



OPEN ACCESS

EDITED BY

Malu Zhang,
National University of Singapore, Singapore

REVIEWED BY

Wenjie Wei,
University of Electronic Science and
Technology of China, China
Yifei Yao,
Shanghai Jiao Tong University, China
Qi Xu,
Dalian University of Technology, China

*CORRESPONDENCE

Yuhang Li
✉ yuhang.li@yale.edu

RECEIVED 01 June 2023

ACCEPTED 25 July 2023

PUBLISHED 14 September 2023

CITATION

Li Y, Yin R, Kim Y and Panda P (2023) Efficient human activity recognition with spatio-temporal spiking neural networks. *Front. Neurosci.* 17:1233037. doi: 10.3389/fnins.2023.1233037

COPYRIGHT

© 2023 Li, Yin, Kim and Panda. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Efficient human activity recognition with spatio-temporal spiking neural networks

Yuhang Li*, Ruokai Yin, Youngeun Kim and Priyadarshini Panda

Department of Electrical Engineering, Yale University, New Haven, CT, United States

In this study, we explore Human Activity Recognition (HAR), a task that aims to predict individuals' daily activities utilizing time series data obtained from wearable sensors for health-related applications. Although recent research has predominantly employed end-to-end Artificial Neural Networks (ANNs) for feature extraction and classification in HAR, these approaches impose a substantial computational load on wearable devices and exhibit limitations in temporal feature extraction due to their activation functions. To address these challenges, we propose the application of Spiking Neural Networks (SNNs), an architecture inspired by the characteristics of biological neurons, to HAR tasks. SNNs accumulate input activation as presynaptic potential charges and generate a binary spike upon surpassing a predetermined threshold. This unique property facilitates spatio-temporal feature extraction and confers the advantage of low-power computation attributable to binary spikes. We conduct rigorous experiments on three distinct HAR datasets using SNNs, demonstrating that our approach attains competitive or superior performance relative to ANNs, while concurrently reducing energy consumption by up to 94%.

KEYWORDS

brain-inspired computing, neuromorphic computing, human activity recognition, spiking neural networks, hardware efficiency

1. Introduction

In recent years, the proliferation of smart devices, such as smartphones and fitness trackers, has led to a growing interest in understanding user activities and behavior for healthcare applications. Human Activity Recognition (HAR) (Lara and Labrador, 2012; Anguita et al., 2013; Vrigkas et al., 2015) is an area of research that aims to identify user activities, with applications spanning sports injury detection, well-being management, medical diagnostics, smart building solutions (Ramanujam et al., 2021), and elderly care (Nweke et al., 2019). To accomplish these objectives, HAR tasks rely on specific input patterns derived from various sensors embedded in smart devices, including accelerometers, gyroscopes, and electroencephalogram (EEG) sensors. As the data collected from wearable sensors are time series in nature, the recognition of temporal patterns in sensor data is crucial for achieving high accuracy and efficiency.

Traditionally, researchers have employed hand-crafted features and straightforward classifiers for HAR tasks. Feature extraction techniques can be broadly categorized into statistical and structural (Figo et al., 2010; Bulling et al., 2014). Statistical features, such as mean, median, time domain, and frequency domain, encapsulate the distribution properties of individual training data samples. In contrast, structural methods account for the interactions between different training data samples, exemplified by techniques like principal component analysis (PCA), linear discriminant analysis (LDA), and empirical cumulative

distribution functions (ECDF) (Abidine et al., 2018). Employing machine learning-based classifiers (Kim and Ling, 2009; Aggarwal and Xia, 2014; Shoaib et al., 2016) in conjunction with hand-crafted features has resulted in reasonably satisfactory performance.

In more recent studies, deep learning techniques have been adopted for end-to-end feature extraction and classification in HAR tasks (Nweke et al., 2018). These approaches employ convolutional layers in Artificial Neural Networks (ANNs) (Mnih et al., 2015; Ignatov, 2018; Wan et al., 2020) and optimize the model using gradient backpropagation. Due to the capacity of gradient descent optimization to automatically determine the most suitable parameters, ANNs have demonstrated proficient performance across diverse datasets. Figure 1 illustrates the process of this algorithm, where the ANN utilizes time series data from wearable sensors to predict human activity.

However, we contend that ANNs, which employ full precision (32-bit) computation and exhibit low sparsity, impose considerable computational complexity and energy consumption on wearable devices. As expounded by Rastegari et al. (2016) and Qin et al. (2020), 32-bit networks necessitate $58\times$ more operations compared to fully 1-bit networks. Furthermore, ANNs rely on ReLU neurons (Krizhevsky et al., 2012) that do not account for temporal correlations. This design choice may be suboptimal, particularly for time series data, as it simply adapts the ANN framework from the image domain.

Due to the high-efficiency demands on wearable devices, reducing the memory and computation cost of HAR models has been a crucial research problem. As an example, Cheng et al. (2022) propose to use an ensemble of a set of experts, where each expert is a simple linear feature extractor. In addition, Tang et al. (2020) use Lego bricks as lower-dimension filters. Although these methods reduce the hardware cost to a certain extent, they lack direct optimization on the hardware side. As we mentioned, an extremely low-bit neural network can reduce hardware costs by an order of magnitude (Li et al., 2019).

Hypothetically, simply adapting the current architecture to 1-bit networks will greatly impact the representation ability, and thus reduce the accuracy of HAR. To address this problem, one has to consider how to extract the temporal information in sensor data more effectively with discrete and limited 1-bit representation.

To address the aforementioned problems, we employ Spiking Neural Networks (SNNs) (Roy et al., 2019; Tavanaei et al., 2019; Deng et al., 2020; Panda et al., 2020; Li et al., 2021b; Xu et al., 2022, 2023; Zhu et al., 2022) in conjunction with convolutional layers for processing time series data in HAR tasks. HAR can benefit from SNNs in two key aspects: (1) SNNs leverage binary spikes (either 0 or 1) for activation, enabling multiplication-free and highly sparse computation, thereby reducing energy consumption for time series data (Zhang et al., 2018, 2021; Wu et al., 2021); (2) SNNs inherently model the temporal dynamics present in time series data, as spiking neurons within SNNs maintain a variable called the membrane potential over time. When the membrane potential surpasses a predefined threshold, the neuron fires a spike in the current time step. Capitalizing on these two advantages, our SNNs exhibit comparable or even superior performance to ANNs.

Additionally, we extend a previous hardware accelerator design to support 1D convolution along the time dimension, making it

suitable for SNN implementation (Yin et al., 2022). We evaluate our SNNs on three widely-used HAR datasets (UCI-HAR Anguita et al., 2013, UniMB SHAR Micucci et al., 2017, HHAR Stisen et al., 2015) and compare them with ANN baselines. Our SNNs achieve the same or higher accuracy than ANNs while reducing energy consumption by up to 94%.

In summary, our contributions are 3-fold:

1. We propose the use of SNNs for HAR tasks, significantly reducing energy consumption while integrating a temporally-evolving activation function.
2. We design a hardware accelerator tailored for deploying SNNs on edge devices.
3. We conduct extensive experiments on three HAR benchmarks, demonstrating that our SNNs outperform ANNs in terms of accuracy while maintaining energy-saving advantages.

2. Materials and methods

2.1. Notations

We use bold lower letters for vector representations. For example, \mathbf{x} and \mathbf{y} denote the input data and target label variables. Bold capital letters like \mathbf{W} denote the matrices (or tensors as clear from the text). Constants are denoted by small upright letters, e.g., a . With bracketed superscript and subscript, we can denote the time dimension and the element indices, respectively. For example, $\mathbf{x}_i^{(t)}$ means the i -th training sample at time step t .

2.2. Background of HAR

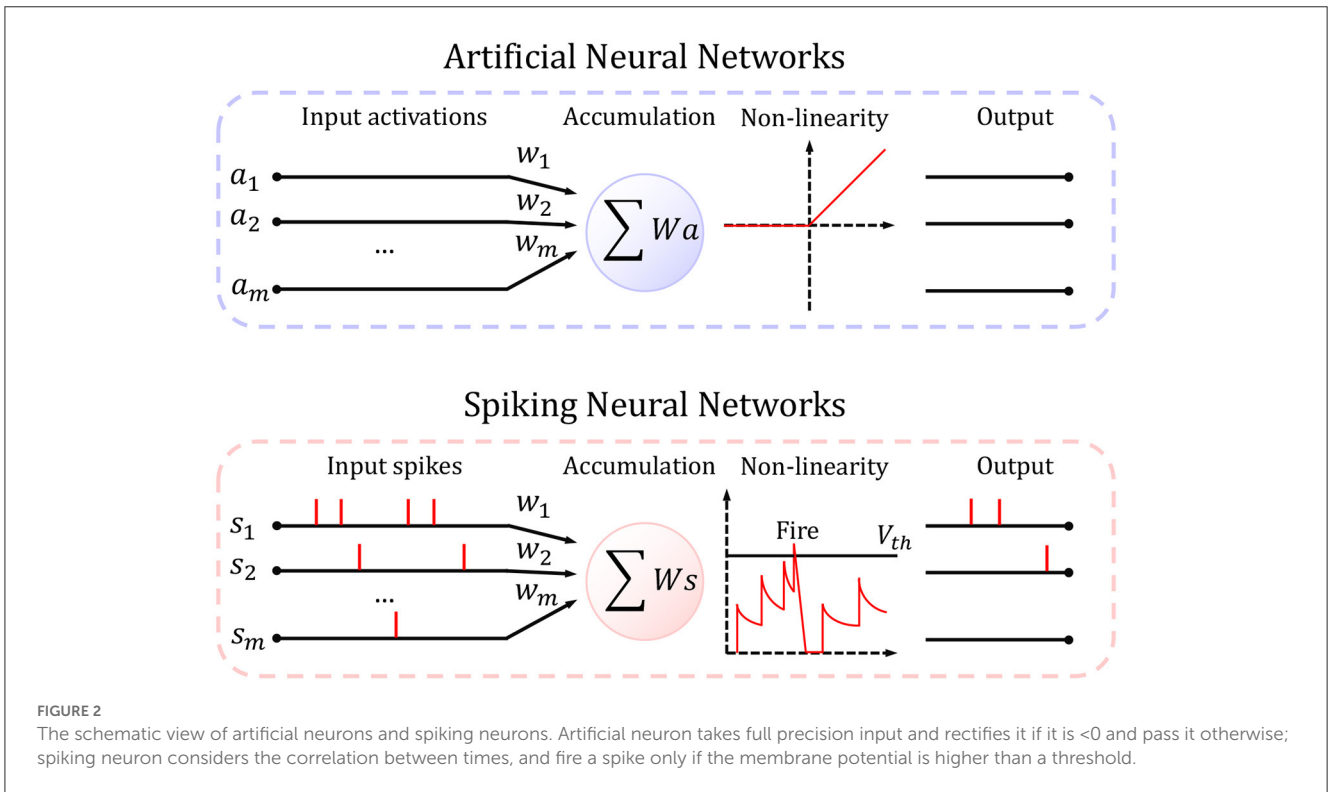
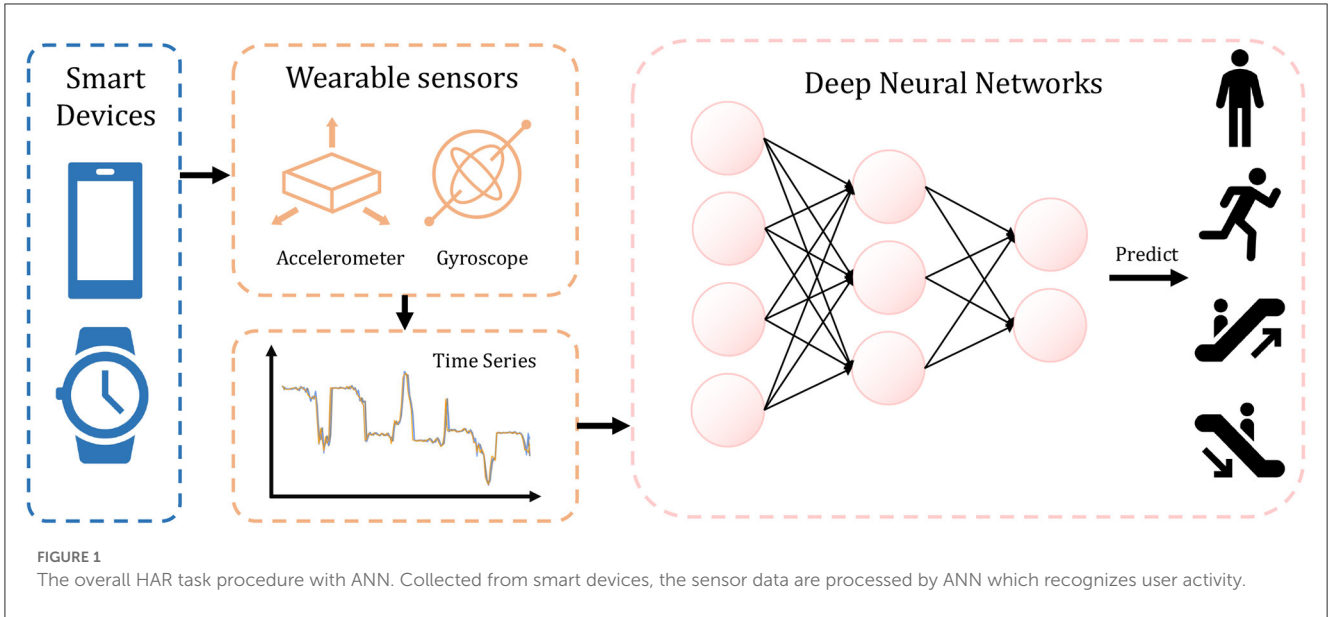
Concretely, we denote the wearable-based sensor dataset with $\{\mathbf{x}_i\}_{i=1}^N$, and each sample $\mathbf{x}_i \in \mathbb{R}^{T \times D}$ is collected when the wearer is doing certain activity \mathbf{y}_i , e.g., running, sitting, lying, standing, etc. Here, data samples are streaming and have T time steps in total. D is the dimension of the sensor's output. As an example, the accelerometer records the acceleration in the (x, y, z) -axis, thus $D = 3$ for the accelerometer data. We are interested in designing an end-to-end model $f(\cdot)$ and optimizing it to predict the activity label \mathbf{y} .

2.3. Spiking neuron

In this section, we introduce the definition of spiking neurons. We adopt the well-known Leaky-Integrate-and-Fire (LIF) neuronal model for spiking neurons (Liu and Wang, 2001), which constantly receives input and fires spikes through time. Formally, the LIF neuron maintains the membrane potential \mathbf{v} through time, and at t -th time step ($1 \leq t \leq T$), the membrane potential receives the pre-synaptic input charge $\mathbf{c}^{(t)}$, given by

$$\mathbf{v}^{(t+1),\text{pre}} = \tau \mathbf{v}^{(t)} + \mathbf{c}^{(t)}, \text{ where } \mathbf{c}^{(t)} = \mathbf{W}\mathbf{s}^{(t)}. \quad (1)$$

Here, τ is a constant between $[0, 1]$ representing the decay factor of the membrane potential as time flows, which controls



the correlation between time steps. $\tau = 0$ stands for 0 correlation and LIF degenerates to binary activation (Rastegari et al., 2016) without temporal dynamics, while $\tau = 1$ stands for maximum correlation and Deng and Gu (2021) and Li et al. (2021a) proves that LIF will become ReLU neuron when T is sufficiently large. $c^{(t+1)}$ is the product between weights W and the spike $s^{(t+1)}$ from previous layer. After receiving the input charge, the LIF neuron will fire a spike if the pre-synaptic membrane potential exceeds some threshold,

given by

$$s^{(t+1)} = \begin{cases} 1 & \text{if } v^{(t+1),pre} > V_{th} \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where V_{th} is the firing threshold. Note that the spike $s^{(t+1)}$ will propagate to the next layer, here we omit the layer index for simplicity.

If the LIF neurons fire a spike, the membrane potential will be reset. This can be done by either soft-reset or hard-reset, denoted by

$$\begin{cases} v^{(t+1)} = v^{(t+1),pre} \cdot (1 - s^{(t+1)}) & \# \text{ Hard-Reset} \\ v^{(t+1)} = v^{(t+1),pre} - s^{(t+1)} \cdot V_{th} & \# \text{ Soft-Reset} \end{cases}, \quad (3)$$

where hard-reset sets $v^{(t+1)}$ to 0, while soft-reset subtracts $v^{(t+1)}$ by V_{th} . We choose LIF neurons because $s^{(t+1)}$ is binary and dependent on input in previous time steps. Figure 2 describes the difference between ANN and SNN in a systematic way. In our experiments, we will conduct ablation studies on the decay factor, the firing threshold, and the reset mechanism.

2.3.1. Integrating spiking neurons into ANN

We first integrate spiking neurons into artificial neural networks by replacing their non-linear activation with LIF. As a result, we can compare the performance between artificial neurons and spiking neurons. Specifically, since the time series data naturally has a time dimension, we also integrate the pre-synaptic potential charge along this time dimension. For instance, suppose $\mathbf{a} \in \mathbb{R}^{n \times c \times T}$ is a pre-activation tensor, where n, c, T represent the batch size, channel number, and total time steps, respectively. We set the charge in each time step for LIF as the pre-activation in the corresponding time step, i.e., $\mathbf{c}^{(t)} = \mathbf{a}_{:, :, t}$. Then, we stack the output spikes along the time dimension again, i.e., $\mathbf{S} = \text{stack}(\{s^{(t)}\}_{t=1}^T)$, for calculating the pre-activation in next layer.

2.4. Optimization

Although LIF neurons manage to model the temporal features and produce binary spikes, the firing function (Equation 2) is discrete and thus produces zero gradients almost everywhere, prohibiting gradient-based optimization. Particularly, the gradient of loss (denoted by L) w.r.t. weights can be computed using the chain rule:

$$\frac{\partial L}{\partial \mathbf{W}} = \sum_{t=1}^T \frac{\partial L}{\partial s^{(t)}} \frac{\partial s^{(t)}}{\partial v^{(t),pre}} \left(\frac{\partial v^{(t),pre}}{\partial \mathbf{c}^{(t)}} \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{W}} + \sum_{t'=1}^{t-1} \frac{\partial v^{(t),pre}}{\partial v^{(t')}} \frac{\partial v^{(t')}}{\partial v^{(t'),pre}} \frac{\partial v^{(t'),pre}}{\partial \mathbf{c}^{(t')}} \frac{\partial \mathbf{c}^{(t')}}{\partial \mathbf{W}} \right). \quad (4)$$

Here, all terms can be differentiated except $\frac{\partial s^{(t)}}{\partial v^{(t),pre}}$ which brings zero-but-all gradients. To circumvent this problem, we use the surrogate gradient method. In detail, we use the triangle surrogate gradient, given by

$$\frac{\partial s^{(t)}}{\partial v^{(t),pre}} = \max \left(0, 1 - \left| \frac{v^{(t),pre}}{V_{th}} - 1 \right| \right). \quad (5)$$

As a result, SNNs can be optimized with stochastic gradient descent algorithms, as shown in Figure 3.

2.5. Hardware implementation

Finally, we introduce the hardware platform that we design for carrying out the experiments on energy efficiency. We extend the overall architecture and PE design from Yin et al. (2022) to support the necessary computation and data movement for our SNNs in HAR tasks. Owing to the 1D convolution and temporal dynamics that are naturally embedded in the time series data, the complexity of the hardware design has been largely reduced.

As shown in Figure 4, our systolic-array-based hardware platform equips one PE array and two global buffers for holding the weights and spikes. The size of the PE array and global buffers are configurable according to different network structures. In this work, we set the number of PEs to 128, weight (W) buffer to 32 KB, and spike (S) buffer to 576 bytes, for matching with the dataflow used in Yin et al. (2022). We briefly explain the computation and data movement flow below.

In Figure 4, at step 1, the entire weights for the layer are fetched into the global buffer from DRAM. The weights and the spikes will be written into the scratchpads inside PEs at step 2. At step 3, the accumulation is carried out for computing the $W_s^{(t)}$ and the partial sum result is added with the residual membrane potential from time step $t - 1$ at step 4. The latest membrane potential for time step t is then sent to the LIF unit at step 5 to generate the output spike $s^{(t)}$. According to the dataflow in Yin et al. (2022), each PE will only focus on working on one output neuron, and the PE array processes the whole output feature map in parallel. After that, spikes for the next layer will be written into the S buffer at step 6, and the whole process will repeat. Note that we can directly apply the input spike to skip the accumulation and weight scratchpad access if the input is equal to zero. We show the energy cost for each operation on the PE level in Table 1 for the reader's reference. Here, E_{mac} is the energy cost for a single multiply-accumulate (MAC) operation (note that the multiplication with spikes becomes logical AND operation between spikes and weights); E_{spa} is the energy cost for handling spike sparsity; E_{LIF} is the energy of a LIF operation; E_{ispad} and E_{wspad} are the single access energy to the input and weight scratchpad separately. All of the values are normalized by the energy cost of a MAC operation in ANN.

3. Experiments

In this section, we verify the effectiveness and efficiency of our SNNs on three popular HAR benchmarks. We first briefly provide the implementation details of our experiments and then compare our method with ANN baselines. Finally, we conduct ablation studies to validate our design choices.

3.1. Implementation details

We implement our SNNs and existing ANNs with the PyTorch framework (Paszke et al., 2019). For all our experiments, we use Adam optimizer (Kingma and Ba, 2014). All models are trained for 60 epochs, with batch size 128. The only flexible hyper-parameter is the learning rate, which is selected from $\{1e - 4, 3e - 4, 1e - 3\}$




FIGURE 3 An example of the forward and backward process of LIF neurons in 3 time steps. \rightarrow , forward; \rightarrow , backward; ①, potential charge; ②, fire; ③, reset; ④, integrate and decay.




FIGURE 4 The illustration of the hardware design we used for the experiment.

TABLE 1 Normalized energy cost for each operation on the PE level.

| Operation | Normalized energy cost |
|-------------|------------------------|
| E_{mac} | 0.175 |
| E_{spa} | 0 |
| E_{LIF} | 0.383 |
| E_{lspad} | 0.107 |
| E_{Wspad} | 1.712 |

The energy is normalized with the energy cost for one MAC operation in the ANN.

with the best validation accuracy. We use Cosine Annealing Decay for the learning rate schedule. For all three HAR datasets, we split them to 64% as the training set, 16% as the validation set, and 20% as the test set. We report test accuracy when the model reaches the best validation accuracy. Note that these datasets only have one label for each input sample, therefore top-1 accuracy is the same as the F-1 score. Similar to the SNN in image recognition tasks (Kim et al., 2022), the last layer of our SNN architecture is a fully connected layer. Therefore, we simply integrate all the membrane potentials in this layer for the softmax class prediction. We use the vanilla cross-entropy loss function rather than other specific loss functions (Deng et al., 2022; Zhu et al., 2022) to optimize our model. The dataset descriptions are shown below:

UCI-HAR (Anguita et al., 2013) contains 10.3 k instances collected from 30 subjects. It involves 6 different activities including

walking, walking upstairs, walking downstairs, sitting, standing, and lying. The sensors are the 3-axis accelerometer and 3-axis gyroscope (both are 50 Hz) from Samsung Galaxy SII.

UniMB SHAR (Micucci et al., 2017) contains 11.7 k instances collected from 30 subjects. It involves 17 different activities including 9 kinds of daily living activities and 6 kinds of fall activities. The sensor is the 3-axis accelerometer (maximum 50 Hz) from Samsung Galaxy Nexus I9250.

HHAR (Stisen et al., 2015) contains 57 k instances collected from 9 subjects. It involves 6 daily activities including biking, sitting, standing, walking, stair up, and stair down. The sensors are accelerometers from 8 smartphones and 4 smart watches (sampling rate from 50 to 200 Hz).

3.2. Comparison with ANNs

For ANN baselines, we select Convolutional Neural Networks (CNN) (Avilés-Cruz et al., 2019), DeepConvLSTM (DCL) (Mukherjee et al., 2020), Long Short Term Memory (LSTM) (Wang and Liu, 2020), and Transformer (Vaswani et al., 2017) architectures. We replace the ReLU neurons with spiking neurons, therefore, we can only integrate them into CNN and DeepConvLSTM since LSTM and Transformer have other activations like tanh and swish. The CNN architecture is marked by C32-MP2-C64-MP2-C64-MP2-FC, where each convolutional layer is a 1-dimensional convolution with a kernel size of 8. For DCL

TABLE 2 Accuracy (%) comparison between different networks on three HAR datasets (DCL, DeepConvLSTM).

| Model | UCI-HAR (Anguita et al., 2013) | SHAR (Mirucci et al., 2017) | HHAR (Straen et al., 2013) |
|-------------|--------------------------------|-----------------------------|----------------------------|
| CNN | 96.29 ± 0.12 | 92.38 ± 0.51 | 96.19 ± 0.14 |
| DCL | 97.87 ± 0.32 | 90.78 ± 1.05 | 97.15 ± 0.17 |
| LSTM | 82.41 ± 4.04 | 83.87 ± 0.96 | 95.59 ± 0.20 |
| Transformer | 96.02 ± 0.27 | 83.19 ± 0.74 | 95.82 ± 0.16 |
| SpikeCNN | 96.40 ± 0.15 | 94.04 ± 0.34 | 96.20 ± 0.09 |
| SpikeDCL | 98.86 ± 0.28 | 92.08 ± 0.77 | 97.52 ± 0.10 |

The bold values indicate that they achieve the highest accuracy under the same architecture family.

TABLE 3 Accuracy (%) comparison between our SNNs with existing ANNs, including CNN, LSTM, and DeepConvLSTM.

| Method | Model | UCI-HAR | SHAR | HHAR |
|-------------------------|----------|---------------------|---------------------|---------------------|
| Ronao and Cho (2016) | CNN | 94.79 | - | - |
| Khan et al. (2018) | CNN | - | - | 78.75 |
| Wang and Liu (2020) | LSTM | 91.65 | - | 85.82 |
| Mukherjee et al. (2020) | DCL | - | 92.30 | - |
| Zhu et al. (2018) | DCL | 97.31 | - | - |
| Ours | SpikeCNN | 96.40 ± 0.15 | 94.04 ± 0.34 | 96.20 ± 0.09 |
| Ours | SpikeDCL | 98.86 ± 0.28 | 92.08 ± 0.77 | 97.52 ± 0.10 |

The bold values indicate that they achieve the highest accuracy under the same architecture family.

architecture, it is marked by C64-C64-C64-C64-LSTM64, where each convolutional layer uses a kernel size of 5. Dropout is applied in Artificial CNN and DCL to reduce redundant activations, but not in spiking CNN and DCL. Each result is averaged from 5 runs (random seeds from 1,000 to 1,004) and includes a standard deviation value.

We summarize the results in Table 2, from which we find that SNNs have higher accuracy than the ANNs. For example, on the UniMB SHAR dataset, SpikeCNN has a 1.7% average accuracy improvement over its artificial CNN counterpart. Even more remarkably, the SpikeDeepConvLSTM (SpikeDCL) on the UCI-HAR dataset reaches 98.86% accuracy, which is 1% higher than DCL. Considering the accuracy is approaching 100%, the 1% improvement would be very significant. For UCI-HAR and HHAR datasets, we find SpikeCNN has similar accuracy to CNN, instead, the SpikeDeepConvLSTM consistently outperforms DeepConvLSTM, indicating that SNNs can be more coherent with the LSTM layer. Regarding the standard deviation of accuracy, we find that SNNs are usually more stable than ANNs, except for only one case, SpikeCNN on UCI-HAR.

We also compare our SNN with existing methods using ANNs on three HAR datasets. The results are summarized in Table 3. It can be found that our method achieves higher accuracy compared to these baselines, demonstrating the effectiveness of our method. For instance, our SpikeCNN has 1.7% higher accuracy than the CNN used in Ronao and Cho (2016) and our SpikeDCL obtains 1.5% higher accuracy than the DCL proposed in Zhu et al. (2018).

3.3. Ablation studies

In this section, we conduct ablation studies with respect to the (hyper)-parameters in the LIF neurons, including decay factor, threshold, and reset mechanism. We test SpikeDCL and SpikeCNN on UCI-HAR and SHAR datasets.

3.3.1. The effect of decay factor

We select 5 fixed decay factors from {0.0, 0.25, 0.5, 0.75, 1.0}. Note that as discussed before $\tau = 0$ indicates no correlation between two consecutive time steps, therefore SNN becomes equivalent to Binary Activation Networks (BAN), while $\tau = 1$ indicates full correlation. Additionally, we add another choice parameterized τ where the decay factor can be learned for each layer. This choice avoids the manual adjustments of the decay factor. Specifically, we initialize $b = 0$ and use $\tau = \text{sigmoid}(b)$ to represent the decay factor. The gradient w.r.t. c is given by

$$\frac{\partial L}{\partial b} = \sum_{t=1}^T \frac{\partial L}{\partial \mathbf{s}^{(t)}} \frac{\partial \mathbf{s}^{(t)}}{\partial \tau} \frac{\partial \tau}{\partial b} + \sum_{t'=1}^{t-1} \frac{\partial \mathbf{v}^{(t),\text{pre}}}{\partial \mathbf{v}^{(t')}} \frac{\partial \mathbf{v}^{(t')}}{\partial \mathbf{v}^{(t'),\text{pre}}} \frac{\partial \mathbf{v}^{(t'),\text{pre}}}{\partial \tau} \frac{\partial \tau}{\partial b}. \quad (6)$$

We provide all results in Table 4. We can find that τ has a huge impact on the final test accuracy. For the UCI-HAR dataset with SpikeDCL, the accuracy of $\tau = 0$ is 94.36% while the accuracy of $\tau = 0.75$ is 98.86%. Additionally, if we compare other $0 < \tau < 1$ cases with $\tau = 0$, we find that $\tau = 0$ always produces a large deficiency. This indicates that considering the temporal correlation with $\tau > 0$ is necessary for the time series tasks. It also verifies our hypothesis in Section 1 that simply using 1-bit without considering temporal information will degrade the accuracy. Moreover, for the SHAR dataset, the $\tau = 1$ case only has 60.55 accuracy while the $\tau = 0.25$ case achieves 91.72% accuracy.

It is also worthwhile to note that different datasets have varying optimal decay factor rates. The UCI-HAR favors 0.75 as its decay factor while the SHAR prefers 0.25. We think the primary reason for this change is that SHAR has sharper variation in its input and has a much larger range than UCI-HAR. Therefore, it should maintain a relatively low τ .

As for the parameterized decay factor, we do not observe its superiority over the fixed decay factor model. The parameterized τ generally achieves decent performance but not the best.

TABLE 4 Ablation study on the decay factor τ .

| Dataset | Model | Decay factor τ | | | | | |
|--------------------------------|----------|---------------------|--------------|-------|--------------|-------|-------|
| | | 0.0 | 0.25 | 0.5 | 0.75 | 1.0 | Param |
| UCI-HAR (Anguita et al., 2013) | SpikeCNN | 95.48 | 95.63 | 95.78 | 96.40 | 95.92 | 96.11 |
| | SpikeDCL | 94.36 | 96.50 | 97.57 | 98.86 | 96.60 | 97.37 |
| SHAR (Micucci et al., 2017) | SpikeCNN | 93.54 | 94.04 | 93.48 | 93.85 | 74.68 | 93.54 |
| | SpikeDCL | 89.53 | 92.08 | 90.93 | 90.10 | 60.55 | 91.53 |

The bold values indicate that they achieve the highest accuracy under the same architecture family.

TABLE 5 Ablation study on the firing threshold V_{th} and the reset mechanism.

| Dataset | Model | Firing threshold V_{th} | | | | Reset | |
|--------------------------------|----------|---------------------------|--------------|-------|-------|-------|--------------|
| | | 0.25 | 0.5 | 0.75 | 1.0 | Hard | Soft |
| UCI-HAR (Anguita et al., 2013) | SpikeCNN | 95.71 | 96.40 | 96.18 | 96.11 | 96.09 | 96.40 |
| | SpikeDCL | 98.27 | 98.86 | 97.60 | 96.81 | 98.53 | 98.73 |
| SHAR (Micucci et al., 2017) | SpikeCNN | 93.91 | 94.04 | 93.89 | 93.87 | 92.75 | 94.04 |
| | SpikeDCL | 91.42 | 92.08 | 91.72 | 91.53 | 91.13 | 92.08 |

The bold values indicate that they achieve the highest accuracy under the same architecture family.

3.3.2. The effect of firing threshold

We next study the effect of the firing threshold. Generally, the firing threshold is related to the easiness of firing a spike. We set the threshold as {0.25, 0.5, 0.75, 1.0} and run the same experiments with the former ablation. Here, through Table 5 we observe that the firing threshold has a unified pattern. SNN reaches its highest performance when the firing threshold is set to 0.5. This result is not surprising since 0.5 is in the mid of 0 and 1, and thus has the lowest error for the sign function. Meanwhile, we find the difference in accuracy brought by the firing threshold is lower than the decay factor. For instance, the largest gap when changing the threshold for SpikeDCL on the SHAR dataset is 0.65%, while this gap can be 32% when changing the decay factor. Therefore, the SNN is more sensitive to the decay factor rather than the threshold.

3.3.3. The effect of reset mechanism

Finally, we verify the reset mechanism for SNNs, namely soft-reset and hard-reset. The results are sorted in Table 5. For all cases, the soft-reset mechanism is better than the hard-reset. We think the reason behind this is that the hard reset will directly set the membrane potential to 0, therefore cutting off the correlation between intermediate time steps. Instead, the soft-reset keeps some previous time step's information on membrane potential after firing.

3.4. Hardware performance evaluation

In this section, we compare the hardware performance between SNN and ANN. Here, we compare two metrics, namely the activation sparsity and the energy consumption. Higher sparsity can avoid more computations with weights in hardware that supports sparse computation. We measure the sparsity either in ReLU (ANNs) or in LIF (SNNs) and visualize them in Figure 5 (blue chart). The ReLU in ANN usually has around 50% sparsity,

an intuitive result since the mean of activation is around 0. LIF neurons, however, exhibit a higher sparsity, approximately 80%, probably due to the threshold for firing being larger than 0. As a result, SNNs can save more operations in inference.

The second metric in hardware performance is energy consumption. We estimate the energy consumption by simulating the proposed hardware design in Section 2.5 together with our ReLU-based ANN baseline through the energy simulator proposed in Yin et al. (2022). The overall energy we consider consists of two parts: computing energy and data-moving energy. SNNs have advantages in computing energy due to their binary activation and higher sparsity. The results are shown in Figure 5 right. It can be seen that SNNs consume up to 94% less energy than ANNs, which could largely promote the battery life in smart devices. However, in the image processing domain, SNNs may have higher data-moving energy because they need to store the membrane potential and access them in the future (Yin et al., 2022, 2023; Moitra et al., 2023). We demonstrate that, in the HAR domain, SNNs have even lower data-moving energy than ANNs. The input data in HAR are augmented multiple times to generate the features in the time dimension. However, the SNNs in HAR do not need to increase the dimension of intermediate features to accommodate the time dimension resulting in lower data-moving costs. In summary, SNNs bring higher task performance due to the LIF neurons, and also energy efficiency due to the binary representation with high sparsity.

3.5. Convergence

In this section, we visualize the convergence curves of both ANN and SNN. We record the training accuracy and validation accuracy during training for the CNN and DCL models. The curves are shown in Figure 6. In the first figure, we can find that the CNN converges faster than the SpikeCNN. The training accuracy of ANN is always higher than the SNN.




FIGURE 5 Hardware costs comparison between ANNs and SNNs on UCI-HAR and SHAR datasets, respectively. We include sparsity and energy consumption.

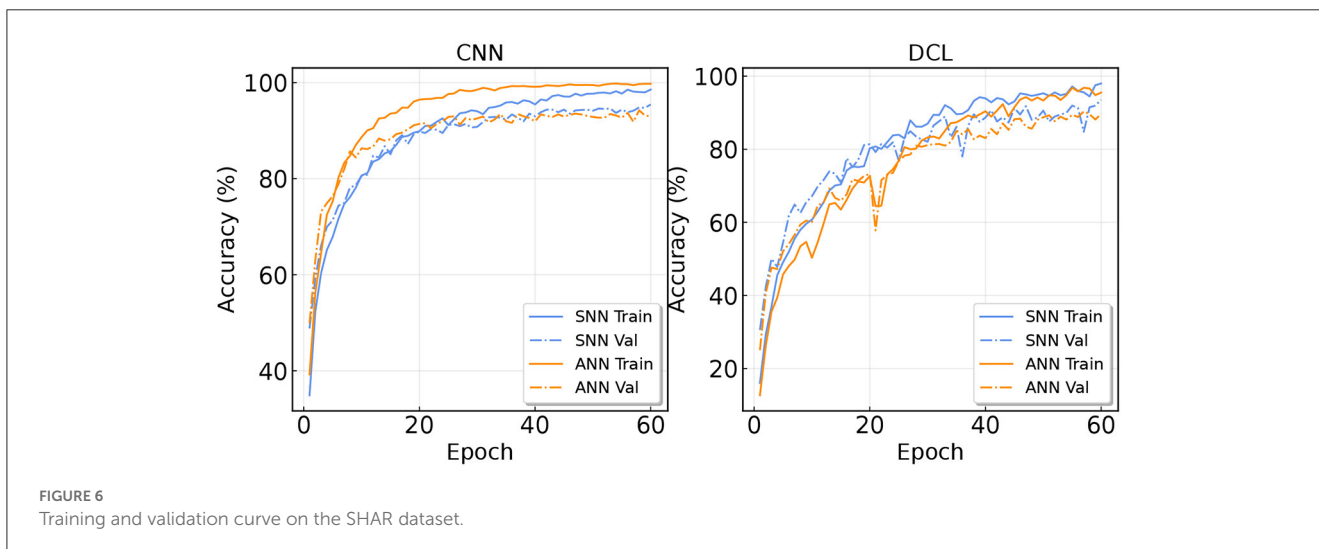


FIGURE 6 Training and validation curve on the SHAR dataset.

The validation accuracy of ANN also maintains its advantages at first, however, the validation accuracy of SNN becomes higher in the later stages. We conjecture that in the pure convolutional architecture, SNN is harder to be optimized than ANN and it may have a smaller generalization gap due to its binary activation.

For the right side of Figure 5, we record the curves of DeepConvLSTM. It can be seen that SNN has faster convergence in this case. The validation accuracy of SNN is always higher than ANN. This result confirms that SNN is more coherent with LSTM layers.

3.6. Representation similarity

In this section, we visualize the similarity between the ANN's and SNN's representation. We use Centered Kernel Alignment (CKA) (Kornblith et al., 2019; Li et al., 2023) to calculate the representation similarity index. We compare CNN and DCL on UCI-HAR and SHAR datasets. We compute the CKA value between convolutional or activation layers, for ReLU and LIF. Therefore, we can construct a heatmap with x, y axes being the layer index, and each entry is the CKA value of layers with those indices. The heatmaps are shown in Figure 7. In general, we find




FIGURE 7
Feature similarity measure between ANN and SNN using CKA.

that the first layer in ANN and SNN produces nearly the same representation. As the network goes deep, the similarity becomes lower implying SNN's latter layers extract different features from the HAR tasks as compared to ANNs. We can tentatively say that the difference in features may be the reason why SNNs and ANNs yield different accuracy on HAR tasks. We also discover that the shallow layers and the deep layers are very different, with a lower than 0.4 CKA value.

4. Conclusion

In this paper, we have shown the supremacy of Spiking Neural Networks (SNNs) over Artificial Neural Networks (ANNs) on HAR tasks, which, to our best knowledge, is the first. Compared to the original ANNs, SNNs utilize their LIF neurons to generate spikes through time, bringing energy efficiency as well as temporally correlated non-linearity. Our results show that SNNs achieve competitive accuracy while reducing energy significantly, and thus demonstrate the advantage of SNNs for low-power wearable devices.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

YL conceptualized the problem, implemented the algorithm, performed the experiments of algorithmic comparison, and wrote the initial manuscript. RY devised the hardware accelerator, performed the experiments of hardware comparison, wrote the hardware section, and gave suggestions to the manuscript. YK assisted with the experiments, gave suggestions to the manuscript, and draw several figures. PP conceptualized the problem, supervised the work, funded the project, and revised the manuscript. All authors contributed to the article and approved the submitted version.

Funding

This work was supported in part by CoCoSys, a JUMP2.0 center sponsored by DARPA and SRC, Google Research Scholar Award, the National Science Foundation CAREER Award, TII (Abu Dhabi), the DARPA AI Exploration (AIE) program, and the DoE MMICC center SEA-CROGS (Award #DE-SC0023198).

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships

References

- Abidine, B., Fergani, L., Fergani, B., and Oussalah, M. (2018). The joint use of sequence features combination and modified weighted svm for improving daily activity recognition. *Patt. Anal. Applic.* 21, 119–138. doi: 10.1007/s10044-016-0570-y
- Aggarwal, J. K., and Xia, L. (2014). Human activity recognition from 3d data: A review. *Patt. Recogn. Lett.* 48, 70–80. doi: 10.1016/j.patrec.2014.04.011
- Anguita, D., Ghio, A., Oneto, L., Parra Perez, X., and Reyes Ortiz, J. L. (2013). “A public domain dataset for human activity recognition using smartphones,” in *Proceedings of the 21th International European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (Bruges), 437–442.
- Avilés-Cruz, C., Ferreyra-Ramírez, A., Zúñiga-López, A., and Villegas-Cortéz, J. (2019). Coarse-fine convolutional deep-learning strategy for human activity recognition. *Sensors* 19, 1556. doi: 10.3390/s19071556
- Bulling, A., Blanke, U., and Schiele, B. (2014). A tutorial on human activity recognition using body-worn inertial sensors. *ACM Comput. Surv.* 46, 1–33. doi: 10.1145/2499621
- Cheng, X., Zhang, L., Tang, Y., Liu, Y., Wu, H., and He, J. (2022). Real-time human activity recognition using conditionally parametrized convolutions on mobile and wearable devices. *IEEE Sensors J.* 22, 5889–5901. doi: 10.1109/JSEN.2022.3149337
- Deng, L., Wu, Y., Hu, X., Liang, L., Ding, Y., Li, G., et al. (2020). Rethinking the performance comparison between snns and anns. *Neural Netw.* 121, 294–307. doi: 10.1016/j.neunet.2019.09.005
- Deng, S., and Gu, S. (2021). Optimal conversion of conventional artificial neural networks to spiking neural networks. *arXiv preprint arXiv:2103.00476*.
- Deng, S., Li, Y., Zhang, S., and Gu, S. (2022). Temporal efficient training of spiking neural network via gradient re-weighting. International Conference on Learning Representations (ICLR). *arXiv preprint arXiv:2202.11946*.
- Figio, D., Diniz, P. C., Ferreira, D. R., and Cardoso, J. M. (2010). Preprocessing techniques for context recognition from accelerometer data. *Pers. Ubiquitous Comput.* 14, 645–662. doi: 10.1007/s00779-010-0293-9
- Ignatov, A. (2018). Real-time human activity recognition from accelerometer data using convolutional neural networks. *Appl. Soft Comput.* 62, 915–922. doi: 10.1016/j.asoc.2017.09.027
- Khan, M. A. A. H., Roy, N., and Misra, A. (2018). “Scaling human activity recognition via deep learning-based domain adaptation,” in *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)* (IEEE), 1–9. doi: 10.1109/PERCOM.2018.8444585
- Kim, Y., Li, Y., Park, H., Venkatesha, Y., and Panda, P. (2022). Neural architecture search for spiking neural networks. *arXiv preprint arXiv:2201.10355*. doi: 10.1007/978-3-031-20053-3_3
- Kim, Y., and Ling, H. (2009). Human activity classification based on micro-doppler signatures using a support vector machine. *IEEE Trans. Geosci. Rem. Sens.* 47, 1328–1337. doi: 10.1109/TGRS.2009.2012849
- Kingma, D. P., and Ba, J. (2014). Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR). *arXiv preprint arXiv:1412.6980*.
- Kornblith, S., Norouzi, M., Lee, H., and Hinton, G. (2019). “Similarity of neural network representations revisited,” in *International Conference on Machine Learning* (Long Beach, CA: PMLR), 3519–3529.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems* 25.
- Lara, O. D., and Labrador, M. A. (2012). A survey on human activity recognition using wearable sensors. *IEEE Commun. Surv. Tutor.* 15, 1192–1209. doi: 10.1109/SURV.2012.110112.00192
- Li, Y., Deng, S., Dong, X., Gong, R., and Gu, S. (2021a). “A free lunch from ann: Towards efficient, accurate spiking neural networks calibration,” in *International Conference on Machine Learning Virtual Conference* (PMLR), 6316–6325.
- Li, Y., Dong, X., and Wang, W. (2019). Additive powers-of-two quantization: An efficient non-uniform discretization for neural networks. International Conference on Learning Representations (ICLR). *arXiv preprint arXiv:1909.13144*.
- Li, Y., Guo, Y., Zhang, S., Deng, S., Hai, Y., and Gu, S. (2021b). “Differentiable spike: Rethinking gradient-descent for training spiking neural networks,” in *Advances in Neural Information Processing Systems* (Virtual Conference) 34, 23426–23439.
- Li, Y., Kim, Y., Park, H., and Panda, P. (2023). Uncovering the representation of spiking neural networks trained with surrogate gradient. Transactions on Machine Learning Research (TMLR). *arXiv preprint arXiv:2304.13098*.
- Liu, Y.-H., and Wang, X.-J. (2001). Spike-frequency adaptation of a generalized leaky integrate-and-fire model neuron. *J. Comput. Neurosci.* 10, 25–45. doi: 10.1023/A:1008916026143
- Miccucci, D., Mobilio, M., and Napolitano, P. (2017). Unimib shar: A dataset for human activity recognition using acceleration data from smartphones. *Appl. Sci.* 7, 1101. doi: 10.3390/app7101101
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature.* 518, 529–533. doi: 10.1038/nature14236
- Moitra, A., Bhattacharjee, A., Kuang, R., Krishnan, G., Cao, Y., and Panda, P. (2023). “Spikesim: An end-to-end compute-in-memory hardware evaluation tool for benchmarking spiking neural networks,” in *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. doi: 10.1109/TCAD.2023.3274918
- Mukherjee, D., Mondal, R., Singh, P. K., Sarkar, R., and Bhattacharjee, D. (2020). Ensemconvnet: a deep learning approach for human activity recognition using smartphone sensors for healthcare applications. *Multim. Tools Applic.* 79, 31663–31690. doi: 10.1007/s11042-020-09537-7
- 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. Applic.* 105, 233–261. doi: 10.1016/j.eswa.2018.03.056
- Nweke, H. F., Teh, Y. W., Mujtaba, G., and Al-Garadi, M. A. (2019). Data fusion and multiple classifier systems for human activity detection and health monitoring: Review and open research directions. *Inf. Fusion* 46, 147–170. doi: 10.1016/j.inffus.2018.06.002
- Panda, P., Aketi, S. A., and Roy, K. (2020). Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. *Front. Neurosci.* 14, 653. doi: 10.3389/fnins.2020.00653
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems* 32.
- Qin, H., Gong, R., Liu, X., Bai, X., Song, J., and Sebe, N. (2020). Binary neural networks: A survey. *Patt. Recogn.* 105, 107281. doi: 10.1016/j.patcog.2020.107281

that could be construed as a potential conflict of interest.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- 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, 13029–13040. doi: 10.1109/JSEN.2021.3069927
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. (2016). “Xnor-net: Imagenet classification using binary convolutional neural networks,” in *European Conference on Computer Vision* (Amsterdam: Springer), 525–542. doi: 10.1007/978-3-319-46493-0_32
- Ronao, C. A., and Cho, S.-B. (2016). Human activity recognition with smartphone sensors using deep learning neural networks. *Expert Syst. Applic.* 59, 235–244. doi: 10.1016/j.eswa.2016.04.032
- Roy, K., Jaiswal, A., and Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. *Nature* 575, 607–617. doi: 10.1038/s41586-019-1677-2
- Shoib, M., Bosch, S., Incel, O. D., Scholten, H., and Havinga, P. J. (2016). Complex human activity recognition using smartphone and wrist-worn motion sensors. *Sensors* 16, 426. doi: 10.3390/s16040426
- Stisen, A., Blunck, H., Bhattacharya, S., Prentow, T. S., Kjærgaard, M. B., Dey, A., et al. (2015). “Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition,” in *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems* (Seoul), 127–140. doi: 10.1145/2809695.2809718
- Tang, Y., Teng, Q., Zhang, L., Min, F., and He, J. (2020). Layer-wise training convolutional neural networks with smaller filters for human activity recognition using wearable sensors. *IEEE Sensors J.* 21, 581–592. doi: 10.1109/JSEN.2020.3015521
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neur. Netw.* 111, 47–63. doi: 10.1016/j.neunet.2018.12.002
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in Neural Information Processing Systems* (Long Beach, CA), 30.
- Vrigkas, M., Nikou, C., and Kakadiaris, I. A. (2015). A review of human activity recognition methods. *Front. Robot. AI* 2, 28. doi: 10.3389/frobt.2015.00028
- 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. Applic.* 25, 743–755. doi: 10.1007/s11036-019-01445-x
- Wang, L., and Liu, R. (2020). Human activity recognition based on wearable sensor using hierarchical deep lstm networks. *Circ. Syst. Signal Proc.* 39, 837–856. doi: 10.1007/s00034-019-01116-y
- Wu, J., Xu, C., Han, X., Zhou, D., Zhang, M., Li, H., et al. (2021). Progressive tandem learning for pattern recognition with deep spiking neural networks. *IEEE Trans. Patt. Anal. Mach. Intell.* 44, 7824–7840. doi: 10.1109/TPAMI.2021.3114196
- Xu, Q., Li, Y., Shen, J., Liu, J. K., Tang, H., and Pan, G. (2023). “Constructing deep spiking neural networks from artificial neural networks with knowledge distillation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Seattle, WA), 7886–7895.
- Xu, Q., Li, Y., Shen, J., Zhang, P., Liu, J. K., Tang, H., et al. (2022). “Hierarchical spiking-based model for efficient image classification with enhanced feature extraction and encoding,” in *IEEE Transactions on Neural Networks and Learning Systems*. doi: 10.1109/TNNLS.2022.3232106
- Yin, R., Li, Y., Moitra, A., and Panda, P. (2023). Mint: Multiplier-less integer quantization for spiking neural networks. *arXiv preprint arXiv:2305.09850*.
- Yin, R., Moitra, A., Bhattacharjee, A., Kim, Y., and Panda, P. (2022). Sata: Sparsity-aware training accelerator for spiking neural networks. *arXiv preprint arXiv:2204.05422*. doi: 10.1109/TCAD.2022.3213211
- Zhang, M., Qu, H., Belatreche, A., Chen, Y., and Yi, Z. (2018). A highly effective and robust membrane potential-driven supervised learning method for spiking neurons. *IEEE Trans. Neural Netw. Learn. Syst.* 30, 123–137. doi: 10.1109/TNNLS.2018.2833077
- Zhang, M., Wang, J., Wu, J., Belatreche, A., Amornpaisannon, B., Zhang, Z., et al. (2021). Rectified linear postsynaptic potential function for backpropagation in deep spiking neural networks. *IEEE Trans. Neur. Netw. Lear. Syst.* 33, 1947–1958. doi: 10.1109/TNNLS.2021.3110991
- Zhu, Q., Chen, Z., and Soh, Y. C. (2018). A novel semisupervised deep learning method for human activity recognition. *IEEE Trans. Ind. Inf.* 15, 3821–3830. doi: 10.1109/TII.2018.2889315
- Zhu, Y., Yu, Z., Fang, W., Xie, X., Huang, T., and Masquelier, T. (2022). “Training spiking neural networks with event-driven backpropagation,” in *Advances in Neural Information Processing Systems* (New Orleans, LA), 35, 30528–30541.