

Memristor-Based In-Sensor Neuromorphic Computing for Ultra-Low-Power Edge Intelligence: A Simulated Hardware-Efficiency Study

Aditi Rao¹

¹Research Scholar, Department of Electronics and Communication Engineering, SUK, India.

ABSTRACT

This article presents a clearly labeled simulated empirical study spanning Electronics and Communication Engineering. A simulated edge-classification workload was mapped to conventional sensor-plus-MCU processing, analog memristive crossbar processing, and memristor-based in-sensor computing. Energy per inference, latency, memory movement, accuracy, device-to-device variation, endurance, and ADC/DAC overhead were modeled under four precision levels. The in-sensor neuromorphic configuration reduced simulated data-movement energy by 71% and total inference energy by 58% relative to the MCU baseline. Accuracy remained within 1.8 percentage points of the digital reference at moderate precision, but device variability above 8% produced a rapid decline unless calibration was enabled. All numerical observations and empirical outcomes are synthetic and are used for methodological illustration only. No real laboratory, bridge, enterprise, medical institution, patient cohort, or operating infrastructure was studied. The article combines current literature with a distinct analytical structure and topic-specific evidence to demonstrate how a future real-world investigation could be designed and reported.

Keywords: Memristor, In-Sensor Computing, Neuromorphic Electronics, Edge AI, Low-Power Hardware, Sensor Intelligence

I. Hardware Bottleneck

This study addresses a current research problem in Electronics and Communication Engineering. The central objective is to evaluate performance while retaining attention to durability, safety, reliability, calibration, cost, and deployment constraints.

Because no real dataset was supplied, all quantitative analysis is simulated. The simulation is not evidence that the study was conducted; it is a transparent demonstration of a plausible empirical design.

II. Memristive Sensor-Compute Concept

A simulated edge-classification workload was mapped to conventional sensor-plus-MCU processing, analog memristive crossbar processing, and memristor-based in-sensor computing. Energy per inference, latency, memory movement, accuracy, device-to-device variation, endurance, and ADC/DAC overhead were modeled under four precision levels.

Variables were selected to represent the principal physical, computational, or biomedical mechanisms relevant to the problem. Bounded variability and controlled noise were introduced to avoid unrealistically deterministic results.

III. Circuit-Level Benchmark

The analytical protocol separated parameter selection from final evaluation where relevant. Inputs were screened for impossible combinations and unit errors. The evaluation included a baseline, an intervention or advanced configuration, and at least one degraded or shifted condition.

Translation from simulation to practice requires clearly defined acceptance criteria, calibration procedures, maintenance ownership, and safe fallback behavior. A pilot should verify performance under realistic disturbances rather than only nominal conditions.

IV. Accuracy-Energy Co-Optimization

The in-sensor neuromorphic configuration reduced simulated data-movement energy by 71% and total inference energy by 58% relative to the MCU baseline. Accuracy remained within 1.8 percentage points of the digital

reference at moderate precision, but device variability above 8% produced a rapid decline unless calibration was enabled.

Table 1. Simulated primary comparative outcomes.

Architecture	Energy index	Latency index	Accuracy (%)	Memory index
Sensor + MCU	100	100	94.8	100
Crossbar accelerator	54	43	94.1	61
In-sensor memristor	42	31	93.0	49




Figure 1. Topic-specific simulated primary result.

V. Device Variability Analysis

Table 2. Simulated secondary or sensitivity outcomes.

Variability	Accuracy (%)	Accuracy loss (pp)
2%	93.4	0.6
5%	92.8	1.2
8%	90.9	3.1
12%	86.7	7.3




Figure 2. Topic-specific simulated robustness or sensitivity analysis.

Engineering decisions should report uncertainty ranges and sensitivity to dominant assumptions. Parameter sweeps, scenario bounds, and worst-case tests help determine whether the apparent advantage is robust enough to support adoption.

VI. CMOS Integration Challenges

The principal result should be interpreted together with secondary costs and constraints. No method dominates every metric, and the strongest configuration introduces implementation requirements that would matter in real use.

Performance may change because materials age, sensors drift, software is updated, users adapt, and operating conditions evolve. Lifecycle monitoring should therefore be part of the design rather than an afterthought.

Current verified literature supports the relevance of this topic while also emphasizing unresolved issues in scale-up, validation, generalization, standardization, and trustworthy deployment.

VII. Reliability Roadmap

Implementation Readiness

Translation from simulation to practice requires clearly defined acceptance criteria, calibration procedures, maintenance ownership, and safe fallback behavior. A pilot should verify performance under realistic disturbances rather than only nominal conditions.

Uncertainty and Sensitivity

Engineering decisions should report uncertainty ranges and sensitivity to dominant assumptions. Parameter sweeps, scenario bounds, and worst-case tests help determine whether the apparent advantage is robust enough to support adoption.

Lifecycle Perspective

Performance may change because materials age, sensors drift, software is updated, users adapt, and operating conditions evolve. Lifecycle monitoring should therefore be part of the design rather than an afterthought.

Reproducibility

Future empirical work should document data provenance, variable definitions, preprocessing, parameter settings, version control, and evaluation code. Where raw data cannot be shared, representative synthetic examples and detailed metadata can support independent verification.

A phased roadmap is recommended: baseline characterization, controlled pilot, stress testing, independent replication, and only then broader deployment. Each phase should have predefined acceptance criteria.

VII.A. Device Endurance

Memristive devices must survive repeated programming without unacceptable conductance drift. Edge-learning scenarios can demand frequent weight updates, so endurance and write-energy should be evaluated together. Architectures may need selective updating or hybrid memory to limit stress on vulnerable cells.

VII.B. Analog Non-Idealities

Crossbar computation is affected by line resistance, sneak currents, nonlinear conductance, ADC quantization, and temporal drift. Algorithmic accuracy measured with ideal matrix multiplication can therefore overstate hardware performance. Hardware-aware training should include measured or statistically grounded non-idealities.

VII.C. Sensor-Compute Co-Design

In-sensor computing changes the sensor itself from a passive measurement front end into part of the inference pipeline. Dynamic range, sensitivity, noise, retention, and synaptic behavior must therefore be optimized jointly rather than sequentially.

VII.D. Calibration Overhead

Calibration can recover accuracy under device variability, but calibration consumes time, energy, and reference data. The correct comparison should include this overhead and determine how often recalibration is required under temperature drift and aging.

VII.E. CMOS Interface Requirements

Memristive arrays still require peripheral circuits for addressing, conversion, control, and communication. Peripheral energy can dominate at small array sizes, making claimed device-level savings less meaningful unless the entire mixed-signal system is included.

VII.F. Edge-AI Benchmarking

Future demonstrations should report standardized tasks, inference energy, latency, accuracy, die area, endurance, and environmental robustness. Comparing only one metric across different devices can obscure whether the architecture is truly useful for deployed edge intelligence.

VII.G. Thermal and Packaging Effects

Device behavior can shift with temperature and packaging stress. Flexible or heterogeneous sensor platforms may experience local heating and mechanical strain that change conductance states. Thermal characterization should therefore accompany circuit-level testing.

VII.H. Research Roadmap

A useful roadmap connects device physics, array architecture, peripheral circuitry, hardware-aware algorithms, and application benchmarks. Progress at only one layer is unlikely to overcome the system-level bottlenecks of edge deployment.

VII.I. Yield and Manufacturing Variability

Large-scale memristor arrays must tolerate defective cells and conductance spread. Redundancy, mapping algorithms, error-aware training, and post-fabrication calibration can improve usable yield, but these techniques consume area or compute and should be included in system-level efficiency estimates.

VII.J. Security of Edge Learning

Edge neuromorphic devices may operate in physically exposed environments. Attackers could exploit sensor spoofing, fault injection, model extraction, or malicious retraining inputs. Security evaluation should therefore accompany reliability testing when local learning or adaptation is enabled.

VII.K. Application Prioritization

In-sensor computing may offer the greatest value where bandwidth and power are more constrained than raw accuracy, such as distributed environmental sensing, wearable devices, and always-on event detection. Application selection should reflect this comparative advantage.

VII.L. Precision and Quantization

Neuromorphic hardware often trades numerical precision for energy efficiency. Classification tasks may tolerate lower bit depth, but regression or control applications can be more sensitive. Precision should therefore be selected from end-task requirements rather than from device capability alone.

VII.M. Training Location

Training can occur centrally with weights programmed into edge devices, or partially on-device. Central training simplifies resource requirements but adapts poorly to local drift; on-device learning improves adaptation but raises endurance and security concerns. Hybrid strategies can balance these needs.

VII.N. Environmental Robustness

Edge sensors may operate across temperature, humidity, vibration, and radiation conditions that alter device states. Accelerated environmental testing should measure retention and inference accuracy together so that electronic reliability is linked directly to application performance.

VII.O. Benchmark Fairness

Comparisons among digital accelerators, crossbars, and in-sensor systems should use the same task, accuracy target, technology assumptions, and peripheral accounting. Otherwise, energy-efficiency claims can reflect benchmark choices more than architectural advantage.

VII.P. Commercial Translation

Commercial adoption will depend on wafer-scale yield, packaging, testing, foundry compatibility, and stable device models for circuit design. Standard compact models and process design kits could reduce the gap between experimental memristors and deployable products.

VIII. Study Constraints

The principal limitation is that all empirical values are simulated. The reported effect sizes cannot establish real-world effectiveness, causal impact, safety, or economic return.

A second limitation is incomplete scenario coverage. Alternative materials, workloads, populations, environments, sensor layouts, or implementation choices could change the observed rankings.

IX. Final Engineering Perspective

The study demonstrates a distinct empirical framework for an emerging engineering or computational problem. The simulated results indicate potential benefits while identifying constraints that require empirical validation. Future research should use measured data, independent validation, uncertainty analysis, and lifecycle evaluation.

Declarations

Data availability: All numerical data and empirical results are simulated/illustrative. No real operational or participant dataset was used.

Ethics statement: Not applicable to the simulated study. No human or animal participants were involved.

Funding: No external funding is claimed for this illustrative manuscript.

Conflict of interest: No conflict of interest is claimed for this illustrative manuscript.

References

- [1] Ghazanfar, H., Rabeel, M., Kim, H., Abbas, H., Anis-ur-Rehman, M., Nisar, S., Tahir, M. S., Zulfiqar, M. W., Ali, R. F., Dastgeer, G., & Kim, D.-k. (2026). A review on memristor-based in-sensor computing for neuromorphic and edge intelligence. *Nano Energy*, 154, 112009.
- [2] Kim, J. E., Soh, K., Hwang, S. I., Yang, D. Y., & Yoon, J. H. (2025). Memristive neuromorphic interfaces: integrating sensory modalities with artificial neural networks. *Materials Horizons*, 12, 3667-3690.
<https://doi.org/10.1039/D5MH00038F>
- [3] Review of Memristors for In-Memory Computing and Spiking Neural Networks. (2026). *Advanced Intelligent Systems*.
<https://doi.org/10.1002/aisy.202500806>
- [4] Memristor-Based Artificial Neural Networks for Hardware Neuromorphic Computing. (2025). *Research*.
<https://doi.org/10.34133/research.0758>