

Melvin Orichi Socana

melvinsorichi@gmail.com | 910.461.9318 | <https://www.linkedin.com/in/melvin-orichisocana/> | anmelus.github.io

PROFESSIONAL SUMMARY

Data & Machine Learning Platform Engineer (Stanford MS, EE) with expertise in architecting scalable ML pipelines and data infrastructure. Specialized in large language models, Airflow orchestration, BigQuery/Vertex AI automation, and infrastructure-as-code. Proven track record in modernizing legacy systems for production-grade ML performance.

EDUCATION

Stanford University, Stanford, CA

Graduated April 2024

Master of Science (MS): Electrical Engineering (Hardware & Software)

EXPERIENCE

Handshake, San Francisco, CA

January 2025 - Present

Machine Learning & Data Platform Engineer

- Architected a centralized LLM Orchestration Platform using Python and FastAPI, enabling asynchronous execution of batch LLM jobs via Pub/Sub. Reduced model onboarding time for internal teams by 30%.
- Increased Data Platform stability by 35% by leading the critical migration of complex Airflow DAGs (Composer to Astronomer) and building new remote execution/deployment workflows with robust CI/CD pipelines.
- Delivered foundational data infrastructure for advanced AI/ML use cases, including architecting BigQuery + Vertex AI batch prediction flows and LLM observability pipelines integrated with Datadog.

Amazon, Seattle, WA

May 2023 – August 2023

Software Development Engineer Intern

- Developed an end-to-end Album Management feature for Amazon Photos using C#/.NET, architecting the backend logic for photo grouping and data persistence.
- Designed and implemented a scalable SQLite schema and paginated queries to handle metadata storage, ensuring high-performance data retrieval from the initial launch.
- Authored a comprehensive Technical Design Document, navigating ambiguous requirements to gain consensus from cross-functional stakeholders (PMs, SDEs) on a final architecture.

Stanford NLP Research, Stanford, CA

November 2023 – February 2024

Graduate AI Research Engineer - Sequence-to-Sequence Semantic Parsing over Wikidata

- Developed a few-shot sequence-to-sequence semantic parser with retrieval-augmented generation (RAG), enabling natural language queries to generate accurate SPARQL queries.
- Enhanced entity-linking pipeline by 20%, leading to increased performance over the prior state-of-the-art.

Active Robotic Sensing Laboratory, Raleigh, NC

August 2021 – May 2022

Machine Learning & Robotics Research Engineer

- Led development on constructing the imaging system for the BenchBot, a robotics device to estimate biomass.
- Designed and implemented a robust script to interface with DepthAI cameras and servos, capturing high-quality plant images to train machine learning models.

TECHNICAL SKILLS

Programming Languages

Python, Go, C++, C#, TypeScript

Cloud & Infrastructure

GCP, Terraform, Docker, Kubernetes, Dataproc, Datadog, Pub/Sub

ML/AI

PyTorch, TensorFlow, NumPy, AnyScale, Vertex AI

Data Engineering

Apache Airflow, dbt, BigQuery, SQL