00 • Major GPA: 3.81/4.00

GPA: 3.81/4.00 (Major), 3.64/4.00 (Cumulative)

AWARDS, GRANTS, & NOMINATIONS

2025 Rising Star in Data Science (Stanford University).

2024 UMD CMNS Graduate Commencement Speaker.

2024 RSA Conference 2024 Security Scholar.

2023 Qualcomm Innovation Fellowship 2023, Abstract Selection, "*SWIFT: Scalable Implementation of On-Device Asynchronous Decentralized Federated Learning.*"

2022 Apple Scholars in AI/ML Nominee (Univ. of Maryland).

2022 ICSSA and Jacob K. Goldhaber Travel Grant (Univ. of Maryland). **Topic:** "On enumeration and computational construction of groups of order 1024."

2020-2024 NSF COMBINE Fellow (DGE-1632976).

2019 NSF GRFP Honorable Mention.

2019 Spotlight on Grad. Research: Seymour Goldberg Memorial Award. Univ. of Maryland.

2015 Distinguished Majors Program (**High Distinction**). Univ. of Virginia. Thesis: *p-adic Numbers and the Hasse-Minkowski Theorem.*

2015 Small Research and Travel Grant (Univ. of Virginia). **Topic:** "3-G Error-Correcting Codes."

2014 Research Grant (Provost's Office and Dept. of Mathematics). **Topic:** "Integer representation by quadratic forms." Univ. of Virginia.

2013 Small Research and Travel Grant (Univ. of Virginia). **Topic:** "Hadamard difference sets."

2012-2015 Echols Scholar. University of Virginia.

EMPLOYMENT

University of Chicago (Chicago, IL, June 2025 - Present)

Postdoctoral Scholar

- Postdoctoral research co-advised by Tian Li and Ce Zhang. My research interests are focused on systems efficiency, distributed learning, and privacy.

Scale AI (San Francisco, CA, Oct 2024 - Present)

Strategy Consultant

- Developed an automated pipeline and taxonomy for detecting and logging 200+ agentic error types.
- Validated and created over 600+ expert-level rubric items for scoring agentic evaluations on the ResearchRubrics benchmark.
- Helped develop the legal/finance benchmark, PRBench. Evaluated human/agent iter-rater agreements to select a judge and tool-calling performance on 30+ models.

Yale University (New Haven, CT, June 2024 - May 2025)

Postdoctoral Research Associate

- Postdoctoral researcher supervised by Dr. Hoon Cho and Dr. Annie Hartley. My research was focused on designs for distributed, private, low-resource clinical systems and protein-drug screening.

Error Corp. (Washington, DC, Jan 2023 - May 2024)

Machine Learning Scientist for Quantum Research

- My work was concerned with learning quantum error-correcting codes and controls. I developed and analyzed the theoretical properties of novel product constructions.

University of Maryland (College Park, MD, Jan. 2018 - Present)

Graduate Research Assistant

- Advisor: Dr. Furong Huang. My research projects have ranged over tensorial neural networks, federated learning, dimensionality-reduction, and generative watermarking.

Georgetown University (Washington, DC, Fall. 2021 - Fall 2023)

Adjunct Lecturer

- Adjunct faculty member, course: MATH035: Calculus 1.

MileMarker (Johns Hopkins University, Baltimore, MD, June 2022 - Nov 2022)

Data Science & ML Intern

- Successfully developed a series of temporally-dependent models able to predict future performance of a surgical resident. Areas: Few-shot learning and explainable ML.

Epic Systems (Madison, WI, Sep. 2015 - May 2017)

Software Developer

- Developer on the Data Courier team, concerned with pointer integrity during EMR exchange.
- Developed a terminal tool for automatically logging and searching through thousands of tickets associated with EMR send & receive errors.
- Implemented support for cross-hospital exchange of an internal EMR data type ubiquitously used by hospital administrators. At the time, a top 5 requested ‘feature’ by the team’s clients.

PUBLICATIONS

* = First author

Conference & Journal

1. A.F. Akyrek, A. Gosai, CBC Zhang, T. Rabbani et al. “PRBench: Large-scale expert rubrics for evaluating high-stakes professional reasoning.” *ACL Main 2026*.
2. M. Sharma, CBC Zhang, C. Bandi, C. Wang, **T. Rabbani**, et al. “Researchrubrics: A benchmark of prompts and rubrics for evaluating deep research agents.” *ICLR 2026*.
3. T. Bossy, J. Vignoud, **T. Rabbani**, JRT Pastoriza, M. Jaggi. “Mitigating unintended memorization with lora in federated learning for llms.” *TMLR 2026*.
4. B. An*, C. Deng*, M. Ding*, **T. Rabbani*** et al. “A Technical Report on Erasing the Invisible: The 2024 NeurIPS Competition on Stress Testing Image Watermarks.” *NeurIPS 2025*.
5. B. An*, M. Ding*, **T. Rabbani***, et al. “WAVES: Benchmarking the Robustness of Image Watermarks.” *ICML 2024*
6. D. Yuan, C Fermler, T Rabbani, F Huang, Y Aloimonos. “A linear time and space local point cloud geometry encoder via vectorized kernel mixture (veckm).” *ICML 2024*.
7. **T. Rabbani***, K. Sang*, F. Huang. “Balancing Label Imbalance in Federated Environments Using Only Mixup and Artificially-Labeled Noise.” (*ICPRAI’24*)
8. **T. Rabbani***, M. Bornstein*, & F. Huang. “Large-Scale Distributed Learning via Private On-Device Locality-Sensitive Hashing.” *NeurIPS 2023*
9. M. Bornstein*, **T. Rabbani***, E. Wang, A. Singh, & F. Huang. “SWIFT: Rapid Decentralized Federated Learning via Wait-Free Model Communication.” *ICLR 2023*.

10. T. Applebaum, J. Clickeman, J. Davis, J. Dillon, J. Jedwab, **T. Rabbani**, K. Smith, & W. Yolland. “Constructions of difference sets in nonabelian 2-groups.” *Journal of Algebra & Number Theory*, Vol. 17, **2** (2023).
11. M. Ding*, **T. Rabbani***, B. An, E. Wang, & F. Huang. “Sketch-GNN: Efficient GNNs with Graph Size-Independent Scalability.” *NeurIPS 2022*
12. **T. Rabbani***, A. Rajkumar, & F. Huang. “Practical and Fast Momentum-Based Power Methods.” *Mathematical and Scientific Machine Learning*. PMLR, Vol. 145 (2021).
13. F. Huang*, **T. Rabbani***, & A. Reustle*. “Fast GPU Convolution for CP-decomposed Tensorial Neural Networks,” *Proceedings of SAI Intelligent Systems Conference*. (2020)
14. K. Smith and **T. Rabbani**. “Nonabelian Orthogonal Building Sets,” *Proceedings of FQ14: The 14th International Conference on Finite Fields and their Applications*. (2020)
15. **T. Rabbani***. “Unique minimal forcing sets and forced representation of integers by quadratic forms,” *Rose-Hulman Undergraduate Journal of Mathematics*, Vol. 17, **1** (2016).
16. **T. Rabbani***. “Improving the Error-Correcting Code Used in 3-G Communication,” *SIAM Undergrad. Research Online (SIURO)*, **8** (2015), 126-137.

Workshop

1. M. Liu*, A. Palnitkar*, **T. Rabbani***, K. Sang* et al. “Hold Onto That Thought: Assessing KV Cache Compression On Reasoning.” *NeurIPS 2025 Workshop on Efficient Reasoning*.
2. **T. Rabbani***, N. Moll. “Assessing the Similarity of Cross-Lingual Seq2Seq Sentence Embeddings Using Low-Resource Spectral Clustering.” *ACL RESOURCEFUL '24*.
3. **T. Rabbani***, K. Sang*, T. Zhou. “Federated Meta-Learning for Low-Resource Translation of Kirundi.” *ACL RESOURCEFUL '24*.
4. (Lead Organizer) B. An*, C. Deng*, M. Ding*, **T. Rabbani***, et al. “Erasing the Invisible: A Stress-Test Challenge for Image Watermarks”. *NeurIPS'24*.
5. (Oral) A. Sallinen*, **T. Rabbani**, et al. “Llama-3-Meditron: An Open-Source Suite of Medical LLMs Based on Llama-3.1.” 4th Annual AAAI Workshop on AI to Accelerate Science and Engineering (AI2ASE). 2025.
6. M. Bornstein*, **T. Rabbani***, Brian Gravelle, F. Huang. “DECLARE: Shrinking the Size of Deep Extreme Multi-Label Classification.” *NeurIPS 2024 Workshop on Machine Learning & Compression*.
7. **T. Rabbani***, M. Bornstein*, F. Huang. “PGHash: Large-Scale Distributed Learning via Private On-Device Locality-Sensitive Hashing.” *ICLR 2023 Workshop on Sparsity in Neural Networks*.
8. **T. Rabbani***, J. Su*, X. Liu, D. Chan, G. Sangston, & F. Huang. “conv_einsum: [...] Fast Evaluation of Multilinear Operations in CTNNs.” *Third Workshop on Seeking Low-Dimensionality in Deep Neural Networks*. (2023)
9. (Oral) M. Bornstein*, **T. Rabbani***, E. Wang, A. Singh, & F. Huang. “SWIFT: Rapid Decentralized Federated Learning via Wait-Free Model Communication.” *FL-NeurIPS'22*.
10. M. Ding, X. Liu, **T. Rabbani**, & F. Huang. “Faster Hyperparameter Search on Graphs via Calibrated Dataset Condensation.” *NeurIPS 2022 Workshop: New Frontiers in Graph Learning*.
11. (Oral) **T. Rabbani***, B. Feng*, Y. Yang, A. Rajkumar, A. Varshney, & F. Huang. “Comfetch: Federated Learning of Large Networks on Memory-Constrained Clients via Sketching.” *FL-AAAI-22*.

TALKS

July 2025 CFAgentic @ ICML'25 (Vancouver, B.C.). Panelist: “Will the future of AI agents be collaborative and federated?”

July 2025 MemFM'25 @ ICML'25 (Vancouver, B.C.). "Mitigating Unintended Memorization with LoRA in Federated Learning for LLMs."

March 2025 AMLPF'24 (Washington, D.C.). "Harnessing AI and ML/LLM Technology."

Dec 2024 GDHF'24 (Nairobi, Kenya). DISCO: An online platform for data-private collaborative machine learning

April 2024 SAP Inspiration Sessions (Virtual). "Benchmarking the Robustness of Image Watermarks."

March 2024 DMV Security Workshop. "How to Attack Generative Image-Based Watermarks."

Jan 2023 Abu Dhabi, UAE @ MBZUAI. Seeking Low Dimensionality in Deep Neural Networks (SlowDNN). "Fast Evaluation of Multilinear Operations in Convolutional Tensorial NNs."

Nov 2022 New Orleans, LA. FL-NeurIPS'22. "SWIFT: Rapid Decentralized Federated Learning via Wait-Free Model Communication."

Nov 2022 New Orleans, LA. GLFrontiers (NeurIPS'22). "Sketch-GNN: Efficient GNNs with Graph Size-Independent Scalability."

Nov 2022 New Orleans, LA. GLFrontiers (with NeurIPS 2022). "Faster Hyperparameter Search for GNNs via Calibrated Dataset Condensation."

June 2022 Mantua, Italy. Combinatorics 2022, "49,487,367,289: On enumeration and computational construction of groups of order 1024."

March 2022 FL-AAAI-22 (Virtual), "Comfetch: Federated Learning of Large Networks on Memory-Constrained Clients via Sketching."

August 2021 MSML21 (Virtual), "Practical and Fast Momentum-Based Power Methods."

June 2019 Vancouver, BC. Finite Fields Conference (FQ14), "New Constructions of Hadamard Difference Sets."

April 2019 Monroe L. Martin Spotlight Talks (Winner), Univ. of Maryland.

Oct. 2017 New York University, Master's Learning Seminar, "The Sum of Squares and Universal Quadratic Forms."

Jan. 2015 San Antonio, TX. Joint Mathematics Meeting, MAA Poster Session, "Improving the Error-Correcting Code Used in 3-G Communication."

Jan. 2014 Baltimore, MD. Joint Mathematics Meeting, MAA Poster Session, "Bent Functions and Difference Sets."

TEACHING & SERVICE

Reviewer: NeurIPS, ICLR, ICML, ICASSP, AISTATS, MSML, IEEE TPDS, FedKDD, ICDCS.

Dec 2024 Neurips'24 Competition (Vancouver, BC). *Lead Organizer*. "Erasing the Invisible: A Stress-Test Challenge for Image Watermarks."

ATTRIB @ NeurIPS'24 Program Committee Member.

ICDCS'24,25 Program Committee Member (FL, Analytics, and Deployment track).

Private ML @ ICLR 2024 Program Committee Member.

FL@FM-NeurIPS23 Program Committee Member.

FL-ICML-23 Program Committee Member.

FL-NeurIPS-22 Program Committee Member.

Fall 2020 - Fall 2021 University of Maryland, COMBINE Director of Undergraduate Research.

Spring 2020 University of Maryland, Discrete Data Structures, Teaching Assistant (No rating).

Fall 2019 University of Maryland, Calculus 1, Teaching Assistant (Rating: 3.64/4.00).

Spring 2019 Directed Reading Program. Student: Samuel Howard. Topic: Error-Correcting Codes.

Spring 2019 University of Maryland, Calculus 1, Teaching Assistant (Rating: 3.80/4.00).

Fall 2018 University of Maryland, Calculus 1, Teaching Assistant (Rating: 3.59/4.00).

Spring 2018 New York University, Math. of Econ II, Recitation Leader.

Spring 2018 New York University, Abstract Algebra, Grader.