John Hodge

jah70@vt.edu | GitHub | Google Scholar | LinkedIn

Summary

Senior Applied Scientist building machine-learning and statistical decision systems for complex physical and compute infrastructure. Developed production multi-source telemetry fusion and repair recommendation systems for AWS Core Compute hardware, combining logs, telemetry, repair history, calibration, and decision-system design to improve post-failure diagnostics, repair yield, and part-replacement decisions. Background spans AWS-scale hardware diagnostics, advanced-accelerator fleet repair recommendations, reliability analysis, multimodal sensing for Just Walk Out, electromagnetics, computer vision, and production applied science. Suited to Senior or Principal-track roles in AI infrastructure, high-value hardware, robotics infrastructure, sensing, reliability, and ML systems where modeling, experimentation, and decision-system design drive measurable production outcomes.

Technical Skills

AI Infrastructure & Hardware Systems: Multi-source telemetry fusion, hardware diagnostics, reliability modeling, repair recommendation systems, advanced-accelerator fleet repair analysis, temporal feature engineering, production ML decision systems, evaluation & benchmarking
ML Techniques: Model calibration, leakage-aware evaluation, ranking & recommender systems, ensemble methods, time-series/log modeling, deep learning, CNNs, vision transformers (ViT), attention-based transformers, GANs, Bayesian inference, distributed training experimentation (PyTorch DDP, JAX SPMD, XLA/NCCL)
NLP & Text ML: Production DistilBERT fine-tuning, BERT-family transformers, log/text modeling, RAG prototyping, prompt engineering
Languages: Python, SQL, Matlab
Libraries & Frameworks: NumPy, SciPy, Pandas, Polars, Scikit-Learn, PyTorch, Hugging Face Transformers, AutoGluon, Matplotlib, Numba, JAX, Flax (NNX), Optax
AWS & Infrastructure: S3, SageMaker, EC2, Athena, Redshift, DynamoDB, Docker, MLflow, Git

Experience

Senior Applied Scientist May 2024 – Present
Amazon Web Services (EC2 Infrastructure Science) *Seattle, WA*

- Own production science strategy for AWS Core Compute hardware diagnostics and repair recommendations, translating ambiguous fleet-recovery constraints into ML decision systems that improve yield, reduce unnecessary part replacement, and scale automated remediation.
- Built production multi-source telemetry fusion pipelines for post-failure hardware component diagnostics, aligning system event logs, memory diagnostics, crash reports, lifecycle events, host snapshots, and repair history into component-level recommendations.
- Architected and deployed production ML models for component diagnosis, slot/part localization, and failure clustering; drove production server repairs at fleet scale, including advanced-accelerator fleet repairs.
- Established leakage-aware temporal evaluation and production calibration for noisy repair-outcome data, including label-taxonomy refinement, component-specific confidence thresholds, and routing policies; improved repair success over the prior baseline across a large set of definitive repair outcomes.
- Authored strategic repair-funnel analysis across the EC2 unsellable-ticket pool, narrowing to the hardware-repair-relevant subset and identifying coverage, part-location, and repair-effectiveness bottlenecks; shaped the 2026 roadmap and quantified a material estimated annualized free-cash-flow impact from production repair recommendations.
- Built an end-to-end Monte Carlo reliability analysis for EC2 Live Update failover planning, modeling deployment schedules, notice-policy semantics, and multi-host workload exposure; delivered methodology, results, and executive-facing artifacts that supported mitigation planning for high-availability customer workloads.
- Drove roadmap alignment across repair operations, hardware engineering, software engineering, and product stakeholders; designed a production stratified A/B testing framework for causal evaluation of ML routing and remediation changes.

Senior Applied Scientist December 2021 – May 2024
Amazon (AWS Just Walk Out Technology) *Seattle, WA*

- Led capacitive sensing science for AWS Just Walk Out and served as senior scientist on RFID-enabled JWO, designing multimodal event-detection algorithms for capacitive shelf sensors and RFID retail/stadium deployments.

Senior Principal RF Engineer August 2014 – December 2021

Northrop Grumman Mission Systems *Baltimore, MD*

- Engineered and optimized a wideband phased array antenna system for a Navy EA-18G program as RF technology lead, validating performance across 150k+ data points.
- Applied machine learning and computer vision to automatic target recognition (ATR) on synthetic aperture radar (SAR) imagery; contributed RF systems expertise across modeling, analysis, and design optimization workflows.
- Led systems model-based engineering efforts and authored proposal sections of a successful $33M engineering contract. Promoted twice: RF Engineer to Principal (2018) and Senior Principal (2020).

Selected AI Infrastructure & Systems

Multi-Source Telemetry Fusion for Hardware Diagnostics | *Production ML & Decision Systems*

- Built a post-failure component diagnostic system that fuses heterogeneous telemetry, logs, lifecycle events, host snapshots, and repair history into calibrated repair recommendations.
- Designed leakage-aware temporal evaluation and confidence-tiered deployment logic, separating high-confidence automated repair actions from uncertain cases requiring additional diagnosis.

EC2 Reliability / Availability Simulation | *Monte Carlo Decision Modeling*

- Modeled operational overlap and workload exposure for EC2 infrastructure maintenance planning, translating schedule, notice-policy, and multi-host dependency assumptions into decision-ready reliability analysis.

Multi-GPU Scaling Study | *PyTorch DDP, JAX SPMD, XLA/NCCL*

- Benchmarked single-node training on an 8-GPU V100 NVLink node, measuring throughput, scaling efficiency, memory use, and accuracy tradeoffs across 1/2/4/8 GPU configurations.

Just Walk Out Sensing | *Capacitive Science Lead & RFID Senior Scientist*

- Led capacitive algorithm design for a thin shelf-sensor pipeline, fusing multi-channel time-series signals into per-lane product hypotheses using CUSUM detection and Bayesian product-capacitance learning; achieved high action/quantity classification accuracy on lab data.
- Drove science roadmap alignment across research, hardware, firmware, and ML stakeholders, identifying open risks in lane mapping, confidence scoring, environmental compensation, and cross-talk.

Selected Publications

- *Index-modulated metasurface transceiver design using reconfigurable intelligent surfaces for 6G wireless networks* (IEEE Journal of Selected Topics in Signal Processing, 2023).
- *RF metasurface array design using deep convolutional generative adversarial networks* (IEEE PAST, 2019): generative CNN/GAN-based design optimization.
- *Multi-discriminator distributed generative model for multi-layer RF metasurface discovery* (IEEE GlobalSIP, 2019): novel multi-discriminator GAN architecture.
- *Performance analysis of spatial and frequency domain index-modulated reconfigurable intelligent metasurfaces* (IEEE ICASSP, 2021).
- *Intelligent time-varying metasurface transceiver for index modulation in 6G wireless networks* (IEEE AWPL, 2019).

Education

Virginia Tech Falls Church, VA

Ph.D. in Electrical Engineering (Generative Deep Learning for Metasurface Design) *December 2021*

Virginia Tech Falls Church, VA

Master of Science (M.S.) in Electrical Engineering *May 2014*

Duke University Durham, NC

Bachelor of Science (B.S.) in Electrical & Computer Engineering with Dual Degree in Physics *May 2012*