

Human activity recognition: suitability of a neuromorphic approach for on-edge AIoT applications

Original

Human activity recognition: suitability of a neuromorphic approach for on-edge AIoT applications / Fra, Vittorio; Forno, Evelina; Pignari, Riccardo; Stewart, Terrence; Macii, Enrico; Urgese, Gianvito. - In: NEUROMORPHIC COMPUTING AND ENGINEERING. - ISSN 2634-4386. - ELETTRONICO. - 2:1(2022), pp. 1-14. [10.1088/2634-4386/ac4c38]

Availability:

This version is available at: 11583/2954499 since: 2023-11-24T12:33:14Z

Publisher:

IOP Publishing

Published

DOI:10.1088/2634-4386/ac4c38

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

PAPER • OPEN ACCESS

Human activity recognition: suitability of a neuromorphic approach for on-edge AIoT applications

To cite this article: Vittorio Fra *et al* 2022 *Neuromorph. Comput. Eng.* **2** 014006

View the [article online](#) for updates and enhancements.

You may also like

- [Progress and perspectives in dry processes for nanoscale feature fabrication: fine pattern transfer and high-aspect-ratio feature formation](#)
Taku Iwase, Yoshito Kamaji, Song Yun Kang et al.
- [Embedding human annoyance rate models in wireless smart sensors for assessing the influence of subway train-induced ambient vibration](#)
Ke Sun, Wei Zhang, Huaping Ding et al.
- [Nondestructive shape process monitoring of three-dimensional, high-aspect-ratio targets using through-focus scanning optical microscopy](#)
Ravi Kiran Attota, Hyeonngon Kang, Keana Scott et al.



PAPER

OPEN ACCESS

RECEIVED

30 November 2021

REVISED

31 December 2021

ACCEPTED FOR PUBLICATION

17 January 2022

PUBLISHED

7 February 2022

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Human activity recognition: suitability of a neuromorphic approach for on-edge AIoT applications

Vittorio Fra¹ , Evelina Forno¹ , Riccardo Pignari¹ , Terrence C Stewart² , Enrico Macii¹ and Gianvito Urgese^{1,*}

¹ Politecnico di Torino, Corso Duca degli Abruzzi, 24, Torino, TO, 10129, Italy

² National Research Council Canada, University of Waterloo Collaboration Centre, Waterloo, Ontario, Canada

* Author to whom any correspondence should be addressed.

E-mail: gianvito.urgese@polito.it

Keywords: neuromorphic engineering, human activity recognition, spiking neural network, artificial intelligence of things, edge computing, Legendre memory unit, raw IMU data

Supplementary material for this article is available [online](#)

Abstract

Human activity recognition (HAR) is a classification problem involving time-dependent signals produced by body monitoring, and its application domain covers all the aspects of human life, from healthcare to sport, from safety to smart environments. As such, it is naturally well suited for on-edge deployment of personalized point-of-care analyses or other tailored services for the user. However, typical smart and wearable devices suffer from relevant limitations regarding energy consumption, and this significantly hinders the possibility for successful employment of edge computing for tasks like HAR. In this paper, we investigate how this problem can be mitigated by adopting a neuromorphic approach. By comparing optimized classifiers based on traditional deep neural network architectures as well as on recent alternatives like the Legendre memory unit, we show how spiking neural networks can effectively deal with the temporal signals typical of HAR providing high performances at a low energy cost. By carrying out an application-oriented hyperparameter optimization, we also propose a methodology flexible to be extended to different domains, to enlarge the field of neuro-inspired classifier suitable for on-edge artificial intelligence of things applications.

1. Introduction

Fast growth and widespread availability of smart devices integrating a high number of sensors have significantly changed the idea of body sensor networks (BSNs) during the last few years, emphasizing the actual feasibility of the concept and successfully leading it towards a simplified and non-invasive monitoring of physiological and activity signals. Thanks to the wearable devices increasingly present in our daily life, BSNs consisting of a number of body-worn sensor nodes wirelessly collaborating can be shrunk down to a single device [1–3]. As a consequence, the coupling of device miniaturization and BSNs dimension reduction opened the way to a wide range of applications for wearable sensors, including, but not limited to: healthcare, elderly assistance, fitness, and gestures recognition [4–7]. Nonetheless, critical challenges still have to be faced, as new constraints must be taken into account when dealing with the limitations imposed by devices intended to be as small and portable as possible [8]. As extensively pointed out in [9], the ultimate goal of edge computing for wearable devices requires a change of paradigm materializing in a reduction of the computational efforts. Typical wearable sensors are indeed affected by severe limitations in terms of power, and the conventional approach—based on data transmission to off-chip, remote servers in charge of processing the acquired signals—introduces an additional limitation on the temporal side. Meeting these challenges would mean setting the scene for effective real-time processing of data, suited to enlarge the range of personalized services to be efficiently and widely deployed on smart edge devices [10–14].

Among the number of possible tasks related to body monitoring, human activity recognition (HAR) stands out for its relevance due to the inherent richness of information and the consequent adaptability to different applications. A reliable and responsive classification of ongoing user activity, besides being useful in monitoring activities of daily living, can also be decisive in safety-critical situations [15]. As a general definition, HAR is the analysis and classification of signals related to human actions, and it can be divided in categories according to the type of devices and sensors employed to acquire those signals [16]. However, due to the wide diffusion of smart devices and to their ease of use with just minor installation constraints, wearable sensors have been attracting the most attention in HAR research over the last decade [17, 18]. At the same time, from the algorithms standpoint, great efforts have been devoted to developing new machine learning (ML) models for more and more accurate classification, employing either traditional ML techniques or deep learning (DL) methods [19–22].

In this race to the best classification result, a vast majority of the proposed solutions have focussed only on the accuracy performances, not taking into account the possibility of shifting towards more biologically inspired models and thus leaving aside the alternative perspective offered by a neuromorphic approach. Nevertheless, adopting such neural-inspired paradigm, spiking neural networks (SNNs) [23], thanks to their event-based asynchronous operations, could represent a valuable candidate for energy-efficient solutions [24]. SNNs are indeed bio-inspired models offering practicable trade-offs between biological conformity and simulation runtimes, which can provide low-power computation as a result of their temporally sparse activity based on binary spikes [25, 26]. Such hallmark, which makes SNNs a clean-cut set of artificial neural networks (ANNs) with increased bio-plausibility, can be directly traced back to the distinguishing feature of biological neural networks with respect to non-spiking ANNs: in the brain, neurons communicate by spreading information in form of spikes, which are referred to as action potentials [27–29]. Relying on spike-based activity, SNNs also feature an intrinsic suitability for temporal information processing, which allows them to treat time as an additional dimension of the input signals [30, 31]. As a consequence of their brain-like, or at least brain-inspired, properties, these network models are of primary interest as a natural programming paradigm in neuromorphic computing.

Complementary to neural computing, whose primary aim is the implementation of ANNs to deal with practical tasks, neuromorphic computing brings the attention to the mimicking of neural processes in new and alternative computer architectures [32, 33], with the direct consequence of driving efforts towards the development of specific neuromorphic hardware [34–41]. Besides these platforms, neuromorphic simulators have also attracted significant interest, resulting in powerful tools able to sustain the neuromorphic computing growth even in cases where dedicated hardware is not readily available [42]. Among them, an intriguing example is represented by Nengo. Based on the Neural Engineering Framework (NEF) [43], it allows to build networks from single neuron models, providing the keys to access low-level neural archetypes to perform high-level functional tasks [44]. Additionally, the front-end API is designed to make it flexible and easy to adapt to specialized neuromorphic platforms, such as Intel's Loihi [45], as well as to deep learning models [46].


In this work, we present a comparative analysis of different models and architectures for the HAR task. We adopted the Wireless Sensor Data Mining (WISDM) smartphone and smartwatch activity and biometrics dataset [47, 48] to investigate a classification approach based on raw data only. Particularly, employing Nengo, we benchmark the beneficial effect of adopting a neuromorphic paradigm alternative to classical deep learning solutions; thus presenting, to the best of our knowledge, the first evaluation of bio-inspired models for HAR directly from raw data.

2. Background

In recent years, a rich literature has been produced in the HAR domain, and deep learning techniques have been extensively applied in a number of works testing them on different datasets [49]. Similarly, neuromorphic computing has been attracting growing interest, leading to an ongoing increase of attention on bio-inspired networks and dedicated hardware [50].

2.1. Human activity recognition

Independently of the targeted application, ranging from healthcare to surveillance, body monitoring focussed on movement can be classified according to the sensors or sensor network characteristics. In this regard, during the last decade, following the increasing amount of data easily recorded and collected through personal and non-invasive devices, a natural selection in favour of wearable sensors, generally embedded in smart devices, started to take place in the HAR field. As a result, many works adopting deep learning techniques, with solutions relying on convolutional neural networks (CNNs), recurrent neural networks (RNNs) or their combinations, have been produced for HAR tasks based on, or intended for, portable platforms [51–57].



In a high number of cases, smartphone sensors, typically inertial measurement units (IMUs), are employed in the gathering of datasets, like in the case of WISDM [58] and UCI-HAR [59] (also referred to as SBHAR [60]). Beside these well-established ones, a newer version of the WISDM dataset [47, 48] is also attracting growing interest, due to its more balanced classes and to the addition of smartwatch signals. Other wearable devices and sensors are then taken into account by datasets like PAMAP2 [61, 62], MHEALTH [63, 64], OPPORTUNITY [65, 66], PUC-Rio [67], WHARF [68], USC-HAD [69], and UTD-MHAD [70].

An example of benchmarking for network architectures on various datasets can be found in [71], where an extensive analysis is carried out involving different techniques and data. Furthermore, interesting insights into the impact of data segmentation on the classification accuracy are given. Identifying an optimal window size for time-varying signals like those treated in HAR is indeed of key importance from a twofold perspective. It can lead to high levels of accuracy, thus providing more reliable classifiers, while in terms of time-to-classify it can be crucial to assess suitability for real-time applications.

In [72], a summary of representative window sizes employed in the HAR task is reported, showing that typical choices fall within 1 s and 10 s. Exceptions can be found in works by Ordoñez *et al* [73], by Wan *et al* [74] and by Xia *et al* [75], where temporal windows down to 0.25 s are employed for different datasets. Mekruksavanich and co-workers [76, 77], as well as Oluwalade *et al* [78] and Ihianle *et al* [79], adopted instead a signal segmentation of 10 s.

2.1.1. WISDM dataset

In 2019, the Wireless Sensor Data Mining (WISDM) Lab published the WISDM smartphone and smartwatch activity and biometrics dataset [47, 48]. Composed of data from 51 subjects performing 18 activities, this dataset collects signals from both the accelerometer and the gyroscope of a smartphone and a smartwatch. Each activity is recorded for 3 min with an acquisition rate of 20 Hz. With respect to the older WISDM dataset [58], this version is not only enriched in number of activities but also improved in terms of class balance, with each activity represented in the dataset with a relative contribution ranging from 5.3% to 5.8% of the 15 630 426 total samples. Additionally, three subsets can be identified within the dataset according to activity type: *on-hand-oriented*, *general hand-oriented* and *eating hand-oriented*. As an example, the kernel density estimation of the 3D smartwatch data from accelerometer and gyroscope of the ‘general, hand-oriented’ subset of the WISDM dataset is shown in figure 1. Here, an overlap between the raw signal values can be appreciated.

2.2. Benchmarks and IoT applications of neuromorphic solutions

The actual chances of success for neuro-inspired and neuromorphic approaches promising energy efficiency improvements are being increasingly tested since the last few years. In [80], Blouw and co-workers benchmark different hardware performing keyword spotting tasks, showing a significantly reduced energy consumption when using Intel’s Loihi chip. Researchers’ achievements using Loihi are summarized in [33], and the platform has also been compared with a SpiNNaker 2 prototype by Yan *et al* [81].

In [82], the effectiveness of an ANN-to-SNN conversion addressing the heartbeat classification task and subsequently deployed on Loihi is evaluated. An extended benchmarking of neuromorphic hardware is then provided by Azghadi *et al* [83], who tested multiple platforms on biomedical applications.

In [84], the benefits of a neuromorphic approach are highlighted, assessing the computational cost reduction provided by SNNs developed in Nengo with respect to architecturally identical deep neural networks (DNNs).

In [85, 86], authors have investigated the advantages of using the SpiNNaker neuromorphic architecture [34] for executing massively parallel general-purpose algorithms such as PageRank and DNA sequence matching, implemented with the MPI paradigm.

Internet of things (IoT) is forecast to be one of the fields which will most benefit from the development of neuromorphic models and technologies. A survey of IoT platforms enabling artificial intelligence (AI) applications has been proposed by Kim *et al* [87], while the impact of neuromorphic systems on Industry 4.0 has been investigated in [88]. The role of edge computing in the artificial intelligence of things (AIoT) field, as well as in healthcare and other smart environments, has been instead reviewed by Chang *et al* [89]. Promising results for event-driven on-edge applications of AI in the IoT field have been shown in [41] presenting the neuromorphic IC μ Brain.

2.3. Nengo

Relying on the Neural Engineering Framework [43] as the guiding principle to build neural models accounting for functional objectives as well as anatomical constraints, Nengo was built as a simulator able to provide sophisticated networks featuring cognitive abilities starting from single neuron models [44].

The three NEF principles, namely representation, transformation and dynamics, are translated by Nengo into the fundamental units for networks construction, defining three core objects called ensemble, node and connection. Their combinations produce two further objects, network and model, while probe is defined as the object allowing to gather data during simulations. Such set of six front-end objects represents the toolkit to build the neural model to be passed to the simulator, which in turn encloses the back-end logic for the network simulation.

A key feature of Nengo is the flexibility of its simulator, ensured by the possibility of adapting it to specific, and possibly specialized, hardware [46]. For instance, NengoLoihi is a specialized backend for running Nengo models on Intel Loihi. Furthermore, as a result of this adaptability, models from different frameworks can be simply integrated through NengoDL's converter, which translates deep learning models by replacing standard activation functions with Nengo's spiking neurons.

2.4. Legendre memory unit

Neural communication relies on complex processes resulting in transmission and filtering of spikes through synapses. These mechanisms can be modelled by means of ordinary differential equations (ODEs) integrated over time, which allow to approximate the behaviour of time cells [90, 91]. The Legendre memory unit (LMU) is a recurrent architecture able to perform such approximation for a continuous-time delay [92]. The main property of the LMU network is the capability of decoding a delayed signal $u(t - \theta')$, contained within a sliding window of length θ , through a high-dimensional projection of the input $u(t)$ that is orthogonalized using the shifted Legendre polynomials [93]. The i th shifted Legendre polynomial is given by equation (1)

$$P_i(r) = (-1)^i \sum_{j=0}^i \binom{i}{j} \binom{i+j}{j} (-r)^j \quad (1)$$


and it is used to delay the input signal through equation (2)

$$u(t - \theta') \approx \sum_{i=0}^{d-1} P_i \left(\frac{\theta'}{\theta} \right) m_i(t) \quad (2)$$

where the highest order $d - 1$ in the series expansion is related to the dimension of the state vector $\mathbf{m}(t)$, defined by the input $u(t)$ as it follows in equation (3)

$$\theta \dot{\mathbf{m}}(t) = \mathbf{A} \mathbf{m}(t) + \mathbf{B} u(t) \quad (3)$$

with \mathbf{A} and \mathbf{B} representing the ideal state-space matrices derived using the Padé approximants through equations (4) and (5).



$$\mathbf{A} = [a]_{ij} \in \mathbf{R}^{d \times d}, \quad a_{ij} = (2i + 1) \begin{cases} -1 & i < j \\ (-1)^{i-j+1} & i \geq j \end{cases} \quad (4)$$

$$\mathbf{B} = [b]_i \in \mathbf{R}^{d \times 1}, \quad b_i = (2i + 1)(-1)^i, \quad i, j \in [0, d - 1]. \quad (5)$$

Although little literature has been produced so far on LMU applications, remarkable results have already been reported, showing state-of-the-art outcomes in terms of accuracy and interestingly small numbers of parameters when performing keyword spotting [94].

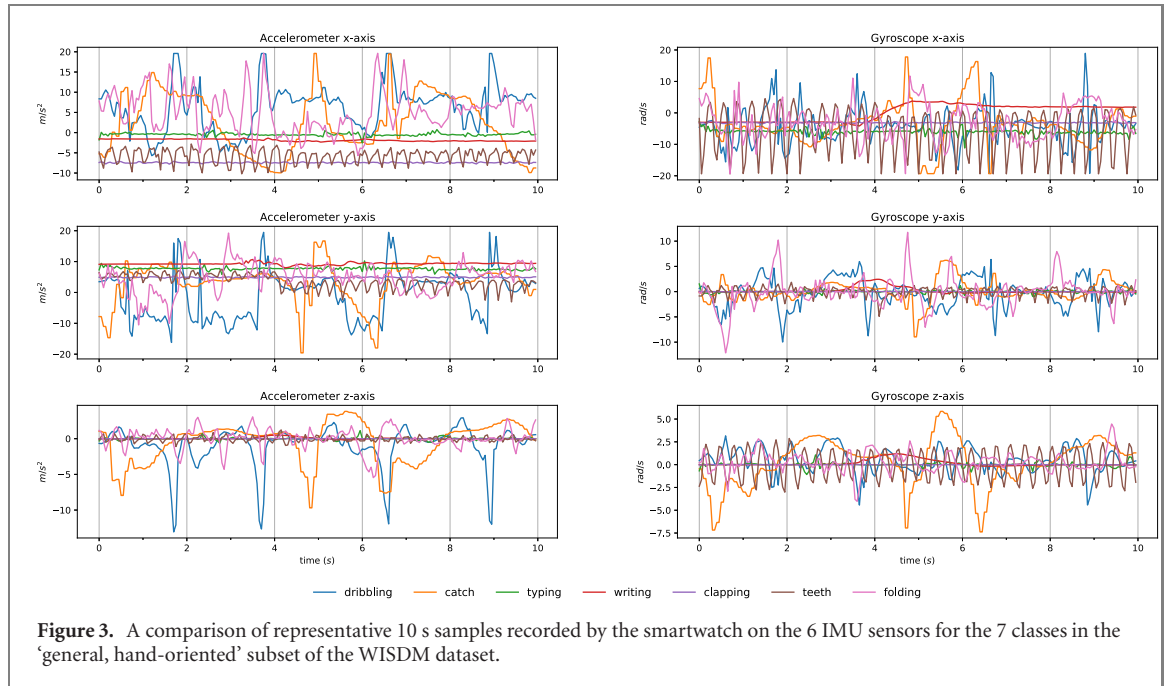
3. Methodology

The unprecedentedly huge amount of data produced in the IoT era is posing renewed challenges to cloud-based solutions based on back-and-forth transmission from end devices. Possible ways to face these difficulties are offered by the so-called fog computing and edge computing. The latter particularly aims at bringing data processing close to the sensors, moving computation down from the application layer to edge devices [95]. To accomplish this goal, the identification of lighter and less demanding computing solutions is of key importance, and the neuromorphic paradigm can provide well-suited tools.

To investigate this aspect, we propose a comparison of different neural networks of both recurrent and convolutional type, spiking and non-spiking. We also adopt neuro-inspired approaches to the HAR task through innovative solutions like the LMU, pointing out the differences between traditional DNNs and SNNs from a twofold perspective: beside the classification performance, we also evaluate the computational effort and memory demand. Such comparison is performed at the end of the optimization pipeline graphically summarized in figure 2. Here, vertical arrows identify preliminary steps, specifically involving dataset selection (a) and design of the optimization experiment (c) and (d), while horizontal arrows depict the subsequent phases along the backbone of the whole study: the neural network architectures selection (b), the hyperparameters optimization (e) and the final achievement of classifiers specifically tailored to HAR (f). In the following subsections we provide more details about each step of our investigation procedure.

3.1. Activity subset and time window

HAR, straightforwardly belonging to classification problems, begins with selection of the data, which can be either acquired on purpose or already collected in a dataset. In this work, we employed the data from smart devices available in the WISDM dataset. Specifically, due the increasing spread of wearable devices and their



suitability for tailored applications in different domains, we decided to focus on smartwatch data, glimpsing the opportunity for future adaptation of the proposed neuro-inspired approach to other wrist-worn devices, possibly employed for personalized point-of-care monitoring or other customized purposes to be brought as close as possible to the user. In this perspective, from the whole dataset we selected (step (a) in figure 2) the subset of general, hand-oriented activities: (1) dribbling in basketball, (2) playing catch with a tennis ball, (3) typing, (4) writing, (5) clapping, (6) brushing teeth and (7) folding clothes. With the aim of reducing as much as possible the required computational effort in view of on-edge deployment, the only preprocessing step performed was segmentation.

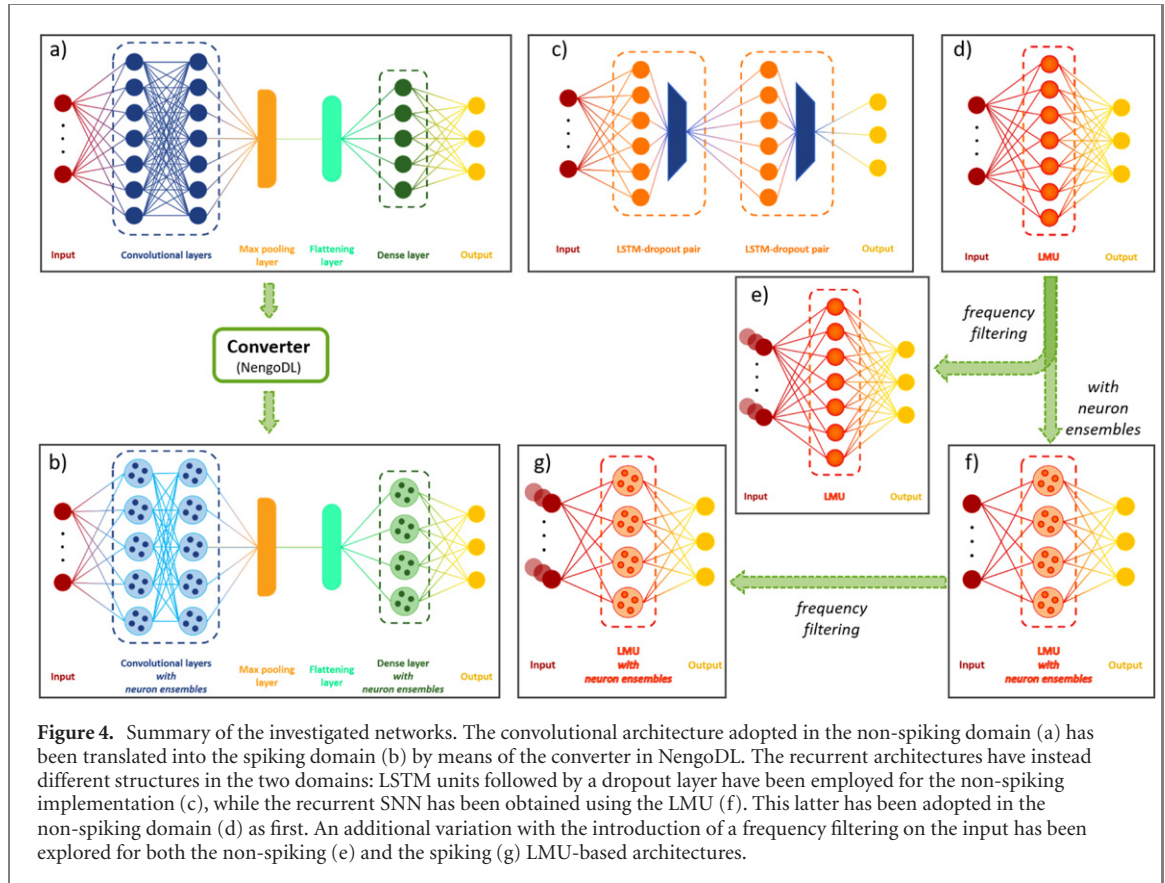
In figure 3, an example of raw 3-axial accelerometer and gyroscope data available in the WISDM dataset. Here, it is highlighted how classification of the raw input signals is not trivial when they are examined in real time, in the absence of any elaboration, filtering or data aggregation. We divided the signals into temporal windows, without overlap, with length of 2 s. Such choice was the result of an initial exploration including longer windows of 5 s and 10 s. With respect to these, the one with length of 2 s offered a valuable trade-off between the need for a sufficiently high number of temporal data for each sample and the goal of providing fast-response classifiers in an anthropocentric definition of real-time. The resulting 36 201 samples, based on raw data only, without any feature extraction, have then been split into training, validation and test set with a 60:20:20 proportion.

3.2. Network architectures

As previously introduced, HAR can be successfully performed employing either convolutional or recurrent architectures. At step (b) in figure 2 we accounted for both of these architecture types in order to benchmark possible alternatives offered by neuro-inspired solutions. The CNN comprises two convolutional layers followed by a max pooling layer, a flattening layer and two dense layers, as it is sketched in figure 4(a). We employed the same structure for both non-spiking and spiking convolutional neural networks (figure 4(b)), in the following referred to also as CNN and sCNN respectively. On the other hand, we implemented a recurrent architecture with a structure consisting of a sequence of two long short-term memory (LSTM) layers, each connected to a dropout layer, followed by a dense layer (figure 4(c)).

Differently from the case of the convolutional architectures, our spiking implementation of recurrent networks does not rely on the same architecture adopted in the non-spiking domain: as it is summarized by figure 4(f), we used the LMU in place of LSTMs, with a single LMU layer instead of the repeated LSTM-dropout pair. To further enrich the network comparison and benchmarking, we employed the LMU in a non-spiking network as well (figure 4(e)). The sCNN, the LMU-based network and its spiking version (sLMU) have been implemented by means of the Nengo neural simulator, employing the NengoDL converter to build the spiking CNN directly from its non-spiking counterpart.

We also worked with Nengo to investigate the impact of a human-inspired feature, borrowed from the auditory system, on networks based on the LMU: by analogy with the cochlea, we introduced a frequency filter (ff) on the input (figures 4(e) and (g)) decomposing the original signals into five channels through the



application, differently from the biological system, of a Butterworth filter bank. In all of the spiking networks under investigation we adopted the rectified integrate and fire neuron model available in Nengo and supported in the Loihi neuromorphic chip.

3.3. Hyperparameter optimization

ANNs can be characterized and described from two complementary perspectives. On the one hand, there is the architecture, namely the number and the type of layers employed and how they are connected to each other; on the other hand, there are the hyperparameters, which specifically identify each network determining its inherent behaviour. Consequently, as it is also pointed out in [96], hyperparameter optimization (HPO) must be accounted for when different network topologies are investigated and compared, especially in cases where unnecessary complexity must be prevented.

Steps from (c) to (e) in figure 2 summarize the procedure for the hyperparameter tuning we performed by means of the neural network intelligence (NNI) toolkit [97], using the built-in annealing algorithm. For each network, we designed an NNI optimization experiment, carried out within a proper search space defined at step (c). Each optimization experiment is composed of 1000 trials, with 4 evenly spaced random re-initializations of the tuner, intended to partially mitigate the problem of local minima affecting the annealing algorithms [98]. At the end of every trial, consisting of 100 training epochs, the weights providing the best training accuracy have been extracted to evaluate the test accuracy, defined to be the optimization objective of the experiment. For all the investigated networks, training has been performed employing Adam optimizer with constant learning rate, including optimization of this latter throughout the experiment trials. All these settings for the NNI experiments are performed in step (d) of figure 2. A summary of the hyperparameters contained in the search spaces employed for the HPO is reported in table 1.

3.4. Comparison criteria

At the end of the proposed pipeline, labelled as step (f) in figure 2, we obtained a trained classifier with optimized hyperparameters for each network architecture. We then set out to compare these classifiers, with the goal of assessing the advantages offered by neuro-inspired approaches without taking the risk of evaluating them from a narrow perspective mainly focussed on accuracy performances.

In order to make a comprehensive comparison between networks which rely not only on different architectures but also on different inherent working principles, we adopted multiple metrics beside the classification accuracy. The number of parameters and the memory footprint have been considered for all the networks. In

Table 1. Summary and description of the optimized hyperparameters. For the spiking networks, all the hyperparameters reported for the corresponding non-spiking implementation are take into account as well.

Network	Hyperparameter	Description
LSTM	units_1	Number of units in the LSTM layers
	units_2	
	dropout_1	Dropout rate between the LSTM layers
	dropout_2	
	l2_2	L2 regularization applied to the recurrent weights matrix in the second LSTM layer
CNN	filters_1	Number of filters in the convolutional layers
	filters_2	
	kernel_size_1	Dimension of the kernel in the convolutional layers
	kernel_size_2	
	dense_1	Number of units in the first dense layer
Spiking CNN	target_rate_1	Target value for neurons firing rates regularization in the convolutional layers
	target_rate_2	
	reg_conv_1	L2-like regularization applied to the neurons firing rates in the convolutional layers
	reg_conv_2	
	scale_firing_rates	Scale factor for the neurons firing rates
	Synapse	Time constant of the synaptic low-pass filter on the output of all the neurons
LMU	n_steps	How long (in simulation time steps ^a) the input is presented to the network
	Units	Size of the LMU kernels
	Order	Number of Legendre polynomials
	Theta	Length of the sliding window
	synapse_in	Time constant of the synaptic low-pass filter on the input connection of the LMU
	synapse_out	Time constant of the synaptic low-pass filter on the output connection of the LMU
	Tau	Time constant of the discretized synaptic low-pass filter on the internal connections to memory
Spiking LMU ^b	n_neurons	In place of units, size of the neuron ensembles (whose number is defined by order)
	synapse_all	Time constant of the synaptic low-pass filter on the connections between neuron ensembles
	max_rate	Firing rate for neuron input equal to 1
All	Batch size	Number of training examples in each learning iteration
	Learning rate	Step size for weights update in each learning iteration

^aThe default value in Nengo of 1 ms is used.

^bAll the hyperparameters for the non-spiking LMU are specifically re-optimized for the spiking implementation.

the case of non-spiking networks, we evaluated the number of floating point operations (FLOPs) and the corresponding estimated energy consumption on Intel's Movidius Neural Compute Stick 2. Whereas, for spiking networks we assessed the number of neurons, the number of synaptic operations (SOPs), and the corresponding estimated energy consumption on Intel's Loihi. Our energy evaluations rely on the results presented in [84].

4. Results and discussion


Neural networks benchmarking, and classifiers comparison in general, is naturally prone to the risk of an oversimplification taking shape in the evaluation of accuracy as the only meaningful metrics; and in the neuromorphic domain, such a simplistic approach can become even more deceptive. Although it certainly plays a decisive role, classification accuracy cannot be considered as independent of other figures of merit like energy consumption or memory footprint. Both of these quantities bring indeed with them crucial information for a deeper evaluation, and knowledge, of neuro-inspired solutions to classification problems. Especially when SNNs achieve classification performances comparable to other non-spiking DNNs, accuracy alone might not be enough to propose a fair comparison and a valuable benchmarking.

Consequently, as already introduced, in this work we accounted for multiple metrics. For each network, they have been evaluated taking into account the optimal hyperparameter configuration provided by specifically designed NNI experiments, each of them carried out performing 1000 trials. Thus guarantying a comparable development effort in optimizing the parameters for the different solutions.

Table 2 summarizes the considered metrics together with the corresponding values for each network. The optimized hyperparameters of the different architectures are instead reported in the supplementary material in tables S1 and S2 for the non-spiking and the spiking networks respectively (available online at <https://stacks.iop.org/NCE/2/014006/mmedia>).


Table 2. Summary of the evaluated metrics. The reported values have been obtained with the optimal hyperparameters configuration for each network.

	LSTM	CNN	sCNN	LMU	LMU (ff)	sLMU	sLMU (ff)
Test accuracy (%)	96.42 ± 0.03	93.81 ± 0.10	92.47 ± 0.08	91.71 ± 0.13	88.16 ± 0.13	94.51 ± 0.15	94.39 ± 0.13
Number of parameters	2125 222	144 899	167 973	76 130	89 014	91 200	132 540
Memory footprint (MB)	8.50	0.58	0.67	0.30	0.36	0.36	0.53
FLOPs ($\times 10^3$)	4249.65	2828.89	/	158.66	197.00	/	/
SOPs ($\times 10^3$)	/	/	10.82	/	/	99.91	127.95
Energy on Movidius (μJ)	3199.99	2130.15	/	119.47	148.34	/	/
Energy on Loihi (μJ)	/	/	5.49	/	/	50.66	64.87



Following its undeniable, although not unrivaled, leading role, classification accuracy is however the first quantity to be taken into account. The best result from this perspective is given by the LSTM-based network (table 2), which scores $(96.42 \pm 0.03)\%$. Interestingly, the second highest accuracy is achieved by its spiking counterpart. The recurrent network based on the spiking implementation of the LMU, relying on rectified integrate and fire neurons, provides indeed a test accuracy of $(94.51 \pm 0.15)\%$; which turns out to overcome the performances of the convolutional architecture regardless of whether it is adopted in the spiking or non-spiking domain. Similarly, the spiking LMU enriched with a frequency filtering inspired by the auditory system outperforms both spiking and non-spiking CNNs, with a test accuracy of $(94.39 \pm 0.13)\%$. To extend these results with reference to the employed dataset, and to complete the picture offered by the classification accuracies, the confusion matrices produced by all the investigated network on the test set are reported in the supplementary material in section S2. Additionally, in section S3 a comparison with other works adopting DL and ML techniques on WISDM dataset is presented, showing that the SNNs investigated in this work can match state-of-the-art results even overcoming them in some cases.

Further along the rows of table 2, the second metrics we considered is the total number of parameters, which directly leads to the memory footprint. From this perspective, what seemed a plausible forecast looking



at the classification accuracy, namely the LSTM-based network as the optimal solution, is overturn. With more than two millions of parameters, this architecture is indeed by far the most demanding in terms of memory footprint, with a size of 8.50 MB. At the opposite end is the network built on the non-spiking LMU, which is more than one order of magnitude smaller with only 0.30 MB of memory footprint. Similar values are found for LMU (ff) and sLMU also, while the spiking LMU with frequency filtering slightly exceeds these values almost reaching those of the convolutional architectures. The relative size of the different networks can be further appreciated from the circles diameter in figure 5.

Combining the information obtained from the results discussed above, it is straightforward to identify the best network in terms of accuracy and the one with the smallest memory footprint. However, these two networks do not coincide and they even are far from each other regarding both metrics, so that only partial conclusions can be drawn unless energy consumption is assessed. Taking such step further, namely quantifying the advantage of adopting a neuromorphic approach for the considered task, a triplet of fundamental quantities is eventually extracted from each network, making the reported benchmarking not only a comparison of values but also a tool to target possible future applications of the proposed neuro-inspired approach. The bottom two rows of table 2 highlight that energy consumption is assessed referring to two different and specialized hardware: for the non-spiking networks, Intel Movidius Neural Compute Stick 2 is considered, while Intel Loihi is taken into account for the spiking networks. In both cases, quantitative evaluations are made through the results of [84], which provide the energy cost for a single operation. In table 2, both the number of operations and the required energy per inference are reported for all the investigated network. The same results are also presented in figure 5, where the energy is on the y -axis and the number of operations defines the circles colour. As it is clearly shown, and expected, all the spiking networks are less computationally expensive, with the spiking CNN providing the lowest value of $5.49 \mu\text{J}$. It is worth also noting that the assessed energy consumption for all the LMU-based networks is one order of magnitude at least smaller than that of CNN and LSTM. Once again, the highest value is provided by the latter, even though in this case, as well as for the memory footprint, it does not correspond to the best result. With more than $3000 \mu\text{J}$, such architecture is indeed almost three orders of magnitude more energy-hungry than the sCNN. Within the range defined by these two opposite ends, the sLMU turns out to be the one with the highest accuracy at low energy cost: $50.66 \mu\text{J}$ to achieve the second best accuracy here reported, which means it is about two orders of magnitude lower in energy but comparable in accuracy with respect to the LSTM-based architecture.

The trade-off between high classification accuracy and small energy consumption offered by the spiking LMU, coupled with its reduced memory footprint, makes this architecture a relevant candidate for possible on-edge applications of neuromorphic classifiers for real-time tasks. In the view of evaluating the different

architectures to address specific tasks and applications, the presented results are also summarized in figure 6, where a radar chart is used to further highlight strengths and weaknesses of the investigated networks.

5. Conclusion

HAR is a time-dependent task whose application domain extends to all the aspects of human life, from health-care to sport, from safety to smart environments. In this paper, starting from the HAR problem, we have proposed an operational strategy to identify the optimal solution given a target application. Specifically, we have shown how a neuromorphic approach can be adopted to deal with time-varying inputs accounting for possible deployment constraints. By performing multiple optimization experiments, we have investigated the characteristics and the performances of multiple neural networks, highlighting advantages and drawbacks of recurrent and convolutional architectures with both spiking and non-spiking implementations. In this regard, we have reported a significant reduction in energy consumption for all the investigated spiking networks with respect to their non-spiking counterparts. In more detail, among these SNNs, the spiking implementation of the LMU has been pointed out as the optimal solution to achieve high classification accuracies with low energy consumption. With the analysis presented in this work, we hence have shown a suitable procedure to evaluate the possible benefits of a neuromorphic classifier for on-edge AIoT applications.

Acknowledgments

The research leading to these results has received funding from European Union Horizon 2020 Programme [H2020/2014-20] under Grant Agreement No. 785907 [HBP-SGA2]. Part of the work has been funded by ECSEL Joint Undertaking (EU H2020) under the Arrowhead Tools research project with Grant Agreement No. 826452.

Conflict of interest

The authors declare no conflict of interest.

Data availability statement

The data that support the findings of this study are openly available. GitHub page of source code used in the experiment: <https://github.com/neuromorphic-polito/NeHAR>. WISDM dataset [47]: <https://archive.ics.uci.edu/ml/datasets/WISDM+Smartphone+and+Smartwatch+Activity+and+Biometrics+Dataset+>.

ORCID iDs

Vittorio Fra  <https://orcid.org/0000-0001-9175-2838>
Evelina Forno  <https://orcid.org/0000-0002-2508-9254>
Riccardo Pignari  <https://orcid.org/0000-0001-6341-8613>
Terrence C Stewart  <https://orcid.org/0000-0002-1445-7613>
Enrico Macii  <https://orcid.org/0000-0001-9046-5618>
Gianvito Urgese  <https://orcid.org/0000-0003-2672-7593>

References

- [1] Gravina R, Alinia P, Ghasemzadeh H and Fortino G 2017 Multi-sensor fusion in body sensor networks: state-of-the-art and research challenges *Inf. Fusion* **35** 68–80
- [2] Kim Y S, Mahmood M, Lee Y, Kim N K, Kwon S, Herbert R, Kim D, Cho H C and Yeo W H 2019 All-in-one, wireless, stretchable hybrid electronics for smart, connected, and ambulatory physiological monitoring *Adv. Sci.* **6** 1900939
- [3] Daneault J-F et al 2021 Accelerometer data collected with a minimum set of wearable sensors from subjects with Parkinson's disease *Sci. Data* **8** 48
- [4] Dami S and Yahaghizadeh M 2021 Predicting cardiovascular events with deep learning approach in the context of the internet of things *Neural Comput. Appl.* **33** 7979–96
- [5] Capela N A, Lemaire E D and Baddour N 2015 Feature selection for wearable smartphone-based human activity recognition with able bodied, elderly, and stroke patients *PLoS One* **10** e0124414
- [6] Allahbakhshi H et al 2020 Using accelerometer and GPS data for real-life physical activity type detection *Sensors* **20** 1–21

- [7] Ceolini E *et al* 2020 Hand-gesture recognition based on EMG and event-based camera sensor fusion: a benchmark in neuromorphic computing *Front. Neurosci.* **14** 1–15
- [8] Forno E *et al* 2020 Techniques for improving localization applications running on low-cost IoT devices 2020 AEIT Int. Conf. of Electrical and Electronic Technologies for Automotive (AEIT AUTOMOTIVE) (IEEE) pp 1–6
- [9] Covi E *et al* 2021 Adaptive extreme edge computing for wearable devices *Front. Neurosci.* **15** 1–27
- [10] Abdullah Al Hafiz Khan Md, Welsh D and Roy N 2018 Firearm detection using wrist worn tri-axis accelerometer signals 2018 IEEE Int. Conf. on Pervasive Computing and Communications Workshops (PerCom Workshops) (IEEE) pp 221–6
<https://ieeexplore.ieee.org/document/8480345/>
- [11] Boudet G *et al* 2019 How to measure sedentary behavior at work? *Front. Public Health* **7** 1–11
- [12] Kumar S, Nehra M, Khurana S, Dilbaghi N, Kumar V, Kaushik A and Kim K-H 2021 Aspects of point-of-care diagnostics for personalized health wellness *Int. J. Nanomed.* **16** 383–402
- [13] Jia G, Zhang G, Yuan X, Gu X, Liu H, Fan Z and Bu L 2021 A synthetical development approach for rehabilitation assistive smart product-service systems: a case study *Adv. Eng. Inf.* **48** 101310
- [14] Dawson L and Akinbi A 2021 Challenges and opportunities for wearable IoT forensics: TomTom spark 3 as a case study *Forensic Sci. Int.: Rep.* **3** 100198
- [15] Frank A E, Kubota A and Riek L D 2019 Wearable activity recognition for robust human-robot teaming in safety-critical environments via hybrid neural networks 2019 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (IEEE) pp 449–54
<https://ieeexplore.ieee.org/document/8968615/>
- [16] Hussain Z, Sheng M and Zhang W E 2019 Different approaches for human activity recognition: a survey *J. Netw. Comput. Appl.* **167** 102738
- [17] Lara O D and Labrador M A 2013 A survey on human activity recognition using wearable sensors *IEEE Commun. Surv. Tutorials* **15** 1192–209
- [18] Ferrari A, Micucci D, Mobilio M and Napolitano P 2021 Trends in human activity recognition using smartphones *J. Reliab. Intell. Environ.* **7** 189–213
- [19] Nweke H F, Teh Y W, Al-garadi M A and Alo U R 2018 Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: state of the art and research challenges *Expert Syst. Appl.* **105** 233–61
- [20] Slim S O *et al* 2019 Survey on human activity recognition based on acceleration data *Int. J. Adv. Comput. Sci. Appl.* **10** 84–98
- [21] Demrozi F, Pravaddelli G, Bihorac A and Rashidi P 2020 Human activity recognition using inertial, physiological and environmental sensors: a comprehensive survey *IEEE Access* **8** 210816–36
- [22] Khan N S and Ghani M S 2021 A survey of deep learning based models for human activity recognition *Wirel. Pers. Commun.* **120** 1593–635
- [23] Maass W 1997 Networks of spiking neurons: the third generation of neural network models *Neural Netw.* **10** 1659–71
- [24] Ranjan J A K, Sigamani T and Barnabas J 2020 A novel and efficient classifier using spiking neural network *J. Supercomput.* **76** 6545–60
- [25] James C D *et al* 2017 A historical survey of algorithms and hardware architectures for neural-inspired and neuromorphic computing applications *Biol. Inspired Cogn. Archit.* **19** 49–64
- [26] Balaji A, Corradi F, Das A, Pande S, Schaafsma S and Catthoor F 2018 Power-accuracy trade-offs for heartbeat classification on neural networks hardware *J. Low Power Electron.* **14** 508–19
- [27] Bohte S M 2004 The evidence for neural information processing with precise spike-times: a survey *Nat. Comput.* **3** 195–206
- [28] Ghosh-Dastidar S and Adeli H 2009 Spiking neural networks *Int. J. Neural Syst.* **19** 295–308
- [29] Maass W 2015 To spike or not to spike: that is the question *Proc. IEEE* **103** 2219–24
- [30] Roy K, Jaiswal A and Panda P 2019 Towards spike-based machine intelligence with neuromorphic computing *Nature* **575** 607–17
- [31] Tavanaei A, Ghodrati M, Kheradpisheh S R, Masquelier T and Maida A 2019 Deep learning in spiking neural networks *Neural Netw.* **111** 47–63
- [32] Indiveri G 2015 Neuromorphic engineering *Springer Handbook of Computational Intelligence* (Berlin: Springer) pp 715–25
http://link.springer.com/10.1007/978-3-662-43505-2%7B%5C_%7D38
- [33] Davies M, Wild A, Orchard G, Sandamirskaya Y, Guerra G A F, Joshi P, Plank P and Risbud S R 2021 Advancing neuromorphic computing with Loihi: a survey of results and outlook *Proc. IEEE* **109** 911–34
- [34] Furber S B, Lester D R, Plana L A, Garside J D, Painkras E, Temple S and Brown A D 2013 Overview of the SpiNNaker system architecture *IEEE Trans. Comput.* **62** 2454–67
- [35] Merolla P A *et al* 2014 A million spiking-neuron integrated circuit with a scalable communication network and interface *Science* **345** 668–73
- [36] Davies M *et al* 2018 Loihi: a neuromorphic manycore processor with on-chip learning *IEEE Micro* **38** 82–99
- [37] Moradi S, Qiao N, Stefanini F and Indiveri G 2018 A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs) *IEEE Trans. Biomed. Circuits Syst.* **12** 106–22
- [38] Frenkel C, Legat J-D and Bol D 2019 A 65 nm 738k-synapse/mm² quad-core binary-weight digital neuromorphic processor with stochastic spike-driven online learning 2019 IEEE Int. Symp. on Circuits and Systems (ISCAS) vol 2019 (IEEE) pp 1–5
<https://ieeexplore.ieee.org/document/8702793/>
- [39] Yang J Q, Wang R, Ren Y, Mao J Y, Wang Z P, Zhou Y and Han S T 2020 Neuromorphic engineering: from biological to spike-based hardware nervous systems *Adv. Mater.* **32** 2003610
- [40] Zhang Y, Qu P and Zheng W 2021 Towards ‘general purpose’ brain-inspired computing system *Tsinghua Sci. Technol.* **26** 664–73
- [41] Stuijt J *et al* 2021 μ Brain: an event-driven and fully synthesizable architecture for spiking neural networks *Front. Neurosci.* **15** 1–15
- [42] Kulkarni S R, Parsa M, Mitchell J P and Schuman C D 2021 Benchmarking the performance of neuromorphic and spiking neural network simulators *Neurocomputing* **447** 145–60
- [43] Eliasmith C and Anderson C H 2003 *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems* (The MIT Press: Cambridge Massachusetts London England) p 356
- [44] Bekolay T *et al* 2014 Nengo: a Python tool for building large-scale functional brain models *Front. Neuroinf.* **7** 1–13
- [45] DeWolf T, Jaworski P and Eliasmith C 2020 Nengo and low-power AI hardware for robust, embedded neurorobotics *Front. Neurobot.* **14** 1–11
- [46] Rasmussen D 2019 NengoDL: combining deep learning and neuromorphic modelling methods *Neuroinformatics* **17** 611–28
- [47] Weiss G M, Yoneda K and Hayajneh T 2019 WISDM smartphone and smartwatch activity and biometrics dataset *UCI Machine Learning Repository: WISDM Smartphone and Smartwatch Activity and Biometrics Dataset Data Set* vol 7 133190–202

- [48] Weiss G M, Yoneda K and Hayajneh T 2019 Smartphone and smartwatch-based biometrics using activities of daily living *IEEE Access* **7** 133190–202
- [49] Ramanujam E, Perumal T and Padmavathi S 2021 Human activity recognition with smartphone and wearable sensors using deep learning techniques: a review *IEEE Sensors J.* **21** 13029–40
- [50] Kendall J D and Kumar S 2020 The building blocks of a brain-inspired computer *Appl. Phys. Rev.* **7** 011305
- [51] Chen Y et al 2016 LSTM networks for mobile human activity recognition *Proc. 2016 Int. Conf. on Artificial Intelligence: Technologies and Applications* (Atlantis Press) <http://atlantispress.com/php/paper-details.php?id=25849464>
- [52] Gumaï A et al 2020 DL-HAR: deep learning-based human activity recognition framework for edge computing *Comput. Mater. Continua* **65** 1033–57
- [53] Agarwal P and Alam M 2020 A lightweight deep learning model for human activity recognition on edge devices *Proc. Comput. Sci.* **167** 2364–73
- [54] Sena J, Barreto J, Caetano C, Cramer G and Schwartz W R 2021 Human activity recognition based on smartphone and wearable sensors using multiscale DCNN ensemble *Neurocomputing* **444** 226–43
- [55] Dua N, Singh S N and Bhaskar Semwal V 2021 Multi-input CNN-GRU based human activity recognition using wearable sensors *Computing* **103** 1461–78
- [56] Kanjilal R and Uysal I 2021 The future of human activity recognition: deep learning or feature engineering? *Neural Process. Lett.* **53** 561–79
- [57] Abdel-Basset M, Hawash H, Chakraborty R K, Ryan M, Elhoseny M and Song H 2021 ST-DeepHAR: deep learning model for human activity recognition in IoHT applications *IEEE Internet Things J.* **8** 4969–79
- [58] Kwapisz J R, Weiss G M and Moore S A 2011 Activity recognition using cell phone accelerometers *ACM SIGKDD Explorations Newsletter* **12** 74–82
- [59] Anguita D et al 2013 A public domain dataset for human activity recognition in free-living conditions *21st European Symp. on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2013* <http://i6doc.com/en/livre/?GCOI=28001100131010>
- [60] Reyes-Ortiz J-L, Oneto L, Samà A, Parra X and Anguita D 2016 Transition-aware human activity recognition using smartphones *Neurocomputing* **171** 754–67
- [61] Reiss A and Stricker D 2012 Introducing a new benchmarked dataset for activity monitoring *2012 16th Int. Symp. on Wearable Computers* (IEEE) pp 108–9 <http://ieeexplore.ieee.org/document/6246152/>
- [62] Reiss A and Stricker D 2012 Creating and benchmarking a new dataset for physical activity monitoring *Proc. 5th Int. Conf. on Pervasive Technologies Related to Assistive Environments—PETRA '12* (ACM) p 1 <http://dl.acm.org/citation.cfm?doid=2413097.2413148>
- [63] Banos O, Garcia R, Holgado-Terriza J A, Damas M, Pomares H, Rojas I, Saez A and Villalonga C 2014 mHealthDroid: a novel framework for agile development of mobile health applications *Proc. 6th Int. Work-Conf. on Ambient Assisted Living an Active Ageing (IWAAL 2014)* pp 91–8
- [64] Banos O, Villalonga C, Garcia R, Saez A, Damas M, Holgado-Terriza J A, Lee S, Pomares H and Rojas I 2015 Design, implementation and validation of a novel open framework for agile development of mobile health applications *Biomed. Eng. Online* **14** S6
- [65] Roggen D et al 2010 Collecting complex activity datasets in highly rich networked sensor environments *2010 7th Int. Conf. on Networked Sensing Systems (INSS)* (IEEE) pp 233–40 <http://ieeexplore.ieee.org/document/5573462/>
- [66] Chavarriaga R, Sagha H, Calatroni A, Digumarti S T, Tröster G, Millán J d R and Roggen D 2013 The opportunity challenge: a benchmark database for on-body sensor-based activity recognition *Pattern Recognit. Lett.* **34** 2033–42
- [67] Ugulino W, Cardador D, Vega K, Velloso E, Milidiú R and Fuks H 2012 Wearable computing: accelerometers' data classification of body postures and movements *Proc. 21st Brazilian Symp. on Artificial Intelligence. Advances in Artificial Intelligence—SBIA 2012* pp 52–61 http://link.springer.com/10.1007/978-3-642-34459-6%7B%5C_%7D6
- [68] Bruno B, Mastrogianni F and Sgorbissa A 2014 A public domain dataset for ADL recognition using wrist-placed accelerometers *The 23rd IEEE Int. Symp. on Robot and Human Interactive Communication* (IEEE) pp 738–43 <http://ieeexplore.ieee.org/document/6926341/>
- [69] Zhang M and Sawchuk A A 2012 USC-HAD: a daily activity dataset for ubiquitous activity recognition using wearable sensors *Proc. 2012 ACM Conf. on Ubiquitous Computing—UbiComp '12* (New York: ACM) p 1036 <http://dl.acm.org/citation.cfm?doid=2370216.2370438>
- [70] Chen C, Jafari R and Kehtarnavaz N 2015 UTD-MHAD: a multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor *2015 IEEE Int. Conf. on Image Processing (ICIP)* vol 2015 (IEEE) pp 168–72 <http://ieeexplore.ieee.org/document/7350781/>
- [71] Mekruksavanich S and Jitpattanakul A 2021 Deep convolutional neural network with RNNs for complex activity recognition using wrist-worn wearable sensor data *Electronics* **10** 1685
- [72] Peppas K et al 2020 Real-time physical activity recognition on smart mobile devices using convolutional neural networks *Appl. Sci.* **10** 1–25
- [73] Ordóñez F and Roggen D 2016 Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition *Sensors* **16** 115
- [74] Wan S, Qi L, Xu X, Tong C and Gu Z 2020 Deep learning models for real-time human activity recognition with smartphones *Mobile Netw. Appl.* **25** 743–55
- [75] Xia K, Huang J and Wang H 2020 LSTM-CNN architecture for human activity recognition *IEEE Access* **8** 56855–66
- [76] Mekruksavanich S and Jitpattanakul A 2020 Smartwatch-based human activity recognition using hybrid LSTM network *2020 IEEE Sensors* vol 2020 (IEEE) pp 1–4 <https://ieeexplore.ieee.org/document/9278630/>
- [77] Mekruksavanich S, Jitpattanakul A, Youplao P and Yupapin P 2020 Enhanced hand-oriented activity recognition based on smartwatch sensor data using LSTMs *Symmetry* **12** 1570
- [78] Oluwalade B et al 2021 Human activity recognition using deep learning models on smartphones and smartwatches sensor data *HEALTHINF 2021—14th Int. Conf. on Health Informatics* (SCITEPRESS—Science and Technology Publications) pp 1–6 <https://scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0010325906450650>
- [79] Ihianle I K, Nwajana A O, Ebeonuwa S H, Otuka R I, Owa K and Orisatoki M O 2020 A deep learning approach for human activities recognition from multimodal sensing devices *IEEE Access* **8** 179028–38
- [80] Blouw P et al 2018 Benchmarking keyword spotting efficiency on neuromorphic hardware (arXiv:1812.01739)

- [81] Yan Y *et al* 2021 Comparing Loihi with a SpiNNaker 2 prototype on low-latency keyword spotting and adaptive robotic control *Neuromorph. Comput. Eng.* **1** 014002
- [82] Buettner K and George A D 2021 Heartbeat classification with spiking neural networks on the Loihi neuromorphic processor 2021 *IEEE Computer Society Annual Symp. on VLSI (ISVLSI)* (IEEE) pp 138–43 <https://ieeexplore.ieee.org/document/9516760/>
- [83] Azghadi M R, Lammie C, Eshraghian J K, Payvand M, Donati E, Linares-Barranco B and Indiveri G 2020 Hardware implementation of deep network accelerators towards healthcare and biomedical applications *IEEE Trans. Biomed. Circuits Syst.* **14** 1138–59
- [84] Peter B and Eliasmith C 2020 Event-driven signal processing with neuromorphic computing systems *ICASSP 2020—2020 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* vol 2020 (IEEE) pp 8534–8 <https://ieeexplore.ieee.org/document/9053043/>
- [85] Forno E, Salvato A, Macii E and Urgese G 2021 PageRank implemented with the MPI paradigm running on a many-core neuromorphic platform *J. Low Power Electron. Appl.* **11** 25
- [86] Urgese G, Barchi F, Parisi E, Forno E, Acquaviva A and Macii E 2019 Benchmarking a many-core neuromorphic platform with an MPI-based DNA sequence matching algorithm *Electronics* **8** 1342
- [87] Kim S *et al* 2019 Survey of IoT platforms supporting artificial intelligence *Proc. Conf. on Research in Adaptive and Convergent Systems* (New York: ACM) pp 65–6 <https://dl.acm.org/doi/10.1145/3338840.3355694>
- [88] An H, Ha D S and Yi Y C 2020 Powering next-generation industry 4.0 by a self-learning and low-power neuromorphic system *Proc. 7th ACM Int. Conf. on Nanoscale Computing and Communication* (New York: ACM) pp 1–6 <https://dl.acm.org/doi/10.1145/3411295.3411302>
- [89] Chang Z, Liu S, Xiong X, Cai Z and Tu G 2021 A survey of recent advances in edge-computing-powered artificial intelligence of things *IEEE Internet Things J.* **8** 13849–75
- [90] Eichenbaum H 2014 Time cells in the hippocampus: a new dimension for mapping memories *Nat. Rev. Neurosci.* **15** 732–44
- [91] Voelker A R and Eliasmith C 2018 Improving spiking dynamical networks: accurate delays, higher-order synapses, and time cells *Neural Comput.* **30** 569–609
- [92] Voelker A R, Kajić I and Eliasmith C 2019 Legendre memory units: continuous-time representation in recurrent neural networks *Advances in Neural Information Processing Systems NeurIPS* vol 32
- [93] Voelker A R and Eliasmith C 2021 Programming neuromorphics using the neural engineering framework *Handbook of Neuroengineering* (Berlin: Springer) pp 1–43 https://link.springer.com/10.1007/978-981-15-2848-4%7B%5C_%7D115-1
- [94] Blouw P *et al* 2020 Hardware aware training for efficient keyword spotting on general purpose and specialized hardware (arXiv:2009.04465v3)
- [95] Zhang J and Tao D 2021 Empowering things with intelligence: a survey of the progress, challenges, and opportunities in artificial intelligence of things *IEEE Internet Things J.* **8** 7789–817
- [96] Suto J 2021 The effect of hyperparameter search on artificial neural network in human activity recognition *Open Comput. Sci.* **11** 411–22
- [97] <https://github.com/microsoft/nni>
- [98] Forno E *et al* 2018 A parallel hardware architecture for quantum annealing algorithm acceleration 2018 *IFIP/IEEE Int. Conf. on Very Large Scale Integration (VLSI-SoC)* vol 2018 (IEEE) pp 31–6 <https://ieeexplore.ieee.org/document/8644777/>