

Adaptive Heart Rate-Based Epileptic Seizure Detection
Using Real-Time User Feedback¹

Thomas De Cooman, Troels Kjaer, Sabine Van Huffel, Helge Sorensen ²

December 2017

Accepted for publication in Physiological Measurement

¹This report is available by anonymous ftp from <ftp.esat.kuleuven.ac.be> in the directory ftp://ftp.esat.kuleuven.be/pub/stadius/tdecooman/reports/SISTA_17_85_submit.pdf

²T. De Cooman and S. Van Huffel are affiliated with KU Leuven, Dept. of Electrical Engineering (ESAT), Research group STADIUS, Kasteelpark Arenberg 10, 3001 Leuven, Belgium and imec, Leuven, Belgium. Troels Kjaer is affiliated with Neurophysiology Center, Zealand University Hospital, Roskilde, Denmark Helge Sorensen is affiliated with Department of Electrical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark

Acknowledgments: Research supported by Bijzonder Onderzoeksfonds KU Leuven (BOF): Center of Excellence (CoE): PFV/10/002(OPTEC), SPARKLE #: IDO-13-0358, C24/15/036, C32/16/00364; Fonds voor Wetenschappelijk Onderzoek-Vlaanderen (FWO): G.0A55.13N (Deep brain stimulation); Agentschap Innoveren & Ondernemen (VLAIO): STW 150466 OSA +, O&O HBC 2016 0184 eWatch; iMinds Medical Information Technologies: ICON: HBC.2016.0167 SeizeIT; Belgian Federal Science Policy Office IUAP #P7/19/ (DYSCO, 'Dynamical systems, control and optimization', 2012-2017); Belgian Foreign Affairs-Development Cooperation VLIR UOS programs (2013-2019); EU: European Union's Seventh Framework Programme (FP7/2007-2013): EU MC ITN TRANSACT 2012: #316679, The HIP Trial: #260777; Erasmus +: INGDIVS 2016-1-SE01-KA203-022114; European Research Council: The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Advanced Grant: BIOTENSORS (n° 339804). Thomas De Cooman is supported by an FWO SBO PhD grant.

Adaptive Heart Rate-Based Epileptic Seizure Detection Using Real-Time User Feedback

Thomas De Cooman^{1,2}, Troels W. Kjær³,
Sabine Van Huffel^{1,2}, Helge B. Sorensen⁴

1 STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Department of Electrical Engineering, KU Leuven, Leuven, Belgium

2 imec, Leuven, Belgium

3 Neurophysiology Center, Zealand University Hospital, Roskilde, Denmark

4 Department of Electrical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark

Thomas De Cooman, Kasteelpark Arenberg 10, 3001 Leuven, Belgium

E-mail: thomas.decooman@esat.kuleuven.be

August 2017

Abstract. Automated seizure detection in a home environment has been of increased interest the last couple of decades. Heart rate-based seizure detection is a way to detect temporal lobe epilepsy seizures at home, but patient-independent classifiers showed to be insufficiently accurate. This is due to the high patient-dependency of heart rate features, whereas this method does not use patient-specific data. Patient-specific classifiers take into account patient-specific data, but often not enough patient data are available for a full robust patient-specific classifier. Therefore a real-time adaptive seizure detection algorithm is proposed here. The algorithm starts with a patient-independent classifier, but gradually adapts to the patient-specific characteristics while they are obtained on-the-run. This is done by using real-time user feedback to annotate previously generated alarms, causing an immediate update to the used support vector machine classifier. Data annotated as seizures are automatically removed from the updating procedure if their detection would lead to too many false alarms. This is done in order to cope with potential incorrect feedback. The adaptive classifier resulted in an overall sensitivity of 77.12% and 1.24 false alarms per hour on over 2833 hours of heart rate data from 19 patients with 153 clinical seizures. This is around 30% less false alarms compared to the patient-independent classifier with a similar sensitivity. This low-complex adaptive algorithm is able to deal well with incorrect feedback, making it ideal for a seizure warning system, that in the future will also include complementary modalities to improve the performance.

Keywords: epilepsy, ECG, seizure detection, patient feedback, real-time adaptation
Submitted to: *Physiol. Meas.*

1. Introduction

Epilepsy is a neurological disease that affects around 1% of the people worldwide. For around 30% of these patients, their disease can not be controlled by anti-epileptic drugs. The question is how the quality of life of these patients can be improved in their everyday life. One solution is to use an automated warning system, which automatically detects an ongoing seizure and alarms the patients' caregivers, so they can give the patient the proper treatment. The key part of such a warning system is the automated seizure detection. In literature, epileptic seizure detection is typically done using the electroencephalogram (EEG), but obtaining EEG outside the hospital for long-term monitoring is not comfortable for the patient (Schulze-Bonhage et al. 2010). Therefore a lot of research has gone to automatically detect seizures using other easier obtainable signals such as accelerometers (ACM), electromyogram (EMG) and heart rate with wearable devices during the last couple of decades (Van de Vel et al. 2013).

Here, only the heart rate is used for seizure detection. Previous studies already showed that ictal heart rate changes could be found during most seizures (Zijlmans et al. 2002, Jansen & Lagae 2010). The ictal heart rate changes happen more frequently in temporal lobe epilepsy (TLE) patients than in other patients (Leutmezer et al. 2003) and are most often seen as expressions of strong sympathetic activations.

Most seizure detection algorithms using heart rate (and in general using all modalities) from literature use so-called patient-independent classifiers (De Cooman et al. 2017, Osorio 2014, Milosevic et al. 2016, Poh et al. 2012). They do not require patient-specific data, making them directly usable in practice. However, due to the large inter-patient variability in heart rate characteristics, these classifiers typically lead to too much false alarms to be used in practice (De Cooman et al. 2017). This issue can be solved by periodically changing the algorithm in order to change it to the patient-specific characteristics, which is nowadays assumed to be done after manual inspection. It is however very time-consuming and expensive to let a clinician evaluate previously recorded data of each patient and annotate potential seizure episodes. These annotations can only be made if they also contain reference video-EEG data, which means that this procedure should be done in a protected environment, lowering the usability in practice again.

Therefore, a real-time adaptive seizure detection algorithm is proposed here. The algorithm starts initially with a patient-independent classifier. Several currently available systems allow the usage of a prewarning, which allow the patient to indicate whether a warning might be correct or not before actually warning the relatives. If the patient can annotate it as a false alarm within a certain time limit (i.e. 10s), it generates no alarm to the relatives, otherwise the relatives are alarmed. This patient feedback can be used as annotations to these alarm data. The algorithm updates itself with this information in real-time, adapting to the patient characteristics over time. As patients might not be aware whether they actually had a seizure or not, extra procedures are added to the updating procedure to increase the robustness against incorrect user feedback. To our knowledge, this is the first paper that discusses real-time adaptive seizure detection using user feedback during full day monitoring. The proposed adaptive classifier can also be applied for seizure detection with other modalities such as ACM and EMG or as part of a multimodal system.

The three implemented approaches (patient-independent, patient-specific and real-time adaptive) are described in section 3. The used datasets are discussed in section 2. The results are shown and discussed in sections 4 and 5.

2. Data Acquisition

Two different datasets are used in this study. The first dataset contains data obtained from the University Hospital of Leuven (UZ Leuven), consisting of 17 patients with TLE and 127 complex partial seizures during in total 918 hours of recording. This dataset is only used for training as will be described in section 3.

The second dataset is the publicly available EPILEPSIAE database, which is used for the evaluation of the different algorithms (Ihle et al. 2012). It contains data from epilepsy centra from Freiburg (Germany), Coimbra (Portugal) and Paris (France). Every patient has at least 5 clinical seizures during recordings of at least 92 hours per patient. In total 19 patients with TLE are used from the EPILEPSIAE dataset containing scalp EEG data, including in total 158 clinical seizures in 2833.4 hours of single-lead electrocardiogram (ECG) data. Five seizure annotations were not used as the ECG was temporarily turned off just before, during or just after the seizure. ECG signals are measured on the chest with a sampling frequency of 256Hz. Seizures were annotated by expert neurologists using the gold video-EEG standard, and grouped in five categories: simple partial (SP), complex partial (CP), secondary generalized (SG), unclassified (UC) and subclinical (SC) seizures. Only the clinical seizures are examined here, as these are of most interest to detect in a home environment. The EEG seizure onsets are used as reference if available, otherwise the clinical onsets are used. An overview of the used datasets is given in the supplementary material.

3. Methodology

Three algorithmic methodologies are evaluated here: a patient-independent, a patient-specific and an adaptive approach. These approaches are now explained in more detail. All mentioned patient-specific data originates from the EPILEPSIAE database and all mentioned patient-independent data originates from the UZ Leuven database.

3.1. Patient-independent seizure detection

The first method is the patient-independent (P-I) algorithm. No patient-specific data are used in order to train the classifier of this algorithm, only data from other patients (from the UZ Leuven dataset) are used. Therefore this type of classifier can be used from the beginning of the monitoring period, making it very useful in practice.

3.1.1. Preprocessing and feature extraction Preprocessing and feature extraction are done as in De Cooman et al. (2017) according to the following steps. First, the heart rate is extracted from the acquired ECG signal. This is done by using a real-time R peak detection algorithm based on a dynamic threshold on the derivative signal. Next, a preprocessing step tries to detect the heart rate increases (HRI) in real-time. This method is based on heart rate gradient analysis and thresholds on the peak heart rate, HRI duration and the absolute increase of heart rate during the HRI. This procedure (called *HRI-EXTRACT*) is completely P-I. If such a strong HRI is detected, multiple heart rate features are extracted from this HRI and from one minute before the HRI. In De Cooman et al. (2014), it was already shown that the maximal heart rate gradient, HRI duration and the pre-HRI Hjorth activity (Hjorth 1970) are the 3 best features to use for automated P-I seizure detection with a linear support vector machine (SVM) classifier. Therefore only these features are used for classification here.

3.1.2. Classification Let $\{x_i, y_i\}$ be the training data with $x_i \in \mathbb{R}^d$ the input vector containing the selected feature data, and $y_i \in \{-1, +1\}$ the corresponding output values. Let classes +1 and -1 correspond to respectively non-seizure and seizure data samples. By restricting ourselves to linear SVM, similarly as in De Cooman et al. (2014), the classifier output \tilde{y}_i of data point x_i is then defined as $\tilde{y}_i = \text{sign}(w^T x_i + b)$ with $w \in \mathbb{R}^d$ the weight vector and $b \in \mathbb{R}$ the scalar bias term. The unknown variables w and b can be found by solving the following optimization problem (with adjustments to remove the class imbalance):

$$\min_{w, b, \xi} \frac{1}{2} w^T w + C^+ \sum_{i=1|y_i=+1}^N \xi_i + C^- \sum_{i=1|y_i=-1}^N \xi_i \quad \text{s.t.} \quad \begin{cases} y_i(w^T x_i + b) \geq 1 - \xi_i \\ \xi_i \geq 0 \end{cases} \quad (1)$$

in which ξ_i indicates the error on the classification of x_i and

$$C^+ = \frac{N^+ + N^-}{2 * N^+}, \quad C^- = \kappa * \frac{(N^+ + N^-)}{2 * N^-} \quad (2)$$

with N^+ and N^- indicating the number of data points belonging to classes +1 and -1. The factor κ (set to 2.5 according to De Cooman et al. (2017)) in C^- increases the impact of misdetecting seizure samples compared to non-seizure samples.

The classifier training procedure depends on the different methods described in the following sections. The P-I classifier is trained on the entire UZ Leuven dataset.

3.2. Patient-specific seizure detection

The (semi-)patient-specific (P-S) approach follows a similar procedure as the P-I method: The same preprocessing procedure and features are used here as in the P-I approach, only a different classifier is used. This classifier is used to indicate what the maximal performance could be for the adaptive approach (see section 3.3).

Because not enough P-S data is available in the EPILEPSIAE database to create a fully P-S classifier, a mixed model is used here. Data from the specific test patient (from the EPILEPSIAE database) is mixed with data from other patients (from the UZ Leuven database) in order to train the classifier. Misclassification of P-S data should have more impact on the optimization function (1), which is adapted into

$$\min_{w, b, \xi} \frac{1}{2} w^T w + C^+ \sum_{i=1|y_i=+1}^N \omega_i \xi_i + C^- \sum_{i=1|y_i=-1}^N \omega_i \xi_i \quad \text{with} \quad \omega_i = \begin{cases} \lambda & : i \in PS \\ 1 & : i \notin PS \end{cases} \quad (3)$$

with PS the collection containing the indices of the P-S data points in the collected dataset. This way, more importance is given to the P-S data points with a factor λ , which is set automatically by using leave-one-patient-out (LOPO) crossvalidation.

Training and testing of the P-S classifier is done using 5-fold crosstesting. The 4 folds of P-S data for training are mixed with P-I data from the UZ Leuven dataset to train the classifier, and is then evaluated on the fifth fold.

3.3. Real-time adaptive seizure detection

The adaptive seizure detection classifier differs from the P-S classifier due to real-time adaptation to the patient-specific characteristics, whereas the P-S classifier mentioned




Figure 1. Overview of the proposed real-time adaptive seizure detector.

in section 3.2 is trained offline with P-S data. The adaptive algorithm originally starts with the P-I classifier described in section 3.1. For each new P-S data point (thus for each HRI detected by *HRI-EXTRACT* extracted in the EPILEPSIAE database) that is detected, the adaptive classifier is immediately updated (see figure 1). Two SVM parameters are updated here: the weight matrix w and the bias term b . From now on we call w_t and b_t ($t \geq 1$) the t -th value of w and b after adding data point x_t to the classifier. The initial values w_0 and b_0 are the parameter values as found by the P-I classifier described in section 3.1. Although there exists an incremental way to update an SVM (Poggio & Cauwenberghs 2001), this method is too complex for real-time updating on limited hardware. Therefore, a simplified heuristic updating approach is used, which adapts to the P-S data equally fast and with lower complexity. This approach also includes procedures in order to lower the impact of incorrect user feedback, which is required if the system is used in practice.

For each new data point, it is first evaluated whether this data point is possibly caused by noise. This can occur if a noisy ECG segment leads to R peak detection errors. This kind of data should not cause a classifier update as it does not represent the real patient’s characteristics. A data point is assumed to be noise-free if less than 10% of all immediate heart rate values in the HRI differ more than 20% compared to the previous immediate heart rate value. This metric thus gives an indication of possible R peak detection errors. Ectopic heartbeats during a HRI happen only rarely during ictal HRIs, so that this procedure did not lead to missing any ictal HRI.

3.3.1. Updating w If the new data point x_t is assumed to be noise-free, w_t can be updated. In SVM, w can be written in the dual formulation as (Vapnik 1999)

$$w = \sum_{i=1}^N \alpha_i y_i x_i. \quad (4)$$

To lower the complexity of the updating procedure, we however stay in the primal SVM problem, and thus do not use P-S α_i (support vector values), but use the same principle. Let w_0 be the weight vector from the P-I classifier. If a P-S data point x_t causes a classifier update, w_t^{nn} (“nn” standing for “not normalized”) is set to

$$w_t^{nn} = w_{t-1}^{nn} + \gamma_t x_t y_t \quad \text{with } \gamma_t = \begin{cases} 1 & : y_t = -1 \\ 0.1 & : y_t = +1 \end{cases} \quad (5)$$

and $w_0^{nn} = w_0$. The γ parameter replaces the original α parameter here. For simplicity, γ_t can only take 2 values, which are chosen to have a similar ratio as the maximum

values the α values can get in the P-I classifier (see equation (2), $\alpha_i \leq C_i$). That way the γ parameters get larger values for the seizure samples as well. Normally only data points causing an error have $\alpha_i \neq 0$ in SVMs. Similarly, the update of w_t only occurs here if $|w_{t-1}^T x_t + b_{t-1}| \leq 1$. This way, also strong outliers will not cause an update of w , increasing the robustness of the SVM updates.

To optimally benefit from the knowledge from the P-I classifier, the adaptive classifier should not deviate too much from it. It can be assumed that each P-S classifier shows sufficient resemblance with the P-I classifier. Therefore, w_t is restricted so that its norm doesn't change more than σ compared to w_0 ($\|w_t - w_0\| < \sigma$):

$$w_t = w_0 + \frac{\sigma}{\|w_t^{nn} - w_0\|} (w_t^{nn} - w_0) \quad (6)$$

This is an assumption that is also used in transfer learning approaches (Yang et al. 2007), which benefit optimally from a reference classifier (here the P-I classifier) in order to get a stable classifier with limited training data. By including it in the adaptive approach, the classifier remains more stable and reliable during all classifier updates. The value of σ is automatically set for each patient using LOPO crossvalidation on all patients of the EPILEPSIAE database except the test patient. If no update occurs after x_t , w_t is set to w_{t-1} .

3.3.2. Updating b Once w is updated, the bias term b also needs to be updated. b_t can be found by choosing the value that leads to the minimal sum of errors ξ_i , with

$$\xi_i = \max(y_i(w_t^T x_i + b'), 0), 1 \leq i \leq N \quad (7)$$

The errors ξ_i are only evaluated for $b' \in [b_{t-1} - \epsilon, b_{t-1} + \epsilon]$ to lower the complexity, preferring a smooth transition from P-I towards P-S classification. Parameter ϵ should be a small value to maintain a smooth classifier transition (empirically set as 0.1). The errors are only evaluated for a fixed number of data points N , which is a mixture of P-I and P-S data. Every time a new P-S data point is annotated, a P-I data point from the same class is removed from memory and computation. This is done in order to limit the computation and memory requirements of the updating procedure.

Let $E_{b'}$ be the averaged error caused by constant term b'

$$E_{b'} = \frac{\sum_{i=1|y_i=+1}^N \omega_i \xi_i}{\sum_{i=1|y_i=+1}^N \omega_i} + 1.5 * \frac{\sum_{i=1|y_i=-1}^N \omega_i \xi_i}{\sum_{i=1|y_i=-1}^N \omega_i} \quad (8)$$

The weights ω_i give extra weight to the P-S data points as in (3). The class imbalance is removed by using weighted averages of both classes. The experimentally found factor 1.5 gives seizure data more impact on the b update. The new b_t is then set as

$$b_t = \arg \min_{b'} E_{b'} \quad (9)$$

After this update, the classification rule for the next data point x_{t+1} then becomes

$$\tilde{y}(x_{t+1}) = \text{sign}(w_t^T x_{t+1} + b_t) \quad (10)$$

The usage of annotations in the updating procedure has two main drawbacks. First, the annotations made by the user might be incorrect, leading to bad classifier

updates. Secondly, short seizures are too hard to detect with heart rate-based seizure detection as they lead to very short HRIs, making it hard to differentiate them from non-epileptic HRIs. They are often also of less importance for real-time seizure detection applications due to their low risk for injuries. Therefore they should not be given the same importance as longer seizures during the update procedure, certainly because the adaptive approach might base its model with just 1 P-S seizure. In order to remove the impact of these issues, an extra procedure is able to detect most of these smaller or incorrectly annotated ictal HRIs. If the output $w_t^T x_i + b_t$ (with annotation $y_i = -1$) of x_i (P-S data point) is larger than 20% of the outputs from the already collected P-S non-seizure data points, ω_i is set to 0 in equation (8) in order to remove it from updating, as it is unclear whether it really represents seizure data, or might be an atypical seizure for that patient. In this case, the data point annotated as a seizure resembles non-epileptic P-S data too much and is therefore unlikely to be accurately detectable with the proposed method. Therefore it is removed from the updating procedure. Incorrectly annotated data should not cause a classifier update, and too short seizures leading to too much false alarms after a potential update should not cause a drastic change in the classifier.

3.3.3. Feedback procedure The real-time SVM updates require the availability of the annotations y_t of each P-S data point x_t . They are made by the patients or their caregivers for all P-S HRIs that are detected by *HRI-EXTRACT*, not only the HRIs that cause an alarm. We assume the feedback procedure to go as follows in practice:


- An alarm is generated: The user gives feedback about the alarm correctness. If no response is given within a given time period (i.e. 10s), the alarm is assumed to be correct as the patient is not able to cancel the alarm during a seizure.
- A seizure is missed: The user can indicate that a seizure is recently missed. In this case, the algorithm can keep information about the HRIs of the last hour. These HRIs are then marked as potential seizure HRIs. The procedure to remove falsely annotated seizure HRIs eliminates most of the falsely annotated seizure HRIs for SVM updating. If no such ictal HRI is found, no classifier update occurs.
- A non-seizure HRI is correctly ignored: If the procedure for a missed seizure is not activated, the HRI is annotated as non-seizure and the classifier is updated.

3.4. Evaluation criteria and additional tests

All classifiers are evaluated only on the EPILEPSIAE database. Four metrics are used for evaluating seizure detection algorithms. The sensitivity (Se) indicates the percentage of detected seizures. The number of false alarms per hour (FP/h) shows how often a FP is generated in time, whereas the positive predictive value (PPV) indicates the percentage of correct detections. The detection delay indicates the time difference between the seizure onset and the moment of detection. A seizure is said to be detected if an alarm is generated between 60 s prior to the seizure onset and 90 s after the seizure onset (De Cooman et al. 2017). Different alarms within one minute are clustered as one alarm. The patient feedback is assumed to be 100% accurate unless specifically mentioned otherwise. To evaluate the impact of incorrect feedback, an additional test is performed when a certain percentage of user feedback is incorrect. For each percentage, 40 runs were simulated in which P-S data was randomly chosen to be incorrectly annotated with the corresponding probability. In order to prove

Table 1. Overview of the results of the different seizure detection classifiers.

Method	Se (%)	FP/h	PPV (%)
Preprocessing procedure (<i>HRI-EXTRACT</i>)	84.31	3.63	1.24
Patient-independent approach	78.43	1.73	2.39
Patient-specific approach	76.47	1.09	3.67
Adaptive approach	77.12	1.24	3.25

**Figure 2.** Sensitivity and FP/h results per patient for the different classifiers.

that the adaptive classifier performs better than the P-I classifier, paired t-tests (with corresponding confidence intervals (CI)) are performed on Se, FP/h and PPV.

4. Results

4.1. Performance of the different seizure detection classifiers

Table 1 gives an overview of the results of the different real-time seizure detection approaches on the EPILEPSIAE database. 84.31% of the seizures were detected with the preprocessing procedure *HRI-EXTRACT*, which indicates the percentage of seizures with ictal strong HRIs. The P-I classifier resulted in a sensitivity of 78.43% with on average 1.73 FP/h. Seizures were detected on average 18s after seizure onset. The P-S approach resulted in only 1.09 FP/h with a similar sensitivity performance. By using the adaptive approach, only on average 1.24 FP/h were found, again with a similar sensitivity as the P-S approach. This is a drop of nearly 30% of false alarms compared to the P-I classifier. Figure 2 shows the results per patient for the 3 approaches. The results for the false alarm rate (FAR) ($p < 10^{-3}$, CI [-0.81,-0.26]) and PPV ($p < 10^{-4}$, CI [0.74%,1.57%]) from the adaptive classifier prove to be significantly better than those of the P-I classifier, but not for Se ($p = 0.38$, CI [-3.71%,1.48%]). Both P-S and adaptive classifier resulted in an average detection delay of around 19s.

The purpose of the adaptive approach is to adapt to the P-S characteristics over time. Figure 3(a) shows the average FAR per 12 hours for the three approaches for two patients. For patient 1, the adaptive classifier starts with a FAR close to the P-I classifier. When more data is collected from the test patient, the results start to converge towards P-S performance, which is accomplished around 5 days after the

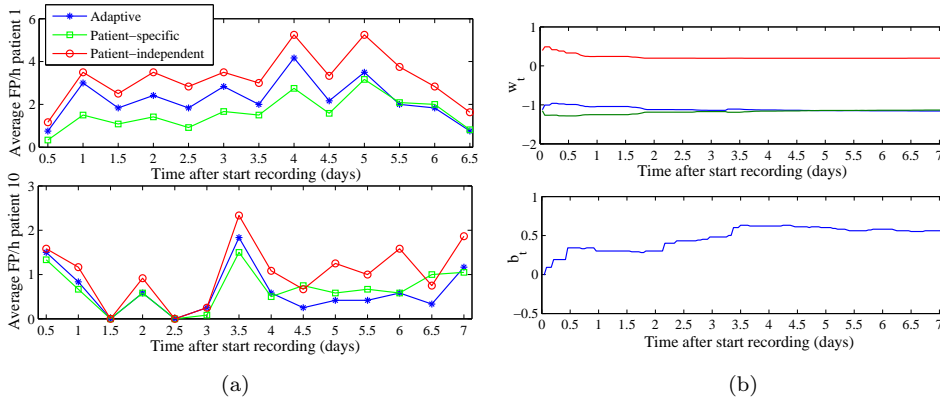


Figure 3. Illustration of convergence for the adaptive classifier. (a) Average FP/h per 12 hours of recording in function of time for patients 1 and 10 for the three different approaches. (b) Example of the SVM parameter convergence for patient 10. The different lines in the w_t figure indicate the terms for the 3 features.

start of the recording. In case of patient 10, a similar pattern can be seen, but the adaptive classifier is able to show even a lower FAR than the P-S version after 4 days. A similar figure is shown in the supplementary material of this paper with the FAR averaged over all patients, illustrating that on average also improvement can be found over time. Figure 3(b) shows the SVM parameter convergence for patient 10, which reaches a nearly stable SVM classifier after 4 days of recording. The time of convergence typically depends on when sufficient P-S seizure data was collected.

4.2. Seizure sensitivity

Figure 4(a) gives an overview of the percentage of detected seizures for the different seizure types for both *HRI-EXTRACT* and the adaptive seizure detection approach. In the results for *HRI-EXTRACT*, it can be seen that $\pm 84\%$ of the SP, CP and UC seizures have significant ictal HRIs. All 4 generalized seizures show ictal strong HRIs.

The overall sensitivity however dropped by using the adaptive classifier. Mostly extra SP and UC seizures are missed by this procedure, resulting in a bigger difference in sensitivity for the CP and SG seizures compared to the SP and UC seizures.

Seizure duration also has an impact on the sensitivity. Figure 4(b) shows the sensitivity of the seizures based on their seizure duration for *HRI-EXTRACT* and the adaptive approach. The longer the seizure lasts, the higher the chance of detecting it with the adaptive classifier. The seizures that are missed extra in comparison with *HRI-EXTRACT* are mostly relatively short seizures ($< 60s$).

4.3. Impact of incorrect user feedback on real-time adaptive seizure detection

For now the assumption was made that all given user feedback was correct. In practice however, this will not be the case. Figure 5(a) shows the effect if only a certain percentage of user feedback is correct. From 10% correct feedback on, the PPV value increases slowly towards optimal PPV performance. The sensitivity does not change much for different feedback percentages, but mainly the FAR changes drastically. However, the overall effect of incorrect user feedback is rather limited. Figure 5(b)




Figure 4. Impact factors on seizure sensitivity for *HRI-EXTRACT* and the proposed adaptive seizure detection classifier. (a) Impact of seizure type on sensitivity. The numbers between brackets indicate the number of seizures per type. (b) Impact of seizure duration on sensitivity. (c) Boxplots of the seizure duration for each seizure type analyzed in the dataset.

shows the impact of incorrect user feedback for both the proposed updating procedure and the standard updating procedure (Poggio & Cauwenberghs 2001). If the standard SVM updating procedure is used with 100% correct feedback, 78.43% sensitivity is achieved with 1.48 FP/h and 2.78% PPV, which is 0.24 FP/h more and 0.47% PPV less than with the proposed updating method. The performance of the standard updating procedure decreases very fast with only 0.1% incorrect user feedback, and with 10% incorrect feedback the PPV is already equally bad as for the preprocessing procedure *HRI-EXTRACT*. The proposed method however only shows a small drop in PPV here, with still a PPV of 3.05% at 10% incorrect user feedback.

5. Discussion

5.1. Comparison between the different approaches

If the P-I classifier is retrained on the UZ Leuven dataset in order to get the same sensitivity as the adaptive classifier on the validation dataset, the P-I classifier results in 77.78% sensitivity with 1.63 FP/h on the EPILEPSIAE database. This is an increase of almost 0.4 FP/h with a nearly identical sensitivity compared to the adaptive classifier. The adaptive classifier also shows results close to the assumed maximal P-S performance of the adaptive classifier, leading to on average only 0.15 FP/h more than the P-S performance. The adaptive approach indeed shows results that come close to P-S optimal performance after less than one week of recorded data with around 3-5 seizures. This shows that the used procedure to adapt to the patient characteristics leads to a nearly as efficient classifier as the P-S classifier, which is normally adapted offline by a neurologist or nurse. This procedure is thus more time and cost effective than the standard way of personalization. The proposed annotation procedure does not cause any extra burden to the patient if it is added to a system that already uses prewarnings. This adaptive approach just takes into account information that is already available in such systems, and will also lead to a lower FAR in practice.

Figure 2 shows the results for the different approaches for every patient. In most cases the FAR for the adaptive classifier lies between the FAR of the P-I and P-S classifiers with similar sensitivity results. For some patients the FAR is even lower for the adaptive classifier than for the P-S classifier. The reason for this is that the P-S classifier does not have a procedure to remove short or abnormal seizures

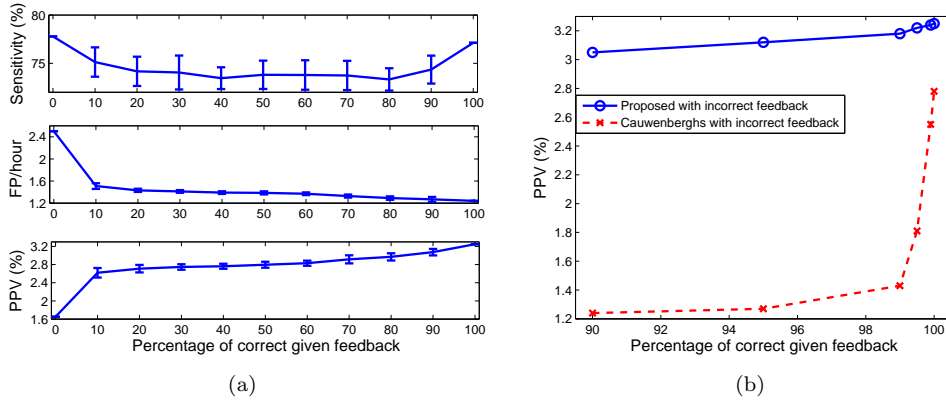


Figure 5. Impact of incorrect feedback on the performance. (a) Effect of different percentages of incorrect user feedback on sensitivity, FP/hour rate and PPV for the proposed adaptive seizure detection approach. The variation in performance between different runs is shown by the variation bars. (b) Comparison of the PPV values for the proposed updating procedure and the traditional SVM updating procedure in case 0, 0.1, 0.5, 1, 5 and 10% of user feedback is incorrect.

from the training set, whereas such a procedure was implemented for the adaptive approach. This procedure removed seizures that were unlikely to be detected from the training set, as adding them to the classifier would lead to an underestimate of the SVM boundary, leading to too many false alarms then strictly necessary. Normally atypical seizure HRIs would only have limited effect during SVM training if sufficient P-S seizures are already collected, but in the incremental approach it can occur that only one (abnormal) seizure might be available for a long time. By making this choice based on the collected non-seizure HRI data, this decision can be made much faster with sufficient accuracy. A second cause for this is that by making the assumption that the adaptive classifier should have a w_t sufficiently similar to the P-I w_0 , a better classifier is found with the adaptive classifier. In some cases, training the P-S classifier led to strong changes in w compared to the P-I w_0 , which often resulted in a slightly worse performance. An example of this can also be seen in figure 3(a), for which after 4 days the adaptive classifier becomes better than the P-S classifier due to overtraining of P-S classifier on several SP seizures which are too hard to detect.

5.2. Seizure sensitivity

Figure 4 shows that the sensitivity changed for both seizure duration and seizure type. In order for a seizure to be accurately detectable using the heart rate, the seizure should indeed last long enough, so that the seizure HRIs can be sufficiently different from non-seizure HRIs (De Cooman et al. 2017, Hampel et al. 2016). Both reasons are however correlated in this database as shown in figure 4(c). Both CP and SG seizures were on average longer than SP and UC seizures, making them easier to detect with heart rate-based seizure detection. This is not only applicable between patients, but also within patients. An example of this can be found in patient 10, for which all CP seizures are detected, but all SP seizures were missed.

There is not only variability between seizure types, but also between patients. Most patients show to have sensitivities of 70-100%, but for 2 patients the sensitivity

was lower than 40% for all approaches. For patient 3, all ictal HRIs showed to be insufficiently strong compared to inter-ictal HRIs, causing all seizures to be removed from the update procedure as they were marked as too short. Therefore this patient has a similar high FAR as the P-I approach. Patient 18 often had ictal bradycardia (sometimes following a small HRI), with an ictal ideoventricular rhythm, which is a potentially life threatening heart condition. Although these seizures were not detected by the algorithm due to the lack of a strong HRI, it shows the benefits of monitoring the heart rate in epileptic patients in a home environment, as it allows to detect heart failure (if additional algorithms are applied) (Van de Vel et al. 2016).

5.3. Impact of incorrect user feedback on performance

Patients might not be aware that they had a seizure, or might not be certain about it. If they do not have somebody accompanying them to indicate whether they indeed had a seizure or not, this could potentially lead to incorrect user feedback.

The impact of incorrect feedback showed to be limited if enough feedback was correct. The robustness against incorrect feedback is obtained by making 2 changes to normal SVM updates. The first one is the use of the restriction $\|w_t - w_0\|$ during the updates. This avoids the adaptive classifier to become worse than the original P-I classifier by assuming that in the P-S approach still most of the information of the P-I approach applies (i.e. a longer HRI leads to a higher seizure probability). That way, the classifier still gives a mediocre performance even if all feedback is incorrect.

The second change was the addition of the rule to remove data points x_i with $y_i = -1$ from updating b_t if the output $w_t x_i + b_{t-1}$ was higher than 20% of the non-seizure P-S data. This not only limits the impact of feedback errors, but also limits the impact of short seizures, again leading to a better performance. As in the real-time adaptive approach only 1 seizure can be available to adapt to, this step can be crucial for a robust but fast adaptation. The impact of annotating a seizure HRI as non-seizure is limited due to the large amount of non-seizure data, leading to a limited change in the classifier update. Figure 5(b) shows these changes did not only make the adaptive classifier more robust against incorrect feedback, it also gave better results than the normal SVM updating procedure in case of 100% feedback accuracy.

By incorporating a heuristic updating approach, the complexity of the updating procedure is much lower than for the standard SVM updating procedure described in (Poggio & Cauwenberghs 2001). Therefore this approach is more usable for real-time updating with limited hardware. Although it was only evaluated here on linear SVM, it can also be applied to other kernels if an estimation of the kernel function is used.

5.4. General discussion and future work

Section 4.3 shows that although giving correct feedback is important for optimal performance, the adaptive approach was able to cope well with large percentages of incorrect feedback. In cases of doubt, patients giving incorrect feedback will not have a drastic effect on the performance, which would be the case with normal updating procedures. A procedure is added here to automatically check if trying to add this seizure data to the SVM model would result in too many false alarms. This way, the model takes care of removing data caused by too short seizures, as they are very hard to detect using only the heart rate and are less important to detect due to their lower risk for injuries. These procedures make the algorithm more useful in practice in case

no video-EEG is available to give true data annotations, which is ideal for usage in a home environment using wearable devices.

The adaptive approach showed to adapt to the P-S heart rate characteristics. Each patient has its own heart rate variability, both ictal and interictal, and incorporating this P-S information leads to an improved performance compared to the P-I performance. However, the results can still be further improved. One of the drawbacks of the system is that we try to get a P-S optimal performance while using features that showed to be good for P-I use. Better results might be obtained by also including features that work better in a P-S setting, such as for example the achieved ictal peak heart rate or the modified cardiac sympathetic index (De Cooman et al. 2015, Jeppesen et al. 2015). This is however left for future work.

Despite the increased performance and usability of the heart rate-based seizure detection approach, the unimodal performance is still too low in order to be used successfully in practice. One of the reasons is the fact that the ECG is sensitive to motion artifacts, making it hard to accurately measure the heart rate from it. In order to accomplish better results, it is required to add additional signal modalities. Previous studies showed that using data from only 1 signal modality for automated seizure detection in a home environment are often insufficient in order to obtain performance that allows usage in practice (Milosevic et al. 2016, Poh et al. 2012). ACM and/or EMG sensors for example can be of great added value for the detection of strong convulsive seizures (Milosevic et al. 2016, Conradsen et al. 2012). Implementation of a similar adaptive approach for these modalities could also increase their unimodal performance. A combination of several adaptive approaches on different modalities will lead to a much stronger detection algorithm which is advised to be used in practice.

6. Conclusion

An adaptive seizure detection algorithm using only heart rate information is proposed. The initial patient-independent classifier was adapted to the patient-specific characteristics by using new patient-specific data points which are annotated with user feedback. The adaptive approach showed to rapidly increase in performance, while also being robust against potential incorrect feedback. This algorithm not only leads to a better performance, but also leads to better usability of such a system in practice. Features from other modalities will also be incorporated in the model in future work.

Acknowledgments

Research supported by Bijzonder Onderzoeksfonds KU Leuven (BOF): Center of Excellence (CoE): PFV/10/002(OPTEC), SPARKLE #: IDO-13-0358, C24/15/036, C32/16/00364; Fonds voor Wetenschappelijk Onderzoek-Vlaanderen (FWO): G.0A55.13N (Deep brain stimulation); Agentschap Innoveren & Ondernemen (VLAIO): STW 150466 OSA +, O&O HBC 2016 0184 eWatch; iMinds Medical Information Technologies: ICON: HBC.2016.0167 SeizeIT; Belgian Federal Science Policy Office IUAP #P7/19/ (DYSCO, 'Dynamical systems, control and optimization', 2012-2017); Belgian Foreign Affairs-Development Cooperation VLIR UOS programs (2013-2019); EU: European Union's Seventh Framework Programme (FP7/2007-2013): EU MC ITN TRANSACT 2012: #316679, The HIP Trial: #260777; Erasmus +: ING-DIVS 2016-1-SE01-KA203-022114; European Research Council: The research leading to these results has received funding from the European Research Council under the

European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Advanced Grant: BIOTENSORS (n° 339804). This paper reflects only the authors' views and the Union is not liable for any use that may be made of the contained information. Thomas De Cooman is supported by an FWO SBO PhD grant.

References

- Conradsen, I., Beniczky, S., Hoppe, K., Wolf, P. & Sorensen, H. B. (2012). Automated algorithm for generalized tonic-clonic epileptic seizure onset detection based on semg zero-crossing rate, *Biomedical Engineering, IEEE Transactions on* **59**(2): 579–585.
- De Cooman, T., Carrette, E., Boon, P., Meurs, A. & Van Huffel, S. (2014). Online seizure detection in adults with temporal lobe epilepsy using single-lead ECG, *Proceedings of the 22nd European Signal Processing Conference (EUSIPCO), 2014*, IEEE, pp. 1532–1536.
- De Cooman, T., Van de Vel, A., Ceulemans, B., Lagae, L., Vanrumste, B. & Van Huffel, S. (2015). Online Detection of Tonic-Clonic Seizures in Pediatric Patients using ECG and Low-Complexity Incremental Novelty Detection, *Proc. of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC2015)*, pp. 5597–5600.
- De Cooman, T., Varon, C., Hunyadi, B., Van Paesschen, W., Lagae, L. & Van Huffel, S. (2017). Online automated seizure detection in temporal lobe epilepsy patients using single-lead ecg, *International Journal of Neural Systems* p. 1750022.
- Hampel, K. G., Jahanbekam, A., Elger, C. E. & Surges, R. (2016). Seizure-related modulation of systemic arterial blood pressure in focal epilepsy, *Epilepsia* .
- Hjorth, B. (1970). EEG analysis based on time domain properties, *Electroencephalography and clinical neurophysiology* **29**(3): 306–310.
- Ihle, M., Feldwisch-Drentrup, H., Teixeira, C. A., Witon, A., Schelter, B., Timmer, J. & Schulze-Bonhage, A. (2012). Epilepsiae—a european epilepsy database, *Computer methods and programs in biomedicine* **106**(3): 127–138.
- Jansen, K. & Lagae, L. (2010). Cardiac changes in epilepsy, *Seizure* **19**(8): 455–460.
- Jeppesen, J., Beniczky, S., Johansen, P., Sidenius, P. & Fuglsang-Frederiksen, A. (2015). Detection of epileptic seizures with a modified heart rate variability algorithm based on lorenz plot, *Seizure* **24**: 1–7.
- Leutmezer, F., Scherthaner, C., Lurger, S., Pötzelberger, K. & Baumgartner, C. (2003). Electrocardiographic changes at the onset of epileptic seizures, *Epilepsia* **44**(3): 348–354.
- Milosevic, M., Van de Vel, A., Bonroy, B., Ceulemans, B., Lagae, L., Vanrumste, B. & Van Huffel, S. (2016). Automated detection of tonic-clonic seizures using 3-d accelerometry and surface electromyography in pediatric patients, *IEEE Journal of Biomedical and Health Informatics* **20**(5): 1333–1341.
- Osorio, I. (2014). Automated seizure detection using EKG, *International Journal of Neural Systems* **24**(02): 1450001.
- Poggio, T. & Cauwenberghs, G. (2001). Incremental and decremental support vector machine learning, *Advances in neural information processing systems* **13**: 409.
- Poh, M.-Z., Loddenkemper, T., Reinsberger, C., Swenson, N. C., Goyal, S., Sabtala, M. C., Madsen, J. R. & Picard, R. W. (2012). Convulsive seizure detection using a wrist-worn electrodermal activity and accelerometry biosensor, *Epilepsia* **53**(5): e93–e97.
- Schulze-Bonhage, A., Sales, F., Wagner, K., Teotonio, R., Carius, A., Schelle, A. & Ihle, M. (2010). Views of patients with epilepsy on seizure prediction devices, *Epilepsy & behavior* **18**(4): 388–396.
- Van de Vel, A., Cuppens, K., Bonroy, B., Milosevic, M., Jansen, K., Van Huffel, S., Vanrumste, B., Lagae, L. & Ceulemans, B. (2013). Non-ecg seizure-detection systems and potential sudep prevention: State of the art, *Seizure* **22**(5): 345–355.
- Van de Vel, A., Smets, K., Wouters, K. & Ceulemans, B. (2016). Automated non-ecg based seizure detection: Do users have a say?, *Epilepsy & Behavior* **62**: 121–128.
- Vapnik, V. (1999). *The nature of statistical learning theory*.
- Yang, J., Yan, R. & Hauptmann, A. G. (2007). Adapting svm classifiers to data with shifted distributions, *Seventh IEEE International Conference on Data Mining Workshops (ICDMW 2007)*, IEEE, pp. 69–76.
- Zijlmans, M., Flanagan, D. & Gotman, J. (2002). Heart rate changes and ECG abnormalities during epileptic seizures: prevalence and definition of an objective clinical sign, *Epilepsia* **43**(8): 847–854.