Addressing Privacy and Cost Challenges in Remote Patient Monitoring with Streamlined 60GHz Radar and Edge Processing

Kevin Chetty[a], Paul Brennan[c], Chong Tang[b], Lai Bun Lok[c], and Fangzhan Shi[a]

[a]Department of Security and Crime Science, University College London, London, UK
[b]School of Electronics and Computer Science, University of Southampton, Southampton, UK
[c]Department of Electronic and Electrical Engineering, University College London, London, UK

ABSTRACT

Remote patient monitoring is critical in elderly care, particularly for detecting incidents like falls and abnormal gait, but current systems face high deployment costs, privacy concerns, and the complexities of continuous data processing. To overcome these challenges, we developed a 60GHz millimetre-wave radar system that processes data on-device, eliminating the need for constant internet access and mitigating privacy risks. This system is optimized for real-time healthcare monitoring in residential and clinical environments. Data captured by the radar are processed on an ATmega328 microcontroller using a quantized Long Short-Term Memory model. The quantization ensures efficient operation under tight resource constraints, enabling accurate classification of movement patterns. The system achieves an inference latency of approximately 300µs, suitable for real-time response. A key innovation is our global confidence mechanism, which reduces false alarms by aggregating predictions over multiple detection frames. This significantly improves reliability in detecting critical events like falls, reducing false positives. The system was tested on five distinct activities: falls, normal walking, irregular walking, standing, and painful standing, achieving over 91% accuracy. Compared to conventional solutions, it provides a cost-effective, privacy-preserving alternative suitable for scalable deployment across healthcare settings. By leveraging on-device machine learning, our approach reduces computational demands and enhances real-time performance without relying on cloud-based processing.

Keywords: Remote Patient Monitoring, 60GHz mm-wave Radar, Real-time Edge Processing, Privacy-preserving Healthcare, Embedded Machine Learning

1. INTRODUCTION

Wireless sensing is a promising solution for remote patient monitoring (RPM) due to its non-invasive nature and strong privacy protection. Unlike camera-based systems that capture identifiable images, millimeter-wave (mm-wave) radar leverages micro-Doppler effects to detect subtle motion variations, enabling the tracking of critical activities such as falls and abnormal gait.[1–3] This makes radar suitable for comprehensive monitoring in both residential and clinical environments. However, many existing RPM systems rely on cloud-based infrastructure for complex data processing, which increases costs, introduces latency, and poses privacy risks.

Embedded machine learning (ML) offers an effective alternative by enabling local data processing on resource-constrained devices.[4] This approach preserves privacy by keeping data on-site, reduces network dependency, and supports real-time responses to critical events. In this work, we present an RPM system that integrates 60GHz mm-wave radar with embedded ML for real-time activity detection. The radar captures motion data, which is processed on-device using a quantized Long Short-Term Memory (LSTM) model.[5] Experimental evaluation shows that our system achieves over 91% activity detection accuracy with minimal response time, offering a scalable, cost-effective solution for healthcare monitoring in both residential and clinical settings. The main contributions of our system include:

Further author information: (Send correspondence to Kevin Chetty, Chong Tang)
Kevin Chetty: E-mail: k.chetty@ucl.ac.uk, Telephone: +44 (0)20 3108 3169
Chong Tang: E-mail: chong.tang@soton.ac.uk, Telephone: +44 (0)77 5913 0383

Radar Sensor Technology XXIX, edited by Abigail S. Hedden, Gregory J. Mazzaro,
Proc. of SPIE Vol. 13471, 134710H · 2025 Published by SPIE
0277-786X · doi: 10.1117/12.3052476

- **Privacy-preserving radar sensing**: Based on a 60GHz radar, our design detects human activities within 0.6m–11m without visual or audio data, solving privacy concerns and enabling deployment in sensitive healthcare settings.
- **Efficient embedded machine learning**: Our system incorporates a quantized LSTM model optimized for deployment on an ATmega328 microcontroller with 300µs inference latency, overcoming resource constraints and enabling real-time responses.
- **Reliable event detection through confidence aggregation**: To reduce false alarms, we implemented a global confidence mechanism that aggregates predictions across multiple frames, improving the accuracy of critical event detection in noisy environments.

2. SYSTEM ARCHITECTURE AND INTEGRATION

This section presents the hardware–software co-design of our system, which aims to preserve privacy while operating efficiently on resource-constrained devices. The design combines a 60GHz radar, an ATmega328 microcontroller and a compact ML pipeline. We first describe the radar module and its signal processing, then explain how edge-based inference and output functionalities are organized, and finally discuss power management and communication mechanisms.

2.1 Radar Sensing Module

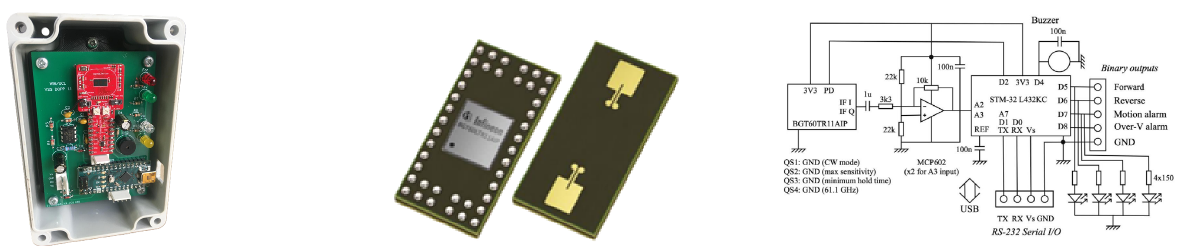

(a) The 60GHz mm-wave radar (b) The BGT60TR11AIP MMIC (c) The mm-wave radar circuit

Figure 1: Radar Sensing Module and Design

The 60GHz radar is built around the BGT60TR11AIP monolithic microwave integrated circuit (MMIC) manufactured by Infineon Technologies (Figure 1a). It features a low-power Doppler radar motion sensor that operates in the 60 GHz ISM band. It consists of two integrated antennas with a centre-to-centre separation of 4.7 mm (approximately one wavelength at 60 GHz): one transmitter (TX) and one receiver (RX) antenna each with 80° half-power beamwidth. It offers the radar designer to flexibility to adjust critical operational parameters that include detection sensitivity, hold time and frequency of operation. The dimensions of the chip are 3.3 x 6.7 x 0.56 mm, which is illustrated in Figure 1b.

This setup allows reliable human activity detection even if a subject is partially occluded by furniture or other objects. The raw radar signals undergo Doppler processing (Doppler-FFT) to extract velocity and amplitude.

2.2 Edge Processing and Output

Data from the radar is processed on an ATmega328 microcontroller that supports a real-time LSTM-based classifier through TensorFlow Lite for Microcontrollers. Before entering the neural network, measurements are screened for invalid samples and normalized to the interval (-1, 1) using empirically derived mean and standard deviation values (0.26m/s and 0.18, respectively). These readings are then segmented into 50ms windows (25 samples at 500Hz), each representing a short temporal snapshot suitable for activity classification.

The quantized LSTM includes 32 hidden units followed by a small fully connected layer for final predictions. Post-training quantization converts both weights and activations to 8-bit integers, significantly reducing the overall model size. At each inference cycle, the LSTM's internal states are reinitialized to prevent gradual drift. Efficient matrix multiplication routines leverage the microcontroller's hardware multiplier and precomputed

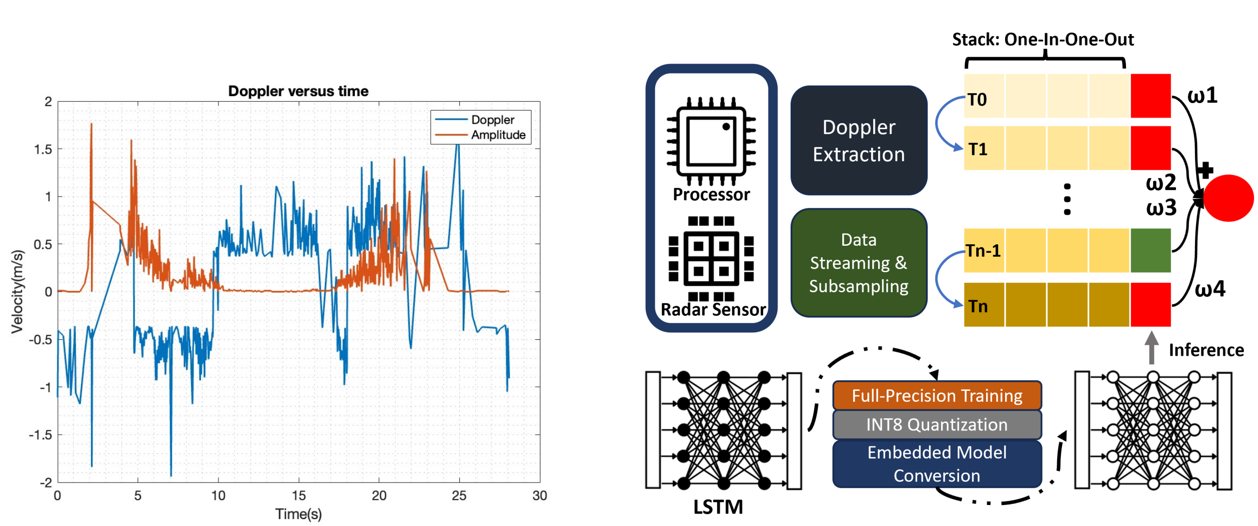

(a) (b)

Figure 2: (a) Measured Doppler velocity and amplitude for a subject walking; (b) Embedded Machine Learning processing pipeline for classifying human activities in real-time

lookup tables for activation functions. This approach achieves an inference latency of approximately 300µs per 50ms window. To mitigate brief anomalies that might mimic falls, the system uses a one-second sliding window (20 frames) and triggers an alert only if at least 16 of these frames (above 80%) indicate a fall. The system streams output data, including timestamps, activity labels, and confidence scores, over TCP or UDP connections. The data format and configuration can be adjusted through JSON files, enabling integration with remote monitoring interfaces. Figure 2 presents example radar outputs and our processing pipeline.

2.3 Power Management and Communication

The system's energy efficiency rests on carefully coordinated duty-cycling of the radar and vigilant use of sleep modes on the microcontroller. By activating the radar only 10% of the time when no motion is detected, the overall power budget decreases significantly. Between inference windows, the microcontroller transitions into low-power sleep modes, drawing minimal current and awakening when a hardware timer interrupt signals the next processing interval. Measurements confirm that a 2000mAh lithium polymer battery can power continuous operation for up to 48 hours, meeting the demands of monitoring scenarios.

Inter-module communication uses USB and TCP/UDP interfaces. Precise timing is maintained by routing the radar's frame synchronization signal to the ATmega328's external interrupt pin, ensuring that data acquisition, inference and notification routines are consistently aligned.

3. EXPERIMENT AND RESULTS

This section presents the evaluation of our system's performance in detecting critical activities. We designed experiments to test classification accuracy across multiple activity types and analyzed latency and reliability in real-time healthcare monitoring scenarios.

3.1 Experimental Setup

The experiments were conducted in a controlled environment to simulate residential and clinical settings. We collected data for five distinct activities: *falls*, *normal walking*, *irregular walking*, *standing*, and *painful standing*. The radar was positioned at a height of 1.5 meters, with 4 subjects performing each activity within 0.6m to 3m. The raw radar data was normalized and segmented into 50ms windows before being fed into a quantized LSTM network.

3.2 Performance Metrics

We evaluated the system using classification accuracy, precision and recall as primary metrics. The model was trained on a dataset comprising 500 labelled samples per activity, with an 80/20 train-test split.

Table 1: Classification Performance Summary

Activity	**Precision (%)**	**Recall (%)**	**Accuracy (%)**
Fall Detection	94.1	92.8	93.5
Normal Walking	91.5	90.3	90.9
Irregular Walking	89.4	88.0	88.7
Standing	92.0	91.7	91.8
Painful Standing	90.8	90.1	90.4

3.3 Discussion and Conclusion

The experimental results demonstrate that our system achieves high classification accuracy across all tested activities, with an average accuracy of 91.1%. Fall detection showed the highest accuracy, validating the system's capability to prioritize critical safety events. The quantized LSTM model effectively handles noisy data by leveraging a confidence aggregation mechanism, which reduces false alarms in real-time operation.

The system's inference latency of approximately 160 to 300µs per window enables rapid response to detected events. This latency, combined with low power consumption and privacy-preserving design, makes the solution suitable for continuous healthcare monitoring. Unlike cloud-based alternatives, our approach reduces costs and network dependency while maintaining robust performance in noisy environments. Future work will explore expanding the dataset to include more activity types and real-world testing in diverse healthcare settings to further validate and enhance system performance.

REFERENCES

[1] Li, W., Bocus, M. J., Tang, C., Vishwakarma, S., Piechocki, R. J., Woodbridge, K., and Chetty, K., "A taxonomy of wifi sensing: Csi vs passive wifi radar," in [*2020 IEEE Globecom Workshops (GC Wkshps*], 1–6, IEEE (2020).

[2] Liu, J., Liu, H., Chen, Y., Wang, Y., and Wang, C., "Wireless sensing for human activity: A survey," *IEEE Communications Surveys & Tutorials* **22**(3), 1629–1645 (2019).

[3] Shah, S. A., Abbas, H., Imran, M. A., and Abbasi, Q. H., "Rf sensing for healthcare applications," *Backscattering and RF Sensing for Future Wireless Communication* , 157–177 (2021).

[4] Biglari, A. and Tang, W., "A review of embedded machine learning based on hardware, application, and sensing scheme," *Sensors* **23**(4), 2131 (2023).

[5] Graves, A., "Long short-term memory. supervised sequence labelling with recurrent neural networks, 2012, 37–45."