

Michael Kuchnik | CV

✉ michaelkuchnik@gmail.com • 🌐 michaelkuchnik.com • 🇺🇸 US Citizen

Education

Carnegie Mellon University

PhD, Computer Science

Advisors: George Amvrosiadis and Virginia Smith

Thesis: Beyond Model Efficiency: Data Optimizations for Machine Learning Systems

Carnegie Mellon University

MS, Computer Science

Georgia Institute of Technology

BS, Computer Engineering

GPA: 3.93/4.00 — Highest Honors

Awards: NDSEG Fellowship (2018), NSF GRFP Honorable Mention (2016, 2017)

Pittsburgh, PA

August 2016–May 2023

Pittsburgh, PA

August 2016–May 2020

Atlanta, GA

August 2012–May 2016

Research Experience

Carnegie Mellon University

Graduate Researcher

Thesis on building systems for efficient machine learning by adding first-class support to the data pipeline (e.g., optimized data loading, data augmentation, testing). Prior work focused on schedulers and filesystems.

Georgia Institute of Technology

Undergraduate Researcher

Researched methods for binary translation from one arbitrary ISA to another (e.g., MIPS to ARM) using superoptimizers and machine learning. Advised by Linda Wills.

Georgia Institute of Technology

Undergraduate Researcher

Led team of 5 in designing a car for children with disabilities capable of steering using only head motion.

Pittsburgh, PA

October 2016–Present

Atlanta, GA

January 2014–May 2016

Atlanta, GA

January 2013–December 2014

Relevant Work Experience

Google

Research Intern

Researched the extent to which accelerators (e.g., TPU) are bottlenecked by I/O. Contributed code to Tensorflow's `tf.data` (e.g., metrics) for fleetwide statistical analysis, which was used to approximate impact of bottlenecks. Built a system, `Plumber`, to automatically diagnose source of I/O bottlenecks in `tf.data` pipelines (published).

Los Alamos National Lab

Graduate Student Researcher

Researched applying machine learning to predict SLURM job runtime properties. Extracted and processed years of scheduler data with `OpenTSDB` and `Apache Spark`. Work contributed to technical report.

Los Alamos National Lab

Graduate Student Researcher

Researched applying machine learning in HPC simulations to reduce I/O. Ran particle physics simulations over hundreds of cores in AWS and HPC clusters. Implemented a database in AWS to store visualization artifacts.

Qualcomm Technologies

Corporate R&D Intern

Wrote C++ and Python code in the Caffe ML framework (for autonomous driving). Added support for processing data augmentation features, neural network quantization/compression, and converting model features into proprietary format. Wrote library to analyze statistical properties of model activations.

Remote

May 2020–August 2020

Los Alamos, NM

May 2018–August 2018

Los Alamos, NM

May 2017–August 2017

San Diego, CA

May 2016–August 2016

Qualcomm Innovation Center

Qualcomm CDMA Technologies Intern

Wrote C++ and Python code for testing face detection software on Android Camera Systems, enabling scale-independent testing. Added support for processing batches of multiple images via ION buffer allocator.

San Diego, CA

May 2015–July 2015

NASA's Kennedy Space Center

Telemetry Engineering Intern

Wrote C++ and Python code allowing extensible simulation and testing of telemetry streams of Orion flight capsule.

Cape Canaveral, FL

June 2014–August 2014

Teaching Experience

Carnegie Mellon University

Advanced Cloud Computing (15-719) TA

Lead TA for project on Apache Spark for Machine Learning ETL. Graded assignments/tests, created and managed course infrastructure to support Spark projects along with 2 MS TAs, and held office hours.

Pittsburgh, PA

January 2022–May 2022

Carnegie Mellon University

Advanced Cloud Computing (15-719) TA

Lead TA for project on cloud infrastructure provisioning on AWS with autoscaling and Terraform. Created serverless machine learning assignment used in 2019 course offering. Graded assignments/tests and held office hours.

Pittsburgh, PA

January 2018–May 2018

Georgia Institute of Technology

Digital Design Lab (ECE 2031) Undergraduate TA

Guided students in implementing projects using an FPGA, digital logic theory, and assembly programming. Graded assignments and tests.

Atlanta, GA

Spring 2014, Fall 2014

Publications

1. **Michael Kuchnik**, Virginia Smith, and George Amvrosiadis. Validating large language models with relm. In *Proceedings of Machine Learning and Systems*, volume 5, 2023. **Outstanding Paper.**
2. **Michael Kuchnik**, Ana Klimovic, Jiří Šimša, Virginia Smith, and George Amvrosiadis. Plumber: Diagnosing and removing performance bottlenecks in machine learning data pipelines. In *Proceedings of Machine Learning and Systems*, volume 4, 2022
3. **Michael Kuchnik**, George Amvrosiadis, and Virginia Smith. Progressive compressed records: Taking a byte out of deep learning data. In *Proceedings of Very Large Databases*, volume 14, pages 2627–2641, 2021
4. Abutalib Aghayev, Sage Weil, **Michael Kuchnik**, Mark Nelson, Gregory R. Ganger, and George Amvrosiadis. The case for custom storage backends in distributed storage systems. *ACM Transactions on Storage*, 2020
5. Abutalib Aghayev, Sage Weil, **Michael Kuchnik**, Mark Nelson, Gregory R. Ganger, and George Amvrosiadis. File systems unfit as distributed storage backends: Lessons from 10 years of ceph evolution. In *Symposium of Operating Systems Principles*, 2019
6. **Michael Kuchnik** and Virginia Smith. Efficient augmentation via data subsampling. In *International Conference on Learning Representations*, 2019

Articles, Technical Reports, and Workshops

1. Abutalib Aghayev, Sage Weil, **Michael Kuchnik**, Mark Nelson, Gregory R. Ganger, and George Amvrosiadis. File systems unfit as distributed storage back ends. *USENIX ;login.*, 2020
2. **Michael Kuchnik**^{*}, Jun Woo Park^{*}, Chuck Cranor, Elisabeth Moore, Nathan DeBardeleben, and George Amvrosiadis. This is why ML-driven cluster scheduling remains widely impractical. Technical report, Carnegie Mellon University Parallel Data Lab, 2019 (^{*} denotes equal contribution)
3. George Amvrosiadis, **Michael Kuchnik**, Jun Woo Park, Chuck Cranor, Gregory R. Ganger, Elisabeth Moore, and Nathan DeBardeleben. The Atlas cluster trace repository. *USENIX ;login.*, 2018

4. Qing Zheng, George Amvrosiadis, Saurabh Kadekodi, **Michael Kuchnik**, Charles Cranor, Garth Gibson, Bradley Settlemyer, Gary Grider, and Fan Guo. Software-defined storage for fast query. PDSW-DISCS '17, 2017

Professional Service

1. (MLSys) MLSys External Reviewer (2023)
2. (TMLR) TMLR Reviewer (2022-present)
3. (NeurIPS) NeurIPS ENLSP Workshop Reviewer (2022)
4. (NeurIPS) NeurIPS Reviewer (2022)
5. (ICLR) ICLR Reviewer (2021)
6. (NeurIPS) NeurIPS Reviewer (2021), **top 8% reviewer**
7. (ICML) ICML Reviewer (2021)
8. CMU Computer Science Department Speaking Skills Reviewer (2020–2023)
9. (NeurIPS) NeurIPS Reviewer (2020), **top 10% reviewer**
10. (CMU) Computer Science Department PhD Admissions Committee (2019)
11. (Georgia Tech) Tau Beta Pi Webmaster (2015–2016)
12. (Georgia Tech) Eta Kappa Nu Treasurer (2015–2016)
13. (Georgia Tech) Eta Kappa Nu Lab Supplies and Scholarship Officer (2014–2015)

Skills, Awards, and Talks

Languages: C/C++*, Python*, SQL, Rust, Java, Matlab, Assembly, Halide, Javascript, VHDL (* denotes primary languages)

Tools: Linux, Git, Makefiles, CMake, OpenMP, MPI, Big Data and Visualization Frameworks, Cloud/DevOps Frameworks, Tensorflow/Jax, PyTorch, Embedded Systems/FPGAs/accelerators, Latex

Mentorship: Mentored five CMU MS students in research projects

Awards: MLSys Travel Grant (2022), Google Cloud Credits Grant (2020), Google TPU Research Cloud Grant (2020), SOSP Student Scholarship (2019), Georgia Tech President's Undergraduate Research Award (PURA) (2014, 2015)

Invited Talks: Plumber: SysML/LSDS Seminar at Imperial College London (2022), 15-719 lecture (2022)