

The Future of AI on Your Wrist

**How Brain-Inspired Computing
Delivers High-Performance
Activity Recognition with
Ultra-Low Power**

Based on the research by V. Fra, et al., Politecnico di Torino
| Neuromorphic Computing and Engineering, 2022.



The AIoT revolution promises smarter, more personalized wearable experiences.

Smart devices are evolving from simple trackers into powerful, personalized companions. On-device AI (Edge AI) is the key to unlocking their true potential for real-time applications in:



Healthcare & Assistance

Continuous monitoring for elderly care, personalized health alerts.



Fitness & Performance

Advanced, real-time feedback for athletes and enthusiasts.

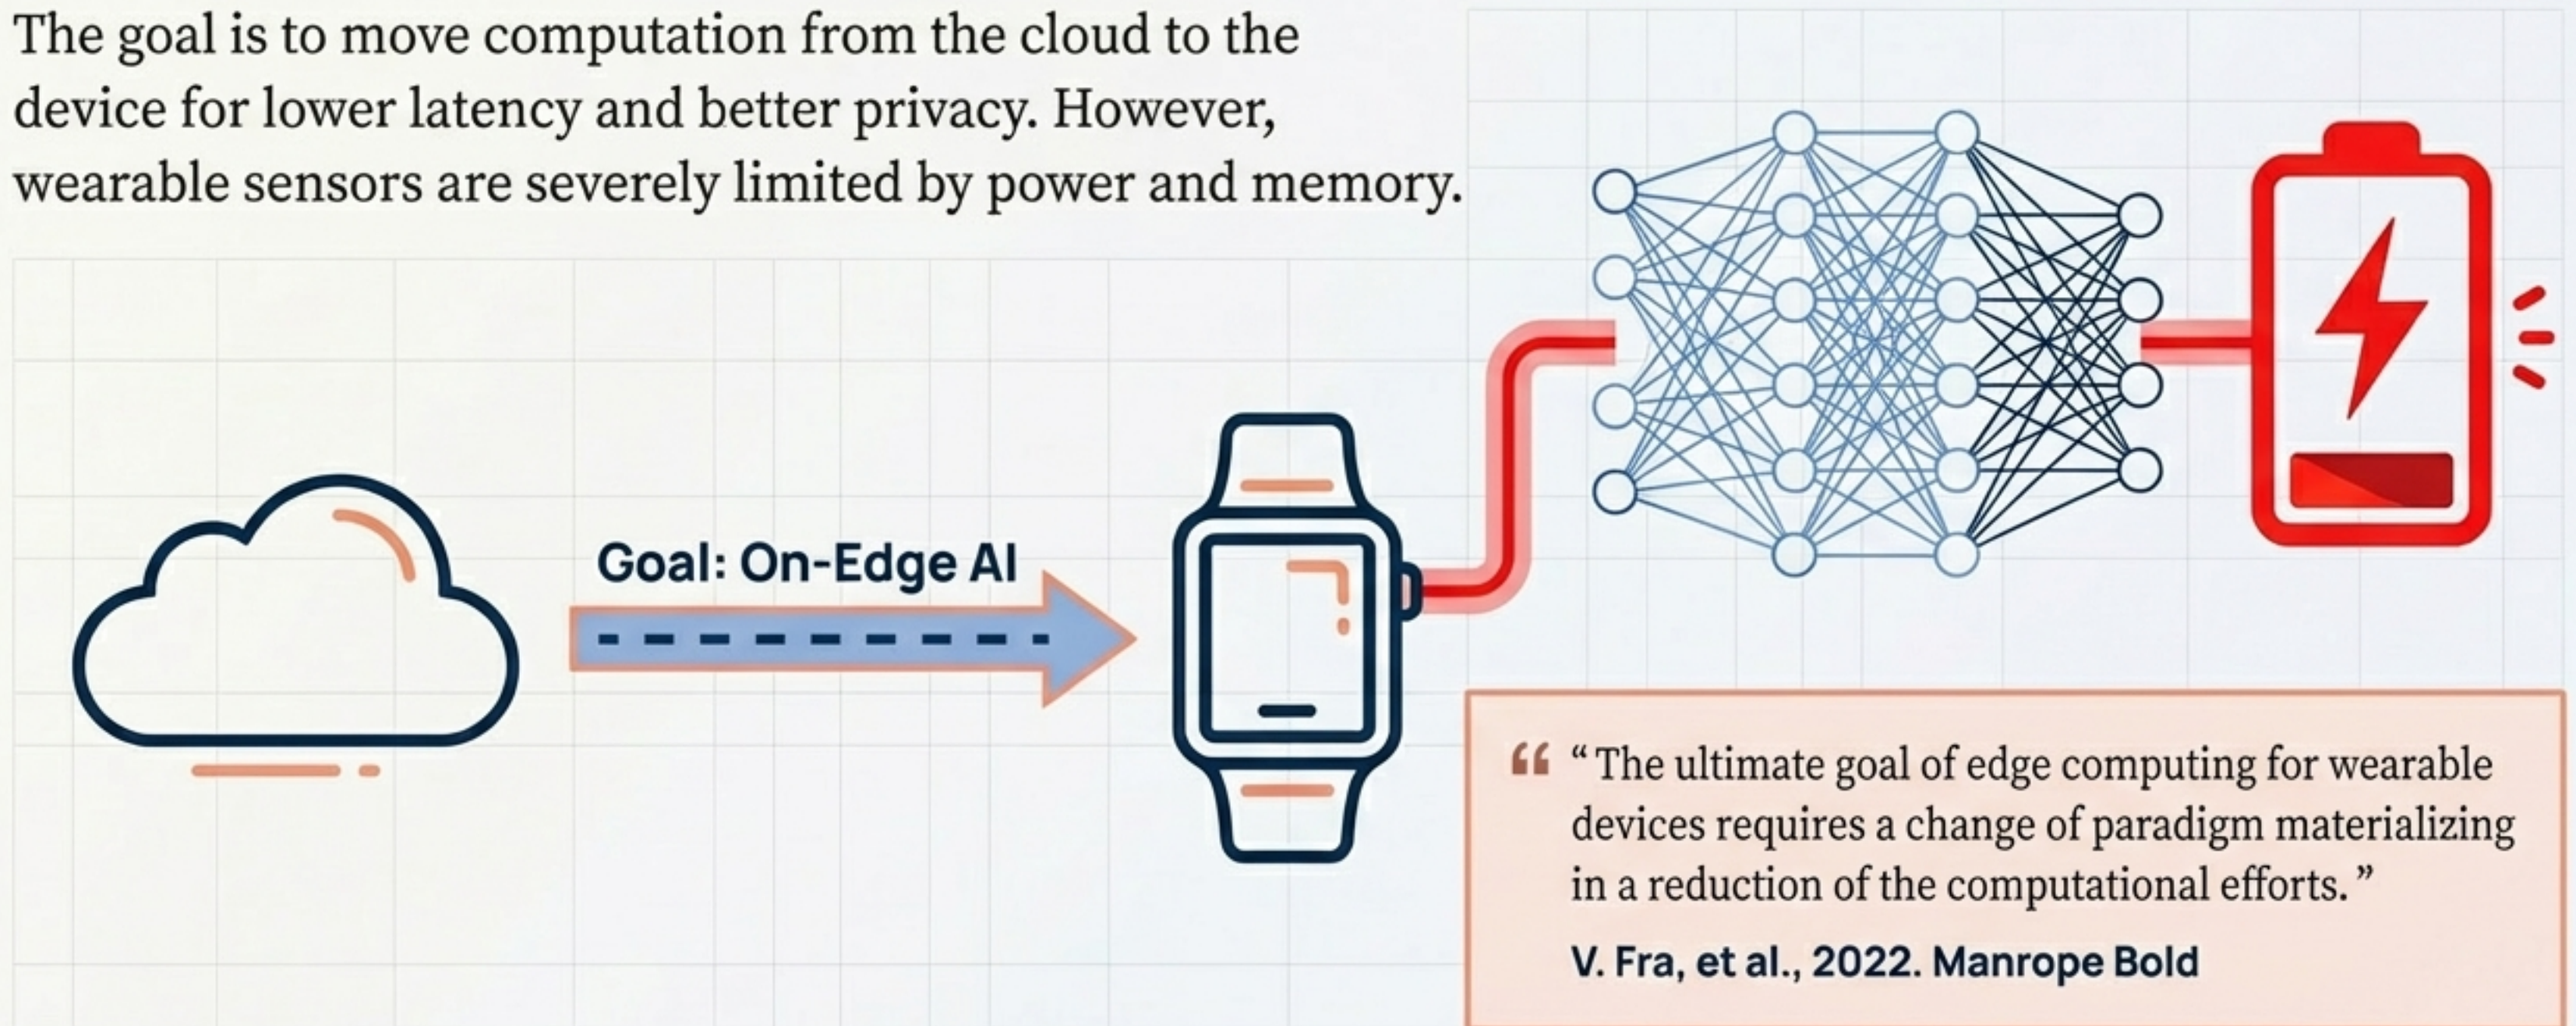


Safety & Smart Environments

Detecting falls or critical events instantly.

But the power demands of traditional AI are a critical bottleneck for edge devices.

The goal is to move computation from the cloud to the device for lower latency and better privacy. However, wearable sensors are severely limited by power and memory.




Can a neuromorphic approach deliver high-accuracy activity recognition without the energy penalty?

The Challenge

To classify complex Human Activities (HAR) using only raw sensor data from a smartwatch, a task that demands sophisticated temporal signal processing.


The Data

The investigation uses the WISDM dataset, focusing on 7 distinct, hand-oriented activities like "dribbling," "typing," and "clapping," which exhibit complex, overlapping signal patterns.




We benchmarked two classes of models: traditional deep learning vs. bio-inspired alternatives.

Traditional Deep Neural Networks (DNNs)



LSTM (Long Short-Term Memory)


A powerful recurrent architecture for time-series data.



CNN (Convolutional Neural Network)


Effective at finding patterns in spatial data, adapted for time-series.

Bio-Inspired & Spiking Networks (SNNs)



LMU (Legendre Memory Unit)


A novel recurrent cell designed to efficiently model time and memory.



sLMU & sCNN

Spiking implementations of the LMU and CNN. SNNs are event-driven and compute sparsely, mimicking how biological neurons communicate via spikes.

A rigorous optimization pipeline was used to ensure a fair and comprehensive comparison.



1. Dataset

Used raw smartwatch IMU data from the WISDM dataset, segmented into 2-second windows. No complex feature extraction.

2. Architectures

Selected both convolutional and recurrent models, in both spiking and non-spiking forms.

3. Optimization


Employed extensive hyperparameter optimization (1,000 trials per model) to find the best possible configuration for each architecture. This ensures “a comparable development effort.”

4. Evaluation

Compared the final, optimized classifiers across a range of metrics beyond just accuracy.

On pure accuracy, the traditional LSTM model sets the highest benchmark.


The LSTM network achieves the best classification accuracy at 96.42%. However, the spiking LMU (sLMU) is a very strong second, outperforming both convolutional architectures.



But the highest accuracy comes at a massive cost in model size.


The LSTM network, while accurate, is over 23 times larger than the sLMU, making it impractical for many resource-constrained edge devices. LMU-based architectures are an order of magnitude smaller.

Memory Footprint (MB)




The neuromorphic advantage: Spiking networks reduce energy consumption by up to 3 orders of magnitude.

When evaluated on specialized hardware (Intel's Loihi for SNNs, Movidius for DNNs), the energy cost per inference reveals a staggering difference. Spiking networks are fundamentally more efficient due to their sparse, event-driven processing.




The optimal trade-off is clear: sLMU delivers high accuracy with minimal resource cost.

While no single model wins on every metric, the sLMU occupies the “sweet spot.” It provides accuracy competitive with the best-in-class models while being orders of magnitude more efficient in terms of energy and memory.



A holistic view shows LMU-based and spiking models dominate on efficiency metrics

A multi-axis comparison highlights the strengths and weaknesses of each architecture. Traditional DNNs excel in accuracy but are outperformed across all metrics related to computational and memory efficiency.



This neuromorphic advantage unlocks the next generation of on-edge AI.

The demonstrated efficiency is not just an incremental improvement; it's an enabling technology. Ultra-low-power models like the sLMU make it feasible to deploy sophisticated, real-time AI on a wide range of battery-powered devices.



Always-on Intelligence

Continuous, complex activity and biometric monitoring without daily charging.



Increased Privacy & Speed

All processing happens on-device, eliminating cloud latency and data transfer.



New Applications

Enables AI in devices previously considered too small or power-constrained.

Three key takeaways from this investigation.

1

The Limit of Brute Force

Traditional deep learning models like LSTMs, while accurate, are often too resource-intensive for practical on-edge deployment in wearables.

2

A Bio-Inspired Breakthrough

Spiking Neural Networks, particularly the sLMU architecture, provide a revolutionary leap in energy and memory efficiency with only a minor trade-off in accuracy.

3

The Future is Efficient

The neuromorphic paradigm represents a crucial path forward for building truly intelligent, autonomous, and long-lasting AIoT and wearable devices.

Data and Resources for Further Exploration

Comparative Metrics for Optimized HAR Classifiers

| Model | Test accuracy (%) | Parameters | Memory (MB) | Operations (x10 ³) | Energy (μJ) |
|-----------|-------------------|------------|-------------|--------------------------------|-------------|
| LSTM | 91.5 | 102400 | 803.2 | 102400 | 15.4 |
| CNN | 95.0 | 3152 | 52.0 | 6811 | 1.0 |
| sCNN | 93.1 | 3152 | 52.0 | 6811 | 1.0 |
| LMU | 90.1 | 128 | 1.3 | 275 | 41.4 |
| LMU (ff) | 88.5 | 64 | 0.8 | 153 | 23.0 |
| sLMU | 92.1 | 128 | 1.3 | 275 | 41.4 |
| sLMU (ff) | 90.8 | 64 | 0.8 | 153 | 23.0 |

Source Code



github.com/neuromorphic-polito/NeHAR

Dataset



WISDM Dataset on UCI Machine Learning Repository

Full Paper



[doi:10.1088/2634-4386/ac4c38](https://doi.org/10.1088/2634-4386/ac4c38)