

Abdelrahman Taeha

916-385-3911 | apply@abdelrahmantaeaha.com | abdelrahmantaeaha.com | linkedin.com/in/abdel-rahman-taeha | github.com/abdeltachass

Education

Georgia Institute of Technology

M.S. Computer Science (Machine Learning Concentration)

Atlanta, GA
Aug. 2026 – May 2028 (Expected)

California State University, Sacramento

B.S. Computer Science

Sacramento, CA
Aug. 2023 – Dec. 2025

Professional Experience

AI Research Engineer | Medical AI Research

August 2025 – Present

- Co-authored [SpineFairBench](#), a counterfactual vision-language model (VLM) benchmark auditing demographic bias in spinal radiology report generation.
 - Evaluated 7,996 paired radiographs across a frozen 9-model panel with radiologist validation, surfacing demographic recommendation drift across all evaluated models.
 - Trained a neutral baseline model and built data collection pipelines for 2 major open-access clinical repositories to isolate demographic bias in VLM report generation.
- Co-developed MedInsider, a FHIR-style benchmark measuring whether medical LLM agents preserve decision integrity when exposed to institutional and financial pressure.
 - Constructed 840 paired-twin scenarios across 5 regulatory failure families and 8 incentive conditions to isolate when and why agents deviate from compliant behavior.
 - Engineered simulated EHR tooling with deterministic action-log scoring, run auditing, and reproducibility manifests to evaluate tool-use agents across 9 medical vision-language models.

Software Engineer | DHCS

May 2025 – Present

- Automated monitoring and reporting pipelines across enterprise systems, reducing manual observability overhead by 40% and cutting mean time to detection (MTTD) for production incidents in half.
 - Transformed backend integrations and validation workflows that surfaced anomalies early, accelerating root-cause analysis by 35%.
- Established internal dashboards visualizing agent-driven operational metrics, reducing incident resolution time by 30%; improved real-time alerting accuracy by 20% through log-processing pipelines.

Software Engineering Intern (Machine Learning) | District Hut

May 2024 – July 2025

- Scaled distributed data pipelines processing 3 TB/week across cloud infrastructure, improving throughput and system reliability in a high-volume production environment.
 - Revamped SQL data-integrity checks, preventing 5 critical production data-corruption issues and automating 75% of validation tasks with 100% test coverage.
- Collaborated with 3 engineers to automate model retraining agents, reducing collective manual effort by 6 hours per week via autonomous performance metrics.

Project Highlights

SENTRY — Vision-Language Incident Report Generator — github.com/abdeltachass/SENTRY

- Developed and deployed an end-to-end vision-language system that generates security-incident reports from camera frames, pairing a CLIP-style image encoder with a LoRA-fine-tuned language model to adapt an open pretrained model efficiently rather than training from scratch.
- Quantified vision-language model hallucination rate against 10,000 manually reviewed outputs, reducing ungrounded claims by 35% and improving overall output reliability.
- Engineered per-sentence Grad-CAM-style grounding that visualizes exactly which pixels the model attended to for each generated claim; shipped as a live, interactive demo on Hugging Face Spaces.

S.A.I.N.T. — Network Intrusion Detection System — github.com/abdeltachass/S.A.I.N.T.

- Partnered with 2 peers to build an AI agent reasoning engine for threat detection, reaching 89% accuracy on NSL-KDD (125K+ samples) via feature engineering across 41 traffic-flow features.
- Designed a low-latency inference agent using Flask + Redis + Plotly, processing 50–100 network connections/sec with an interactive agent-in-the-loop dashboard for analyst review.

PingVault — Network Observability Agent — github.com/abdeltachass/PingVault

- Created a production-grade network monitoring tool spanning 3 protocols (ICMP, SNMP, TCP/UDP) with agent-based metrics collection, Alembic migrations, and a real-time HTML observability dashboard mirroring enterprise SRE tooling.

Skills

Languages: Python, Java, C/C++, SQL, JavaScript, Swift, MATLAB, R

ML & AI: PyTorch, TensorFlow, scikit-learn, pandas, NumPy, Model Evaluation

AI & Agents: Autonomous Agents, LLM Orchestration (LangChain/CrewAI), RAG, Prompt Engineering

Frameworks & Cloud: React.js, Node.js, Spring Boot, Flask, FastAPI, AWS, GCP, Docker, Git, GitHub Actions, CI/CD

Research & Domains: VLM Evaluation, Counterfactual Benchmarking, Medical Imaging, Demographic Bias Auditing, SIEM, Monitoring Pipelines