

Aryan Pandey

716-770-8687 | aryanpan@buffalo.edu | [linkedin.com/in/aryanpan](https://www.linkedin.com/in/aryanpan)

Research Interests: Hardware-efficient machine learning, ML–hardware co-design, in-memory computing, emerging compute architectures.

EDUCATION

University at Buffalo, State University of New York Buffalo, NY
Ph.D. in Neuromorphic Computing Jan 2025 – Present

University at Buffalo, State University of New York Buffalo, NY
B.S. in Computer Engineering (GPA: 3.7, Dean's List) Aug 2018 – May 2022

RESEARCH

Recurrent Differentiable Weightless Networks (RDWN) | *PyTorch, Verilog, Vivado* Feb 2026 – Present

- Added recurrence to DWNs, a LUT-only network family that replaces MACs with table lookups for orders-of-magnitude cheaper FPGA inference than quantized DNNs.
- On Spiking Heidelberg Digits, the recurrent variant reached **82.2%** test accuracy vs. **67%** for the feedforward baseline; **97.75%** on N-MNIST.
- Built PyTorch → Verilog export flow; demonstrated SHD model end-to-end on FPGA at 0.6W.

Hardware-Efficient RF Device Fingerprinting | *PyTorch, Norse* Nov 2025 – Jan 2026

- Built end-to-end SNN pipeline from raw IQ samples using biologically inspired spatiotemporal receptive fields with LIF/LI dynamics; matched CNN baseline on 30-device LoRa classification (99% validation, ~85% held-out).
- Implemented metric-learning embeddings for open-set rogue detection; achieved 0.98 AUC and 92% cross-manufacturer generalization.

RL-Based Multi-Objective Neural Architecture Search for SNNs | *PyTorch, REINFORCE* Jan 2026 – Present

- Built an LSTM-controller NAS pipeline trained with REINFORCE-with-baseline to search SNN architectures under a multi-objective reward combining accuracy, spike sparsity, and a temporal-coding-quality score (CV-ISI, Van Rossum distance).
- Designed a gated multi-objective reward that prevents no-spike trivial-attractor failures.
- On Spiking Heidelberg Digits, the controller converged to a single architecture (96% policy mass, last-16 mean reward **0.47 vs. 0.14** for random search under matched compute.)
- Designed the reward shaping, search space, and held-out evaluation protocol around the multi-objective trade-off; large-scale runs to characterize the accuracy/sparsity/timing Pareto front are in progress.

EXPERIENCE

Research Software Engineering Fellow Jun 2026 – Aug 2026
Princeton University – Princeton Neuroscience Institute Remote

- Contributing to *PsyNeuLink*, an open-source Python library for compositional models of cognition and neural computation.
- Designing and implementing hierarchical Bayesian parameter estimation for simulator-based cognitive models (e.g., Drift Diffusion Model.)
- Building a parameter-estimation composition that uses Neural Likelihood Estimation (NLE) to recover full posterior distributions over individual- and group-level parameters.
- Validating end-to-end on multi-subject DDM data via parameter recovery, posterior predictive checks, and comparison against HDDM analytical baselines; delivering tutorial notebooks and user-facing API documentation.

Teaching Assistant – VLSI, Computer Architecture, Computer Security Aug 2025 – Present
University at Buffalo Buffalo, NY

- Led labs on CMOS standard-cell design; taught Cadence Virtuoso schematic/simulation/layout workflows including parasitic extraction and DRC/LVS sign-off. Mentored teams through full-flow hierarchical mixed-signal IC projects (DPLLs, TRNGs) from spec to tapeout; created assignments and exams.

Software Engineer Jun 2022 – Apr 2024
Axon Remote

- Supported a large-scale archival + ingestion platform on Azure (Temporal, Blob Storage), reducing cloud storage spend by \$1.5M/month and scaling to **250 PB**.
- Operated globally deployed ingestion for body-cams/TASER/CJIS data: on-call ownership for Dockerized services at 99.99%+ availability.
- Built Scala microservices and test coverage (unit/integration/load) in a large production codebase; shipped monthly releases across regions (Kubernetes/Puppet).

TECHNICAL SKILLS

Core: Python, C/C++, CUDA, PyTorch, Triton, Norse, Linux **Hardware/EDA:** Cadence Virtuoso, Xilinx Vivado, Verilog, SPICE, DRC/LVS

Systems: Docker, Kubernetes, Git, CI/CD, OpenMP, SQL **Focus:** ML–hardware co-design, in-memory computing, GPU training optimization, NAS, VLSI