

Mila Aguirre

Machine Learning Architect, Personalization

(206) 555-0148 mila.aguirre@example.com linkedin.com/example/milaaguirre Seattle, WA 98109

STRENGTHS

- Architecture Leadership**
Unified six recommendation surfaces through reusable interfaces; teams gained clearer ownership and faster technical decisions.
- Relevance Judgment**
Balanced recall, precision, NDCG, diversity, conversion, and AOV; product partners gained confidence in releases.
- Cross-Team Influence**
Architecture reviews aligned more than 40 contributors without formal authority; difficult tradeoffs became shared decisions.
- Production Focus**
Investigated serving issues alongside model teams; practical fixes kept latency, freshness, and reliability visible.
- Mentorship**
Research reviews and design critiques helped 12 engineers and scientists make stronger choices under pressure.

SKILLS

Recommender Systems Retrieval

Learning to Rank Reranking

Personalization

Reinforcement Learning

Cold-Start Modeling Python Java

SQL PyTorch TensorFlow

OpenSearch Vector Databases

Feature Stores

LANGUAGES

English Native

SUMMARY

Principal-level **Machine Learning Architect** with 16+ years across software engineering, machine learning, and **data science**, including 11+ years building **recommender systems** for commerce, media, and marketplace products. Brings hands-on leadership across retrieval, ranking, reranking, **real-time serving**, feature stores, vector databases, streaming pipelines, and low-latency APIs. Sets technical direction through architecture RFCs, technical reviews, platform migrations, research prototypes, code reviews, and production investigations. Connects engineering, data science, product, merchandising, and business priorities through measurable decisions involving recall, precision, NDCG, diversity, conversion, and AOV. Ready to help LoomCart make secondhand shopping easier through scalable, relevant, adaptable personalization.

EXPERIENCE

Principal Machine Learning Architect

Northstar Marketplaces January 2022 - Present Seattle, WA

- Owns architecture for personalized marketplace experiences across six surfaces, guiding engineers, data scientists, product leaders, and merchandising partners. Shapes reusable recommendation platforms, investment decisions, service standards, and production practices for an 18-million-item catalog.
- Established reusable **retrieval** interfaces across six marketplace surfaces, reducing fragmented recommendation **architecture**.
 - Authored **architecture RFCs** for feature ownership, deployment standards, and service-level objectives across 40+ partners.
 - Built OpenSearch candidate retrieval with vector search and streaming features for 18 million active items.
 - Created evaluation gates linking recall, **precision**, NDCG, diversity, **conversion**, and AOV for release decisions.
 - Migrated batch features toward continuous refresh, improving **cold-start** coverage for newly listed products and sellers.
 - Mentored engineers and scientists through design critiques, research reviews, incident analysis, and production debugging.

Staff Machine Learning Engineer

BrightLoop Media June 2018 - December 2021 New York, NY

- Led technical direction for personalized discovery across web, mobile, and partner applications. Coordinated retrieval, ranking, experimentation, feature platforms, and stakeholder decisions across engineering, data science, product, and commercial groups.
- Coordinated retrieval and **ranking** services supporting personalized discovery across web, mobile, and partner applications.
 - Built two-stage retrieval and **reranking** with learning-to-rank models, nearest-neighbor search, and online features.
 - Connected NDCG and **recall** changes with completion rate, conversion, and AOV through online experimentation.
 - Migrated teams toward shared feature definitions and **streaming pipelines**, making freshness visible across products.
 - Converted **catalog** quality, **merchandising**, and revenue goals into model requirements and service interfaces.
 - Mentored 12 engineers and scientists through architecture reviews, incident retrospectives, and implementation planning.

Senior Data Scientist, Recommendations

HarborNine Commerce January 2015 - May 2018 Austin, TX

Developed recommendation models and evaluation systems for a growing commerce catalog. Partnered with product and merchandising teams on experiments, relevance tradeoffs, inventory constraints, and measurable commercial outcomes.

MY CAREER



- Principal Machine Learning Architect at Northstar Marketplaces (4.7 Years)
- Staff Machine Learning Engineer at BrightLoop Media (3.5 Years)
- Senior Data Scientist, Recommendations at HarborNine Commerce (3.3 Years)
- Machine Learning Engineer at VectorPath Systems (2.4 Years)
- Software Engineer at Cedar Ridge Software (1.9 Years)

- Developed personalized **retrieval** and **ranking** with implicit feedback, item attributes, session context, and seller signals.
- Built reusable datasets monitoring precision, recall, **NDCG**, coverage, and **diversity** across model releases.
- Implemented inventory-aware **reranking** with category constraints, price bands, and merchandising priorities.
- Designed experiments linking recommendation quality with conversion, basket composition, and average order value.
- Prototyped **Kafka**-based feature generation and service caching for recent views, searches, and purchases.
- Presented model tradeoffs to technical and business audiences, shaping sequenced delivery plans.

Machine Learning Engineer

VectorPath Systems 📅 July 2012 - December 2014 📍 Chicago, IL

Implemented recommendation pipelines and retrieval services for retail and media clients. Supported model evaluation, OpenSearch integration, deployment practices, data quality, and client-facing technical delivery.

- Implemented batch recommendation pipelines with collaborative filtering, content features, and scheduled model scoring.
- Built **retrieval** evaluation tools comparing candidate-generation strategies against precision, recall, and **ranking** benchmarks.
- Integrated **OpenSearch** indexes with **Python** services for fast item lookup and catalog-aware experiments.
- Added validation and lineage checks across clickstream, catalog, and transaction feeds before model jobs.
- Containerized inference components with documented deployment, rollback, and troubleshooting procedures.
- Delivered client demonstrations clarifying **cold-start** behavior, system limits, and expected business outcomes.

Software Engineer

Cedar Ridge Software 📅 July 2010 - June 2012 📍 Raleigh, NC

Built Java and Python services for search, catalog metadata, customer activity, and distributed commerce workflows. Contributed testing, ingestion, relevance, defect investigation, and an early personalized browse prototype.

- Developed **Java** and Python services for search, catalog metadata, and customer activity processing.
- Created batch ingestion jobs with schema validation and retry handling for dependable analytics feeds.
- Improved search relevance through **OpenSearch** analyzers, filters, and query templates.
- Built API tests and performance checks, documenting latency findings for senior reviewers.
- Investigated production defects with QA, product, and operations partners, translating workflows into maintainable fixes.
- Presented a personalized browse prototype that launched subsequent recommendation research and experimentation.

PROJECTS

Real-Time Resale Personalization Platform 📅 2025

Created an adaptive recommendation concept for resale catalogs, combining streaming behavior, vector retrieval, reranking, cold-start signals, and low-latency serving. Focused decisions on relevance, diversity, conversion, and AOV.

Open-Source Ranking Toolkit 📅 2019

Maintained reusable ranking components and evaluation utilities for retrieval research. Contributions supported clearer comparisons across recall, precision, NDCG, and reranking approaches.

LEADERSHIP & AWARDS

- ACM RecSys Industry Impact Recognition, 2023
- Internal Distinguished Technical Leadership Award, 2021
- Open-source Contributor Recognition for Ranking Toolkit, 2019

EDUCATION

Master of Science in Computer Science

North Carolina State University 🎓 GPA: 4.0 📅 2010 📍 Raleigh, NC

Coursework: Machine Learning, Distributed Systems, Information Retrieval, Data Mining, Software Engineering

CERTIFICATIONS

- AWS Certified Machine Learning - Specialty 📅 2024
- Google Professional Machine Learning Engineer 📅 2025

TECHNICAL SKILLS

- **Recommendation Models:** Collaborative Filtering, Learning to Rank, Reranking
- **Search Platforms:** OpenSearch, Elasticsearch, Query Templates
- **Machine Learning Frameworks:** PyTorch, TensorFlow, Scikit-learn
- **Programming Languages:** Python, Java, SQL
- **Feature Infrastructure:** Feature Stores, Online Features, Feature Engineering
- **Vector Retrieval:** Vector Databases, Approximate Nearest Neighbor, Embeddings
- **Streaming Systems:** Kafka, Streaming Pipelines, Event Processing
- **Distributed Systems:** Distributed Services, High Throughput, Scalability
- **Serving Systems:** Real-Time Serving, Low-Latency APIs, Model Deployment
- **Evaluation Methods:** Recall, Precision, NDCG
- **Business Metrics:** Conversion, Average Order Value, Diversity
- **Cloud Platforms:** AWS, Containers, Cloud Services
- **Experimentation:** A/B Testing, Online Experiments, Offline Evaluation
- **Architecture Practice:** Architecture RFCs, Technical Reviews, Platform Migrations
- **Data Operations:** Clickstream Data, Catalog Data, Data Lineage

PROFESSIONAL AFFILIATIONS

- ACM RecSys Conference Volunteer Reviewer, 2022 - Present
- Women in Machine Learning Mentor, 2018 - Present
- Open-source Ranking Systems Maintainer, 2019 - Present

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST