Self-Aware Anomaly-Detection for Epilepsy Monitoring on Low-Power Wearable Electrocardiographic Devices

Farnaz Forooghifar*, Amin Aminifar†, Tomas Teijeiro*,
Amir Aminifar‡, Jesper Jeppesen+, Sandor Beniczky+, David Atienza*
*Embedded Systems Laboratory, EPFL, Switzerland, †Western Norway University of Applied Sciences, Norway
‡Lund University, Sweden, +Aarhus University Hospital, Aarhus, Denmark

Abstract—Low-power wearable technologies offer a promising solution to pervasive epilepsy monitoring by removing the constraints concerning time and location, on one hand, and fulfilling long-term tracking, on the other hand. In the case of epileptic seizures, as the attacks infrequently occur, using an anomaly detection approach reduces the need to record long hours of data for each patient before detecting the successive coming seizures. In this work, by combining the concepts of self-aware system and anomaly detection, we propose an energy-efficient system to detect epileptic seizures on single-lead electrocardiographic signals, which is personalized after analyzing the first seizure of the patient. This system, then, uses a simple anomaly-detection model, whenever the model is deemed reliable, and uses a more complex model otherwise. We show that after the personalization, the number of patients, for which the method provides high sensitivity, can reach 26 out of 43 patients with the false alarm rate (FAR) of 4 alarms/day. Thus, the number of responders to the system is increased by 24%, while the FAR is only increased by one alarm/day, compared to the system that just uses the simple model. This benefit occurs while the system complexity decreases by 27.7% compared to the complex model. After adding the two-level (simple and complex) anomaly-detection, the complexity is tuned between 72.3% and 37.6% of the complex model. Similarly, the sensitivity is tuned between 66.5% and 60.3%.

Index Terms—self-awareness, anomaly detection, low-power, wearable devices, epileptic seizures.

I. Introduction and Related Work

Wearable devices are integrated into our everyday life, monitoring our bio-signals continuously, mainly to gather lifestyle statistics. However, these devices are also aiming towards providing real-time notifications to patients, doctors and caregivers. One of the most common neurological disorders that can benefit from wearable devices for constant monitoring is epilepsy [1]. However, a major limitation for such devices is their battery life and the need to get recharged frequently. At the same time, a significant obstacle in the development of epileptic seizure detection algorithms is the lack of seizure data specifically in developing patient-specific models.

A solution to compensate for the lack of seizure data is using anomaly detection methods, which develop their detection models based only on the non-pathological data [2]. Different methods have been developed to perform anomaly detection, including 1-class SVM [3], isolation forest (iForest) [4], neural networks [5], etc. There are several studies, which have used anomaly detection methods in epileptic seizure detection systems. In [6], the 1-class SVM method is used on EEG data, and the energy-based statistics are employed as features. Although the sensitivity and detection latency are highly satisfying for this work, they have a FAR of 1.6 alarms/hour, which is not acceptable in real scenarios. In [7], the authors have performed patient-specific transfer learning on ECG data. They have focused on FAR reduction, but did not discuss the complexity of their detection algorithm. More recently, in [8], anomaly detection is done on heart rate variability, which again only focuses on accuracy and not on energy consumption.

As we are embedding the detection method on a wearable device, a low-complexity algorithm is necessary to be able to perform constant monitoring of the patient. Thus, we have decided to perform the anomaly detection using iForest algorithm on ECG signal. We train a patient-specific anomaly-detection model, but since not all anomalies correlate to seizures, we exploit the data from the global population to get a first estimator of the anomalies corresponding to seizures. Then, the model is further tuned using the information from the first confirmed seizure of each patient.

In addition to anomaly detection, our solution for reducing energy consumption relies on the notion of self-awareness [9], which equips the system with control units to facilitate monitoring the performance, adapting to changes, and improving autonomously. Although the concept of self-awareness has been used previously for real-time seizure detection [10], [11], it has been applied to classification models, assuming that we have enough seizure data to train a binary classifier. In contrast, here, we are combining this concept with anomaly detection methods to develop our system for datasets with a significant imbalance between normal and seizure data. By utilizing self-awareness, complex anomaly-detection is performed only if necessary, reducing the overall complexity of the calculations.

To validate our approach, we used the dataset and the validation methodology proposed in [12], which is based on the notion of responder subjects, defined as those for which the method provides a sensitivity over $\frac{2}{3}$. To summarize, our main contributions in this paper are as follows:

- In order to avoid large FAR for the patients with high seizure thresholds, we use the information from the first seizure of each patient to decide which detection model should be used. In this way, after applying our method on the ePatch dataset, compared to our complex model, we have reduced the average FAR from 1 alarm per hour to 1 alarm per each 6 hours. Also, the complexity is reduced by 27.7%. At the same time, in comparison with the simple model, the number of responders is increased by 24%, while the FAR is only increased by one alarm/day.

- We utilize the notion of self-awareness in anomaly detection models, to further decrease the energy consumption in the epilepsy detection systems designed for datasets with insufficient seizure samples. By proposing our two-level anomaly-detection model, we offer a trade-off between quality and energy consumption caused by tuning the complexity of the overall system. According to ePatch dataset results, the complexity of detection part is tuned between 72.3% and 37.6% of the complex model, while the sensitivity is tuned between 66.5% and 60.3%.

II. Multi-level self-aware anomaly-detection system

In this section, we describe our seizure detection approach. First, it focuses on an anomaly detection model trained on the non-pathological data of the patient, which is all the data before the first confirmed seizure, or the first 24 hours of recording without a seizure [12]. This model uses leave-one-patient-out cross-validation estimation of the anomaly threshold for seizures. Second, to gain more information and improve the system for each specific patient, we observe the performance of the detection model on the first seizure of the patient. According to this observation, we subsequently use two different self-aware approaches to improve the quality of detection, as well as to decrease the energy consumption of the system. The general overview of our detection system is shown in Figure 1. In the following sections, we describe its different parts.

A. Noise Filtering and Feature Extraction

As the first step, to reduce the noise, a low-pass filter is applied to remove the frequencies above 60 Hz. Then, as shown in Figure 1, a first set of features are extracted from the filtered segment. As detailed in Subsection II-C, feature extraction is tightly coupled with the self-awareness mechanism. Thus, only the minimum set of features required by the simplest model (f_1) are calculated in this step.

B. Anomaly Detection

To overcome the lack of seizure data, we use the iForest algorithm to perform anomaly detection based on the non-pathological data of the patients. The iForest model consists of several decision trees, which output a score for each input indicating how easily this sample can be separated from other samples. The overall score of each sample is then calculated as the average of the scores given by all the trees of the iForest model. Next, according to a threshold score that is chosen in the training phase, the test input is marked as a normal sample or as an anomaly. Thus, we decide the threshold based on two factors: the target false alarm rate (r) and other patients' threshold score. We set the threshold of patient i as the average of other patients' threshold scores for the target FAR:

$$S_{Th}^i = \frac{\sum_{n=1, n \neq i}^{N} (S_{Th}^n | FAR = r)}{N - 1}, \tag{1}$$

where N is the total number of patients, and S_{Th}^n is the score threshold of patient n.

C. Self-Awareness

In this work, self-awareness is defined as a method to decide or switch between different levels of the detection models. This module can be split into two parts:

1) Personalized threshold selection: In this step, we decide whether to use the simplest model ($Model_1$) or our multi-level detection module, described in the next part. To gain information about the complexity of seizure detection in each patient, we use information from the output score of its first observed seizure obtained from passing features f_1 through $Model_1$. Here, we assume that if this seizure is not automatically detected, the patient will have a mechanism to report it to the system. We test the first seizure data with the iForest model, which is trained on the non-pathological data that is gathered before this seizure ($Model_1$). If the scores are higher than the previously defined threshold score ($S_{seiz} > S_{th}$), the detection of the seizures for this patient is assumed to be complex. As a result of this complexity, the threshold should be set to a high value, which can result in many false alarms. To keep the FAR within an acceptable range for these patients, we use the result from the simplest detection model ($Decision_1$). For the other patients, we use our multi-level anomaly-detection module to increase the system's overall sensitivity.

2) Multi-level anomaly-detection: There usually exists a trade-off between the complexity of a model and the accuracy of its decisions. To reduce the runtime complexity, we propose a new self-aware anomaly-detection technique, which relies on a multi-level scheme that combines several models. Then, to reduce the energy consumption of the system, we train different detection models with different accuracies and energy-consumptions ($Model_{1 \rightarrow M}$). The simplest model ($Model_1$) consumes the least energy, but at the same time supports fewer patients in terms of detection quality, compared to other models. As we move towards more complex models, the accuracy is improved, but the energy consumption is also increased. Combining these models enables us to control the trade-offs between accuracy and energy.

We first invoke the simplest anomaly-detection model ($Model_1$), and check the confidence in this first decision ($Conf > Th_1$). If the result's confidence is satisfactory, we adopt the decision ($Decision_1$) as the final result; otherwise,
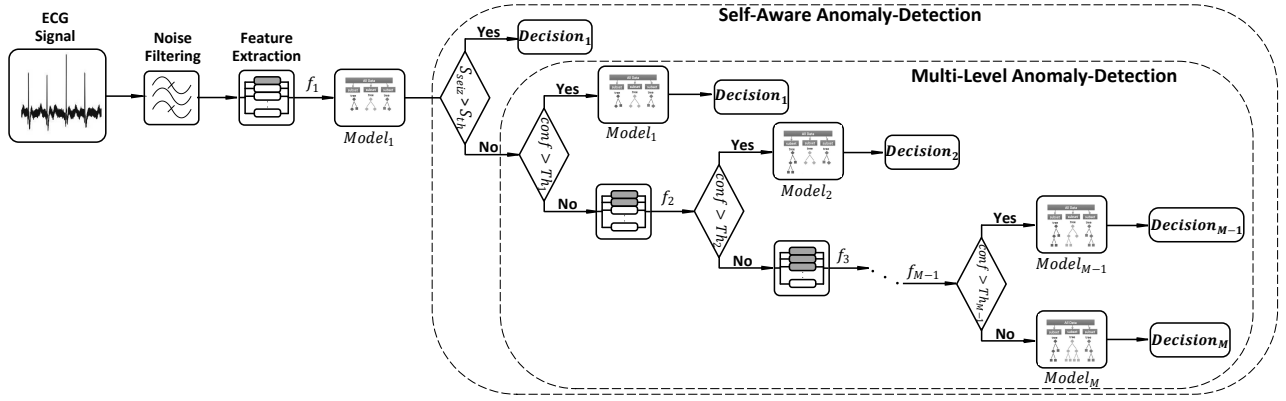
Figure 1: Overview of our proposed self-aware anomaly-detection system

we move one level forward. Then, we calculate several additional features (f_2) and also estimate the confidence of the second model ($Model_2$). If this model's decision ($Decision_2$) is deemed confident ($Conf > Th_2$), it is chosen as the final result. For each input, we continue this procedure until we find a confident model to make the decision on the specific input sample. In the case that no model give a sufficient confidence, the system will rely on the most complex one ($Model_M$).

The confidence of each model's decision is evaluated based on the score consistency of the model. We calculate the confidence threshold based on the average standard deviation of the scores provided by different trees of the iForest for the normal data of the patient before the first seizure, as follows:

$$Th_i = mean(std(score)) + \alpha_i \times std(std(score)), \quad (2)$$

where $score$ is a matrix provided by the iForest model, containing the scores given by all the trees for all training samples. The parameter α_i determines how often the $Model_i$ is invoked. To decide whether the model is confident about the new sample, we suppose that the deviation of the scores of different trees should be less than the threshold calculated in the training phase for this model. So, if this deviation is less than the above confidence threshold, the model's decision is considered confident, and the next models are not triggered.

In our self-aware anomaly-detection technique, for many (test) samples, the simpler models yield confident decisions, meaning that the more complex models are not required. This situation leads to an improved battery life. We formulate the expected complexity as follows:

$$\mathbb{E}(C_{ML}) = \sum_{i=1}^{M} P_i \cdot C_i, \quad (3)$$

where C_{ML} and C_i are the overall complexity of the multi-level detection system and the complexity of $Model_i$, respectively. The parameter P_i captures the probability of invoking $Model_i$, which determines the expected complexity of the proposed technique.

III. EXPERIMENTAL SETUP AND RESULTS

A. Dataset

We have used an ECG dataset collected at the Aarhus University Hospital in Denmark [12], including 43 patients

with an average recording length of 74 hours per patient. The ECG signal is recorded in single-lead mode at a frequency of 512 Hz using a wearable ePatch device. Seizure onset was marked by trained experts as the first clinical or EEG sign of the seizure (whichever came first), and the seizure termination was considered as the last clinical or EEG sign of the seizure (whichever came last). In [12], three basic HRV analysis methods with two different time-lengths of sliding window with maximum overlapping are evaluated (sliding from heartbeat to heartbeat). These basic methods are the followings: Heart Rate differential (HR-diff) method, Cardiac Sympathetic Index (CSI), and Modified CSI (ModCSI). We have used the features, which are calculated for a window of 100 RR intervals, as they achieve a higher detection quality.

B. Two-level Anomaly-Detection

In this study, we limited the multi-level detection system of Figure 1 to two models and designed a two-level anomaly-detection model based on the iForest algorithm. In particular, we trained two iForest models. The first model only considers a limited set of features, which we refer to as the simple model ($Model_1$). For this model, we have used the three features achieving highest number of responders, namely, $CSI \times slope$, $CSI\_filtered \times slope$, and $MCSI\_filtered \times slope$. The second-level model uses all the ten features calculated for 100 RRIs for training and detection and is referred to as the complex model. Since the simple model considers a limited subset of features, it requires less computation for detection compared to the complex model. On the other hand, given the number of features, the complex model is generally more accurate than the simple one.

C. Performance Results

As the provided ECG signal contains continuous recording of the patients, we have used sensitivity and false alarm rate (FAR) performance metrics to evaluate our method. We also use the same concept of responders as introduced in the ePatch paper, thus calculating the performance metrics for both responders (FAR_resp and $sens\_resp$) and the whole dataset (FAR_all and $sens\_all$). The performance results are shown in Table I, where we compare the modes of using only

Table I: Comparison among the performance metrics of different proposed systems, r = 2 alarms/day

Method	P_{simple} %	responders	FAR_resp (/day)	sens_resp %	FAR_all (/day)	sens_all %	complexity (% of $C_{complex}$)
simple		21	2.85	97.7	3.04	59.5	30.0
self-aware	82	22	3.51	96.3	3.99	60.3	37.6
	53	22	4.41	97.8	4.16	61.8	49.9
	16	24	4.99	95.2	5.29	63.4	65.6
	0	26	3.99	93.4	4.46	66.5	72.3
complex		32	24.17	92.1	27.39	77.9	100

the simple model, using only the complex model, and using a self-aware model with different probabilities of employing the simple model. The three different probabilities (P_{simple}) shown in this table results from using the values -1, 0, and 1 for α_i in Eq. (2). In fact, in our self-aware model, we always have the first form of self-awareness, which results in using the simple model for 17 patients. Thus, the numbers presented in the second column are the probabilities of using the simple model for the 26 remaining patients.

As shown in Table I, there is a trade-off between the number of responders (or the overall sensitivity) and the average FAR. If we use the first form of self-awareness, the number of responders is increased by 24% with respect to always using the simple classifier. At the same time, the FAR is increased from 3 to 4 alarms per day. Then, by adding the two-level anomaly-detection method, with different probabilities of using the simple model, the range of responders is between 21 and 26. Then, the number of FAR varies between 3 and 5 alarms per day. Moreover, although the number of responders is very high when we only use the complex model, the FAR is about one alarm per hour, which is not acceptable.

D. Complexity Results

Considering the aforementioned features, whose definitions are available in [12], we observe that the computational complexity for all of these features is O(N), where N is the number of RRIs in one window. Thus, the complexity of complex features calculation is $10/3$ times of the simple features calculation complexity ($\frac{C_{complex}}{C_{simple}} = \frac{10}{3}$). As a result, for the total complexity ($C_{self-aware}$) the following is obtained:

$$\frac{C_{self-aware}}{C_{complex}} = \frac{N_s \times C_{simple} + N_c \times C_{complex} \times E(C_{ML})}{N \times C_{complex}},$$

where N_s and N_c are the number of patients only using the simple model and the number of patients using the complex one totally or partially, respectively. Then, the term $E(C_{ML})$ is the expected complexity calculated as Eq. (3) with $M = 2$. Thus, if we only use the first part of self-awareness, the complexity is reduced to 72.3%. If we add the second part, then the self-aware system complexity for different values of P_{simple}, as shown in Table I, becomes 37.6 %, 49.9 %, and 65.6% of the complex system. To compare the simple and the complex models on an actual embedded device, according to the latest ECG-based smart wearables designs [13], we have implemented both models on the NVIDIA Jetson Nano Developer Kit [14]. For processing 1 hour of data, the simple model takes 6.08 seconds and consumes 0.202 mAh, while the complex model takes 16.28 seconds and consumes 0.575 mAh.

As a result, the energy consumption of the complex model is 2.85 times higher than the simple model on this device, which is in line with the complexity results shown in Table I.

IV. CONCLUSIONS

In this paper, we have presented a self-aware anomaly-detection system for seizure detection on ECG signals. This system enables us to trade-off detection quality with energy and complexity, by using a new multi-level anomaly-detection approach. For our seizure detection case study trained on the ePatch dataset, we have shown that after applying personalization, the number of responders can reach 26 out of 43 patients with the FAR of 4 alarms/day instead of 3 alarms/day, compared to the system that only uses the simple model. For this minimal increase of FAR, the computational complexity of the system decreases by 27.7% compared to the complex model. Then, after adding the two-level anomaly-detection the complexity is tuned between 72.3% and 37.6% of the complex model, while the sensitivity is tuned between 66.5% and 60.3%. These experimental results show that the proposed methodology can provide significant benefits independently of the specific wearable platform where it is implemented, as it relies mainly on the biosignal characteristics and number of calculated features for each model.

REFERENCES

[1] W. H. Organization. (2016) Epilepsy. [Online]. Available: http://www.who.int/mental_health/neurology/epilepsy/en/index.html
[2] V. Chandola *et al.*, "Anomaly detection: A survey," *ACM computing surveys (CSUR)*, vol. 41, no. 3, pp. 1–58, 2009.
[3] B. Schölkopf *et al.*, "Estimating the support of a high-dimensional distribution," *Neural computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
[4] F. T. Liu *et al.*, "Isolation forest," in *2008 eighth ieee international conference on data mining*. IEEE, 2008, pp. 413–422.
[5] C. Zhou *et al.*, "Anomaly detection with robust deep autoencoders," in *23rd ACM SIGKDD*, 2017.
[6] A. B. Gardner *et al.*, "One-class novelty detection for seizure analysis from intracranial eeg." *Journal of Machine Learning Research*, 2006.
[7] T. De Cooman *et al.*, "Semi-supervised one-class transfer learning for heart rate based epileptic seizure detection," in *CinC*. IEEE, 2017.
[8] T. Yamakawa *et al.*, "Wearable epileptic seizure prediction system with machine-learning-based anomaly detection of heart rate variability," *Sensors*, vol. 20, no. 14, p. 3987, 2020.
[9] P. R. Lewis *et al.*, *Self-aware Computing Systems*. Springer, 2016.
[10] F. Forooghifar *et al.*, "A self-aware epilepsy monitoring system for real-time epileptic seizure detection," in *MONET*. ACM/Springer, 2019.
[11] F. Forooghifar *et al.*, "Resource-aware distributed epilepsy monitoring using self-awareness from edge to cloud," *IEEE TBioCAS*, 2019.
[12] J. Jeppesen *et al.*, "Seizure detection based on heart rate variability using a wearable electrocardiography device," *Epilepsia*, 2019.
[13] R. Braojos *et al.*, "Ultra-low power design of wearable cardiac monitoring systems," in *DAC*. IEEE, 2014.
[14] *Jetson Nano Developer Kit:*. [Online]. Available: https://developer.nvidia.com/embedded/jetson-nano-developer-kit