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

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

Vittorio Fra[1], Evelina Forno[1], Riccardo Pignari[1], Terrence C. Stewart[2], Enrico Macii[1], Gianvito Urgese[1]

[1]Politecnico di Torino, 10138, Torino (TO), ITALY
[2] National Research Council Canada, University of Waterloo Collaboration Centre, Waterloo, Ontario, CANADA

E-mail: gianvito.urgese@polito.it

January 2022

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 (POC) 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 (DNN) architectures as well as on recent alternatives like the Legendre Memory Unit (LMU), we show how spiking neural networks (SNNs) 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 (AIoT) applications.

Keywords: Neuromorphic Engineering, Human Activity Recognition, Spiking Neural Network, Artificial Intelligence of Things, Edge Computing, Legendre Memory Unit, Raw IMU Data.

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, 2, 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, 5, 6, 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, 11, 12, 13, 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 (ADL), can also be decisive in safety-critical situations [15]. As a general definition, human activity recognition 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, 20, 21, 22].

In this race to the best classification result, a vast majority of the proposed solutions have focused 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, 28, 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, 35, 36, 37, 38, 39, 40, 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 human activity recognition 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 human activity recognition 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 focused 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 human activity recognition tasks based on, or intended for, portable platforms [51, 52, 53, 54, 55, 56, 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 human activity recognition 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 coworkers [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 minutes 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 gyrospcope of the "general, hand-oriented" subset of the

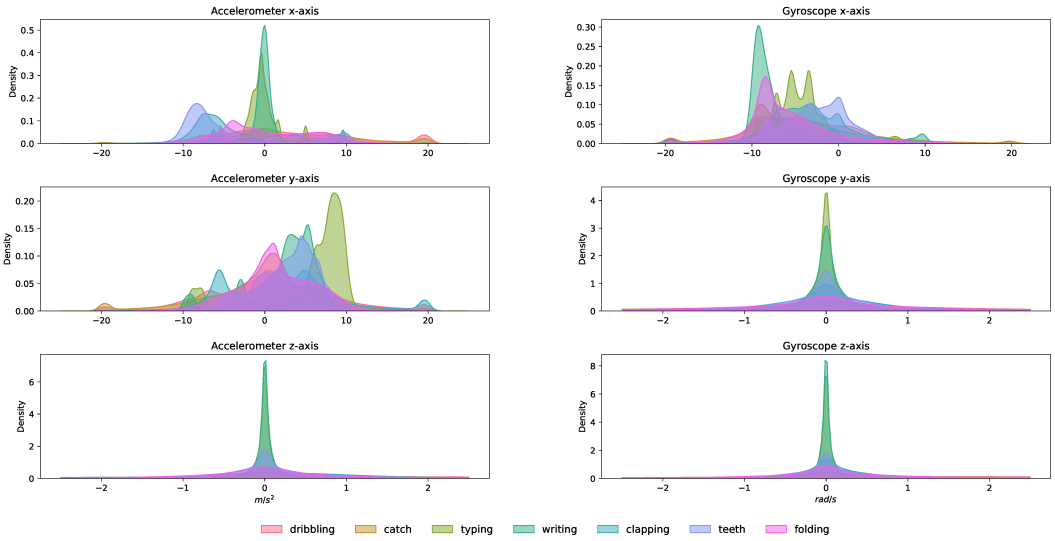

Figure 1. Kernel density estimation of the values recorded from smartwatch on the 6 IMU sensor axes for the 7 classes in the "general, hand-oriented" subset of the WISDM dataset.

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 coworkers 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 DNNs.

In [85] and [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 [90] presenting the neuromorphic IC μBrain.

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 [91, 92]. The *Legendre Memory Unit (LMU)* is a recurrent architecture able to perform such approximation for a continuous-time delay [93]. 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 [94]. The i^{th} shifted Legendre polynomial is given by *Eq. 1*

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

and it is used to delay the input signal through *Eq. 2*

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

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 *Eq. 3*

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

with $\mathbf{A}$ and $\mathbf{B}$ representing the ideal state-space matrices derived using the Padé approximants through *Eq. 4* and *Eq. 5*.

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

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

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 [95].

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 [96]. 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, 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 human activity recognition (f). In the following subsections we provide more details about each step of our investigation procedure.

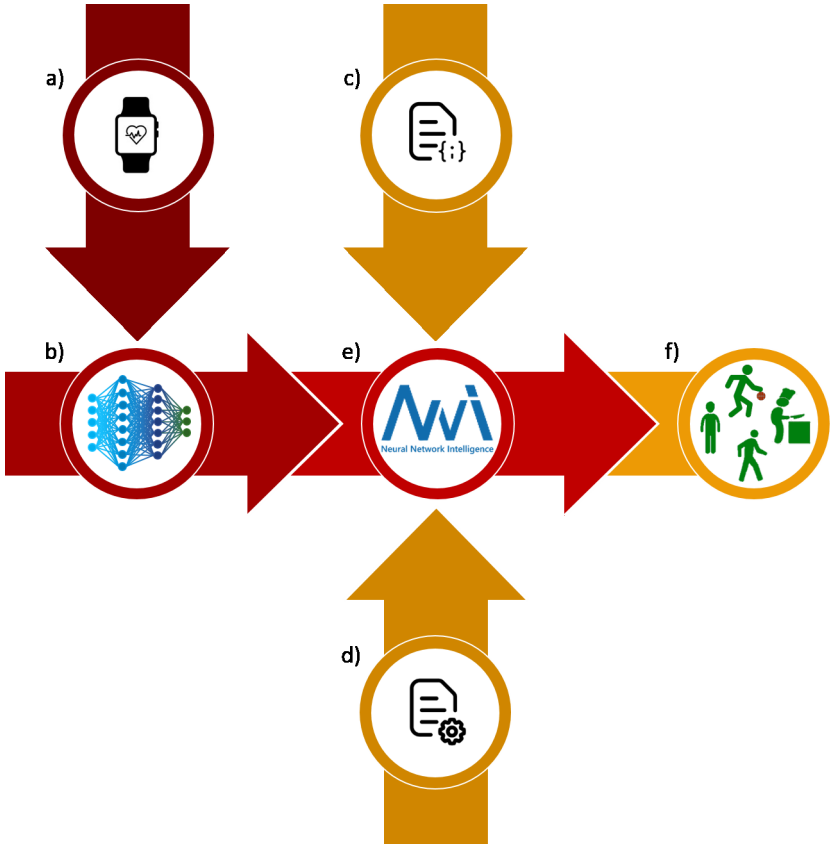

Figure 2. The procedure we adopted can be divided into two complementary phases. On the one hand, depicted by the vertical arrows, there are the preliminary steps: dataset selection in a), hyperparameters search space definition in c) and optimization experiment configuration in d). On the other hand, there is the backbone of the pipeline, summarized by the horizontal arrows: neural network architectures selection in b), hyperparameters optimization in e) and classifiers evaluation in f).

3.1. Activity subset and time window

Human activity recognition, 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 (POC) 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

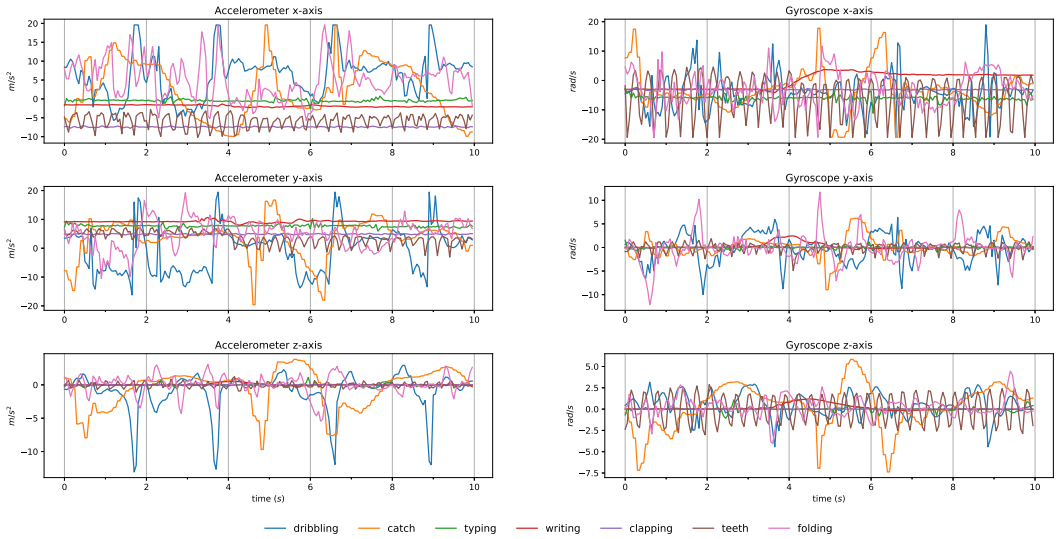

Figure 3. A comparison of representative 10-second samples recorded by the smartwatch on the 6 IMU sensors for the 7 classes in the "general, hand-oriented" subset of the WISDM dataset.

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 Legendre Memory Unit (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 (*Figure 4.e* and *Figure 4.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

Artificial neural networks 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 [97], 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 [98], using the built-in Annealing algorithm. For each network, we designed a NNI optimization experiment, carried out within a proper search space defined at step **c**. Each optimization experiment is composed of 1,000 trials, with 4 evenly spaced random re-initializations of the tuner, intended to partially mitigate the problem of local minima affecting the annealing algorithms [99]. 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

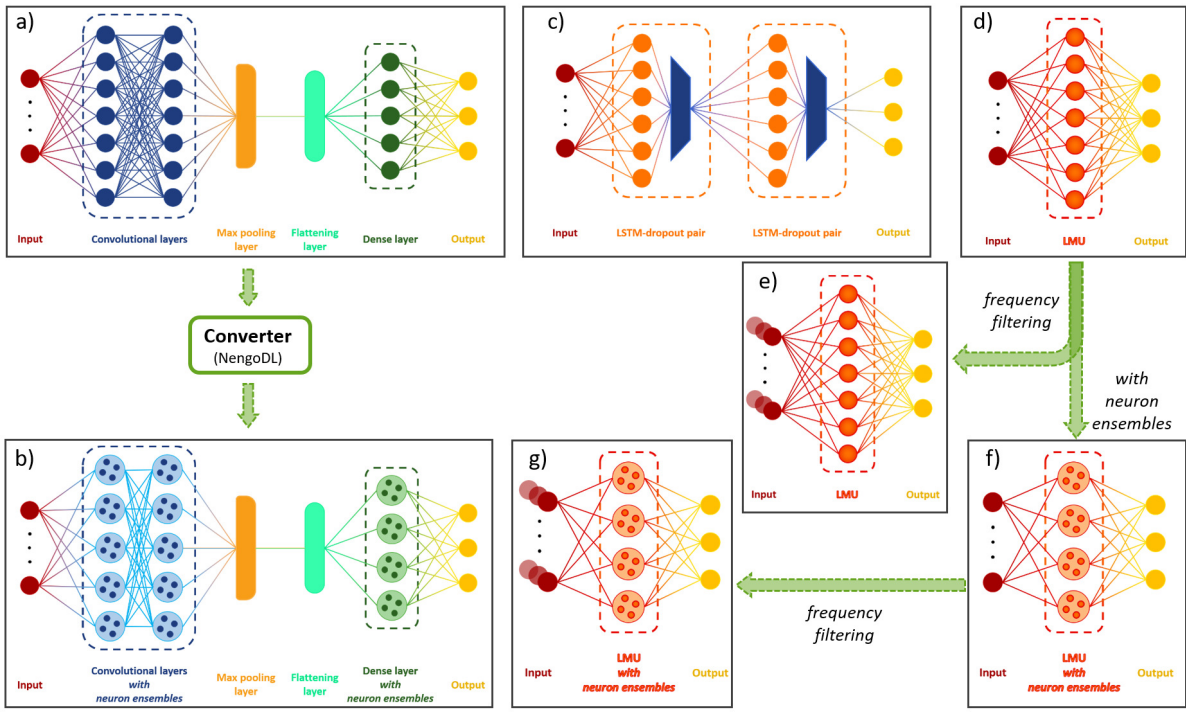

Figure 4. Summary of the investigated networks. The convolutional architecture adopted in the non-spiking domain (a) has been translated into the spiking domain (b) by means of the Converter in NengoDL. The recurrent architectures have instead different structures in the two domains: LSTM units followed by a dropout layer have been employed for the non-spiking implementation (c), while the recurrent SNN has been obtained using the Legendre Memory Unit (f). This latter has been adopted in the non-spiking domain (d) as first. An additional variation with the introduction of a frequency filtering on the input has been explored for both the non-spiking (e) and the spiking (g) LMU-based architectures.

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 focused 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 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

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
	n_steps	How long (in simulation time steps *) the input is presented to the network
LMU	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 **	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

* The default value in Nengo of 1 ms is used
** All the hyperparameters for the non-spiking LMU are specifically re-optimized for the spiking implementation

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 spiking neural networks 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 1,000 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 *Table S1* and *Table S2* for the non-spiking and the spiking networks respectively.

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 hardwares: 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

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	2,125,222	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 ($\times 10^3$)	4,249.65	2,828.89	/	158.66	197.00	/	/
SOPs ($\times 10^3$)	/	/	10.82	/	/	99.91	127.95
Energy on Movidius (μJ)	3,199.99	2,130.15	/	119.47	148.34	/	/
Energy on Loihi (μJ)	/	/	5.49	/	/	50.66	64.87

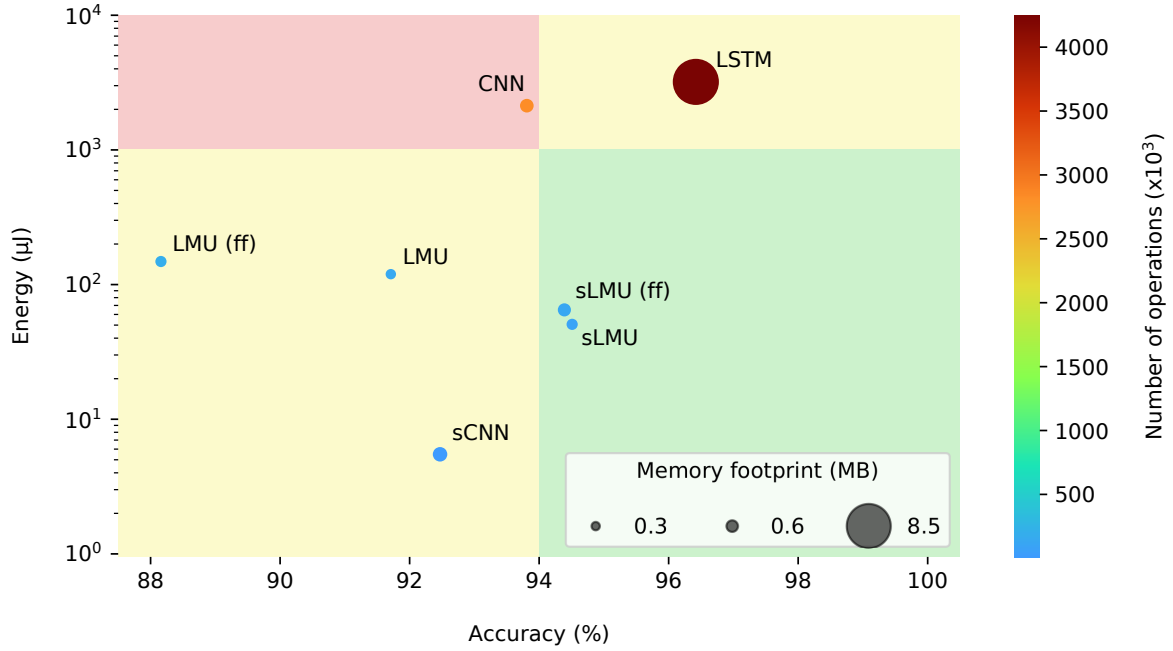

Figure 5. In order to highlight the advantages of a neuromorphic approach to temporal signal classification, an *energy vs. accuracy* diagram can be employed to effectively show how relevant is the gain in terms of energy reduction with respect to a possible drop in classification accuracy. The results here presented highlight that all the investigated SNNs and all the LMU-based networks ensure an energy consumption one order of magnitude at least smaller than that of traditional DNNs. Concerning the memory footprint, similar conclusion can be drawn, with CNN and LSTM turning out to be largest networks. From the accuracy standpoint, instead, spiking LMUs provide performances comparable to both CNN and LSTM, even overcoming the former.

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 color. 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 μJ. 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 μJ, 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 μJ 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

Human activity recognition is a time-dependent task whose application domain extends to all the aspects of human life, from healthcare 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

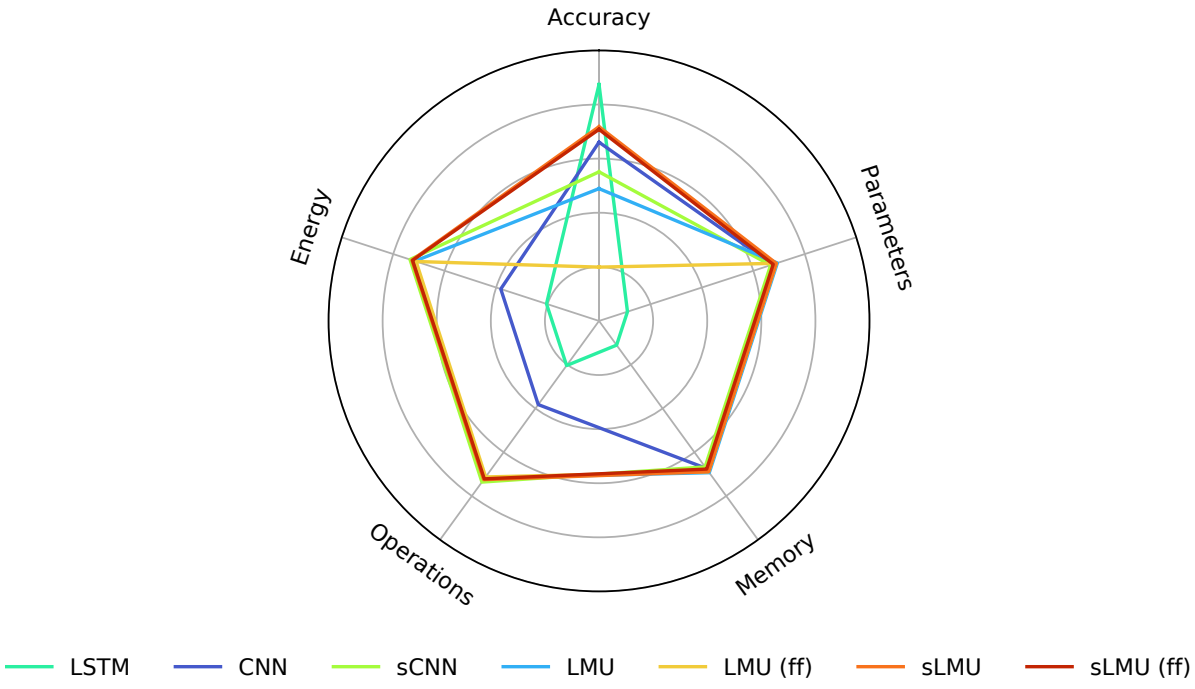

Figure 6. A radar chart of the presented results allows to easily compare the investigated networks focusing on each of the evaluated metrics. It is worth noting that the traditional DNN architectures here considered, namely the LSTM and the CNN, are outperformed by the alternative ones, based on the LMU, in all the metrics related to energy and memory. Similarly, the spiking CNN also provides significant improvements with respect to LSTM and non-spiking CNN in terms of both energy and memory.

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.

Conflicts of interest

The authors declare no conflict of interest.

Data availability

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+`

References

[1] Raffaele Gravina et al. "Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges". In: *Information Fusion* 35 (May 2017), pp. 68–80. ISSN: 15662535. DOI: 10.1016/j.inffus.2016.09.005. URL: http://dx.doi.org/10.1016/j.inffus.2016.09.005%20https://linkinghub.elsevier.com/retrieve/pii/S156625351630077X.

[2] Yun-Soung Kim et al. "All-in-One, Wireless, Stretchable Hybrid Electronics for Smart, Connected, and Ambulatory Physiological Monitoring". In: *Advanced Science* 6.17 (Sept. 2019), p. 1900939. ISSN: 2198-3844. DOI: 10.1002/advs.201900939. URL: https://onlinelibrary.wiley.com/doi/10.1002/advs.201900939.

[3] Jean-Francois Daneault et al. "Accelerometer data collected with a minimum set of wearable sensors from subjects with Parkinson's disease". In: *Scientific Data* 8.1 (Dec. 2021), p. 48. ISSN: 2052-4463. DOI: 10.1038/s41597-021-00830-0. URL: http://www.nature.com/articles/s41597-021-00830-0.

[4] Sina Dami and Mahtab Yahaghizadeh. "Predicting cardiovascular events with deep learning approach in the context of the internet of things". In: *Neural Computing and Applications* 33.13 (July 2021), pp. 7979–7996. ISSN: 0941-0643. DOI: 10.1007/s00521-020-05542-x. URL: https://doi.org/10.1007/s00521-020-05542-x%20https://link.springer.com/10.1007/s00521-020-05542-x.

[5] Nicole A Capela, Edward D Lemaire, and Natalie Baddour. "Feature Selection for Wearable Smartphone-Based Human Activity Recognition with Able bodied, Elderly, and Stroke Patients". In: *PLOS ONE* 10.4 (Apr. 2015). Ed. by Christian Lovis, e0124414. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0124414. URL: www.nserc-crsng.gc.ca%20https://dx.plos.org/10.1371/journal.pone.0124414.

[6] Hoda Allahbakhshi et al. "Using accelerometer and GPS data for real-life physical activity type detection". In: *Sensors (Switzerland)* 20.3 (2020), pp. 1–21. ISSN: 14248220. DOI: 10.3390/s20030588.

[7] Enea Ceolini et al. "Hand-Gesture Recognition Based on EMG and Event-Based Camera Sensor Fusion: A Benchmark in Neuromorphic Computing". In: *Frontiers in Neuroscience* 14.August (Aug. 2020), pp. 1–15. ISSN: 1662-453X. DOI: 10.3389/fnins.2020.00637. URL: https://www.frontiersin.org/article/10.3389/fnins.2020.00637/full.

[8] Evelina Forno et al. "Techniques for improving localization applications running on low-cost IoT devices". In: *2020 AEIT International Conference of Electrical and Electronic Technologies for Automotive (AEIT AUTOMOTIVE)*. IEEE. 2020, pp. 1–6.

[9] Erika Covi et al. "Adaptive Extreme Edge Computing for Wearable Devices". In: *Frontiers in Neuroscience* 15.May (May 2021), pp. 1–27. ISSN: 1662-453X. DOI: 10.3389/fnins.2021.611300. arXiv: 2012.14937. URL: https://www.frontiersin.org/articles/10.3389/fnins.2021.611300/full.

[10] Md Abdullah Al Hafiz Khan, David Welsh, and Nirmalya Roy. "Firearm Detection Using Wrist Worn Tri-Axis Accelerometer Signals". In: *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, Mar. 2018, pp. 221–226. ISBN: 978-1-5386-3227-7. DOI: 10.1109/PERCOMW.2018.8480345. URL: https://ieeexplore.ieee.org/document/8480345/.

[11] Gil Boudet et al. "How to Measure Sedentary Behavior at Work?" In: *Frontiers in Public Health* 7.JUL (July 2019), pp. 1–11. ISSN: 2296-2565. DOI: 10.3389/fpubh.2019.00167. URL: https://www.frontiersin.org/article/10.3389/fpubh.2019.00167/full.

[12] Sandeep Kumar et al. "Aspects of Point-of-Care Diagnostics for Personalized Health Wellness". In: *International Journal of Nanomedicine* Volume 16 (Jan. 2021), pp. 383–402. ISSN: 1178-2013. DOI: 10.2147/IJN.S267212. URL: https://www.dovepress.com/aspects-of-point-of-care-diagnostics-for-personalized-health-wellness-peer-reviewed-article-IJN.

[13] Guo Jia et al. "A synthetical development approach for rehabilitation assistive smart product-service systems: A case study". In: *Advanced Engineering Informatics* 48.May (Apr. 2021), p. 101310. ISSN: 14740346. DOI: 10.1016/j.aei.2021.101310. URL: https://doi.org/10.1016/j.aei.2021.101310%20https://linkinghub.elsevier.com/retrieve/pii/S1474034621000641.

[14] Liam Dawson and Alex Akinbi. "Challenges and opportunities for wearable IoT forensics: TomTom Spark 3 as a case study". In: *Forensic Science International: Reports* 3.April (July 2021), p. 100198. ISSN: 26659107. DOI: 10.1016/j.fsir.2021.100198. URL: https://doi.org/10.1016/j.fsir.2021.100198%20https://linkinghub.elsevier.com/retrieve/pii/S2665910721000293.

[15] Andrea E. Frank, Alyssa Kubota, and Laurel D. Riek. "Wearable activity recognition for robust human-robot teaming in safety-critical environments via hybrid neural networks". In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Nov. 2019, pp. 449–454. ISBN: 978-1-7281-4004-9. DOI: 10.1109/IROS40897.2019.8968615. URL: https://ieeexplore.ieee.org/document/8968615/.

[16] Zawar Hussain, Michael Sheng, and Wei Emma Zhang. "Different Approaches for Human Activity Recognition: A Survey". In: *Journal of Network and Computer Applications* 167 (June 2019), p. 102738. ISSN: 10848045. arXiv: 1906.05074. URL: http://arxiv.org/abs/1906.05074%7B%5C%7D0Ahttp://dx.doi.org/10.1016/j.jnca.2020.102738%20https://linkinghub.elsevier.com/retrieve/pii/S1084804520302125%20http://arxiv.org/abs/1906.05074%20http://dx.doi.org/10.1016/j.jnca.2020.102738.

[17] Oscar D. Lara and Miguel A. Labrador. "A Survey on Human Activity Recognition using Wearable Sensors". In: *IEEE Communications Surveys & Tutorials* 15.3 (2013), pp. 1192–1209. ISSN: 1553-877X. DOI: 10.1109/SURV.2012.110112.00192. URL: http://ieeexplore.ieee.org/document/6365160/.

[18] Anna Ferrari et al. "Trends in human activity recognition using smartphones". In: *Journal of Reliable Intelligent Environments* 7.3 (Sept. 2021), pp. 189–213. ISSN: 2199-4668. DOI: 10.1007/s40860-021-00147-0. URL: https://doi.org/10.1007/s40860-021-00147-0%20https://link.springer.com/10.1007/s40860-021-00147-0.

[19] Henry Friday Nweke et al. "Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges". In: *Expert Systems with Applications* 105 (Sept. 2018), pp. 233–261. ISSN: 09574174. DOI: 10.1016/j.eswa.2018.03.056. URL: https://doi.org/10.1016/j.eswa.2018.03.056%20https://linkinghub.elsevier.com/retrieve/pii/S0957417418302136.

[20] Salwa O. Slim et al. "Survey on Human Activity Recognition based on Acceleration Data". In: *International Journal of Advanced Computer Science and Applications* 10.3 (2019), pp. 84–98. ISSN: 21565570. DOI: 10.14569/IJACSA.2019.0100311. URL: http://thesai.org/Publications/ViewPaper?Volume=10%7B%5C&%7DIssue=3%7B%5C&%7DCode=ijacsa%7B%5C&%7DSerialNo=11.

[21] Florenc Demrozi et al. "Human Activity Recognition Using Inertial, Physiological and Environmental Sensors: A Comprehensive Survey". In: *IEEE Access* 8.Dl (2020), pp. 210816–210836. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.3037715. arXiv: 2004.08821. URL: https://ieeexplore.ieee.org/document/9257355/.

[22] Nida Saddaf Khan and Muhammad Sayeed Ghani. "A Survey of Deep Learning Based Models for Human Activity Recognition". In: *Wireless Personal Communications* 0123456789 (May 2021). ISSN: 0929-6212. DOI: 10.1007/s11277-021-08525-w. URL: https://doi.org/10.1007/s11277-021-08525-w%20https://link.springer.com/10.1007/s11277-021-08525-w.

[23] Wolfgang Maass. "Networks of spiking neurons: The third generation of neural network models". In: *Neural Networks* 10.9 (Dec. 1997), pp. 1659–1671. ISSN: 08936080. DOI: 10.1016/S0893-6080(97)00011-7. URL: https://linkinghub.elsevier.com/retrieve/pii/S0893608097000117.

[24] Joshua Arul Kumar Ranjan, Titus Sigamani, and Janet Barnabas. "A novel and efficient classifier using spiking neural network". In: *The Journal of Supercomputing* 76.9 (Sept. 2020), pp. 6545–6560. ISSN: 0920-8542. DOI: 10.1007/s11227-019-02881-y. URL: https://doi.org/10.1007/s11227-019-02881-y%20http://link.springer.com/10.1007/s11227-019-02881-y.

[25] Conrad D. James et al. "A historical survey of algorithms and hardware architectures for neural-inspired and neuromorphic computing applications". In: *Biologically Inspired Cognitive Architectures* 19 (Jan. 2017), pp. 49–64. ISSN: 2212683X. DOI: 10.1016/j.bica.2016.11.002. URL: http://dx.doi.org/10.1016/j.bica.2016.11.002%20https://linkinghub.elsevier.com/retrieve/pii/S2212683X16300561.

[26] Adarsha Balaji et al. "Power-Accuracy Trade-Offs for Heartbeat Classification on Neural Networks Hardware". In: *Journal of Low Power Electronics* 14.4 (Dec. 2018), pp. 508–519. ISSN: 1546-1998. DOI: 10.1166/jolpe.2018.1582. URL: https://www.ingentaconnect.com/content/10.1166/jolpe.2018.1582.

[27] Sander M. Bohte. "The evidence for neural information processing with precise spike-times: A survey". In: *Natural Computing* 3.2 (2004), pp. 195–206. ISSN: 1567-7818. DOI: 10.1023/B:NACO.0000027755.02868.60. URL: http://link.springer.com/10.1023/B:NACO.0000027755.02868.60.

[28] Samanwoy Ghosh-Dastidar and Hojjat Adeli. "Spiking Neural Networks". In: *International Journal of Neural Systems* 19.04 (Aug. 2009), pp. 295–308. ISSN: 0129-0657. DOI: 10.1142/S0129065709002002. URL: https://www.worldscientific.com/doi/abs/10.1142/S0129065709002002.

[29] Wolfgang Maass. "To Spike or Not to Spike: That Is the Question". In: *Proceedings of the IEEE* 103.12 (Dec. 2015), pp. 2219–2224. ISSN: 0018-9219. DOI: 10.1109/JPROC.2015.2496679. URL: http://ieeexplore.ieee.org/document/7329679/.

[30] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. "Towards spike-based machine intelligence with neuromorphic computing". In: *Nature* 575.7784 (Nov. 2019), pp. 607–617. ISSN: 0028-0836. DOI: 10.1038/s41586-019-1677-2. URL: http://dx.doi.org/10.1038/s41586-019-1677-2%20http://www.nature.com/articles/s41586-019-1677-2.

[31] Amirhossein Tavanaei et al. "Deep learning in spiking neural networks". In: *Neural Networks* 111 (Mar. 2019), pp. 47–63. ISSN: 08936080. DOI: 10.1016/j.neunet.2018.12.002. arXiv: 1804.08150. URL: https://doi.org/10.1016/j.neunet.2018.12.002%20https://linkinghub.elsevier.com/retrieve/pii/S0893608018303332.

[32] Giacomo Indiveri. "Neuromorphic Engineering". In: *Springer Handbook of Computational Intelligence*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 715–725. DOI: 10.1007/978-3-662-43505-2_38. URL: http://link.springer.com/10.1007/978-3-662-43505-2%7B%5C_%7D38.

[33] Mike Davies et al. "Advancing Neuromorphic Computing with Loihi: A Survey of Results and Outlook". In: *Proceedings of the IEEE* 109.5 (2021), pp. 911–934. ISSN: 15582256. DOI: 10.1109/JPROC.2021.3067593.

[34] Steve B. Furber et al. "Overview of the SpiNNaker System Architecture". In: *IEEE Transactions on Computers* 62.12 (Dec. 2013), pp. 2454–2467. ISSN: 0018-9340. DOI: 10.1109/TC.2012.142. URL: http://ieeexplore.ieee.org/document/6226357/.

[35] Paul A. Merolla et al. "A million spiking-neuron integrated circuit with a scalable communication network and interface". In: *Science* 345.6197 (Aug. 2014), pp. 668–673. ISSN: 0036-8075. DOI: 10.1126/science.1254642. URL: https://www.sciencemag.org/lookup/doi/10.1126/science.1254642.

[36] Mike Davies et al. "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning". In: *IEEE Micro* 38.1 (Jan. 2018), pp. 82–99. ISSN: 0272-1732. DOI: 10.1109/MM.2018.112130359. URL: http://ieeexplore.ieee.org/document/8259423/.

[37] Saber Moradi et al. "A Scalable Multicore Architecture With Heterogeneous Memory Structures for Dynamic Neuromorphic Asynchronous Processors (DYNAPs)". In: *IEEE Transactions on Biomedical Circuits and Systems* 12.1 (Feb. 2018), pp. 106–122. ISSN: 1932-4545. DOI: 10.1109/TBCAS.2017.2759700. arXiv: 1708.04198. URL: https://ieeexplore.ieee.org/document/8094868/.

[38] Charlotte Frenkel, Jean-Didier Legat, and David Bol. "A 65-nm 738k-Synapse/mm2 Quad-Core Binary-Weight Digital Neuromorphic Processor with Stochastic Spike-Driven Online Learning". In: *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*. Vol. 2019-May. 5. IEEE, May 2019, pp. 1–5. ISBN: 978-1-7281-0397-6. DOI: 10.1109/ISCAS.2019.8702793. URL: https://ieeexplore.ieee.org/document/8702793/.

[39] Jia-Qin Yang et al. "Neuromorphic Engineering: From Biological to Spike-Based Hardware Nervous Systems". In: *Advanced Materials* 32.52 (Dec. 2020), p. 2003610. ISSN: 0935-9648. DOI: 10.1002/adma.202003610. URL: https://onlinelibrary.wiley.com/doi/10.1002/adma.202003610.

[40] Youhui Zhang, Peng Qu, and Weimin Zheng. "Towards 'general purpose' brain-inspired computing system". In: *Tsinghua Science and Technology* 26.5 (2021), pp. 664–673. ISSN: 18787606. DOI: 10.26599/TST.2021.9010010.

[41] Jan Stuijt et al. "μBrain: An Event-Driven and Fully Synthesizable Architecture for Spiking Neural Networks". In: *Frontiers in Neuroscience* 15.May (May 2021), pp. 1–15. ISSN: 1662-453X. DOI: 10.3389/fnins.2021.664208. URL: https://www.frontiersin.org/articles/10.3389/fnins.2021.664208/full.

[42] Shruti R. Kulkarni et al. "Benchmarking the performance of neuromorphic and spiking neural network simulators". In: *Neurocomputing* 447 (Aug. 2021), pp. 145–160. ISSN: 09252312. DOI: 10.1016/j.neucom.2021.03.028. URL: https://doi.org/10.1016/j.neucom.2021.03.028%20https://linkinghub.elsevier.com/retrieve/pii/S0925231221003969.

[43] Chris Eliasmith and Charles H. Anderson. "Neural engineering : computation, representation, and dynamics in neurobiological systems". In: (2003), p. 356.

[44] Trevor Bekolay et al. "Nengo: A Python tool for building large-scale functional brain models". In: *Frontiers in Neuroinformatics* 7.JAN (2014), pp. 1–13. ISSN: 16625196. DOI: 10.3389/fninf.2013.00048.

[45] Travis DeWolf, Pawel Jaworski, and Chris Eliasmith. "Nengo and Low-Power AI Hardware for Robust, Embedded Neurorobotics". In: *Frontiers in Neurorobotics* 14.October (Oct. 2020), pp. 1–11. ISSN: 1662-5218. DOI: 10.3389/fnbot.2020.568359. URL: https://www.frontiersin.org/article/10.3389/fnbot.2020.568359/full.

[46] Daniel Rasmussen. "NengoDL: Combining Deep Learning and Neuromorphic Modelling Methods". In: *Neuroinformatics* 17.4 (Oct. 2019), pp. 611–628. ISSN: 1539-2791. DOI: 10.1007/s12021-019-09424-z. URL: http://link.springer.com/10.1007/s12021-019-09424-z.

[47] Gary M. Weiss. "WISDM Smartphone and Smartwatch Activity and Biometrics Dataset". In: *UCI Machine Learning Repository: WISDM Smartphone and Smartwatch Activity and Biometrics Dataset Data Set* 7 (2019), pp. 133190–133202.

[48] Gary M. Weiss, Kenichi Yoneda, and Thaier Hayajneh. "Smartphone and Smartwatch-Based Biometrics Using Activities of Daily Living". In: *IEEE Access* 7 (2019), pp. 133190–133202. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2019.2940729. URL: https://ieeexplore.ieee.org/document/8835065/.

[49] E. Ramanujam, Thinagaran Perumal, and S. Padmavathi. "Human Activity Recognition with Smartphone and Wearable Sensors using Deep Learning Techniques: A Review". In: *IEEE Sensors Journal* 21.12 (2021), pp. 13029–13040. ISSN: 15581748. DOI: 10.1109/JSEN.2021.3069927.

[50] Jack D. Kendall and Suhas Kumar. "The building blocks of a brain-inspired computer". In: *Applied Physics Reviews* 7.1 (Mar. 2020), p. 011305. ISSN: 1931-9401. DOI: 10.1063/1.5129306. URL: http://aip.scitation.org/doi/10.1063/1.5129306.

[51] Yuwen Chen et al. "LSTM Networks for Mobile Human Activity Recognition". In: *Proceedings of the 2016 International Conference on Artificial Intelligence: Technologies and Applications*. Atlantis Press, 2016. ISBN: 978-94-6252-162-9. DOI: 10.2991/icaita-16.2016.13. URL: http://www.atlantis-press.com/php/paper-details.php?id=25849464.

[52] Abdu Gumaei et al. "DL-HAR: Deep Learning-Based Human Activity Recognition Framework for Edge Computing". In: *Computers, Materials & Continua* 65.2 (2020), pp. 1033–1057. ISSN: 1546-2226. DOI: 10.32604/cmc.2020.011740. URL: http://www.techscience.com/cmc/v65n2/39861.

[53] Preeti Agarwal and Mansaf Alam. "A Lightweight Deep Learning Model for Human Activity Recognition on Edge Devices". In: *Procedia Computer Science* 167.2019 (2020), pp. 2364–2373. ISSN: 18770509. DOI: 10.1016/j.procs.2020.03.289. arXiv: 1909.12917. URL: https://doi.org/10.1016/j.procs.2020.03.289%20https://linkinghub.elsevier.com/retrieve/pii/S1877050920307559.

[54] Jessica Sena et al. "Human activity recognition based on smartphone and wearable sensors using multiscale DCNN ensemble". In: *Neurocomputing* 444 (July 2021), pp. 226–243. ISSN: 09252312. DOI: 10.1016/j.neucom.2020.04.151. URL: https://doi.org/10.1016/j.neucom.2020.04.151%20https://linkinghub.elsevier.com/retrieve/pii/S0925231220317823.

[55] Nidhi Dua, Shiva Nand Singh, and Vijay Bhaskar Semwal. "Multi-input CNN-GRU based human activity recognition using wearable sensors". In: *Computing* 103.7 (July 2021), pp. 1461–1478. ISSN: 0010-485X. DOI: 10.1007/s00607-021-00928-8. URL: https://doi.org/10.1007/s00607-021-00928-8%20https://link.springer.com/10.1007/s00607-021-00928-8.

[56] Ria Kanjilal and Ismail Uysal. "The Future of Human Activity Recognition: Deep Learning or Feature Engineering?" In: *Neural Processing Letters* 53.1 (Feb. 2021), pp. 561–579. ISSN: 1370-4621. DOI: 10.1007/s11063-020-10400-x. URL: https://doi.org/10.1007/s11063-020-10400-x%20http://link.springer.com/10.1007/s11063-020-10400-x.

[57] Mohamed Abdel-Basset et al. "ST-DeepHAR: Deep Learning Model for Human Activity Recognition in IoHT Applications". In: *IEEE Internet of Things Journal* 8.6 (Mar. 2021), pp. 4969–4979. ISSN: 2327-4662. DOI: 10.1109/JIOT.2020.3033430. URL: https://ieeexplore.ieee.org/document/9238036/.

[58] Jennifer R. Kwapisz, Gary M. Weiss, and Samuel A. Moore. "Activity recognition using cell phone accelerometers". In: *ACM SIGKDD Explorations Newsletter* 12.2 (Mar. 2011), pp. 74–82. ISSN: 1931-0145. DOI: 10.1145/1964897.1964918. URL: https://dl.acm.org/doi/10.1145/1964897.1964918.

[59] Davide Anguita et al. "A Public Domain Dataset for Human Activity Recognition in Free-Living Conditions". In: *21th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2013*. April. Aug. 2013. ISBN: ISBN 978-2-87419-081-0. URL: http://www.i6doc.com/en/livre/?GCOI=28001100131010.

[60] Jorge-L. Reyes-Ortiz et al. "Transition-Aware Human Activity Recognition Using Smartphones". In: *Neurocomputing* 171 (Jan. 2016), pp. 754–767. ISSN: 09252312. DOI: 10.1016/j.neucom.2015.07.085. URL: http://dx.doi.org/10.1016/j.neucom.2015.07.085%20https://linkinghub.elsevier.com/retrieve/pii/S0925231215010930.

[61] Attila Reiss and Didier Stricker. "Introducing a New Benchmarked Dataset for Activity Monitoring". In: *2012 16th International Symposium on Wearable Computers*. IEEE, June 2012, pp. 108–109. ISBN: 978-0-7695-4697-1. DOI: 10.1109/ISWC.2012.13. URL: http://ieeexplore.ieee.org/document/6246152/.

[62] Attila Reiss and Didier Stricker. "Creating and benchmarking a new dataset for physical activity monitoring". In: *Proceedings of the 5th International Conference on PErvasive Technologies Related to Assistive Environments - PETRA '12*. February. ACM Press, 2012, p. 1. ISBN: 9781450313001. DOI: 10.1145/2413097.2413148. URL: http://dl.acm.org/citation.cfm?doid=2413097.2413148.

[63] Oresti Banos et al. "mHealthDroid: a novel framework for agile development of mobile health applications". In: *Proceedings of the 6th International Work-conference on Ambient Assisted Living an Active Ageing (IWAAL 2014)*. 2014, pp. 91–98.

[64] Oresti Banos et al. "Design, implementation and validation of a novel open framework for agile development of mobile health applications". In: *BioMedical Engineering OnLine* 14.Suppl 2 (2015), S6. ISSN: 1475-925X. DOI: 10.1186/1475-925X-14-S2-S6. URL: http://biomedical-engineering-online.biomedcentral.com/articles/10.1186/1475-925X-14-S2-S6.

[65] Daniel Roggen et al. "Collecting complex activity datasets in highly rich networked sensor environments". In: *2010 Seventh International Conference on Networked Sensing Systems (INSS)*. IEEE, June 2010, pp. 233–240. ISBN: 978-1-4244-7911-5. DOI: 10.1109/INSS.2010.5573462. URL: http://ieeexplore.ieee.org/document/5573462/.

[66] Ricardo Chavarriaga et al. "The Opportunity challenge: A benchmark database for on-body sensor-based activity recognition". In: *Pattern Recognition Letters* 34.15 (Nov. 2013), pp. 2033–2042. ISSN: 01678655. DOI: 10.1016/j.patrec.2012.12.014. URL: http://dx.doi.org/10.1016/j.patrec.2012.12.014%20https://linkinghub.elsevier.com/retrieve/pii/S0167865512004205.

[67] Wallace Ugulino et al. "Wearable Computing: Accelerometers' Data Classification of Body Postures and Movements". In: *Proceedings of 21st Brazilian Symposium on Artificial Intelligence. Advances in Artificial Intelligence - SBIA 2012*. 2012, pp. 52–61. DOI: 10.1007/978-3-642-34459-6_6. URL: http://link.springer.com/10.1007/978-3-642-34459-6%7B%5C_%7D6.

[68] Barbara Bruno, Fulvio Mastrogiovanni, and Antonio Sgorbissa. "A public domain dataset for ADL recognition using wrist-placed accelerometers". In: *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, Aug. 2014, pp. 738–743. ISBN: 978-1-4799-6765-0. DOI: 10.1109/ROMAN.2014.6926341. URL: http://ieeexplore.ieee.org/document/6926341/.

[69] Mi Zhang and Alexander A. Sawchuk. "USC-HAD: A Daily Activity Dataset for Ubiquitous Activity Recognition Using Wearable Sensors". In: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing - UbiComp '12*. New York, New York, USA: ACM Press, 2012, p. 1036. ISBN: 9781450312240. DOI: 10.1145/2370216.2370438. URL: http://dl.acm.org/citation.cfm?doid=2370216.2370438.

[70] Chen Chen, Roozbeh Jafari, and Nasser Kehtarnavaz. "UTD-MHAD: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor". In: *2015 IEEE International Conference on Image Processing (ICIP)*. Vol. 2015-Decem. IEEE, Sept. 2015, pp. 168–172. ISBN: 978-1-4799-8339-1. DOI: 10.1109/ICIP.2015.7350781. URL: http://ieeexplore.ieee.org/document/7350781/.

[71] Sakorn Mekruksavanich and Anuchit Jitpattanakul. "Deep Convolutional Neural Network with RNNs for Complex Activity Recognition Using Wrist-Worn Wearable Sensor Data". In: *Electronics* 10.14 (July 2021), p. 1685. ISSN: 2079-9292. DOI: 10.3390/electronics10141685. URL: https://www.mdpi.com/2079-9292/10/14/1685.

[72] Konstantinos Peppas et al. "Real-time physical activity recognition on smart mobile devices using convolutional neural networks". In: *Applied Sciences (Switzerland)* 10.23 (2020), pp. 1–25. ISSN: 20763417. DOI: 10.3390/app10238482.

[73] Francisco Ordóñez and Daniel Roggen. "Deep Convolutional and LSTM Recurrent Neural Networks for Multimodal Wearable Activity Recognition". In: *Sensors* 16.1 (Jan. 2016), p. 115. ISSN: 1424-8220. DOI: 10.3390/s16010115. URL: http://www.mdpi.com/1424-8220/16/1/115.

[74] Shaohua Wan et al. "Deep Learning Models for Real-time Human Activity Recognition with Smartphones". In: *Mobile Networks and Applications* 25.2 (Apr. 2020), pp. 743–755. ISSN: 1383-469X. DOI: 10.1007/s11036-019-01445-x. URL: http://link.springer.com/10.1007/s11036-019-01445-x.

[75] Kun Xia, Jianguang Huang, and Hanyu Wang. "LSTM-CNN Architecture for Human Activity Recognition". In: *IEEE Access* 8 (2020), pp. 56855–56866. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.2982225. URL: https://ieeexplore.ieee.org/document/9043535/.

[76] Sakorn Mekruksavanich and Anuchit Jitpattanakul. "Smartwatch-based Human Activity Recognition Using Hybrid LSTM Network". In: *2020 IEEE SENSORS*. Vol. 2020-Octob. IEEE, Oct. 2020, pp. 1–4. ISBN: 978-1-7281-6801-2. DOI: 10.1109/SENSORS47125.2020.9278630. URL: https://ieeexplore.ieee.org/document/9278630/.

[77] Sakorn Mekruksavanich et al. "Enhanced Hand-Oriented Activity Recognition Based on Smartwatch Sensor Data Using LSTMs". In: *Symmetry* 12.9 (Sept. 2020), p. 1570. ISSN: 2073-8994. DOI: 10.3390/sym12091570. URL: https://www.mdpi.com/2073-8994/12/9/1570.

[78] Bolu Oluwalade et al. "Human Activity Recognition using Deep Learning Models on Smartphones and Smartwatches Sensor Data". In: *HEALTHINF 2021 - 14th International Conference on Health Informatics*. SCITEPRESS - Science and Technology Publications, 2021, pp. 1–6. ISBN: 978-989-758-490-9. DOI: 10.5220/0010325906450650. URL: https://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0010325906450650.

[79] Isibor Kennedy Ihianle et al. "A deep learning approach for human activities recognition from multimodal sensing devices". In: *IEEE Access* 8 (2020), pp. 179028–179038. ISSN: 21693536. DOI: 10.1109/ACCESS.2020.3027979.

[80] Peter Blouw et al. "Benchmarking Keyword Spotting Efficiency on Neuromorphic Hardware". In: *arXiv* (Dec. 2018). ISSN: 23318422. arXiv: 1812.01739. URL: http://arxiv.org/abs/1812.01739.

[81] Yexin Yan et al. "Comparing Loihi with a SpiNNaker 2 prototype on low-latency keyword spotting and adaptive robotic control". In: *Neuromorphic Computing and Engineering* 1.1 (Sept. 2021), p. 014002. ISSN: 2634-4386. DOI: 10.1088/2634-4386/abf150. URL: https://iopscience.iop.org/article/10.1088/2634-4386/abf150.

[82] Kyle Buettner and Alan D. George. "Heartbeat Classification with Spiking Neural Networks on the Loihi Neuromorphic Processor". In: *2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, July 2021, pp. 138–143. ISBN: 978-1-6654-3946-6. DOI: 10.1109/ISVLSI51109.2021.00035. URL: https://ieeexplore.ieee.org/document/9516760/.

[83] Mostafa Rahimi Azghadi et al. "Hardware Implementation of Deep Network Accelerators Towards Healthcare and Biomedical Applications". In: *IEEE Transactions on Biomedical Circuits and Systems* 14.6 (Dec. 2020), pp. 1138–1159. ISSN: 1932-4545. DOI: 10.1109/TBCAS.2020.3036081. arXiv: 2007.05657. URL: https://ieeexplore.ieee.org/document/9250613/.

[84] Peter Blouw and Chris Eliasmith. "Event-Driven Signal Processing with Neuromorphic Computing Systems". In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Vol. 2020-May. IEEE, May 2020, pp. 8534–8538. ISBN: 978-1-5090-6631-5. DOI: 10.1109/ICASSP40776.2020.9053043. URL: https://ieeexplore.ieee.org/document/9053043/.

[85] Evelina Forno et al. "PageRank Implemented with the MPI Paradigm Running on a Many-Core Neuromorphic Platform". In: *Journal of Low Power Electronics and Applications* 11.2 (2021), p. 25.

[86] Gianvito Urgese et al. "Benchmarking a Many-Core Neuromorphic Platform With an MPI-Based DNA Sequence Matching Algorithm". In: *Electronics* 8.11 (2019), p. 1342.

[87] Seoyeon Kim et al. "Survey of IoT platforms supporting artificial intelligence". In: *Proceedings of the Conference on Research in Adaptive and Convergent Systems*. New York, NY, USA: ACM, Sept. 2019, pp. 65–66. ISBN: 9781450368438. DOI: 10.1145/3338840.3355694. URL: https://dl.acm.org/doi/10.1145/3338840.3355694.

[88] Hongyu An, Dong Sam Ha, and Yang Cindy Yi. "Powering next-generation industry 4.0 by a self-learning and low-power neuromorphic system". In: *Proceedings of the 7th ACM International Conference on Nanoscale Computing and Communication*. New York, NY, USA: ACM, Sept. 2020, pp. 1–6. ISBN: 9781450380836. DOI: 10.1145/3411295.3411302. URL: https://dl.acm.org/doi/10.1145/3411295.3411302.

[89] Zhuoqing Chang et al. "A Survey of Recent Advances in Edge-Computing-Powered Artificial Intelligence of Things". In: *IEEE Internet of Things Journal* 8.18 (Sept. 2021), pp. 13849–13875. ISSN: 2327-4662. DOI: 10.1109/JIOT.2021.3088875. URL: https://ieeexplore.ieee.org/document/9453402/.

[90] Jan Stuijt et al. "μBrain: An Event-Driven and Fully Synthesizable Architecture for Spiking Neural Networks". In: *Frontiers in Neuroscience* 15.May (May 2021), pp. 1–15. ISSN: 1662-453X. DOI: 10.3389/fnins.2021.664208. URL: https://www.frontiersin.org/articles/10.3389/fnins.2021.664208/full.

[91] Howard Eichenbaum. "Time cells in the hippocampus: a new dimension for mapping memories". In: *Nature Reviews Neuroscience* 15.11 (Nov. 2014), pp. 732–744. ISSN: 1471-003X. DOI: 10.1038/nrn3827. URL: http://www.nature.com/articles/nrn3827.

[92] Aaron R Voelker and Chris Eliasmith. "Improving Spiking Dynamical Networks: Accurate Delays, Higher-Order Synapses, and Time Cells". In: *Neural Computation* 30.3 (Mar. 2018), pp. 569–609. ISSN: 0899-7667. DOI: 10.1162/neco_a_01046. URL: http://direct.mit.edu/neco/article-pdf/30/3/569/1046671/neco%7B%5C_%7Da%7B%5C_%7D01046.pdf%20https://direct.mit.edu/neco/article/30/3/569-609/8369.

[93] Aaron R. Voelker, Ivana Kajic, and Chris Eliasmith. "Legendre memory units: Continuous-time representation in recurrent neural networks". In: *Advances in Neural Information Processing Systems* 32.NeurIPS (2019). ISSN: 10495258.

[94] Aaron R Voelker and Chris Eliasmith. "Programming Neuromorphics Using the Neural Engineering Framework". In: *Handbook of Neuroengineering*. Singapore: Springer Singapore, 2021, pp. 1–43. DOI: 10.1007/978-981-15-2848-4_115-1. URL: https://link.springer.com/10.1007/978-981-15-2848-4%7B%5C_%7D115-1.

[95] Peter Blouw et al. "Hardware Aware Training for Efficient Keyword Spotting on General Purpose and Specialized Hardware". In: *arXiv* 1.1 (Sept. 2020). ISSN: 23318422. arXiv: 2009.04465v3. URL: http://arxiv.org/abs/2009.04465.

[96] Jing Zhang and Dacheng Tao. "Empowering Things With Intelligence: A Survey of the Progress, Challenges, and Opportunities in Artificial Intelligence of Things". In: *IEEE Internet of Things Journal* 8.10 (May 2021), pp. 7789–7817. ISSN: 2327-4662. DOI: 10.1109/JIOT.2020.3039359. arXiv: 2011.08612. URL: https://ieeexplore.ieee.org/document/9264235/.

[97] Jozsef Suto. "The effect of hyperparameter search on artificial neural network in human activity recognition". In: *Open Computer Science* 11.1 (June 2021), pp. 411–422. ISSN: 2299-1093. DOI: 10.1515/comp-2020-0227. URL: https://www.degruyter.com/document/doi/10.1515/comp-2020-0227/html.

[98] URL: https://github.com/microsoft/nni.

[99] Evelina Forno et al. "A Parallel Hardware Architecture For Quantum Annealing Algorithm Acceleration". In: *2018 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC)*. Vol. 2018-Octob. IEEE, Oct. 2018, pp. 31–36. ISBN: 978-1-5386-4756-1. DOI: 10.1109/VLSI-SoC.2018.8644777. URL: https://ieeexplore.ieee.org/document/8644777/.