EdgeAI Models for Human Activity Recognition on Low-Power Devices

Vinamra Sharma*, Danilo Pau†, José Cano*
*University of Glasgow, Scotland, UK, †STMicroelectronics, Italy

Human Activity Recognition (HAR) is a crucial technology for healthcare monitoring and smart living environments, but its widespread adoption hinges on running robust models directly on low-power devices [1]–[3]. This work addresses that challenge by applying Tiny Machine Learning (TinyML) optimization techniques to enable efficient on-device HAR without sacrificing accuracy. The research leverage model compression methods such as quantization and pruning, to drastically reduce neural network size and computational load while maintaining high recognition performance [4]. By minimizing dependence on cloud computing, our on-device approach inherently preserves user privacy and data security, which is vital for mobile and wearable health applications. We developed deep learning HAR models (using the benchmark UCI-HAR dataset) and deployed them on three representative edge platforms: the ST SensorTile.box PRO which uses an Arm-Cortex-M33 processor, a Raspberry Pi with quad-core ARM Cortex-A72 processor, and an Arduino with an 8-bit AVR microcontroller. These devices span a range from relatively powerful (Raspberry Pi with 512MB memory) to severely resource-constrained (Arduino with 1MB memory). The HAR models were optimized through 8-bit post-training quantization and other TinyML techniques to fit the tight memory and compute budgets. This cross-platform evaluation demonstrates the feasibility of run-anywhere activity recognition, and it illuminates the trade-offs between model complexity, accuracy, and energy consumption in practice.

Experimental Results: The optimized HAR models achieved substantial reductions in size and energy usage with only minimal impact on accuracy. The model footprint was compressed by roughly 4× on average, incurring only about a 4% drop in classification accuracy compared to the uncompressed models shown in the table I. While, table II shows how on the memory-limited Arduino device we fit a tiny 150KiB model that still reached about 89% activity recognition accuracy. The SensorTile.box PRO ran a 320 KiB model at 92.5% accuracy, while the Raspberry Pi could accommodate a larger 600KiB model, achieving up to 94.8% accuracy. Energy measurements further highlight the benefits of TinyML optimizations: an inference on the Arduino consumes only 26uJ, which is roughly half the per-inference energy of the Raspberry Pi (47uJ) and less than half that of the SensorTile (59uJ). Among the three platforms, the Arduino implementation was the most energy-efficient, whereas the Raspberry Pi yielded the highest accuracy, illustrating a clear accuracy–efficiency trade-off. Overall, this work demonstrates a practical on-device learning solution for HAR that bridges the gap between accuracy and efficiency on low-power hardware. By keeping all computations at the edge, our approach enables continuous and real-time activity monitoring without offloading data, thereby ensuring privacy and reliability in sensitive healthcare and ubiquitous sensing scenarios. The results show that through EdgeAI, it is possible to deploy state-of-the-art HAR models on affordable microcontroller-based devices, achieving a favorable balance between model accuracy, memory footprint, and energy consumption. This paves the way for more secure and resource-efficient machine learning in wearable health monitoring systems and other smart environment systems, making advanced HAR technology more accessible and scalable in real-world deployments.

TABLE I
Model Accuracy (%) Comparison Across Hardware. Different results per model are due to different randomizations to extract the test sets.

Hardware Platform	CNN INT8	LSTM FP32	Hybrid FP32-FP8
SensorTile.box PRO	**92.5**	91.0	88.4
Raspberry Pi	**94.8**	93.2	90.2
Arduino	**89.4**	86.4	82.3

TABLE II
Comparison of CNN Model Sizes & Parameter Counts.

Hardware Platform	Model Size KiB	Trainable Parameters
SensorTile.box PRO	320	16,560
Raspberry Pi	600	23,750
Arduino	**150**	**2,360**

Acknowledgment

This work was partially supported by the EU Project dAIEDGE (Grant Agreement Nr 101120726).

References

[1] S. Bian, M. Liu, B. Zhou, and P. Lukowicz, "The state-of-the-art sensing techniques in human activity recognition: A survey," *Sensors*, vol. 22, no. 12, 2022.
[2] S. Gupta, S. Jain, B. Roy, and A. Deb, "A TinyML approach to human activity recognition," in *Journal of Physics: Conference Series*, vol. 2273, no. 1, 2022.
[3] F. J. Gruber, "Sensor based human activity recognition using edge computing and tinyml in combination with convolutional neural networks," Ph.D. dissertation, UAS, 2023.
[4] V. Sharma, D. Pau, and J. Cano, "Efficient tiny machine learning for human activity recognition on low-power edge devices," in *RTSI*, 2024.