

KEVIN COYLE
SOFTWARE ENGINEER - MACHINE LEARNING

kevincoyle.xyz | kevin.s.coyle@gmail.com

SUMMARY

I am a software engineer, and I typically focus on ML applications. I am comfortable with single-node and distributed applications. I have built and maintained machine learning and generative AI based software for Fortune 100 companies. My work includes data warehouse design and data warehouse migrations, deploying machine learning models at scale, and implementing bespoke applications.

SKILLS

Front End Software Stack:

JavaScript | React | NextJS | Husky | Tailwind | HTMX | Vite

Back End / Database Software Stack:

Python | NodeJS | ExpressJS | Apollo | Graph | PostgreSQL | neo4j | MongoDB | Redis | PineconeDB | PGVector | ChromaDB | FAISS
| Selenium | Flask

Analytics Software Stack:

Python | Scala | R | SQL | NumPy | Pandas | SciKit-Learn | SciPy | PyTorch | MLflow | Databricks | Apache Spark | Delta Lake | Feature Store | Snowflake | Amazon Redshift | Airflow | Tableau | Looker | Periscope | OpenAI | Llama 3

General Tech Stack: Docker | Amazon AWS Suite | EC2 | S3 | Git | Terraform | Kubernetes | Azure

Other Skills:

Salesforce Marketing Cloud | Salesforce Data Cloud | Microsoft Office Suite | Google Office Suite | Comfortable in Ubuntu, Mac OS, and Microsoft

PROFESSIONAL EXPERIENCE

Manager AI and Analytics

A&MPLIFY | Los Angeles, CA | 2023 - Present

Lead generative AI research and development. Lead Senior Data Scientists and Junior Data Scientists with thought-leadership, solution architecture, and code guidance. Notable projects: generative AI personalized recommender for hotel chain, resulting in 30% reduction in call center costs; dashboard built in Next.js with generative AI, deployed on Kubernetes with Terraform; IC for agriculture client - built a Typescript AWS application to track cross-breeding; AWS infrastructure rebuild for eCommerce company, creating \$750k in annual savings; internal initiatives including: technical coding environment for recruiting; thought pieces on LLMs.

Solution Architect

Databricks | San Francisco, CA | 2020 - 2023

Designed and deployed machine learning models and data pipelines using the Databricks platform with customer specs and requirements. Solved customer queries and issues with distributed and single-node machine learning and deep learning. Notable projects: end-to-end design of Data Engineering learning pipeline for Fortune 50 company. Optimized Apache Spark code to solve issues with skew and spill, saving several customers > \$1.2mn in compute costs.

Other notable relevant experience:

Senior Data Scientist

Capgemini | Burbank, CA | 2018 - 2020

Data Scientist

Alo Yoga | Commerce, CA | 2017 - 2018

Data Science Lead Instructor

General Assembly | Santa Monica, CA | 2017 - 2022 (part-time)

EDUCATION

General Assembly, Santa Monica, CA

Data Science Immersive

California State University, Sonoma, Rohnert Park, CA

Bachelor of Arts, International Relations - Political Science