

MAAZ M.

mmaaz.ca ◊ m.maaz@mail.utoronto.ca ◊ github.com/mmaaz-git

EDUCATION

University of Toronto, Ph.D. in Operations Research

Sep 2022 - Aug 2025

Interests: optimization, machine learning, formal verification

Awards: Vanier Scholarship (\$150,000) – given to 55 top PhD students in Canada

EXPERIENCE

Thinking Machines

Member of Technical Staff

Apr 2026 -

San Francisco, CA

- Leading posttraining data efforts to train frontier AI models
- Building synthetic data pipelines to systematically address failure modes and improve reasoning and robustness
- Worked across research and product teams to drive model performance through data sourcing

Anthropic

Research Fellow

Jun 2025 - Feb 2026

San Francisco, CA

- Worked on cybersecurity with Nicholas Carlini and Zac Hatfield-Dodds
- Built agent that uses property-based testing via Hypothesis and found 100s of bugs in popular Python libraries
- Studied generalization of reinforcement learning (GRPO) for algorithmic reasoning

Amazon

Applied Scientist (intern)

Jun 2022 - Sep 2022

Toronto, ON

- Developed outlier detection system for Amazon Alexa, using a CNN developed in Keras/Tensorflow
- Devised a novel unsupervised performance estimation technique, which yielded a 75% reduction in error compared to state-of-the-art methods in experiments – paper featured on the [Amazon Science website](#)

SELECTED PUBLICATIONS

- *Agentic property-based testing: finding bugs across the Python ecosystem* NeurIPS 2025. [arxiv:2510.09907](#)
- *Formal verification of Markov chains with learned parameters.* ICML 2025. [arxiv:2501.15767](#)
- *Siamese graph neural networks for drug discovery.* Poster at MoML 2023, Mila - Quebec AI Institute.
- *On orderings of probability vectors and unsupervised performance estimation.* IJCAI 2023. [arxiv:2306.10160](#)

SELECTED SOFTWARE

markovml

- Python package for constructing Markov processes with embedded machine learning models and verifying their properties
- Supports Sklearn and Pytorch ([Github link](#))

Citizen Gemini

- Developed an AI chatbot for citizens to ask questions about Ontario legislature (meetings, bills, etc.)
- Uses Google's Gemini LLMs with long-context, multiple models for retrieval and answer generation, and a dynamic caching approach to reduce latency ([Github link](#))

Misc. open-source contributions

- Fixed bug in NumPy: developed a more numerically stable algorithm for generating Wald samples, thus ensuring samples are non-negative
- Contributed to SymPy: first solver for polynomials with algebraic coefficients
- First implementation of cylindrical algebraic decomposition (CAD) in Python
- Utility that converts SymPy systems to Z3 in order to check feasibility

SKILLS

Python (pandas, scikit-learn, SymPy, PyTorch, Tensorflow, Keras) • R (tidyverse, ggplot2) • Javascript (React)
• Z3 • Lean • Gurobi