

Neuro-Morphic Hardware for Real-Time Embodied Cognition in Humanoid Robots

Jay Pak *

Shenzhen Merchiston international school, Shenzhen, Guangdong, China

* Corresponding author Email: jayblabla@163.com

Abstract: Embodying intelligence in humanoid robots requires continuous, low-latency integration of multimodal streams of sensory input along with adaptive control policies. Typical von Neumann architecture is ineffective at meeting hard power, latency, and parallelism constraints imposed by on-board, real-time embodied intelligence. This work outlines an integrated research study of neuromorphic hardware and spiking neural network architecture for providing real-time embodied cognition over a humanoid robotic platform. We architected a neuromorphic processing unit (NPU) based on a spiking architecture and co-designed sensor interfaces, learning rules, and motor control policies in order to capitalize upon event-driven computing. The system was evaluated in a series of closed-loop tasks centered upon real-time perception–action loops: reactive balance recovery, tactile-guided object manipulation, and fast visuomotor reaching under environmental perturbation. Performance was compared against a benchmark conventional controller running optimized deep learning and control stacks in low-power CPU/GPU hardware. We discover that the neuromorphic system facilitates equivalent task performance, decreased end-to-end perception-to-actuation latency, and decreased energy consumption by wide margins. More, the spiking networks enable continuous on-line adaptation based on local plasticity rules, enabling fast recovery from novel perturbations without off-line retraining. These outcomes indicate a direction for a potential use of neuromorphic hardware, when co-designed alongside embodied control algorithms, as providing a direction for energy-efficient, low-latency cognitive behavior in humanoid robots.

Keywords: Neuromorphic Hardware; Spiking Neural Networks; Embodied Cognition; Humanoid Robots; Real-Time Control; On-Line Adaptation; Event-Driven Perception.

1. Introduction

Robots that are humanoid and function in dynamic, uncertain worlds need to perceive, decide, and act in closed-loop cycles constantly, where delay or excessive power consumption is possible, leading to a loss in safety and performance. Embodied cognition is a concept that underscores how intelligence is generated from close coupling among sensory input, motor output, and physical properties of the body. Embedding embodied cognition in mobile humanoid robots raises significant computational challenges: sensory bandwidths from cameras, event-based vision sensors, tactile, and proprioceptive sensors need to be combined and processed in real time; control policies need to adapt in real time to dynamic changes; and energy constraints require very efficient computing to enable extended autonomous behavior.

Classic deep learning and model-predictive control approaches have improved robot vision and control but generally typify batch-processed, centrally calculated latencies and power costs not compatible with fully onboard continuous processing in mobile humanoids. Neuromorphic computing, rooted in biological neural nets, envisions a completely different substrate for computing: event-driven, highly parallel processing of sparse spike-like input-output signals with local compute and co-located memory[1]. These properties promise orders-of-magnitude reductions in energy per operation and lower latencies in sensorimotor loops.

Previous work has shown neuromorphic superiority in closed-loop perception experiments, event-based visual processing, and low-level motor control building blocks, yet systematic deployments of neuromorphic hardware supporting embodied cognition in humanoid robots are scarce.

Essential open questions are how to design end-to-end sensing-to-actuation pipelines over neuromorphic platforms, how to facilitate on-line adaptive learning in accordance with local rules of plasticity, and how to measure trade-offs in latency, energy, and robustness relative to traditional platforms.

Here, we present a systematic exploration of neuromorphic hardware for real-time embodied cognition in a humanoid robot. We architected spiking neural network (SNN) controllers and percept modules designed for an NPU and coupled them to the robot's sensor array and actuators. Our work has three parts: (1) an engineering proof-of-concept of end-to-end embodied cognition pipelines running on neuromorphic modules; (2) experimental benchmarks contrasting latency, energy, and task performance relative to a standard low-power CPU/GPU reference platform; and (3) evaluation of on-line local adaptive behavior mechanisms supporting rapid adaptation without global retraining. The rest of the paper outlines experimental approaches, quantitative findings, and consequences for next robot cognitive architecture.

2. Methods

Robotic platform and sensors

Robotic

Experimental work took place on a mid-scale humanoid platform that had 20 actuated DOF based on compliant leg and arm actuators. The sensor suite included a stereo event-based camera providing high-temporal-resolution visual input, a matrix of distributed tactile sensors over both hands, joint encoders serving for proprioception, as well as an inertial measurement unit (IMU)[2]. Buses for power and

communication were brought out to enable a neuromorphic processing module to execute locally on-board and trade low-latency event streams with sensors and motor command modules.

Neuromorphic chip and standard baseline

The event-driven, on-chip synaptic memory-based spiking architecture was applied in a neuromorphic processing unit. The NPU included programmable time constants for synapses, local plasticity to perform spike-timing-dependent plasticity (STDP), and reward-modulated gating. For comparison purposes, we implemented an optimized classical stack in a low-power CPU, along with an energy-efficient GPU embedded in a chip. The classical stack employed frame-based percept pipelines and compact neural nets for policy inference, optimized for minimum latency and power consumption under hardware limitations[3].

Design of spiking networks and sensor interfaces

NPU perception modules took streams of events directly: event-driven visual stimuli were transformed to spiking inputs by address-event representation, tactile arrays sent spike trains in proportion to contact event rates, and IMU/proprrioceptive inputs were encoded in terms of a population spike code[4]. Visual feature detection used a hierarchical spiking convolutional network with temporal filters and spatial feature detectors[5]. Motor policies were as SNN controllers converting proprioceptive and perceptual spike patterns to motor command spike trains; a low-latency spike-to-torque interface was applied to convert these to actuator commands.

We formulated local adaptation rules specific to embodied behavior: unsupervised STDP for initial sensory feature discovery, and a three-factor, reward-dependent adaptation rule for control adaptation. The reward was extracted from subsystem-specific performance indicators calculated locally (e.g., error in balance, grasp succeed indicator), transmitted as neuromodulatory spikes to gating synapses.

Task and evaluation metric

Three closed-loop experiments emphasized different facets of embodied cognition:

Recovery from reactive balance loss: The robot stood in upright position and was provided with randomized lateral perturbations. Parameters consisted of time-to-recovery, maximum center-of-mass deviation, and corrective steps.

Tactile-guided manipulation of objects: The robot manipulated and grasped objects of variable compliance and orientation in response to tactile input to adjust grip and direction of motion. Performance measures included succeeding rate, grasp stability, and manipulation latency.

Rapid visuomotor reaching during perturbation: The robot performed visually guided pointing and reaching to moving objects as external perturbations were applied to the platform. Measures included endpoint accuracy, latency from movement of target to correction, and energy per trial.

We defined actuation command release time from sensory event arrival as end-to-end latency. The power consumption of processing equipment (NPU and CPU/GPU baseline) was measured during the entire time. The adaptability was estimated by introducing new disturbances or object properties during experimentation and measuring recovery without off-line retraining.

Experimental protocol

C Each experiment was performed in matched trials for the neuromorphic and standard systems. For the neuromorphic experiment, SNNs were seeded with light pretraining of

sensory filters only; control policies were generic initially and free to adapt online through local plasticity. For the standard experiment, perception and policy networks were pre-trained off-line and used for inference only; adaptation was based on off-line parameter updates carried out between trials and was thereby switched off during real-time comparison. Trials were randomized and replicated over several sessions and robot instances to evaluate repeatability. Statistical comparisons used task measures, latencies, and energy between experiments.

3. Results

The neuromorphic pipeline exhibited markedly lower end-to-end latencies in every task relative to the traditional baseline. Median latency from sensor inputs to actuation in event-driven touch and vision loops in the NPU was several times lower, permitting corrective action in tens of milliseconds or sooner than the baseline's frame-based processing latency. Active-task power was several times lower for the neuromorphic chip relative to the CPU/GPU stack, corresponding to lower energy in successful trials. The energy-latency trade was in favor of the neuromorphic approach, particularly in high-frequency sensorimotor loops where event-driven sparse computation skipped unnecessary processing.

Neuromorphic and conventional system success rates in tasks were comparable in the three tasks, in which neuromorphic controllers achieved similar or slightly superior endpoint accuracy in fast reaching and higher grasp stability in tactile manipulation. For reactive recovery from loss of balance, both systems recovered to upright following nominal perturbations, but where there were sudden, high-frequency perturbations, the neuromorphic system recovered sooner as it reduced peak center-of-mass excursions and corrective step numbers in such situations.

A characteristic strength of the neuromorphic system was its ability to rapidly adapt on-line through local plasticity. When faced with a rapid compliance transition of manipulated objects during trials, the NPU-based controller adapted grip parameters in a few trials and recovered high success levels without centralized retraining. Correspondingly, when robot joint damping was transitioned rapidly, the neuromorphic controller recovered control gains through reward-modulated synaptic adaptation and recovered stable balance faster than the non-adapting conventional controller. These adoptions were achieved at minimal added latency and negligible energetic cost, validating the possibility of continuous learning in embodied situations.

The event-based SNN architecture scaled well to larger sensor bandwidths: as event rates for visual or tactile input went up proportionally with environmental complexity, computation scaled an equivalent amount to event sparsity, avoiding linear scaling overheads in frame-based pipelines. The neuromorphic controller exhibited robustness to sensor noise by leveraging sparse, temporally precise spike codes and local plasticity mechanisms to suppress spurious input influence. The conventional baseline, by contrast, required additional filtering and additional computation to scale robustness in noisier conditions.

4. Discussion

Our experiments demonstrate that co-designed neuromorphic hardware, along with spiking network

architecture and local learning rules, is able to deliver power-efficient, low-latency embodied cognition in humanoid robots, while still permitting on-line adaptive behavior. The event-driven operation of neuromorphic processors has a natural complement in the sparse, temporally-wealthy information streams output by modern sensors, permitting close perception-action loops necessary for reactive and adaptive behavior.

This reduction in end-to-end latency is most significant for safety-critical reactions such as recovery from a loss of balance, where millisecond reductions are a decrease in falls and reductions in mechanical stress. The reduction in energy extends operational endurance for untethered humanoid robots, allowing longer missions or assigning power budget to other subsystems such as locomotion or communication.

On-line adaptation by local plasticity is a major strength relative to static, off-line learned policies. Biological nervous systems employ local learning for fast, context-specific calibration; neuromorphic controllers may emulate this capability at the hardware level. Importantly, adaptation shown herein did not entail centralized supervision or high-bandwidth updates of its parameters, and is thus applicable to distributed, embodied agents.

Nonetheless, there are a few remaining challenges. First, programming and debugging of spiking networks and neuromorphic hardware need novel toolchains and design patterns, and there is no mature development ecosystem available as there is for traditional deep learning. Second, realizing in spiking substrates the representational richness found in dense neural networks may demand attentive encoding and possible multi-stage hybrid architecture incorporating hybrid front-ends of neuromorphic flavor and back-ends of symbolic or optimization-based nature. Third, although local plasticity facilitates fast adaptation, long-term stability and preventing catastrophic interference during lifelong operation require meta-learning approaches and regulation of plasticity. From a systems perspective, successful deployment in humanoid robots requires tight co-design of sensors, encoders, learning rules, and motor interfaces[6]. Event-based sensors and SNNs are a transparent combination, but most existing robotic sensors are frame-based; closing these modalities in an efficient manner is an engineering problem. Safety and verification of adaptive properties in the neuromorphic controller also require new methods of certifying behavior in response to variability in internal state. It will be necessary for future work to explore hybrid infrastructure where high-frequency, low-latency loops are mapped to neuromorphic computing components and higher-level planning is reserved for traditional processors, allowing a stratified cognitive architecture. Scalability to diversely shaped humanoid bodies and mission profiles requires research in neuromorphic-friendly learnings

that generalize across tasks and bodies preferably through meta-plasticity or intrinsic homeostatics. Lastly, broad community software toolchain, benchmark, and standardized task creation will accelerate adoption and rigorous comparison across platforms[7].

5. Conclusion

This work offers a full-scale proof that embodied cognition in humanoid robots is possible in real time using neuromorphic hardware, and it reaches matching task performance at a fraction of latency and power consumption, and facilitates rapid on-line adaptation via local plasticity. These findings imply that neuromorphic methods are a promising direction towards energy-efficient, low-latency robotic thinking, especially in cases where fast sensorimotor integration and continuous adaptation are involved. Further co-design of hardware, algorithms, and sensors will be necessary in order to reach full potential of neuromorphic systems in highly complex humanoid applications.

References

- [1] Aitsam, M., Davies, S., & Di Nuovo, A. (2022). Neuromorphic computing for interactive robotics: A systematic review. *Ieee Access*, 10, 122261-122279.
- [2] Qu, J., Wang, W., Ren, X., Zhang, Y., Bu, L., & Liu, L. (2024). Embodied Neuromorphic Intelligence in Healthcare: Evaluating Pose-Matching Interaction Using fNIRS and Behavioral Data. *IEEE Internet of Things Journal*.
- [3] Rast, A. D., Adams, S. V., Davidson, S., Davies, S., Hopkins, M., Rowley, A., ... & Cangelosi, A. (2018). Behavioral learning in a cognitive neuromorphic robot: an integrative approach. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12), 6132-6144.
- [4] Shankar, S., Pan, Y., Jiang, H., Liu, Z., Darbandi, M. R., Lorenzo, A., ... & Liu, T. (2025). Bridging Brains and Machines: A Unified Frontier in Neuroscience, Artificial Intelligence, and Neuromorphic Systems. *arXiv preprint arXiv:2507.10722*.
- [5] Indiveri, G., & Sandamirskaya, Y. (2019). The importance of space and time for signal processing in neuromorphic agents: the challenge of developing low-power, autonomous agents that interact with the environment. *IEEE Signal Processing Magazine*, 36(6), 16-28.
- [6] Chicca, E., Stefanini, F., Bartolozzi, C., & Indiveri, G. (2014). Neuromorphic electronic circuits for building autonomous cognitive systems. *Proceedings of the IEEE*, 102(9), 1367-1388.
- [7] Richter, C., Jentzsch, S., Hostettler, R., Garrido, J. A., Ros, E., Knoll, A., ... & Conradt, J. (2016). Musculoskeletal robots: scalability in neural control. *IEEE Robotics & Automation Magazine*, 23(4), 128-137.