International Journal of
Technology and Emerging Research
ISSN: 3068-109X (online) | ISSN: 3068-1995 (print) | https://www.ioro.org/ijter

International Conference on Intelligent Technologies for Tomorrow (ICITT-2026) · 18 August 2026

Neuromorphic Computing: Current Progress and the Future of Brain-Inspired Computing

Jisna C Jeejo [1], **Habeeba M A** [2]

[1] *MSC Computer Science, Ansar Women's College, Thrissur, India*

[2] *MSC Computer Science, Ansar Women's College, Thrissur, India*

* **Corresponding author:** jisnacjeejo@gmail.com

Abstract

Neuromorphic computing borrows its design logic from the nervous system rather than from the von Neumann architecture that has dominated computing for seventy years. Instead of shuttling data back and forth between separate memory and processing units, it favors event-driven communication, computation that happens close to (or inside) memory, and massive parallelism across simple processing elements. This paper takes stock of where the field currently stands: spiking neural networks, digital and analog processors, memristive and other emerging devices, the software ecosystems that support them, and application areas ranging from robotics and edge intelligence to biomedical monitoring and event-based vision. What emerges from the recent literature is a field that has largely moved past small proof-of-concept chips and is now building larger, more programmable platforms with tighter hardware-algorithm integration. Even so, real obstacles remain around training methods, benchmarking practices, programmability, device variability, and fabrication, and it is still unclear how much of the field's energy advantage survives contact with general-purpose workloads. The view taken here is that brain-inspired computing is heading toward a hybrid future: conventional digital processors will keep doing what they do best, while event-driven and in-memory accelerators take over the workloads where they have a genuine edge. Neuromorphic computing is therefore unlikely to displace mainstream AI hardware outright, but it looks well positioned to become a key ingredient in low-power, adaptive, real-time intelligence at the edge.

Keywords: *neuromorphic computing; brain-inspired computing; spiking neural networks; memristors; edge AI; in-memory computing*

1. Introduction

Artificial intelligence keeps getting better, but that progress has quietly exposed a weakness in the hardware it runs on. Training and running today's neural networks means moving enormous amounts of data between memory and processing units, and that movement often costs more energy than the arithmetic it supports. The problem is sharpest at the edge, where devices have to work within tight battery, thermal, and latency budgets. Neuromorphic computing responds to this by looking at how nervous systems handle information: sparsely, as discrete events, with computation happening right where memory lives.

None of this means the brain is a blueprint to be copied wholesale into silicon; nobody is trying to fabricate a cortex. What the brain offers instead is a set of design principles worth stealing: massive parallelism, local communication, event-driven activity, ongoing adaptation, and a memory-computation relationship that isn't split down the middle. Neuromorphic engineering is the work of translating those principles into circuits, devices, algorithms, and full systems. Increasingly, researchers treat it not as a single technology but as a broad design space spanning spiking processors, in-memory computing, analog neural circuits, emerging memories, and hybrid architectures [1]–[4].

This paper looks at where brain-inspired computing stands today and where it appears to be heading. It works through the basic concepts, the parallel development of hardware and algorithms, recent research directions, practical applications, the challenges that remain unresolved, and the priorities likely to shape the next few years of work. The argument threaded through all of this is that the field has reached a meaningful turning point: it is moving from isolated demonstrations toward platforms built with real applications in mind, and its success from here will hinge less on hardware breakthroughs alone and more on co-designing hardware, algorithms, and software together.

2. Literature Review

The literature on neuromorphic computing has grown along several threads that only partly overlap, and it helps to pull them apart before looking at where the field is headed. One thread concerns foundational hardware and the memory-processing relationship. Indiveri and Liu [1] laid out early principles for how memory and information processing could be co-located in neuromorphic circuits, and their framing still shows up in how researchers justify event-driven, in-memory designs today. Davies et al. [5] took that thinking into large-scale silicon with Loihi, a manycore processor supporting on-chip learning, and it remains a reference point that later digital platforms are measured against.

A second thread takes a systems-level view of the field rather than focusing on individual chips. Kudithipudi et al. [6] make the case, in a widely cited Nature piece, that neuromorphic computing needs to be evaluated at scale — as complete systems with real workloads — rather than through isolated chip-level benchmarks that flatter the hardware. Related surveys on general-purpose brain-inspired computing [4] and on brain-inspired computing as a systematic research area [9] echo this concern, arguing that the field's credibility depends on standardized, reproducible evaluation rather than favorable task selection. A 2025 editorial in Frontiers in Neuroscience [11] and a topical review in Wiley Interdisciplinary Reviews [12] add further weight to this consensus, both emphasizing that neuromorphic research is maturing from a neuroscience-adjacent curiosity into an engineering discipline with its own standards of evidence.

A third thread works at the device and materials level. Reviews of memristive and other emerging synaptic devices [7] and of neuromorphic devices more broadly [8] converge on a similar diagnosis: the physics of memristors, phase-change memory, resistive RAM, and related technologies is promising, but endurance, retention, and fabrication variability are still far from solved. These sources are notably more cautious than the systems-level papers about near-term deployment, which is a useful counterweight when reading the more optimistic hardware-scaling literature.

A fourth, smaller thread pushes the field toward its more speculative edges. Work on modeling macroscopic brain dynamics with brain-inspired computing [3] and on organoid intelligence as a computing substrate [2] treats neuromorphic principles as tools for neuroscience itself, not just as an engineering target. This literature is scientifically interesting but sits further from deployment, and it raises questions — about reproducibility, ethics, and what "intelligence" even means in this context — that the hardware-focused literature mostly avoids. Finally, application-specific work such as the integration of brain-computer interfaces with neuromorphic processing for human digital twins [10] shows the field reaching outward into biomedical and human-interface domains, where low latency and event-driven processing offer a concrete advantage over conventional pipelines.

Read together, these strands suggest a field in transition: strong agreement that event-driven, co-located memory-and-compute architectures are worth pursuing, but real disagreement about how close any of it is to displacing conventional hardware, and a persistent gap between device-level optimism and systems-level caution. That gap is a large part of the motivation for this paper — it is easier to find a strong result for a single chip or a single benchmark than to find a convincing account of end-to-end system performance, and the sections that follow try to hold both views in view at once.

3. Fundamental Principles

From biological neurons to artificial systems

A biological neuron collects signals at its synapses, integrates them over time, and fires an electrical spike once its membrane potential crosses a threshold. Artificial spiking neurons approximate this with mathematical models, the leaky integrate-and-fire neuron being the most common. A synapse, meanwhile, stores a connection weight and can adjust that weight in response to activity. Simple as these mechanisms are, together they support temporal computation, adaptation, and communication that stays sparse rather than constant.

Conventional neural networks represent values as continuous activations pushed through dense matrix multiplication. Spiking neural networks represent information differently — through the timing, rate, or pattern of discrete spikes — which can cut down computation substantially when the input events themselves are sparse. It also makes time a first-class dimension of the computation itself, which turns out to matter a great deal for speech, gesture recognition, event-based vision, and other sensor-driven tasks where timing carries information.

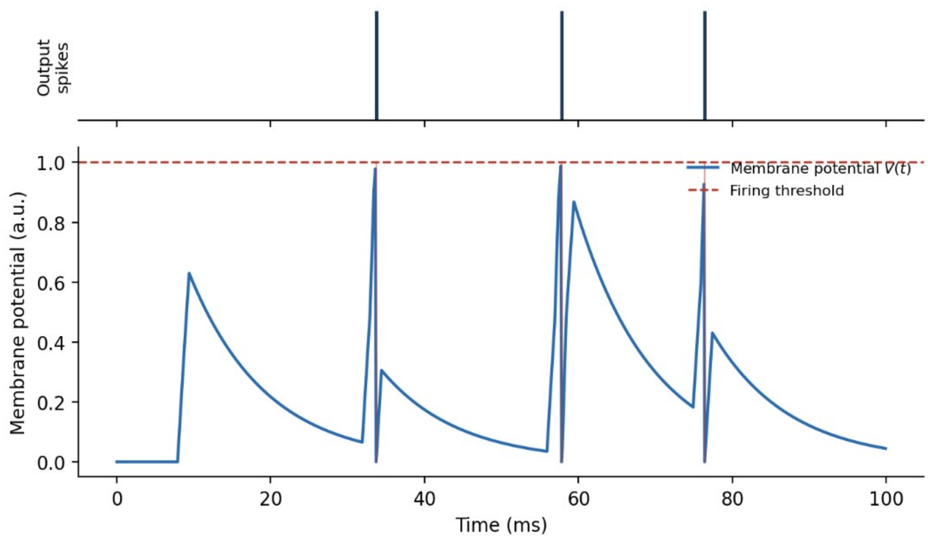

Figure 1. Simulated dynamics of a leaky integrate-and-fire neuron. The membrane potential (bottom) rises with incoming input current and decays exponentially between events; each time it crosses the firing threshold, the neuron emits a spike (top) and resets. This threshold-and-reset behavior is the core computational primitive behind spiking neural networks.

4. Current Hardware Progress

Neuromorphic hardware has grown along several complementary paths rather than converging on one dominant design. Digital architectures offer programmability and dependable operation; analog and mixed-signal circuits can model neuronal dynamics directly, often at lower energy cost; and emerging memory devices try to fuse storage and computation so that weights no longer need to travel between separate memory and processing units. If there's a single headline here, it's not any one chip — it's how much the space of viable engineering options has widened.

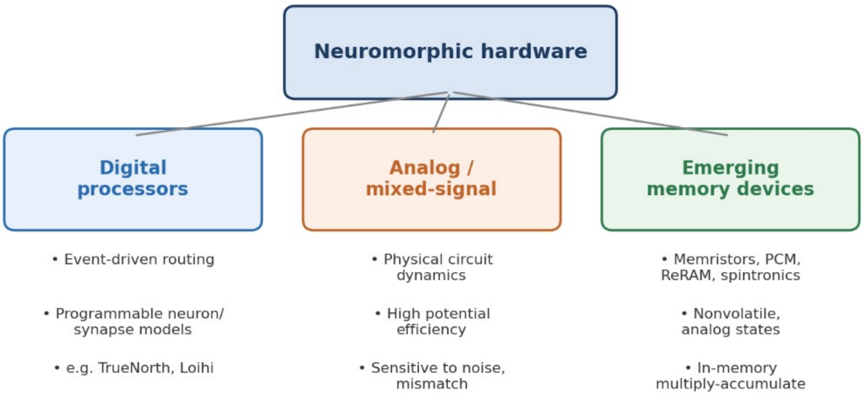

Figure 2. A simplified taxonomy of the neuromorphic hardware landscape, organized around three complementary design paths: digital processors, analog/mixed-signal circuits, and emerging memory devices. Each path trades off programmability, energy efficiency, and reliability differently.

Digital neuromorphic processors

Digital neuromorphic processors implement networks of artificial neurons and synapses on conventional semiconductor technology, generally relying on asynchronous communication and event routing so that only active neurons consume meaningful resources. Research platforms such as IBM TrueNorth and Intel Loihi showed that large populations of spiking neurons could be integrated into genuinely programmable systems. More recent work has shifted toward scaling these systems up, improving multi-chip communication, adding learning support, building better developer tools, and connecting them to conventional AI workflows [5][6].

The strongest argument for digital neuromorphic systems is control: designers can reproduce specific neuron models, reshape network topology, and implement learning rules in software or firmware without redesigning silicon. The trade-offs show up in memory capacity, communication overhead, and the simple fact that on dense workloads, conventional GPUs and CPUs are already extremely well optimized — leaving less room for a neuromorphic advantage than one might expect.

Analog and mixed-signal circuits

Analog circuits represent membrane potentials, currents, and synaptic dynamics directly as physical voltages or currents. Because the computation is carried out by the physics of the circuit itself rather than by a sequence of digital operations, analog implementations can reach high energy efficiency and reproduce time dynamics naturally. Mixed-signal architectures split the difference, pairing analog neuron and synapse blocks with digital event routing, configuration, and learning control.

The catch is that analog systems are sensitive to process variation, temperature, noise, and component mismatch, which makes calibration and error tolerance essential rather than optional. Recent research increasingly treats this variability as a design constraint to be engineered around — through robust learning and adaptive circuits — rather than as a defect to be eliminated, which is a fairly significant shift in framing from earlier work.

Emerging memories and devices

Memristors, phase-change memory, resistive RAM, ferroelectric devices, spintronic elements, and two-dimensional materials are all being explored as artificial synapses. What makes them attractive is non-volatile storage combined with analog conductance states, plus the possibility of performing multiplication or accumulation right where the data is stored rather than shipping it elsewhere first. Recent reviews place these devices at the center of the broader movement toward neuromorphic and in-memory computing [7] [8].

That said, progress at the device level remains largely experimental. Endurance, retention, write variability, nonlinear updates, fabrication yield, and CMOS integration are all still open concerns. If reliability and manufacturing issues can be solved, though, these technologies could enable dense synaptic arrays and local learning at a scale that's difficult to reach with digital or purely analog approaches alone.

5. Algorithms and Software

Hardware progress only matters if algorithms can actually make use of it. Spiking neural networks are the dominant software paradigm in this space, but they come with training difficulties that conventional deep learning doesn't have to deal with. The threshold operation that produces a spike is non-differentiable, so ordinary backpropagation can't be applied directly without some kind of approximation. Researchers have worked around this with surrogate gradients, local plasticity rules, reinforcement learning, and conversion methods that turn already-trained conventional networks into spiking ones.

Learning methods

Surrogate-gradient training swaps in a smooth approximation for the zero or undefined derivative of a spike, which allows gradient-based optimization during learning while still preserving discrete spikes at inference time. Local learning methods, spike-timing-dependent plasticity chief among them, are more biologically plausible and tend to be friendlier to hardware, but they're harder to scale and often fall short of fully supervised methods on accuracy.

A middle path is ANN-to-SNN conversion: train a conventional neural network first, then transform it into a spiking network. This can produce strong results, but it often needs many time steps to do so, and the energy advantage can evaporate if the converted network ends up generating too many spikes to reach the original accuracy. The field's center of gravity is shifting toward directly trained, hardware-aware SNNs and hybrid models that use spiking layers only where they earn their keep.

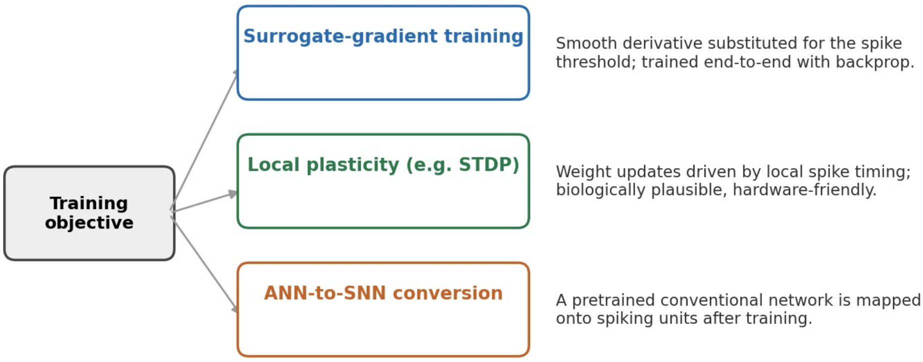

Figure 3. Three main routes to a trained spiking neural network. Surrogate-gradient training and local plasticity both train spiking dynamics directly, while ANN-to-SNN conversion repurposes a conventional pretrained network. Each route makes a different trade-off between training difficulty, accuracy, and hardware efficiency.

Toolchains and benchmarks

None of this works without reliable software tools and benchmarks that measure what actually matters. Frameworks such as NEST, Brian, Lava, and other research environments support simulation or deployment, but portability across different chips is still limited — a model that performs well in simulation can behave quite differently on real hardware once quantization, limited precision, routing constraints, and device noise enter the picture.

Accuracy alone is not enough for a fair benchmark. Useful evaluation needs to report energy per inference, latency, memory footprint, spike count, training cost, robustness, and the overhead introduced by sensors and data conversion. Recent work on neuromorphic computing at scale is blunt about this: without system-level measurement, it's hard to know whether a claimed efficiency advantage is real or an artifact of how the comparison was set up [6].

6. Recent Research Directions

A few directions stand out in the recent literature. Researchers are increasingly building general-purpose brain-inspired systems rather than narrow, task-specific circuits. Hybrid architectures that combine conventional deep learning with spiking or event-driven modules are becoming more common, as teams look for ways to get neuromorphic efficiency without giving up the accuracy of mainstream deep learning. There's also a growing effort to connect neuromorphic processors to event-based sensors, robots, biomedical interfaces, and brain-modeling pipelines. And running underneath all of this is continued materials research into how emerging devices might implement memory, learning, and nonlinear dynamics directly in hardware, rather than approximating them in software [2][3][9].

7. Applications

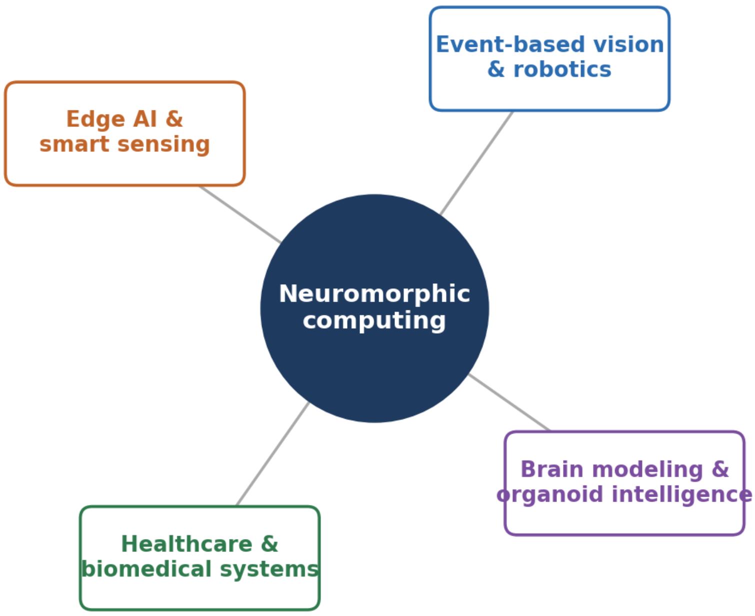

Figure 4. Four application domains where neuromorphic computing's sparse, event-driven processing offers a genuine advantage over conventional architectures: event-based vision and robotics, edge AI and smart sensing, healthcare and biomedical systems, and brain modeling.

Event-based vision and robotics

Event cameras report brightness changes rather than capturing full frames at fixed intervals, which produces sparse, low-latency data that maps naturally onto spiking processors. Neuromorphic vision systems built on this data can support object tracking, optical flow estimation, gesture recognition, and collision avoidance. In robotics specifically, the combination of fast response times and low energy use is especially valuable for drones, autonomous vehicles, and other mobile platforms that can't afford either slow reactions or heavy power draw.

Edge artificial intelligence

Plenty of edge devices need to run continuously on limited battery and with unreliable connectivity. Neuromorphic systems can perform local inference on audio, vibration, temperature, or visual signals without having to ship all the raw data to the cloud. That opens the door to applications like keyword detection, predictive maintenance, smart sensors, security monitoring, and agricultural sensing — all cases where sending everything to a server simply isn't practical.

Healthcare and biomedical systems

Biomedical signals tend to be temporal and often sparse, which makes them a natural fit for event-driven processing. Neuromorphic hardware is being explored for electroencephalography, electrocardiography, prosthetic control, and wearable monitoring. The integration of brain-computer interfaces with neuromorphic processing is an emerging area in its own right, particularly for systems that need to interpret neural signals with very low latency [10]. Before any of this reaches broad deployment, though, safety, privacy, clinical validation, and explainability will all need to be addressed — these are not incidental concerns in a medical context

Brain modeling and organoid intelligence

Brain-inspired computing feeds back into neuroscience as well as drawing from it. Neuromorphic models can simulate neural dynamics and test hypotheses about how information is actually processed in biological systems. More recent research has pushed toward organoid intelligence — using biological substrates themselves as a computing frontier inspired by the brain [2]. These directions are scientifically exciting, but they also raise real questions about reproducibility, ethics, measurement, and even what it means to call something "intelligent" in the first place.

8. Evaluation of Current Progress

Table 1 summarizes where the field has made real headway and where the open problems remain, organized by the four dimensions used throughout this paper.

Dimension	Progress	Remaining problem	Near-term priority
Hardware	More programmable chips and emerging memories	Variability and integration	Reliable heterogeneous platforms
Algorithms	Better SNN training and hybrid models	Limited portability	Hardware-aware learning
Applications	Strong potential at the edge and in robotics	Few broad deployments	End-to-end demonstrations
Evaluation	More energy and latency reporting	Inconsistent benchmarks	Standardized metrics

Table 1. Snapshot of progress, open problems, and near-term priorities across four dimensions of neuromorphic computing.

9. Challenges and Limitations

The first challenge is programmability. Conventional AI has the benefit of mature frameworks, optimized libraries, and programming models that most developers already understand. Neuromorphic systems, by contrast, work with asynchronous events, temporal states, hardware-specific neuron models, and memory organizations that don't map onto standard abstractions. That means a steeper learning curve, and often a need to redesign algorithms from the ground up rather than simply port existing code over.

The second is measurement. An accelerator can look extremely efficient if only the core chip is measured while sensor interfaces, host processors, memory transfers, cooling, and training costs are left out of the accounting. A fair comparison has to define the whole workload and hold neuromorphic and conventional systems to the same boundaries. Without that discipline, claims of brain-like efficiency are hard to reproduce and, frankly, hard to trust.

The third is reliability. Emerging devices can drift, switch stochastically, wear out faster than expected, and vary meaningfully from one device to the next. Biological systems tolerate this kind of noise through redundancy and adaptation, but engineered systems need explicit mechanisms — calibration, error correction, robust learning — to do the same job, and those mechanisms can eat into the idealized benefits that analog or in-memory approaches promise on paper.

A fourth limitation is simply that the brain itself is not fully understood. Brain-inspired engineering has to decide which biological principles are worth reproducing and which can be safely ignored. A system can be useful without being biologically accurate, and a biologically realistic system can be too complex or too hard to manufacture to be useful at all. The most sensible way to judge these ideas is by how well they perform in practice, not by how closely they resemble biology.

10. Future of Brain-Inspired Computing

Heterogeneous and hybrid systems

The most realistic future for this field is a heterogeneous one. CPUs and GPUs aren't going anywhere — they will keep handling dense computation and large-scale training — while neuromorphic modules take over the sparse, temporal workloads they're actually good at. A single system might use a conventional processor for model initialization, a spiking accelerator for event-based inference, and an analog memory array for selected matrix operations. Splitting the work this way is likely to deliver better overall efficiency than insisting on one architecture for every task, however elegant that might sound.

Continual and adaptive intelligence

Brains learn continuously as their environment changes, and future neuromorphic systems may be able to do something similar — supporting continual learning, online adaptation, and local personalization without constantly phoning home to a remote server. That would be a real benefit for robots and devices operating in conditions the original training data never anticipated. The main risks are catastrophic forgetting, unstable learning, and the practical difficulty of verifying safety after a system has adapted itself in the field.

Open standards and reproducibility

Getting the field to a mature state will require shared datasets, open hardware interfaces, common energy-measurement procedures, and programming abstractions that actually port across chips. Reproducible benchmarks are what separate genuine architectural advantages from improvements that only look good because of favorable task selection or incomplete accounting. This systems-level perspective is becoming more central to how general-purpose brain-inspired computing research frames itself [4][9].

11. Conclusion

Neuromorphic computing has made real progress across hardware, algorithms, materials, and applications. Digital spiking processors have shown that programmable, event-driven computation works at meaningful scale; analog and mixed-signal circuits offer efficient neuronal dynamics; and emerging memories point toward a possible route to dense in-memory learning. On the software side, surrogate-gradient training, conversion methods, local plasticity, and hybrid architectures are steadily widening the range of workloads these systems can handle.

At the same time, the evidence doesn't support treating neuromorphic computing as a wholesale replacement for conventional computing — and it probably shouldn't be read that way. The clearest opportunities are workloads with sparse, temporal, real-time data and tight energy budgets. The biggest obstacles are toolchain maturity, reliable fabrication, fair evaluation, limited generality, and the practical difficulty of integrating neuromorphic hardware into complete systems rather than isolated demos.

Where this leaves the field is at a point best defined by co-design rather than by any single breakthrough. Neuroscience can keep supplying useful principles, materials research can keep opening up new physical mechanisms, computer architecture can provide scalable organization, and machine learning can provide practical training methods. If these strands continue to converge, neuromorphic systems stand a reasonable chance of making a real contribution to sustainable AI, edge intelligence, robotics, healthcare, and scientific computing — not by replacing what already works, but by covering the ground conventional computing was never well suited to in the first place.

Data Availability Statement

No original dataset was generated for this review paper. The discussion is based on the cited literature.

References

[1] G. Indiveri and S.-C. Liu, "Memory and information processing in neuromorphic systems," Proceedings of the IEEE, vol. 103, no. 8, pp. 1379–1397, 2015.

[2] A. S. Bhat, T. S. A. Kumar, and S. M. I. S. R. authors, "Computing inspired by the brain: a journey from algorithms to organoid intelligence," Nature Computational Science, 2026.

[3] G. Li, Y. Zhang, J. Wang, et al., "Modeling macroscopic brain dynamics with brain-inspired computing," Nature Communications, 2025.

[4] W. Zhang, S. Ma, X. Ji, X. Liu, Y. Cong, and L. Shi, "The development of general-purpose brain-inspired computing," Nature Electronics, vol. 7, pp. 954–965, 2024.

[5] M. Davies, N. Srinivasa, T.-H. Lin, G. N. Chunchu, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, et al., "Loihi: A neuromorphic manycore processor with on-chip learning," IEEE Micro, vol. 38, no. 1, pp. 82–99, 2018.

[6] D. Kudithipudi, C. Schuman, C. M. Vineyard, T. Pandit, C. Merkel, R. Kubendran, J. B. Aimone, G. Orchard, C. Mayr, R. Benosman, J. Hays, C. Young, C. Bartolozzi, A. Majumdar, S. G. Cardwell, M. Payvand, S. Buckley, S. Kulkarni, H. A. Gonzalez, G. Cauwenberghs, C. S. Thakur, A. Subramoney, and S. Furber, "Neuromorphic computing at scale," Nature, vol. 637, pp. 801–812, 2025.

[7] A. Basu, J. Acharya, M. Bhattacharya, and S. K. Nandy, "Brain-inspired neural networks: neuromorphic devices and their applications," Nanotechnology, 2025.

[8] S. Furber, C. Frenkel, and A. Rubino, "Boosting AI with neuromorphic computing," Nature Computational Science, 2025.

[9] G. Li, Y. Zhang, X. Liu, and J. Chen, "Brain-inspired computing: a systematic survey and future trends," Proceedings of the IEEE, 2025.

[10] A. R. Alimardani, M. A. H. Hassan, and S. S. M. Rahman, "Integrating brain–computer interface and neuromorphic computing for human digital twins," arXiv preprint arXiv:2410.23639, 2024.

[11] M. Payvand, C. Frenkel, and G. Indiveri, "Brain-inspired computing: from neuroscience to neuromorphic electronics," Frontiers in Neuroscience, editorial, 2025.

[12] G. Indiveri, B. Linares-Barranco, R. Legenstein, and S.-C. Liu, "Brain-inspired computing: from neuroscience to neuromorphic electronics for new forms of artificial intelligence," Frontiers research topic, 2023–2025.