

Open-source RISC-V platforms for embedded medical grade EMG processing: Are we there yet?[☆]

Victor Kartsch^{a,*,}, Simone Benatti^{b,c}, Fiorenzo Artoni^d, Silvestro Micera^{e,f}, Luca Benini^{a,b}

^a ETH Zurich, Switzerland

^b DEL, University of Bologna, Italy

^c University of Modena e Reggio Emilia, Italy

^d Department of Basic Neurosciences, University of Genève, Switzerland

^e Bertarelli Foundation Chair in Translational Neuroengineering, Center for Neuroprosthetics and Institute of Bioengineering, School of Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

^f Department of Excellence in Robotics and AI, Scuola Superiore Sant'Anna, Italy

ARTICLE INFO

Keywords:

Biomedical embedded processing
Benchmark
EMG
EMG sensor interface
HMI applications
Parallel processing
Ultra-low-power
Wearable
RISC-V
open-source
Artificial intelligence

ABSTRACT

Open hardware, which includes both silicon and embedded hardware platforms, has fueled progress in multiple sectors. The RISC-V instruction set architecture is at the leading edge of this transformation thanks to its flexibility and vast adoption. The efficiency of RISC-V-based processors and accelerators enables artificial intelligence-based biomedical applications to operate in real-time, safeguarding user privacy while ensuring long battery life. These capabilities are critical in human-machine interaction (HMI) systems to bridge the gap between complex high-end systems and low-power, low-cost wearable hardware to provide real-time and long-term consumer interfaces with medical-grade performance. To shed more light on this topic, we examine different processing platforms in the context of electromyography (EMG)-based gesture recognition. Based on experimental EMG data, we also extend analysis and profiling to sensor interfaces, feature extraction methods, and classification algorithms to optimize accuracy, energy efficiency, and cost in various application scenarios, thereby providing a resource for researchers and practitioners interested in open-source design at the system level, enabling them to select the most appropriate solution that achieves an optimal balance between performance and complexity. Our results indicate that for applications such as gaming and consumer object control, where cost and time-to-market outweigh precision and power constraints, low-complexity algorithms (e.g., linear discriminant analysis) implemented on low-cost microcontroller-based platforms offer sufficient classification accuracy. On the other hand, medical-grade performance for continuous monitoring in wearable devices requires the next generation of ultra-low-power RISC-V architectures, capable of executing high-performance and computationally demanding algorithms. These results highlight the potential of open hardware in biomedical applications, not only for the performance of current RISC-V architectures but also for their flexibility, which allows designers to modify current integrated circuits and system designs for specific applications (without incurring licensing costs), further reducing the gap with high-end systems.

1. Introduction

Open-source hardware is emerging as a key enabler in many application domains, offering a cost-effective, and customizable alternative to proprietary solutions while fostering community-driven innovation [1]. Its impact is evident across a broad spectrum of applications, from high-performance computing systems to ultra-low-power devices, accelerating progress in research and industry [2–4].

RISC-V is at the forefront of this movement, an open-standard, royalty-free ISA that is transforming the Internet of Things (IoT), ultra-low-power computing, and edge AI. By enabling efficient local processing, RISC-V architectures reduce reliance on cloud computing, lower latency, and enhance privacy while also supporting the current trend to create self-learning systems [5]. These advantages are particularly relevant in biomedical applications, where real-time signal processing, energy efficiency, and data security are critical [6–8].

[☆] This article is part of a Special issue entitled: 'EuroMicro' published in Microprocessors and Microsystems.

* Corresponding author.

E-mail address: victor.kartsch@iis.ee.ethz.ch (V. Kartsch).

The growing demand for health monitoring fuels innovation in wearable and implantable systems, enabling disease prevention, rehabilitation, occupational performance optimization, and seamless human-machine interaction (HMI). Among the many emerging HMI technologies, electromyography (EMG)-based gesture recognition systems are gaining traction in consumer technology [9], rehabilitation and diagnostics [10], augmented and virtual reality [11,12], and robot control [13].

Besides commercial devices [14–16], open-source architectures for EMG-based HMIs have also gained traction as accessible and cost-effective solutions for medical, research, and consumer applications [17–19], opening the possibility to develop wearable systems for many application scenarios. Still, achieving medical-grade EMG performance in wearable form factor remains an open challenge across multiple HMI building blocks.

Computational power and memory remain a major constraint, as available embedded processing platforms cannot cope with the current deep learning models requirements [20,21]. Similarly, as wearable EMG devices must continuously operate throughout the day, energy efficiency is another important challenge. Achieving a proper balance requires efficient processing architectures, aggressive power management strategies, and energy-aware AI algorithms that minimize computational overhead without compromising classification accuracy. In addition, form factor and integration pose significant design challenges as well, as wearable EMG systems must be compact, lightweight, and unobtrusive, ensuring user comfort without compromising signal fidelity, requiring integrating state-of-the-art EMG sensors within a small footprint, optimizing electrode placement, and implementing robust amplification and noise filtering techniques to maintain signal quality under real-world conditions.

While many studies address individual challenges in hardware, software, and signal acquisition [22–24], there is a lack of comprehensive research that examines these components as an integrated system. A holistic analysis is essential to accurately assess computational efficiency, power consumption, signal quality, and overall system integration. Such an approach would provide valuable insights for optimizing future designs and improving the development of more efficient and well-balanced systems.

To address these challenges, this study presents a multi-parametric EMG decoding analysis, systematically evaluating the key building blocks of EMG-based HMIs. By integrating processing platforms, sensor interfaces, feature extraction methods, and classification algorithms, and experimental data acquired in the different hardware configurations, this work explores the trade-offs between computational efficiency, power consumption, and classification accuracy across different application scenarios. Specifically, we characterize the computational complexity of several AI-based gesture recognition algorithms on three embedded platforms, each of them representing distinct costs, power efficiency, and real-time performance. Additionally, the impact of sensor hardware choices on system accuracy and complexity is assessed, providing a more comprehensive perspective on wearable EMG system design.

This analysis offers a structured framework for optimizing the design of EMG-based HMIs, providing insights into how different hardware-software configurations align with application-specific constraints. By examining case studies ranging from low-cost, consumer-oriented systems to medical-grade continuous monitoring solutions, this work highlights how emerging open-source architectures, such as RISC-V-based platforms, can bridge the gap between high-performance computing and low-power embedded implementations.

The work is organized as follows: Section 2 summarizes related works in the literature, Section 3 introduces the material and methods for our analysis, Section 4 presents our experimental results and briefly describes our findings, Section 5 offers a discussion and outlines general guidelines for future deployments and finally, Section 6 report our conclusions.

2. Related work

Table 1 summarizes State-of-the-Art (SoA) approaches to combine EMG signal acquisition and online processing towards wearable-like HMIs. It is important to note that many of the most recent platforms presented in Table 1 are based on RISC-V System-on-Chips (SoCs), highlighting the significance of both the architecture and the open-source movement as key enablers of research and development. Early attempts include [25], introducing a system for the real-time classification of three EMG gestures running on a small single-board computer (based on the proprietary Cortex-A8 core). Based on a digital signal controller board, [26] presented an embedded system for gesture classification based on the Multilayer Perceptron Network. Still, these early platforms' (including [27], where only signal acquisition is performed) excessive power consumption (> 1 watt) rendered them unsuitable for wearable long-term monitoring due to the large batteries required to ensure several hours of operation.

More recent advancements in architecture design have led to highly optimized low-power systems compatible with embedded applications, which have been leveraged in several studies. By using an ARM Cortex-M4, authors in [28] provided a system for real-time processing of EMG signals while also introducing further improvements in Digital Signal Processing (DSP) and keeping a power envelope fully compatible with wearable long-term monitoring. RISC-V-based architectures [35] allowed further efficiency and computational power in recent applications, as presented in [29,36], where the authors implemented an EMG classification system on a Parallel Ultra-low Power (PULP)-based wearable architecture (2x4 cm) while providing more than 24hs of uninterrupted operation.

Together with efficient processing and battery life, EMG decoding performance constitutes another optimization parameter of EMG systems. A typical EMG decoding pipeline includes a feature extraction stage followed by a classification algorithm. Techniques such as Waveform Length (WL), Root Mean Square (RMS), and Discrete Wavelet Transform (DWT) have gained recognition for their low computational requirements and robust performance [29,37–39].

For pattern recognition, probabilistic algorithms such as Linear Discriminant Analysis (LDA) and Support Vector Machines (SVM) remain popular for embedded applications, offering classification accuracies above 80%–90% while maintaining computational efficiency compatible with resource-constrained architectures [40]. For instance, [30] employed an SVM combined with a genetic algorithm for parameter optimization to classify seven gestures, achieving an accuracy of approximately 91%. Similarly, [31] compared the performance of LDA and Artificial Neural Networks (ANN) for EMG signal classification, demonstrating that while ANN outperformed LDA, the latter still provided comparable accuracy suitable for many applications with significantly lower computational complexity.

Advanced machine learning (ML) and artificial intelligence (AI) algorithms have been explored to enhance performance and incorporate additional features for EMG classification. For instance, [32] introduced a wearable EMG gesture recognition system based on the hyperdimensional computing (HDC) paradigm, optimized for a parallel ultra-low-power (PULP) architecture. This system features a “one-shot” online learning and provides a real-time classification (8 ms of processing latency), achieving 85% accuracy for 11 gestures while consuming only $83.2 \mu\text{J}$ per inference, enabling extended operational autonomy. Moreover, [33] presented TEMPONet, a Temporal Convolutional Network (TCN) specifically designed for robust, session-independent real-time EMG classification. TEMPONet achieves a classification accuracy of 93.7% when deployed on a low-power multicore IoT processor (GAP8), performing real-time inference in under 13 ms with an energy consumption of only 0.9 mJ per classification. These advancements demonstrate the potential for integrating newer ML/AI techniques into wearable systems when combined with SoA processing architectures.

Table 1
Summary of the features of HMI systems found in literature.

	Contrib	Emb ^a	Plt ⁱ	# Gst ⁱ	Ftrs ⁱ	Csfr ⁱ	Acc	AFE	Chs	E. Type	Pwr
Zhang [25] :	App	no	Cortex-A8	3	Multiple ^d	LDA	n/s	MA-300 ^e	4	Gel-based ^f	1.3 W ^c
Duran [26] :	App	no	C2000 (TI)	5	FFT ^k	MLP ^l	n/s	TL084	n/s	Gel-based	>1 W ^h
Champaty [27] :	App	Partially	Arduino	5	n/s	Acq. Only	n/s	AD620	4 ^b	n/s	1.3 W ^c
Benatti [28] :	App	yes	Cortex-M4	7	Time-domain	HMM ^k SVM	0.9	13E200 ^g	8	Dry	29.7 mW
Kartsch [29] :	App	yes	RISC-V	11	GA ^j	HDC	0.84–0.99	ADS1298	8	Gel-based	11.8 mW
Aviles [30] :	Alg	no	n/a	7	GA	SVM	0.91	USB-6002	4	n/s	n/s
Saeed [31] :	Alg	no	n/a	4	Time/freq	ANN	0.91–0.97	n/s	6	n/s	n/a
Benatti [32] :	App+Alg	yes	RISC-V	11	RMS	HDC	0.85	ADS1298	8	Gel-based	10.4 mW
Zanghieri [33] :	App+Alg	yes	RISC-V	9	raw	TCN	0.93	ADS1298	8	Gel-based	60 mW

^a Emb=Embedded.

^b The system allows up to 8 channels. Demo app only employs 4 channels.

^c Estimated.

^d Mean absolute value, zero crossings, waveform length, and number of slope sign changes.

^e From Motion Labs Inc [34].

^f Inferred. The amplifier employed requires gel-based electrodes.

^g Coupled with the internal Analog-to-Digital Converter (ADC) of the MCU.

^h Estimated considering the power consumption of both the DSP processor and the AFE.

ⁱ Plt=Platform, Gst=Gestures, Ftrs=Features, Csfr=Classifier, Pwr=Power.

^j Genetic algorithm.

^k HMM: Hidden Markov Model.

^l MLP: Multilayer Perceptron.

Still, although the previous works provide a reference for designing HMIs, there is still a lack of information on the effects that the sensor interface parameters, such as electrode configuration (single-ended, bipolar), the type (hydrogel-based or dry), and the electrical interface to the front end (passive, active), can have on classification performance and overall hardware complexity of the HMI device. Although it is possible to find few works benchmarking some of these configurations based on SNR or electrode-to-skin impedance [41,42], it is necessary to consider that the final performance and complexity of an HMI system also depends on the algorithm, as well as the hardware and software platform.

Hence, the performance of an HMI system must be assessed comprehensively, considering every component — sensor interface, data processing algorithm, and computing platform — as each contributes significantly to the system's overall functionality. Based on experimental data, this work provides a detailed comparison matrix to evaluate classification performance that identifies the most suitable algorithms for specific electrode configurations and processing chains, ensuring optimal results. Furthermore, we examine the computational cost and power consumption of deploying these algorithms on various embedded platforms, offering valuable insights into the challenges of implementing systems designed for long-term, real-time operation. Through several application scenarios, we also demonstrate that achieving long-term monitoring with medical-grade performance necessitates next-generation RISC-V parallel Microcontroller Unit (MCU) platforms, which will enable the deployment of advanced AI-based algorithms while maintaining low-cost and energy efficiency, making them ideal for wearable applications.

3. Material and methods

The following section details each area of the EMG processing pipeline considered in this work (as summarized in Fig. 1) and introduces the metrics considered for the forthcoming analysis in the result section.

3.1. Signal acquisition setup

The signal acquisition setup is critical in ensuring the reliability and accuracy of EMG signal processing. It encompasses all hardware

and configuration choices, ranging from the electrode configuration to the analog-front-end (AFE) specifications. This section outlines the configurations and components employed in this study.

3.1.1. Electrode configuration

We focus on two main electrode configurations, namely, single-ended and bipolar. Single-end is the most straightforward electrode arrangement, where all electrodes are single-referenced. With higher hardware complexity, the bipolar configuration aims at increasing signal quality (removing common-mode signal components) by computing the signal differentially between two adjacent electrodes.

These two configurations can be presented in different *flavors* when coupled with dry/gel-based electrodes or when using an active (or passive) connection to the AFE. The active electrode circuitry employed in this work is based on [43], which includes a low-power, low-noise, rail-to-rail Operational Amplifier (AD8603 by Analog Devices). Dry electrodes are based on [44] (Sahara, g.Tec), while the gel-based electrodes are based on [45] (Covidien).

Specifically, we offer an analysis of the following combinations:

- *Wet Passive Bipolar (WPB)*: 17 gel-based passive electrodes (8 Channels x 2 Electrodes + Voltage Biasing electrode).
- *Wet Passive Single-ended (WPS)*: consisting of 10 gel-based passive electrodes (8 Channels x 1 Electrodes + Common reference electrode + Voltage Biasing electrode).
- *Wet Active Bipolar (WAB)*: 17 gel-based electrodes (8 Channels x 2 Electrodes + Voltage Biasing electrode) with active circuitry.
- *Wet Active Single-ended (WAS)*: 10 gel-based electrodes (8 Channels x 1 Electrodes + Common reference electrode + Voltage Biasing electrode) with active circuitry.
- *Dry Active Bipolar (DAB)*: 17 dry electrodes (8 Channels x 2 Electrodes + Voltage Biasing electrode) with active circuitry.
- *Dry Active Single-ended (DAS)*: 10 dry electrodes (8 Channels x 1 Electrodes + Common reference electrode + Voltage Biasing electrode) with active circuitry.

As the dry electrodes employed for this research typically require active circuitry to overcome the effects of the high electrode-to-skin impedance [41,44], the passive dry configuration has been excluded.

The analysis offered in the following sections focuses on two primary parameters to assess the electrode configuration performance: the




Fig. 1. Block diagram of the complete tests presented in this work. This study includes the effects of different electrode configurations, the DSP strategy (feature extraction and classification), and the efficiency of various embedded platforms for the execution of the various DSP chains.

Table 2
Available AFEs for HMIs.

	ADS1298	AD7768	MAX30001
IRF ^a	0.40 μ V	1.31 μ V	0.77 μ V
Channels	8	8	1
Sample rate	250–32kSPS	1k–32kSPS	125–512SPS
Power/channel ^b	0.75 mW	9.5 mW	0.3 mW
Max Gain	12x	n/a	20x
CMRR	–115 dB	–110 dB	–115 dB
Crosstalk	–127 dB	–120 dB	n/a
Resolution	24 bit	24 bit	18 bit
Price	32EUR	10EUR	6EUR
Manufacturer	Texas Inst.	Analog D.	Maxim

^a Input Referred Noise in V_{RMS} .

^b During sampling.

gesture classification accuracy and the complexity cost. The latter is calculated as follows:

$$C = N_e \times T_{setup} \times N_{failurePoints} \quad (1)$$

where C is the cost, N_e is the number of electrodes, T_{setup} is the setup time of the full electrode configuration, and $N_{failurePoints}$ is the number of failure points associated with the use of passive/active electrodes. In this work, we assume a T_{setup} of $5min/N_e$ and $1min/N_e$ for the gel-based and dry electrodes, respectively. For $N_{failurePoints}$, we assume a value of 2 (cables and contact points) and 13 (electronic components, cables and contact points) for passive and active electrodes, respectively.

3.1.2. Analog-front-end (AFE)

Although the design of the AFE is out of scope, Table 2 presents a brief comparison among several commercial and research AFEs. The ADS1298 provides the best signal quality with high versatility (eight channels, Programmable Gain Amplifier (PGA), and 250–32kSPS operational sampling rates). In terms of power consumption (<0.75mW/Channel), it is only outmatched by the MAX30001, which only provides a single channel per package. For these reasons and given the device's popularity, our experiments are based solely on ADS1298.

3.2. Digital signal processing

Digital signal processing (DSP) is essential for EMG decoding. This section outlines the preprocessing techniques, feature extraction methods, and classification algorithms used in this work.

Table 3
Preprocessing Filter Employed in this work.

	Type	Order	Cut-off
Low-pass	IIR	3	300 Hz
High-pass	IIR	5	2 Hz
Notch	IIR	3	49–51 Hz

3.2.1. Pre-processing

Preprocessing of EMGs is widely associated with filtering and aims to keep only the usable EMG bands. High-pass filters eliminate DC and low-frequency components (0–2Hz) while providing a zero mean signal. Low-pass filters remove high-frequency components that might be present as a consequence of external noise. A notch filter is also frequently used to remove the Power Line Interference (PLI) present on the signal (50 Hz/60Hz). In this work, the filters summarized in Table 3 have been employed for all tests.

3.2.2. Feature extraction

Feature selection can significantly affect the classification performance, regardless of the classification mechanism implemented [46]. Although many feature extraction mechanisms exist, as delineated in the previous sections, we focused here on the three most commonly used methods [47] as their reported performance is among the highest [48]. From the time domain, these include the Root-Mean-Square (RMS) envelope and the Waveform Length (WL) and, from the frequency domain, the Discrete Wavelet Transform (DWT).

RMS provides temporal information about the signal envelope by converting the signal's highly variable components to a stable and smooth set of points. Waveform length extracts the cumulative length of the waveform of the signal [49]. The DWT provides time–frequency information by decomposing the input signal into a bank of low-pass and high-pass filters to create a series of coefficients called “Detailed Coefficients” ($E_{D(n)}$). Fig. 2 shows the typical output after the feature extraction for these algorithms.

In this work, RMS and WL are calculated over a 60-sample window, while the DWT relies on a 4-level decomposition over a window of 256 samples. The performance of these feature extraction algorithms is directly measured through classification accuracy.

3.2.3. Classification

Our study includes three widely adopted classification algorithms for EMG classification: Linear Discriminant Analysis (LDA), Support Vector Machines (SVM), and High Dimensional Computing (HDC).




Fig. 2. Signal output after applying the different feature extraction mechanisms included in this work. After transformation, signals highlight specific features, as denoted in the examples on the right, that can potentially reduce the complexity of the classification strategy while increasing its performance.

LDA is a probabilistic algorithm that predicts a class based on the Bayesian Theorem. In LDA, each class is associated with a discriminant function that builds the feature space's decision boundaries. The implementation of the LDA in this work is based on [50].

SVM is a robust supervised statistic algorithm that separates two different classes using a hyperplane that allows the maximum separation among the various regions (or class clusters) of the feature space [51]. Given its optimal accuracy and favorable trade-off between inference complexity and performance, it has been implemented heavily for biosignal classification, especially for real-time embedded applications [52]. The SVM implementation in this work relies on [53]. For all experiments, we use the RBF (Radial Basis Function) kernel with $C=1500$ (identified as the optimal value for our experiments based on preliminary tests on smaller subsets of the data).

HDC is a brain-inspired algorithm that computes with points in the HD space (hypervectors) as an alternative to numbers [54]. The hypervectors are dense binary (pseudo)-random vectors composed of randomly placed 0s and 1s. The core of the learning is the associative memory (AM) block that contains the vectors assigned to specific classes. During inference, unseen data is encoded into an n -gram (query) hypervector and compared with all the prototype hypervectors in AM to determine the class of the point. In this work, the HDC implementation relies on [55] and it is based on 10k hypervectors.

To derive the performance results (accuracy) for all algorithms, we employed 5-fold cross-validation.

3.3. Embedded platforms

The DSP chains introduced previously have been implemented on the following architectures: ATmega-based Arduino single-board, the ARM-based STM32 MCU, and the PULP-based Mr. Wolf MCU. Table 4 summarizes the MCU platforms. Performance evaluation focuses on computational complexity and power consumption/latency. Computational complexity is defined as the number of cycles required to complete a task. The power consumption/latency is obtained considering the platform energy efficiency (power consumption at run and sleep modes) to execute a given algorithm (hence, based on the computational load) targeting a specific output latency.

3.3.1. Arduino

Arduino is a single-board microcontroller used mostly for sensor data acquisition and forwarding. However, it has also been used to implement more complex systems, even for biosignal acquisition [56]. In this work, we employed the ATmega328P-based Arduino Uno, which offers several peripherals (SPI, UART, I2C) and operates at the maximum frequency of 16 MHz. It includes 32Kb of Flash memory, 2Kb of Static RAM (SRAM), and 1Kb of Electrically Erasable Programmable Read-Only Memory (EEPROM).

Table 4

Embedded platforms included in this work.

	ISA	Max freq	Memory ^a	Power ^{b,c}
Arduino	8-bit RISC 1-core	16MHz	2kB/33kB	85mW/0.1mW
STM32	32-bit RISC 1-core	168MHz	260kB/1Mb	152mW/39mW ^c
Mr. Wolf	32-bit RISC-V 8-core	450MHz	578kB	20mW/108μW ^c

^a Random Access Memory (RAM)/Retentive memory values, otherwise only RAM.

^b Specified as Run/Sleep power consumption.

^c For: STM32@168MHz, Arduino@16MHz, and Mr. Wolf@100MHz.

3.3.2. STM32F4

The STM32F4 processor is a single core with a three-stage pipeline and branch speculation, implementing 32-bit RISC instruction set architecture. It also features a floating-point unit for fast floating-point operations, 192 KB of SRAM, and 1 MB of non-volatile flash memory. Its rich set of peripherals (SPI, I2C, UART) allows flexibility to communicate with external devices such as specialized ADC for biopotentials. It features non-maskable interrupts (NMIs) and up to 240 physical interrupts with different priority levels (8 to 256). It also offers several power modes, allowing for a reduction in the overall power consumption of the system. In this work, we evaluate the performance of the STM32F407 on [57], a previously developed HMI system, also featuring the ADS1298 for signal acquisition and a BLE module to communicate with an external device.

3.3.3. Mr. Wolf

As a representative of the more capable class of multicore MCUs, Mr. Wolf is a multi-core programmable SoC, combining a 12 Kgates RISC-V processor with a parallel cluster of eight RISC-V processors equipped with two floating-point units and flexible and powerful DSP extensions [58]. The SoC features a full set of peripherals, including SPI, I2C, I2S, UART, and many GPIOs, managed by a dedicated DMA channel (μDMA) for autonomous data transfer. The system features a single-cycle latency multi-banked L1 memory (64 kB) for fast sharing among the cluster's processors and an L2 memory for general-purpose storage (512 kB). In this work, we use BioWolf [59], an advanced biosensing device for data streaming and real-time DSP based on Mr Wolf (PULP), an ARM-SoC MCU (Nordic) for Bluetooth Low Energy (BLE) communications and the ADS1298 for bio-signal acquisition.

4. Experimental results

4.1. Dataset collection setup

EMG data are collected from 5 healthy subjects (2 male, 3 female, aged 25–35 years) and comprise six sessions, each with a different




Fig. 3. Threshold selection and label output of an EMG gesture. (A) shows the averaged EMG signal in the WL feature domain, indicating the different sections of the signal (rest, transient, gesture) to be extracted. (B) shows the density distribution of the values in (A), where the local minima represent optimal points for threshold selection. Finally, (C) shows a typical example of filtered EMG signals (Ch1 to Ch8) during an isometric contraction (gesture), where the transient and gesture labels have been identified using the thresholds derived from the local minima found in (B).

electrode configuration. During each session, subjects perform three contractions per gesture, each separated by five seconds of rest to reduce muscular fatigue. Once finished, ten seconds of rest are left. The gestures (eight + rest) included a typical set compatible with HMIs (open hand, fist, index pointing, power grip, thumb up, wrist supination, wrist pronation, horns, shaka, rest). Subjects gave their informed consent for inclusion before they participated in the study.

EMG was collected around the forearm muscles using eight equally-distributed channels, arranged in two rows for bipolar channels (positive and negative inputs, respectively) and a single row for single-ended electrodes, with reference (common negative) placed next to the first channel. The complete data acquisition relied on [57], based on the ADS1298.

The acquired data was labeled automatically based on the data points' amplitude density distribution on the WL feature domain. In this process, raw data containing the gesture is first converted to the WL domain to produce labels, as presented in Fig. 3A. Here, it is noticeable that transients have a variable amplitude in a given range (5–60 μ V) but appear less frequently in the dataset. Hence, transients have a lower density, which is reflected as a local minimum when visualizing the occurrence of the data points for different amplitudes, as indicated in Fig. 3B. The value corresponding to the local minimum is then used as a threshold to determine a gesture's presence or involuntary movements.

We excluded involuntary movements using the expected contraction duration (five seconds) as a reference. Fig. 3C presents an example of the labeled data.

4.2. Algorithm and electrode configuration performance

Fig. 4 summarizes the performance results of the different EMG classification chains and all electrode configurations. The heatmap figure (Fig. 4A) shows the mean performance (among all test subjects) of the various electrode configurations using every combination of feature and classification methods. In contrast, the figures in the left column show average performance results for the electrode configurations (Fig. 4B) and DSP chains (Fig. 4C).

The results on the overall performance of the processing chains denote that the SVM (shaded in blue) offers the best overall accuracy (91.82%), with the lowest variability among different features ($\sim 4\%$). The LDA performance (shaded in orange) is also above 90% but exhibits a more considerable sensitivity to the feature selection ($\sim 7\%$). Finally, the HDC (shaded in green) provides the lowest accuracy in the group ($\sim 84\%$), with $\sim 4\%$ variability among the different feature extraction algorithms.

This figure also provides results for the feature extraction mechanism. For instance, it is noticeable that WL is robust as it outperforms all other features for most classification algorithms studied. Specifically, its performance for the lowest and highest accuracy values (95%–86%) is the best among the group. The RMS (93%–82%) follows the same trend (slightly lower performance) for the SVM and LDA classification. DWT (92%–83%) shows its best overall performance when coupled with HDC. It is also worth noticing that DWT allows achieving the best performance for all electrode configurations but dry single-ended when combined with SVM.




Fig. 4. Performance of the electrode configurations and the algorithms implemented. The figure on the right (A) presents the performance of each DSP chain for each electrode configuration, while the figures on the left show the average performance for the electrode configurations (B) computed for all DSP chains (and subjects) and the average performance of the DSP chains (C) for all electrode configurations (and subjects). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 5. Comparison between the electrode configuration complexity and the error. The size of the circles represent the error's variance.

Results for the electrode performance show that all wet-based electrodes (both active and passive) offer the best performance (90%–91%). The dry active group delivers a slightly lower performance (84%), primarily due to the lower performance of the single-ended version. Still, the latter can achieve compliant performance with a proper processing chain (SVM+WL, with up to 92%). It is also to note that all bipolar configurations in all three groups outperform single-ended versions. The most significant differences are for the dry active group, with up to 11% of performance loss to the single-ended. The performance drop can be associated with the expected lower common-mode rejection ratio.

Although these results delineate a clear trend, they still must be contextualized to the system complexity (defined in Section 3.1.1). Such contrast is highlighted in Fig. 5, which reports the calculated costs for all the electrode configurations to the classification's error ($1 - \text{accuracy}$). The size of the data points ratio represents the error variance.

First, looking at the complexity axis, it is noticeable that single-ended configurations hold the lowest cost as they require fewer electrodes and cables. WPS has the lowest complexity from this group, followed by the active version (DAS), which, although having a shorter setup time, requires electronic components. On the other hand, WAB and WAS, both gel-based active configurations, hold the highest complexity due to the setup time and electronics needed.

Including the error in this analysis, it is noticeable that the complexity cost is the largest for configurations with the lowest error, such as WAB, although marginally. Nevertheless, others, such as WPS, WPB, and DAB, also hold low error values while keeping low complexity.

4.3. Computational load

This section focuses on the execution of the algorithms for feature extraction and classification presented in Section 3.2, implemented onto the platforms described in Section 3.3.

Table 5 shows the number of cycles required to extract the Waveform Length (window size=60), Root-Mean-Square (window size=60), and the discrete Wavelet Transform (4-levels with a 256 sample window), where for all platforms, WL is the least demanding. The table also denotes that the RMS, on average, requires more than 2x the processing cycles to the WL (highest, 6.76x). Finally, the DWT is the heaviest, demanding more from 17x to 34x of the processing cycles, depending on the platform.

Table 6 presents the number of cycles to execute the different classification algorithms. When comparing the LDA with the other algorithms, it is noteworthy that the SVM, when computed with 300 support vectors (average number of support vectors generated after training for all tests), can be from 6.5x to 12.41x more computationally intensive and HDC, from 45x to 59x (using 10k hypervectors), excluding Arduino. When coupled with PULP's parallel computing capabilities and bit manipulation instructions [55], the HDC breach can be reduced to 9x, which is in the SVM range (7.25x).

For the classification algorithms, we also extend the analysis to the memory requirements (not platform dependent), where the SVM ranges from 16KB (WL/RMS) to 64KB (DWT).¹ Closely following, the HDC requires 48KB (WL/RMS) to 78KB (DWT). Finally, the LDA poses the least requirements, ranging from 0.5KB (WL/RMS) to 5.2KB (DWT). It is important to note, however, that the SVM model size is not fixed a priori: it is proportional to the number of support vectors, which depends on the specific training data and on the chosen hyperparameters (e.g., the regularization parameter C). As a consequence, when the classifier is retrained on data from different subjects or under different operating conditions, the resulting model may require more memory

¹ For each channel, DWT produces four feature vectors.

Table 5

Cycle count for feature extraction (in KCycles)

Kernel ^b	COTS ^a		Mr. Wolf	
	STM32	Arduino	PULP-1C	PULP8C ^c
WL	0.95	21.76	0.72	0.23
RMS	6.43	38.59	1.43	1.33
DWT	30.04	517.50	24.76	3.92
Computational complexity relative to WL				
RMS	6.76x	1.77x	1.98x	6.56x
DWT	31.62x	23.78x	34.38x	17.04x

^a COTS: Components of the Shelf.^b For a single data point (single feature vector).^c Values obtained using the data-parallelism paradigm. For a complete window (i.e., including the complete feature vector space), values can be further optimized using the task-parallelism paradigm, where, for instance, the computational cost of RMS can be reduced by $\sim 5x$.**Table 6**

Cycle count per Classification Algorithm (in KCycles)

Kernel ^a	STM32	Arduino	PULP1C	PULP8C
SVM	149	1851	91	29
HDC	711	906400 ^b	632	36
LDA	12	193	14	4
Computational complexity relative to LDA				
SVM	12.41x	9.59x	6.5x	7.25x
HDC	59.25x	4696x	45.14x	9x

^a For an 8-element input feature space. When using the DWT (32-element feature vector), values can increase by a factor ranging from 2x to 6x, depending on the kernel.^b Due to memory constraints on the device, this value has been calculated through linear regression [55].

than the values reported here and, in the worst case, exceed the available memory on a resource-constrained platform. On the other hand, for both LDA and HDC, the memory required is always constant, which might be more suitable for an embedded platform when performing online/real-time training.

As it may already be evident, executing these algorithms on different platforms can result in dissimilar computational loads. Fig. 6 presents the stacked CPU load for SVM, LDA, and HDC for all architectures to amplify these effects. From here, it is evident that Mr Wolf, in both 8-core and single-core versions, allows for the smallest CPU load (~ 10 and ~ 12 to PULP1C and STM32, respectively). It is also worth noting that STM32 and Mr Wolf single-core offer a comparable CPU load (1.18x in favor of PULP1C). Nevertheless, when considering the device energy consumption, Mr Wolf significantly outperforms STM32 as its energy efficiency is around 18MFlops/mW, against the 1.2MFlops/mW for the STM32, as will be more evident in the following section. Arduino always provides the most considerable CPU load of all the platforms presented.

4.4. Power consumption, energy, and output latency

The computational load plays a significant role in energy efficiency. Nevertheless, the output latency (i.e., the time required for the DSP pipeline to provide an output) ultimately determines the system's overall power consumption. Fig. 7 presents the power consumption of the processing platforms included in this work for the execution of all DSP chains.

We first focus on Arduino, which does not allow the execution of all the algorithms studied (within the presented time frame). Still, when combined with LDA+(WL/RMS), it delivers latencies below 50 ms with a maximum power budget of about 50 mW (51 mW@25ms, 1.2 mJ/inference, or 20 mW@50ms, 1mJ/inference). Nevertheless, it is important to consider that LDA does not provide the highest accuracy.

As denoted by the line plots, using SVM+(RMS/WL) is also possible, although at a higher power consumption (~ 70 mW) and higher latency (>110 ms), with a total energy consumption of ~ 7 mJ/inference.

The STM32 instead supports the execution of all studied algorithms, even for very low latencies. Compared with Arduino when performing LDA + (WL/RMS), it allows reducing the output latency to the minimum value included in our estimation (1 ms) while employing about half of the power (~ 30 mW, requiring only 32uJ/inference). Furthermore, when executing SVM at the Arduino's minimum achievable latency (~ 120 ms), the STM32 provides a power reduction of 22x.

Similar to STM32, PULP1C allows the execution of all the algorithms. Nevertheless, the main differences reside in the average power. Such distinction is more evident when comparing the execution of the SVM+RMS (@1ms of latency), denoting a power reduction of $\sim 22x$ (265.70 mW to 11.69mW, or 265.7uJ/inference to 11.69 uJ/inference).

PULP8C follows the same trend presented for PULP1C while providing a significant boost in power efficiency due to its parallel architecture. When considering latencies above 30 ms, PULP8C can run all algorithms at a sub-milliwatt power budget, with the lowest (LDA+WL) requiring less than 200 uW (or, 4uJ/inference).

4.5. Validation

To fully validate the previously introduced results in terms of accuracy and power consumption, we implemented and tested a real-time processing pipeline using a fixed configuration. Specifically, the STM32 was used as a computing platform, and the processing chain depicted in Fig. 8 was used for the EMG decoding, which includes a pre-processing step, a feature extraction (WL) and a classification algorithm (SVM). Regarding the electrode interface, the bipolar wet configuration was selected.

The online tests consisted of two sessions, denoted as training and inference. The data resulting from the training was used to create the SVM model on a desktop computer, later exported to the microcontroller. For inference (performed in real-time with an output latency of 25 ms), an android application designed as part of this work stored raw data and classification results on the MCU. The application was also used to notify users of the type of gesture and duration of the contractions. Three subjects participated in the online testing under the same experimental conditions introduced in Section 4.1. Both accuracy ($\sim 88\%$) and power consumption (~ 25 mW) were reported within the range of our estimations.

5. Discussion

5.1. HMI decoding benchmark for embedded real-time processing

From our multi-parametric EMG decoding benchmark, the first subject of analysis focused on the performance of different electrode configurations based on the accuracy performance computed through all the proposed processing frameworks, where active gel-based electrodes showed the best overall performance. The results also show that single-ended and bipolar configurations provide similar classification accuracy when coupled with gel-based electrodes ($<1\%$ on average), demonstrating that highly complex bipolar systems (2x number of electrodes) can be dropped without significant accuracy degradation. Still, bipolar configurations are of significant importance when the signal quality is compromised, such as when using dry electrodes (due to poor mechanical contact). Here, up to 11% of accuracy improvements can be obtained.

In terms of DSP, our analysis suggests that, when classical ML approaches are used, the feature extraction providing the best overall performance is the Waveform Length, followed by the RMS. Still, it is worth noting that DWT, when coupled with a proper classification algorithm (SVM) and electrode configuration, can achieve the highest




Fig. 6. Stacked cycle count for the execution of all classification and feature extraction algorithms for different platforms.




Fig. 7. Latency vs power consumption trade-off for the execution of all DSP chains while executed on the platforms targeted on this work.

accuracy among the group (97%). In turn, following the heatmap presented in Fig. 4, it is evident that SVM provides superior performance

(dark red). Specifically, it provides 5% and 9% performance boosts to LDA and HDC, respectively.




Fig. 8. Complete hardware–software setup for system validation.

The results above are more significant when factoring in the computational power required. The WL shows a significant advantage when considering the feature extraction mechanism as the computational load is small ($\sim 2\times$ and $\sim 30\times$ to RMS and DWT). As for classification algorithms, the SVM poses itself in the middle, with the LDA being $\sim 12\times$ lighter and the HDC $\sim 4\times$ more computationally expensive.

The selection of the processing platform plays an even more critical role when considering power consumption and output latency. Here is evident that a programmable multicore architecture, such as PULP, features the lowest power consumption and the lowest latency. Employing a multicore platform, allows executing all algorithms at the lowest latency in our scale (1 ms) within a power envelope below < 11 mW (10.38uJ/inference for the heaviest workload). It is worth mentioning that the multicore platform (i.e., Mr. Wolf) is not commercially available but comes from an open-source project, thanks to which it is possible to replicate or modify the design. Furthermore, we observe that similar SoCs are emerging [60,61].

Also, Arduino, an open-source platform, fits some requirements for EMG processing. For instance, when combined with the lightest processing presented in this work (LDA+WL, 90% of accuracy), it can still provide a low classification output latency with a power envelope of around 50 mW (@25ms). Additionally, it can also process EMG signals through the SVM ($\sim 93\%$ accuracy), although at a higher power consumption and latency (~ 70 mW @110ms, employing 8.21mJ/inference).

Finally, it is essential to emphasize the significance of open-source movements in the development of new platforms for medical, industrial, and research applications. This is evident in the performance, flexibility, and availability of SoCs, such as RISC-V, and systems like Arduino, as well as those discussed in Section 2. These platforms offer designers a rapid and accessible means to advance the state of the art without the constraints of closed-source systems, thereby reducing development time and costs while enhancing the performance of HMI systems.

5.2. Application scenarios

Table 7 presents an overview of five application scenarios and their building blocks, selected based on the information provided in this manuscript.

5.2.1. Low-cost system in the absence of strict requirements for performance

For this type of application, time-to-market (TTM) and cost are the most critical parameters, where accuracy and power consumption can be relaxed to reduce the cost and complexity. Considering these conditions, Arduino might be a starting option to minimize design time as the system is already offered in a designed PCB. Arduino can be coupled with LDA+WL, with a minimum latency of about 50 ms, fitting the need for several general applications, including control of wheelchairs, home appliances, and simple gaming, while providing up to 14 h of battery life. To simplify the electrode system and reduce costs, the system can be coupled with a passive single-ended sensor interface (89% accuracy), achieving an estimated price under 50EUR.

5.2.2. Wearable application for general use

Devices in this category aim for a small footprint and a low weight, which requires designing a custom PCB board that increases costs and time to market. While targeting general applications, latency is not a significant constraint and can be used as leverage to reduce the device's power consumption or obtain higher accuracy. Hence, our selection includes the STM32 as a processing platform with SVM and WL as the DSP chain. Coupling the system with single-ended versions of wet or dry electrodes can significantly reduce system complexity and intrusiveness. With an estimated bulk price of 100EUR, this configuration will deliver from 92% accuracy with a battery life of 35 h. Some applications include wearable devices (bracelets) for remote control, including smartphones, TVs, prosthetics, gaming and fitness tracking.

5.2.3. Wearable application for EMG monitoring and prosthetics control

We assume that achieving the highest accuracy on a wearable form factor is the main goal for this category. Still, latency could be spared to boost battery life, such as long-term monitoring. For this, our analysis suggests a bipolar configuration with wet electrodes that, together with SVM and WL as the DSP chain, can deliver more than 95% accuracy. Under such conditions, while running on STM32, the device can operate up to 28hs (@50ms of latency). By employing the electrode configuration with the highest performance, this device also guarantees the best signal quality when used for data logging in a long-term monitoring study. The proposed system's features come with a high price, estimated at 200EUR. Such devices can target the fine control of prosthetic devices and long-term (LT) signal monitoring to study neuromuscular disorders such as Amyotrophic Lateral Sclerosis and Muscular Dystrophy, and gait analysis, to name a few.

5.2.4. High-performance, low-latency, and low-power HMI

For this, we target optimizing all markers simultaneously. This can be achieved currently by employing the SVM and DWT as DSP chains with wet active electrodes. At the same time, power consumption can be significantly reduced when using Mr Wolf (8C). In this configuration, the system can achieve more than 96% accuracy with a battery duration of 22hs. Some applications requiring the performance exerted by this configuration include drone control, gaming, and prosthetics.

5.2.5. High-performance embedded application with online training and inference

The previously introduced system uses the SVM to achieve the best performance. Nevertheless, as stated earlier, the SVM requires offline data analysis to produce accurate results and a model within the memory bounds of a given platform. Hence, SVM is not a valid option for a fully embedded stand-alone device. Instead, LDA and HDC are both parametric regarding the memory footprint, where the latter outstands as it allows for higher classification performance. HDC (together with DWT) can be both deployed on STM32 (up to 2 h of battery) or Mr Wolf (> 50 hs of battery) with an output latency of 50 ms while delivering up to 90% of classification accuracy when coupled with bipolar wet electrodes.

Table 7

List of application scenarios meeting specific performance targets.

	Min accuracy	Min latency	Power	E. conf	DSP chain	Platform	Cost	Battery life ^a	Application
Low TTM/Cost:	89%	50 ms	25 mW (1.26 mJ/inf)	SE+WET	LDA+WL	Arduino	<50 EUR	14 h	Wheelchairs, Home Appliances, Gaming
Wearable General User:	92%	~100 ms	11 mW (1.1 mJ/inf)	SE+WET	SVM+WL	STM32	<100 EUR	35 h	Wearable Remote Control, Gaming, Fitness Tracking
Medical:	95%	~50 ms	11 mW (0.55 mJ/inf)	BI+WET	SVM+WL	STM32	<200 EUR	28 h	Prosthetics, LT Monitoring
High-Perf:	96%	Lowest (1 ms)	16 mW (16 uJ/inf)	BI+WET	SVM+DWT	PULP8C	<500 EUR	22 h	RT-Control (Drones, Games, Advanced Prosthetics)
Stand-alone:	90%	50 ms	179 mW (8.95 mJ/inf) or 7 mW (0.35mJ/inf) ^b	BI+WET	HDC+DWT	STM32 PULP8C	<500 EUR	2 h or 50 h ^b	Online Re-trainable HMI

^a Including the power consumption of the ADS1298, estimated.^b STM32/PULP8C.

5.3. Limitations of this work

In this work, we deliberately adopted overall classification accuracy as a compact and widely used metric in EMG-based HMI studies, given the approximately balanced nature of our dataset and the focus on design-space exploration across hardware platforms, feature sets, and classifiers [62–64]. A more comprehensive evaluation, incorporating additional measures such as precision, recall, F1 score, and ROC/AUC, is likely to provide further insight into classifier behavior, particularly in scenarios involving unbalanced datasets or asymmetric costs. Integrating such a multi-metric analysis into the proposed benchmarking framework is an important direction for future work.

Beyond the choice of evaluation metrics, another important limitation of the present study is the relatively small cohort size (five subjects), which constrains the extent to which our findings can capture the full inter-subject variability encountered in real-world clinical populations. In our experiments, classifiers are trained subject-specifically with multiple repetitions per gesture and evaluated using cross-validation, resulting in high and consistent accuracies. This setup makes severe underfitting unlikely within the scope of the collected data. Nevertheless, a larger and more heterogeneous cohort may reveal additional sources of variability and could affect both absolute performance and the relative ranking of the examined algorithms and platforms. The proposed framework is therefore intended as a methodological and benchmarking tool that can be extended in future work to larger-scale studies as part of a dedicated clinical validation effort.

6. Conclusions

This work presented a complete HMI decoding benchmark for embedded real-time processing. It provided novel and detailed experimental information about sensor interface performance, feature extraction, and classification algorithms, with a detailed analysis of their computational cost and power consumption, on several embedded platforms. We later described five application scenarios sweeping through low-cost and low complexity to high-performance ultra-low-power implementations. This analysis demonstrates how devices and algorithms can be efficiently integrated for different applications. Specifically, our study shows that low-complexity algorithms (e.g., LDA) can be efficiently deployed on low-cost MCUs while achieving sufficient accuracy

for general applications. On the other hand, achieving a medical-grade performance while targeting long-term monitoring in a wearable form factor requires the new generation of advanced ultra-low-power parallel MCU platforms, allowing the execution of high-performance algorithms in real-time while still delivering a battery life (for the complete HMI) ranging from 22 to 50 h, further closing the gap between medical-grade high-end devices to low-power embedded systems. Ultimately, this study underscores the transformative potential of open-source RISC-V architectures in biomedical applications. Their flexibility, cost-effectiveness, and growing ecosystem facilitate the development of next-generation ultra-low-power wearable systems, bridging the gap between traditional high-end medical devices and embedded low-power solutions.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Victor Javier Kartsch Morinigo reports financial support was provided by Swiss National Science Foundation. Simone Benatti reports financial support was provided by PNRR - Ecosister. Victor Javier Kartsch Morinigo reports financial support was provided by Horizon Europe. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This project was supported by the Swiss National Science Foundation (Project PEDESITE) under grant agreement 193813 and the Horizon 2020 Foresight project (GA:833673). S.B. acknowledges financial support from PNRR MUR project ECS_00000033_ECOSISTER.

Data availability

Data will be made available on request.

References

- [1] Knut Blind, Mirko Böhm, Paula Grzegorzewska, Andrew Katz, Sachiko Muto, Sivan Patsch, Torben Schubert, The Impact of Open Source Software and Hardware on Technological Independence, Competitiveness, and Innovation in the EU Economy, Tech. Report, European Commission, Directorate-General for Communications Networks, Content and Technology, 2021, <http://dx.doi.org/10.2759/430161>.
- [2] Francesco Conti, Angelo Garofalo, Davide Rossi, Giuseppe Tagliavini, Luca Benini, Open-source heterogeneous SoCs for AI: The PULP platform experience, 2024, Open-source resources available at <https://pulp-platform.org/>, arXiv preprint arXiv:2412.20391.
- [3] Emad Ebeid, Martin Skriver, Kristian Husum Terkildsen, Kjeld Jensen, Ulrik Pagh Schultz, A survey of open-source UAV flight controllers and flight simulators, *Microprocess. Microsyst.* 61 (2018) 11–20.
- [4] Gerrit Niezen, Parisa Eslambolchilar, Harold Thimbleby, Open-source hardware for medical devices, *BMJ Innov.* 2 (2) (2016) 78–83, This is an Open Access article distributed under the terms of the Creative Commons Attribution (CC BY 4.0) license.
- [5] Lan Mei, Thorir Mar Ingolfsson, Cristian Cioflan, Victor Kartsch, Andrea Cossetti, Xiaying Wang, Luca Benini, An ultra-low power wearable BMI system with continual learning capabilities, *IEEE Trans. Biomed. Circuits Syst.* (2024) 1–12.
- [6] Marcello Zanghieri, Simone Benatti, Alessio Burrello, Victor Kartsch, Francesco Conti, Luca Benini, Robust real-time embedded EMG recognition framework using temporal convolutional networks on a multicore IoT processor, *IEEE Trans. Biomed. Circuits Syst.* 14 (2) (2020) 244–256.
- [7] Thorir Mar Ingolfsson, Xiaying Wang, Michael Hersche, Alessio Burrello, Lukas Cavigelli, Luca Benini, ECG-TCN: Wearable cardiac arrhythmia detection with a temporal convolutional network, in: 2021 IEEE 3rd International Conference on Artificial Intelligence Circuits and Systems, AICAS, 2021, pp. 1–4.
- [8] Sebastian Frey, Marco Guermandi, Simone Benatti, Victor Kartsch, Andrea Cossetti, Luca Benini, Biogap: a 10-core FP-capable ultra-low power IoT processor, with medical-grade AFE and BLE connectivity for wearable biosignal processing, in: 2023 IEEE International Conference on Omni-Layer Intelligent Systems, COINS, 2023, pp. 1–7.
- [9] T. Scott Saponas, Desney S. Tan, Dan Morris, Ravin Balakrishnan, Demonstrating the feasibility of using forearm electromyography for muscle-computer interfaces, CHI '08, Association for Computing Machinery, New York, NY, USA, 2008, pp. 515–524.
- [10] Bojan Milosevic, Simone Benatti, Elisabetta Farella, Design challenges for wearable EMG applications, in: Design, Automation & Test in Europe Conference & Exhibition, DATE, 2017, IEEE, 2017, pp. 1432–1437.
- [11] Thalmic MYO, 2016, <https://developerblog.myo.com/tag/emg/>.
- [12] Varun Kohli, Utkarsh Tripathi, Vinay Chamola, Bijay Kumar Rout, Salil S. Kanhere, A review on virtual reality and augmented reality use-cases of brain computer interface based applications for smart cities, *Microprocess. Microsyst.* 88 (2022) 104392.
- [13] Roberto Meattini, Simone Benatti, Umberto Scarcia, Luca Benini, Claudio Melchiorri, Experimental evaluation of a semg-based human-robot interface for human-like grasping tasks, in: 2015 IEEE International Conference on Robotics and Biomimetics, ROBIO, IEEE, 2015, pp. 1030–1035.
- [14] BTS Bioengineering, BTS bioengineering, 2025, (Accessed 24 February 2025).
- [15] Delsys Europe, Trigno EMG, 2025, (Accessed 24 February 2025).
- [16] Noraxon, sEMG systems, 2025, (Accessed 24 February 2025).
- [17] Jakob Karolus, Francisco Kiss, Caroline Eckerth, Nicolas Viot, Felix Bachmann, Albrecht Schmidt, Paweł W. Woźniak, EMBody: A data-centric toolkit for EMG-based interface prototyping and experimentation, *Proc. ACM Human-Computer Interact.* 5 (EICS) (2021) 195:1–195:29, Open-source toolkit available at <https://github.com/HCUM/embody>.
- [18] Jayden Chapman, Anany Dwivedi, Minas Liarokapis, A wearable, open-source, lightweight forcemyography armband: On intuitive, robust muscle-machine interfaces, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2021, pp. 4138–4143, Open-source design files and software available at <https://github.com/newdexterity/FMGInterface>.
- [19] V. J. Kartsch Morinigo, Simone Benatti, Luca Benini, A high SNR, low-latency dry EMG acquisition system for unobtrusive HMI devices, in: Proceedings of the IEEE Biomedical Circuits and Systems Conference, BioCAS, 2022, pp. 544–548, Open-source hardware design to be released.
- [20] Zhaoyun Zhang, Jingpeng Li, A review of artificial intelligence in embedded systems, *Micromachines* 14 (5) (2023) 897, Open-access article.
- [21] Angelo Garofalo, Giuseppe Tagliavini, Francesco Conti, Luca Benini, Davide Rossi, XpulpNN: Enabling energy efficient and flexible inference of quantized neural networks on RISC-V based IoT end nodes, *IEEE Trans. Emerg. Top. Comput.* 9 (3) (2021) 1489–1505.
- [22] Mingde Zheng, Michael S. Crouch, Michael S. Eggleston, Surface electromyography as a natural human-machine interface: A review, *IEEE Sensors J.* 22 (10) (2022) 9198–9214.
- [23] Lian Cheng, Jun Li, Aiying Guo, Jianhua Zhang, Recent advances in flexible non-invasive electrodes for surface electromyography acquisition, *Npj Flex. Electron.* 7 (39) (2023).
- [24] Marcello Zanghieri, Pierangelo M. Rapa, Mattia Orlandi, Étienne Buteau, Félix Chamberland, Benoit Gosselin, Luca Benini, Simone Benatti, Wearable high-density sEMG processing with class activation maps with an embedded temporal convolutional network, in: 2024 IEEE Biomedical Circuits and Systems Conference, BioCAS, 2024, pp. 1–5.
- [25] Xiaorong Zhang, He Huang, Qing Yang, Real-time implementation of a self-recovery EMG pattern recognition interface for artificial arms, in: Engineering in Medicine and Biology Society (EMBC), 2013 35th Annual International Conference of the IEEE, IEEE, 2013, pp. 5926–5929.
- [26] Cristhian Manuel Duran Acevedo, Javier Eduardo Jauregui Duarte, Development of an embedded system for classification of EMG signals, in: 2014 III International Congress of Engineering Mechatronics and Automation, CIIMA, IEEE, 2014, pp. 1–5.
- [27] Md Asif Ahamed, Md Asraf-Ul Ahad, Md Hanif Ali Sohag, Mohiuddin Ahmad, Development of low cost wireless biosignal acquisition system for ECG EMG and EOG, in: 2015 2nd International Conference on Electrical Information and Communication Technologies, EICT, IEEE, 2015, pp. 195–199.
- [28] M. Rossi, S. Benatti, E. Farella, L. Benini, Hybrid EMG classifier based on HMM and SVM for hand gesture recognition in prosthetics, in: 2015 IEEE International Conference on Industrial Technology, ICIT, 2015, pp. 1700–1705.
- [29] Victor Kartsch, Marco Guermandi, Simone Benatti, Fabio Montagna, Luca Benini, An energy-efficient IoT node for HMI applications based on an ultra-low power multicore processor, in: 2019 IEEE Sensors Applications Symposium, SAS, IEEE, 2019, pp. 1–6.
- [30] Marcos Aviles, Luz-María Sánchez-Reyes, Rita Q. Fuentes-Aguilar, Diana C. Toledo-Pérez, Juvenal Rodríguez-Reséndiz, A novel methodology for classifying EMG movements based on SVM and genetic algorithms, *Micromachines* 13 (12) (2022) 2108.
- [31] Bushra Saeed, Muhammad Zia-ur Rehman, Syed Omer Gilani, Faisal Amin, Asim Waris, Mohsin Jamil, Muhammad Shafique, Leveraging ANN and LDA classifiers for characterizing different hand movements using emg signals, *Arab. J. Sci. Eng.* 46 (2021) 1761–1769.
- [32] Simone Benatti, Fabio Montagna, Victor Kartsch, Abbas Rahimi, Davide Rossi, Luca Benini, Online learning and classification of EMG-based gestures on a parallel ultra-low power platform using hyperdimensional computing, *IEEE Trans. Biomed. Circuits Syst.* 13 (3) (2019) 516–528.
- [33] Marcello Zanghieri, Simone Benatti, Alessio Burrello, Victor Kartsch, Francesco Conti, Luca Benini, Robust real-time embedded emg recognition framework using temporal convolutional networks on a multicore iot processor, *IEEE Trans. Biomed. Circuits Syst.* 14 (2) (2019) 244–256.
- [34] Motion Lab Inc, https://www.motion-labs.com/prod_preamp.html.
- [35] PULP Platform Team, PULP platform, 2024, <https://pulp-platform.org/>. (Accessed 21 November 2024).
- [36] Marcello Zanghieri, Simone Benatti, Alessio Burrello, Victor Javier Kartsch Morinigo, Roberto Meattini, Gianluca Palli, Claudio Melchiorri, Luca Benini, sEMG-based regression of hand kinematics with temporal convolutional networks on a low-power edge microcontroller, in: 2021 IEEE International Conference on Omni-Layer Intelligent Systems, COINS, IEEE, 2021, pp. 1–6.
- [37] Derya Karabulut, Faruk Ortes, Yunus Ziya Arslan, Mehmet Arif Adli, Comparative evaluation of EMG signal features for myoelectric controlled human arm prosthetics, *Biocybern. Biomed. Eng.* 37 (2) (2017) 326–335.
- [38] Francesco Riillo, Lucia Rita Quitadamo, Francesco Cavrini, Emanuele Gruppioni, Carlo Alberto Pinto, N. Cosimo Pastò, Laura Sbernini, Lorenzo Albero, Giovanni Saggio, Optimization of EMG-based hand gesture recognition: Supervised vs. unsupervised data preprocessing on healthy subjects and transradial amputees, *Biomed. Signal Process. Control.* 14 (2014) 117–125.
- [39] B. Azzerboni, M. Carpentieri, F. La Foresta, F.C. Morabito, Neural-ICA and wavelet transform for artifacts removal in surface EMG, in: 2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No. 04CH37541), vol. 4, IEEE, 2004, pp. 3223–3228.
- [40] Simone Benatti, et al., Multiple biopotentials acquisition system for wearable applications, in: BIODEVICES, 2015, pp. 260–268.
- [41] C. Pylatiuk, M. Muller-Riederer, A. Kargov, S. Schulz, O. Schill, M. Reischl, G. Brethauer, Comparison of surface EMG monitoring electrodes for long-term use in rehabilitation device control, in: 2009 IEEE International Conference on Rehabilitation Robotics, IEEE, 2009, pp. 300–304.
- [42] T. Finni, Min Hu, P. Kettunen, T. Vilavuo, S. Cheng, Measurement of EMG activity with textile electrodes embedded into clothing, *Physiol. Meas.* 28 (11) (2007) 1405.
- [43] Marco Guermandi, Roberto Cardu, Eleonora Franchi Scarselli, Roberto Guerrieri, Active electrode IC for EEG and electrical impedance tomography with continuous monitoring of contact impedance, *IEEE Trans. Biomed. Circuits Syst.* 9 (1) (2015) 21–33.
- [44] g.tec GmbH, g.SAHARA.

- [45] COVIDIEN, Covidien Kendall, 2015.
- [46] Anish C. Turlapaty, Balakrishna Gokaraju, Feature analysis for classification of physical actions using surface EMG data, *IEEE Sensors J.* 19 (24) (2019) 12196–12204.
- [47] Zainal Arief, Indra Adji Sulistijono, Roby Awal Ardiansyah, Comparison of five time series EMG features extractions using Myo Armband, in: 2015 International Electronics Symposium, IES, IEEE, 2015, pp. 11–14.
- [48] Nurhazimah Nazmi, Mohd Azizi Abdul Rahman, Shin-Ichiro Yamamoto, Siti Anom Ahmad, Hairi Zamzuri, Saiful Amri Mazlan, A review of classification techniques of EMG signals during isotonic and isometric contractions, *Sensors* 16 (8) (2016) 1304.
- [49] Victor Kartsch, Simone Benatti, Mattia Mancini, Michele Magno, Luca Benini, Smart wearable wristband for EMG based gesture recognition powered by solar energy harvester, in: 2018 IEEE International Symposium on Circuits and Systems, ISCAS, IEEE, 2018, pp. 1–5.
- [50] Alaa Tharwat, Linear vs. quadratic discriminant analysis classifier: a tutorial, *Int. J. Appl. Pattern Recognit.* 3 (2) (2016) 145–180.
- [51] Shifei Ding, Zhibin Zhu, Xiekai Zhang, An overview on semi-supervised support vector machine, *Neural Comput. Appl.* 28 (5) (2017) 969–978.
- [52] Alberto Dellacasa Bellingegni, Emanuele Gruppioni, Giorgio Colazzo, Angelo Davalli, Rinaldo Sacchetti, Eugenio Guglielmelli, Loredana Zollo, NLR, MLP, SVM, and LDA: a comparative analysis on EMG data from people with trans-radial amputation, *J. Neuroeng. Rehabil.* 14 (1) (2017) 82.
- [53] Chih-Chung Chang, Chih-Jen Lin, LIBSVM: A library for support vector machines, *ACM Trans. Intell. Syst. Technol.* 2 (2011) 27:1–27:27, Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [54] Pentti Kanerva, Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors, *Cogn. Comput.* 1 (2) (2009) 139–159.
- [55] Fabio Montagna, Abbas Rahimi, Simone Benatti, Davide Rossi, Luca Benini, PULP-HD: Accelerating brain-inspired high-dimensional computing on a parallel ultra-low power platform, in: 2018 55th ACM/ESDA/IEEE Design Automation Conference, DAC, IEEE, 2018, pp. 1–6.
- [56] Parth Gargava, Karan Sindwani, Sumit Soman, Controlling an arduino robot using brain computer interface, in: Proceedings of 3rd International Conference on Reliability, Infocom Technologies and Optimization, IEEE, 2014, pp. 1–5.
- [57] Victor Javier Kartsch, Simone Benatti, Pasquale Davide Schiavone, Davide Rossi, Luca Benini, A sensor fusion approach for drowsiness detection in wearable ultra-low-power systems, *Inf. Fusion* 43 (2018) 66–76.
- [58] P. Davide Schiavone, et al., Slow and steady wins the race? A comparison of ultra-low-power RISC-V cores for internet-of-things applications, in: PATMOS, 2017, pp. 1–8.
- [59] Victor Kartsch, Giuseppe Tagliavini, Marco Guermandi, Simone Benatti, Davide Rossi, Luca Benini, BioWolf: A sub-10-mw 8-channel advanced brain-computer interface platform with a nine-core processor and BLE connectivity, *IEEE Trans. Biomed. Circuits Syst.* 13 (5) (2019) 893–906.
- [60] Eric Flamand, Davide Rossi, Francesco Conti, Igor Loi, Antonio Pullini, Florent Rotenberg, Luca Benini, GAP-8: A RISC-v SoC for AI at the edge of the IoT, in: 2018 IEEE 29th International Conference on Application-Specific Systems, Architectures and Processors, ASAP, IEEE, 2018, pp. 1–4.
- [61] Global Industry Analysts, Inc. (GIA), 2018, <https://greenwaves-technologies.com/>.
- [62] Ali Moin, Andy Zhou, Abbas Rahimi, Simone Benatti, Alisha Menon, Senam Tamakloe, Jonathan Ting, Natasha Yamamoto, Yasser Khan, Fred Burghardt, Luca Benini, Ana C. Arias, Jan M. Rabaey, An EMG gesture recognition system with flexible high-density sensors and brain-inspired high-dimensional classifier, in: 2018 IEEE International Symposium on Circuits and Systems, ISCAS, 2018, pp. 1–5.
- [63] Jose Manuel Fajardo, Orlando Gomez, Flavio Prieto, EMG hand gesture classification using handcrafted and deep features, *Biomed. Signal Process. Control.* 63 (2021) 102210.
- [64] Afroza Sultana, Farruk Ahmed, Md Shafiul Alam, A systematic review on surface electromyography-based classification system for identifying hand and finger movements, *Heal. Anal.* 3 (2023) 100126.



Victor Kartsch received the Ph.D. degree in electrical engineering and computer science from the University of Bologna, in 2020 (Ph.D. Advisor Prof. Luca Benini). During the Ph.D., he worked on the hardware–software design of fully embedded human–machine interaction (HMI) systems with a full-stack perspective, targeting both EMG and EEG signals. One of the most important systems developed by Dr. Kartsch includes BioWolf, an ultralow-power HMI for signal acquisition and real-time processing of computationally-heavy algorithms, which has been adopted



by many research institutions in the field. His work experience also includes data analysis and optimization of signal processing and machine learning algorithms for embedded systems. He has published several papers in international peer-reviewed conferences and journals. Currently, he is a Research Fellow with the Integrated System Laboratory, ETH Zürich, where he also works on designing wearable systems for several applications. At ETH, he is also working on the design of hardware and software solutions for UAVs, and, by extension, on integrating such systems with HMIs for advanced control.

Simone Benatti received the Ph.D. degree in electronics, telecommunications and information technologies from the University of Bologna, under the supervision of Prof. Luca Benini. During the Ph.D., he was a Visiting Fellow with the BWRC - University of California, Berkeley (Supervisor Prof. Jan Rabaey). In 2023, he has been appointed as a Visiting Professor with the EFCL-ETHZ. Currently, he serves as an Assistant Professor with the University of Modena e Reggio Emilia, while pursuing his collaboration with IIS-ETH in Zurich. His research interests include energy-efficient embedded systems for IoT and biomedical applications. This includes hardware/software codesign to efficiently address performance, as well as advanced algorithms. He works on designing and optimizing energy-efficient embedded systems for biopotential (ExG) acquisition and processing, and biosignal-based HMIs. In this field, he has published more than 100 papers in international peer-reviewed conferences and journals. He has ongoing collaborations with several international research institutes, such as ETHZEFCL, EPFL, TU Graz, FBK and Politecnico di Torino. He is the recipient of the GHAI Grant (H2020-MSCA-RISE-2017, G.A. 777822).



Fiorenzo Artoni is a Maître Assistant (Senior Researcher) at the University of Genève, Campus Biotech – Switzerland. Previously he was a Post-Doc at the Biorobotics Institute, Sant’Anna School of Advanced Studies (SSSA) and Ecole Polytechnique Federale de Lausanne (EPFL) - Switzerland, after being awarded a H2020 Marie Curie Individual Fellowship (project “BIREHAB”) and the NCCR robotics Spin Funds Grant. He obtained a Ph.D. in Biorobotics at SSSA in 2014, after securing a dual M.Sc. in Automation Engineering and Industrial Engineering and a dual B.Sc. in Biomedical Engineering and Industrial Engineering in 2011 both from the University of Pisa and SSSA. His research interest include the development of novel noninvasive biomarkers to optimize post-stroke rehabilitation treatments and the development of methods for EEG/EMG analysis including Independent Component Analysis and syntax of microstates. He is particularly active in the fields of tactile feedback restoration for neuroprostheses, speech decoding, tracking of anesthesia, and Mobile Brain/Body Imaging



Silvestro Micera is currently Professor of Biomedical Engineering at the Scuola Superiore Sant’Anna (SSSA, Pisa, Italy) and at the Ecole Polytechnique Federale de Lausanne (Lausanne, Switzerland) where he is holding the Bertarelli Foundation Chair in Translational NeuroEngineering. He received the University degree (Laurea) in Electrical Engineering from the University of Pisa, in 1996, and the Ph.D. degree in Biomedical Engineering from the Scuola Superiore Sant’Anna, in 2000. From 2000 to 2009, he has been an Assistant Professor of BioRobotics at the Scuola Superiore Sant’Anna. In 2007, he was a Visiting Scientist at the Massachusetts Institute of Technology, Cambridge, USA with a Fulbright Scholarship. From 2008 to 2011 he was the Head of the Neuroprosthesis Control group and Group Leader at the Institute for Automation, Swiss Federal Institute of Technology, Zurich, CH. He was the recipient of the “Early Career Achievement Award” and of the “Technical Achievement Award” of the IEEE Engineering in Medicine and Biology Society in 2009 and 2021, respectively. Dr. Micera’s research interests include the development of neuroprostheses based on the use of implantable neural interfaces with the central and peripheral nervous systems to restore sensory and motor function in disable persons. In particular, he is currently involved in translational experiments for hand prosthesis control in amputees, and the restoration

of vestibular function, grasping and locomotion in different neurological disorders. He is author of more than 300 WoS peer-reviewed papers and several international patents. He is currently Associate Editor of IEEE Transactions on Neural Systems and Rehabilitation Engineering and of IEEE Transactions on Medical Robotics and Bionics. He is also member of the Editorial Boards of the Journal of Neuroengineering and Rehabilitation, and of Journal of Neural Engineering.



Luca Benini received the Ph.D. degree from Stanford University. He holds the Chair of Digital Circuits and Systems with the ETH Zurich, Zurich, Switzerland, and is a Full Professor with the University of Bologna, Italy. His research interests include energy-efficient parallel computing systems, smart sensing micro-systems, and machine learning hardware. He is a Fellow of the ACM and a Member of the Academia Europaea. He is the recipient of the 2016 IEEE CAS Mac Van Valkenburg Award, the 2020 EDAA Achievement Award, the 2020 ACM/IEEE A. Richard Newton Award, and the 2023 IEEE CS E.J. McCluskey Award.