

Hardware-Aware Mathematical and Numerical Evaluation of LIF, Izhikevich, and Digital Spiking Neuron Models for Low-Cost Neuromorphic Implementation

JESUS RODRIGUEZ-FLORES^{1,a,*}, EDGART MARCHAN-MILLAN^{2,b},
ALEX MOYOTA-PAGUAY^{3,c}

¹Department of Mechatronics, Faculty of Engineering and Industrial Sciences,
Universidad UTE,
Mariscal Sucre Av. s/n and Mariana de Jesus, Quito,
ECUADOR

²Independent Researcher,
Calle La Diligencia No. 36, 2nd Floor, Door A, 50019 Zaragoza,
SPAIN

³Department of Applied Mathematics,
Escuela Superior Politécnica de Chimborazo (ESPOCH),
Morona Santiago Campus, Macas,
ECUADOR

^aORCID: <https://orcid.org/0000-0001-6254-2348>

^bORCID: <https://orcid.org/0009-0001-6802-8654>

^cORCID: <https://orcid.org/0009-0008-4864-7643>

**Corresponding Author*

Abstract: - Low-cost neuromorphic computing requires neuron models that combine relevant spiking dynamics with numerical stability and technological feasibility. This work compares three representative spiking-neuron families: Leaky Integrate-and-Fire (LIF), Izhikevich, and a simplified digital ShiftLIF-type approximation. A unified, reproducible, and hardware-oriented framework was developed using a synthetic stimulus bank, homogeneous MATLAB simulations, and controlled variations in temporal discretization, noise, quantization, and parametric perturbations. The evaluation combined dynamic, numerical, and hardware-aware metrics with multicriteria aggregation across application scenarios. Izhikevich achieved the lowest mean first-spike error (0.809 ms) and the highest noise robustness (0.981), whereas LIF attained the lowest normalized-state RMSE (0.017) and the lowest mean computation time (1.875 ms). Although ShiftLIF did not lead the local temporal-accuracy metrics, it showed the lowest discretization sensitivity (0.177), parametric sensitivity (0.0195), quantization sensitivity (0.3084), and arithmetic cost per iteration (4 operations). It also obtained the highest score in the balanced multicriteria scenario (0.691) and ranked first in all four evaluated scenarios. By contrast, under 8-bit quantization, LIF completely lost spiking activity in 70.6% of the runs. Overall, no model was universally superior; however, within the adopted framework, ShiftLIF provided the most favorable trade-off for low-cost digital neuromorphic implementations.

Key-Words: - Neuromorphic computing, Spiking neural networks, Leaky integrate-and-fire model, Izhikevich model, ShiftLIF, Hardware-aware evaluation, Multicriteria benchmarking, Low-cost neuromorphic hardware.

Received: April 5, 2026. Revised: June 8, 2026. Accepted: July 8, 2026. Published: October 5, 2026.

1 Introduction

1.1 Context and Motivation of the Problem

The recent expansion of artificial intelligence has renewed attention to a structural limitation of conventional computing: rising energy costs, latency, and data movement between memory and processing units, especially on edge, embedded, and autonomous platforms with restricted power, thermal capacity, and bandwidth. Neuromorphic computing has therefore regained relevance as an approach based on parallelism, local processing, and sparse communication for efficiency-critical tasks [1], [2], [3]. Within this paradigm, Spiking Neural Networks (SNNs) are central because they process information through discrete temporal events, enabling asynchrony, sparsity, and temporal coding for dynamic perception, sensing, neuromorphic vision, and real-time control [4], [5], [6].

This technological relevance depends on the possibility of implementing neuromorphic principles on lightweight, reconfigurable, and economically viable hardware. In this context, neuron-model selection becomes decisive, since arithmetic cost, memory demand, numerical stability, and tolerance to reduced precision may scale significantly from isolated neurons to larger systems. Thus, model choice must jointly consider dynamic expressiveness, compact implementation, efficient inference, and local event processing [7], [8], [9].

1.2 State of the Problem and Research Gap

SNN research has produced a broad family of neuron models, ranging from simple, interpretable formulations such as the LIF model to more expressive dynamics involving artificial synapses, STDP-inspired plasticity, and emerging non-volatile memories [10], [11], [12]. The challenge is not the lack of alternatives, but the existence of models with different trade-offs among biological realism, numerical stability, and technological feasibility. Richer models may reproduce excitability patterns more faithfully, but usually require more state variables, more operations per iteration, and greater sensitivity to discretization or quantization; simpler models may be easier to implement, but less expressive. Recent studies on photonic, mixed-signal, processor-based, and memristive neuromorphic platforms confirm that a model's usefulness depends on the balance between functional behavior and implementation economy [13], [14], [15].

However, many comparisons remain partial or heterogeneous, emphasizing devices, architectures,

applications, or training methods rather than neuron models evaluated under identical excitation, noise, discretization, quantization, and measurement conditions. This limits the availability of reproducible criteria for deciding which model family is preferable for low-cost, scalable, and functionally sufficient neuromorphic hardware. Therefore, a homogeneous comparative framework is still needed to distinguish model-dependent behavior from effects associated with the platform, task, or training strategy.

1.3 Research Problem, Objective, and Contribution

The central problem addressed in this work is the difficulty of selecting, under multiple constraints and comparable evidence, a suitable spiking neuron model for low-cost neuromorphic hardware. This selection depends on the joint behavior of temporal precision, numerical stability, computational complexity, robustness to noise and quantization, memory requirements, and ease of implementation, rather than on biological plausibility alone.

The objective of the article is to compare three representative neuron families—LIF, Izhikevich, and a simplified digital ShiftLIF-type approximation—under a unified, reproducible, and hardware-oriented methodological framework. The models are evaluated using homogeneous excitation, temporal discretization, noise, quantization, and parametric perturbation conditions, so that dynamic, numerical, and hardware-aware metrics can be assessed jointly. The study is deliberately comparative and *in silico*: it does not claim physical validation on an FPGA, nor direct measurements of energy consumption or logic resource usage, but provides a controlled basis for anticipating comparative behavior before physical implementation.

The main contribution is therefore not a new neuron model or a universal ranking, but a theoretical-methodological framework that integrates dynamic response, numerical sensitivity, arithmetic cost, and technological convenience within a single protocol. These criteria are articulated through a multicriteria reading across application scenarios, supporting model selection beyond general claims of “realism” or “simplicity”.

1.4 Delimitations and Organization of the Article

The validation level adopted here is comparative, computational, and oriented toward technological plausibility. Consequently, the results should not be

interpreted as final evidence of physical performance on a specific platform, nor extrapolated directly to area, power, maximum operating frequency, or measured energy consumption. The comparison is restricted to explicit dynamic, numerical, and hardware-aware metrics, including first-spike latency, trajectory RMS error, firing rate, temporal-step sensitivity, noise robustness, parametric sensitivity, arithmetic cost, memory per neuron, quantization sensitivity, ease of implementation, and dynamic richness.

The remainder of the manuscript is organized as follows. Section 2 presents the comparative methodology and benchmark design. Section 3 reports the results. Section 4 discusses their mathematical and technological implications. Finally, Section 5 summarizes the main conclusions.

2 Methodology

The methodology was designed as a *computational, reproducible, multicriteria, and hardware-oriented* protocol to compare, under homogeneous excitation and evaluation conditions, three spiking-neuron families: the *Leaky Integrate-and-Fire* (LIF) model, the Izhikevich model, and a simplified digital *ShiftLIF*-type approximation. This design addresses a recognized need in neuromorphic computing: model comparison should not rely solely on dynamic fidelity but also on temporal precision, numerical stability, computational cost, quantization tolerance, and implementation feasibility. Recent reviews and benchmarking frameworks support integrated evaluations combining hardware-dependent and hardware-independent criteria under realistic edge-computing conditions [16], [17], [18], [19].

This study was developed from an original computational protocol specifically designed for the present comparison. All manuscript text, simulation scripts, computational framework components, synthetic stimulus data, result tables, figures, multicriteria indices, and numerical outputs were created specifically for this study. No previously published text passages, code fragments, framework components, datasets, figures, tables, or numerical results from the authors' prior work or from third parties were reused or recycled. Established neuron models and numerical methods were implemented independently and are cited to their original or relevant sources.

Computational traceability was ensured by explicitly defining the model equations, using a fixed global random seed, maintaining homogeneous simulation conditions, and

systematically recording scenarios, metrics, and numerical outputs. Therefore, the reported results derive directly from the described protocol and should be interpreted as reproducible computational evidence within the *in silico* scope of the study.

The study followed six linked phases: i) definition of the synthetic stimulus bank and experimental conditions; ii) homogeneous mathematical formulation of the three models; iii) temporal discretization and reproducible numerical simulation in Matlab; iv) controlled incorporation of noise, quantization, and parametric perturbations; v) extraction of dynamic, numerical, and hardware-aware metrics; and vi) multicriteria integration of results to establish scenario-dependent rankings. This structure ensured that the reported evidence emerged from an explicit and repeatable computational protocol.

2.1 General Design of the Computational Evaluation Benchmark

The research was developed as a comparative *in silico* study with full control over excitation, discretization, and evaluation conditions. It was not intended as FPGA validation or as experimental characterization of a built neuromorphic circuit, but as a synthetic benchmark plausibly aligned with low-cost digital implementation constraints. This strategy is consistent with recent works that use simulation frameworks, algorithm–hardware co-design, and quantitative evaluation to anticipate behavior before physical implementation [20], [21], [22].

To ensure repeatability, a global random seed was fixed, the stimulus families were kept constant, and the same temporal horizon was applied to all models. The simulation horizon was $T_{\text{end}} = 250$ ms, with a fine reference step $\Delta t_{\text{ref}} = 0.02$ ms and a family of simulation steps

$$\Delta t \in \{0.10, 0.25, 0.50, 1.00\} \text{ ms}, \quad (1)$$

which made it possible to study temporal precision and sensitivity to discretization simultaneously.

Additionally, relative noise levels in the input were considered as

$$\eta \in \{0, 0.02, 0.05\}, \quad (2)$$

and the adopted quantization regimes were defined by the set

$$\mathcal{B} = \{\infty, 16, 12, 8\} \text{ bits}, \quad (3)$$

where $B \in \mathcal{B}$ denotes the selected quantization regime and ∞ denotes double-precision simulation. Baseline runs used nominal parameters. Parametric uncertainty was analyzed separately through

controlled variations of 0%, 5%, and 10% in the characteristic parameters of each model, thus separating nominal behavior from internal robustness.

2.2 Stimulus Bank and Input Construction

Model excitation was defined through five input families: Step, Pulse, PulseTrain, Periodic, and NoisySignal. The goal was not to reproduce the full diversity of biological currents, but to provide a minimal yet sufficiently varied set capable of eliciting integration regimes, isolated spikes, event trains, and responses to periodic and perturbed signals.

In general form, the input applied to each model can be written as

$$I(t) = I_{\text{base}}(t) + \eta \sigma_I \xi(t), \quad (4)$$

where $I_{\text{base}}(t)$ represents the deterministic component of the stimulus, η is the relative noise level, σ_I is a scale associated with the maximum stimulus amplitude, and $\xi(t)$ is a zero-mean random sequence. For NoisySignal, an internally perturbed oscillatory component was added. In all runs, the current was saturated within the benchmark's operating interval to avoid numerically extreme values and preserve comparability among models.

The PulseTrain family was used as the reference condition for temporal-response and raster figures because it clearly reveals differences in excitability pattern, first-spike latency, and event regularity, in line with recent event-driven evaluation practices [18], [22], [23].

2.3 Mathematical Formulation of the Compared Models

2.3.1 LIF Model

The LIF model was defined through the classical ordinary differential equation of leaky integration:

$$\tau_m \frac{dV}{dt} = -(V - V_{\text{rest}}) + RI(t) \quad (5)$$

where V is the membrane potential, V_{rest} is the resting potential, R is an effective resistive gain, and τ_m is the membrane time constant. When $V(t)$ reaches the threshold V_{th} , a spike is registered, the variable is reset to V_{reset} , and a short refractory period is activated.

For the discrete implementation, a forward Euler scheme was employed:

$$V_{k+1} = V_k + \Delta t \frac{-(V_k - V_{\text{rest}}) + R g_I I_k}{\tau_m}, \quad (6)$$

where g_I denotes the input gain used in the simulation. The spike event was defined by the rule if $V_{k+1} \geq V_{\text{th}}$, $s_{k+1} = 1$, $V_{k+1} \leftarrow V_{\text{reset}}$. (7)

The LIF model was included as a parsimonious reference for contrasting computational cost and dynamic expressiveness [18], [23], [24].

2.3.2 Izhikevich Model

The second model was the Izhikevich model, selected for its richer excitability dynamics at moderate cost relative to more detailed biophysical formulations. Its continuous formulation was expressed as

$$\frac{dv}{dt} = 0.04v^2 + 5v + 140 - u + g_I I(t) \quad (8)$$

$$\frac{du}{dt} = a(bv - u) \quad (9)$$

where v is the membrane variable and u is the recovery variable. When v reaches the spiking threshold, the reset condition is applied

$$\text{if } v \geq 30 \text{ mV, } v \leftarrow c, \quad u \leftarrow u + d. \quad (10)$$

Its numerical discretization was also performed using explicit Euler:

$$v_{k+1} = v_k + \Delta t (0.04v_k^2 + 5v_k + 140 - u_k + g_I I_k) \quad (11)$$

$$u_{k+1} = u_k + \Delta t a(bv_k - u_k) \quad (12)$$

For this nonlinear model, global stability of the explicit Euler scheme is not claimed, because the voltage equation is nonlinear and the threshold-reset mechanism defines a hybrid dynamics. Nevertheless, local stability can be stated on bounded subthreshold domains. Let $F(v, u; I)$ denote the continuous-time vector field. Its Jacobian with respect to the state variables is

$$J(v, u) = \frac{\partial F}{\partial (v, u)} = \begin{bmatrix} 0.08v + 5 & -1 \\ ab & -a \end{bmatrix}. \quad (13)$$

Around a stable subthreshold equilibrium (v^*, u^*) , if μ_i are the eigenvalues of $J(v^*, u^*)$, the linearized explicit Euler update is locally stable when

$$|1 + \Delta t \mu_i| < 1, \quad i = 1, 2. \quad (14)$$

Equivalently, for eigenvalues with negative real part, a sufficient step-size restriction is

$$0 < \Delta t < \min_i \frac{-2\text{Re}(\mu_i)}{|\mu_i|^2}. \quad (15)$$

This condition is local and does not constitute a global convergence proof for the hybrid reset system.

This model provided a reference with greater phenomenological richness, useful for examining the trade-offs among dynamic realism, arithmetic cost, memory, and reduced-precision tolerance [18], [19], [24].

2.3.3 Simplified Digital ShiftLIF-type Model

The third family was a simplified digital *ShiftLIF*-type approximation, intended to preserve integration-and-threshold logic while replacing costly operations with dynamics better suited to minimal-arithmetic implementation. Its discrete update was formulated as

$$x_{k+1} = x_k + \lambda_{hw}(x_{rest} - x_k) + g_I I_k, \quad (16)$$

where x_k denotes the normalized internal state associated with the spiking logic of the model, and the effective leakage factor was obtained from:

$$\lambda = \frac{\Delta t}{\tau_m}, \quad (17a)$$

$$\lambda_{hw} = 2^{\text{round}(\log_2 \lambda)}. \quad (17b)$$

That is, the decay parameter was approximated as a power of 2 to favor future implementation using binary shifts. The spiking rule was

$$\text{if } x_{k+1} \geq x_{th}, \quad s_{k+1} = 1, \quad x_{k+1} \leftarrow x_{reset}. \quad (18)$$

The preceding update can also be interpreted as an affine subthreshold map:

$$x_{k+1} = (1 - \lambda_{hw})x_k + \lambda_{hw}x_{rest} + g_I I_k. \quad (19)$$

This expression does not modify the simulation model; it only makes explicit the mathematical structure of the same update rule used in the benchmark. Under the operating condition $0 < \lambda_{hw} \leq 1$, the subthreshold map is contractive with respect to the state variable, since:

$$|x_{k+1} - y_{k+1}| = |1 - \lambda_{hw}| |x_k - y_k|, \quad |1 - \lambda_{hw}| < 1, \quad (20)$$

for two trajectories x_k and y_k driven by the same input sequence. Therefore, differences due only to initial conditions decay geometrically while the state remains below threshold.

The same structure also supports bounded-input bounded-state behavior. If the input sequence is bounded, $|I_k| \leq I_{max}$, then the subthreshold state satisfies

$$|x_k - x_{rest}| \leq |1 - \lambda_{hw}|^k |x_0 - x_{rest}| + \frac{|g_I| I_{max}}{1 - |1 - \lambda_{hw}|} \quad (21)$$

Thus, bounded inputs generate bounded subthreshold states. In addition, the reset operation

preserves boundedness because every threshold crossing maps the state to the finite value x_{reset} .

For the noisy input already defined in the benchmark, $I_k = I_{base,k} + \eta \sigma_I \xi_k$, with bounded non-vanishing perturbation $|\xi_k| \leq \xi_{max}$, exact deterministic convergence to a single fixed point is not generally claimed. Instead, the deviation induced by the persistent noise remains ultimately bounded, with

$$\limsup_{k \rightarrow \infty} |e_k| \leq \frac{|g_I| \eta \sigma_I \xi_{max}}{1 - |1 - \lambda_{hw}|}, \quad (22)$$

where e_k denotes the difference between the noisy and noiseless subthreshold trajectories associated with the same deterministic input component. Hence, under persistent bounded noise, the *ShiftLIF* state is expected to evolve within a bounded operating tube rather than converge exactly to a deterministic fixed point. This provides a formal interpretation of the boundedness, subthreshold contraction, and input-to-state stability properties of the same *ShiftLIF* update used in the numerical experiments.

This formulation enabled explicit study of the trade-off between digital simplicity and loss of continuous dynamic richness, consistent with recent neuromorphic co-design trends centered on low-precision arithmetic and model adaptation to the implementation substrate [22], [25], [26].

2.4 Discretization, Quantization, and Numerical Reproducibility

The three models were integrated in discrete time over the same temporal grid. The comparison maintained the same stimulus structure for each combination $(m, s, \eta, \Delta t, B)$, where m denotes the model, s the stimulus, η the noise level, Δt the temporal step, and $B \in \mathcal{B}$ the selected quantization regime. The reference simulation was obtained with $\Delta t_{ref} = 0.02$ ms and double precision, and was used to compute temporal errors and aggregated metrics.

The quantization of inputs and states was implemented by uniform quantization over bounded operating intervals. For a scalar variable $x \in [x_{min}, x_{max}]$ and finite B , the employed operator was

$$Q_B(x) = x_{min} + \Delta_Q \text{round}\left(\frac{x - x_{min}}{\Delta_Q}\right), \quad (23)$$

$$\Delta_Q = \frac{x_{max} - x_{min}}{2^{B-1}}.$$

For $B = \infty$, the corresponding run was interpreted as a double-precision simulation without fixed-point quantization. This procedure made it possible to represent the effect of reduced precision

on first-spike timing, firing rate, spike loss, and trajectory shape, in agreement with recent studies highlighting the role of quantization and algorithm-hardware co-design in resource-constrained SNN deployment [19], [25], [26].

Reproducibility was further ensured by fixing a global seed, storing complete result tables, exporting figures, and recording each scenario, so that the evaluation rested on auditable and numerically aggregable outputs rather than on isolated visual inspection.

2.5 Dynamic and Numerical Evaluation Metrics

The evaluation relied on complementary metrics to avoid drawing conclusions from a single observable.

2.5.1 First-spike Latency

The temporal precision of the first spike was defined as

$$e_m^{(1)}(\Delta t) = |t_{m,1}(\Delta t) - t_{m,1}(\Delta t_{\text{ref}})|, \quad (24)$$

where $t_{m,1}(\Delta t)$ is the time of the first spike of the model m under a given temporal step, and $t_{m,1}(\Delta t_{\text{ref}})$ is the fine-discretization reference. This metric summarizes the combined effect of subthreshold dynamics, excitation, and temporal discretization.

2.5.2 Normalized-state RMSE

Since the models do not share the same internal-state semantics or operating range, trajectory comparison was performed using a normalized variable. If $x_m(k)$ denotes the comparable state magnitude of model m - membrane voltage in LIF and Izhikevich, and functional state in ShiftLIF -, the following definition was adopted:

$$\tilde{x}_m(k) = \frac{x_m(k) - x_{m,\min}}{x_{m,\max} - x_{m,\min}}. \quad (25)$$

With this normalization, the root mean square error was calculated as:

$$\text{RMSE}_m = \sqrt{\frac{1}{N} \sum_{k=1}^N (\tilde{x}_m(k) - \tilde{x}_{m,\text{ref}}(k))^2}, \quad (26)$$

where $\tilde{x}_{m,\text{ref}}(k)$ is the reference trajectory interpolated over the analyzed grid. This definition is the basis of the normalized-state RMSE reported later.

2.5.3 Mean Firing Rate and Temporal Variability

For each run, the mean firing rate was estimated as

$$\text{FR}_m = \frac{N_{\text{spikes}}}{T_{\text{end}}/1000}, \quad (27)$$

and, when the spike train was sufficiently long, the coefficient of variation of the inter-spike intervals was also calculated:

$$\text{CV}_{\text{ISI}} = \frac{\sigma(\Delta \text{ISI})}{\mu(\Delta \text{ISI})}. \quad (28)$$

This second measure helped distinguish more regular from more irregular responses and contributed to the dynamic richness index.

2.5.4 Computation Time

The execution time per run was obtained directly in MATLAB for each experimental combination. Although environment-dependent, its relative comparison within the same benchmark remains informative for identifying trends in arithmetic load.

2.6 Robustness Evaluation: Discretization, Noise, and Parametric Sensitivity

Stability with respect to the temporal step size was evaluated using an aggregate index. If $q \in \{e^{(1)}, \text{RMSE}, \text{FR}\}$ denotes a representative observable of the model m when varying Δt , the following was defined:

$$\text{CV}_{q,m} = \frac{\sigma(q_m(\Delta t))}{\mu(q_m(\Delta t)) + \varepsilon}, \quad (29)$$

and subsequently, the global discretization sensitivity index as

$$S_{\Delta t}^{(m)} = \frac{1}{3} (\text{CV}_{e^{(1)},m} + \text{CV}_{\text{RMSE},m} + \text{CV}_{\text{FR},m}), \quad (30)$$

where $\varepsilon > 0$ is a small constant used to avoid indeterminate divisions. Lower values of $S_{\Delta t}^{(m)}$ indicate lower sensitivity to changes in the temporal grid.

Noise robustness was estimated by comparing degradation at the minimum and maximum noise levels. If r_m represents the mean firing rate and ρ_m the RMSE at both extremes of the noise set, an average degradation and an associated robustness were defined as:

$$D_{\eta}^{(m)} = \frac{1}{2} \left(\frac{|r_m^{\max} - r_m^{\min}|}{\max(|r_m^{\min}|, 1)} + \frac{|\rho_m^{\max} - \rho_m^{\min}|}{\max(\rho_m^{\min}, \varepsilon)} \right), \quad (31)$$

$$R_{\eta}^{(m)} = \frac{1}{1 + D_{\eta}^{(m)}}. \quad (32)$$

Thus, lower degradation implies greater robustness.

Parametric sensitivity was studied separately by applying controlled perturbations of 0%, 5%, and 10% to the characteristic parameters of each model. For each variation level, multiple replicates were performed, and three observables were compared:

first spike, mean rate, and total number of events. If δ denotes the perturbation level and r the corresponding replicate, the average score per variation level was obtained as

$$P_m(\delta) = \frac{1}{3R} \sum_{r=1}^R \left(\Delta e_{m,r,\delta}^{(1)} + \Delta FR_{m,r,\delta} + \Delta N_{sp,m,r,\delta} \right), \quad (33)$$

where each term $\Delta(\cdot)$ represents a normalized relative change with respect to the nominal reference. This distinction separated robustness to external noise from robustness to internal parameter uncertainty.

2.7 Hardware-aware Metrics

Beyond dynamic and numerical behavior, the benchmark included implementation-oriented metrics. For each model, the approximate arithmetic cost per iteration, the number of internal states, and the memory required per neuron were counted. The memory per neuron was estimated by

$$M_m = \frac{n_{states,m} b_{eff} + 1}{8} \text{ [bytes]}, \quad (34)$$

where $n_{states,m}$ is the number of state variables of the model m , b_{eff} corresponds to the effective number of bits per state, and the additional term represents spike-event storage. For the comparative summary reported in the results, $b_{eff} = 25$ effective bits was adopted.

Quantization sensitivity was calculated by comparing the double-precision regime with the lowest-precision regime through relative changes in mean firing rate and normalized-state RMSE:

$$Q_m = \frac{1}{2} \left(\frac{|\overline{FR}_m^{(\infty)} - \overline{FR}_m^{(8)}|}{\max(|\overline{FR}_m^{(\infty)}|, 1)} + \frac{|\overline{RMSE}_m^{(\infty)} - \overline{RMSE}_m^{(8)}|}{\max(\overline{RMSE}_m^{(\infty)}, \varepsilon)} \right). \quad (35)$$

Complementarily, for the 8-bit regime, complete spike loss was recorded through:

$$L_m^{(8)} = \frac{N_{no\ spike,m}^{(8)}}{N_{runs,m}^{(8)}} \times 100. \quad (36)$$

This criterion is especially relevant because, in low-cost neuromorphic platforms, precision reduction is usually a practical design condition rather than a secondary detail [17], [25], [26].

Ease of implementation was assigned an ordinal score informed by the algebraic structure of each model, privileging formulations based mainly on accumulation, comparison, reset, and simple operations. The purpose was not to replace hardware synthesis, but to incorporate an explicit variable related to future minimal implementation.

2.8 Dynamic Richness Index and Multicriteria Aggregation

To avoid reducing the comparison to a single isolated magnitude, a bounded dynamic-richness index was defined from two complementary components: i) the diversity of response classes observed under the different stimuli and ii) the dispersion of the mean firing rates. If C_m denotes the number of behavior classes identified for model m , $C_{max} = \max_m C_m$, $\sigma_{r,m}$ denotes the dispersion of the mean firing rates for model m , and $\sigma_{r,max} = \max_m \sigma_{r,m}$, the dynamic-richness operator was defined as

$$\Gamma_m = \alpha \frac{C_m}{C_{max}} + (1 - \alpha) \min \left\{ \frac{\sigma_{r,m}}{\sigma_{r,max} + \varepsilon}, 1 \right\}, \quad 0 \leq \alpha \leq 1, \varepsilon > 0. \quad (37)$$

In this study, $\alpha = 0.6$ was adopted to give slightly greater relevance to the diversity of response classes than to rate dispersion. Since both components are normalized in the interval $[0,1]$ and the coefficients define a convex combination, the resulting descriptor satisfies

$$0 \leq \Gamma_m \leq 1. \quad (38)$$

Thus, Γ_m is not treated as an arbitrary score, but as a bounded composite operator that translates qualitative differences in excitability into a normalized quantity compatible with the subsequent multicriteria aggregation.

All remaining metrics were also mapped into normalized utilities in the interval $[0,1]$. For benefit criteria, where larger values indicate better performance, and for cost criteria, where smaller values indicate better performance, the normalized utilities were respectively defined as

$$z_{m,k}^+ = \frac{\mathcal{Y}_{m,k} - \mathcal{Y}_{k,min}}{\mathcal{Y}_{k,max} - \mathcal{Y}_{k,min} + \varepsilon}, \quad (39a)$$

$$z_{m,k}^- = \frac{\mathcal{Y}_{k,max} - \mathcal{Y}_{m,k}}{\mathcal{Y}_{k,max} - \mathcal{Y}_{k,min} + \varepsilon}. \quad (39b)$$

After normalization, each utility was restricted to the closed interval $[0,1]$, so that

$$0 \leq z_{m,k} \leq 1. \quad (40)$$

If $z_{m,k}$ denotes the normalized utility of model m under criterion k , the global score in scenario s was defined as the convex aggregation operator

$$S_m^{(s)} = \sum_{k=1}^K w_k^{(s)} z_{m,k}, \quad w_k^{(s)} \geq 0, \quad \sum_{k=1}^K w_k^{(s)} = 1. \quad (41)$$

Therefore, the global score also satisfies:

$$0 \leq S_m^{(s)} \leq 1. \quad (42)$$

This establishes the multicriteria score as a bounded, monotone, scenario-dependent utility operator. Its monotonicity follows from the non-negativity of the weights: improving any normalized criterion $z_{m,k}$ while keeping the others fixed cannot decrease $S_m^{(s)}$. Consequently, the ranking is structurally interpretable and remains explicitly linked to the selected scenario weights rather than to an unconstrained empirical formula.

Four scenarios were considered: balanced multicriteria scenario, maximum simplicity, dynamic richness, and lightweight hardware. In operational terms, the balanced multicriteria scenario distributed weights evenly; the maximum-simplicity scenario prioritized computational cost, memory, and ease of implementation; the dynamic-richness scenario emphasized temporal precision, response variability, and the bounded dynamic-richness index; and the lightweight-hardware scenario privileged operations per iteration, memory, and quantization tolerance. This procedure is consistent with contemporary recommendations favoring context-sensitive benchmarking over rigid universal hierarchies [16], [17], [20].

2.9 Methodological Traceability with Respect to the Results

The results reported later derive directly from this protocol: temporal curves and rasters come from the PulseTrain reference condition; first-spike errors were computed against the fine reference run; RMSE was calculated on normalized states to preserve comparability among models with different internal semantics; sensitivity to Δt , noise, and parameter perturbation was obtained by systematic aggregation across scenarios; and the hardware-aware indicators arose from the discretized equations and quantization scheme. Accordingly, the reported evidence should be read as the coherent outcome of a unified simulation, measurement, and multicriteria aggregation framework designed to compare dynamic plausibility and technological convenience within the same benchmark [16], [17], [19].

3 Results

This section summarizes the experimental evidence by integrating dynamic response, numerical performance, computational cost, and *hardware-aware* feasibility. The comparison among LIF, Izhikevich, and ShiftLIF is therefore interpreted as a criterion-dependent analysis rather than as a search for an absolute “best model,” in agreement with

recent views that spiking-model evaluation should jointly consider dynamic fidelity, computational efficiency, scalability, and implementation feasibility [17], [18], [19], [21], [24].

In aggregate, Izhikevich provides the best first-spike precision and the greatest robustness to noise; LIF yields the lowest normalized-state RMSE and the lowest mean computation time; and ShiftLIF becomes the most favorable option when arithmetic simplicity, low memory, reduced parametric sensitivity, and ease of implementation are weighted jointly [17], [19], [20], [21], [22]. Table 1 and Table 2 in Appendix summarize the quantitative basis for the subsequent analysis.

Together, the two tables expose the central trade-off of the study. Izhikevich leads temporal precision and noise robustness but requires the highest arithmetic load and memory footprint. LIF minimizes normalized-state RMSE and mean computation time, but is strongly penalized under severe quantization. ShiftLIF does not minimize local temporal-error metrics, but offers the best combination of low arithmetic cost, low memory, reduced parametric sensitivity, and ease of implementation, consistent with current neuromorphic benchmarking trends [17], [19], [20], [22], [27].

3.1 Comparative Dynamic Response

3.1.1 Temporal Evolution of the Membrane Potential

Figure 1 shows that, under the same *pulse train* stimulus, the internal trajectories differ substantially across models. LIF exhibits classical leaky-integration behavior without reaching threshold in the representative condition; Izhikevich shows abrupt peaks and recovery, reflecting richer excitability; and ShiftLIF operates in a bounded digital regime whose normalized state should not be interpreted as biological voltage.




Fig. 1: Representative temporal responses of the three models to a pulse-train stimulus. In ShiftLIF, the normalized state associated with the spiking logic is shown.

Source: created by the authors

This confirms that the variables being compared are not semantically equivalent. In continuous models, waveform morphology retains interpretive value, whereas in shift-based models, the emphasis shifts to event timing, stability, and implementation cost [18], [21], [23], [24].

3.1.2 Obtained Spiking Patterns

Figure 2 complements the temporal comparison by focusing on event timing. In the illustrated condition, LIF remains subthreshold, Izhikevich generates isolated events, and ShiftLIF produces repeated packets of spikes within each excitation window. These patterns reflect distinct coding logics: richer event morphology in Izhikevich, interpretable leaky integration in LIF, and operational regularity in ShiftLIF.




Fig. 2: Spike raster for a representative condition
Source: created by the authors

This does not imply that denser or sparser rasters are intrinsically superior; rather, each model expresses a spiking regime consistent with its design objective. For low-cost neuromorphic computing, shift-based representations are attractive because they replace costly operations with comparisons, shifts, and accumulations [19], [22], [23].

3.1.3 Most Significant Dynamic Differences

The main dynamic difference lies in how excitability is distributed over time. Izhikevich clearly separates subthreshold integration, spike peak, and recovery, which favors local morphological expressiveness. LIF remains an intermediate leaky-integrator reference, whereas ShiftLIF regularizes spiking into a logic more favorable for digital computation. This contrast must be interpreted together with the aggregated metrics and the intended technological objective [17], [18], [19], [20].

3.2 Numerical Performance

3.2.1 Temporal Precision of Spiking

The temporal precision of the first spike was evaluated through

$$e_m^{(1)}(\Delta t) = |t_{m,1}(\Delta t) - t_{m,1}(\Delta t_{\text{ref}})|, \quad (43)$$

where $t_{m,m}$ represents the time of the first spike of model m , and Δt_{ref} is the finest reference temporal step. Figure 3 shows monotonic degradation as Δt increases, though the slopes differ across models. In the representative condition, LIF exhibits the lowest visible error, Izhikevich remains within the sub-millisecond-to-millisecond range, and ShiftLIF shows the largest displacement of the first event.




Fig. 3: Temporal error of the first spike as a function of the temporal step
Source: created by the authors

The aggregated reading differs from this local observation. Once noise, parametric perturbations, and quantization are integrated, the lowest mean first-spike error corresponds to Izhikevich, as reported in Table 1 (Appendix) 0.809 ms versus 32.154 ms for LIF and 84.770 ms for ShiftLIF. This distinction avoids extrapolating a single representative figure to the entire experimental space [17], [18], [20].

3.2.2 Stability with Respect to the Temporal Step

Sensitivity to the temporal step, summarized in Table 1 (Appendix), was 0.419 for LIF, 0.268 for Izhikevich, and 0.177 for ShiftLIF. Thus, although ShiftLIF presents the largest first-spike error in the representative figure and in the global mean of that specific indicator, it is the least sensitive model when several outputs are aggregated. This broadens the notion of numerical stability: a model may be less precise in one observable and still more robust at the operational level [17], [20], [27].

3.2.3 Parametric Sensitivity and Numerical Robustness

Parametric sensitivity is summarized in Table 3 (Appendix). For a 5% variation, the mean sensitivity was 0.1341 for LIF, 0.0489 for Izhikevich, and 0.0017 for ShiftLIF. For a 10% variation, the values rose to 0.4650, 0.0781, and 0.0372, respectively. Thus, sensitivity increases with perturbation amplitude in all models, while the relative order remains unchanged, favoring ShiftLIF under controlled parameter variations.

Noise robustness follows a different order in Table 1 (Appendix) Izhikevich reaches 0.981, followed by ShiftLIF at 0.846 and LIF at 0.741. This shows that ShiftLIF is especially robust to parametric perturbation, whereas Izhikevich is more resistant to external noise [19], [20], [24].

3.3 Computational Complexity and Memory

3.3.1 Arithmetic Load Per Iteration

Figure 4 and Table 2 (Appendix) show the relation between arithmetic load and execution time. Izhikevich requires 14 operations per iteration and a mean time of 3.684 ms; LIF requires 6 operations and 1.875 ms; and ShiftLIF requires 4 operations and 2.145 ms. Thus, LIF was the fastest model in the MATLAB benchmark, whereas ShiftLIF achieved the lowest arithmetic count.




Fig. 4: Mean computational cost by model

Source: created by the authors

These results indicate that per-neuron arithmetic complexity explains the higher cost of Izhikevich, but does not fully determine the difference between LIF and ShiftLIF, where implementation details of the update logic also matter [19], [22], [27].

3.3.2 Number of States and Storage Cost

The memory footprint confirms the same general trend. As shown in Table 2 (Appendix), Izhikevich uses 6.375 bytes per neuron, whereas LIF and ShiftLIF require 3.250 bytes. This difference follows directly from model structure: Izhikevich requires more state variables and therefore more storage per functional unit.

In large-scale neuromorphic systems, this difference propagates approximately linearly with the number of neurons and synapses, affecting occupied area, memory traffic, and energy consumption. This observation agrees with recent work emphasizing the importance of efficient neuronal and synaptic structures for large-scale or resource-limited edge implementations [21], [28].

3.3.3 Theoretical Scalability

If N denotes the number of neurons and T the discrete temporal horizon, the total simulation cost can be approximated as:

$$C_m \propto NT\kappa_m, \quad (44)$$

where κ_m represents the per-update cost constant of model m . The three models grow linearly with N and T , but differences in κ_m lead to different practical scalability. Because Izhikevich combines more operations and greater memory, its effective growth is less favorable for massive deployment. ShiftLIF minimizes arithmetic complexity per iteration and is attractive for large-scale networks or embedded accelerators, although LIF yielded the lowest mean execution time in the present simulation environment [19], [22], [27].

3.4 Hardware-aware Evaluation

3.4.1 Sensitivity to Quantization and Fixed-Point Arithmetic

Sensitivity to quantization was one of the most discriminating criteria of the study. According to Table 2 (Appendix), the aggregated values were 3.9891 for LIF, 0.4745 for Izhikevich, and 0.3084 for ShiftLIF, where lower values indicate greater tolerance to reduced precision. This order agrees with Figure 5, where ShiftLIF dominates the normalized quantization-tolerance axis, and LIF is clearly penalized.

Table 4 (Appendix) reinforces this finding through complete spike loss in the 8-bit regime. LIF lost the spike in 127 of 180 runs, equivalent to 70.6%. ShiftLIF lost the event in 72 of 180 runs, or 40.0%. Izhikevich, by contrast, exhibited no complete loss of spikes under the same conditions. This result is practically relevant for low-cost

platforms based on fixed-point arithmetic [19], [20],

3.4.2 Estimated Implementation Complexity

The assigned ease of implementation was 5 for LIF, 2 for Izhikevich, and 5 for ShiftLIF, as reported in Table 2 (Appendix). This reflects their algebraic structure: LIF and ShiftLIF can be implemented through accumulation, comparison, and reset operations, whereas Izhikevich requires additional states and nonlinear dynamics. Therefore, implementation difficulty remains a relevant comparative criterion when future migration to neuromorphic hardware is considered [21], [22].

3.4.3 Technological Feasibility of Each Model Family

Figure 5 summarizes the technological comparison across operations, memory, quantization sensitivity, and ease of implementation. ShiftLIF combines the lowest arithmetic cost, the same memory footprint as LIF, the best aggregate tolerance to quantization, and the highest implementation score. LIF preserves simplicity and the lowest observed mean execution time, but its vulnerability to severe quantization penalizes it. Izhikevich retains advantages in temporal precision and noise robustness while preserving spiking activity across all 8-bit runs, albeit at higher arithmetic and memory costs [19], [20], [21], [22].




Fig. 5: Normalized hardware-aware comparison
Source: created by the authors

3.5 Final Index and Comparative Ranking

3.5.1 Global Result of the Proposed Index

The scenario-dependent multicriteria score used in the results follows the bounded convex aggregation operator defined in Section 2.8. For consistency, it is written as:

[23].

$$S_m^{(s)} = \sum_{k=1}^K w_k^{(s)} z_{m,k}, \quad z_{m,k} \in [0,1], \quad w_k^{(s)} \geq 0, \quad \sum_{k=1}^K w_k^{(s)} = 1. \quad (45)$$

Thus, each score $S_m^{(s)}$ is bounded in $[0, 1]$ and represents a normalized scenario-dependent utility rather than an unconstrained empirical index. As shown in Table 5 (Appendix), ShiftLIF obtained the highest score in the balanced scenario (0.691), followed by LIF (0.615) and Izhikevich (0.502).

3.5.2 Ranking by Application Scenario

Across the four scenarios, the ranking remained stable. ShiftLIF ranked first in the balanced, maximum-simplicity, dynamic-richness, and lightweight-hardware scenarios, with scores of 0.691, 0.851, 0.696, and 0.786, respectively. LIF consistently ranked second, while Izhikevich ranked third. The advantage of ShiftLIF was broader when simplicity or lightweight hardware was prioritized, and narrower in the dynamic-richness scenario, indicating that the result is stable but still dependent on the intended application context [17], [20].

3.5.3 Ranking Robustness

Figure 6 visually confirms the stable order ShiftLIF > LIF > Izhikevich across the evaluated scenarios. This supports one of the main findings of the study: the final ranking does not depend on accidental weighting, but on the joint effect of arithmetic simplicity, memory efficiency, implementation ease, and acceptable tolerance to perturbations and discretization.




Fig. 6: Multicriteria ranking by scenario
Source: created by the authors

Overall, no model is universally superior. Izhikevich retains the best aggregated first-spike precision and noise robustness; LIF yields the

lowest normalized-state RMSE and mean computation time; and ShiftLIF emerges as the most consistent option for a plausible, reproducible, and hardware-oriented digital approximation for low-cost neuromorphic hardware [17], [19], [20].

4 Discussion

The results confirm that comparison among LIF, Izhikevich, and ShiftLIF cannot be reduced to a single metric, because each model family expresses a different balance among dynamic realism, numerical stability, and implementation convenience. This supports a *hardware-aware* interpretation of spiking-neuron selection, in which functional fidelity, computational efficiency, reduced-precision tolerance, and deployment context must be considered jointly [16], [17], [19].

4.1 Scientific Interpretation of the Results

From a mathematical standpoint, the observed differences follow directly from the state structure of each model. LIF is governed by first-order leaky-and-reset dynamics, which makes it simple and interpretable but limits its internal expressiveness. Izhikevich introduces a quadratic nonlinearity and a recovery variable, producing richer transitions among integration, spike generation, and recovery. ShiftLIF replaces part of the continuous decay with a binary-shift-adapted update, sacrificing phenomenological detail in exchange for operational regularity and digital simplicity. Thus, the greater local expressiveness of Izhikevich and the greater implementation simplicity of ShiftLIF are expected consequences of their constitutive equations [18], [24].

The numerical behavior is also consistent with the algebraic form of each update. Explicit Euler provides a homogeneous comparison framework, but not identical internal sensitivity. For Izhikevich, the stability statement must be interpreted locally, under bounded subthreshold operation and admissible step sizes, rather than as global convergence of the hybrid reset system. For ShiftLIF, the formal analysis in Section 2.3.3 shows that, under $0 < \lambda_{hw} \leq 1$, the subthreshold map is contractive, bounded inputs generate bounded states, and persistent bounded noise leads to an ultimately bounded operating tube rather than exact deterministic convergence. Therefore, algebraic simplicity does not imply universal robustness, but the shift-based structure gives ShiftLIF favorable aggregate behavior under discretization and reduced precision [17], [20], [25].

Overall, expressiveness and simplicity define a design frontier. Izhikevich provides higher local dynamic richness and better aggregated first-spike precision, but requires more operations, more memory, and a less implementation-friendly structure. LIF remains a clear and efficient intermediate reference, with low normalized-state RMSE and the lowest mean execution time in the MATLAB benchmark. ShiftLIF, although less expressive, converts its structural simplicity into lower arithmetic cost, lower parametric sensitivity, lower aggregate discretization sensitivity, and better implementation attributes. Model selection therefore depends on the intended balance between dynamic richness and computational economy [18], [19], [24].

4.2 Computational and Technological Meaning

A key methodological point is that local figures and global metrics must not be confused. The temporal and raster plots correspond to a representative *PulseTrain* condition, whereas the main conclusions are drawn from aggregating across stimuli, temporal steps, noise, quantization, and parametric perturbations. This explains why LIF may appear favorable in a local first-spike figure, while Izhikevich achieves the best aggregated temporal precision, and why ShiftLIF must be interpreted not as the most biologically expressive model, but as a regularized spiking logic oriented toward low-cost implementation [17], [20].

For large-scale simulation and lightweight hardware, the preferred model depends on the dominant criterion. If software execution time in the present benchmark is prioritized, LIF is the fastest alternative. If structural scalability, arithmetic cost, memory footprint, and future migration to minimal digital hardware are emphasized, ShiftLIF becomes the most favorable option because it uses fewer operations per iteration and supports implementations based on shifts, accumulation, comparison, and reset. Its advantage is reinforced by its low aggregate quantization sensitivity and high ease-of-implementation score. Nevertheless, ShiftLIF is not invulnerable: under 8-bit quantization, it still exhibited complete spike loss in part of the runs, whereas Izhikevich preserved spiking activity but at substantially higher arithmetic and memory cost. Thus, ShiftLIF should be understood as the best technological trade-off, not the most precise or robust model on every criterion [17], [19], [21], [22], [25].

4.3 Relationship with the State of the Art and Scope of the Contribution

The study aligns with recent trends favoring contextual benchmarking, algorithm–hardware co-design, and the integrated evaluation of neuromorphic models under edge-computing constraints [16], [17], [19], [20], [27]. Its methodological distinction lies in exposing the three models to the same stimulus families, temporal steps, noise levels, quantization regimes, and scenario-based aggregation. In the revised formulation, the multicriteria stage is expressed as a bounded convex aggregation operator, so that the final scores remain normalized, monotone with respect to the criteria, and explicitly linked to the selected application scenario. This transforms the comparison into a decision-support framework rather than a purely descriptive classification [17], [20].

The contribution is therefore theoretical-methodological and *hardware-aware*. The article does not claim validation on a real FPGA, ASIC, or embedded device, nor direct measurement of physical energy consumption, occupied area, communication latency, or maximum operating frequency. Instead, it provides a reproducible *in silico* framework that predicts how the relative advantages of LIF, Izhikevich, and ShiftLIF change when implementation constraints are accounted for. The bounded multicriteria formulation, the local Euler-stability clarification for Izhikevich, and the formal stability interpretation of ShiftLIF strengthen this contribution without changing the computational scope of the evidence [17], [19], [20].

4.4 Limitations and Outlook

The main limitations are the restricted model set and the validation level. The comparison does not include other relevant SNN families, such as AdEx, QIF, or additional hybrid approximations, which could occupy intermediate positions between expressiveness and simplicity [18], [24]. Moreover, although the protocol is reproducible as an *in silico* benchmark, there is still no verification on a real FPGA, ASIC, or specific embedded platform; therefore, physical metrics such as energy, LUT usage, latency, clock frequency, or consumption per spike remain outside the empirical scope of the study [17], [22].

Future work should implement the framework on reconfigurable hardware and compare the synthetic metrics with physical measurements. It should also extend the benchmark to additional neuron models, plasticity mechanisms, and network-level tasks to determine whether the observed advantage of

ShiftLIF persists when moving from isolated-neuron analysis to collective behavior and concrete neuromorphic applications [20], [22], [25].

5 Conclusion

This study developed a reproducible, multicriteria, and hardware-oriented framework to compare LIF, Izhikevich, and ShiftLIF under homogeneous conditions of excitation, discretization, noise, quantization, and parametric perturbation. The results confirm that model suitability cannot be determined by a single indicator but by the joint balance among dynamic expressiveness, numerical precision, robustness, computational cost, and technological feasibility.

Izhikevich showed the greatest local dynamic richness, the best aggregated first-spike precision, and the highest robustness to noise, making it suitable when expressive neuronal behavior is prioritized. LIF achieved the lowest normalized-state RMSE and the lowest mean computation time in the MATLAB benchmark, preserving its value as a simple, interpretable, and computationally efficient reference. ShiftLIF did not lead the pointwise precision indicators, but offered the best overall trade-off when arithmetic cost, memory, parametric sensitivity, discretization sensitivity, quantization tolerance, and ease of implementation were considered jointly. Its advantage remained stable across the multicriteria scenarios, supporting its suitability for low-cost neuromorphic hardware where sufficient functional plausibility must be achieved with reduced technological requirements [16], [17], [20].

The contribution of this work is therefore theoretical-methodological and hardware-aware: it provides an objective procedure for early neuron-model selection before physical validation on FPGA, ASIC, or other reconfigurable platforms [19], [25]. Future work should extend the framework toward real hardware validation, additional neuron models and plasticity mechanisms, and applications in event-based vision, embedded control, and neuromorphic sensing.

Acknowledgement:

The authors acknowledge the academic and scientific community whose previous contributions to neuromorphic computing inspired this work.

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

The authors wrote, reviewed and edited the content as needed and they have not utilised artificial intelligence (AI) tools. The authors take full responsibility for the content of the publication.

References:

- [1] D. R. Muir and S. Sheik, The Road to Commercial Success for Neuromorphic Technologies, *Nature Communications*, Vol.16, 2025, Art. no. 3586, doi: 10.1038/s41467-025-57352-1.
- [2] A. Ahmadi, S. Gomar, and M. Ahmadi, Editorial: From Theory to Practice: The Latest Developments in Neuromorphic Computing Applications, *Frontiers in Neuroscience*, Vol.18, 2024, Art. no. 1511987, doi: 10.3389/fnins.2024.1511987.
- [3] K. Kim, M. S. Song, H. Hwang, S. Hwang, and H. Kim, A Comprehensive Review of Advanced Trends: From Artificial Synapses to Neuromorphic Systems With Consideration of Non-Ideal Effects, *Frontiers in Neuroscience*, Vol.18, 2024, Art. no. 1279708, doi: 10.3389/fnins.2024.1279708.
- [4] L. Deng, H. Tang, and K. Roy, Editorial: Understanding and Bridging the Gap Between Neuromorphic Computing and Machine Learning, Volume II, *Frontiers in Computational Neuroscience*, Vol.18, 2024, Art. no. 1455530, doi: 10.3389/fncom.2024.1455530.
- [5] C. Zhou, H. Zhang, L. Yu, Y. Ye, Z. Zhou, L. Huang, Z. Ma, X. Fan, H. Zhou, and Y. Tian, Direct Training High-Performance Deep Spiking Neural Networks: A Review of Theories and Methods, *Frontiers in Neuroscience*, Vol.18, 2024, Art. no. 1383844, doi: 10.3389/fnins.2024.1383844.
- [6] C. Cimorelli, J. A. Millan-Romera, H. Voos, and J. L. Sanchez-Lopez, Hardware, Algorithms, and Applications of the Neuromorphic Vision Sensor: A Review, *Sensors*, Vol.25, No.19, 2025, Art. no. 6208, doi: 10.3390/s25196208.
- [7] H. Li, D. Kumar, and N. El-Atab, A Neuromorphic Event Data Interpretation Approach With Hardware Reservoir, *Frontiers in Neuroscience*, Vol.18, 2024, Art. no. 1467935, doi: 10.3389/fnins.2024.1467935.
- [8] A. Renner, F. Sheldon, A. Zlotnik, L. Tao, and A. Sornborger, The Backpropagation Algorithm Implemented on Spiking Neuromorphic Hardware, *Nature Communications*, Vol.15, 2024, Art. no. 9691, doi: 10.1038/s41467-024-53827-9.
- [9] A. Siddique, M. I. Vai, and S. H. Pun, A Low Cost Neuromorphic Learning Engine Based on a High Performance Supervised SNN Learning Algorithm, *Scientific Reports*, Vol.13, 2023, Art. no. 6280, doi: 10.1038/s41598-023-32120-7.
- [10] S. Daddinounou and E.-I. Vatajelu, Bi-Sigmoid Spike-Timing-Dependent Plasticity Learning Rule for Magnetic Tunnel Junction-Based SNN, *Frontiers in Neuroscience*, Vol.18, 2024, Art. no. 1387339, doi: 10.3389/fnins.2024.1387339.
- [11] S. Wang, S. Gao, C. Tang, et al., Memristor-Based Adaptive Neuromorphic Perception in Unstructured Environments, *Nature Communications*, Vol.15, 2024, Art. no. 4671, doi: 10.1038/s41467-024-48908-8.
- [12] J. Im, S. Pak, S.-Y. Woo, W. Shin, and S.-T. Lee, Flash Memory for Synaptic Plasticity in Neuromorphic Computing: A Review, *Biomimetics*, Vol.10, No.2, 2025, Art. no. 121, doi: 10.3390/biomimetics10020121.
- [13] I. Tsilikas, A. Tsirigotis, G. Sarantoglou, S. Deligiannidis, A. Bogris, and C. Posch, Photonic Neuromorphic Accelerators for Event-Based Imaging Flow Cytometry, *Scientific Reports*, Vol.14, 2024, Art. no. 24179, doi: 10.1038/s41598-024-75667-9.
- [14] J. Bartels, O. Gallou, H. Ito, M. Cook, J. Sarnthein, G. Indiveri, and S. Ghosh, Event Driven Neural Network on a Mixed Signal Neuromorphic Processor for EEG-Based Epileptic Seizure Detection, *Scientific Reports*, Vol.15, 2025, Art. no. 15965, doi: 10.1038/s41598-025-99272-6.
- [15] A. Dahiya, S. Rani, S. Kumar, and T. Prodromakis, Low-Voltage-Driven Memristor Framework for Efficient Neuromorphic Computation With STDP Functionality, *Frontiers in Nanotechnology*, Vol.7, 2025, Art. no. 1723433, doi: 10.3389/fnano.2025.1723433.
- [16] C. D. Schuman, S. R. Kulkarni, M. Parsa, J. P. Mitchell, P. Date, and B. Kay, Opportunities for Neuromorphic Computing Algorithms and Applications, *Nature Computational Science*, Vol.2, No.1, 2022, pp. 10–19, doi: 10.1038/s43588-021-00184-y.
- [17] J. Yik, K. Van den Berghe, D. den Blanken, Y. Bouhadjar, M. Fabre, P. Hueber, W. Ke, M. A. Khoei, D. Kleyko, N. Pacik-Nelson, et al., The NeuroBench Framework for Benchmarking Neuromorphic Computing

- Algorithms and Systems, Nature Communications, Vol.16, 2025, Art. no. 1545, doi: 10.1038/s41467-025-56739-4.
- [18] B. Ayasi, C. J. Carmona, M. Saleh, and A. M. García-Vico, A Practical Tutorial on Spiking Neural Networks: Comprehensive Review, Models, Experiments, Software Tools, and Implementation Guidelines, Eng, Vol.6, No.11, 2025, Art. no. 304, doi: 10.3390/eng6110304.
- [19] P. M. Ferreira, S. Wang, Y. Gao, and A. Benlarbi-Delai, A Comparative Review of Deep and Spiking Neural Networks for Edge AI Neuromorphic Circuits, Frontiers in Neuroscience, Vol.19, 2025, Art. no. 1676570, doi: 10.3389/fnins.2025.1676570.
- [20] E. Gemo, S. Spiga, and S. Brivio, SHIP: A Computational Framework for Simulating and Validating Novel Technologies in Hardware Spiking Neural Networks, Frontiers in Neuroscience, Vol.17, 2024, Art. no. 1270090, doi: 10.3389/fnins.2023.1270090.
- [21] S. Al Abdul Wahid, A. Asad, and F. Mohammadi, A Survey on Neuromorphic Architectures for Running Artificial Intelligence Algorithms, Electronics, Vol.13, No.15, 2024, Art. no. 2963, doi: 10.3390/electronics13152963.
- [22] S. Geng, Z. Wang, Z. Liu, M. Zhang, X. Zhu, and Y. Dan, Hardware Implementation of FPGA-Based Spiking Attention Neural Network Accelerator, PeerJ Computer Science, Vol.11, 2025, Art. no. e3077, doi: 10.7717/peerj-cs.3077.
- [23] K. Alammari, M. Heidarpur, M. Ahmadi, and A. Ahmadi, LIF Neuron—A Memristive Realization, Frontiers in Electronics, Vol.5, 2024, Art. no. 1366299, doi: 10.3389/felec.2024.1366299.
- [24] Sanaullah, S. Koravuna, U. Rückert, and T. Jungeblut, Exploring Spiking Neural Networks: A Comprehensive Analysis of Mathematical Models and Applications, Frontiers in Computational Neuroscience, Vol.17, 2023, Art. no. 1215824, doi: 10.3389/fncom.2023.1215824.
- [25] J. Li, M. Xu, H. Dong, B. Lan, Y. Liu, H. Chen, Y. Zhuang, Y. Xie, and L. Chen, Balancing Accuracy and Efficiency: Co-Design of Hybrid Quantization and Unified Computing Architecture for Spiking Neural Networks, Frontiers in Neuroscience, Vol.19, 2025, Art. no. 1665778, doi: 10.3389/fnins.2025.1665778.
- [26] A. Stanojevic, S. Wozniak, et al., High-Performance Deep Spiking Neural Networks With 0.3 Spikes per Neuron, Nature Communications, Vol.15, 2024, Art. no. 6793, doi: 10.1038/s41467-024-51110-5.
- [27] M. Vieth, A. Rahimi, A. G. Mohammadi, J. Triesch, and M. Ganjtabesh, Accelerating Spiking Neural Network Simulations With PymoNNto and PymoNNtorch, Frontiers in Neuroinformatics, Vol.18, 2024, Art. no. 1331220, doi: 10.3389/fninf.2024.1331220.
- [28] P. Urbizagastegui, A. van Schaik, and R. Wang, Memory-Efficient Neurons and Synapses for Spike-Timing-Dependent-Plasticity in Large-Scale Spiking Networks, Frontiers in Neuroscience, Vol.18, 2024, Art. no. 1450640, doi: 10.3389/fnins.2024.1450640.

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

Jesus Rodriguez-Flores developed the conceptual approach of the study, carried out the mathematical analysis, designed the numerical comparison framework, interpreted the results, and prepared the manuscript.

Edgart Marchan-Millan contributed to the analytical discussion of the study, supported the interpretation of the comparative results, reviewed the technical consistency of the manuscript, and participated in the revision and refinement of the final text.

Alex Moyota-Paguay contributed to the technical review of the manuscript, supported the assessment of the comparative results from an implementation-oriented perspective, and participated in the critical revision and improvement of the final version of the article.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflict of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0)

This article is published under the terms of the Creative Commons Attribution License 4.0

https://creativecommons.org/licenses/by/4.0/deed.en_US

APPENDIX

Table 1. Aggregated summary of dynamic and numerical performance by model.

Model	Mean first-spike error [ms]	Normalized-state RMSE	Sensitivity to Δt	Mean computation time [ms]	Robustness to noise
LIF	32.154	0.017	0.419	1.875	0.741
Izhikevich	0.809	0.064	0.268	3.684	0.981
ShiftLIF	84.770	0.522	0.177	2.145	0.846

Source: created by the authors

Table 2. Aggregated summary of computational complexity and hardware-aware attributes.

Model	Ops/iter per neuron	Memory [bytes]	Parametric sensitivity	Sensitivity to quantization	Ease of implementation
LIF	6	3.250	0.2995	3.9891	5
Izhikevich	14	6.375	0.0635	0.4745	2
ShiftLIF	4	3.250	0.0195	0.3084	5

Source: created by the authors

Table 3. Mean parametric sensitivity under controlled perturbations.

Model	Variation level	Mean sensitivity	Standard deviation
LIF	0.05	0.1341	0.1165
LIF	0.10	0.4650	0.3101
Izhikevich	0.05	0.0489	0.0579
Izhikevich	0.10	0.0781	0.0876
ShiftLIF	0.05	0.0017	0.0060
ShiftLIF	0.10	0.0372	0.0461

Source: created by the authors

Table 4. Complete spike loss under 8-bit quantization.

Model	Evaluated runs	Runs with no spike	Proportion with no spike [%]
LIF	180	127	70.6
Izhikevich	180	0	0.0
ShiftLIF	180	72	40.0

Source: created by the authors

Table 5. Multicriteria ranking by application scenario.

Model	Balanced multicriteria scenario	Maximum simplicity	Dynamic richness	Lightweight hardware
LIF	0.615	0.711	0.685	0.658
Izhikevich	0.502	0.294	0.427	0.397
ShiftLIF	0.691	0.851	0.696	0.786

Source: created by the authors