

Rudra Khakhar

San Luis Obispo, CA | rudrak2906@gmail.com | linkedin.com/in/rudrakhakhar | rudrakhakhar.com

SUMMARY

Computer engineer who ships production AI and stress-tests it: **LLM agents**, **RAG pipelines**, and evaluation harnesses on **AWS Bedrock**, embedded next to the people who have the problem, from requirements through post-launch support. Owns the eval loop end to end, from golden datasets and precision/recall/F1 threshold tuning to **drift monitoring** in production. Hired full-time within a year of joining as an intern for turning unstructured operational data into systems executives rely on.

EDUCATION

California Polytechnic State University, San Luis Obispo

B.S. Computer Engineering

San Luis Obispo, CA

Sept 2022 – June 2026

EXPERIENCE

Cal Poly — Environmental Health & Safety and Risk Management

San Luis Obispo, CA

Data Systems & Cloud Infrastructure Engineer, Cal Poly EHS & Risk Management

July 2026 – Present

- Lead technical contributor for the department's cloud-native transition under the **Amazon DxHub** program: built the **DynamoDB** data layer and **Amazon Bedrock** LLM services behind campus-wide safety assessment workflows.
- Deployed NLP classification (**AWS Comprehend**) that triages incoming safety data by risk in real time, paired with **custom LLM agents** that automate multi-step safety protocols; shipped the department's **React/Next.js** front ends end to end, including a password-gated fall-protection registry.
- Built an in-browser **document-intelligence** app (JavaScript/pdf.js) that parses State Fire Marshal inspection PDFs and auto-classifies violations across **campus buildings**, replacing manual tracking with a formula-driven compliance dashboard.

AI/ML Intern, Environmental Health & Safety (Cal Poly Corporation)

July 2025 – June 2026

- Architected a context-aware **RAG pipeline** (AWS Lambda, Bedrock/Claude) grounding an AI-powered Enterprise Risk Management portal, with **pgvector** hybrid search on Amazon RDS combining metadata filters with semantic retrieval.
- Owned the **model evaluation** loop end to end: curated labeled datasets, tuned decision thresholds via **precision/recall/F1** analysis, and served models behind REST endpoints with **drift monitoring**; productionized OCR document-understanding pipelines, eliminating hundreds of hours per term of manual entry.

Risk Management Intern

June 2025 – January 2026

- Operationalized AI governance: authored **NIST AI RMF** (ISO/IEC 23894) safeguards with human-in-the-loop review gates that new AI tools must meet before deployment.
- Led the Cal Maritime merger exposure analysis, isolating **\$700K+** in open claims, and built a dynamic residual-risk scoring system that sped up high-priority claim identification by up to **50%**.

AI Digital Transformation Intern

June 2024 – September 2024

Meera Services

Remote

- Re-architected a production billing-event pipeline (Python) with API-based validation and idempotent retries; integrated **PagerDuty** alerting and authored ops playbooks, cutting manual interventions **30%**.

PROJECTS

Cal Poly EHS Assistant, “Rudy”

July 2026 – Present

- Live, embeddable safety chatbot that answers only from a verified knowledge base; passes **301 adversarial checks** with zero hallucinated emergency numbers, running at **\$0/month**.

OFSM Inspection Tracker

July 2026 – Present

- Diagnosed a failed prototype and rebuilt it as a real app, replacing vision-model extraction with a deterministic parser once clean PDF text layers were found. Tracks the team's full inspection archive matching the master spreadsheet row for row, **32/32 tests** passing, **\$0/month**.

Shockwatch — Renewal Case-Control Layer

July 2026

- Built as forward-deployed-engineer proof of work for a YC-backed insurance brokerage: a fail-closed verification layer that binds every client-facing number to hash-checked evidence. **114 passing tests** including adversarial cases, in a synthetic-data sandbox.

The Sentinel Protocol: Cyber-Physical Integrity Verification (Capstone, team project)

Sept 2025 – June 2026

- Built with a team a containerized, edge-native security system detecting **sensor-spoofing attacks** on maritime navigation stacks; designed the **V&V harness** benchmarking detection robustness against simulated red-team attacks.

High-Altitude Telemetry Platform (Mesosphere) — OWL Integrations

Jan 2025 – June 2025

- Led a four-engineer team building resilient **LoRaWAN** mesh networking with a **YOLOv12** edge-AI camera; adaptive power strategies improved mission duration **~25%** and cut power consumption **~30%** in field tests.

TECHNICAL SKILLS

AI/LLM & Agents: LLM agents & multi-step agentic workflows, RAG & hybrid retrieval (pgvector), prompt engineering, AWS Bedrock & Comprehend, model evaluation & benchmarking (golden datasets, precision/recall/F1, drift monitoring), NIST AI RMF, adversarial/red-team testing

ML & Data: PyTorch, Scikit-Learn, unsupervised clustering, computer vision (YOLO), NLP, pandas/NumPy, Power BI

Cloud & Backend: AWS (Lambda, DynamoDB, RDS, Bedrock, Comprehend), Vercel, Supabase, Neon Postgres, Docker, Git, REST APIs, Microsoft Power Platform

Languages: Python, TypeScript, JavaScript, C, Java, SQL, System Verilog, HTML/CSS