STM JOURNALS

Current Trends in
Signal Processing

https://journals.stmjournals.com/ctsp

ISSN: 2277-6176 (Online)
ISSN: 2321-4252 (Print)
Volume 16, Issue 2, 2026
May–August
DOI (Journal): 10.37591/CTSP

Review CTSP

Neuromorphic Spiking Neural Networks and Event-Driven Hardware: Architectures, Learning Paradigms, and Experimental Realizations

Diksha Rajput[1], Lakshya Vijayvargiya[2], Anuradha[3,*]

Abstract

With the current computational boom, the research community is seeking more sustainable energy efficient i.e. biologically inspired models of conventional artificial neural networks (ANNs). Spiking neural networks (SNNs) known as the third generation of neural network models, provide a revolutionary approach by mimicking the asynchronized event-driven and temporally accurate signaling of the mammalian brain. Whereas conventional deep learning models operate with real-valued activations and dense matrix multiplications, SNNs use binary spikes to encode and transmit information. This allows Spiking Neural Networks (SNNs) to be event driven and to be implemented with low-power neuromorphic hardware. This comprehensive research paper analyzes major neural models, learning techniques like spike timing dependent plasticity a surrogate gradient, and modern neuromorphic hardware used for event driven computing. Comparative experimental data show event-driven neuromorphic systems can achieve up to 1000× better energy efficiency during the inference than standard GPU platforms. This analysis addresses key implementation bottlenecks, including back propagation limitations, software tooling gaps and conversion accuracy losses. Ultimately, this analysis shows that combining SNNs with neuromorphic hardware has strong potential for future low power and real-time edge intelligent systems across the different fields enabling computing architectures for next-generation applications globally that are scalable, adaptable, intelligent, sustainable, autonomous, and energy-aware.

Keywords: Spiking neural networks, neuromorphic engineering, event-driven computing, low-power AI, and robotics

INTRODUCTION

The explosion of artificial intelligence technology has brought about an unforeseen crisis in terms of computational efficiency and sustainability of the environment. Deep learning frameworks have achieved remarkable results in areas such as natural language processing and computer vision; however, they have started suffering because of the "von Neumann bottleneck" [1],which refers to the physical gap between computing and memory storage. At current parameter scales in the trillions, the energy used in moving data between the two often exceeds the energy used in the computation itself. Moreover, the industry faces the problem of "Dark Silicon" where power density issues mean that not all transistors can be used together without triggering a disastrous heat threshold. This has made it necessary to move away from the synchronous frame-based approach of the second generation towards more biological frameworks of computational parsimony.

***Author for Correspondence**
Anuradha
E-mail: anuradha.cse@jecrc.ac.in

[1,2]Student, Department of Computer Science and Engineering, Jaipur Engineering College and Research Centre, Jaipur, Rajasthan, India
[3]Assistant Professor, Department of Computer Science and Engineering, Jaipur Engineering College and Research Centre, Jaipur, Rajasthan, India

Received Date: July 22, 2026
Accepted Date: August 01, 2026
Published Date: August 08, 2026

Citation: Diksha Rajput, Lakshya Vijayvargiya, Anuradha. Neuromorphic Spiking Neural Networks and Event-Driven Hardware: Architectures, Learning Paradigms, and Experimental Realizations. Current Trends in Signal Processing. 2026; 16(2): 39–46p.

The human brain can be considered the best example of efficient intelligence, consuming approximately 20 watts of energy and exceeding supercomputers in complex and unstructured tasks. The principles responsible for the efficiency of the brain consist of parallel processing, the locality and above all, the spike-based communication. Spiking Neural Networks (SNNs) are algorithms based on the functioning of biological neural systems/models [1]. In Spiking Neural Networks (SNNs), the information is represented in the form of spikes transmitted between the neurons instead of the floating-point values. The network can operate in an inactive state when there is no input data and minimizes the energy consumption in dynamically active elements of the system.

The combination of spiking neural networks with event-driven hardware is no longer a research novelty but a reality [2]. Developments in neuromorphic engineering have led to the creation of dedicated processing units [1], where memory and processing are integrated, overcoming the usual separation of memory and processing that results in a bottleneck and allowing real-time, on-chip learning through synaptic plasticity. Architectures such as Intel's Loihi 2 [3] and IBM's NorthPole [4] can process high-dimensional sensory data with microsecond latency and nanowatt idle power.

LITERATURE REVIEW

It can be noted that there has always been a tension within the field of spiking neural network studies between biological accuracy and computational simplicity. Previously, SNNs were divided into two main areas of research: computational neuroscience (which works with in-depth biophysical models) and machine learning (which works with differentiable continuous functions for optimization). The renewed interest in SNNs is a result of bridging together computational neuroscience and machine learning via advanced software and hardware infrastructure [2].

Development of SNN Software Frameworks

The development of software frameworks has been the major factor of democratization of SNN research. Initially, the most common simulators like NEURON and NEST were developed for high fidelity modeling of ion channels and synaptic dynamics. They provided essential support for the needs of neuroscientists but could not be applied efficiently in large scale ml because of their heavy computational requirements. PyTorch and JAX-based libraries have dramatically changed the situation [5] and offer means to implement SNN behavior in dl models.

Spiking neural network (SNN) software frameworks, such as snnTorch and SpikingJelly, incorporate the technique of surrogate gradient descent, which makes it possible to apply backpropagation through time (BPTT) to non-differentiable Spiking Neural Networks. As a result, deep SNN models, including Spiking ResNet Transformers, have become trainable, thus competing with ANN based solutions in solving tasks such as ImageNet and CIFAR-100. Table 1 compares the major SNN software frameworks synapses within a 65 mW power consumption [6]. Unfortunately, the lack of weight adjustment capabilities in TrueNorth means that it cannot support adaptable applications. Intel’s Loihi, released in 2018, and Loihi 2 [3] overcame this limitation by including programmable dynamics and learning abilities using local plasticity laws. By 2025, the market will have expanded to include specialty edge processors such as BrainChip Akida 2.0 [7] and IBM NorthPole [4], each targeting specific use cases from computer vision to real-time speech recognition.

Table 1. SNN software frameworks.

Framework	Learning rules	Characteristics
SpikingJelly	Surrogate Gradients, ANN-SNN Conversion	Fast GPU; CUDA required
snnTorch	BPTT, Surrogate Gradients, RTRL	Flexible; easy to use; slower than CUDA
BindsNET	STDP, R-STDP, Unsupervised Learning	Biologically realistic; complex usage
Norse	SuperSpike	Memory-efficient; complex state handling
Spyx	Surrogate Gradients	JIT optimized (TPU/GPU); newer ecosystem

ISSN: 2277-6176 (Online), ISSN: 2321-4252 (Print)

METHODOLOGY

An event-driven intelligence architecture requires a multilayered implementation methodology involving neuron modeling, information encoding, and training processes. In neuromorphic systems, the "program" is defined not by a sequence of instructions but by the connectivity, synaptic weights, and temporal constants of the neural network.

Spiking Neuron Models and Mathematical Representations

The behavior of a single neuron in an SNN can be mathematically modeled as a dynamic system using ordinary differential equations (ODES). The choice of model dictates the balance between biological realism and hardware area efficiency.

- *The leaky integrate and fire (LIF) Model* is widely used by the majority of neuromorphic implementations due to its simplicity and computational efficiency. It treats the neuron as a capacitor (C_m) in parallel with a resistor (R_m). The evolution of the membrane potential v(t) is given by

One of the key lessons learned through the recent trend in benchmarking is the performance-flexibility balance. Although frameworks such as snnTorch offer the highest usability for designing general-purpose SNN, frameworks such as SpikingJelly and Norse provide significant speedups on particular neuron models using customized CUDA kernels along with state-of-the-art compiler optimizations. This hints at an industry move towards compiled solutions that can take advantage of particular instruction sets on neuromorphic hardware.

Historical Development of Neuromorphic Hardware

Neuromorphic hardware represents a significant departure from the architecture introduced by John von Neumann in 1945 [1]. The field was inaugurated in the late 1980s by Carver Mead, who realized that the subthreshold operation of CMOS transistors could be used to implement the exponential and logarithmic functions inherent to the biological neural network. Although the analog design was efficient but suffered from noise and variations in the components used.

In the 2010s, there was a paradigm shift towards digital neuromorphic hardware which provided the precision and flexibility required for machine learning. IBM's TrueNorth, launched in 2014, was a revolutionary product packing 1 million neurons and 256 million

$$\tau \frac{dV}{dt} = = -(V - V_{\{rest\}}) + RI$$

The Izhikevich and Hodgkin-Huxley models: For systems desiring complex spiking patterns, such as bursting or adaptation, the Izhikevich model provides an effective yet computationally inexpensive solution.

$$\frac{dv}{dt} = a(bv - u)$$

Where, *(u)* is the membrane recovery variable. The Hodgkin-Huxley model represents the most detailed biophysical description, including four nonlinear ODEs to describe sodium and potassium channel gating. Although it is too sophisticated for most AI tasks, it is used in highly accurate neuroscience simulations on platforms such as SpiNNaker.

Encoding of Spatio-Temporal Information

For an event-driven system, the conversion of real-value sensory data into discrete spike trains is a critical step. The choice of encoding method impacts the accuracy and the energy efficiency of the network.

- *Rate encoding*: Information is represented by the average number of spikes in a given time window. While keeping the signal clear, it requires long simulation times and high spike counts, potentially negating the energy advantages of SNNs.
- *Temporal encoding (TTFS):* Time-to-First-Spike coding uses the latency of the first spike to represent input intensity. Brighter pixels or stronger signals trigger earlier spikes, enabling ultrafast inference with only one spike per neuron.
- *Latency coding*: Similar to TTFS, latency coding relies on the discovery that tells how sensory events cause upstream neurons to fire earlier, allowing for rank-order processing.
- *Differential coding*: An approach introduced in 2025 that transmits only changes in rate information rather than the rates directly, thereby reducing spike counts in image processing tasks.

Training Methods

Solving the Problem of the Non-Differentiability of Spike Activation Function Derivative. The derivative of the spike function is the Dirac Delta function, which is zero everywhere except on the threshold, where it becomes infinitely large, making traditional backpropagation unsuitable. Three main approaches for dealing with this problem have been devised. Earlier supervised training frameworks, such as SLAYER, addressed this challenge by enabling error backpropagation through SNNs [8].

- *Surrogate gradient learning (SGL):* uses a continuous derivative instead of the Heaviside function discontinuous derivative when going backwards. If we define the overdrive as UOD =v-Vth, a common surrogate is the fast sigmoid:

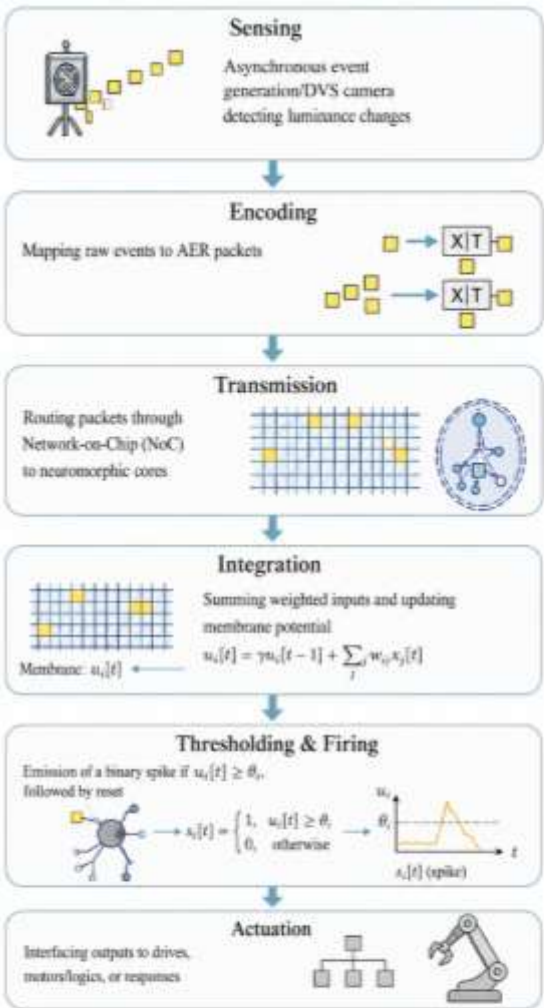

Figure 1. L SNN system architecture and event-driven processing flow.

$\partial^{S}=\partial v$ (1)

$(k\ |U|+1)^2$

OD

where *(S)* determines how close an approximation of the derivative will be, e.g., AdaLi (Adaptive and Lightweight backpropagation) dynamically adjust gradient update boundaries to mitigate vanishing and explosion issues.

- *ANN to SNN conversion*: The method involves the maturity of ANN training by mapping pretrained weights to a spiking architecture. Mathematical similarity between ReLU units and firing rate neurons makes it possible to obtain nearly identical results without fine-tuning thanks to the direct transfer. Despite its initial slow speed, by 2025 the process became fast due to advancements in threshold tuning and differential encoding techniques.
- *Local plasticity (STDP):* Spike-Timing-Dependent Plasticity is an unsupervised learning algorithm in which synaptic weights are adjusted based on the millisecond-scale temporal correlation between pre- and post-synaptic spikes. If a presynaptic spike arrives shortly before the postsynaptic neuron fires, the synapse is potentiated; otherwise, it is depressing. STDP is ideal for on-chip adaptation in edge robotics where centralized retraining is impossible.

The overall event driven processing flow is illustrated in Figure 1.

Results and Current Applications

The practical implementation of SNNs on event-driven hardware has moved beyond synthetic benchmarks, such as MNIST, to high-dimensional real-world tasks. The primary metrics for success in these applications are accuracy, latency, and energy consumption per inference (J/OP).

Performance Comparison of Neuromorphic Hardware

Table 2 presents the technical specifications and energy efficiency of the leading neuromorphic platforms currently available or announced as of 2025.

As shown in Table 3, the transition of SNNs to high-performance tasks is evident in the 2025 results for ImageNet-1K, where models such as Max-Former achieve 82.39% top 1 accuracy while consuming only a fraction of the energy required by standard Vision Transformers. Furthermore, the development of I2E conversion pipelines allows for an unprecedented 92.5% accuracy on real-world DVS-CIFAR10 benchmarks.

Table 2. Comparative specifications of prominent neuromorphic hardware.

Platform	Type	Neuro ns	Power	Key highlight
Intel Lohihi 2	Digital	1M	~1W	Programmable dynamics; high energy efficiency
IBM NorthPol e	Digital	—	Variab le	No off-chip DRAM; high CNN efficiency
SpiNNak er 2	Digital (ARM)	152K	W-scale	Large-scale neural simulation
IBM TrueNort h	Digital	1M	65–70 mW	Ultra-low power; inference-only
BrainChi p Akida 2.0	Digital	1.2M	500 mW	Hybrid SNN/CNN; edge AI optimized
DYNAP-SE2	Mixed signal	1K	<100 mW	Analog-style multi-timescale dynamics

Table 3. SNN Performance benchmarks (2024–2025).

Model	Dataset	Accuracy	Latenc y	Energy savings vs ANN
STAA-SNN	ImageN et-1K	70.40%	4steps	20x lower energy
Ma x-For mer	ImageN et	82.39%	4steps	70% Reduction
ResNet-18	CIFA R-10	94.75%	2steps	High Efficiency

Healthcare and Biosignal Monitoring: In medical electronics, the low thermal footprint of SNNs is their most valuable asset. Neuromorphic microcontrollers like Pulsar and Innatera are being used for always-on bio signal monitoring [9]. These systems can detect epileptic seizures from EEG data or identify cardiac anomalies from ECG signals while consuming less than 35mW, enabling wearable devices with multi-week battery lives. The use of STDP for online learning allows these devices to adapt to the unique physiological signatures of individual patients without uploading sensitive data to the cloud.

Digital NPUs, such as Loihi 2 [3] and NorthPole [5], represent the most scalable path for enterprise AI, whereas analog and mixed-signal systems, such as DYNAP-SE2, achieve extreme efficiencies (as low as 1.2 fJ/OP) for specialized sensing tasks. The integration of 224MB of on-chip memory in NorthPole effectively "obliterates" the von Neumann bottleneck, allowing for high-throughput inference without the latency of external memory calls.

CASE STUDY

Autonomous Robotics and Drone Navigation

One of the the most compelling applications of SNNs is in high-speed autonomous navigation. Research conducted at the University of Zurich utilizing a Prophesee EVK4 event camera and an Intel Loihi 2 chip has produced racing drones capable of navigating obstacle courses at 80 km/h. The event-based sensor detects obstacles in 2 ms, providing the SNN controller with sufficient time to adjust the flight trajectory, resulting in a speed increase of 30% over frame-based systems [9, 10].

DISCUSSIONS

The empirical success of SNNs highlights a profound shift in the technological landscape; however, several systemic challenges must be addressed before widespread commercial adoption can occur.

The Deployment Paradox and Hardware-Software Co-Design

The "Deployment Paradox" describes a scenario in which the theoretical energy gains of sparse SNN algorithms are negated by the architectural constraints of general-purpose hardware, such as GPUs. Because GPUs are optimized for synchronous, dense execution, the irregular firing patterns of SNNs lead to "control-flow divergence" and non-contiguous memory requests, forcing the hardware to run inefficiently. The solution lies in hardware-software co-design, where the algorithm is built with the physical constraints of the chip in mind. Recent strategies, such as Ternary-8-bit Hybrid Weight Quantization (T8HWQ), reconcile the heterogeneous computational patterns of SNNs, preserving accuracy while simplifying the hardware requirements for FPGA-based edge deployment.

Toolchain Maturity and Talent Availability

The ecosystem for traditional AI is mature, featuring robust tools for model versioning, A/B testing, and debugging. In contrast, the SNN toolchain remains fragmented and is primarily academic. Interpretability is a particular concern; while SHAP values and feature importance are well understood for ANNs, explaining a decision based on the temporal dynamics of spike trains is significantly more difficult. Furthermore, finding engineers capable of navigating the intersection of neuroscience, materials science, and computer science is a critical barrier to enterprise integration.

Security and Robustness in the Spiking Domain

Initial research suggests that the discrete and stochastic nature of SNNs may offer inherent protection against adversarial attacks. However, 2025 studies have demonstrated that SNNs are susceptible to Membership Inference Attacks (MIAs), in which an adversary can determine whether a specific sample is part of the training set. The susceptibility to these attacks increases with inference latency, highlighting a new dimension of the accuracy-efficiency trade-off. This necessitates the development of privacy-preserving training techniques, such as spiking differential privacy, to ensure the trustworthiness of edge-deployed systems. The complete SNN processing workflow is shown in Figure 2.

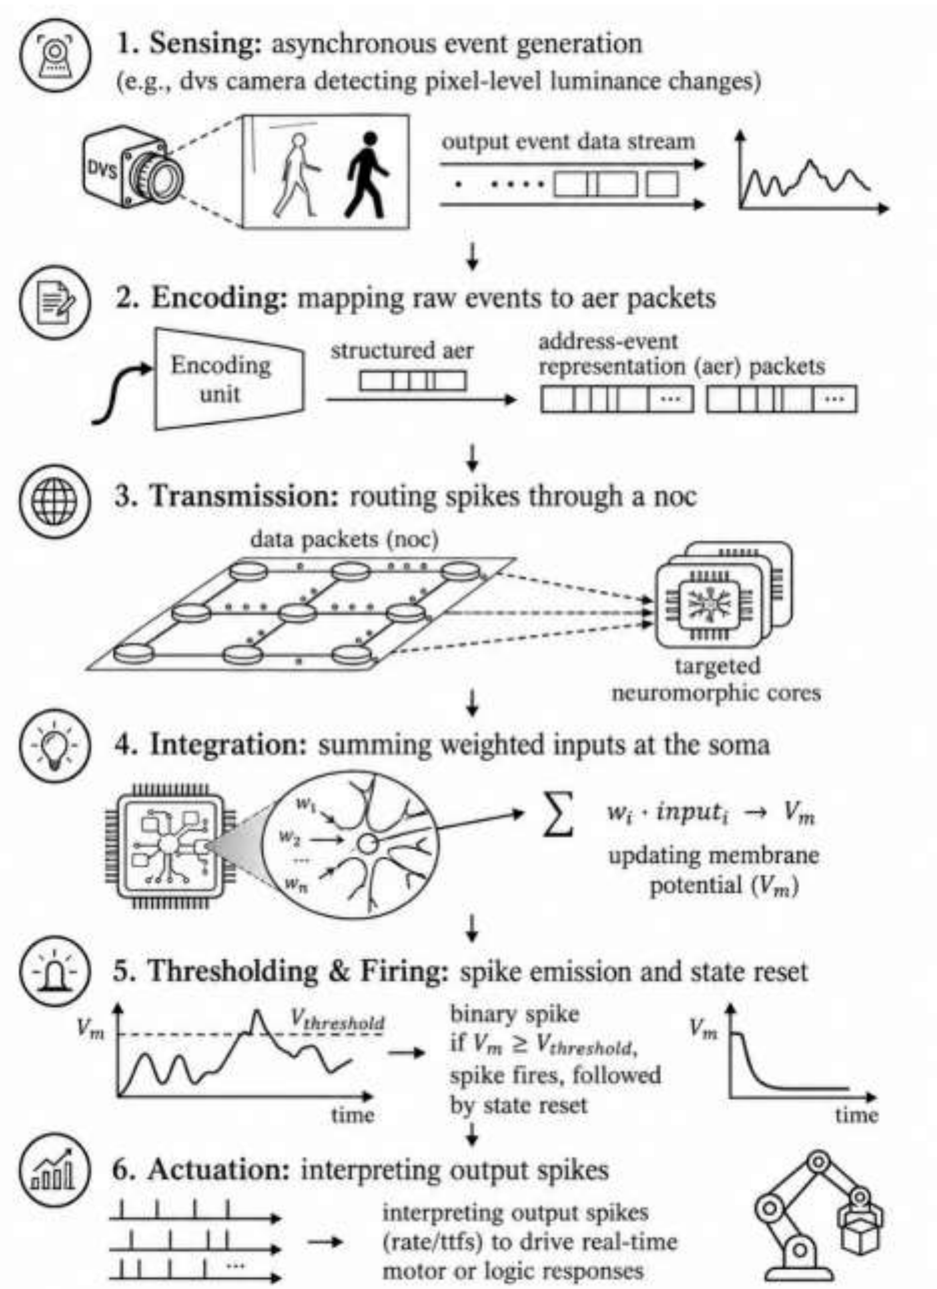

Figure 2. Event-driven anomaly detection system workflow.

CONCLUSION

As an AI architecture based on the fundamental nature of biology, SNNs, together with event-driven computing, offer a viable solution for creating sustainable and autonomous intelligent systems. The adoption of the biological concept of neural spiking allows SNNs to surpass the energy limitations of conventional computing by integrating energy-efficient characteristics into the network framework. The evolution of ideas behind models such as LIF and STDP to the creation of Loihi 2 chips and NorthPole platforms marks a milestone in the evolution of computation paradigms.

The synthesized evidence from 2024 and 2025 research confirms that SNNs have achieved competitive accuracy in complex vision and robotics tasks while providing orders of magnitude of energy efficiency. Simultaneously, the establishment of frameworks for collaborative edge-cloud computing and image-to-event transformation allows the application of new concepts for today's industries to be implemented successfully. As technology progresses through breakthroughs related to the development of 3D stacks and novel memristive devices, neuromorphic computing is expected to approach the complexity of biological processes in terms of performance. Hardware-aware training algorithms and robust privacy frameworks, ensuring that the next generation of intelligent machines is both efficient and trustworthy.

REFERENCES

1. Davies M, Srinivasa N, Lin TH, Chinya G, Cao Y, Choday SH, et al. Loihi: a neuromorphic manycore processor with on-chip learning. IEEE Micro. 2018;38(1):82-99.
2. Merolla PA, Arthur JV, Alvarez-Icaza R, Cassidy AS, Sawada J, Akopyan F, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. Science. 2014;345(6197):668-73.
3. Fang W, Chen Y, Ding J, Yu Z, Masquelier T, Chen D, et al. SpikingJelly: an open-source machine learning infrastructure platform for spike-based intelligence. Sci Adv. 2023;9(40):eadi1480.
4. Math S, Madhuri HD, Yadhav V, Patil M, Selvakumar P. Quantum-inspired or neuromorphic discrete-event computing paradigms for future ubiquitous systems. In: Ubiquitous Computing for Discrete Event Systems 2026. IGI Global Scientific Publishing; 2026. p. 99-138.
5. Modha DS, Akopyan F, Andreopoulos A, Appuswamy R, Arthur JV, Cassidy AS, et al. Neural inference at the frontier of energy, space, and time. Science. 2023;382(6668):329-35.
6. Zenke F, Vogels TP. The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks. Neural Comput. 2021;33(4):899-925.
7. Shrestha SB, Orchard G. Slayer: spike layer error reassignment in time. In: Advances in Neural Information Processing Systems. 2018;31.
8. Indiveri G, Liu SC. Memory and information processing in neuromorphic systems. Proc IEEE. 2015;103(8):1379-97.
9. Cheng Y, Wang M, Hao Z, Buyya R. Brain-inspired AI for edge intelligence: a systematic review [preprint]. arXiv. 2026 Mar 19. arXiv:2603.26722. doi:10.48550/arXiv.2603.26722.
10. Patiño-Saucedo A, Rostro-Gonzalez H, Serrano-Gotarredona T, Linares-Barranco B. Event-driven implementation of deep spiking convolutional neural networks for supervised classification using the SpiNNaker neuromorphic platform. Neural Netw. 2020;121:319-328.

