Aditya Samalla

(480) 758-7754 · aditya.samalla@asu.edu · adityaron.github.io · linkedin.com/in/adityasamalla

Summary

Engineering manager and Principal ML Engineer with 10+ years building production AI. I lead LLM, agentic (MCP), and detection-ML work for a cloud-security platform: 40% fewer false positives, 40% more detections at 47% lower cost, and 4 patents (2 issued).

Core Competencies

- LLMs & Small Language Models, RAG, Agentic / MCP
- LLM Evaluation, Red-teaming & AI Safety
- Transfer Learning, Deep Learning, MLOps
- Anomaly & Behavior Detection, Graph ML
- Production ML Systems & Data Platforms
- Engineering Management, Roadmaps, Mentoring
- **Languages:** Python, SQL, Scala, Go, R
- **AI/ML:** Claude (Bedrock), Azure OpenAI, LangChain, PyTorch, TensorFlow
- **Data:** Spark, Apache Iceberg/Lakehouse, Snowflake, Postgres, Neo4j, Airflow
- **Infra:** Kubernetes, Docker, AWS/GCP/Azure, Argo CD, Grafana

Work Experience

Fortinet (formerly Lacework) **Aug 2022 – Present**

Software Engineering Manager / Principal Machine Learning Engineer *Sunnyvale, California*

Promoted Senior → Staff Machine Learning Engineer → current role (3 levels, 2022–2026)

- Manage and mentor a team across detection, GenAI, and platform; set technical direction for MCP auth and consent, the composite-alert redesign, and the SLM detectors; own the roadmap.
- Led four customer-facing launches in six months: AI Incident Report and Remediation tabs, click-to-filter alert exploration, the intrusion graph, and the MCP server.
- Founded and lead AI Assist: an LLM + RAG system on Claude (Amazon Bedrock) that turns cloud-security alerts into prioritized incident reports and grounded remediation, with an alert-context API for MCP / agentic workflows; **30% higher analyst engagement**.
- Designed and shipped the FortiCNAPP MCP server: customers query live alerts, compliance, and vulnerabilities from an AI assistant with per-user, read-only, revocable access and results that declare when data is partial.
- Built the LLMOps and AI-safety practice: LLM-as-judge evals, red-teaming (**5,000+ prompts**), a report-quality A/B harness that can fail a release, prompt caching in all five prod regions, production SLOs, and per-request cost observability.
- Founding engineer and primary architect of the correlation engine behind Composite Alerts: **400+ detections** across **140+ MITRE ATT&CK techniques** fused into single high-fidelity incidents, **cutting false positives 40%**; led the redesign of the alert experience (an intrusion graph showing the evidence behind every link, plus AI Incident Report and Remediation tabs).
- Re-architected the detection pipeline so adding a detection adds no query load: **40% more detections at 47% lower cost**, on top of **$1.5M+ in earlier annualized savings**.
- Designed an ephemeral graph-database microservice to surface attack paths: **recall +15%, latency 10× lower, $1M/year less warehouse cost**.
- Own 10+ production ML detectors, including two fine-tuned small language models: **+10% detection coverage**; shipped the customer-facing detection catalog.
- Built the human-in-the-loop labeling tool and weak-supervision pipeline behind the alert-scoring models; precision gains contributed to a **$1.2M enterprise deal**.
- Conceived and designed an LLM threat-brief pipeline with the security research team (source-checked MITRE tags, CISA KEV status, analyst review); its briefs power the MCP server's threat checks.

Salesforce **Jun 2021 – Aug 2022**

Senior Data Scientist *San Francisco, California*

- Built and deployed supervised models (AWS SageMaker) forecasting customer subscription renewals for Tableau Online, **reducing churn 1.5%** at product scale.
- Applied topic modeling to autonomously group unstructured community ideas and stakeholder feedback for feature ranking, **increasing enhancement velocity 10%**.
- Designed supervised models using SHAP analysis to assess and improve Tableau visualization performance and user experience.
- Built data pipelines and models ingesting public census data for the Tableau Data Equity Hub — Tableau Foundation's data-for-good initiative helping organizations use data to advance equity and social justice.

Hewlett Packard Enterprise **Jun 2017 – Jun 2021**

Senior Staff Data Scientist *San Jose, California*

- Built core ML for HPE InfoSight, HPE's AIOps platform, including a workload anomaly-detection engine over storage-performance metrics with visual performance timelines.
- Engineered a storage recommendation engine that diagnosed performance issues and prescribed fixes, **automatically resolving 86% of customer cases**.
- Built a virtual-environment recommendation system with cross-stack latency attribution (host, storage, network) that surfaces and prioritizes exceptions for triage, **resolving 54% of issues beyond the storage layer**; created the "Impact Score" to quantify severity.
- Applied dimensionality reduction and unsupervised learning to high-dimensional telemetry to find natural variability and summarize virtual environments into meaningful topologies; built diagnostic signatures with support, performance engineering, and product.
- Built capacity-forecasting models and a workload-based sizing tool that recommends HPE products for data-driven sales: **15% revenue and 10% sales-efficiency gains**; architected a lead-generation integration enabling **$3M+ in opportunities**.

Nimble Storage **Sep 2016 – May 2017**

Data Scientist *San Jose, California*

- Built statistical and machine-learning models to proactively predict storage issues and performance bottlenecks — part of Nimble's predictive-analytics work that became HPE InfoSight.
- Designed executive data-visualization dashboards (Tableau / R Shiny) and built Flask APIs feeding downstream systems — driving KPI tracking and data-driven decisions across marketing, sales, and pricing with cross-functional partners.
- Ran exploratory data analysis on operational, sales, and support metrics and automated the supporting data pipelines, **improving support satisfaction 12.5%**.

Accomplishments

- **Patent (Issued, USPTO):** US 12,627,687 — Generative AI-Based Multi-stage Composite Event Processing — co-inventor.
- **Patent (Granted):** US 9,065,952 B2 — robust determination of the significance of performance issues in IT environments. **Patent (Pending):** ML-based detection of compromised AWS credentials. **Patent (Filing Approved):** LLM-based triage of cloud alerts using historical context and customer-provided guidance.
- **Publication / blog:** "FortiCNAPP AI Assist: Accelerating the Investigation and Remediation of Cloud Threats" (Fortinet Community, 2026); "Performance Analysis of Noise Cancellation in Speech Signals Using LMS, FT-LMS and RLS Algorithms."
- Member, **Beta Gamma Sigma** — International Honor Society (highest scholastic recognition). Presented ML features at HPE Techcon, Discover, and storage conferences.

Education

W. P. Carey School of Business, Arizona State University — *M.S. in Analytics, GPA 4.00* **2016**

Jawaharlal Nehru Technological University — *B.Tech, Electronics & Communication Engineering* **2011**