

AMAN SHARMA

Atlanta, GA | +1 (470) 890-8006 | amananytime07@gmail.com
arcaman07.github.io | github.com/arcAman07 | LinkedIn | Google Scholar

EDUCATION

Master of Science in Computer Science | **Georgia Institute of Technology** | Aug 2026 - May 2028 (Expected)
Graduate Student Researcher, advised by Prof. Chao Zhang

Bachelor of Technology in Computer Science and Engineering | **Vellore Institute of Technology** | Oct 2020 - Apr 2025
Coursework: Data Structures and Algorithms, Deep Learning, NLP, Applied Linear Algebra

EXPERIENCE

Lossfunk AI Lab

Research Engineer

Bangalore, India
Jun 2025 - Aug 2026

- Designed and open-sourced EsoLang-Bench (400 task-language instances), revealing frontier-model accuracy dropping from **100% on Python/JavaScript to 0-11%** on equivalent problems in esoteric languages.
- Identified metaprogramming strategies and distilled reusable Python helpers; helpers plus guidance raised Sonnet 4.6's solve rate from **15% to 80%** on EsoLang-Bench's Brainfuck suite.
- Developed entropy-based early stopping from token logprobs, reducing reasoning-token use and computation by **25-50% at matched accuracy** on AIME 2024/2025 and GPQA Diamond reasoning benchmarks.
- Introduced inverse-entropy voting, matching or exceeding six baselines (including majority voting) in **96.7% of sequential configurations** across 5 models on AIME and GPQA Diamond reasoning benchmarks.

MIT Media Lab (Large Population Models)

Research Engineer, advised by Dr. Ayush Chopra

Remote | Part-time
Feb 2024 - Aug 2026

- Post-trained **Qwen2.5-1.5B** coordinators with LoRA, SFT and GRPO for multi-agent scheduling and routing; measured **28% lower sequential execution depth** versus a sequential multi-agent baseline on APEX tasks.
- Co-developed Ripple Effect Protocol with Cisco for decentralized resource allocation, reducing simulated supply-chain cost by **41.8% versus A2A** in the Beer Game benchmark.
- Engineered AgentTorch's vectorized PyTorch runner for distributed COVID-19 and predator-prey simulations, saving **5-10 milliseconds per step** in a **100,000-agent COVID-19 model**.

Ikigai Labs

Machine Learning Engineer

Bangalore, India
May 2023 - Oct 2023

- Benchmarked LLMs against traditional tabular models for text classification, measuring **93% versus 65% accuracy** to guide production model selection.

NumFOCUS (Google Summer of Code)

Open Source Developer

Remote
Jun 2022 - Sep 2022

- Contributed **12+ automatic-differentiation rules** for Julia image registration and stitching, all merged upstream; optimized differentiable operations for a **3x speedup**.

TECHNICAL SKILLS

Languages & Frameworks: Python, C++, Java, Go, Julia, SQL, PyTorch, JAX, TensorFlow, FastAPI, Flask

ML Libraries & Tools: scikit-learn, Hugging Face Transformers, PEFT, TRL, vLLM, SGLang, CUDA, OpenCV

Research Interests: RL, post-training, LLM efficiency, multi-agent systems, alignment, model evaluation, benchmark design

Data & Infrastructure: NumPy, pandas, Docker, AWS, Google Cloud Platform (GCP), Linux, Git

SELECTED PUBLICATIONS

Aman Sharma et al. *Frontier Coding Agents Use Metaprogramming*. **NeurIPS 2026 Main Track** (accepted).

Aman Sharma, Paras Chopra. *EsoLang-Bench*. **NeurIPS 2026 Evaluations and Datasets Track** (accepted).

Aman Sharma, Paras Chopra. *Think Just Enough*. NeurIPS 2025 FoRLM Workshop.

Aman Sharma, Paras Chopra. *The Sequential Edge*. NeurIPS 2025 ER Workshop.

Ayush Chopra, Aman Sharma, et al. *Ripple Effect Protocol*. Preprint.

PROJECTS

TexGuardian

- Built and published a Python assistant for LaTeX papers with **13,800+ PyPI downloads**, integrating LLM-generated fixes, vision-model PDF review and citation validation in a **7-stage pipeline**.

Memer.AI

- Built a FastAPI meme-generation application with ResNet/LSTM models for image understanding and captions; served **3,000+ users** with **sub-second response times**.