

RESEARCH INTERESTS

My research interests are at the intersection of machine learning, code optimization, and computer architecture. These include:

- Hardware-aware autonomous performance engineering
- ML for systems infrastructure optimizations
- Energy-efficient distributed ML systems

My current research is in LLM-driven GPU kernel optimization. I am exploring synthetic data generation and inference-time approaches to enable autonomous kernel optimizations that utilize the underlying hardware and kernel porting between different GPU architectures.

EDUCATION

Harvard University

Cambridge, MA

- Ph.D. in Computer Science, Supported by NSF GRFP
- Expected graduation in May 2029.
- Advisor: Prof. Vijay Janapa Reddi

2024 – 2029

Duke University

Durham, NC

- B.S. with double major in Electrical Engineering & Computer Science
- Graduated with Distinction in 3 years
- Advisor: Prof. Dan Sorin

2021 – 2024

RESEARCH EXPERIENCE

Harvard University, Edge Computing Lab Graduate Research Assistant

Cambridge, MA
Aug 2024 – Present

- Researching hardware-aware autonomous GPU kernel optimization. Enabling LLMs to make hardware-specific kernel optimizations and use machine feedback as reinforcement learning training rewards to port kernels between GPU architectures.
- Led power measurement and energy-efficiency benchmarking frameworks for distributed ML systems in collaboration with MLCommons, published in IEEE HPCA.
- Advisor: Prof. Vijay Janapa Reddi

AMD, Research and Advanced Development Research Intern

Santa Clara, CA
May 2025 – Aug 2025

- Led work on profiling-guided autonomous GPU kernel (Triton) optimization for ML and science workloads and contributed to open source AMD IntelliPerf repository.
- Achieved up to 2.1× speedup and 70% L2 hit rate improvement with scheduling optimizations for disaggregated GPU architectures.
- Selected as a Spotlight Presentation at the NeurIPS ML for Systems Workshop.
- Advisors: Muhammad Awad, Ryan Swann, Ganesh Dasika

Duke University, Sorin Lab Undergraduate Research Assistant

Durham, NC
Aug 2022 – May 2024

- Worked on rigorous evaluation of computer architectures with statistical guarantees in realistic environments, published in IEEE MICRO.
- Contributed to research on simulating and evaluating microarchitectural side-channels for high-security, low-overhead defense schemes.
- Advisor: Prof. Dan Sorin

Duke University, APEX Lab Undergraduate Research Assistant

Durham, NC
Jan 2023 – May 2023

- Collaborated with Duke University's Donald Lab to develop a hardware accelerator for energy minimization calculations in computational protein design.

- Deployed the design to FPGA for 3.4× faster cyclic coordinate descent algorithm execution over existing GPU implementations.
- Advisors: Prof. Lisa Wu Wills, Prof. Bruce Donald

Princeton Plasma Physics Laboratory
Engineering Intern

Princeton, NJ
Jan 2021 – Jun 2021

- Semester-long intern with the Princeton University and U.S. Department of Energy national laboratory working on efficient ML algorithms for science workloads.
- Led work in designing and implementing new algorithms to model and predict changes in ocean currents from raw NASA satellite data.
- Advisor: Eliot Feibush

INDUSTRY
EXPERIENCE

Yext
Software Engineering Intern

New York, NY
May 2023 – Aug 2023

- Intern on consumer data team at Yext (NYSE: YEXT) during the company's first deployments of ML search optimization as a service products.
- Developed backend daemon to asynchronously log entity updates and plug into customer-facing ML serving infrastructure.

Ribbon Home
Software Engineering Intern

New York, NY
May 2022 – Aug 2022

- Intern at Ribbon Home, a Series C Greylock-backed real estate technology startup.
- Created an internal TypeScript administrative dashboard to view, change, and deploy new transaction fees to new contracts.
- Startup acquired by EasyKnock shortly after internship

PUBLICATIONS

Complete List: Google Scholar [SLVf-nMAAAAJ]

1. **A. Tschand**, K. Ramakrishnan, M. Awad, R. Swann, J. Ma, K. Lowery, G. Dasika, and V. J. Reddi, "Swizzleperf: Hardware-aware LLMs for GPU kernel Performance Optimization," *NeurIPS ML for Systems Workshop Spotlight*, 2025.
2. **A. Tschand**, A. Rajan, S. Idgunji, A. Ghosh, J. Holleman, C. Kiraly, P. Ambalkar, R. Borkar, R. Chukka, T. Cockrell, et al., "MLPerf Power: Benchmarking the Energy Efficiency of Machine Learning Systems from μ Watts to MWatts for Sustainable AI," in *2025 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2025.
3. F. Mazurek, **A. Tschand**, Y. Wang, M. Pajic, and D. Sorin, "Rigorous Evaluation of Computer Processors with Statistical Model Checking," in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture*, 2023.
4. S. Prakash, A. Cheng, J. Yik, **A. Tschand**, R. Ghosal, I. Uchendu, J. Quaye, J. Ma, S. Grampurohit, S. Giannuzzi, et al., "QuArch: A Question-Answering Dataset for AI Agents in Computer Architecture," *IEEE Computer Architecture Letters*, 2025.
5. S. Prakash, A. Cheng, O. Kindgren, A. Ahamed, G. Knight, J. Kufel, F. Rodriguez, **A. Tschand**, D. Kong, et al., "Lifetime-Aware Design of Item-Level Intelligence," *arXiv:2509.08193*, 2025.
6. W. Li, P. Crowley, **A. Tschand**, Y. Wang, M. Pajic, and D. Sorin, "Rigorous Evaluation of Microarchitectural Side-Channels with Statistical Model Checking," *arXiv:2510.02475*, 2025.
7. **A. Tschand**, "Semi-Supervised Machine Learning Analysis of Crop Color for Autonomous Irrigation," *Smart Agricultural Technology*, 2023.

SKILLS

Programming Python, C/C++, CUDA, PTX, Triton, HIP, Verilog, Java, Assembly
Tools NVIDIA system profiling, AMD system profiling, PyTorch, DSPy, AWS, Docker, Git

AWARDS

- NSF Graduate Research Fellowship Program (GRFP) Recipient 2025
- NeurIPS ML for Systems Workshop Spotlight Presentation 2025
- Marie Foote Reel Award (Top Undergraduate Research in Duke ECE Department) 2024
- ISEF Top Award, Craig R. Barrett Award for Innovation 2021
- ISEF 1st Overall Category Award, Engineering Mechanics 2021
- Science Talent Search Scholar 2021