

# HARNode: A Time-Synchronised, Open-Source, Multi-Device, Wearable System for Ad Hoc Field Studies

Philipp Lepold

philipp.lepold@kit.edu

Karlsruhe Institute of Technology

Karlsruhe, Germany

Tobias Röddiger

tobias.roeddiger@kit.edu

Karlsruhe Institute of Technology

Karlsruhe, Germany

Michael Beigl

michael.beigl@kit.edu

Karlsruhe Institute of Technology

Karlsruhe, Germany

## Abstract

Human activity recognition (HAR) research often lacks accessible, comprehensive field data. Commercial systems are rarely open source, hard to expand, and limited by issues like node synchronisation, data throughput, unclear sensor placement, complexity, and high cost. As a result, researchers typically use only a few intuitively placed sensors and conduct limited field trials. HARNode overcomes these challenges with a fully open-source hardware and software platform. Each node includes an ESP32-S3 module (AtomS3), a 9-axis IMU (Bosch BMX160), pressure and temperature sensors (Bosch BMP388), a display, and an I<sup>2</sup>C port. Data is streamed via Wi-Fi, with NTP-based time synchronisation achieving 1 ms accuracy. The system runs for up to 8 hours and is built using off-the-shelf parts, a simple online PCB service, and a compact 3D-printed housing with Velcro straps, enabling flexible and scalable body placement while requiring little hardware knowledge. In a study with ten subjects wearing eleven HARNodes each, setup took under five minutes per person. A random forest classifier distinguished walking from stair-climbing transitions, showing the benefits of sensor-overprovisioning: Seven nodes achieved  $\approx 98\%$  accuracy, matching the performance of all eleven. These findings confirm HARNode's value as a fast-deploying, scalable tool for field-based HAR research and optimised sensor placement.

## CCS Concepts

- Computer systems organization → Sensors and actuators;
- Human-centered computing → Ubiquitous and mobile devices;
- Computing methodologies → Machine learning.

## Keywords

Human Activity Recognition, Wearable Sensors, Sensor Nodes, Open-Source Hardware, Multi-Device Systems, Stair Walking Prediction, Exoskeleton

## ACM Reference Format:

Philipp Lepold, Tobias Röddiger, and Michael Beigl. 2025. HARNode: A Time-Synchronised, Open-Source, Multi-Device, Wearable System for Ad Hoc Field Studies. In *Companion of the 2025 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '25), October 12–16, 2025, Espoo, Finland*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3714394.3756255>



This work is licensed under a Creative Commons Attribution 4.0 International License.

*UbiComp Companion '25, Espoo, Finland*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1477-1/2025/10

<https://doi.org/10.1145/3714394.3756255>

## 1 Introduction and Related Work

Wearable inertial sensors have become extremely compact and affordable, leading to their widespread use in mobile and wearable computing research. In particular, inertial measurement units (IMUs) are commonly used for human activity recognition (HAR), including classification of activities such as level walking and stair climbing [23, 29]. A major limitation in human activity detection research and use is the availability of a large amount of collected data, especially in very specific application domains where meaningful data can only be collected in the field. In our application example, HAR using IMUs is essential for real-time gait transition detection and adaptive assistance in wearable exoskeletons to support elderly and mobility-impaired individuals [17, 25].

Currently, due to technical limitations - e.g. node synchronisation time and throughput - and practical considerations - setup time and cost - we mostly rely on a limited number of wearable sensors on the body, typically positioned at one or two body locations [21]. This is suboptimal, as the correct positioning of the wearable on the body is based more on intuition than evidence.

In this paper, we present a fully open-source, extendable wearable system that addresses these issues and is tuned for fast and easy setup (including a fast setup procedure) and data collection in the field with an (almost) arbitrarily large number of highly synchronised sensor nodes.

**Table 1: Comparison of adaptable wearable IMU sensor systems.** Open = open-source; Max = maximum number of IMU nodes; BT = Bluetooth; prop. = proprietary; 7 nodes are a typical max. for BT

System	Open	Max	Connectivity	Features
OpenSenseRT [6]	Yes	6–8+	Wired sensors	Low cost; easy assembly
movisens EdaMove 4 [2]	No	(7)	BT, USB	EDA; steps; waterproof
Xsens Awinda [3]	No	17	prop.	High accuracy
APDM Opal [5]	No	24	prop.	Gait analysis
Noraxon [4]	No	(7)	prop. +BLE	EMG integration
BioX Bands [27]	No	(7)	BT SPP	FSR for muscle-contraction
<b>HARNode (ours)</b>	Yes	100+	Wi-Fi (802.11 b/g/n)	Fast field setup; low cost; extendable

Existing open-source alternatives, such as OpenSenseRT [24], are constrained by wired configurations, external battery requirements, and complex setup procedures, or are commercial. Table 1 gives a short overview of various systems and compares them to this paper's approach. In addition to these works, specific devices have been reported, for example for mass sensing on the body [26]. We have excluded such work, as it is not intended to be easily usable or reproducible by researchers.

Central to the solution presented in this paper is the use of off-the-shelf components, including an ESP32-based AtomS3 microcontroller with integrated display for status monitoring and sensor

identification. Each node has a built-in 9-axis IMU (accelerometer, gyroscope, magnetometer), as well as barometric pressure and temperature sensors. All the hardware is housed in a compact 3D-printed case, which also fits the battery and is secured to the body with flexible Velcro straps, making the nodes lightweight, wireless, and easy to attach to different parts of the body. Data is streamed over Wi-Fi, and we use a high-precision NTP-based method for time synchronisation of the nodes. All elements of HARNode are completely open source, including hardware, firmware, and the server.

In addition to describing the hardware, software, system, and its use to set up field data collection, we also report a small proof-of-concept study that demonstrates the capabilities of the system. Here, we conducted a study with 10 participants, each wearing 11 HARNodes positioned on different parts of the body. The goal was to collect motion data during stair approach to systematically investigate which sensor placements are most informative for detecting transitions relevant to exoskeleton control.

## 2 Challenges and Design Requirements

Current work on human activity recognition (HAR) in the wild is constrained by data scarcity – often addressed using additional synthetic data (e.g. Hong et al. [15]), communication and synchronisation limits [10], high deployment cost [16], and ad hoc sensor placement [14][19]. Translating these observations into explicit design requirements is essential for scalable, evidence-based sensing systems. We identified the following requirements.

- (1) **Time alignment and constrained throughput.** Multi-node wearables drift by several milliseconds on commodity Bluetooth links, corrupting feature fusion and frequency analysis [9].  
**Requirement:** Best-of-the-art commodity sub-millisecond time-stamping, high ( $\gg 50$  kbit/s) throughput, local buffering with burst upload
- (2) **Deployment overhead.** Every additional node inflates hardware, fitting and maintenance costs; user surveys list donning time, comfort and aesthetics as primary barriers [18, 22].  
**Requirement:** single-board MCUs with integrated IMUs, tool-less Velcro or textile mounts, semi-automatic node configuration to cut setup below 20 seconds per node.
- (3) **Evidence-free sensor placement.** Recognition accuracy varies sharply with body location; mis-placement can halve F-scores [11].  
**Requirement:** Standardised setup for systematic placement campaigns, with high number of nodes to find optimal placements, and modular housings enabling quick repositioning.
- (4) **Wearability aspects.** Wearability aspects must be taken into account, as summarised by Zeagler [28].  
**Requirement:** Hardware and housing should be appropriate for attachment, taking into account e.g. attachment points and other wearability factors and practical factors such as battery life.

## 3 Hard- and Software

The HARNode Wearable system consists of three main elements: The HARNode hardware, the software (HARNode device firmware

and server software), and an off-the-shelf computer or laptop as a server. The hardware and software developed in this work are open-source and tailored for ease of use, also outside the laboratory environment.

### 3.1 Hardware

We developed a PCB and platform based on the M5Stack AtomS3<sup>1</sup>. The PCB snaps into a custom 3D-printed housing that also fits the battery and has Velcro straps for easy attachment. Overall lifetime with a single fully charged battery is designed to be between 2-8 hours for the below given sampling rates and power consumption numbers (variations depend on battery size and Wi-Fi/RF situation).

**3.1.1 AtomS3.** The core of the HARNode is an AtomS3 microcontroller, whose schematic and source code is also open-source under the MIT licence. It is based on the ESP32-S3 architecture (Xtensa LX7 dualcore, 300  $\mu$ s/min RTC deviation [13]) and has an integrated display and button, as well as Wi-Fi (IEEE 802.11 b/g/n) and Bluetooth 5 (LE) connectivity. In our setup while using a sample rate of 166.67 Hz, the average power consumption is 350 mA @ 3.3 V (Wi-Fi + CPU).

**3.1.2 PCB.** An overview of the PCB and its schematic can be seen in Figure 1. The components include:

- *Bosch BMX160*, a 9-axis inertial measurement unit (IMU): accelerometer (up to 1600 Hz, 180  $\mu$ A), gyroscope (up to 6400 Hz, 850  $\mu$ A), magnetometer (12 Hz, 600  $\mu$ A)
- *Bosch BMP388* pressure sensor (up to 200 Hz, 700  $\mu$ A, 2 cm resolution in height difference) and temperature sensor (200 Hz,  $\pm 0.5$  °C)
- *NanJing Top Power ASIC TP4056X* linear lithium-ion charging controller
- *Winsok Semiconductors WSTDW01* battery protection circuit
- *Texas Instruments TPS61023* boost converter: to supply the AtomS3 with 5 V, from 3 V - 4.2 V from the battery
- *Jauch LP523450JU or LP103450JH* 950 mAh / 5.5 mm thickness / 25 g or 1900 mAh / 10 mm / 35 g 3.7 V lithium-polymer battery (connected to the board)
- On/Off switch
- USB-C charging port
- 2.54 mm header pins to plug in the AtomS3 (and additional sensor boards), providing power and access to the sensors via an I<sup>2</sup>C bus

With the schematic and Gerber files provided<sup>2</sup>, the assembled PCB can be ordered online, e.g. at JLCPCB<sup>3</sup>. Only the battery needs to be soldered to the PCB; apart from that, no hardware knowledge is required.

**3.1.3 3D-printed enclosure.** The 3D-printed housing, depicted in Figure 2A, fits the PCB and provides, by default, 35 x 55 x 5.5 mm space for a battery, e.g. a Jauch LP523450JU 950 mAh 3.7 V lithium-polymer battery. Like the PCB, the housing is open-source and its STL file is provided, for online production, or by printing yourself<sup>4</sup>. The battery is installed by soldering its terminals to the pads on the

<sup>1</sup><https://docs.m5stack.com/en/core/AtomS3>

<sup>2</sup><https://github.com/teco-kit/HARNode>

<sup>3</sup><https://jlcpcb.com>

<sup>4</sup><https://github.com/teco-kit/HARNode>

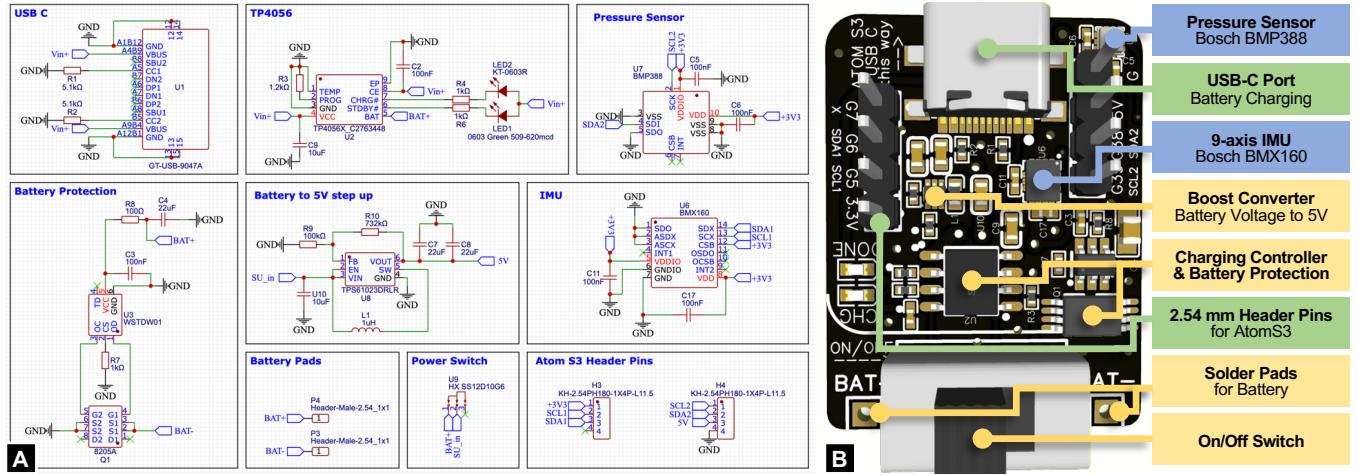


Figure 1: (A) Schematic diagram of the sensor node, showing the sensors, and power circuitry; (B) Assembled sensor node PCB with labelled components.

PCB. This housing provides a tight fit for the PCB; it slides right in and stays in place, requiring no adhesives. To the left and right of the housing, there are slots where Velcro straps are attached to mount the HARNode to the body, as explained below.

**3.1.4 Velcro Straps.** To attach the HARNode to the body, a flexible 50 mm Velcro loop band is used, depicted in Figure 2B. The narrower slot in the housing is where the strap is stapled, and the broader side features a pass-through slot for adjusting the strap. After threading the strap through this slot, the Velcro hook strip at the end can be secured to the loop band. For this, 5 cm of 50 mm self-adhesive Velcro hook band is glued onto the end of the loop band.

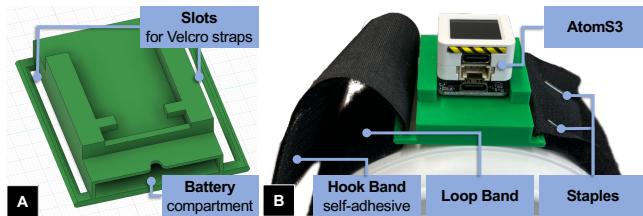


Figure 2: (A) Rendering of the 3D-printed HARNode housing; (B) Assembled HARNode with Velcro straps and labelled components.

## 3.2 Assembly

The assembly of a HARNode is straightforward and requires only five simple steps. Anyone who is comfortable soldering two wires to the PCB is able to assemble the HARNode.

**3.2.1 Prerequisites.** To build the HARNode, the items listed in Table 2 are needed. A stapler and a soldering iron are also required (and a 3D printer if the case is self-printed). The cost of one HARNode depends on the quantity of required items ordered but does not exceed approximately 48 €. The price of 20 assembled PCBs

including shipping was 217.95 € at JLCPCB, which is less than 11 € per PCB.

Table 2: Items needed to build a HARNode, including approximate prices.

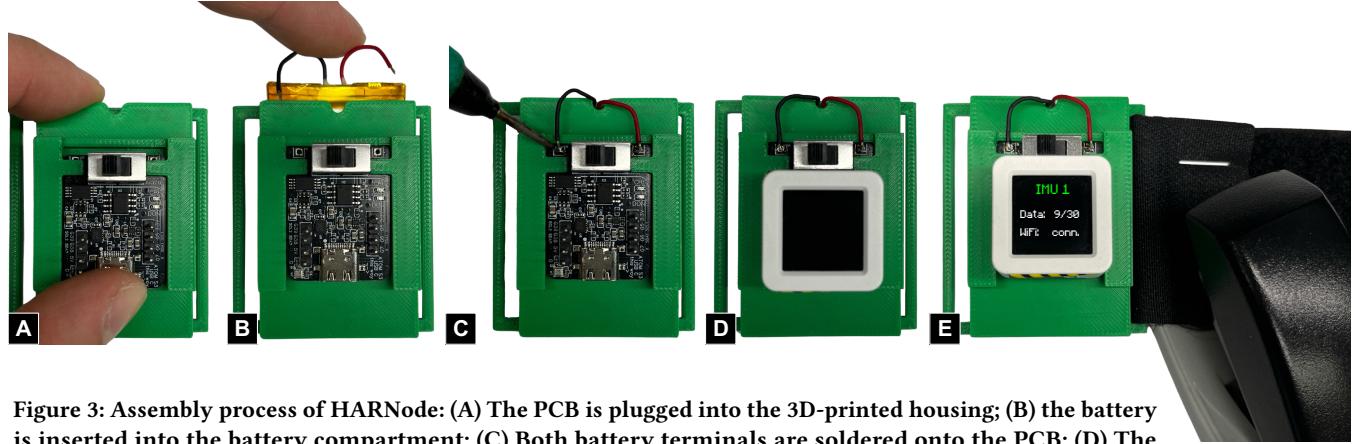
Item	Approx. Price (€)
Assembled PCB	11
3D printed Housing	<1
Suitable Battery	10
50 mm Loop Band (per meter)	3
50 mm Adhesive Hook Band (per meter)	3
M5Stack Atom3	20
<b>Total</b>	<b>48</b>

**3.2.2 Assembly Process.** The assembly process is depicted in Figure 3 and consists of five steps: 1) First, the PCB is plugged into the 3D printed housing by pushing it through the mounting brackets of the housing. The housing ensures a tight fit; no adhesives are required. 2) Then, the battery is inserted into the battery compartment underneath, and 3) both terminals are soldered onto the corresponding pads on the PCB. 4) Now the AtomS3 can be plugged in, onto the PCB. 5) Then, the Velcro strap loop band is attached by inserting it into the narrower slot of the housing, using a stapler to secure it. Finally, about 5 cm of self-adhesive hook band is glued onto the end of the loop band.

## 3.3 Software

The software needed to operate HARNodes includes the firmware running on the AtomS3's as well as the server that receives the UDP packets sent by the HARNodes. Both have been custom-developed and are released as open-source software<sup>5</sup>.

<sup>5</sup><https://github.com/teco-kit/HARNode>



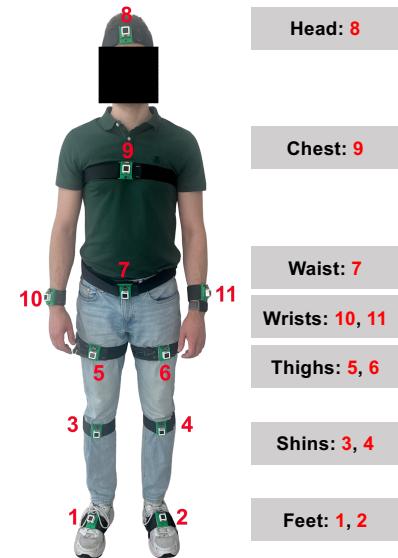
**Figure 3: Assembly process of HARNode:** (A) The PCB is plugged into the 3D-printed housing; (B) the battery is inserted into the battery compartment; (C) Both battery terminals are soldered onto the PCB; (D) The AtomS3 is plugged in; (E) The loop band is attached by stapling, and a piece of hook band is glued to its end.

**3.3.1 Firmware.** The firmware of the nodes is fully customisable, to tweak the sampling rate and sensor configuration that are used. It is written in C++ using the Arduino framework and is deployed via PlatformIO<sup>6</sup>, making it easy to adapt to your needs. Each node connects through a list of predefined Wi-Fi SSIDs that must be provided in the firmware beforehand. In this way, we reduce the setup time while maintaining flexibility to run parallel studies. The node first sends three NTP-requests towards the local server (and interpolates the received times to accommodate for round trip time), to synchronise each node to about 1 ms in average (5 ms max.) [20] for standard field setups with one access point and multiple HARNodes. NTP requests are repeated every minute to keep the deviation of the RTC clock in the microsecond level. Once the NTP response is received and the node is synchronised, it starts sending UDP packets with a default of 30 sensor readings per packet, with an adjustable default sampling rate of 166.67 Hz (read every 6 milliseconds), which we found to be a good compromise between motion detection, battery life, and achievable time synchronisation while preventing UDP packet fragmentation [12].

**3.3.2 Server.** The NodeJS server provides the NTP responses that each HARNode requests, receives, and stores all UDP data packets, and provides a web-based GUI to start recordings and synchronise a video feed to the HARNodes, provided by the webcam or a connected phone.

### 3.4 Operation

To use HARNodes, a Wi-Fi access point is required. Any access point can be used, including smartphone-based hotspots or other mobile access points, enabling ad hoc setups outside the laboratory. The amount of HARNodes usable is limited only by the router's and server's hardware used. For example, phone hotspots are often limited to a few connected devices [8], but ordinary Wi-Fi access points can typically handle many more connections. In our setup, each HARNode sent a packet with 1082 bytes of data at 5.6 Hz (166.67 Hz sampling rate, 30 datapoints of 9-axis 32-bit float IMU data + 8 bit for ID + 8 bit for the number of data points in one



**Figure 4: Sensor node placement on participant.** For our validation study, we used 11 HARNodes, positioned on the feet (1, 2), shins (3, 4), thighs (5, 6), waist (7), head (8), chest (9), and both wrists (10, 11).

packet), which is approximately 6.06 kB/s or  $\approx 48.48$  kbit/s. Even with very conservative estimates of only 5 Mbit/s effective data rate for a 2.4 GHz 802.11 n network (in theory, the maximum possible minimal data rate is 72.2 Mbit/s, up to 600 Mbit/s [1]) would easily enable more than 100 nodes:  $\frac{5 \text{ Mbit/s}}{48.48 \text{ kbit/s}} \approx 103$  nodes.

The Wi-Fi network credentials must be entered in the firmware before flashing. After turning on HARNode, it automatically connects to the entered Wi-Fi network and starts streaming. The server must be connected to the same Wi-Fi network to receive the sensor streams.

<sup>6</sup><https://platformio.org>

## 4 Field Setup and Deployment

A key feature of our system is the standardised deployment procedure that enables a fast, error-minimised setup in field environments. Before a study session, each sensor node is charged and flashed with the appropriate firmware. The experimenter begins by powering on all HARNodes; each device will show a default identifier on its display. The following steps are then performed to instrument the participant:

- (1) **Assigning position via the display:** Taking one node at a time, the researcher first scrolls through the numbers of body positions using the built-in button (see Figure 4). The chosen label appears on the screen, confirming that this node is intended for that location. Then we cycle to the specific position on the body (front, back, left, right (seen from the front)), which currently makes a total of 44 possible positions.
- (2) **Attach node to participant:** The node is placed at the corresponding body location and secured by tightening its Velcro strap. Thanks to the one-size-fits-most strap design, this process only takes a few seconds per sensor and does not require precise alignment or calibration fixtures - the goal is to place the IMU roughly over the intended body segment. The display helps to ensure correct orientation: as long as the text on the screen is the right way up, the sensor node is correctly placed.
- (3) **Repeat for all sensors:** Steps 1 - 2 are repeated for each sensor node until the participant is wearing the full set of required IMUs (in our example case: 11 nodes, covering major limbs and torso). The integrated displays help to quickly verify that every intended body location has a sensor and each sensor is correctly identified. It also shows Wi-Fi and server connectivity.
- (4) **Initialise recording:** Once all nodes are in place, the researcher uses the server's web interface to start the synchronised recording. At this moment, a broadcast command is used to signal the nodes to note the start of an official trial. All sensor data streams are timestamped against the NTP-synchronised clocks as they arrive at the server. The GUI can also initiate a simultaneous video recording if needed. The participant can then perform the experimental tasks while the system logs all the data in the background.

Using this procedure, the entire setup (from donning sensors to beginning data capture) can be completed in just a few minutes, even with a large number of sensors. The protocol is highly repeatable and does not require special technical skills, which facilitates consistent deployment across different operators or study sites. By minimising setup time and complexity, our system lowers the barrier to conduct multi-sensor field studies in realistic environments, enabling researchers to gather rich motion data with ease and reliability.

## 5 Validation Study

To validate the capabilities of the HARNodes and to deliver a contribution to systematically finding relevant positions for IMUs in exoskeletons, we conducted a study to train a machine learning

model to precisely detect the moment when somebody walks towards stairs, before their feet even touch the first step.

### 5.1 Setup

Ten participants (mean weight  $74.55 \text{ kg} \pm 22.46$ ; mean height  $1.75 \text{ m} \pm 0.10$ ; 10 right-footed, 1 left-footed) each wore 11 HARNodes and completed 10 repetitions of walking towards stair ascent, 10 towards descent, and 2 minutes of level walking. The distance from start to the stairs was 2 metres. HARNodes were strategically positioned on key body locations to capture comprehensive movement data (on both feet, shins, thighs and arms, as well as on the waist, chest and head), as shown in 4. IMU data were sampled at 166.67 Hz and streamed via Wi-Fi with 30 data points per UDP packet, per HARNode, towards a laptop (MacBook Pro M3) which ran the server software. Each stair-walking activity was recorded on video for manual labelling.

### 5.2 Pre-Processing

The raw sensor data of all 11 HARNodes were aligned and interpolated using the precise, NTP-based, timestamps of each node, and then normalised across all sensors and participants. Each session was labelled either "*walking*" or "*towards stairs*", based on the video recordings as ground truth. We define the start of stair-approach intervals as the moment when both legs and feet are aligned in a parallel position, immediately prior to the leading leg initiating the final step before stair ascent; the end of the stair-approach interval as the moment immediately preceding foot placement on the first step in stair ascent, or just before the leading foot moves below the stance foot to contact the step in stair descent.

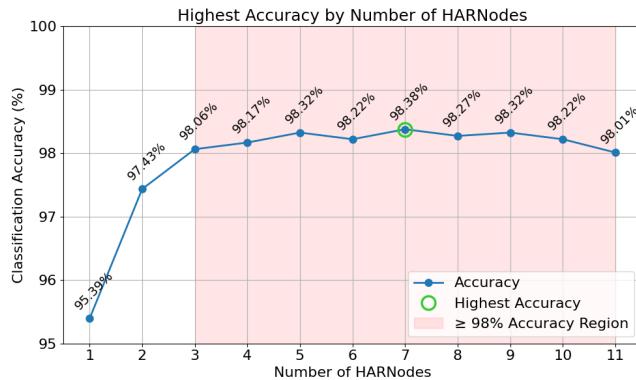
The continuous data streams were then segmented into fixed-length windows of 25 samples with 75% overlap to facilitate feature extraction. From these segmented windows, the mean, standard deviation, minimum, maximum, range, median, kurtosis and skew were extracted as features for each sensor and axis (ax, ay, az, gx, gy, gz, mx, my, mz; a: accelerometer, m: magnetometer, g: gyroscope). Left- or right-footedness was also included as a feature for each participant. Balancing of class distribution in the dataset was always performed before model training, by randomly sampling from the under-represent class. This allows us to compare accuracy as a metric, because it is not skewed.

### 5.3 Classification

For the classification task, a random forest machine learning model was used, applying a 70/30 train/test split. Random forests deliver good performance even for our limited amount of training data, which would make it unfeasible to use state-of-the-art deep-learning models such as CNNs and Transformers [7].

## 6 Evaluation

To determine both the number and the placement of sensor nodes required for good classification performance, we employed sensor overprovisioning and trained random forest models for binary classification ("*walking*" and "*towards stairs*") across all possible combinations of sensor positions ( $2^{11} = 2048$  models). For every amount of sensor (1-11), the model with the highest accuracy was calculated, depicted in Figure 5.



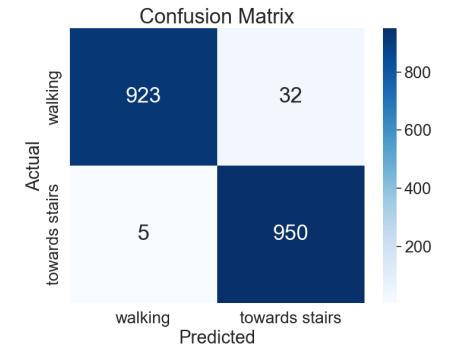
**Figure 5: Plot showing number of HARNodes for the classification task vs highest classification accuracy of any combination of n sensors. Using 3-11 nodes results in nearly the same accuracy of  $\approx 98\%$ .**

The best model with all 11 sensor nodes achieves 98.01% accuracy (Leave-one-subject-out cross-validation mean accuracy:  $90.99\% \pm 5.62$ ), while the highest overall accuracy of 98.38% (LOOCV mean accuracy:  $90.53\% \pm 3.74$ ) is achieved with only 7 nodes. Even only using three sensor nodes (*left foot, left wrist and waist*), results in 98.06% accuracy (LOOCV mean accuracy:  $88.74\% \pm 5.62$ ), surpassing 11 sensor nodes. The detailed precision and recall values for all classifications can be seen in Table 3.

**Table 3: Classification accuracy, as well as precision and recall for the binary classification *walking* vs *walking towards stairs*; Acc. = Accuracy; Prec. = Precision; Rec. = Recall.**

# Sensors	Acc. (%)	Walking (%)		Towards Stairs (%)	
		Prec.	Rec.	Prec.	Rec.
1	95.39	98.33	92.36	92.79	98.43
2	97.43	98.81	96.02	96.13	98.85
3	98.06	99.25	96.86	96.93	99.27
4	98.17	99.36	96.96	97.03	99.37
5	98.32	99.46	97.17	97.24	99.48
6	98.22	99.36	97.07	97.13	99.37
7	98.38	99.46	97.28	97.34	99.48
8	98.27	99.57	96.96	97.04	99.58
9	98.32	99.57	97.07	97.14	99.58
10	98.22	99.68	96.75	96.85	99.69
11	98.01	99.67	96.34	96.45	99.69

This clearly shows that only using a small subset of all sensors achieves comparable performance, proving the importance of screening for overprovisioning of sensors when building a human activity recognition system. When trying to find the best location for IMUs this is especially useful, for example, in exoskeletons. Not only is the performance important, but also the specific location and amount of IMU's required. We therefore also looked at five specific body positions with practical interest, with a maximum of 4 sensor nodes and computed their classification accuracies for the stair detection task, shown in Table 4. For example, a sensor



**Figure 6: Confusion matrix for binary classification (*walking* vs *towards stairs*) of the model with the highest classification accuracy (7 HARNodes).**

node on the wrist is similar to a watch and socially acceptable; feet, legs, and waist are already covered by the exoskeleton, and a sensor on the head could be hidden inside a hat, for instance. The results show that with decreasing numbers of sensor nodes on the body, the accuracy drops, but not significantly. Only when relying on only the foot as single sensor placement, we see a sharper decline in accuracy.

**Table 4: Five HARNode locations of interest, due to their practicality and their respective classification accuracies.**

Sensor Locations	Accuracy (%)
right foot, right thigh, right shin, waist	97.07
right foot, right thigh, right wrist	96.34
right foot, right thigh	95.39
right foot, right wrist, head	96.49
right foot	90.26

## 7 Conclusion

The HARNode system addresses key challenges in the collection of field data for human activity recognition (HAR), in particular the need for high node synchronisation (approx. 1 ms) and the easy handling of a large number of sensors. Based on open-source hardware (ESP32-S3, Bosch BMX160 IMU), it enables quick setup in the field and streaming of data via Wi-Fi. A proof-of-concept study with 10 participants and 11 HARNodes validated the system for recognising the transition to stair climbing and demonstrated the practical value of rapid sensor setup and sensor overprovisioning (11 sensors = 3 selected sensors,  $\approx 98\%$  classification accuracy).

Future endeavours include developing a revised body attachment model with more attachment points, automating the calculation of optimal sensor combinations, providing more diverse attachment options beyond Velcro straps, further minimising setup time, and integrating online annotation features during field data collection. All resources about the HARNode system are publicly available under the MIT licence on <https://github.com/teco-kit/HARNode>.

## Acknowledgements

This work was funded by the Carl Zeiss Foundation through the JuBot project.

## References

- [1] 2009. IEEE Standard for Information technology— Local and metropolitan area networks— Specific requirements— Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput. *IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009)* (2009), 1–565. doi:10.1109/IEEESTD.2009.5307322
- [2] 2025. Movisens GmbH. <https://www.movisens.com/>. Accessed: May 21, 2025.
- [3] 2025. MVN Awinda. <https://www.movella.com/products/motion-capture/xsens-mvn-awinda>. Accessed: May 21, 2025.
- [4] 2025. Noraxon: Innovative Biomechanics Tools for Research & Clinical Use. <https://www.noraxon.com/>. Accessed: May 21, 2025.
- [5] 2025. The Opal: a research-grade wearable sensor built for ultimate control. <https://biofitechcs.com/apdm-wearable>. Accessed: May 21, 2025.
- [6] 2025. Wearable and Real-time Kinematics Estimates with OpenSense. <https://opensimconfluence.atlassian.net/wiki/spaces/OpenSim/pages/53084280/Wearable+and+Real-time+Kinematics+Estimates+with+OpenSense>. Accessed: May 21, 2025.
- [7] Laith Alzubaidi, Jinshuai Bai, Aiman Al-Sabaawi, Jose Santamaría, Ahmed Shihab Albarhi, Bashar Sami Nayyef Al-Dabbagh, Mohammed A Fadhel, Mohamed Manoufali, Jinglan Zhang, Ali H Al-Timemy, et al. 2023. A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications. *Journal of Big Data* 10, 1 (2023), 46.
- [8] Apple Inc. 2024. *How to set up a Personal Hotspot on your iPhone or iPad*. <https://support.apple.com/en-us/111785> Accessed: 2025-05-21.
- [9] Giorgio Biagetti, Michele Sulis, Laura Falaschetti, and Paolo Crippa. 2025. High-Accuracy Clock Synchronization in Low-Power Wireless sEMG Sensors. *Sensors* 25, 3 (2025), 756. doi:10.3390/s25030756
- [10] Chiara Contoli, Valerio Freschi, and Emanuele Lattanzi. 2024. Energy-aware human activity recognition for wearable devices: A comprehensive review. *Pervasive and Mobile Computing* (2024), 101976.
- [11] Anis Davoudi, Mamoun T. Mardini, Parisa Rashidi, and Todd Mainini. 2021. The Effect of Sensor Placement and Number on Physical Activity Recognition and Energy Expenditure Estimation in Older Adults. *JMIR mHealth and uHealth* 9, 5 (2021), e23681. doi:10.2196/23681
- [12] Lars Eggert, Godred Fairhurst, and Greg Shepherd. 2017. UDP Usage Guidelines. <https://datatracker.ietf.org/doc/html/rfc8085>. RFC 8085.
- [13] Espressif Systems. 2021. *ESP32-S3 Series Datasheet*. [https://www.espressif.com/sites/default/files/documentation/esp32-s3\\_datasheet\\_en.pdf](https://www.espressif.com/sites/default/files/documentation/esp32-s3_datasheet_en.pdf) Accessed: 2025-05-26.
- [14] Manuel Gil-Martín, Javier López-Iniesta, Fernando Fernández-Martínez, and Rubén San-Segundo. 2023. Reducing the impact of sensor orientation variability in human activity recognition using a consistent reference system. *Sensors* 23, 13 (2023), 5845.
- [15] Zhiqing Hong, Zelong Li, Shuxin Zhong, Wenjun Lyu, Haotian Wang, Yi Ding, Tian He, and Desheng Zhang. 2024. Crosshar: Generalizing cross-dataset human activity recognition via hierarchical self-supervised pretraining. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 2 (2024), 1–26.
- [16] Zawar Hussain, Michael Sheng, and Wei Emma Zhang. 2019. Different approaches for human activity recognition: A survey. *arXiv preprint arXiv:1906.05074* (2019).
- [17] Akim Kapsalyamov, Prashant K Jamwal, Shahid Hussain, and Mergen H Ghayesh. 2019. State of the art lower limb robotic exoskeletons for elderly assistance. *IEEE Access* 7 (2019), 95075–95086.
- [18] Ahsan Raza Khan, Habib Ullah Manzoor, and Ahmed Zoha. 2022. A Privacy and Energy-Aware Federated Framework for Human Activity Recognition. *Sensors* 23, 23 (2022), 9339. doi:10.3390/s23239339
- [19] Kai Kunze and Paul Lukowicz. 2014. Sensor placement variations in wearable activity recognition. *IEEE Pervasive Computing* 13, 4 (2014), 32–41.
- [20] David L Mills. 2002. Internet time synchronization: the network time protocol. *IEEE Transactions on Communications* 39, 10 (2002), 1482–1493.
- [21] Luís Moreira, Joana Figueiredo, João Cerqueira, and Cristina P Santos. 2022. A review on locomotion mode recognition and prediction when using active orthoses and exoskeletons. *Sensors* 22, 19 (2022), 7109.
- [22] Jianyuan Ni, Hao Tang, Syed T. Haque, Yan Yan, and Anne H. H. Ngu. 2024. A Survey on Multimodal Wearable Sensor-Based Human Action Recognition. *arXiv e-prints* (2024). arXiv:2404.15349
- [23] Mattia Pesenti, Giovanni Invernizzi, Julie Mazzella, Marco Boccilone, Alessandra Pedrocchi, and Marta Gandolla. 2023. IMU-based human activity recognition and payload classification for low-back exoskeletons. *Scientific Reports* 13, 1 (2023), 1184.
- [24] Patrick Slade, Ayman Habib, Jennifer L Hicks, and Scott L Delp. 2021. An open-source and wearable system for measuring 3D human motion in real-time. *IEEE Transactions on Biomedical Engineering* 69, 2 (2021), 678–688.
- [25] Matteo Sposito, Tommaso Poliero, Christian Di Natali, Marianna Semprini, Giacinto Barresi, Matteo Laffranchi, Darwin Gordon Caldwell, Lorenzo De Micheli, and Jesús Ortiz. 2022. Exoskeletons in elderly healthcare. In *Internet of Things for Human-Centered Design: Application to Elderly Healthcare*. Springer, 353–374.
- [26] Akihisa Tsukamoto, Naoto Yoshida, Tomoko Yonezawa, Kenji Mase, and Yu Enokibori. 2023. Where Are the Best Positions of IMU Sensors for HAR?—Approach by a Garment Device with Fine-Grained Grid IMUs. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*, 445–450.
- [27] <https://www.bioxbgroup.dk/products-biox-armband/>. 2020. BioXBands. Accessed: May 21, 2025.
- [28] Clint Zeagler. 2017. Where to wear it: functional, technical, and social considerations in on-body location for wearable technology 20 years of designing for wearability. In *Proceedings of the 2017 ACM International Symposium on Wearable Computers*. 150–157.
- [29] Shibo Zhang, Yaxuan Li, Shen Zhang, Farzad Shahabi, Stephen Xia, Yu Deng, and Nabil Alshurafa. 2022. Deep learning in human activity recognition with wearable sensors: A review on advances. *Sensors* 22, 4 (2022), 1476.