Rohan Sinha

rs2748@cornell.edu | 848-667-2205 | linkedin.com/in/rohansinha23 | github.com/newrohansinha

Education

Cornell University — Ithaca, NY
Computer Science — *Expected May 2029*
- Relevant Coursework: Data Structures, Object-Oriented Programming, Linear Algebra, Calculus I, Calculus II

Experience

Reinforcement Learning Engineer — Sep 2026 – Present
Cornell University — *Ithaca, NY*
- Reduced simulated pinball drag 35–42% versus uncontrolled flow by training PPO and testing on held-out states.
- Built PPO/GPPO comparisons across 4 configurations using JAX to differentiate policies through fluid simulations.
- Cut benchmarked GPPO update time 27% by optimizing checkpointing while preserving numerical gradients.
- Scaled RL across 4 NERSC A100 GPUs using Slurm pipelines to save and resume policy weights and optimizer state.

Machine Learning Engineer — Feb 2026 – Aug 2026
Temple Allen Industries — *Rockville, MD*
- Reduced actuator position error 38% by training SAC policies on a custom Gymnasium actuator simulator.
- Improved aircraft-part detection mAP 12.4% by tuning PyTorch CV models with OpenCV and IoU analysis.
- Built a Python actuator-analysis harness with NumPy/pandas to measure PID gain sensitivity and recovery time.
- Parsed 100 Hz telemetry streams with pandas to compute p95 latency, settling error, and performance regressions.
- Automated MLflow tracking for 37 checkpoints, standardizing IoU, precision, recall, and actuator metrics.

Computational Number Theory Researcher — Apr 2025 – May 2026
University of Maryland — *College Park, MD*
- Co-authored ACM SIGACT News publication on empirical congruence monoids and empirical prime densities.
- Cut Miller–Rabin per-check latency 33% across 1B+ candidates by optimizing Numba primality test routines.
- Validated 2 prime-density models with SciPy nonlinear optimization, matching empirical results at full scale.
- Achieved 3% empirical MAPE by developing custom Numba sieves for large-scale prime-density estimation.

NLP Researcher — Nov 2024 – Jun 2026
The George Washington University — *Washington, DC*
- Analyzed pre/post-COVID semantic drift with Python NLP pipeline for 22,200 Reddit Q&A pairs 2019–2023.
- Generated 133,200 calibrated ratings using batched PyTorch inference across 6 semantic/personality traits.
- Improved human-model agreement 15% after scikit-learn calibration on annotated validation and test datasets.
- Tested observed semantic shifts with statsmodels, placebo windows, Newey–West errors, and forecast checks.

Projects

Adaptive gRPC Compression with Contextual Bandits — Jan 2026 – July 2026
- Reduced p95 RPC latency 61.7% across 161K calls by training LinUCB to select gzip/zstd from runtime context.
- Saved 863.7M payload bytes (56.7%) across 161K RPCs by adapting compression to entropy and CPU load.
- Achieved 1.38 μs decision latency by embedding bounded LinUCB inference in Go and streaming gRPC interceptors.

ML CI Test Conflict Detection Scheduler — Feb 2026 – June 2026
- Instrumented 280 pytest cases for CI conflicts by tracing files, sockets, ports, DBs, and processes with Rust eBPF.
- Scheduled CI across 8 workers with 3 risk policies using conflict probability, overlap duration, and LPT placement.
- Built Rust eBPF/Python/Go/Postgres pipeline to trace 280 tests, predict conflicts, and schedule workers

Supply Chain Intelligence Engine — April 2026 – May 2026
- Architected React/TypeScript + Supabase/Postgres platform with RLS, Edge Functions, and source-audit ledger.
- Built Deno ETL normalizing 2K+ BLS/EIA/FRED/UN records with Gemini-assisted claim extraction pipelines.
- Mapped external shocks to calibrated exposure scores and mitigation measures across 5+ supply-chain workflows.

Technical Skills

Languages: Python, Java, C, C++, JavaScript, TypeScript, SQL, HTML/CSS, R, Go
ML/AI: PyTorch, JAX, Flax, scikit-learn, Hugging Face Transformers, OpenCV, Gymnasium, Ray RLlib, MLflow
Technologies: NumPy, pandas, SciPy, Slurm, Git, Linux, Docker, Postgres, Jenkins, Maven, Deno, Vue.js