

Keerthana Gurushankar

[Email](#) | [Personal Website](#) | [GitHub](#) | [LinkedIn](#)

EDUCATION

Carnegie Mellon University

PhD Student, Computer Science

Aug 2022 – Present

Research focus: *Performance Modeling for Machine Learning Systems*

B.S., M.S., Mathematics, GPA: 3.80

Aug 2017 – May 2021

WORK EXPERIENCE

- **Machine Learning Intern, CapSen Robotics** May 2023 - Aug 2023
Computer Vision models and software for robotics applications
 - Built and tuned segmentation-based ML models for detecting known objects in 3D images for robot applications
 - Contributed performance improvements to production-level project, reducing failure rate from 1/1000 to 1/3000
 - Designed & developed new calibration method, eliminating hardware changes & reducing calibration time 6-fold
- **Pre-doctoral Researcher, CMU** May 2021 – May 2022
Neural modeling with Information theoretic decompositions
 - Performed research project modeling neural data with information theory, leading to publication in top conference
 - Deduced tractability results using probability, linear algebra, algorithmic methods & data visualization
 - Collaborated with and presented work to researchers of diverse technical backgrounds

RESEARCH PROJECTS

- **Data Integrity in Clinical Trials - Citadel Global PhD Datathon 3rd Place Winner** Nov 2024 - Nov 2024
 - Produced a novel large-scale statistical analysis for data integrity in clinical trials (using Python/pandas/scipy)
 - Identified 1k trials with low integrity data among 22k published in 2014-24, using chi-square/Benford's law tests
 - Best among 30 submissions, 1k+ applicants worldwide, awarded \$2.5k for "rigorous & highly practical insights"
- **Caching with Delayed Hits** Feb 2024 - Aug 2024
 - Developed novel analysis of latency in high-throughput caches with queueing models, simulation & traces
 - Implemented cache policies to get 10-40% improvement over LRU, proved LRU is at most 4x worse than optimal
- **Deep Learning Library Implementations** Aug 2022 – Dec 2022
 - Built a complete deep learning library from scratch, with basic implementations of PyTorch and Numpy
 - Wrote efficient CPU/GPU backends using C++/CUDA, autodiff support, modules for optimizers & data loaders

SKILLS

Programming Languages: Python, C/C++, SML, MATLAB, Mathematica

SELECTED COURSEWORK

Graduate Artificial Intelligence	Deep Learning Systems	Performance Modeling	Computer Networks
Stochastic Calculus	Probability & Computing	Game Theory	Programming Languages

HONORS/AWARDS

2022 Irwin Mark Jacobs & Joan Klein Jacobs Presidential Fellowship - presented prestigious fellowship/funding at MIT
2019 D. E. Shaw Discovery Fellowship - top 30 from over 200 applicants to attend prestigious fellowship programme
2019 CMU Quantathon - member of winning team solving quantitative finance problem in university-wide competition
2018 Putnam Competition - ranked in top 500 nationally
2017 International Physics Olympiad Selection Camp - top 35 students in Physics nationally
2017 NIOS Senior Secondary Board Examination - Highest score among 300,000 students