

KARTHIK RAJA ANANDAN

669-325-0235 | kanandan@ucsc.edu | linkedin.com/in/kitrakrev | github.com/kitrakrev | Pleasanton, CA

EDUCATION

University of California, Santa Cruz Sep 2024 – Dec 2025
M.S. in Artificial Intelligence, GPA: 3.9/4.0 | SemEval 2025 Task 4 Winner (Best Paper Award)
Santa Cruz, CA

Anna University (SSNCE) Aug 2019 – Aug 2023
B.E. in Computer Science, GPA: 3.9/4.0 | ICPC Top 300 | Smart India Hackathon Winner
Chennai, India

EXPERIENCE

AI/ML Engineer (Founding Engineer) Jan 2026 – Present
Elas Pleasanton, CA

- Built **160-scenario test automation framework** to validate agentic AI system correctness — automated eval suite with **golden replay harness** covering routing, tool selection, and code generation. Caught **regressions early** by injecting quality feedback before deployment.
- Designed routing engine with **100% test coverage** and **CI-enforced quality gates** — feature-aware **integration tests** verifying correctness across **10+ specialized agents**. Implemented SKIP_AUTH elimination with CI guards to prevent unsafe code in production.
- Led development across **10+ repos (800+ commits)** — authored **168-page** developer docs covering **57 service flows**, built **automated release pipelines (Argo CI/CD)**, established code review ensuring **system-wide stability**. Deployed on **AWS ECS, Kubernetes**.

AI Research Engineer Sep 2025 – Dec 2025
Samsung Research Capstone Santa Cruz, CA

- Built autonomous GUI agent for **long-horizon web task validation** — designed event capture system to record, replay, and verify multi-step **agentic workflows**. Created **integration tests** ensuring agent correctness across diverse environments.
- Developed **automated test harnesses** for quantized model inference — validated **1.5-bit** quantized **32B Qwen** maintained correctness across **2x RTX 3070 GPUs** with systematic **regression testing**.
- Designed **validation pipelines** for agent behavior — automated correctness checks on tool selection, scheduling, and task completion across long-horizon autonomous workflows.

Software Development Engineer Aug 2023 – Sep 2024
Citi Corp Chennai, India

- Architected event-driven microservices processing **millions of events daily** — built **integration tests** validating system-level correctness. Developed **CI/CD pipelines (Jenkins, Docker)** improving deployment velocity by **20%**.
- Designed automated data validation pipelines (**Hadoop, Spark**) cutting processing time by **25%**. Built **developer tooling** including internal DSL for UI validation. Established **code review** practices ensuring system stability.

NLP Engineering Intern Jun 2023 – Aug 2023
MultiOn (Stanford Startup) Palo Alto, CA

- Shipped validated integration merged into **LangChain** (PR #12392). Enhanced agentic system accuracy by **40%**. Open-source: **Google Tunix #1552, LangChain #12392 (merged), Ivy #13280 (merged)**.

RESEARCH

Graduate Researcher — UCSC (Chenguang Wang's Lab) Sep 2024 – Dec 2025
Multi-agent interpretability | AutoGen/Chi Wang | DeepMind | Google Tunix PR #1552 Santa Cruz, CA

- Contributing to **Google Tunix, Ivy** (merged), massgen, and rllm. **Samsung** capstone on agent testing. **SemEval 2025 Task 4 Winner** (Best Paper). **4 published ML papers**. Open source.

Research Assistant — HPC Lab, Anna University (SSNCE) May 2021 – Sep 2023
CLIP-based multimodal routing | Autonomous robotics | HuggingFace benchmark Chennai, India

- Built multimodal routing with **automated validation (0.90 confidence)**. Designed 6-DOF robotic arm with test-driven optimization (**73%** accuracy). Published benchmark on **HuggingFace**.

TECHNICAL SKILLS

Languages: Python, C++, Java, Go, JavaScript/TypeScript, Rust, SQL, Bash

Testing & Validation: Test Automation Frameworks, Integration Testing, Golden Replay Harness, CI/CD (Argo, Jenkins), Automated Releases, Code Review, Performance Validation, Regression Testing

Agentic AI & Dev Tools: Multi-Agent Orchestration, LLMs, vLLM, PyTorch, LangChain, MCP Servers, RAG, Eval Pipelines, DeepEval, Langfuse, GRPO, DDP

Infrastructure: Docker, Kubernetes, AWS (ECS, SQS), Redis, PostgreSQL, Kafka, Distributed Computing, Data Storage, Networking, NLP

ACHIEVEMENTS

NVIDIA GTC 2026 (GPU Kernel Demo) • **Google Tunix** Contributor (PR #1552) • **LangChain** Merged PR #12392 • **Ivy** Merged PR #13280 • **100% Test Coverage** • **SemEval 2025 Task 4 Winner** (Best Paper) • Smart India Hackathon Winner • ICPC Top 300