

Ryan Russon

Senior Manager, Machine Learning Engineering

PROFILE

AWS Machine Learning Certified data scientist, engineer, and former Naval officer with experience in data science and engineering design, encompassing a large array of problems with proven results. Enthused about leveraging data and web technologies to provide custom and automated predictive analytics or data delivery for customers. A dedicated leader and an excellent communicator with an ability to manage and convey project-critical information to direct reports, peers, and high-level executives. Highly motivated and goal-oriented with skills in Python programming, TensorFlow, MLOps, data visualization, predictive modeling, deep learning, problem identification, root cause analysis, and process improvement.

PROFESSIONAL EXPERIENCE

Senior Manager, Machine Learning Engineering

07/2022 - Present

Capital One

Manage disparate across intents, ranging from large-scale, model development to smaller-scale, exploration environments for thousands of Capital One scientists, engineers, and analysts.

- Owns the premier platform for orchestrating large-scale model training jobs, built on Kubernetes and OSS with various distributed compute demands including Spark, Dask, and PyTorch (GPU). The platform scales to thousands of distributed nodes across different clusters.
- Leads intent to improve MLOps in a highly regulated financial environment including vulnerability remediation, cluster access, and security while maintaining a smooth UX for model developers.
- Delivers sandbox environments with guardrails to allow internal customers to experiment with the latest advances in LLMs.

Manager, Data Science and MLOps

11/2020 - 07/2022

Maven Wave Partners

- Led the implementation of new training and serving solutions for a global bank utilizing Kubeflow on AWS, including custom SDKs, CI/CD, and testing procedures for production model training pipelines on terabytes of data for real-time, fraud detection models (XGBoost).
- Implemented and designed a fuzzy matching model (Decision Tree + Connected Components) and supporting platform architecture for record deduplication at scale on AWS. Utilized various AWS services such as Glue, Lambda, DynamoDB, RDS, SNS, and SQS to resolve matches against billions of records, reducing execution time by 96%.
- Built an end-to-end model training platform for traceability and repeatability on terabytes of data for a delivery forecasting system on GCP, using Kubeflow Pipelines to orchestrate Vertex AI, BigQuery, and GCS with Stackdriver logging. Reduced model training time by 80% and improved model prediction performance by 4%-10% over the client model.
- Delivered POC for topic modeling solution that leveraged Streamlit, AWS Transcribe, SageMaker Notebooks, and S3 to derive major topics for call-center conversations to inform leadership on recurring and high-profile help desk services rendered.

Senior Data Scientist

07/2015 - 11/2020

Micron Technology / Intel Corp

- Developed deep learning models (VGG-16, ResNet-152, transfer learning) for anomaly detection with automated reactions, saving thousands of dollars per week in sourcing and preventing quality issues.
- Led cross-functional teams to provide an in-house, custom web interface that allowed for quick image classification before model building (human-in-the-loop), enabling thousands of images to be manually classified in hours instead of days.
- Combined various relational data sources (MSSQL/Oracle/Postgres/Hive) and NoSQL sources (Spark/HDFS/HBase) to provide new insights and analytics from manufacturing systems.
- Created a custom web service in Django that solves constrained optimization for automatic process control models reducing human error by 98% and saving over +\$1M per year in human-caused product mistakes.
- Managed several ETL jobs transferring data from various Micron and Intel worldwide SQL and NoSQL data sources to local and private cloud targets which improved fab engineer data accessibility and cleansing tasks by over 300%.
- Built a custom web interface using PHP/Python/JavaScript to source wafer defects saving 60 engineering hours per month. Allowed users to create and submit landmark image data to Apache HBase that was later used in conjunction with a bespoke search algorithm.
- Developed methodology to derive process control models from inline data saving hundreds of hours of tool time and thousands of dollars in resources over traditional Design-of-Experiment methods.

Technical Training Program Manager and Instructor

05/2011 - 07/2015

US Navy

- Managed the command's internal audit and continuous improvement program comprised of more than 60 individual programs, 30 divisions, and 4,000 personnel and provided oversight for all technical curricula and administrative procedures. Developed new training programs and policies.
- Designated Navy Master Training Specialist (MTS), led an advanced course on theory and design of the complex electrical and mechanical support systems for nuclear power plant operation for more than 300 personnel with a 97% course completion rate.

Undergraduate Research Assistant

01/2009 - 05/2011

University of Utah

- Applied statistical methods (i.e. partial least squares, bootstrapping, and ANOVA) to predict the onset of AD from hundreds of MRI scans and cognitive tests.

EDUCATION

M.S., Electrical and Computer Engineering

Purdue University

Automatic Control (MPC, Fuzzy Logic, PID), Project Management, System Optimization

B.S., Biomedical Engineering

University of Utah

Digital Image Processing, Engineering Design, Statistics

LINKS

- [Personal Site](#)
- [Media](#)
- [GitHub](#)
- [LinkedIn](#)

CERTIFICATIONS

AWS Certified Machine Learning – Specialty

SKILLS

Python

MLOps

Kubernetes

Kubeflow

Amazon Web Services

Predictive Modeling

Spark

Google Cloud Platform

TensorFlow, Scikit-Learn, PyTorch

OpenCV