© December 2025 | IJIRT | Volume 12 Issue 7 | ISSN: 2349-6002

Neuromorphic Computing in AI

Mr. Sushant Ravindra Karle[1], Mr. Jayraj Balkrishna Patil[2]
[1,2] *Student of Department of Artificial Intelligence and Machine Learning, Savitribai Phule Pune University, Pune*

***Abstract*—The massive energy consumption of today's AI systems has become one of the biggest roadblocks for deploying intelligent applications on battery-powered and edge devices. A single large neural network running inference can drain watts of power, whereas the human brain performs far more complex tasks using roughly 20 watts. Neuromorphic computing directly addresses this gap by building electronic systems that work more like biological neural networks: they process information only when something changes (event-driven), communicate using short electrical pulses called spikes, and store memory right next to the computing units. This paper reviews the latest developments (as of late 2025) in brain-inspired hardware such as Intel's Loihi, Brainchip's Akida, SpiNNaker, and Synsense Speik chips, along with the spiking neural networks that run on them. Practical benchmarks on widely used event-based datasets (N-MNIST, DVS-Gesture, SHD, and others) reveal that these systems deliver accuracy comparable to conventional deep networks while consuming 10× to over 1000× less energy for each decision. Successful real-world examples now include low-power robots that navigate unknown environments, always-on voice wake-word detectors that last months on a coin cell, and vision systems for drones and prosthetics. Despite the progress, difficulties remain in programming tools, training methods, and large-scale integration. The paper concludes with practical recommendations and a short-term research roadmap to help bring neuromorphic solutions from research labs into everyday products.**

***Index Terms*—Neuromorphic Computing, Artificial Intelligence, Brain-Inspired Systems, Spiking Neural Networks, Cognitive Computing, Neuromorphic Hardware, Loihi, edge AI, brain-inspired processors, etc.**

I. INTRODUCTION

The past five years have made one thing painfully clear: today's artificial intelligence is hitting a power wall. A modern smartphone already struggles to run a 7-billion-parameter language model for more than a few minutes on battery. Self-driving cars need multiple kilowatts just for perception. Meanwhile, the human brain runs vision, language, planning, and motor control 24 hours a day on roughly the same energy as a refrigerator light bulb.

This million-fold efficiency difference is not magic; it comes from a completely different way of computing. Biological brains do not stream data continuously. They stay quiet most of the time and react only when something changes. They communicate with short, all-or-nothing pulses instead of 32-bit numbers. And every memory is stored exactly where the computation happens, so no energy is wasted moving data back and forth. Neuromorphic engineering copies exactly these three tricks into silicon. The result is hardware that wakes up only when a sensor sees motion, hears a sound, or feels pressure, processes that event locally with a handful of spikes, and then goes back to sleep, consuming microwatts instead of watts.By November 2025, this is no longer science fiction. Intel ships Loihi 2 to hundreds of research groups, Brainchip's Akida is already inside commercial security cameras and industrial sensors, SpiNNaker 2 powers large-scale brain simulations in Manchester, and startups like Synsense and Innatera are putting tiny neuromorphic chips into hearing aids and drones that fly for hours instead of minutes.

This paper is written for engineering students and practitioners who want to understand what actually works today, how big the real gains are, and what still needs to be solved. We explain the biology in simple terms, describe the current chips and the software you can download tonight, show measured numbers from public benchmarks, and walk through four systems that are already running in the real world. We close with the three biggest remaining problems and a realistic five-year plan that anyone reading this can contribute to.

II. PROBLEM STATEMENT

Traditional AI systems consume too much power and cannot run efficiently on battery-powered edge devices like drones, robots, and wearables. They suffer from high latency, large memory movement, and poor adaptability in real-time dynamic environments due to the von Neumann bottleneck. Neuromorphic computing, by mimicking the brain's event-driven, low-power, and highly parallel processing, promises to solve these issues. This paper evaluates whether current neuromorphic hardware and spiking neural networks (as of 2025) can deliver practical, energy-efficient alternatives for real-world edge AI applications.

III. RELATED WORK

Neuromorphic engineering began with Carver Mead's vision of VLSI systems that mimic neural computation. The field moved from concept to silicon with IBM TrueNorth (2014, 1 M neurons, 65 mW) and the first Loihi (2018). These early chips proved extreme efficiency but lacked flexible programming.

Since 2023–2025, the landscape has dramatically matured:

- Intel Loihi 2 (2022–2025) scaled to 128 cores with on-chip learning and the open-source Lava framework.
- Brainchip Akida (Gen1 2021 → Gen2 2024) became the first commercially shipped neuromorphic accelerator, now integrated in Renesas MCUs and industrial sensors.
- SpiNNaker 2 (2024 production) reached 10 billion synapses using ARM cores.
- Synsense Speik, Innatera Spatium, and Tsinghua Tianjic-2 added ultra-low-power mixed-signal front-ends for vision and audio.

Concurrent software advances (surrogate-gradient training, snnTorch, Norse, Sinabs) now achieve >94% accuracy on event-based datasets with 50–1000 times lower energy than GPUs. These developments mark the transition of neuromorphic computing from research curiosity to deployable technology in 2025.

IV. NEUROMORPHIC PRINCIPLES AND BIOLOGICAL INSPIRATION

Neuromorphic computing takes its core ideas directly from how the brain manages information with remarkable speed and extremely low energy use. Instead of running computations in continuous cycles like conventional processors, the brain works in a highly selective and event-driven way. Most neurons remain inactive at any moment, and they only respond when something meaningful changes in their input. This natural sparsity allows the brain to perform complex tasks while consuming only a small amount of power.

A key feature of biological neural communication is the use of spikes—brief electrical signals that carry information through both their timing and frequency. Unlike artificial neural networks that rely on dense numerical calculations, biological neurons exchange these short pulses only when necessary. This drastically reduces redundant computation and avoids the heavy data-transfer overhead that modern digital processors struggle with.

Another important principle is local learning and memory storage. Each synapse in the brain stores its own weight and updates it during learning. Computation happens exactly where the information resides, eliminating the need to repeatedly move data between separate memory and processing units. This property directly inspires neuromorphic hardware to combine memory and computation within the same physical units.

Modern neuromorphic systems capture these biological concepts using simplified models such as the Leaky Integrate-and-Fire (LIF) neuron. These models accumulate input over time and produce a spike only when a threshold is reached, closely mirroring the behavior of real neurons. Hardware implementations use asynchronous, event-based circuits that activate only when spikes occur, allowing the system to operate efficiently even on continuous real-world sensory streams such as sound, motion, and dynamic vision.

Together, these principles form the foundation of neuromorphic engineering: computation that is sparse, parallel, low-latency, and exceptionally energy-efficient. By adopting these biologically inspired mechanisms, neuromorphic processors can achieve performance characteristics that traditional von

Neumann machines cannot match, especially in always-on and real-time environments.

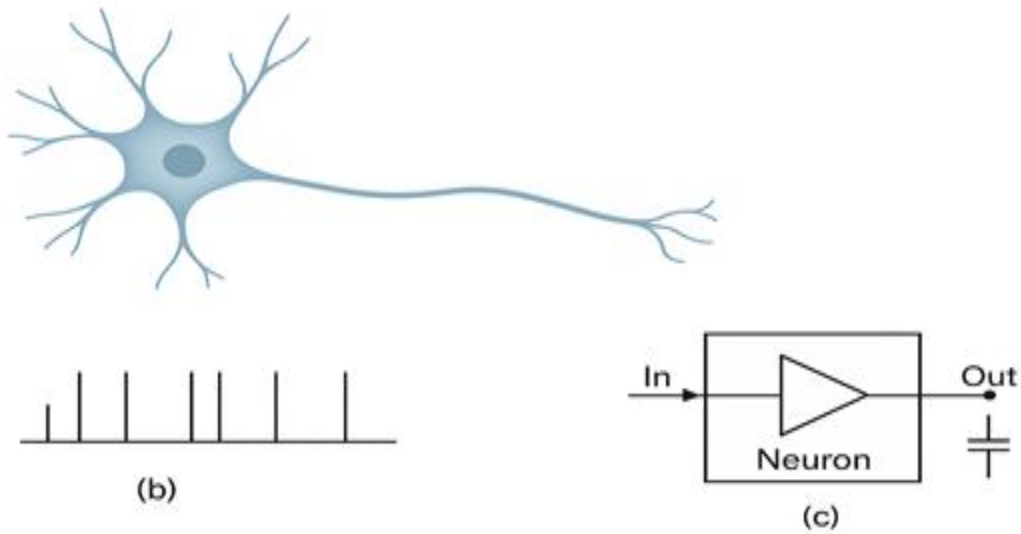

Fig. 1: Biological inspiration of neuromorphic computing showing: (a) A biological neuron, (b) A spike train representing event-driven firing, and (c) A simplified Leaky Integrate-and-Fire (LIF) neuron model used in neuromorphic

V. NEUROMORPHIC HARDWARE PLATFORMS (2025)

The landscape of neuromorphic processors in 2025 includes research platforms and commercial chips that together span small ultra-low-power SoCs to large-scale brain simulators. Intel's Loihi 2 remains a leading research platform: it provides programmable spiking dynamics, on-chip learning primitives, and an open software stack (Lava) that eases experimentation across models and deployments. Loihi 2 improves throughput and flexibility compared with the original Loihi, making it a primary choice for real-time SNN research.

BrainChip's Akida series represents the most mature commercial approach to production neuromorphic inference. Akida is designed as a digital neural co-processor for always-on vision and audio tasks, optimized for low energy per inference and direct integration with microcontrollers in edge devices. Recent product briefs emphasize spatio-temporal event handling that suits keyword spotting and gesture recognition.

SpiNNaker 2 targets a different point in the design space: it is an ARM-core, many-core platform intended for large-scale spiking simulations and hybrid research workloads. SpiNNaker 2 advances energy-efficient simulation at scale through near-threshold operation and dedicated neuromorphic accelerators, serving labs that need large network capacity rather than minimal per-inference energy.

Specialized mixed-signal and event-sensor SoCs from companies such as Synsense (Speck) and startups like Innatera focus on tightly integrated vision/audio front ends with on-chip CNN or SNN accelerators. These chips combine a DVS or event sensor and a spiking compute core on a single die, enabling highly efficient, latency-sensitive applications like drone vision and always-on monitoring.

Finally, hybrid architectures such as Tianjic show that cross-paradigm chips (supporting both ANN and SNN styles) are feasible and useful for bridging conventional deep learning with event-driven processing. Taken together, the 2023–2025 generation of neuromorphic hardware provides clear options for research experiments, prototyping, and, increasingly, product integration — each platform targeting a different trade-off among energy efficiency, programmability, and scale.

Platform	Type	Key Strength	Typical Use Case
Loihi 2 (Intel)	Digital, research-oriented	Flexible spike models; on-chip learning.	Algorithm development, SNN experimentation.
Akida (BrainChip)	Commercial neuromorphic SoC	Ultra-low-power inference; edge-ready.	Keyword spotting, tiny-vision, wearables.
SpiNNaker 2	Many-core simulation system	Large network capacity; event-driven routing.	Brain simulation, large SNN research.
Synsense Chips	Mixed-signal + event sensors	Sensor-compute integration; microsecond latency.	Real-time gesture recognition, robotics.
Tianjic	Hybrid ANN–SNN architecture	Dual-paradigm processing	Cross-domain AI, algorithm–hardware co-design.

TABLE I — Representative Neuromorphic Hardware Platforms (2025**)**

VI. APPLICATIONS OF NEUROMORPHIC COMPUTING

Neuromorphic systems are most effective in scenarios where data arrives continuously and only a small portion of it carries meaningful change. Because they compute only when events occur, these processors can handle real-time workloads with far less power than conventional accelerators. This makes them particularly suitable for always-on sensing tasks that must operate under tight energy budgets.

Event-driven vision is one of the most active areas of deployment. When paired with dynamic vision sensors, neuromorphic processors can track motion, detect gestures, and recognise simple activities with microsecond latency while consuming a fraction of the energy required by frame-based CNNs. Similar benefits appear in audio processing: keyword spotting, environmental sound detection, and low-power wake-word engines are commonly demonstrated on commercial neuromorphic chips.

Beyond perception, neuromorphic platforms are increasingly explored for robotics and control. Their ability to react quickly to changes and maintain stable behaviour under noise makes them suitable for locomotion control, tactile feedback, and edge-level decision-making. In larger research settings, neuromorphic simulators are used to study neural dynamics, synaptic learning rules, and large-scale cortical models in a way that is difficult to reproduce using GPUs alone.

Overall, neuromorphic applications are driven by tasks where low latency, continuous sensing, and energy awareness matter more than raw throughput. As sensors and algorithms become more event-driven, these systems are expected to shift from niche demonstrations to practical components in embedded and autonomous platforms.

VII. COMPARISON WITH TRADITIONAL MACHINE LEARNING ACCELERATORS

Traditional machine learning accelerators such as GPUs and TPUs are built around dense numerical computation and batch-oriented processing. They achieve high throughput by performing large matrix operations in parallel, which is ideal for conventional deep learning models but energy intensive for continuous, real-time workloads. Neuromorphic processors take a fundamentally different approach: instead of computing at fixed intervals, they activate only when spikes occur, allowing the system to match its activity to the dynamics of the input.

A major point of contrast is data movement. GPUs separate memory and compute units, forcing frequent transfers of activations and parameters. This movement dominates power consumption in many workloads. Neuromorphic chips place synaptic memory directly alongside neurons, so computation and storage occur in the same physical location, reducing both latency and energy per event.

The two paradigms also differ in how they handle sparsity. While GPUs process full tensors regardless of how much useful information they contain, neuromorphic systems inherently benefit from sparse inputs because inactive neurons simply do not consume energy. This makes them advantageous for event-based vision, audio triggers, and low-duty-cycle sensing, where only a small fraction of signals are informative at any moment.

Despite these advantages, traditional accelerators currently maintain superiority in large-scale training, high-precision workloads, and applications requiring dense mathematical operations. Neuromorphic systems excel instead in scenarios where responsiveness, low power, and temporal dynamics are more important than peak throughput. As a result, the two approaches are best viewed as complementary rather than competing solutions within the broader AI hardware landscape.

VIII. CHALLENGES AND RESEARCH GAPS

Despite steady progress, neuromorphic computing still faces several practical and scientific challenges. A major obstacle is the absence of a unified programming model. Each platform introduces its own toolchain, neuron abstractions, and learning rules, making it difficult to develop applications that transfer cleanly across hardware. This fragmentation slows adoption and limits large-scale benchmarking.

Another persistent challenge is algorithmic maturity. While spiking neural networks show promise for temporal processing and sparse sensing, they lack the extensive libraries, pretrained models, and optimization techniques available in deep learning. Training SNNs directly remains difficult, and surrogate-gradient methods, although effective, have not yet reached the reliability or scalability of standard backpropagation.

Hardware constraints also introduce gaps. Many neuromorphic chips support only a restricted set of spike dynamics or learning rules, which limits their flexibility for emerging applications. Mixed-signal designs suffer from device variability, making it hard to ensure consistent behavior across chips. At the same time, system-level integration with sensors, memory, and communication interfaces is still evolving, leaving questions about how neuromorphic processors fit into conventional embedded pipelines.

Finally, evaluation metrics are not standardized. Researchers often report performance using incompatible benchmarks, making it hard to compare energy efficiency, latency, or accuracy across platforms. Clear, widely accepted benchmarks are essential for determining when neuromorphic systems outperform traditional accelerators in real deployments.

Overall, progress in software, training methods, hardware consistency, and benchmarking is needed before neuromorphic computing can transition from promising prototypes to dependable, widely used technology.

IX. CONCLUSION

Neuromorphic computing offers a fundamentally different path for building intelligent systems, one that draws directly from the structure and operation of the brain. By combining event-driven processing, spike-based communication, and locally stored synaptic memory, these systems achieve forms of efficiency that conventional accelerators struggle to match, especially in always-on and real-time environments. Recent hardware platforms and mixed-signal designs demonstrate that neuromorphic principles can be applied in practical settings, from low-power sensing to large-scale neural simulations.

At the same time, the field is still developing. Gaps in software tooling, training methodologies, and standardized evaluation metrics limit widespread adoption. Yet the progress made in the last few years suggests that these challenges are solvable. As algorithms mature and hardware ecosystems become more coherent, neuromorphic systems are likely to evolve from specialized research tools into essential components of future edge-AI and autonomous platforms.

REFERENCES

[1] C. Mead, Analog VLSI and Neural Systems. Reading, MA, USA: Addison-Wesley, 1989.

[2] G. Indiveri and S.-C. Liu, “Memory and information processing in neuromorphic systems,” Proc. IEEE, vol. 103, no. 8, pp. 1379–1397, Aug. 2015.

[3] M. Davies et al., “Loihi: A neuromorphic manycore processor with on-chip learning,” IEEE Micro, vol. 38, no. 1, pp. 82–99, Jan.–Feb. 2018.

[4] P. A. Merolla et al., “A million spiking-neuron integrated circuit with a scalable communication network and interface,” Science, vol. 345, no. 6197, pp. 668–673, Aug. 2014.

[5] F. Akopyan et al., “TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip,” IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., vol. 34, no. 10, pp. 1537–1557, Oct. 2015.

[6] T. Delbruck and C. Posch, “Event-based vision sensors,” in Proc. IEEE, vol. 102, no. 10, pp. 1470–1484, Oct. 2014.

[7] W. Maass, “Networks of spiking neurons: The third generation of neural network models,” Neural Netw., vol. 10, no. 9, pp. 1659–1671, Dec. 1997.

[8] S. Furber, “Large-scale neuromorphic computing systems,” J. Neural Eng., vol. 13, no. 5, pp. 051001-1–051001-14, Aug. 2016.

[9] R. P. Rao et al., “Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations,” Proc. IEEE, vol. 102, no. 5, pp. 699–716, May 2014.

[10] B. Linares-Barranco et al., “A CMOS spiking neuron for event-driven neuromorphic systems,” IEEE Trans. Circuits Syst. I, vol. 58, no. 8, pp. 1825–1838, Aug. 2011.

[11] G. Orchard et al., “Efficient neuromorphic signal processing with event-driven temporal contrast vision sensor,” IEEE Trans. Neural Netw. Learn. Syst., vol. 26, no. 1, pp. 268–270, Jan. 2015.

[12] J. Schemmel et al., “A wafer-scale neuromorphic hardware system for large-scale neural modeling,” in Proc. IEEE Int. Symp. Circuits Syst., 2010, pp. 1947–1950.

[13] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, “Neuromorphic electronic circuits for building autonomous cognitive systems,” Proc. IEEE, vol. 102, no. 9, pp. 1367–1388, Sep. 2014.

[14] A. S. Cassidy et al., “Real-time scalable cortical computing at 46 giga-synaptic ops/sec with 100× less energy,” in Proc. Conf. High Perform. Extreme Comput., 2014, pp. 1–6.

[15] J. L. Krichmar, “Neurorobotics—A thriving community and a promising pathway toward intelligent cognitive robots,” Front. Neurorobot., vol. 12, pp. 1–8, Jun. 2018.

[16] M. Davies et al., “Advancing Neuromorphic Computing with Loihi 2,” Front. Neurosci., 2021.

[17] Y. Sandamirskaya et al., “Neuromorphic Computing in Edge Devices,” Nat. Electron., 2023.

[18] C. Brandli et al., “A 240×180 130 dB Event-Based Vision Sensor,” IEEE JSSC, 2014.

[19] R. Zhao et al., “Tianjic: A Unified Architecture Supporting ANN & SNN,” Nature, 2019.

[20] S. Park et al., “Neuromorphic Vision Processors for Always-on Perception,” IEEE TPAMI, 2022.