Arul Kolla

arulk@mit.edu • arulkolla • arulkolla

Education

Massachusetts Institute of Technology (MIT) — **Cambridge, MA**
B.S. in Computer Science (Artificial Intelligence and Decision Making), GPA: ***5.0/5.0*** — *August 2023 – May 2027*
Selected coursework: Deep Learning (grad); Advanced Algorithms (grad); Quantitative Methods for Natural Language Processing (grad); Hardware Architecture for Deep Learning (grad); Representation, Inference, and Reasoning in AI; Computation Structures.

Experience

MIT Energy-Efficient Circuits and Systems Group — **Cambridge, MA**
Researcher, Efficient AI Inference with Nonlinear Kernel Approximation | Python — *Aug 2025 – present*

- Designed and implemented approximation algorithms for nonlinear functions in transformer architectures, aiming to reduce their computational cost by up to 40% while maintaining minimal accuracy loss via polynomial approximation.
- Prototyped hardware support for approximated nonlinear functions on GPU kernels, aiming to compare to baseline implementations, using PyTorch integration, RTL modeling, and optimization techniques.

MIT FutureTech Lab — **Cambridge, MA**
Researcher, ML Esolangs Project | Python — *Aug 2024 – present*

- Pioneered research on esoteric languages to highlight systemic gaps in large language model performance in zero-shot scenarios, advancing understanding of LLM limitations in unconventional programming paradigms.
- Drove measurable gains via "self-scaffolding" in-context augmentation of over 15% without any fine-tuning.

Amazon — **Bellevue, WA**
Software Development Engineering Intern (Amazon AGI) | Python — *June 2025 – August 2025*

- Drove cost savings projected over $1,000,000 by identifying and archiving 95% of model checkpoints that are never restored, with retry-tolerant orchestration of long-latency Glacier operations.
- Orchestrated a resilient end-to-end workflow with Lambda and Step Functions to monitor, retry, and finalize multi-day S3 operations across thousands of models.

OpenAI — **San Francisco, CA (Remote)**
Computational Linguistics Specialist (Contract via Mercor) | Python — *Dec 2024 – Jan 2025*

- Analyzed 100+ complex linguistics problems and multi-step reasoning chains for RLHF, contributing to a significant increase in model factual accuracy and a reduction in logical hallucinations.
- Benchmarked model reasoning across 10+ types of tasks, providing expert-level feedback on adversarial test cases that improved model alignment with human solutions in zero-shot scenarios.

Projects

Un-watermarking Neural Decoders | *Python, PyTorch* — *Dec 2025*

- Built a 4-stage LoRA training pipeline around an 8B-parameter Llama-3.1 decoder (baseline probe, radioactive fine-tuning, Dt construction, adversarial fine-tuning), with automated computation of compression advantages on both radioactive and generic datasets.
- Evaluated two curricula (random vs. top-k high-loss selection) and empirically showed that only high-loss fine-tuning simultaneously increases generic compression advantage and decreases radioactive compression advantage, demonstrating an elastic un-watermarking effect that random continued fine-tuning fails to achieve with 90% efficacy.

Semantic Contamination Detection in LLMs | *Python, TensorFlow* — *Jan 2025*

- Developed robust detection pipelines leveraging perplexity, entropy, and min-$k\%$ token probabilities to analyze how semantically similar yet lexically distinct training examples can lead to misleading evaluation performance, achieving 93% accuracy for reliable detection of contamination signals.
- Created the TOFU-PARA and TOFU-WIKI datasets to rigorously benchmark semantic contamination.

Awards and Honors

- Silver and bronze medalist at the International Linguistics Olympiad (**IOL**) in 2023.
- International Master rank on **Codeforces** (top 0.5%) in 2025.
- Qualifier for the USA Mathematical Olympiad (**USAMO**) in 2023.
- First place internationally in the North American Computational Linguistics Olympiad (**NACLO**) in 2023.

Skills

- **Technical:** C++, Python, Java, C, LaTeX, HTML, JavaScript, CSS, TypeScript, React, PHP, Git