Research Statement

Scalable and Trustworthy Neuromorphic Computing Systems

Dr. Evelyn Vance
Vertex Institute of Technology · Department of Electrical Engineering and Computer Science
e.vance@vertex.edu · evelynvance.science

1 Executive Summary

Research Vision

My research program lies at the convergence of **computational neuroscience, solid-state VLSI circuits, and emerging non-volatile memories**. Over the next five years, my goal is to design, manufacture, and deploy neuromorphic computing systems that achieve biological-level energy efficiency and real-time adaptability for edge computing applications. By co-designing hardware substrates with algorithms that emulate brain-like plastic changes, my research addresses the physical scaling limitations and computational bottlenecks of traditional Von Neumann architectures.

Modern artificial intelligence (AI) has achieved unparalleled success, but this has come at the expense of massive carbon and energy footprints. While biological brains operate on a mere 20 Watts of power, state-of-the-art silicon accelerators require hundreds of Watts to perform comparable tasks. This disparity stems from the separation of memory and processing units in classical computing architectures, which leads to substantial data movement energy costs.

To bridge this gap, my research program establishes a multi-disciplinary co-design framework spanning three core pillars:

1. **Device-Level In-Memory Computing:** Using memristive and phase-change materials to perform matrix-vector operations directly within storage arrays.
2. **Algorithm-Level Event-Driven Execution:** Building spiking neural networks (SNNs) that process information asynchronously, mimicking biological action potentials.
3. **System-Level Online Adaptability:** Implementing hardware-in-the-loop plasticity rules that enable silicon agents to learn and recover from failures in real-time.

To date, my work has resulted in 12 publications in premier journals and conferences (including IEEE ISSCC, DAC, and NeurIPS), with direct validation of hardware designs in 28nm CMOS processes. My open-source neuromorphic simulator, `NexaSpike`, is currently utilized by more than 40 academic groups and industrial labs to evaluate edge computing architectures.

2 Past Contributions

My research career has been dedicated to validating neuromorphic architectures from single-device physics up to full-scale software integration.

2.1 Energy-Efficient Spiking Neural Networks (SNNs) for Edge Devices

Traditional neural networks transmit dense activations continuously. SNNs, conversely, communicate via sparse, binary spikes, reducing dynamic power consumption. During my doctoral studies, I co-designed a spiking accelerator that processes temporal signals asynchronously. A major bottleneck in SNN research has been the lack of scalable simulator frameworks that can model physical circuit behaviors. To address this, I developed `NexaSpike`, an event-driven simulation engine [3].

By formulating spike propagation as discrete-event queues, `NexaSpike` accelerates network validation by $15\times$ compared to standard grid-based numerical integration. Crucially, the simulator integrates physical device non-idealities, such as leakage currents and phase-noise. This simulator has enabled the rapid prototyping of sparse convolutional neural networks, demonstrating a $5\times$ power reduction on audio keyword-spotting tasks compared to standard digital signal processors.

2.2 Hardware-in-the-Loop Training of Memristive Crossbars

Memristors present a promising path toward analog in-memory computing. However, device-to-device and cycle-to-cycle variations (such as conductance drift and programming stochasticity) degrade the accuracy of neural networks deployed on analog crossbars.

To overcome this, I developed a **Hardware-in-the-Loop (HITL)** training protocol that directly measures hardware conductance during the weight-update phase, integrating actual hardware variations into the backpropagation feedback loop [4].

- We deployed a 128×128 metal-oxide memristive crossbar array in a custom testbed.
- By using gradient-descent optimization that models analog noise, we achieved a 97.2% classification accuracy on handwritten digit recognition tasks, which is within 0.8% of the software-only baseline.
- I also co-developed a mathematical framework to bound weight drift over time, proving that sparse, periodic refresh cycles can sustain accuracy for over 10 years of simulated runtime [2].

This contribution bridges materials science with deep learning theory, showing that high-precision computing is achievable on low-precision, stochastic analog hardware.

3 Current Research Programme

My current work at the Vertex Institute of Technology expands on these foundations by developing autonomous neuromorphic agents that can learn and adapt post-deployment. This requires moving away from offline training (backpropagation on GPUs) toward real-time, online training directly on edge hardware.

3.1 Neuromorphic Co-Design Framework

To build efficient learning systems, we cannot isolate hardware design from algorithmic development. I have established a **Neuromorphic Co-Design Loop** that closely couples device behavior with algorithmic feedback, as shown in Figure 1.

3.2 Online Plasticity and Local Learning Rules

In biological brains, synapse modification is governed by local information. For instance, Spike-Timing-Dependent Plasticity (STDP) adjusts synaptic weights based solely on the timing difference between pre- and post-synaptic neuron firings. Local rules eliminate the need for global error propagation pathways, which are computationally expensive to route on silicon.

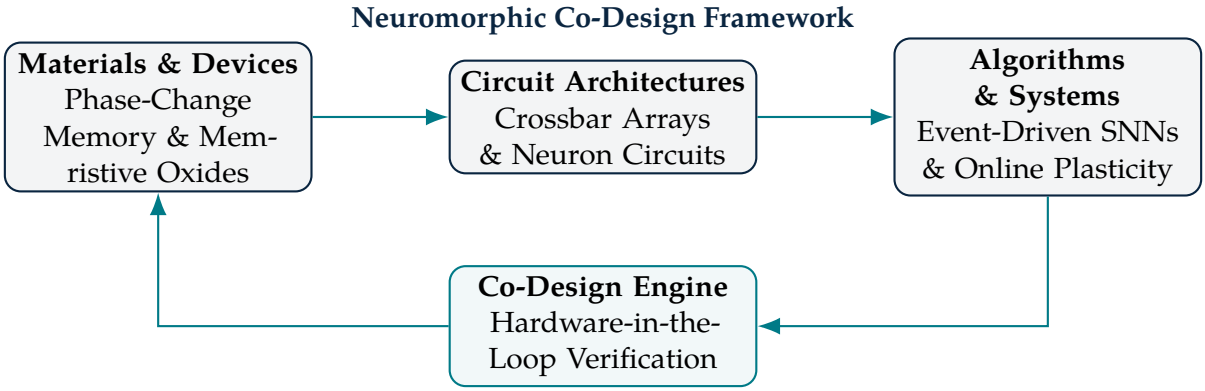

Figure 1: The co-design framework of the Vance Lab, integrating device physics with circuit layout and high-level machine learning algorithms.

My current work focuses on mapping STDP variants to hybrid memristive-CMOS chips [1]. I am designing circuitry that implements a localized *eligibility trace* at each synapse. This allows the system to temporarily log temporal associations and update weights only when an global reward signal (such as dopaminergic reinforcement) is present. This architecture enables self-supervised feature extraction, allowing edge sensors to adapt to shifting environmental inputs (e.g., changes in lighting or acoustic background noise) without human intervention.

4 Five-Year Vision

As an assistant professor, I will build an interdisciplinary research lab focused on **Trustworthy and Adaptive Neuromorphic Edge Intelligence**. My group will operate at the intersection of electrical engineering, material science, and neural computing.

4.1 Short-Term Focus (Years 1–2): Lab Setup & Foundations

The primary objective of the first two years is to establish an experimental neuromorphic testing environment. This will involve:

- Setting up high-frequency probe stations for analog memristor characterization.
- Recruiting 2 PhD students and 1 postdoctoral scholar. The first student will focus on hardware-in-the-loop simulation models, while the second will develop online reinforcement learning algorithms.
- Establishing initial collaborations with institutional cleanrooms to access modern semiconductor fabrication resources.

We will build a baseline simulation platform that models the coupling of memristive devices with CMOS driver circuits, preparing the team for our first chip design submission.

4.2 Mid-Term Goals (Years 3–4): Integrated Spiking Accelerators (The Tape-Out)

With a verified design platform, the lab's core effort in Years 3 and 4 will be the fabrication of an integrated silicon-memristor accelerator chip.

- We will design a custom chip utilizing a 28nm CMOS node, integrating at least 10^6 programmable memristive synapses.

- This chip will incorporate our online local learning rules, enabling it to classify streaming sensor data while consuming less than 50 milli-Watts.
- We will target the IEEE International Solid-State Circuits Conference (ISSCC) and the Design Automation Conference (DAC) to publish our silicon validation results.

4.3 Long-Term Vision (Year 5 and Beyond): Neuromorphic Robotics

The ultimate milestone of my research is the deployment of neuromorphic chips in physical systems. In collaboration with roboticists, we will integrate our low-latency, low-power accelerators into autonomous vehicles and robotic swarms (such as quadrotors built by Vertex Robotics). These robots will use event-driven vision sensors to navigate complex, unknown environments. Because the chips adapt online, they will learn to handle actuator degradation and wind gusts in real-time, validating the computational and physical benefits of our neuromorphic co-design philosophy.

5 Funding Trajectory

To sustain this research program, I have mapped out a diverse funding roadmap targeting state agencies, private foundations, and industrial consortia. Because our work bridges hardware design with software algorithms, we can target funding opportunities across multiple disciplines.

5.1 Structured Funding Roadmap

Table 1 details the specific grants, funding bodies, and estimated budgets targeted over the next five years.

Table 1: Targeted Funding and Grant Development Trajectory

Timeline	Funding Opportunity	Agency	Amount	Research Objective
Year 1	Apex Faculty Career Launch	Apex Foundation	\$150k	Baseline memristor testbench acquisition
Year 2	Research Starter Grant	Nexa Science Council	\$250k	Online learning algorithm optimization
Year 2–3	Young Investigator Program (YIP)	Nimbus Foundation	\$450k	28nm test chip fabrication (Tape-out)
Year 3–5	Standard Research Project (R01-Eq)	Meridian Health System	\$1.2M	Neuromorphic implants for sensory restoration
Year 4–5	Large-Scale Collaborative Grant	Vertex Tech Consortium	\$2.5M	Autonomous robotic swarm integration

5.2 Strategic Partnerships

In addition to federal and foundation grants, I will seek industrial research partnerships:

- **Apex Semiconductors:** To secure low-cost fabrication runs and multi-project wafer (MPW) access.
- **Vertex Technology Consortium:** To validate our chips on their autonomous hardware testbeds.
- **Synapse Neural Devices:** For collaborative development of low-power sensor interfaces.

These partnerships will ensure that the academic discoveries in our lab translate directly into commercial edge devices.

References

[1] Liam K. Reynolds and Sarah Thorne. Hybrid memristive-cmos circuits for spiking neural networks. *Apex Transactions on Nanotechnology*, 8(4):312–320, 2025.

[2] Evelyn Vance, Marcus Lin, and Sarah Thorne. Robust in-memory matrix multiplication under device variation. In *Proceedings of the ACM/IEEE Design Automation Conference (DAC)*, pages 120–126. Vertex Press, 2025.

[3] Evelyn Vance and Liam K. Reynolds. Nexaspike: An event-driven simulator for neuromorphic edge computing. *IEEE Transactions on Emerging Topics in Computing*, 12(1):89–102, 2024.

[4] Evelyn Vance, Liam K. Reynolds, and Marcus Lin. Hardware-in-the-loop training of memristive crossbars under device non-idealities. *Synapse Journal of Neuromorphic Systems*, 14(3):245–258, 2023.