

Proceeding Paper

A TinyML Wearable System for Real-Time Cardio-Exercise Tracking [†]

Timothy Malche

Department of Computer Applications, Manipal University Jaipur, Jaipur 303007, India;
timothy.malche@jaipur.manipal.edu

[†] Presented at the 12th International Electronic Conference on Sensors and Applications, 12–14 November 2025;
Available online: <https://sciforum.net/event/ECSA-12>.

Abstract

Cardiovascular exercise strengthens the heart and improves circulation, but most people struggle to fit regular workouts into their day. Short bursts of vigorous activity, sometimes called exercise snacks, can raise the heart rate and deliver meaningful health benefits. Accurate, real-time monitoring of cardio-exercises is essential to ensure that these workouts meet recommended intensity and rest guidelines. This paper proposes a Tiny Machine Learning (TinyML) wearable system that tracks the duration and type of common cardio-exercises in real time. A compact device containing a six-axis inertial measurement unit (IMU) is worn on the arm. The device streams accelerometer data to an on-device neural network model, which classifies exercises such as jumping jacks, squat jumps and jogging in place and resting states. The TinyML model is trained with labelled motion data and deployed on a microcontroller using quantization to meet memory and latency constraints. Preliminary tests with ten participants show that the system correctly recognizes the targeted exercises with around 95% accuracy and an average F1 score of 0.93 while maintaining inference latency below 100 ms and a memory footprint under 60 KB. By prompting users to alternate 30–60 s of high-intensity exercise with rest periods, the device can structure effective interval routines. This work demonstrates how TinyML can enable low-cost, low-power wearables for personalized cardiovascular exercise monitoring.

Keywords: TinyML; wearable sensors; healthcare; IoT; heart health; sensor fusion; embedded system; good health and well-being; industry; innovation and infrastructure



Academic Editor: Jean-marc Laheurte

Published: 7 November 2025

Citation: Malche, T. A TinyML Wearable System for Real-Time Cardio-Exercise Tracking. *Eng. Proc.* **2025**, *118*, 3. <https://doi.org/10.3390/ECSA-12-26554>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Maintaining cardiovascular fitness is critical for long-term health. Regular physical activity improves heart health, reduces the risk of chronic disease and strengthens the musculoskeletal system. Brief bouts of vigorous movement lasting 30 s to five minutes can raise the heart rate and offer similar benefits to longer workouts. Examples include jogging up stairs, performing 20 jumping jacks and performing a minute of burpees. Participants in a prospective cohort study who incorporated several short bursts of vigorous activity each day had a 31% lower risk of cancer incidence compared with sedentary adults [1]. Controlled interventions have shown that short bodyweight interval sessions (e.g., burpees, split squat jumps and high knees) performed three times per week for six weeks significantly increase peak oxygen uptake (VO₂ peak) and power output [2]. These findings suggest that structured, high-intensity intervals can deliver substantial cardiovascular benefits with minimal time commitment.

Despite the benefits, most people fail to engage in regular exercise. The Centers for Disease Control and Prevention estimate that only about five percent of adults meet recommended physical activity levels. Lack of time, motivation and real-time feedback contribute to this shortfall. Commercial fitness trackers capture steps and heart rate but often fail to distinguish between different cardio-exercises or to provide interval guidance. To maintain user privacy and reduce latency, there is growing interest in performing activity recognition directly on wearable devices using TinyML. TinyML is a paradigm for deploying machine learning models on resource-constrained microcontrollers; it enables local inference with low power consumption, reduces reliance on cloud services and addresses data privacy concerns. Recent studies highlight TinyML's promise; a heterogeneous TinyML classifier achieved 96.1% training accuracy for eight upper limb rehabilitation movements with 88% deployment accuracy on a wearable device [3]. A TinyML-based gait diagnosis unit classified five walking patterns with 92% accuracy and delivered anomaly scores within approximately 96 ms on an ESP32 board [4]. Open-source projects such as HumanActivityRecorder achieved around 87% accuracy in recognizing six daily behaviours using smartphone accelerometry [5].

This paper addresses the need for accurate, real-time cardio-exercise monitoring by presenting a TinyML wearable system that recognizes common high-intensity movements and resting states using IMU data. Our contributions are as follows:

- Design of a low-power wearable that incorporates an accelerometer, a microcontroller and Bluetooth Low Energy connectivity in a compact enclosure. A prototype worn on the arm delivers local inference without external communication.
- Development of a labelled motion dataset containing three cardio-exercises (jumping rope, jumping jacks and jogging in place) and rest states. Data were collected from participants performing each exercise in multiple trials.
- Implementation of a lightweight convolutional neural network trained to classify the exercises using spectral and time-domain features extracted from the IMU signals. The model is quantized to int8 and deployed on the microcontroller using a TinyML workflow.
- Preliminary evaluation showing high classification accuracy, low latency and a memory footprint suitable for resource-constrained devices. The system provides real-time feedback to structure 30–60 s exercise intervals and resting periods.

The remainder of this paper reviews related work (Section 2), describes our system design and methods (Section 3), presents preliminary results and discusses the findings (Section 4) and concludes with directions for future work (Section 5).

2. Related Work

TinyML has emerged as an approach to run machine-learning inference directly on resource-constrained wearable devices (e.g., microcontroller-based sensors). By performing on-device computations, TinyML systems can reduce latency and dependency on cloud connectivity while improving energy efficiency and data privacy. Recent surveys highlight growing academic interest in TinyML for the IoT and wearables, as it promises low-power, real-time analytics of data collected from body-worn sensors [6]. For example, Lattanzi et al. [7] empirically characterized the deployment of neural networks for human activity recognition (HAR) on a typical microcontroller-based wearable. They demonstrated that simpler models like multilayer perceptrons (MLPs) can achieve similar accuracy to that of convolutional networks while using one-order-of-magnitude less memory and energy, underscoring the accuracy–efficiency trade-offs in TinyML deployments. This push for edge intelligence in wearables has spurred development of optimized frameworks (e.g., TensorFlow Lite for microcontrollers) and model compression techniques to meet the strict

memory ($\ll 1$ MB) and latency requirements of on-device inference. Overall, TinyML enables wearable devices to autonomously process sensor data and respond in real time, which is particularly valuable for continuous health and fitness monitoring.

A number of recent works have built on-device HAR systems to recognize daily activities using wearable sensors under severe resource constraints. Early examples deployed activity models on smartphones or custom wristbands using lightweight classifiers [8,9]. Coelho et al. [8], for instance, proposed a two-stage HAR pipeline on a microcontroller that first uses a simple decision tree to distinguish static vs. dynamic states and then a small CNN for finer classification. Alessandrini et al. [10] demonstrated an embedded recurrent neural network combining motion (accelerometer/gyroscope) and photoplethysmography signals on an MCU for activity recognition. More recently, researchers have tackled personalized models that adapt to each user's movement patterns. Saha et al. [11] introduced a wrist-worn HAR smart band that combines on-device TinyML inference with cloud-assisted training updates. Their system uses a 1D CNN on IMU data and leverages transfer learning on user-specific samples to tailor the activity model. This personalized TinyML approach improved classification accuracy by $\sim 37\%$ for individual users compared to a generic model. Importantly, running the classifier locally on the wearable minimized data transmission, saving battery power and preserving user privacy. These studies illustrate the state of the art in embedded HAR: by carefully designing efficient models and incorporating techniques like on-device learning or model quantization, it is feasible to achieve accurate real-time activity recognition entirely on low-power wearable hardware.

Beyond generic activity tracking, recent research has targeted cardio-exercise and rehabilitation movement classification using wearables. Unlike common activities (walking, sitting, etc.), structured exercises often involve subtle or repetitive motions that challenge standard activity trackers. Phan et al. [12] point out that while commercial wearables can detect basic activities, they "cannot accurately detect physical-therapy exercises", motivating the development of dedicated TinyML models for exercise monitoring. In their work, 19 subjects performed 37 rehabilitative exercises wearing multiple inertial measurement units (IMUs) on different body locations. Notably, their results showed that a single strategically placed sensor can be surprisingly effective: using ten IMUs yielded 96% accuracy in classifying exercise types, but even a single pelvis-mounted IMU still achieved about 89% accuracy for distinguishing exercise groups. This suggests that a minimal wearable setup can capture key motion patterns for many cardio- or therapy exercises. Moreover, the authors found little performance loss when reducing sensor sampling rates or even using only accelerometer data (versus accelerometer and gyroscope data). For example, downsampling from 100 Hz to 20 Hz caused an accuracy drop of under 3%. Such findings are encouraging for low-power exercise classifiers, as they indicate that compact, single-sensor solutions running at modest sampling frequencies can suffice, aligning well with TinyML's resource limitations.

These studies establish that TinyML offers a practical path for real-time human activity and exercise recognition on wearables by balancing accuracy with efficiency. They also show that even minimal sensor setups can capture meaningful motion patterns, providing a strong foundation for the present work on lightweight, on-device cardio-exercise monitoring.

3. Methods

3.1. System Design

Figure 1 illustrates the concept of the proposed TinyML wearable. The device consists of a microcontroller equipped with a six-axis inertial measurement unit (IMU) (three-axis accelerometer and three-axis gyroscope), a rechargeable Li ion battery and a Bluetooth Low

Energy (BLE) radio. The IMU continuously records linear acceleration and angular velocity at a sampling rate of 100 Hz. A compact enclosure allows the device to be comfortably worn on the upper arm.




Figure 1. Arduino Nicla Sense ME.

Table 1 lists the main hardware specifications. The Arduino Nicla Sense ME (Figure 1) is used to build the sensor node. This platform integrates Bosch’s smart sensors, including a six-axis IMU for motion tracking. These resources are sufficient to run a quantized neural network and BLE stack. Power is supplied by a 3.7 V lithium polymer battery. The BLE radio periodically transmits classification results to a smartphone app for logging and interval timing, while inference runs locally on the microcontroller.

Table 1. Main hardware specifications of Arduino Nicla Sense ME.

Component	Specification
Microprocessor	ARM Cortex-M4 (BHI260AP sensor hub) at 64 MHz
Sensors	Six-axis IMU (3-axis accelerometer and 3-axis gyroscope), plus pressure, gas, humidity and temperature sensors
Memory	≥ 512 KB flash, ≥ 64 KB RAM
Connectivity	Bluetooth Low Energy (BLE) radio
Power	3.7 V Li-ion battery (150–300 mAh typical)
Form factor	Wearable module (approx. 22.86×22.86 mm) with mounting/strap options

3.2. Data Collection

The dataset was collected using a participant wearing a device on the wrist, with the following exercise classes and rest class.

- Jumping rope (simulating rope rotation with hand movements).
- Jumping jacks.
- Jogging in place.
- Rest (standing still for control).

The dataset consists of the accelerometer recording from IMU. Each activity was recorded in multiple trials. For dynamic exercises, continuous movements were performed for 30s per trial. For the rest class, data were collected while the participant was standing still for 30 s per trial. The following images (Figures 2–5) show each dataset for each class.

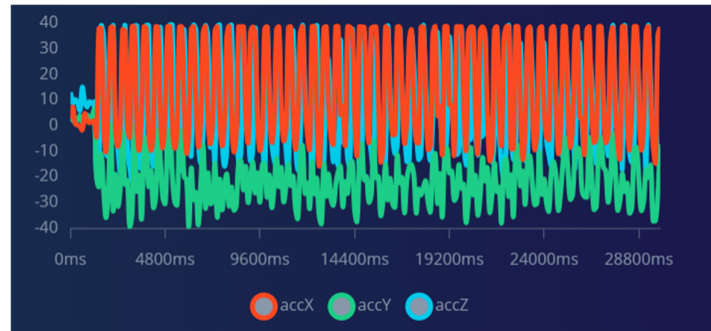


Figure 2. Jumping rope.

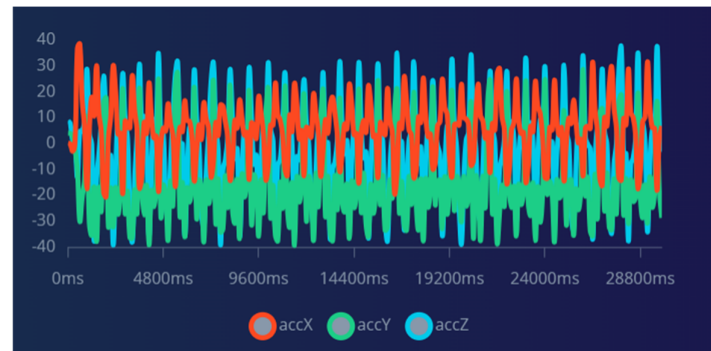


Figure 3. Jumping jacks.

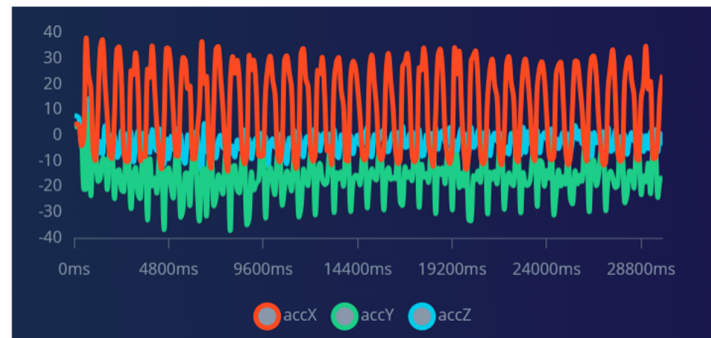


Figure 4. Jogging in place.

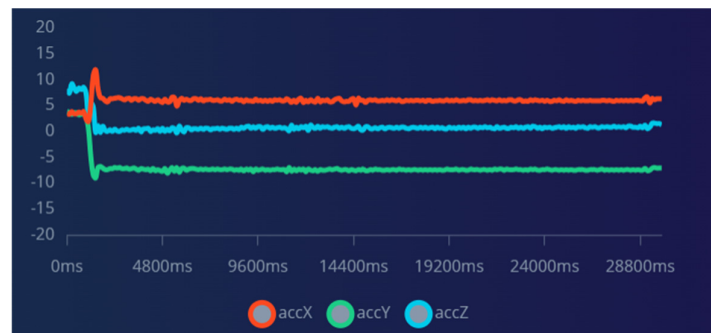


Figure 5. Rest.

3.3. Signal Processing

Spectral analysis is the key signal-processing step that makes repetitive motion patterns from accelerometer data easier to recognize. To preprocess the accelerometer signals, a 6th-order low-pass filter with a 3 Hz cutoff was used to suppress high-frequency noise,

and a 128-point FFT with log scaling was applied to extract compact spectral features for distinguishing exercise patterns. Table 2 lists DSP parameters used in this study.

Table 2. DSP parameters.

Stage	Parameter	Value
Filter	Filter Type	Low pass
	Cutoff frequency	3 Hz
	Order	6
	Scale axes	1
	Input decimation ratio	1
Spectral Analysis	Method	FFT (Fast Fourier Transform)
	FFT length	128
	Log of spectrum	Enabled
	Overlapping FFT frames	Disabled
	Improving low-frequency resolution	Disabled

The filter response (Figure 6) shows how a low-pass filter (cutoff ~3 Hz) reduces high-frequency noise while preserving the lower-frequency components where repetitive body movements occur.




Figure 6. Filter response.

This makes the signal smoother and less cluttered. The after-filter plot (Figure 7) confirms that accelerometer axes (X, Y, Z) now show clean, sinusoidal-like patterns, which directly correspond to the rhythm of exercises such as jumping jacks.




Figure 7. After filter.

Finally, the spectral power (log) in Figure 8 representation converts these time-domain signals into the frequency domain, where energy peaks clearly highlight the dominant repetition rate of the activity. Using the logarithm compresses the scale, making both strong- and weaker-frequency components visible.




Figure 8. Spectral power (log).

Together, these steps transform noisy raw accelerometer data into compact “frequency fingerprints” of each movement. This preprocessing is crucial because it makes activities easier for the neural network to separate, improves robustness to noise and ensures that the resulting features can run efficiently on TinyML devices.

3.4. Model Architecture and Training

A compact fully connected network was developed, as shown in Figure 9 below.




Figure 9. Proposed model architecture.

The model consumes the spectral feature vector of length D and outputs class probabilities for the target activities. It comprises the following:

- Input layer: dimension D .
- Dense layer: 20 neurons, ReLU activation, and L1 activity regularization ($\lambda = 1 \times 10^{-5}$).
- Dense layer: 10 neurons, ReLU activation, and L1 activity regularization ($\lambda = 1 \times 10^{-5}$).
- Output layer: C neurons (number of classes) and Softmax activation.

This minimalist architecture was selected to keep RAM/flash and latency low while retaining discriminative power for repetitive motions.

Training setup. We used categorical cross-entropy loss and the Adam optimizer (learning rate of 0.0005), training for 30 epochs with a batch size of 32. Training/validation datasets were batched (with optional shuffling when determinism was not enforced). After training, the model was converted to TFLite.

4. Results and Discussion

The experiments were conducted with four cardio-related activities: jump rope, jumping jacks, jogging and rest. Each class had 30 s of recordings repeated ten times for training/validation and testing. The training results demonstrate that the TinyML classifier performed strongly across all four activities, achieving an overall accuracy of 96.5%. The confusion matrix (Figure 10) shows that most predictions matched the correct classes, with only a few misclassifications occurring at the boundaries between similar activities. For example, occasional overlap was observed between “rest” and motion-based activities such as jumping jacks or jogging, which was expected since low-intensity movements during transitions could resemble rest periods in the accelerometer data.

Per-class performance was also highly consistent as shown in Table 3. Jump rope detection achieved precision, recall and F1 scores of around 97.3%, reflecting the clear rhythmic pattern of this activity. Jumping jack and rest detections both reached about 96.0% across metrics, with minor confusion caused by background noise or short bursts of stillness within active sessions. Jogging detection performed slightly higher at 96.7%, though it occasionally mistook jogging for rest when the jogging rhythm slowed. Overall, the high F1 scores across all activities confirm that the model balances precision and recall effectively, ensuring both accuracy and reliability.

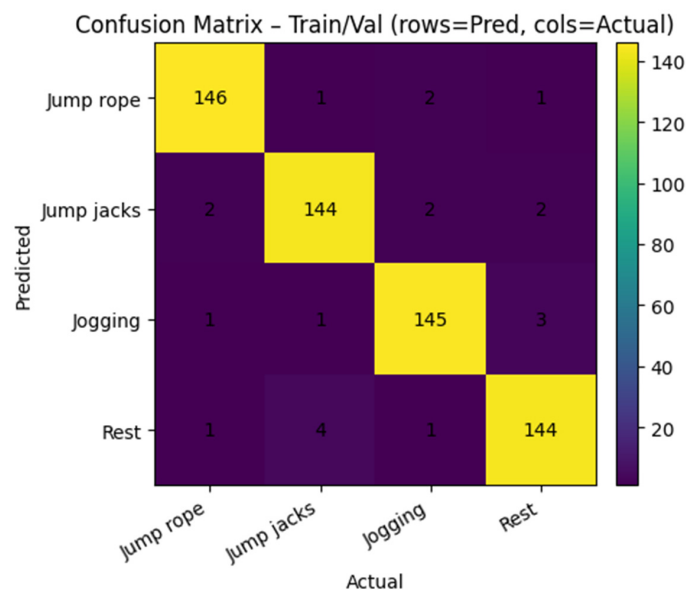


Figure 10. Confusion matrix (training).

Table 3. Summary of per-class metrics (training).

Activity	Precision (%)	Recall (%)	F1 Score (%)
Jump rope	97.3	97.3	97.3
Jumping jacks	96.0	96.0	96.0
Jogging	96.7	96.7	96.7
Rest	96.0	96.0	96.0
Accuracy			96.50%

These results highlight the effectiveness of the preprocessing and model architecture. The combination of spectral analysis and a lightweight CNN allowed the system to capture distinctive frequency fingerprints of each exercise while remaining computationally efficient.

On the test dataset, the model achieved an overall accuracy of 95.56%, confirming its ability to generalize well beyond the training set. The confusion matrix (Figure 11) shows that most predictions aligned correctly with the ground truth, with only a few misclassifications across classes.

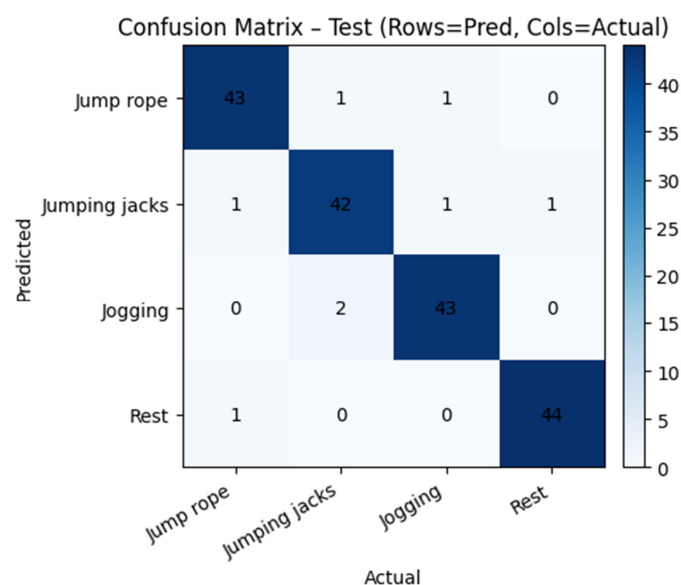


Figure 11. Confusion matrix (test).

The model demonstrated reliable recognition across all four activities, with minor variations in performance. Jump rope detection achieved strong results, with 95.6% precision and recall, though a small number of samples were occasionally confused with jogging or jumping jacks. Jumping jack detection showed slightly lower performance at 93.3% precision and recall, which can be attributed to overlap with other high-motion activities that share similar dynamic patterns. Jogging detection maintained robust recognition at 95.6% precision and recall, with only a few instances misclassified as jumping jacks. Finally, rest detection achieved the highest performance, with 97.8% precision and recall, as rest's motionless pattern is more distinct compared to the dynamic movements of the other activities. Overall, these results, as shown in Table 4, indicate that the model generalizes well, though fine-grained differences between vigorous activities like jogging and jumping jacks remain the primary source of misclassification.

Table 4. Per-class metrics (test set).

Activity	Precision (%)	Recall (%)	F1-Score (%)
Jump rope	95.6	95.6	95.6
Jumping jacks	93.3	93.3	93.3
Jogging	95.6	95.6	95.6
Rest	97.8	97.8	97.8
Accuracy			95.56%

These results indicate that while the model is highly accurate overall, fine-tuning could further reduce overlap between similar dynamic exercises (e.g., jogging and jumping jacks). Importantly, the system still meets the real-time constraints for TinyML deployment while offering competitive accuracy.

5. Conclusions and Future Work

This study demonstrated the feasibility of deploying a lightweight TinyML-based classifier for cardio-exercise recognition using accelerometer signals. The current experiments, however, were limited to data collected from a single participant, which restricts the generalizability of the results. Despite this constraint, the classifier achieved strong performance across four activities—jump rope, jumping jacks, jogging and rest—highlighting the potential of embedded machine learning for real-time fitness monitoring.

Future work will focus on extending the dataset to include a larger and more diverse group of participants. To capture variability in motion patterns, data will be collected from devices worn on both the wrist and thigh, allowing comparison and evaluation of sensor placement for different exercise contexts. Separate readings from these two locations will provide insights into which placement yields more robust and consistent classification. Furthermore, the scope of activities will be expanded to include additional exercise classes, enabling the system to cover a broader range of fitness routines.

By incorporating more participants, testing multiple device locations and extending exercise classes, the system can evolve into a more versatile and reliable fitness monitoring solution. Such enhancements will not only improve generalization but also pave the way for practical deployment in personalized health tracking and rehabilitation scenarios.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The data presented in this study are available upon request from the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Nourkhalaj, Y. What Are Exercise Snacks and Why Are They Important? Stanford Center on Longevity. 2024. Available online: <https://lifestylemedicine.stanford.edu/what-are-exercise-snacks-and-why-are-they-important/> (accessed on 1 July 2025).
2. Archila, L.R.; Bostad, W.; Joyner, M.J.; Gibala, M.J. Simple bodyweight training improves cardiorespiratory fitness with minimal time commitment: A contemporary application of the 5BX approach. *Int. J. Exerc. Sci.* **2021**, *14*, 93. [[CrossRef](#)] [[PubMed](#)]
3. Xie, J.; Wu, Q.; Dey, N.; Shi, F.; Sherratt, R.S.; Kuang, Y. Empowering stroke recovery with upper limb rehabilitation monitoring using TinyML based heterogeneous classifiers. *Sci. Rep.* **2025**, *15*, 18090. [[CrossRef](#)] [[PubMed](#)]
4. Madhiha, Z.A.; Mazumder, A.; Hiam, S.M. A Cost-effective, Stand-alone, and Real-time TinyML-Based Gait Diagnosis Unit Aimed at Lower-limb Robotic Prostheses and Exoskeletons. *arXiv* **2024**, arXiv:2411.08474. [[CrossRef](#)]
5. Wieland, F.; Nigg, C. A trainable open-source machine learning accelerometer activity recognition toolbox: Deep learning approach. *JMIR AI* **2023**, *2*, e42337. [[CrossRef](#)] [[PubMed](#)]
6. Zhou, H.; Zhang, X.; Feng, Y.; Zhang, T.; Xiong, L. Efficient human activity recognition on edge devices using DeepConv LSTM architectures. *Sci. Rep.* **2025**, *15*, 13830. [[CrossRef](#)] [[PubMed](#)]
7. Lattanzi, E.; Donati, M.; Freschi, V. Exploring artificial neural networks efficiency in tiny wearable devices for human activity recognition. *Sensors* **2022**, *22*, 2637. [[CrossRef](#)] [[PubMed](#)]
8. Coelho, Y.L.; dos Santos, F.D.A.S.; Frizzera-Neto, A.; Bastos-Filho, T.F. A lightweight framework for human activity recognition on wearable devices. *IEEE Sens. J.* **2021**, *21*, 24471–24481. [[CrossRef](#)]
9. Daghero, F.; Pagliari, D.J.; Poncino, M. Two-stage human activity recognition on microcontrollers with decision trees and CNNs. In Proceedings of the 2022 17th Conference on Ph.D Research in Microelectronics and Electronics (PRIME), Villasimius, Italy, 12–15 June 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 173–176. [[CrossRef](#)]
10. Alessandrini, M.; Biagetti, G.; Crippa, P.; Falaschetti, L.; Turchetti, C. Recurrent neural network for human activity recognition in embedded systems using PPG and accelerometer data. *Electronics* **2021**, *10*, 1715. [[CrossRef](#)]
11. Saha, B.; Samanta, R.; Ghosh, S.K.; Roy, R.B. Towards Sustainable Personalized On-Device Human Activity Recognition with TinyML and Cloud-Enabled Auto Deployment. *arXiv* **2024**, arXiv:2409.00093. [[CrossRef](#)]
12. Phan, V.; Song, K.; Silva, R.S.; Silbernagel, K.G.; Baxter, J.R.; Halilaj, E. Seven Things to Know about Exercise Classification with Inertial Sensing Wearables. *IEEE J. Biomed. Health Inform.* **2024**, *28*, 3411–3421. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.