

Javon Kitson

Washington, DC Metro (Germantown, MD) | (302) 312-8594 | javonkitson@gmail.com

[linkedin.com/in/javonkitson](https://www.linkedin.com/in/javonkitson) | javonkitson.com | github.com/noface-0

Experience

NVIDIA

AI/ML HPC Cluster Engineer

Seattle, WA

May 2026 – Present

AeroVironment

Research Engineer III

Germantown, MD

May 2025 – May 2026

- Architect and lead administrator for AI/ML GPU cluster serving 100+ researchers, optimizing SLURM scheduling and resource quotas to achieve 95%+ utilization across distributed training workloads.
- Engineered multi-node training pipelines and deployed reproducible YOLOX benchmarking suite across 48 A6000 GPUs, enabling data-driven capacity planning.
- Built enterprise-scale dataset/model distribution system managing 500TB+ of training data, reducing model deployment time and accelerating new researcher onboarding from weeks to days.

BlueHalo (Acquired by AeroVironment)

Research Engineer III

Germantown, MD

February 2023 – May 2025

- Administered 48-GPU AI/ML cluster serving 100+ researchers, implementing Prometheus/Grafana observability stack and optimizing SLURM to maintain 30-100% utilization based on workload demands.
- Developed standardized distributed training templates adopted by 10+ teams, reducing experiment setup time by days to weeks and enabling seamless multi-GPU scaling across the full cluster infrastructure.
- Fine-tuned and evaluated deep learning models using MLflow experiment tracking, leveraging the full cluster for significant speedups before containerizing and delivering to clients.

Intelligent Automation Inc. (Acquired by BlueHalo)

Software Engineer I/II

Rockville, MD

July 2020 – February 2023

- Delivered full-stack features in React and Python APIs (FastAPI), owning UI components, REST endpoints, and data contracts used by multiple internal teams.
- Built GitLab CI pipelines with unit, integration, and end-to-end tests, linters, and static security scans to improve release reliability.

Projects

Geoestimation (Python)

- Implemented a custom data pipeline to normalize, resample, and preprocess large satellite images for training.
- Developed and trained an EfficientNetB7-based UNet model in TensorFlow to accurately extract building footprints from RGB-PanSharpen imagery.
- Implemented automated area/perimeter calculation system using contour analysis computing real-world building dimensions (m²/meters) with sub-meter accuracy and visualization overlays.

Simple-Neural Architecture Search (Python)

- Engineered a neural architecture search framework implementing ENAS, and a hybrid PNAS+ENAS approach.
- Built modular architecture space supporting 8 neural network types (CNN, MLP, ResNet, MobileNet, DenseNet, ShuffleNetV2, EfficientNet) with dynamic layer configurations and visualization tools.
- Created interactive Streamlit interface for visualizing search results and architecture performance metrics across the complete NAS pipeline.

Deep Reinforcement Learning Stock Trading (Python)

- Built an ensemble trading system combining Genetic Algorithm portfolio selection, an FNN price predictor, and DRL agents (SAC, PPO), trained on First Rate Data's multi-year 10,120 ticker dataset.
- In paper-trading backtests, PPO achieved +52% ROI and SAC +89% ROI vs initial capital.

Skills

Python, C, Java, JavaScript/TypeScript | PyTorch, TensorFlow, Transformers, LoRA/QLoRA, MLflow | torchrun/DDP, NCCL, SLURM | FastAPI, S3, NGINX, MongoDB, SQL | AWS, Docker, Kubernetes | Prometheus, Grafana | React, Tailwind CSS

Education

University of San Diego, San Diego, CA

Master of Science in Applied Artificial Intelligence, Graduated May 2024

Loyola University Maryland, Baltimore, MD

Bachelor of Arts in Computer Science, Minor Biomedical Physics, Graduated May 2020