

ALEX DAVID ORDONEZ

AI ENGINEER | GENERATIVE AI | RAG | AGENTIC SYSTEMS | AWS

+1 (248) 222-9261 | alex.ordonez.eng@gmail.com | [Portfolio](#) | [LinkedIn](#) | [GitHub](#)

PROFESSIONAL SUMMARY

AI and Data Engineer with 5+ years of experience across data engineering, analytics, cloud technologies, and AI product delivery. Hands-on experience building LLM applications, RAG pipelines, multi-agent workflows, and NLP systems, with technical leadership spanning architecture, evaluation, and deployment. Built a multi-tenant AWS data lake supporting financial analytics and downstream AI and currently develops healthcare automation and enterprise data solutions. Combines a strong foundation in statistical modeling and machine learning with expertise in reliable data pipelines and scalable AI systems.

TECHNICAL SKILLS

Generative AI, LLMs & RAG	OpenAI GPT-5-mini and GPT-4.1-mini/nano, Google Gemini, Hugging Face Transformers, LangChain, LangGraph, RAG, embeddings, semantic/vector search, prompt engineering, grounded generation, structured extraction, multi-agent orchestration, LLM/agent evaluation
Machine Learning, NLP & Analytics	scikit-learn, TensorFlow, PyTorch, sentiment analysis, emotion classification, BERT/RoBERTa-based NLP, predictive/statistical modeling, classification, clustering, forecasting, time-series analysis, feature engineering, model evaluation; Power BI, Excel, Matplotlib, KPI dashboards, data visualization
Programming & AI Applications	Python, SQL, PySpark / Apache Spark, Pandas, NumPy, JavaScript, TypeScript, Java, R, SAS, C++, Bash; FastAPI, Flask, REST APIs, JSON, React, Angular / AngularJS, Streamlit
Vector Search & Data Platforms	MongoDB, MongoDB Atlas Vector Search, Qdrant, FAISS, PostgreSQL, DynamoDB, Oracle DB, SQL Server, Snowflake, Databricks; structured and semi-structured data
AWS, Cloud & AI Infrastructure	AWS Lambda, SQS, API Gateway, S3, Glue, Athena, Step Functions, EventBridge, SageMaker Unified Studio, CloudWatch, IAM, Lake Formation; GCP Cloud Run, Cloud Storage, Vertex AI
Data Engineering for AI	ETL/ELT, AWS data lakes, bronze/silver/gold layers, batch pipelines, web/document data extraction, ingestion, schema design/evolution, partitioning, Parquet, data quality/validation, mock/test data, model-ready datasets, data lineage
MLOps, DevOps & Evaluation	Docker, Git, GitHub Actions, CI/CD, pytest, unit/integration testing, prompt/retrieval evaluation, LangGraph agent orchestration, asynchronous processing, AWS Step Functions/EventBridge/SQS, API integrations, retry/failure handling, OpenLineage, Datadog, CloudWatch, Agile/Scrum

PROFESSIONAL EXPERIENCE

Data Systems & Digital Marketing Coordinator | Pediatric & Young Adult Medicine (PYAM) | MN | Aug 2026 - Present

- **Analyze legacy PDF medical intake forms**, extracting questions, answer types, validation rules, and age/visit logic into **structured digital form specifications**.
- **Build automated digital intake forms** for pediatric and young-adult workflows, routing content by **patient age and visit type** and standardizing captured responses.
- **Develop and maintain a modern clinic website** with responsive patient-facing navigation, clear content structure, and digital conversion paths.
- **Design Google Ads and Meta Ads campaigns** and analyze **website traffic, user interactions, conversions, and campaign performance** to optimize acquisition funnels and digital outreach.
- **Video Generation Pipeline:** Designed and implemented an AI-powered video production pipeline for PYAM, using Nano Banana to generate and refine visual assets before producing polished, visually consistent videos with Veo and Runway.

Data Operations Engineer | Kansul Holdings | United States (Remote) | Part-time | Jun 2026 - Present

- **Implemented Azure-based application and data solutions**, using PostgreSQL to support reliable operational workflows and analytics-ready data.
- **Built React web applications and Looker dashboards** to surface operational metrics and support data-driven decision-making.

Data Engineer | Linker Finance | United States (Remote) | May 2026 - Jul 2026

- Built and validated a **multi-tenant AWS data lake** using **Amazon S3, AWS Glue/Athena, Lambda, and Parquet**, organizing provider data into bronze, silver, and gold layers and creating a governed data foundation for **downstream analytics and an AI implementation**.
- Developed **Python ETL pipelines** for **Plaid, Sardine, Amplitude, and internal RDS sources**, standardizing schemas, transforming structured/semi-structured financial data, handling encrypted fields, and preparing partitioned, AI-ready datasets for efficient downstream consumption.
- Implemented **data-quality validation and automated pytest coverage** for ETL workflows, including required-column checks, schema validation, failure-path testing, and **mock datasets** that reproduced provider responses before AWS deployment.
- Created **tenant-aware FastAPI endpoints** deployed through **AWS Lambda and API Gateway** to expose curated lake data to backend applications and downstream AI services while supporting secure multi-tenant access patterns.
- Designed **deployment and orchestration patterns** using **AWS Step Functions and EventBridge**, and supported Glue Catalog/Lake Formation configuration to improve repeatability, discoverability, and governed access to data used by applications and AI workflows.
- Configured **SageMaker Unified Studio and OpenLineage** to improve data-lineage visibility, while building **CloudWatch/Datadog dashboards and monitors** for pipeline health, service usage, latency, errors, and operational performance.

Research Assistant - Data Science | Oakland University Business School | Michigan, US | Sep 2025 - May 2026

- **Medical Survey Analytics:** Built ML/LLM-driven pipelines for large-scale medical survey data, transforming unstructured responses into structured insights through preprocessing, feature engineering, statistical analysis, NLP/LLM-assisted interpretation, and visualization.
- **AMI Meter Forecasting & Grid Analytics:** Developed SAS/Python pipelines for 15-minute AMI electricity data, including ingestion and validation, time-series forecasting, phase-level analysis, anomaly/imbalance detection, scenario simulation, and LLM-assisted interpretation of analytical results.
- **Hazard Analysis:** Built scalable Python pipelines combining geospatial, temporal, and impact data to identify multi-hazard relationships, exposure and risk patterns, root-cause drivers, and critical-infrastructure risk; produced dashboards and decision-support outputs.

AI Engineer / Technical Lead | Product Manager Accelerator | Michigan, US | May 2025 - Sep 2025

- Served as **Technical Lead for AI product delivery**, translating business requirements into AI system architecture and Agile implementation tasks, coordinating contributors, reviewing technical outputs, and guiding integration, testing, and production handoff.
- Designed **RAG pipelines** using **MongoDB Atlas Vector Search and Qdrant**, embeddings, semantic retrieval, and grounded LLM generation to support context-aware responses and reusable enterprise knowledge access.
- Built **web/document data scraping, extraction, and ingestion workflows** that cleaned and transformed unstructured data into **structured JSON, searchable embeddings, and agent-ready context** for downstream AI applications and analytics.
- Implemented **multi-agent AI applications** with **LangChain/LangGraph, OpenAI, FastAPI, MongoDB, AWS Lambda/SQS, Docker, and GitHub Actions**, including asynchronous processing, API/tool integrations, structured outputs, prompt/retrieval evaluation, and production-oriented deployment workflows.

Data Analyst | Laureate Education | Honduras | Dec 2021 - Nov 2023

- Designed and maintained **ETL pipelines using Python, SQL, PostgreSQL, MongoDB, and DynamoDB** to integrate enterprise data for analytics, reporting, and downstream application use.
- Developed **Power BI dashboards and automated KPI reporting**, improving visibility into customer and operational performance by **30%**.
- Built a **CSAT data collection and reporting platform using Python and AngularJS**, enabling structured analysis of customer feedback and service-quality trends.

Data Engineer | Claro Enterprise Solutions | Honduras | Jan 2021 - Nov 2021

- Automated **ETL workflows for large-scale telecom datasets** using **Python, Java, SQL, and Oracle DB**, improving access to operational data for technical reporting and analytics.
- Built **data retrieval, transformation, validation, and reporting workflows** for high-volume structured datasets and supported query optimization for recurring analytical use cases.

SELECTED AI & MACHINE LEARNING PROJECTS

MyHandyAI - Multi-Agent DIY Repair Assistant | LangChain/LangGraph, FastAPI, AWS Lambda/SQS, MongoDB, RAG

- Led development of a **multimodal, multi-agent RAG assistant** combining image analysis, semantic retrieval, clarification, tool selection, safety guidance, repair planning, cost estimation, and context-specific image generation.

ClientSignalEQ - Customer Communication Risk & Coaching AI | GPT, Hugging Face, Embeddings, FastAPI, AWS, MongoDB

- Built an **AI early-warning and coaching platform** using GPT generation, BERT/roBERTa-family sentiment/emotion models, MiniLM embeddings, semantic retrieval, response-delay/escalation signals, and MongoDB to produce structured risk metrics and actionable coaching.

Analytics AI - Conversational Data Analyst Platform | GPT-5-mini, LangChain, FastAPI, React, MongoDB, GCP

- Built a **chat-driven analytics application** with GPT-5-mini, LangChain, FastAPI, React, MongoDB, Cloud Run, and Cloud Storage for natural-language EDA, forecasting, visualization, summaries, and persisted project context.

Databricks Sentiment Analysis Pipeline | Databricks, Spark/PySpark, Python, NLP

- Built a **Spark/PySpark NLP pipeline** for scalable text ingestion, preprocessing, sentiment analysis, and structured AI/ML outputs in Databricks.

EDUCATION

Master of Science in Business Analytics | Oakland University, Rochester, MI | Sep 2024 - Dec 2025

Bachelor's in Computer Systems Engineering | Universidad Tecnológica Centroamericana, Honduras | Jan 2017 - Mar 2022

CERTIFICATIONS & AWARDS

Databricks AI Agents Fundamentals

Databricks Fundamentals

AI Engineer Certification - Product Manager Accelerator

SAS Workbench & Sustainability Champion - SAS Hackathon

Languages: English (Fluent) | Spanish (Fluent)