Evaluating Spiking Neural Network On Neuromorphic Platform For Human Activity Recognition

Sizhen Bian
sizhen.bian@pbl.ee.ethz.ch
ETH Zürich
Zürich, Switzerland

Michele Magno
michele.magno@pbl.ee.ethz.ch
ETH Zürich
Zürich, Switzerland

ABSTRACT

Energy efficiency and low latency are crucial requirements for designing wearable AI-empowered human activity recognition systems, due to the hard constraints of battery operations and closed-loop feedback. While neural network models have been extensively compressed to match the stringent edge requirements, spiking neural networks and event-based sensing are recently emerging as promising solutions to further improve performance due to their inherent energy efficiency and capacity to process spatiotemporal data in very low latency. This work aims to evaluate the effectiveness of spiking neural networks on neuromorphic processors in human activity recognition for wearable applications. The case of workout recognition with wrist-worn wearable motion sensors is used as a study. A multi-threshold delta modulation approach is utilized for encoding the input sensor data into spike trains to move the pipeline into the event-based approach. The spikes trains are then fed to a spiking neural network with direct-event training, and the trained model is deployed on the research neuromorphic platform from Intel, Loihi, to evaluate energy and latency efficiency. Test results show that the spike-based workouts recognition system can achieve a comparable accuracy (87.5%) comparable to the popular milliwatt RISC-V bases multi-core processor GAP8 with a traditional neural network (88.1%) while achieving two times better energy-delay product (0.66 µJ s vs. 1.32 µJ s).

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**.

KEYWORDS

neuromorphic computing, human activity recognition, spiking neural networks, workouts recognition, Loihi

1 INTRODUCTION

Human activity recognition(HAR) is an active research topic in ubiquitous computing and human-computer interaction that aims to digitalize human behaviors and provide feedback for a better understanding and assistance, especially in healthcare [4, 38], sports science [7, 25, 39], and rehabilitation [23, 29]. With the increasing popularity of wearable devices, more complex HAR tasks have been moved from the laboratory to real-life scenarios powered by local AI solutions [3, 5, 36]. However, processing data with machine learning and neural network models on wearables still presents many research challenges due to limited computational resources and power constraints. This has led to exploring various neural networks compressions strategies, such as pruning and quantization, to have tiny machine learning models that work with a kilobyte of memory and Giga operations per second in the best case [18, 24].

While exploring the energy and latency performance of the edge tiny neural networks, the spiking neural network (SNN), a network that mimics the biological neurons in the brain, is also emerging as a promising technology for energy-efficient edge computing. The increase in interest is due to the event-driven nature and its ability to process spatio-temporal data in real-time with low latency and energy efficiency. In traditional neural networks, the activation function is typically differentiable, allowing the backpropagation to train a network. However, in SNNs, the firing of a neuron is a discrete event rather than a differentiable continuous function; thus, traditional backpropagation cannot be used to train SNNs. Two approaches have been explored in the past years to overcome this challenge, direct SNN training with surrogate gradient [26, 30] and ANN-to-SNN by first accumulating the events [20, 21, 28, 37]. While falling behind in reaching state-of-the-art accuracy, direct training is more biologically similar and preserves the temporal resolution of the spikes. To fully make use of the biological plausibility and validate the envelope of energy and latency efficiency of SNN in HAR, this work will use the surrogate gradient approach for direct event training of a workout recognition dataset, as a case study of neuromorphic solution for HAR. The neuromorphic chip Loihi will be used to profile the SNN performance taking the state-of-the-art performance from the general ANN on advanced edge processors with dedicated hardware accelerators as the baseline.

2 RELATED WORK AND CONTRIBUTION

The previous application works with SNN are mostly focused on vision tasks with dynamic vision sensors (DVS) [1, 27], where the input data is a sequence of visual events. A few edge explorations with DVS were also presented to validate the performance of end-to-end neuromorphic platforms from sensing to computing [31]. Compared with vision applications, exploring SNN on low-dimensional

Table 1: SNN explorations in ubiquitous computing

Authors/Application Year	Sensor/ Dataset	Neuromor- phic platform	Encoding	Training	Accuracy (vs ANN)	Inference Latency [a] (ms)	Inference Energy (mJ)	Energy-Delay Product (µJ s)
[15]- 2019	ECG MIT-BIH	DYNAP	Delta Modu- lation	SVM+ rSNN	95.6% (94.2%)	NA[c]	NA	NA
[32]- 2021	EEG iEEG- HFO	DYNAP- SE	Delta Modu- lation	Direct SNN	78.0% (67.0%)[d]	NA	NA	NA
[14]- 2020	EMG Customized	Loihi	Delta Modu- lation	Direct SNN	55.7% (68.1%)	5.9 (3.8 on GPU)	0.173 (25.5 on GPU)	1.0 (97.3 on GPU)
[11]- 2019	Audio Customized	Loihi	Rate Encod- ing	ANN-to- SNN	97.9% (97.9%) [e]	3.38 (1.30 on GPU, 2.4 on Jetson)	0.27 (29.8 on GPU, 5.6 on Jetson)	0.91 (38.7 on GPU, 13.44 on Jetson)
[12]- 2023	Audio QUT- NOISE	Xylo	Power band bin to spike	Direct SNN	98.0% (97.9%)	100[f]	0.0093 [g] (0.25: MAX78000, 11.2: Cortex)	0.93
Ours	**IMU, Ca- pacitive RecGym**	**Loihi**	**MT Delta Modulation**	**Direct SNN**	**87.5% (88.1%)**	4.4 (3.2 on GAP8)	0.15 (0.41 on GAP8)	0.66 (1.31 on GAP8)

[a] Time elapsed between the end of the input and the classification.
[b] Only dynamic energy is considered in Loihi.
[c] Not Available.
[d] HFO was detected with morphology detector [19].

[e] True Positive.
[f] Median classification latency (from the onset of an audio sample until the first spike from the correct class output neuron).
[g] Dynamic energy consumption.

data like audio and sensor signals is relatively new and much less explored [17]. However, recent studies have shown promising results in using SNNs for ubiquitous computing with low-dimensional signal sensors, which can provide essential insights into related topics like HAR. Table 1 lists several recent studies that explore SNNs for ubiquitous computing with low dimensional signals and their resulting performance in different applications. Kyle et al. [13] and Federico et al. [15] explored the heartbeat classification with SNN with different training strategies and validated it on two neuromorphic processors with competitive accuracy. The energy-delay product (EDP) on Loihi shows over twenty-eight times more efficiency than the inference on a CPU. Enea et al. [14] run a direct-trained SNN on Loihi with customized EMG and DVS data set for hand gesture recognition. Similar to [13], the EMG results on Loihi outperform in EDP compared with the results on GPU by ninety-seven times more efficiently. Besides the biological signals, audio signals were also explored with SNN [11]. In [12], a fresh edge neuromorphic processor, Xylo, was used to classify ambient audios. An impressive inference energy was reported on Xylo with only 9.3µJ, over twenty-six times less energy than the edge IoT processor MAX78000 owning a convolutional hardware accelerator. One common result of those SNNs on low-dimensional signals is that SNN supplies state-of-the-art inference energy and impressive EDP compared with ANN on CPU and GPU. Besides this, the ANN-to-SNN training approach often results in competitive accuracy while the direct trained SNN shows incompetence in accuracy compared with the ANN result[14]. The reason is that while encoding the signal to spikes, for example, the delta modulation, information loss is happening, especially for fast and huge signal variations [2, 14, 35].

In this work, we bring the following contributions:

(1) We demonstrated the feasibility of using SNN for sensor-based HAR tasks pursuing latency and energy efficiency with a direct-trained SNN on the neuromorphic platform Loihi. The first Spiking-IMU dataset and the corresponding direct-trained SNN are released for benchmarking of HAR with the neuromorphic solution [1].
(2) With spike trains generated by a multi-threshold delta modulation approach, a comparable accuracy (87.5%) is achieved compared with the ANN approach on the novel IoT processor GAP8 (88.1%), which has a dedicated RISC-V cluster for hardware acceleration and presented the state of the art edge AI performance in a rich of applications.
(3) The latency and energy efficiency of the neuromorphic approach HAR and the mainstream approach HAR were compared in this case study, and it showed that the neuromorphic approach of HAR using SNNs on Loihi outperforms the ANN method in terms of inference energy on GAP8 while falling behind lightly in latency. However, the neuromorphic approach shows nearly two times the energy-delay product (0.66 µJ s vs. 1.31 µJ s).

3 SYSTEM ARCHITECTURE

Figure 1 depicts the pipeline of the proposed SNN for HAR applications, including three key steps: spike encoding from sensor data, off-line SNN training, and on line SNN inference on neuromorphic processor. To have a fair comparison with neural networks and low power digital processors, in this work we use a public data

[1]https://github.com/zhaxidele/HAR-with-SNN

Figure 1: Human activity recognition using the spiking neural network where the network is processed on neuromorphic platforms pursuing energy and latency efficiency

Figure 2: Encoded spike trains of signal Acc_x from the workout of BenchPress with thresholds of [0.00005, 0.0001, 0.0002, 0.0004, 0.0008]

set, RecGym [8], as a case study. However, it is important to notice that the proposed approach can be used with other HAR-related data sets from various sensing modalities. The data set records ten volunteers' gym sessions with a sensing unit composed of an IMU sensor and a Body Capacitance sensor[6, 10]. The sensing units were worn at three positions: on the wrist, pocket, and calf. Twelve gym activities are recorded, including eleven workouts like Arm-Curl, LegPress, and StairsClimber, and a "Null" activity when the volunteer hangs around between different workout sessions. Each participant performed the selected workouts for five sessions in five days. Altogether, fifty sessions of gym workout data are presented in this data set. In this study, we only focus on the motion signals with the sensing unit worn on the wrist.

3.1 Spike encoding

To directly train an SNN, traditional numerical sequential value needs to be transformed into spike streams that carry both temporal and spatial knowledge of the original signals. Different encoding approaches have been explored mostly for vision data transformation [22], like latency encoding, rate encoding, delta modulation, etc. Each has advantages and limitations and has been adopted in different works [2]. For example, latency encoding normally achieves the best processing latency and energy consumption with fewer synaptic operations while being more susceptible to noise. Rate coding is demonstrated to exist in sensory systems like the visual cortex and motor cortex [34], showing the best resilience to input noise while limited by a lengthy processing period. In this work, we used the delta modulation approach due to the optimal trade

of complexity and latency from both the firmware and hardware implementation. Moreover, the analog sensory information can be directly encoded to the spike train at the front end. To address the accuracy degradation caused by information loss during the encoding, we gave multiple thresholds for spike train generation. The relationship between the continuous signal $s(t)$ and its spiking counterpart $\hat{s}(t)$ is given by Equation 1.

$$\hat{s}_i(t) = \begin{cases} 1, & \text{if } s(t) - s(t-1) > \epsilon_i \\ -1, & \text{if } s(t) - s(t-1) < -\epsilon_i \end{cases} \tag{1}$$

where ϵ_i is the threshold empirically chosen for spike encoding, and i ($0 \sim 4$)represents the index of applied thresholds list. Figure 2 depicts the five spike train channels encoded from the X-axis of the accelerometer, where the fast and large signal trend gives spikes in more spike train channels. For the inertial data, the thresholds were empirically set to 0.00005 x ($i + 1$), while 0.0000125 x ($i + 1$) for the capacitance data, and the threshold element number was empirically set to five. In future work, a systematical exploration of choosing the best threshold value and element number will be explored. As seven continuous signals were collected in the data set, we got thirty-five (7x5) spike trains for SNN training and inferring. With two seconds time window as a classification instance, we got 81291 spiking samples for building the SNN model.

3.2 Spiking neural network and Loihi

One of main contribution of this paper is the design of a SNN model and its evaluation on the Intel Loihi research platform. Loihi [16] is an asynchronous neuromorphic digital processor mainly for research. The processor consists of a many-core mesh of 128 neuromorphic cores for spike processing and three synchronous Lakemont x86 cores to monitor and configure the network and assist with the injection and recording of input/output spikes. Each neuro core in Loihi can access its local memories independently without needing to share a global memory bus and can implement up to 1024 current-based leaky integrate and fire neurons. Among other research platforms, Loihi has been selected because it includes a software SDK to design and profile proposed SNN.

Our proposed SNN is composed of two convolutional and two dense layers (32C64C128D12D) with a kernel size of three, as Table 2 lists. The threshold for neuron spiking was empirically selected. The current and voltage decay constants for the leaky integrate and fire neurons were set to 1024 (32 ms) and 128 (4 ms), respectively. Before spike encoding, the data set was interpolated to 1 kHz using

the univariate spline method, aiming to maximally approach the biological behaviors of the brain regarding information feeding. Each sample contains two seconds length of spike trains. The model was trained offline on GPU with weighted classes and leave-one-user-out cross-validation, and the trained weights and delays were then used to configure the network on Loihi hardware for inference purposes.

To fully make use of the biological plausibility of SNN, we used the framework SLAYER [33] for direct training, aiming to pursue the envelope of energy and latency efficiency of SNN. SLAYER evaluates the gradient of the convolutional and linear layers in an SNN by a temporal credit assignment policy, which distributes the error credit of an error back both through layers and in time, as a spiking neuron's current state depends on its previous states. Then a probability density function is used to estimate the change in the neuron state, thus approximating the derivative of the spike function. With SLAYER, the synaptic weights and axonal delays can be trained, and some state-of-the-art performances have been achieved on neuromorphic datasets like the NMNIST and ibmDVS-Gesture [33]. SLAYER supports Loihi-specific implementation for neuron model and weight quantization.

Table 2: SNN model for the spiking RecGym dataset

	Type	Size	Feature Size	Features	Stride
0	Input	7x5x2	-	-	-
1	Conv	7x5x32	3x3	32	1
2	Conv	7x5x64	3x3	64	1
3	Dense	2240	-	128	-
4	Dense	128	-	12	-

4 EXPERIMENTAL EVALUATION

Table 3 lists the workouts classification performance with the trained SNN on Loihi. In comparison, we selected an ANN model using the same data set and being deployed on two different IoT processors, presented in [9]. Such a comparison has seldom been made, as previous SNN evaluations mostly used ANN deployed on GPU/CPU as the baselines. The result will be meaningful for developing ubiquitous neuromorphic edge computing by supplying a straightforward comparison with the state-of-the-art using mainstream solutions. The multi-threshold spike encoding approach results in an accuracy of 87.5% with the directly trained SNN, which is much better than the single-threshold encoding result (below 60%) and acceptable compared with the accuracy from the ANN approach considering that the accuracy of direct-trained SNN on spike streams degrades in most cases. The inference latency of SNN on Loihi implies the time elapsed between the end of the input and the classification output and is reported as 4.4 ms, which is also much better than the latency on general IoT processors like STM32 with Cortex-M7 core but falls behind slightly to the GAP8, which features 8 RISC-V cores for dedicated hardware acceleration. However, the neuromorphic pipeline outperforms in dynamic energy consumption (0.15 mJ), benefitting from the sparsity of the spike trains and the in-memory computing of Loihi, which results in an EDP of 0.66 µJ s, while the EDP on GAP8 and STM32 are almost two

Table 3: Classification profiling vs. general edge solutions

Hardware	Model	Accuracy	Latency (ms)	Energy (mJ)	Energy-Delay Product (µJ s)
Loihi (Neuromorphic)	SNN	87.5%	4.4	0.15	0.66
GAP8 (RISC-V)	ANN	88.1%	3.2	0.41	1.31
STM32 (Cortex-M7)	ANN	89.3%	20.88	8.07	168.5

times and over two hundred times higher, respectively. The energy reported here is the dynamic energy on Loihi, which is measured by enabling the energy probe during inference as the difference between the total energy consumed by the network and the static energy when the chip is idle. We have to acknowledge that Loihi is not for edge computing specifically. Instead, it is designed more for general-purpose neuromorphic research. Thus there is still space for raising the neuromorphic performance, for example, the spike injection speed (the primary x86 core always waits 1ms before allowing the system to continue to the next timestep). To have a more fair comparison, end-to-end solutions of neuromorphic and traditional approaches should be developed, adopting the newly released edge neuromorphic processors [12, 31].

5 CONCLUSION

This work explored the neuromorphic solution of human activity recognition with a typical case study of workout recognition. Neuromorphic solutions, mainly inferring the SNN associated with the neuromorphic processor, have been emerging benefiting from its latency and energy efficiency. We started with a multi-threshold delta modulation to encode the raw motion sensor signal into multiple spike trains, aiming to reduce the information loss during spike generation. A shallow SNN model was then used to train the spike-form workouts signal with the SLAYER framework. The model runs on Loihi showed a comparable accuracy of 87.5% and an impressive energy-delay-product of 0.66 µJ s, compared with the state-of-the-art ANN solution on GAP8. This work demonstrates the efficiency of neuromorphic solutions in ubiquitous computing that pursues latency and energy efficiency. For future work, we will focus on new features in neuromorphic solutions that exceed the traditional edge solutions, for example, learn on the fly that can adapt the SNN models for specific subjects and environments, boosting the inference accuracy. We will also explore the newly released edge neuromorphic platforms and Loihi2, which has redesigned asynchronous circuits supplying faster speed and enhanced learning capabilities, featuring multiple times performance boosting compared with its predecessor.

ACKNOWLEDGMENTS

REFERENCES

[1] Mohamed Akrarai, Nils Margotat, Gilles Sicard, and Laurent Fesquet. 2020. A Novel Event Based Image Sensor with spacial and temporal redundancy suppression. In *2020 18th IEEE International New Circuits and Systems Conference (NEWCAS)*. IEEE, 238–241.

[2] Daniel Auge, Julian Hille, Etienne Mueller, and Alois Knoll. 2021. A survey of encoding techniques for signal processing in spiking neural networks. *Neural Processing Letters* 53, 6 (2021), 4693–4710.

[3] Abu Bakar, Rishabh Goel, Jasper de Winkel, Jason Huang, Saad Ahmed, Bashima Islam, Przemysław Pawełczak, Kasım Sinan Yıldırım, and Josiah Hester. 2022. Protean: An energy-efficient and heterogeneous platform for adaptive and hardware-accelerated battery-free computing. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*. 207–221.

[4] Sarnab Bhattacharya, Rebecca Adaimi, and Edison Thomaz. 2022. Leveraging sound and wrist motion to detect activities of daily living with commodity smartwatches. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 2 (2022), 1–28.

[5] Sizhen Bian and Paul Lukowicz. 2021. Capacitive sensing based on-board hand gesture recognition with TinyML. In *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers*. 4–5.

[6] Sizhen Bian and Paul Lukowicz. 2021. A systematic study of the influence of various user specific and environmental factors on wearable human body capacitance sensing. In *EAI International Conference on Body Area Networks*. Springer, 247–274.

[7] Sizhen Bian, Vitor F Rey, Peter Hevesi, and Paul Lukowicz. 2019. Passive capacitive based approach for full body gym workout recognition and counting. In *2019 IEEE International Conference on Pervasive Computing and Communications (PerCom*. IEEE, 1–10.

[8] Sizhen Bian, Vitor Fortes Rey, Siyu Yuan, and Paul Lukowicz. 2022. The Contribution of Human Body Capacitance/Body-Area Electric Field To Individual and Collaborative Activity Recognition. *arXiv preprint arXiv:2210.14794* (2022).

[9] Sizhen Bian, Xiaying Wang, Tommaso Polonelli, and Michele Magno. 2022. Exploring Automatic Gym Workouts Recognition Locally On Wearable Resource-Constrained Devices. In *2022 IEEE 13th International Green and Sustainable Computing Conference (IGSC)*. IEEE, 1–6.

[10] Sizhen Bian, Siyu Yuan, Vitor Fortes Rey, and Paul Lukowicz. 2022. Using human body capacitance sensing to monitor leg motion dominated activities with a wrist worn device. In *Sensor-and Video-Based Activity and Behavior Computing: Proceedings of 3rd International Conference on Activity and Behavior Computing (ABC 2021)*. Springer, 81–94.

[11] Peter Blouw, Xuan Choo, Eric Hunsberger, and Chris Eliasmith. 2019. Benchmarking keyword spotting efficiency on neuromorphic hardware. In *Proceedings of the 7th annual neuro-inspired computational elements workshop*. 1–8.

[12] Hannah Bos and Dylan Muir. 2023. Sub-mW Neuromorphic SNN audio processing applications with Rockpool and Xylo. *Embedded Artificial Intelligence: Devices, Embedded Systems, and Industrial Applications* (2023), 69.

[13] Kyle Buettner and Alan D George. 2021. Heartbeat classification with spiking neural networks on the loihi neuromorphic processor. In *2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 138–143.

[14] Enea Ceolini, Charlotte Frenkel, Sumit Bam Shrestha, Gemma Taverni, Lyes Khacef, Melika Payvand, and Elisa Donati. 2020. Hand-gesture recognition based on EMG and event-based camera sensor fusion: A benchmark in neuromorphic computing. *Frontiers in Neuroscience* (2020), 637.

[15] Federico Corradi, Sandeep Pande, Jan Stuijt, Ning Qiao, Siebren Schaafsma, Giacomo Indiveri, and Francky Catthoor. 2019. ECG-based heartbeat classification in neuromorphic hardware. In *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.

[16] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro* 38, 1 (2018), 82–99.

[17] Mike Davies, Andreas Wild, Garrick Orchard, Yulia Sandamirskaya, Gabriel A Fonseca Guerra, Prasad Joshi, Philipp Plank, and Sumedh R Risbud. 2021. Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proc. IEEE* 109, 5 (2021), 911–934.

[18] Lei Deng, Guoqi Li, Song Han, Luping Shi, and Yuan Xie. 2020. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proc. IEEE* 108, 4 (2020), 485–532.

[19] Tommaso Fedele, Sergey Burnos, Ece Boran, Niklaus Krayenbühl, Peter Hilfiker, Thomas Grunwald, and Johannes Sarnthein. 2017. Resection of high frequency oscillations predicts seizure outcome in the individual patient. *Scientific Reports* 7, 1 (2017), 1–10.

[20] Vittorio Fra, Evelina Forno, Riccardo Pignari, Terrence C Stewart, Enrico Macii, and Gianvito Urgese. 2022. Human activity recognition: suitability of a neuromorphic approach for on-edge AIoT applications. *Neuromorphic Computing and Engineering* 2, 1 (2022), 014006.

[21] Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. 2019. End-to-end learning of representations for asynchronous event-based data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5633–5643.

[22] Wenzhe Guo, Mohammed E Fouda, Ahmed M Eltawil, and Khaled Nabil Salama. 2021. Neural coding in spiking neural networks: A comparative study for robust neuromorphic systems. *Frontiers in Neuroscience* 15 (2021), 638474.

[23] Shane Halloran, Lin Tang, Yu Guan, Jian Qing Shi, and Janet Eyre. 2019. Remote monitoring of stroke patients' rehabilitation using wearable accelerometers. In *Proceedings of the 23rd International Symposium on Wearable Computers*. 72–77.

[24] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149* (2015).

[25] Alexander Hoelzemann, Julia Lee Romero, Marius Bock, Kristof Van Laerhoven, and Qin Lv. 2023. Hang-Time HAR: A Benchmark Dataset for Basketball Activity Recognition using Wrist-worn Inertial Sensors. *arXiv preprint arXiv:2305.13124* (2023).

[26] Eimantas Ledinauskas, Julius Ruseckas, Alfonsas Juršėnas, and Giedrius Buračas. 2020. Training deep spiking neural networks. *arXiv preprint arXiv:2006.04436* (2020).

[27] Chenghan Li, Luca Longinotti, Federico Corradi, and Tobi Delbruck. 2019. A 132 by 104 10μm-Pixel 250μW 1kefps dynamic vision sensor with pixel-parallel noise and spatial redundancy suppression. In *2019 Symposium on VLSI Circuits*. IEEE, C216–C217.

[28] Yuhang Li, Ruokai Yin, Hyoungseob Park, Youngeun Kim, and Priyadarshini Panda. 2022. Wearable-based human activity recognition with spatio-temporal spiking neural networks. *arXiv preprint arXiv:2212.02233* (2022).

[29] Huajun Long, Jie Li, Rui Li, Xinfeng Liu, and Jingyuan Cheng. 2023. Dual-modality Smart Shoes for Quantitative Assessment of Hemiplegic Patients' Lower Limbs' Muscle Strength. *arXiv preprint arXiv:2305.13977* (2023).

[30] Emre O Neftci, Hesham Mostafa, and Friedemann Zenke. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine* 36, 6 (2019), 51–63.

[31] Georg Rutishauser, Robin Hunziker, Alfio Di Mauro, Sizhen Bian, Luca Benini, and Michele Magno. 2023. ColibriES: A Milliwatts RISC-V Based Embedded System Leveraging Neuromorphic and Neural Networks Hardware Accelerators for Low-Latency Closed-loop Control Applications. *arXiv preprint arXiv:2302.07957* (2023).

[32] Mohammadali Sharifshazileh, Karla Burelo, Johannes Sarnthein, and Giacomo Indiveri. 2021. An electronic neuromorphic system for real-time detection of high frequency oscillations (HFO) in intracranial EEG. *Nature communications* 12, 1 (2021), 3095.

[33] Sumit B Shrestha and Garrick Orchard. 2018. Slayer: Spike layer error reassignment in time. *Advances in neural information processing systems* 31 (2018).

[34] Kyle H Srivastava, Caroline M Holmes, Michiel Vellema, Andrea R Pack, Coen PH Elemans, Ilya Nemenman, and Samuel J Sober. 2017. Motor control by precisely timed spike patterns. *Proceedings of the National Academy of Sciences* 114, 5 (2017), 1171–1176.

[35] Antonio Vitale, Elisa Donati, Roger Germann, and Michele Magno. 2022. Neuromorphic Edge Computing for Biomedical Applications: Gesture Classification Using EMG Signals. *IEEE Sensors Journal* 22, 20 (2022), 19490–19499.

[36] Xiaying Wang, Michele Magno, Lukas Cavigelli, and Luca Benini. 2020. FANN-on-MCU: An open-source toolkit for energy-efficient neural network inference at the edge of the Internet of Things. *IEEE Internet of Things Journal* 7, 5 (2020), 4403–4417.

[37] Hao Wu, Yueyi Zhang, Wenming Weng, Yongting Zhang, Zhiwei Xiong, Zheng-Jun Zha, Xiaoyan Sun, and Feng Wu. 2021. Training spiking neural networks with accumulated spiking flow. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 10320–10328.

[38] Ruidong Zhang, Jihai Zhang, Nitish Gade, Peng Cao, Seyun Kim, Junchi Yan, and Cheng Zhang. 2022. EatingTrak: Detecting fine-grained eating moments in the wild using a wrist-mounted IMU. *Proceedings of the ACM on Human-Computer Interaction* 6, MHCI (2022), 1–22.

[39] Bo Zhou, Mathias Sundholm, Jingyuan Cheng, Heber Cruz, and Paul Lukowicz. 2016. Never skip leg day: A novel wearable approach to monitoring gym leg exercises. In *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 1–9.