

# Arpit Ramani

Atlanta, GA | (404) 528-0051 | [arpit.ramani07@gmail.com](mailto:arpit.ramani07@gmail.com) | [github.com/ArpitRamani](https://github.com/ArpitRamani) | [eidola.me](https://eidola.me) | [LinkedIn](#)

## EDUCATION

### Georgia Institute of Technology

*B.S. in Computer Science, GPA: 4.00*

Atlanta, GA

*Expected May 2029*

- **Relevant Coursework:** Object-Oriented Programming, Discrete Mathematics, Linear Algebra, Differential Equations

### Wheeler Magnet School of Advanced Mathematics and Sciences

*GPA: 4.71, SAT: 1580*

Atlanta, GA

*Sep 2022 – May 2026*

- **Awards:** Regeneron International Science and Engineering Fair 3rd Place Grand Award in Computational Biology and 2x Finalist; Society for Neuroscience and American Academy of Neurology presentations; Y Combinator Startup School 2026

## EXPERIENCE

### FinalDose (Y Combinator P26)

*Machine Learning Engineer Intern*

London, UK

*Jun 2026 – Present*

- Trained a scikit-learn random-forest classifier to 0.852 AUROC over 1,368 SpCas9 positions, beating a PyTorch ESM protein-language-model baseline at 0.777 and the prior in-house baseline at 0.826. [REPOSITORY](#)
- Built the feature pipeline behind the AUROC gain in Python and NumPy, pairing mean-field direct-coupling analysis over a 5,254-sequence alignment with graph-topology descriptors on the coevolution network, then showed the remaining ceiling is label noise rather than the model.
- Rebuilt the free-energy pipeline around a 31-window alchemical cycle with MBAR on AWS EC2 A10G instances, cutting GPU time per estimate roughly 15x, from 2,300 ns of sampling to 150 ns.

### Emory University, Dr. Jingjing Yang Lab

*Computational Genomics Research Intern*

Atlanta, GA

*Jun 2024 – Present*

- Ran a Bayesian spike-and-slab xWAS pipeline over 931 transcriptomic and 716 proteomic ROSMAP samples across 1,000+ hours of Emory HPC compute, surfacing 26 Alzheimer's risk genes invisible to cis-only methods (PrediXcan, FUSION, TIGAR) among 130 TWAS and 36 PWAS genes mapped.
- Packaged the Bayesian xWAS pipeline for reuse across distributed Slurm job arrays, with parameterized configs keeping data-use-agreement paths out of version control and synthetic fixtures covering schema validation. [REPOSITORY](#)

### Emory University, Dr. Blaine Roberts Lab

*Computational Proteomics Research Intern*

Atlanta, GA

*May–Aug 2026, Jul–Dec 2025*

- Built a differential-expression pipeline in R over a sparse 34-patient ALS cohort, chaining TAMPOR batch-effect normalization, random-forest imputation, and empirical-Bayes moderated tests under FDR control to recover statistical power. [REPOSITORY](#)
- Deployed a distributed molecular dynamics pipeline on AWS ParallelCluster, fanning replicas across Slurm GPU arrays with OpenMM and CUDA under an AMBER ff19SB/OPC force field at a 4 fs timestep. [REPOSITORY](#)
- Mapped 1,038 ALS and 1,628 Alzheimer's CSF nitropeptide sites against the FragPipe search database in Python, filtering to 506 and 479 high-confidence residues for conservation analysis. [REPOSITORY](#)

## PROJECTS

[Eidola](#) | *Next.js, FastAPI, Postgres/pgvector, Docker, Cloud Run, Stripe, Swift*

Mar 2026 – Present

- Launched a multi-tenant memory app now serving 106 active users, built on a Next.js and FastAPI stack shipped through Cloud Build CI/CD onto Cloud Run with row-level-security isolation, Stripe billing, 149 endpoints, 103 test suites, and a 20.5k-line Swift iOS/watchOS client.
- Engineered the retrieval-augmented generation layer: an HNSW index over 3,072-dimension half-precision embeddings in pgvector, per-user HDBSCAN topic clustering, and answers streamed token by token over server-sent events.
- Earned \$25,000 in cloud compute credits from Anthropic, Google Cloud, and AWS through Y Combinator Startup School.

[Ferromilab](#) | *Python, AlphaFold, MODELLER, Rosetta, HADDOCK, CHARMM QM/MM*

May 2024 – May 2025

- Automated an end-to-end structure-prediction and docking pipeline (AlphaFold and MODELLER refinement, HADDOCK epitope mapping, 1,352 grafted Rosetta CDRs) for bispecific antibody design against the hepcidin-ferroportin interface.
- Ran the pipeline on Google Cloud Compute A100 GPUs, Google Colab notebooks, and AWS EC2, with CHARMM QM/MM partitioned onto polar-contact residues to keep the quantum region tractable.

## TECHNICAL SKILLS

**Languages:** Python, TypeScript, JavaScript, Swift, R, SQL, Bash

**ML & Data:** PyTorch, scikit-learn, gradient boosting, pandas, NumPy, RAG, vector search

**Tools & DevOps:** Docker, Kubernetes, GKE, AWS, Google Cloud Run/Build, CI/CD, Slurm, CUDA, Git, Linux

**Web & Backend:** Next.js, React, FastAPI, PostgreSQL, REST APIs, Stripe, Supabase