

Shubh Saxena

AI Engineer

shubh.saxena2012@gmail.com | +91 7703901127 | [LinkedIn](#)

EDUCATION

IIT Roorkee

2023 – 2028

Integrated M.Tech in Geophysical Technology | Minor in Economics

EXPERIENCE

Polymath (YC) | *RL Engineer*

Jul 2026 – Present

- **Architected** the task distribution for **4 long-horizon coding-agent RL environments**, specifying tool-use observation/action spaces, state-transition and reset semantics, and parameterized scenarios across checkout, release, failover, and voice recovery.
- **Designed** verifier-derived reward functions with shaped intermediate signals and terminal invariant checks, improving long-horizon credit assignment while rejecting reward hacking, no-op policies, and shallow state edits.
- **Ran distributed GRPO policy optimization** and ablations over reward normalization, KL control, rollout sampling, and curriculum difficulty; diagnosed reward distributions and failure trajectories to stabilize learning and guide environment iteration.

Decompute | *Remote MLE Intern*

Apr 2026 – June 2026 | *Cupertino, CA*

- **Enabled** date-aware recall across **3 memory classes** by architecting a memory-aware RAG layer over structured temporal records, semantic user profiles, and episodic chat history for profile-aware conversations.
- **Grounded** profile-aware responses in retrieved user evidence by fusing metadata-filtered BM25 and dense search with reciprocal-rank fusion, and specifying cross-encoder reranking with CRAG-style relevance correction for ambiguous memory queries.
- **Improved** resilience by **60%** and protected subscription revenue by combining reverse-proxy security, rate limiting, circuit breakers, and message queues for crash reporting with fail-closed Stripe webhook verification, maintaining accurate tier-gating for **100%** of users.

Department of Management Studies, IIT Roorkee | *AI Engineer Intern*

Feb 2026 – Mar 2026

- **Cut** manual compliance work by **80%** by orchestrating agentic RAG over 2,400 IPR pages, routing consumer-authenticity and seller-eligibility queries through retrieval, verification, and citation-grounded response tools.
- **Delivered >95% Recall@10** and **91.3% cache hits** through MPNet and Pinecone HNSW candidate retrieval, cross-encoder reranking, and three-tier Redis caching for grounded generation.
- **Automated 70%** of verification at **99.1% accuracy** by serving 17.4ms ResNet-50 inference as an agent tool and independently scaling retrieval to six replicas and GPU inference to eight across seven EKS services.

HCLTech | *DevOps Intern*

May 2025 – July 2025

- **Decreased** release cycles by **40%** with zero manual intervention through Jenkins → ArgoCD GitOps, Trivy/SonarQube gates, and an **80% coverage threshold** for Kubernetes microservices.
- **Saved \$55K+** annually (**96% infrastructure cost reduction**) by consolidating 100 NAT Gateways into 2 with a Terraform Transit Gateway hub-and-spoke architecture.
- **Accelerated** multi-account onboarding to **under 30 minutes** with Terraform modules for 50+ isolated AWS environments.

PROJECTS

Autonomous Swarm Intelligence & SecOps Pipeline | [GitHub](#)

- **Enabled** zero-lag drone telemetry anomaly detection across **4 health dimensions** by implementing a stateful PyFlink data pipeline over Redpanda streams integrated with **DVC** and **MLflow** for absolute ML reproducibility.
- **Maintained** stable, sub-100ms multi-agent swarm orchestration by deploying a **Ray Serve** GPU inference cluster tied to a **LangGraph StateGraph** hierarchy, and cut incident root-cause time by **65%** via **OpenTelemetry** tracing exported to **Jaeger**, centralized **ELK** logs, and Prometheus/Grafana on Kubernetes.
- **Hardened** the LLM-to-actuator command path against unsafe maneuvers by enforcing a deterministic policy gate exposed through a typed **MCP** tool boundary and validated by a deterministic **agentic evaluation harness** that replays mission-safety scenarios in CI before any command reaches the drone.

Intelligent Document AI Pipeline | [GitHub](#) | [Certificate](#)

- **Secured** $\geq 93\%$ invoice-field accuracy at **\$0.01/document** with an automated three-tier adjudication extraction ladder.
- **Reduced** production compute by **30%** through isolated Docker parsing microservices and selective SLM routing.

Autonomous Agent Evaluation & Orchestration Framework | [GitHub](#)

- **Cut** manual evaluation overhead by **4×** and increased throughput by **3×** by modeling Terminal-Bench-2 tasks as an OpenEnv-compatible FastAPI/WebSocket RL environment with GRPO reward shaping.
- **Reduced** model evaluation latency by **60%** through distributed Docker/Kubernetes runners.

RESEARCH

MACE-RL: Meta-Adaptive Curiosity-Driven Exploration | *Accepted, AAAI-26 Student Abstract Program* | [Zenodo](#)

2026

TECHNICAL SKILLS AND CERTIFICATIONS

AI/ML: PyTorch, DDP, Ray Serve, MLflow, DVC, GRPO, PPO, Contrastive Learning

RAG & Agents: MPNet, Pinecone HNSW, BM25, RRF, Cross-Encoders, LangGraph, LangSmith, MCP

AI Systems: Python, FastAPI, AWS, Kubernetes, Docker, Redis, PostgreSQL, Redpanda, Apache Flink, OpenTelemetry

Certifications: [Introduction to FinOps](#) (FinOps Foundation)