

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Towards Long-term Non-invasive Monitoring for Epilepsy via Wearable EEG Devices

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Towards Long-term Non-invasive Monitoring for Epilepsy via Wearable EEG Devices / Ingolfsson T.M.; Cossettini A.; Wang X.; Tabanelli E.; Tagliavini G.; Ryvlin P.; Benini L.; Benatti S.. - ELETTRONICO. - (2021), pp. 01-04. (Intervento presentato al convegno 2021 IEEE Biomedical Circuits and Systems Conference, BioCAS 2021 tenutosi a Berlin, Germany nel 07-09 October 2021) [10.1109/BioCAS49922.2021.9644949].

Availability:

This version is available at: <https://hdl.handle.net/11585/871078> since: 2022-02-27

Published:

DOI: <http://doi.org/10.1109/BioCAS49922.2021.9644949>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

T. M. Ingolfsson et al., "Towards Long-term Non-invasive Monitoring for Epilepsy via Wearable EEG Devices," 2021 IEEE Biomedical Circuits and Systems Conference (BioCAS), Berlin, Germany, 2021, pp. 01-04

The final published version is available online at
<https://dx.doi.org/10.1109/BioCAS49922.2021.9644949>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Towards Long-term Non-invasive Monitoring for Epilepsy via Wearable EEG Devices

Thorir Mar Ingolfsson*, Andrea Cossetti†, Xiaying Wang*, Enrico Tabanelli†, Giuseppe Tagliavini†, Philippe Ryvlin§, Luca Benini*†, Simone Benatti†‡

*Integrated Systems Laboratory, ETH Zürich, Zürich, Switzerland

†DEI, University of Bologna, Bologna, Italy

‡University of Modena and Reggio Emilia, Modena, Italy

§Lausanne University Hospital (CHUV), Switzerland

Abstract—We present the implementation of seizure detection algorithms based on a minimal number of EEG channels on a parallel ultra-low-power embedded platform. The analyses are based on the CHB-MIT dataset, and include explorations of different classification approaches (Support Vector Machines, Random Forest, Extra Trees, AdaBoost) and different pre/post-processing techniques to maximize sensitivity while guaranteeing no false alarms. We analyze global and subject-specific approaches, considering all 23-electrodes or only 4 temporal channels. For 8s window size and subject-specific approach, we report zero false positives and 100% sensitivity. These algorithms are parallelized and optimized for a parallel ultra-low power (PULP) platform, enabling 300h of continuous monitoring on a 300 mAh battery, in a wearable form factor and power budget. These results pave the way for the implementation of affordable, wearable, long-term epilepsy monitoring solutions with low false-positive rates and high sensitivity, meeting both patient and caregiver requirements.

Index Terms—healthcare, time series classification, smart edge computing, machine learning, deep learning.

I. INTRODUCTION

Epilepsy is a brain disease that affects more than 50 million people worldwide [1]. Conventional treatments are mostly pharmacological, but they can require surgery or invasive neurostimulation in the case of drug-resistant subjects. To personalize patient treatments, continuous monitoring of the brain activity is required, and it can be performed in epilepsy monitoring units (EMUs), where patients are monitored by a video surveillance system synchronized with 32-64 channels EEG caps. EMU monitoring helps clinicians to refine the diagnosis of the seizures and to monitor the effects of therapies in the short term (i.e., 1–2 weeks). However, continuous monitoring of the patients outside the EMU is paramount to control and refine therapies in the medium period.

The golden diagnostic standard is represented by Electroencephalography (EEG) systems, which unfortunately are cumbersome and can make patients uncomfortable because of perceived stigmatization. Thus, both patients and caregivers would benefit from the availability of wearable long-term EEG monitoring devices, which can also detect sub-clinical seizures with high levels of sensitivity and low false detection rate [2].

This solution poses several challenges. Firstly, traditional EEG caps have 16–32 channels, providing full coverage of the scalp. Hence, it must be understood if it is possible to

effectively detect seizures with a reduced electrode set, to be placed into minimally intrusive devices, such as glasses or behind-the-ear systems [3]–[6]. The second challenge is linked to the performance of the system, especially for what concerns specificity. False alarms represent one of the major causes of stress in seizure detection frameworks, because they raise the level of anxiety of the subject, which has been identified to be a potential trigger for seizures [7]. Recent works [8] provide epilepsy detection algorithms with very low false-alarms, but they are based on computationally demanding algorithms and are used in implantable systems for pre-surgical monitoring.

To address the first challenge, investigations about applying machine learning (ML) algorithms on reduced EEG montages are needed. In particular, the focus should be put on meeting the limited computation resources of embedded platforms. As regards the second challenge, parallel and ultra-low-power (PULP) embedded platforms represent a key enabling technology for the successful deployment of such algorithms into wearable long-term monitoring solutions [9], since they overcome the limits of commercial devices in terms of power consumption, while providing high computational capabilities.

Within these considerations, this work presents a seizure detection framework designed for wearable unobtrusive systems. Specifically, we propose the following contributions:

- A comparison over several state-of-the-art (SoA) detection algorithms (namely, Support Vector Machines, Random Forest, Extra Trees, and AdaBoost), on a benchmark EEG dataset, showing how performance changes between full-coverage (23 channels) and minimal (4 channels over the temporal lobe, *temporal channels* as a shorthand in the remaining text) setup.
- An investigation on the impact of different pre- and post-processing approaches to maximize sensitivity and minimize the false-positive (FP) rate. This includes analyzing the performance improvements when relying on subject-specific training instead of global models.
- The implementation and the performance tuning of the above framework on real PULP chip target, namely Mr. Wolf [9], designed in 40nm FDSOI technology and widely used in ultra-low-power wearable nodes for biomedical applications [10].

Experimental results demonstrate that a high-accuracy low-FP rate framework for continuous epilepsy monitoring can




Fig. 1. Schematic representation of the workflow. Input data (all 23 channels or 4 temporal channels) are pre-processed by means of DWT. Classification is performed with SVM, RF, ET, or AB algorithms. The output of the classifier is post-processed with a majority voting.

be designed in a wearable long-lifetime system with a low electrode count.

II. MATERIALS & METHODS

Our investigations are based on the CHB-MIT dataset from the Children’s Hospital of Boston and MIT [11], [12]. Fig. 1 shows a sketch of the overall workflow of our investigations, whose details are elaborated in the following sections.

A. Pre-processing and feature extraction

Discrete wavelet transform (DWT) is a signal decomposition technique for time-frequency analyses [13], widely used in preprocessing stages of ML algorithms. In contrast with other signal transforms (e.g., Fourier transforms), DWT captures both frequency and temporal features of a given signal by applying a cascade series of filters, shaped as *Mother Wavelets*. Input signals are passed iteratively through low-pass and a high-pass filters, resulting in a series of *Approximation* and *Detail* coefficients, respectively. Both filters can be implemented as convolutions between the signal and pre-computed kernels, followed by downsampling by two. Approximation coefficients of a given decomposition level represent the input of the next level; following this pattern, each level halves the time resolution but doubles the frequency resolution.

In the context of bio-signal analysis, DWT has been used in state-of-the-art research works to extract features from physiological data since they provide a good trade-off between performance and signal-to-noise ratio [14]. We use DWT with three different window lengths: 2s, 4s, or 8s. Our multi-core platform (Sec. II-D) allows to process these different window sizes in a fast, energy efficient, and accurate manner, performing the 4-level DWT on two cores for each channel.

B. Classification

We target the deployment of memory efficient and low computational algorithms to enable lightweight epileptic seizure detection at the edge. Thus, we estimate the performance of 4 supervised learning methods: a Support Vector Machine (SVM), Random Forest (RF), Extra Trees (ET), and AdaBoost (AB) classifiers. Supervised learning relies on two phases: the training, where a model is tuned by using labeled data (we use ictal and interictal EEG samples to create the models), and the inference, where, exploiting the trained model, the algorithm determines which class a sample belongs to. In the following, we give a brief overview of the algorithms mentioned above.

1) *SVM* is a supervised ML method that offers high generalization performance. By learning a set of hyperplanes

in high-dimensional space, SVM separates the feature space and classifies instances. To make data linearly separable, we map them into high-dimensional feature spaces deploying the Radial Basis Function (RBF). The RBF kernel consists of a real-valued function relying on the distance between the input and a fixed point. The distance metric is usually the Euclidean Distance. During inference the SVM classifies instances based on which side of the hyperplane they lie. Referring to the UCI epileptic dataset [15], a 98.18% classification accuracy was achieved by [16] when applying SVM.

2) *RF* is an easy-to-use widely-adopted ML algorithm that relies on an ensemble of low-correlation tree-structured classifiers [17]. By randomly picking a feature space subset at each tree splitting node, the training algorithm searches for the most performing thresholds to split training instances until reaching a leaf node. By aggregating trees’ votes, the model inference returns the input prediction picking up the class with the largest number of votes. A 99.78% accuracy was achieved by [16] on the UCI dataset.

3) *ET* is a tree-based ensemble classification method that leverages stronger randomization to achieve a favorable bias-variance trade-off [18]. ET picks a random feature space subset at each node while it randomly generates a threshold for each candidate feature, selecting the optimal one as the node splitting rule. The additional randomness reduces the complexity of the induction process, increases the training speed, and weakens correlation. ET follows the same approach as RF for model inference. An ET-based epileptic signal classification algorithm was proposed by [19], reaching up to 89.49% seizure detection accuracy on the CHB-MIT dataset.

4) *AB* is a statistical classification meta-algorithm that combines a set of weak learners through a weighted vote to reduce bias and variance in supervised learning tasks [20]. The boosting technique enables learning an ensemble algorithm with higher performance than single learners and is less prone to overfitting. During inference, AB sums up the weighted vote of the weak learners to classify between instances. An automatic seizure detection algorithm based on AB was shown to reach 99.08% accuracy on the UCI epileptic dataset [21].

As in [22], during the training phase we heavily penalize the false alarm rate by assigning different weights to *seizures* and *non-seizures*. We define the class weights as the inverse of the frequency of the occurrence of the two classes. In this configuration, the weight ratios are “balanced”. We perform classifications considering the following scenarios:

- training/testing on all patients, using all 23 channels
- training/testing on all patients, with 4 temporal channels
- subject-specific training/testing, using all 23 channels
- subject-specific training/testing, with 4 temporal channels

C. Post-processing

Once occurring, epileptic seizures are expected to last multiple seconds. As in [23], we apply moving-averages to the labels with a window of 4s, performing majority voting, thus effectively smoothing out the predictions. The low-pass filtering effect of the moving average gives a great benefit in terms FP reduction, since it basically eliminates the fluctuations of the classifier output. In the following, we consider moving averages over 3 successive classifications.

D. Embedded platform

The seizure detection framework is implemented on the BioWolf wearable ExG device [10] (Fig. 2). BioWolf includes an 8-channel analog front-end (AFE) for biopotential acquisition, the PULP Mr. Wolf multicore processor, and a wireless Bluetooth low-energy (BLE) link enabled by a Nordic nRF8232 chip. By virtue of the multicore digital platform, BioWolf outperforms by at least one order of magnitude the performance achievable with conventional single-core low-power processors, such as ARM CORTEX M4 or MSP 430 based devices, with a comparable power budget. The BioWolf system is based on a fully programmable processor, which combines versatility of a programmable system, where algorithms can be tuned and improved, with the efficiency of a computational 8-core cluster with 2 shared floating point units. It should also be noted that a hardwired application specific integrated circuit (IC) solution is more efficient (at the price of lack of flexibility) only in applications where the power envelope is very critical (e.g., implantable devices), and typically can not exceed 1–2 mW total power.




Fig. 2. Left: main building blocks of the BioWolf [10] embedded platform. Right: wearable form factor of BioWolf, compared to a 2 euros coin.

E. Power Measurements

We perform the measurements on Mr. Wolf considering cluster domain with the processors running at 100 MHz (frequency of maximal energy efficiency [9]). The measurements are performed using a Keysight N6705C power analyzer, with 0.2048 ms sampling interval.

III. RESULTS & DISCUSSION

Table I summarizes the accuracy metrics reported for the four models considered in this paper, applied on the CHB-MIT data. We compare a global and a subject-specific approach to test the models. The reported numbers correspond to a balanced weighting ratio. In the global approach, SVM and AB have too high FPs/h. RF and ET have fewer FPs/h, but lack the sensitivity to be considered reliable approaches. These conclusions apply both when considering all 23 channels (10-20 EEG standard) and when considering only the 4 temporal ones (F7-T7, T7-P7, F8-T8, T8-P8).

On the contrary, in subject-specific training, specificity and sensitivity increase for all models, providing at the same time a reduced FP rate. Fig. 3 shows that the performance of AB increases (both in terms of specificity and sensitivity) at increasing temporal window lengths in the feature pre-processing; the other algorithms provide a similar behavior (Tab. I), where for the subject-specific training we approach 0 FP/h, while retaining high sensitivity values. Fig. 3 also




Fig. 3. Sensitivity/Specificity plot for subject-specific AdaBoost. Left: global model approach, demonstrating the dependencies on the weighting ratio and on the choice of the window length (2 s, 4 s, or 8 s). Increasing the weight ratio results in greater sensitivity, at the cost of a reduced specificity. Longer windows improve specificity, at the cost of a reduced sensitivity. With weight ratio = 100, the obtained FP rate is 133 FP/h, 36 FP/h, and 11 FP/h, for the 2 s, 4 s, and 8 s windows, respectively. Right: subject-specific approach, demonstrating the benefit of post-processing of output. For both plots the electrodes considered are 4 temporal channels.

shows the effect of post-processing (smoothing) the classification outputs: both specificity and sensitivity are increased. We notice that, at least for data analyzed in this experiments, reducing the number of electrodes does not affect significantly classification performance. As a result, the deployment of a minimally invasive setup, which can be integrated in smart-glasses or other wearable devices with low stigmatization, seems to be feasible.

Regarding the embedded deployment of the classification algorithms, the BioWolf platform benefits of the multicore computing cluster of the PULP chip. Tab. II summarizes energy measurements and parallel speedup concerning sequential execution for the four algorithms evaluated on Mr. Wolf. RF, ET, and AB achieve near-ideal speedups between $6.89\text{--}7.21\times$, only limited by architectural factors, such as contentions due to the concurrent accesses to tightly coupled data memory (TCDM) banks and instruction misses on the instruction cache, which does not scale linearly with the cores. Instead, SVM consists of a workload parallelizable up to around 97%, leading to an ideally achievable speedup of $6.53\times$. Therefore, it reaches a lower speedup of $5.84\times$, also bounded by the aforementioned architectural limits. RF, ET, and AB then show much better energy usage than the SVM, requiring only around $1\mu\text{J}$ per inference. This stems from the much less resource-intensive calculations needed than for decision-tree-based inference. In comparison, assuming an SoA commercial AFE for biosignal acquisition such as T.I. ADS1298, which requires 0.75 mW per channel, and considering a battery of 300 mAh, our approaches implemented on Mr. Wolf would last for roughly 300 hours of data acquisition and classification at 4 s intervals. Comparing to [24], which implemented a seizure classification method on an STM32L476 ARM Cortex-M4 microcontroller, we see a $7.34\times$ improvement in battery life.

The implementations of RF, EF, and AB on a multi-core edge platform, such as Mr. Wolf, require a power envelope of only around 22 mW with sub-100 μs processing time, thus with lower energy requirements than the AFE and BLE of the complete platform that enables multi-day functionality [10]. These results pave the way to the desing of unobtrusive EEG wearable devices that can help patient and clinicians in long

TABLE I
COMPARISON BETWEEN ALGORITHMS

		All Channels				Temporal Channels											
		2s window, w/o smooth.				2s window, w/o smooth.				8s window, w/o smooth.				8s window, w smooth.			
		SVM	RF	ET	AB	SVM	RF	ET	AB	SVM	RF	ET	AB	SVM	RF	ET	AB
Global	Specificity [%]	95.9	99.9	99.9	94.0	96.3	99.9	99.9	98.5	96.3	99.9	99.9	98.1	97.1	99.9	99.9	99.1
	Sensitivity [%]	72.7	39.0	34.8	84.6	75.1	41.2	35.2	84.3	75.1	48.5	42.3	77.3	81.1	49.5	47.3	85.3
	FP/h	73.5	1.8	1.6	107.8	65.9	2.2	1.6	27.2	16.5	0.4	0.09	8.3	13.1	0.4	0.09	3.6
Subject-Specific	Specificity [%]	97.9	99.9	99.9	99.4	99.1	99.8	99.8	99.4	98.9	99.9	99.8	99.9	99.4	100	100	100
	Sensitivity [%]	84.2	63.6	65.1	87.8	88.9	86.7	89.2	88.9	100	100	100	100	100	100	100	100
	FP/h	37.8	2.2	1.8	10.3	17.1	3.8	2.5	10.2	5.1	0.5	1.0	0.4	2.7	0	0	0

TABLE II
COMPARISON BETWEEN ADABOOST, RANDOM FOREST, EXTRA TREES AND SUPPORT VECTOR MACHINE ON MR WOLF.

Classifier	SVM	RF	ET	AB
Time/inference [ms]	0.20	0.052	0.052	0.057
Power [mW]	25.41	21.89	21.89	22.44
Energy/inference [μ J]	5.08	1.14	1.14	1.28
Parallel speedup	5.84 \times	7.21\times	7.21\times	6.89 \times

term monitoring of epileptic seizures in out-of-the-hospital environments.

IV. CONCLUSION

This work presents a systematic analysis on the feasibility of long-term epilepsy monitoring with minimal EEG setups. We used DWT for signal pre-processing, in combination with four different classification algorithms (SVM, Random Forest, Extra Trees, AdaBoost), to assess the impact of limiting the number of EEG channels to only 4 covering the temporal lobe, enabling a discrete and user comfortable setup. Analyses are done on the CHB-MIT dataset, and we demonstrate that with proper weighting, selection of the observation window size, and smoothing of predictions, 100% sensitivity and zero false-positives can be achieved, when the algorithms are trained and tested on the same subjects. We further deploy the algorithms on a PULP platform, resulting in minimal energy requirements ($\approx 1 \mu$ J per inference), outperforming competing commercial devices. These results demonstrate that a PULP system is indeed one of the best candidates for future wearable epilepsy monitoring systems based on minimal EEG setups to avoid stigmatization. Future work will focus also on expanding the proposed techniques to include sensor fusion from other data sources, such as electrocardiograms [25] or inertial sensors, to be realized in a wearable setting of a body area network and on testing the whole system in an ambulatory and domestic environments.

ACKNOWLEDGMENT

This project was supported by the Swiss National Science Foundation (Project PEDESITE) under grant agreement 193813.

REFERENCES

[1] W. H. Organization *et al.*, “Epilepsy: a public health imperative,” 2019.
[2] E. Bruno *et al.*, “Seizure detection at home: Do devices on the market match the needs of people living with epilepsy and their caregivers?” *Epilepsia*, vol. 61, pp. S11–S24, 2020.

[3] R. Asif *et al.*, “Epileptic seizure detection with a reduced montage: A way forward for ambulatory EEG devices,” *IEEE Access*, vol. 8, pp. 65 880–65 890, 2020.
[4] N. Pham *et al.*, “WAKE: a behind-the-ear wearable system for microsleep detection,” in *Proc. Int. MobiSys*, 2020, pp. 404–418.
[5] D. Sopic *et al.*, “e-glass: A wearable system for real-time detection of epileptic seizures,” in *IEEE ISCAS*, 2018, pp. 1–5.
[6] M. Guermandi *et al.*, “A wearable device for minimally-invasive behind-the-ear eeg and evoked potentials,” in *IEEE BioCAS*, 2018, pp. 1–4.
[7] H. R. McKee *et al.*, “Stress as a seizure precipitant: Identification, associated factors, and treatment options,” *Seizure*, vol. 44, pp. 21–26, 2017, 25th Anniversary Issue.
[8] A. Burrello *et al.*, “Laelaps: An energy-efficient seizure detection algorithm from long-term human iecg recordings without false alarms,” in *IEEE DATE*, 2019, pp. 752–757.
[9] A. Pullini *et al.*, “Mr. wolf: An energy-precision scalable parallel ultra low power soc for iot edge processing,” *IEEE J. Solid-State Circuits*, vol. 54, no. 7, pp. 1970–1981, 2019.
[10] V. Kartsch *et al.*, “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.*, vol. 13, pp. 893–906, 2019.
[11] A. H. Shoeb, “Application of machine learning to epileptic seizure onset detection and treatment,” Ph.D. dissertation, MIT, 2009.
[12] A. L. Goldberger *et al.*, “Physiobank, physiobank, and physionet: components of a new research resource for complex physiologic signals,” *circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
[13] S. G. Mallat, “A theory for multiresolution signal decomposition: the wavelet representation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, 1989.
[14] S. Benatti *et al.*, “Towards EMG control interface for smart garments,” in *Proc. ACM ISWC*, 2014, pp. 163–170.
[15] R. Andrzejak *et al.*, “Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: dependence on recording region and brain state,” *Phys Rev E Stat Nonlin Soft Matter Phys*, vol. 64, pp. 1–6, 2001.
[16] S. Raut *et al.*, “Comparative study on machine learning classifiers for epileptic seizure detection in reference to eeg signals,” in *Proc. Int. Conf. AIAI*. Springer, 2021, pp. 185–194.
[17] M. Tomasin *et al.*, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
[18] P. Geurts *et al.*, “Extremely randomized trees,” *Machine Learning*, vol. 63, no. 1, pp. 3–42, 2006.
[19] D. Hu *et al.*, “Epileptic signal classification based on synthetic minority oversampling and blending algorithm,” *IEEE Trans. Cogn. Devel. Syst.*, pp. 1–1, 2020.
[20] Y. Freund *et al.*, “A decision-theoretic generalization of on-line learning and an application to boosting,” *J. Comput. Syst. Sci.*, vol. 55, no. 1, pp. 119–139, 1997.
[21] A. R. Hassan *et al.*, “Epilepsy seizure detection using complete ensemble empirical mode decomposition with adaptive noise,” *Knowledge-Based Systems*, vol. 191, p. 105333, 2020.
[22] V. Shah *et al.*, “Validation of temporal scoring metrics for automatic seizure detection,” in *Proc. IEEE SPMB*, vol. 1, no. 1, 2020.
[23] U. Pale *et al.*, “Systematic assessment of hyperdimensional computing for epileptic seizure detection,” *arXiv preprint arXiv:2105.00934*, 2021.
[24] R. Zanetti *et al.*, “Robust epileptic seizure detection on wearable systems with reduced false-alarm rate,” in *Proc. IEEE EMBC*, 2020, pp. 4248–4251.
[25] T. M. Ingolfsson *et al.*, “Ecg-tcn: Wearable cardiac arrhythmia detection with a temporal convolutional network,” in *IEEE AICAS*, 2021.