

Aaron Rassiq

Product-Minded AI / ML Engineer

+1 (925) 323-9638 | aaronrassiq@berkeley.edu | [linkedin.com/in/aaronrassiq](https://www.linkedin.com/in/aaronrassiq) | github.com/arassiq

EDUCATION

University of California, Berkeley

Master of Information and Data Science

Berkeley, CA

present

University of San Francisco, GPA: 3.81

Analytics, Computer Science

San Francisco, CA

Aug. 2021 – May 2025

WORK EXPERIENCE

Data and AI Engineering Intern

Feb 2026 - Present

Alation

Hybrid

- Conducted user interviews and workflow research with internal business intelligence teams, translating pain points into product requirements, evaluation criteria, and an AI roadmap before committing to system architecture; shipped tools supporting 150+ users.
- Designed production RAG and multi-agent systems on AWS Bedrock by decomposing workflows into deterministic services and model-driven decisions, incorporating grounded retrieval and tool calling for reliable enterprise automation.
- Built end-to-end data and retrieval pipelines across Salesforce, Snowflake, Slack, PostgreSQL, and FAISS, enabling AI agents to reason over current enterprise context and automate reporting workflows.
- Owned reliability and user experience across FastAPI, React, and TypeScript services, diagnosing model and retrieval failure modes while using AWS SQS, SNS, Fargate, and real-time streaming for resilient, low-latency execution.

Machine Learning/Cloud Engineer

Oct 2025 - Feb 2026

Mefferdi

Remote

- Architected an AWS-based data lakehouse (S3, Glue) to support scalable ETL workflows, enabling ingestion and schema design for evolving telemetry data across sensor and model pipelines.
- Built data and simulation pipelines to process batch telemetry streams and stress-test perception systems under varying environmental conditions, identifying failure modes and improving system reliability prior to deployment.

Machine Learning Engineer

June 2024 - Feb 2026

uGRIDD

San Francisco, CA

- Trained and fine-tuned computer vision models (YOLO, SAM) for segmentation and detection tasks, building end-to-end systems and iterating on performance in real-world deployments.
- Improved detection performance (precision/mAP) by 22% across lane, shoulder, and barrier classes through feature engineering, dataset balancing, and evaluation on real-world edge cases.
- Built data processing pipelines to convert model outputs into structured formats (JSON/CSV) for downstream analytics, integration, and queryable workflows.

Machine Learning Researcher

May 2024 - May 2025

University of San Francisco

San Francisco, CA

- Translated the need for customizable lane representations into a multimodal product prototype, combining instance segmentation with a text-based vectorizer for user-directed feature extraction.
- Designed a geospatial image-processing pipeline using OpenCV's watershed algorithm to segment 300+ satellite images, achieving approximately 90% accuracy in detecting regional vegetation loss.
- Reduced segmentation latency by 40% by moving performance-critical processing to a C++-accelerated backend, enabling scalable batch execution.

TECHNICAL SKILLS

Languages: Python, TypeScript, SQL, C++, Java, JavaScript, Node.js

Applied AI/ML: LLM Applications, RAG, Agentic Systems, Tool Calling, Prompt Engineering, Embeddings, Vector Search, Model Evaluation, Computer Vision, PyTorch, OpenCV, Hugging Face, scikit-learn

Product Engineering: User Research, Product Discovery, Requirements Definition, Rapid Prototyping, Success Metrics, Experimentation, Failure Analysis

Data & Backend: ETL/ELT Pipelines, Data Modeling, Schema Design, FastAPI, REST APIs, PostgreSQL, Snowflake, FAISS, Batch and Real-Time Systems

Cloud & Infrastructure: AWS (Bedrock, ECS, Fargate, S3, Glue, Lambda, SQS, SNS, DynamoDB), Docker, CI/CD, Google Cloud

Frontend & Tools: React, Git, Linux, Pandas, NumPy, Salesforce, LangChain, LangGraph, MCP