Trevor Hollis

Seattle, WA · trevor.hollis@example.com · (206) 555-0197 · linkedin.com/in/example-trevor-hollis

My Key Qualifications

- **5 years building and deploying production ML systems** for e-commerce and consumer platforms, from OfferUp in 2021 to Nordstrom today.
- Shipped **recommendation, ranking, and demand forecasting models** at Nordstrom, matching the core use cases this role owns.
- **Precision@10 up 11%** on a personalized product recommendation pipeline built with Spark features, MLflow, and offline evaluation on click and purchase history.
- Cut **SKU-level demand forecast MAPE from 18.4% to 14.9%** with Airflow-orchestrated XGBoost training jobs for seasonal categories.
- **Python, SQL, Spark, Airflow, Docker, Kubernetes, and AWS** (S3, EC2, EKS, Lambda, SageMaker) used daily across both roles.
- Ran a search-ranking **A/B test that lifted add-to-cart rate 3.2%**, owning offline evaluation, the containerized inference service, and experiment readouts.
- Built **model monitoring dashboards** tracking latency, throughput, and data drift, cutting failed-inference detection from 4 hours to 35 minutes.

Experience

Machine Learning Engineer — Nordstrom

Mar 2023 – Present · 3 yrs 8 mos

- Built and productionized a personalized product recommendation pipeline for logged-in web and app traffic, improving precision@10 by 11% with candidate-ranking models trained on Spark feature pipelines, MLflow experiment tracking, and offline evaluation on click and purchase history.
- Shipped a daily SKU-level demand forecasting workflow for seasonal categories, cutting MAPE from 18.4% to 14.9% using Airflow-orchestrated training jobs and XGBoost gradient-boosted models in Python and SQL.
- Refactored PySpark transformations and repartitioned large event tables on AWS S3 for a customer segmentation feature set used by lifecycle marketing and merchandising, cutting feature-refresh runtime from 95 minutes to 38 minutes.
- Implemented model monitoring dashboards for batch and near-real-time inference services, tracking latency, throughput, data drift, and job status in CloudWatch; mean time to detect failed inference jobs dropped from 4 hours to 35 minutes.
- Supported online experimentation for a search ranking update: defined offline evaluation criteria, shipped a containerized inference service on Docker and Kubernetes, and partnered with product on experiment readouts, measuring a 3.2% lift in add-to-cart rate in the winning A/B test variant.

Sample resume made with Specific Resume · specificresume.com

Data Scientist — **OfferUp**, a consumer marketplace platform

Nov 2021 – Feb 2023 · 1 yr 3 mos

- Built a buyer-seller engagement propensity model for lifecycle messaging in Python and scikit-learn, lifted targeted message click-through rate by 9%.
- Trained and calibrated XGBoost listing-quality scores used in search and moderation review queues, cutting manual review volume by 18 hours per week.
- Automated recurring experiment and model-performance reporting by scheduling SQL and Python jobs in Airflow, reducing analyst reporting turnaround from 2 days to 4 hours.

Education

- **University of Washington** — Master of Science in Information Management · 2019 – 2021

Skills

- **ML & Data Processing:** Python, SQL, Spark, PySpark, scikit-learn, XGBoost
- **MLOps & Orchestration:** Airflow, MLflow, Docker, Kubernetes
- **Cloud & AWS:** AWS, S3, EKS, SageMaker, CloudWatch
- **ML Domains:** Recommender Systems, Demand Forecasting

Sample resume made with Specific Resume · specificresume.com