

Problem Statement

Modern AI systems consume significant power and suffer from memory bandwidth bottlenecks, spending substantial time and energy moving data between memory and computation centers. This becomes a bigger issue in edge applications where power and latency matter. Neuromorphic computing offers a lower-power, event-driven alternative. This project studies whether model behavior can be evaluated more realistically through a controlled comparison framework before physical neuromorphic hardware is built.

Approach

This project compares conventional artificial neural networks (ANNs), digital spiking neural networks (SNNs), and hardware-aware SNNs. The hardware-aware branch uses AIHWKIT to model RRAM non-idealities such as write noise, conductance drift, and endurance effects.




Figure 1. Example weighted neural structure

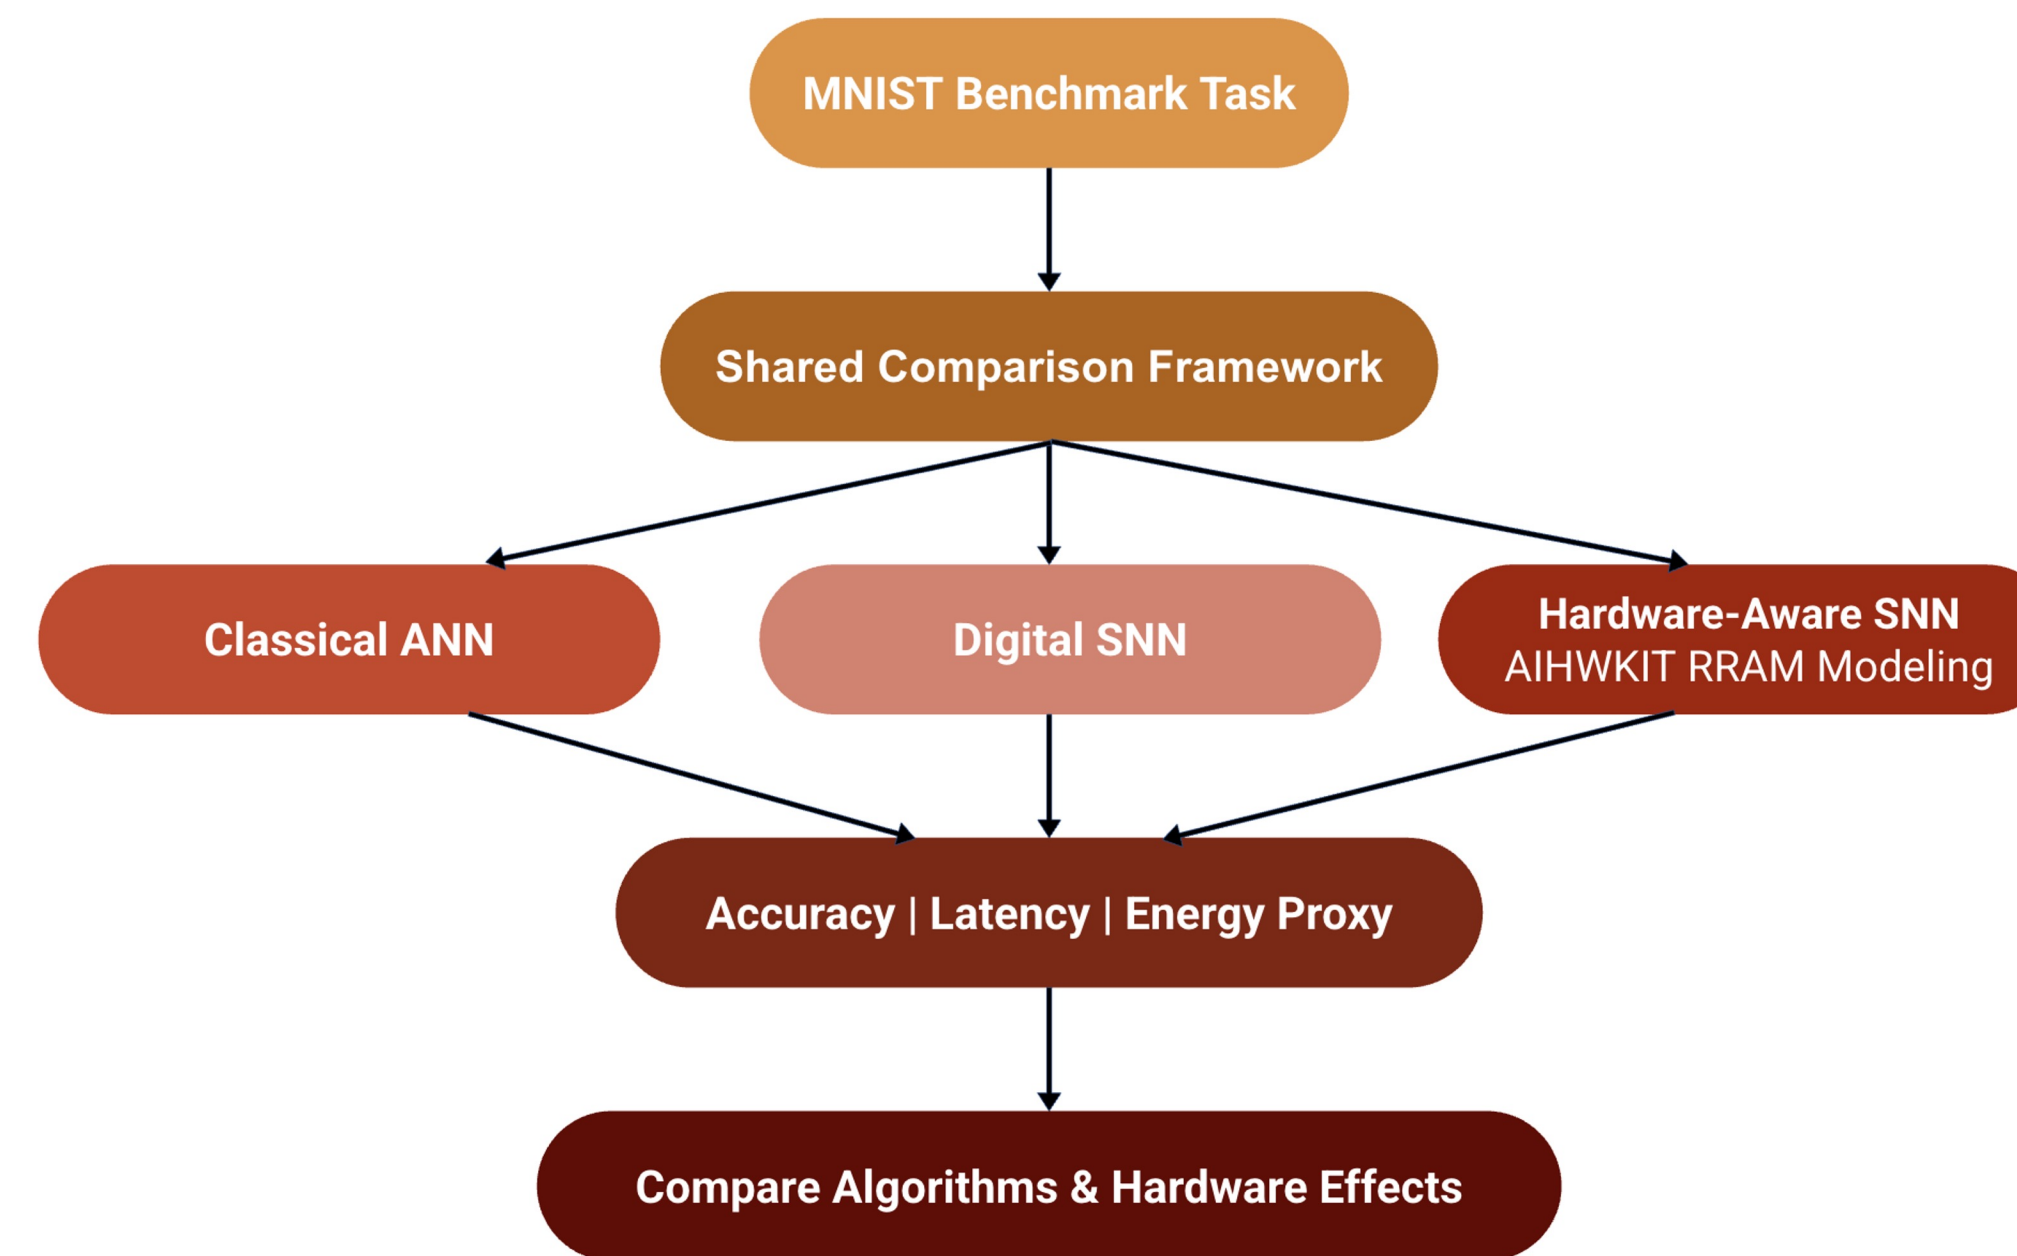


Figure 2. Framework Diagram

Results

Tested on the MNIST dataset using matched model branches. Earlier fully connected comparisons reached about 97% accuracy. Later analog-training investigations revealed larger AIHWKIT performance gaps and helped identify practical training limits.

- MNIST used as the benchmark task
- Early fully connected comparisons stayed close in accuracy
- AIHWKIT enabled hardware-aware modeling of RRAM effects
- Later analog-training runs revealed practical limits in the current setup
- Spike activity and latency were tracked as efficiency-related proxy metrics

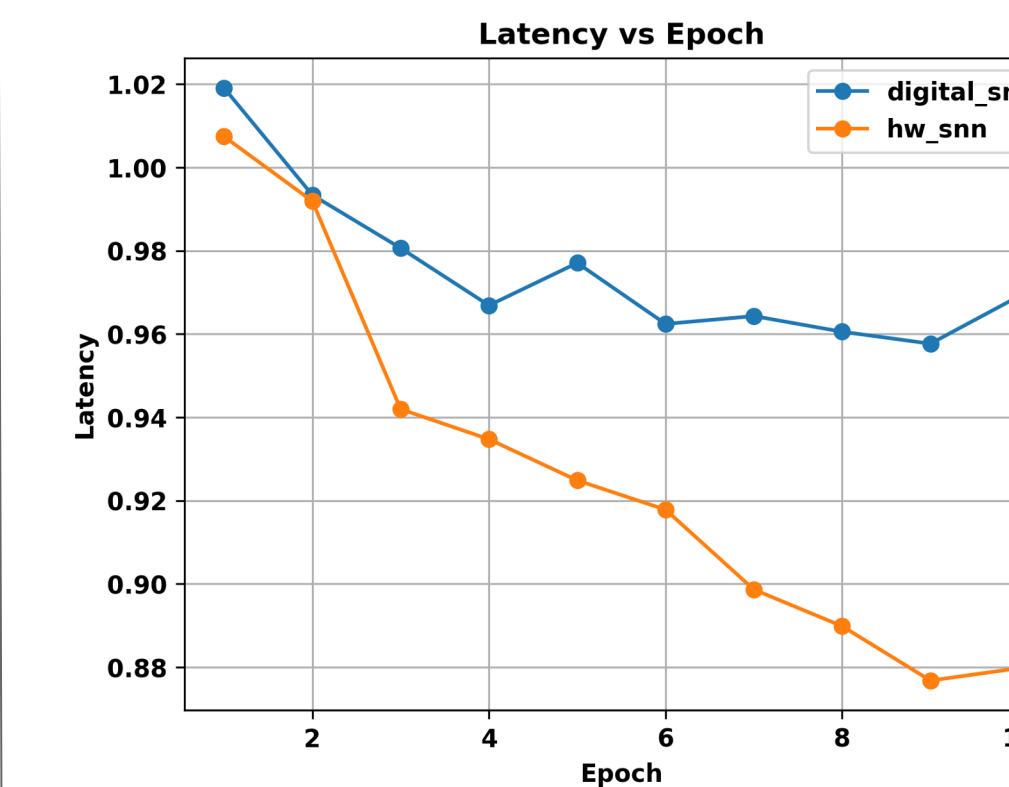


Figure 3. Latency Proxy vs Epoch

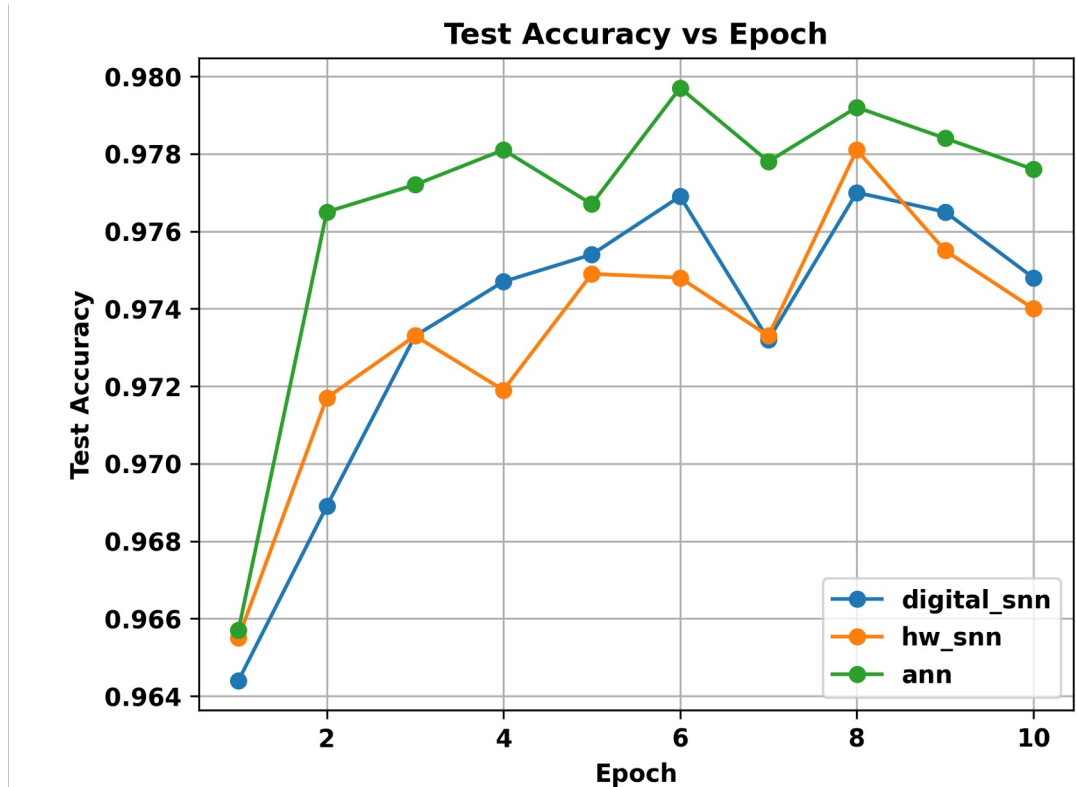


Figure 4. Test Accuracy vs Epoch

Model	Acc	Energy Proxy	Latency	Main Takeaway
ANN FC	97.94%	--	--	Strong baseline
SNN FC	96.75%	18.80 spikes	0.79 steps	Close to ANN
AIHWKIT Digital FC	97.01%	19.28 spikes	0.73 steps	HW-aware, still close
AIHWKIT Analog FC	60.07%	6.18 spikes	5.68 steps	Reveals training limits
ANN Conv	98.94%	--	--	Strong conv baseline
SNN Conv	98.68%	17.53 spikes	1.30 steps	Close to ANN
AIHWKIT Digital Conv	98.52%	16.62 spikes	1.19 steps	HW-aware, still close
AIHWKIT Analog Conv	64.41%	5.54 spikes	4.12 steps	Large analog gap

Table 1. Model Benchmarks

Impact & Next Steps

This project provides a practical software framework for studying neuromorphic AI before physical hardware is available. The framework helps expose hardware-performance tradeoffs before fabrication.

Next steps: refine analog-training conditions to better understand the tradeoffs between realism and model performance.