

Energy-Efficient Neuromorphic Chips for Real-Time Robotic Control: A Review

Shuming Liu

*School of Engineering, University of Birmingham, Birmingham, United Kingdom
liushuming653@gmail.com*

Abstract. Neuromorphic computing has gained increasing attention as a bio-inspired solution to the limitations of traditional computing systems in power and real-time constraints. In robotic control tasks, real-time processing and energy efficiency are crucial, especially for autonomous and mobile systems. Neuromorphic chips mimic the structure and operation of biological neurons, enabling low-latency and low-power processing. This review examines the architectures, performance benchmarks, and application cases of mainstream neuromorphic hardware platforms used in robotic perception and control. It also compares their energy consumption and latency with conventional control platforms. Furthermore, this paper identifies key challenges in system integration and suggests future directions including improved scalability, online learning, and the combination with edge AI frameworks. The paper finds that neuromorphic chips can significantly reduce energy consumption and improve real-time performance in robotic control. However, challenges in algorithm adaptation, standardization, and hardware integration remain. This study aims to provide researchers with insights into the practical implementation of neuromorphic control systems in energy-sensitive robotic applications.

Keywords: Neuromorphic computing, Spiking neural network, Real-time control, Energy efficiency, Robotics

1. Introduction

Neuromorphic computing, inspired by the human brain, has emerged as a promising paradigm to overcome the limitations of traditional computing systems in terms of latency and energy consumption. In the field of robotics, real-time control is a core requirement, especially for tasks such as navigation, obstacle avoidance, and sensory feedback processing. Traditional processors such as Central Processing Unit (CPUs) and Graphics Processing Unit (GPUs), while powerful, often lead to excessive power draw and processing delays, which are unsuitable for energy-sensitive autonomous systems.

Recent research has focused on neuromorphic hardware platforms such as Intel's Loihi, IBM's TrueNorth, and SpiNNaker, which emulate spiking neural networks (SNNs) to achieve energy-efficient computation. According to Davies et al., Loihi supports event-driven computation and on-chip learning, making it suitable for mobile robotic applications [1]. Furber et al. demonstrate that

SpiNNaker can efficiently model large-scale neural systems for real-time control [2]. These works highlight the growing interest and potential of neuromorphic control in practical robotics.

This paper adopts a literature review method to explore and evaluate the current state of neuromorphic chips applied to robotic control tasks. The review categorizes major chip platforms, compares their energy and latency performance, and discusses open challenges and future research directions. The paper contributes to the understanding of how neuromorphic systems can reshape low-power autonomous control and guide future designs of robotic computing architectures.

2. Neuromorphic chip architectures for robotic control

Neuromorphic chips emulate the spiking behavior of biological neurons to achieve energy-efficient and real-time computation. Three representative platforms are frequently cited in the literature: IBM's TrueNorth, Intel's Loihi, and the University of Manchester's SpiNNaker. Each adopts a different architectural strategy. TrueNorth implements a hardwired SNN core optimized for static pattern recognition, while Loihi introduces programmable learning rules and on-chip plasticity, making it suitable for adaptive robotic tasks. SpiNNaker, on the other hand, is a massively parallel digital architecture ideal for modeling large-scale SNNs. Table 1 summarizes their key specifications and differences. These architectural distinctions directly affect their suitability for control tasks involving adaptability, learning, and energy sensitivity.

Table 1. Comparative summary of neuromorphic hardware platforms: Loihi, TrueNorth, and SpiNNaker [1-4]

Platform	Process	Compute Cores	Max Neurons / Synapses	On-chip Learning	Interconnect/ I/O	Representative Power/ Notes
Intel Loihi (2018)	14 nm; ~60 mm ²	128 neuro-cores + 3 x86 management cores	~1.3×10 ⁵ neurons; ~1.3×10 ⁸ synapses (literature approx.)	Supported (programmable synaptic plasticity; online learning)	On-chip asynchronous NoC; chip-to-chip links (board-level scaling)	Event-driven; energy per synaptic event (task-dependent)
IBM TrueNorth (2014/2015)	28 nm	4096 neurosynaptic cores	1×10 ⁶ neurons; 2.56×10 ⁸ synapses	Not supported (offline training/mapping)	On-chip 2D event-routing network	~65 mW (representative chip power) SpiNNaker (v1)
SpiNNaker (v1)	UMC 130 nm; ~102 mm ²	18× ARM968E-S per chip + on-chip router	Software neurons (depends on mapping and system scale)	Supported (STDP etc. implemented in software on cores)	Multicast packet routing; chip-to-chip links; Ethernet host interface	~1 W peak per chip (180 MHz, all cores)

3. Application of neuromorphic chips in robotic control

Neuromorphic chips have been deployed across core robotic-control tasks that demand low latency and low power. This section groups representative use cases into path planning, obstacle avoidance, and sensor fusion, highlighting how Loihi, SpiNNaker, and TrueNorth are used in practice.

3.1. Path planning and navigation

SNN-based controllers can implement spatial representations and reactive policies for navigation. Loihi has been combined with event-based vision to realize real-time, energy-efficient navigation, where on-chip learning and event-driven computation reduce latency and power draw compared with conventional platforms [1,3,5,6]. These properties make Loihi-style architectures suitable for onboard planning on mobile robots operating under tight energy budgets.

3.2. Obstacle avoidance

Event-driven perception paired with spiking controllers supports fast reflex-like avoidance. Using dynamic vision sensors as input to SNNs yields microsecond-level response to motion, enabling rapid detection and evasive maneuvers at low power. TrueNorth prototypes have also been used for recognition-driven control pipelines, where low-power pattern recognition informs downstream avoidance or gating actions [4].

3.3. Sensor fusion and motor control

SNNs naturally fuse asynchronous modalities (e.g., vision, IMU/proprioception) via population codes and spikes. On SpiNNaker, large SNNs have been mapped to multi-core fabrics for closed-loop locomotion and multi-sensor control, demonstrating scalable real-time performance under strict power limits [2]. Mixed-signal neuromorphic processors have further shown adaptive motor control with online learning in the loop, illustrating how spiking controllers can maintain stable behavior while adapting to perturbations [3,7].

In summary, across path planning, obstacle avoidance, and sensor fusion, neuromorphic implementations process sensor streams in real time and issue motor commands at substantially lower energy than CPU/GPU baselines. Loihi enables on-chip adaptation, SpiNNaker scales large SNN controllers across many cores, and TrueNorth provides ultra-low-power recognition pipelines that can drive control decisions. These properties make neuromorphic chips well suited for energy-constrained autonomous platforms.

4. Comparative energy and latency performance

Compared to traditional GPU or CPU-based robotic systems, neuromorphic chips show a significant advantage in terms of energy-delay product. Loihi, for example, achieved over three orders of magnitude improvement in energy efficiency when solving optimization problems compared to CPUs under equivalent process and area constraints [1]. A 2022 study reported a custom SNN accelerator achieving 0.72mW power consumption and 1.47 μ s latency, far surpassing traditional embedded controllers [6]. Moreover, neuromorphic architectures benefit from event-driven operations, meaning they consume power only when computation occurs, unlike synchronous platforms. These characteristics make them ideal candidates for battery-powered robots and edge AI systems.

4.1. Energy efficiency

Across control-related workloads, neuromorphic processors report orders-of-magnitude energy savings over CPU/GPU baselines. For instance, Loihi demonstrates over three orders of magnitude higher energy efficiency than CPUs on selected optimization problems when normalized for

process/area, illustrating the benefit of sparse, event-driven computation [1]. Mixed-signal and digital neuromorphic prototypes also show sub-milliwatt operation in closed-loop settings, highlighting suitability for battery-powered platforms. In contrast to fixed-throughput synchronous processors, neuromorphic cores naturally idle with no activity, improving average energy per task.

4.2. Latency and real-time response

Neuromorphic controllers achieve microsecond-to-millisecond control latencies depending on the task and sensing pipeline. A 2022 study reports a custom SNN accelerator running closed-loop control at 0.72 mW with 1.47 μ s latency, surpassing traditional embedded controllers under comparable tasks. Event-based perception feeding SNNs further reduces end-to-end delay by avoiding frame aggregation, which is crucial for high-speed reflexes in navigation and obstacle avoidance [6].

4.3. Event-driven advantage and EDP

Because spikes occur only when informative events happen, computation and memory access scale with activity, not wall-clock time. This yields lower energy per inference at similar or better response times, improving EDP relative to synchronous CPU/GPU pipelines—especially under sparse sensory input typical of mobile robots.

Neuromorphic chips deliver substantial energy savings and competitive (often superior) latency for closed-loop robotic control, leading to favorable EDP for edge deployment (see Table 2).

Table 2. Energy and latency metrics reported for representative neuromorphic systems [1,6]

Platform / Study	Benchmark / Task	Energy Metric	Latency Metric	Baseline / Relative
Intel Loihi (2018)	Optimization problems (event-driven SNN)	$\geq 1000\times$ energy efficiency vs CPU (under equivalent process/area)	Task-dependent (event-driven; no fixed TDP)	CPU baseline; qualitative latency advantage
Custom SNN accelerator (2022)	Closed-loop control micro-benchmarks	0.72 mW	1.47 μ s	Surpasses traditional embedded controllers

5. Challenges and future trends

Despite strong promise, several practical gaps still limit large-scale deployment of neuromorphic control. First, software toolchains and training remain immature: coding, mapping, and debugging SNN controllers is fragmented across platforms (e.g., NxSDK/Lava for Loihi, PyNN for SpiNNaker, legacy Corelet/Compass for TrueNorth), and current training routes (surrogate gradients, ANN-to-SNN conversion) are not yet turnkey for robotics. A near-term direction is to unify front ends with PyTorch-compatible APIs and PyNN back ends, provide ROS/ROS2 integration packages and examples, and standardize model exchange/mapping reports so experiments are reproducible and portable across chips. Second, lack of standardization and interoperability—from model formats to event-sensor interfaces—raises integration cost with mainstream robotics stacks [8]. Community efforts should converge on common event-stream formats (for DVS/IMU), ROS2 messaging, and a minimal SNN IR so the same controller can be compiled to multiple targets with predictable behavior. Third, benchmarking is inconsistent, with papers mixing chip vs. board power, different

event rates, and open- vs. closed-loop settings, making cross-platform claims weak. It is recommended to report full experimental conditions, providing paired CPU/GPU baselines, and using energy–delay product (EDP) alongside energy and latency on a small suite of canonical robotics tasks. Fourth, generalization and on-chip adaptation are limited: many SNN controllers are task-specific and brittle under domain shifts; on-chip learning exists but is constrained by resources and stability. A practical path is safety-bounded adaptation (e.g., gain tuning/reflex layers) on chip while keeping a frozen policy backbone, complemented by offline meta-learning or few-shot updates, as shown in mixed-signal demonstrations of adaptive motor control [1, 6, 7]. Finally, system-level integration and scaling remain non-trivial: sensor-to-chip I/O, memory locality, and multi-chip communication can introduce latency/jitter that harms real-time loops. Co-design of perception and control with activity-aware routing/memory placement, priority-aware inter-chip links (e.g., SpiNNaker fabric, TrueNorth 2D network, Loihi chip-to-chip links), and standardized power domains/telemetry for battery operation are promising remedies [9]. Overall, tighter coupling with edge-AI stacks, stronger and safer on-chip learning, and community benchmarks/standards are likely to turn today's demos into repeatable, portable controllers for autonomous robots .

6. Conclusion

This paper reviews the current state and future potential of neuromorphic chips in real-time robotic control systems, with a particular focus on energy efficiency. Through an examination of representative neuromorphic platforms such as Intel's Loihi, IBM's TrueNorth, and SpiNNaker, we summarized their architectural differences, control applications, and performance in terms of energy consumption and latency. The review demonstrates that neuromorphic computing offers significant advantages in low-power robotic applications, particularly in scenarios requiring real-time responsiveness and edge-level deployment.

However, current research faces limitations. The absence of standardized development tools and training frameworks for spiking neural networks (SNNs) hinders widespread adoption. Additionally, inconsistent performance benchmarking across platforms complicates direct comparisons and the selection of suitable hardware for specific applications. These challenges point to the need for unified programming environments, modular architectures, and real-world validation.

Looking forward, neuromorphic chips are poised to play a crucial role in the evolution of intelligent and autonomous systems. Integrating neuromorphic hardware with edge AI and developing hybrid control frameworks could enhance efficiency and adaptability. Future advancements may also focus on low-power learning algorithms, flexible hardware, and cross-domain benchmarking, establishing neuromorphic computing as a promising frontier for energy-efficient, intelligent robotic systems.

References

- [1] Davies, M., Srinivasa, N., Lin, T. H., China, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., & Liao, Y. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1), 82–99. doi: 10.1109/MM.2018.112130359.
- [2] Furber, S. B., Galluppi, F., Temple, S., & Plana, L. A. (2014). The SpiNNaker project. *Proceedings of the IEEE*, 102(5), 652–665. doi: 10.1109/JPROC.2014.2304638.
- [3] Aitsam, M., Di Nuovo, A., & Ahmad, S. (2022). Neuromorphic computing for interactive robotics: A systematic review. *IEEE Access*, 10, 113952–113970. doi: 10.1109/ACCESS.2022.3219440.
- [4] Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., & Plank, J. S. (2017). A survey of neuromorphic computing and neural networks in hardware. *arXiv preprint*. Available: <https://arxiv.org/abs/1705.06963>

- [5] Yang, Y., Wang, M., et al. (2023). Neuromorphic electronics for robotic perception and navigation: Current trends and future outlook. *Neurocomputing*, 533. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0952197623010229>
- [6] Sandamirskaya, Y., Kaboli, M., Conradt, J., & Celikel, T. (2022). Neuromorphic computing hardware and neural architectures for robotics. *Science Robotics*, 7(67), eabl8419.
- [7] Glatz, S., Pedersen, B. U., Mayr, C., & Schmuker, M. (2018). Adaptive motor control and learning in a spiking neural network realised on a mixed-signal neuromorphic processor. *arXiv preprint*. Available: <https://arxiv.org/abs/1810.10801>
- [8] Putra, R. V. W., Shibata, M., & Yamashita, Y. (2024). Embodied neuromorphic artificial intelligence for robotics: Perspectives, challenges, and research development stack. *arXiv preprint*. Available: <https://arxiv.org/abs/2404.03325>
- [9] Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197), 668–673. doi: 10.1126/science.1254642.