

JULIAN MERCER

MACHINE LEARNING ENGINEER

San Francisco, CA | (415) 555-0148 | julian.mercer@email.com | linkedin.com/in/julianmerc

PROFESSIONAL SUMMARY

Machine Learning Engineer with 8+ years shipping ranking, forecasting, and language models into production. Designs reproducible training systems, low-latency inference, and monitored deployment pipelines that improve model quality without sacrificing reliability or cost control.

CORE SKILLS

Python | PyTorch | TensorFlow | Machine learning | MLOps | Model deployment | Feature stores | Spark | Kubernetes | MLflow | Airflow | Experimentation | Model monitoring | AWS | SQL

PROFESSIONAL EXPERIENCE

Senior Machine Learning Engineer Airbnb, San Francisco, CA

Jul 2022 - Present

- Led a learning-to-rank system for search retrieval serving 90M monthly queries, increasing booking conversion 6.8% in controlled experiments.
- Built feature pipelines on Spark and Feast for 340 signals, reducing training-serving skew incidents 73% across quarterly model releases.
- Deployed GPU inference services on Kubernetes, cutting p95 latency from 146 to 71 milliseconds while lowering compute cost 28%.
- Introduced shadow evaluation, drift alerts, and rollback gates, reducing model-related severity incidents from nine to two annually.
- Created offline-to-online evaluation standards covering bias, calibration, and segment performance for six production models.
- Mentored five engineers through design reviews and experimentation plans, increasing independently owned launches across the team.

Machine Learning Engineer Spotify, New York, NY

May 2019 - Jun 2022

- Developed sequence models for playlist continuation, increasing seven-day listening retention 3.9% across 18 markets.
- Automated distributed training and model registration with Airflow and MLflow, reducing release preparation from five days to six hours.
- Designed counterfactual evaluation for recommendation changes, preventing two launches that weakened long-tail artist discovery.
- Partnered with platform engineers on reusable deployment templates adopted by four recommendation teams.

Data Scientist Etsy, Brooklyn, NY

Aug 2016 - Apr 2019

- Built gradient-boosted fraud models across 42M listings, reducing manual review volume 31% at unchanged loss rates.
- Created reusable Python feature tooling and validation tests, shortening experiment setup 36% for seven analysts.
- Partnered with Trust Operations to tune thresholds by risk tier, improving high-risk case precision 14 points.

EDUCATION

M.S. in Computer Science, Georgia Institute of Technology, Atlanta, GA

CERTIFICATIONS

AWS Certified Machine Learning – Specialty | Databricks Machine Learning Professional