

PARKER MOESTA

Senior Machine Learning Scientist

Ann Arbor, MI 48912 | parkermoesta@gmail.com | (313) 397-5184 | parkermoesta.ai

PROFESSIONAL SUMMARY

Senior Machine Learning Scientist with 5+ years of industry experience applying machine learning, data science, and production engineering to high-impact business problems. Expertise spans recommender systems, deep learning, large-scale feature engineering, experimentation, A/B testing, and causal inference, with a track record of delivering models and ML-driven systems tied to \$100M+ in measurable incremental revenue. Owns end-to-end ML development, from research to distributed data pipelines to production architecture, model deployment, and low-latency serving systems.

PROFESSIONAL EXPERIENCE

Domino's Pizza Ann Arbor, MI

Senior Machine Learning Scientist Apr 2026 – Present

- Lead applied ML and production architecture for eCommerce personalization, translating research ideas into scalable ML systems and model-driven digital experiences across web and mobile channels.
- Design low-latency ML service architecture for real-time recommendation serving, including model interface patterns, feature access strategies, logging requirements, and deployment considerations for production software ecosystems.
- Architect and build end-to-end recsys components spanning candidate generation, ranking/scoring workflows, feature integration patterns, and offline-to-online evaluation pipelines.
- Design production rollout plans for online A/B testing, including experiment setup, detailed logging requirements, model feedback loops, and iterative optimization workflows.
- Partner with data engineering, software engineering, and product stakeholders to define production-ready ML patterns that support reliable deployment, monitoring, and iteration across digital channels.

Data Scientist III, Personalization Jan 2025 – Mar 2026

- Led end-to-end development and productionization of the Customer-360 Feature Store, a distributed Spark pipeline processing 1B+ records over rolling one-year behavior windows on a weekly cadence and generating 400+ features across behavioral, transactional, temporal, coupon, loyalty, and product-affinity dimensions
- Established reusable feature infrastructure powering 12+ production models and analytics workflows, improving consistency, scalability, and speed-to-deployment for downstream ML initiatives.
- Built and iterated on the company's first production recsys models, advancing from association-rule mining and collaborative filtering toward multi-tower architectures within a two-stage recommendation framework
- Developed reusable training, scoring, and audience-generation pipelines for personalized marketing and product recommendation use cases, operationalizing model outputs for campaign activation and segmented personalization.
- Led early R&D for a transformer-based customer sequence model, pretraining GPT-style architecture on order histories for next-order prediction and personalized recommendation.
- Became a Data Science subject-matter expert for scalable Spark pipelines, distributed feature engineering, model development patterns, and production-ready ML workflows.
- Delivered production ML models attributing to \$30M in annual incremental revenue through improved targeting and personalized experiences

Data Scientist, Digital Experience / eCommerce Dec 2023 – Jan 2025

- Led analysis for 12+ high-impact digital experiments across web and mobile channels, applying A/B testing, PSM, DiD, and synthetic control methods to evaluate product, marketing, and customer-experience initiatives.
- Developed a deep-learning cart-abandonment prediction model using LSTM networks in PyTorch on Adobe clickstream data, driving a 6% reduction in abandonment through timely automated customer interventions validated via holdout control groups
- Translated complex causal inference findings into executive-ready recommendations, informing national business decisions with measurable financial impact and no observed degradation in customer retention.
- Refactored legacy A/B testing instrumentation into a productionized CI/CD workflow, improving reliability, repeatability, and scalability of analytics delivery.
- Podium finisher (3rd place) in the Databricks Global GenAI Hackathon; project featured in a Databricks technical blog

Coupa Software Ann Arbor, MI

Data Scientist, Sales Strategy & Optimization 2020 – 2023

- Engineered an ML-powered product offering recsys using collaborative filtering (SVD matrix factorization) to surface personalized add-on suggestions, contributing to a \$50M increase in incremental add-on revenue
- Designed and prototyped an internal Q&A chatbot using LangChain and LlamaIndex integrated with the ChatGPT API, implementing a RAG architecture for intelligent document retrieval and context-aware responses
- Built a Customer Whitespace Opportunity Scoring model combining ML predictions with business heuristics, deployed via an interactive Tableau dashboard enabling sales teams to identify \$20M+ in untapped revenue opportunities
- Partnered with sales, strategy, and executive stakeholders to translate model outputs into account-planning workflows, renewal strategy, and expansion recommendations.

Oracle Boston, MA

Enterprise Account Manager 2016 – 2018

- Led end-to-end POC implementation of an Oracle Digital Assistant chatbot for a major retail partner, architecting an internal call-deflection solution that reduced support-center volume by 25% and improved employee self-service
- Cultivated strategic C-level relationships with Fortune 500 enterprise clients, driving contract renewals and expansion through data-driven insights and consultative engagement

EDUCATION

University of Michigan, Ann Arbor 2021 – 2024

M.S., Data Science | GPA: 3.95 / 4.0

University of Michigan, Ann Arbor 2012 – 2016

B.A., Economics

TECHNICAL SKILLS

Languages & Engineering: Python, SQL, PySpark, Git, Docker, CI/CD, unit testing, integration testing, APIs, production coding practices

ML / AI: Recommender Systems, Deep Learning, NLP, Transformers, Embeddings, RAG, Causal Inference, A/B Testing, Time Series, Bayesian Methods, Statistical Modeling

MLOps & Data Platforms: Databricks, MLflow, Airflow, AWS, GCP, Feature Stores, Model Serving, Distributed Data Pipelines

Libraries: PyTorch, TensorFlow, scikit-learn, HuggingFace, LangChain, LlamaIndex, Pandas, NumPy, statsmodels, EconML, spaCy

Analytics & Visualization: Tableau, Jupyter, Experimentation Analytics, Executive Storytelling

PUBLICATIONS

Pethers, C., Ruesga, M., Moesta, P., & Williams, E. J. (2021). *Simulation Improves Service and Operations at a Franchise Resale Store*. Proceedings of the 4th North American International Conference on Industrial Engineering and Operations Management (IEOM), Monterrey, Mexico.