

A Micropower Support Vector Machine Based Seizure Detection Architecture for Embedded Medical Devices

Ali Shoeb^{1,2}, Dave Carlson², Eric Panken², Timothy Denison²

¹Massachusetts Institute of Technology, Boston, MA, 02139, USA (ashoeb@mit.edu)

²Medtronic Neuromodulation Technology, Minneapolis, MN, 55410, USA (timothy.denison@medtronic.com)

Abstract—Implantable neurostimulators for the treatment of epilepsy that are capable of sensing seizures can enable novel therapeutic applications. However, detecting seizures is challenging due to significant intracranial EEG signal variability across patients. In this paper, we illustrate how a machine-learning based, patient-specific seizure detector provides better performance and lower power consumption than a patient non-specific detector using the same seizure library. The machine-learning based architecture was fully implemented in the micropower domain, demonstrating feasibility for an embedded detector in implantable systems.

I. INTRODUCTION

Epilepsy is a disorder of the central nervous system that predisposes individuals to experiencing seizures. Despite advances in pharmacologic treatments, approximately 1 in 3 patients continues to experience frequent seizures [1]. Today, these individuals can supplement their pharmacologic regimen with electrical stimulation of the brain to better control their seizures. Much research has shown that periodic, electrical excitation of various cerebral sites impacts the frequency of seizures [2].

An implantable neurostimulator capable of detecting and reacting to the onset of a seizure can facilitate applications that are not possible using devices that passively deliver a periodic stimulus. The enhanced stimulator could 1) Alert patients of the electrical onset of a seizure before the development of clinical symptoms 2) Align a stimulus with the onset of a seizure to suppress the seizure 3) Maintain an account of seizure activity so that physicians can objectively determine the efficacy of stimulator settings.

Performing rapid and reliable seizure onset detection on an implantable neurostimulator is challenging. Algorithms for detecting seizure onset within Intracranial EEG (IEEG) signals must be capable of detecting seizures given substantial variability of seizure characteristics across patients [3], and similarity between seizure and non-seizure activity both within and across patients. At the same time, such algorithms need to consume a small amount ($\sim 50\mu\text{W}$) of the total power budget of the neurostimulator [4].

In this paper we address the computational and implementation challenges associated with detecting seizure onset on an implantable device. We demonstrate how machine learning techniques, specifically Two-Class Support Vector Machines (SVM) [5], can be used to synthesize *patient-specific* detectors that outperform a *patient non-specific* detector based on [6]. We also discuss

Reduced Set Methods [7-8], a model-order reduction technique that enables efficient implementation of the discriminant functions produced by the Support Vector Machine learning algorithm. Finally, we compare the power consumption associated with implementing the patient-specific and patient non-specific algorithms on a neural integrated circuit (IC) [4] capable of low power extraction of IEEG signal features key to both algorithms.

II. SEIZURE ONSET DETECTION METHODS

A. Intracranial EEG

Seizure onset detection is complicated by variability of seizure characteristics across patients [3] and similarity between seizure and non-seizure activity within and across patients. For example, Figure 1 shows how seizures from patients A and B differ both in IEEG signal amplitude and spectral content; while Figure 2 shows the degree to which these characteristics can be similar for seizure and non-seizure activity recorded from the same individual.




Figure 1: Seizure of Patient A (top panel following dotted line) differs from seizure of Patient B (bottom panel) in signal amplitude and spectral content.




Figure 2: Spectral content of Non-seizure activity (bottom panel) recorded from Patient C overlaps spectral content of seizure activity (top panel) recorded from the same patient.

For a given patient, however, seizure and non-seizure activity tend to be stereotypical [3]; meaning an individual's

seizure looks similar whenever it reoccurs. This is demonstrated in Figure 3 which shows a second seizure from patient A.




Figure 3: The second seizure from Patient A resembles the first seizure in signal amplitude and spectral content.

The variability in seizure characteristics across patients and their consistency within a patient motivated us to adopt a patient-specific approach to seizure detection. Such an approach uses pre-recorded seizure and non-seizure signals to “teach” an algorithm to differentiate between an individual’s seizure and non-seizure activity. The success of this approach depends on the extent to which an individual’s seizure and non-seizure activity are similar and on the consistency of an individual’s seizure signature.

B. Patient-Specific Seizure Onset Detection

The architecture of the patient-specific detector is illustrated in Figure 4. The detector extracts, from each channel, features that can be used to infer the presence of seizure activity. Since spectral energy has been shown to be useful in the context of intracranial seizure detection [6, 7], the chosen features were the energy within the frequency bands 0-16Hz and 15-37Hz. The spectral features extracted from each channel are then concatenated to form a 4-dimensional feature vector in order to capture dependencies between the input channels. Finally, a support-vector machine (SVM) trained to differentiate between a patient’s seizure and non-seizure activity classifies the feature vector.

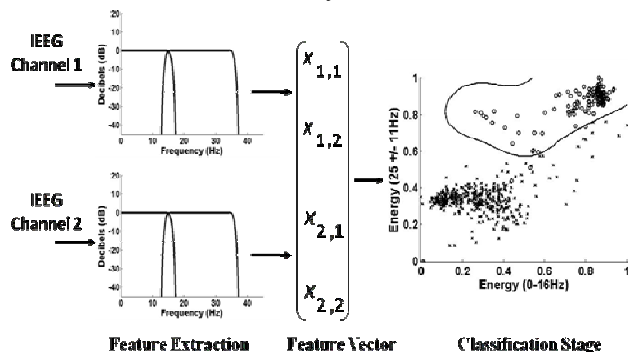


Figure 4: The Patient-specific detector architecture.

The SVM determines whether an observed feature vector is representative of seizure or non-seizure activity based on which side of a decision boundary the feature vector falls. In Figure 4, we plot a nonlinear decision boundary separating seizure feature vectors (circles) from non-seizure feature vectors (crosses) in a two-dimensional space for the purpose of visualization; in the actual system classification occurs in a four-dimensional space. The decision boundary computed by the SVM learning algorithm can also be linear as illustrated in Figure 5.




Figure 5: Linear and Nonlinear SVM decision boundaries.

Mathematically, the discriminant function that establishes whether a feature vector X falls within the region of the feature space defined by seizure activity is expressed as follows in the case of a linear boundary

$$W^T X + \beta > 0$$

Where the vector W and the bias term β are parameters determined by the SVM learning algorithm. In the case of a nonlinear boundary, determined using a radial basis kernel, the discriminant function takes the form shown below. The N coefficients α_i , support-vectors X_i , and bias term β are determined by the SVM learning algorithm, and the parameter $\gamma = 0.1$ is user defined.

$$\left\{ \sum_{i=1}^N \alpha_i \exp(-\gamma \|X - X_i\|^2) \right\} + \beta > 0$$

The number of support-vectors N is partly governed by the complexity of the classification task. As the similarity between a patient’s seizure and non-seizure activity increases more support-vectors are needed in order to define a more complex decision boundary. For our dataset, N was found to be within the range of 50-300, which makes implementation very costly.

The application of Reduced Set Methods [8], a model order reduction technique, allows the nonlinear discriminant function to be expressed using $M \ll N$ support-vectors. As an example, Figure 6 shows the superposition of a nonlinear boundary requiring $N=50$ terms (solid line) and an approximation of that boundary using $M=8$ terms (dashed line). Our evaluation of nonlinear decision boundaries used $M=8$ terms.




Figure 6: Approximation of a nonlinear SVM boundary using Reduced Set Methods

C. Patient Non-specific Seizure Onset Detection

The patient non-specific detector evaluated in this paper is a simplification of the Osorio-Frei seizure detector [4,6]; the architecture of this detector is shown in Figure 7. The detector monitors, for each channel, the ratio of current and background energy in the 15-37Hz frequency band; the current energy is computed using a short time window of 2 seconds while the background energy is computed using a long time window of 30 minutes. Whenever the ratio of foreground to background energy, on any given channel, exceeds a threshold γ for T seconds a seizure is declared.

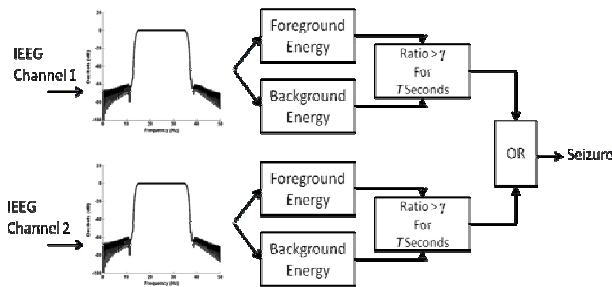


Figure 7: The Patient non-specific detector architecture.

III. TESTING METHODOLOGY AND RESULTS

A. Data Set

The data set used to evaluate the performance of the two detectors included 81 hours of intracranial EEG collected from 17 adult subjects. On average, 4.5 hours of recording time containing 3 seizures were available per patient. For each patient, an electroencephalographer identified the onset time of all seizures as well as the two IEEG channels demonstrating the earliest signs of seizure activity; these two channels were the only channels processed by both algorithms and varied across patients.

B. Testing Methodology

Due to the small number of seizures available for each patient, a leave-one-out testing methodology was adopted. Consider a patient recording consisting of K ten-minute blocks of IEEG data containing L seizures. The patient-specific detector would be trained on $K/2$ data blocks containing $L-1$ seizures. Next, the performance of both the patient-specific and patient non-specific detectors is assessed on the remaining $K/2$ blocks containing the L^{th} seizure. This is repeated L times so that the ability of both algorithms to detect each of the patient's L seizures is tested. Once all trials are run, we note the number of seizures detected; the average latency with which seizures are detected; and the number of false detections across the L trials for both algorithms. A false detection is defined as a detection made during any window of time noted by an electroencephalographer to be free of seizure activity.

C. Performance Comparison

The performance of the patient-specific detector using both linear and nonlinear SVMs was compared to the

performance of the patient non-specific detector set to declare a seizure following $T=3$ seconds and $T=10$ seconds.

Figure 8 illustrates the average latency with which the patient-specific and non-specific detectors declare the onset of a seizure for each of the 17 subjects. The patient-specific detector that used a nonlinear SVM (reduced to use $M=8$ terms) detected 60/61 seizures within 8.5 ± 5.1 seconds; the same detector using a linear SVM detected 60/61 seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm detected 55/61 seizures within 10.5 ± 8.7 seconds. Note that the patient non-specific detector failed to detect the seizures of Patient 3 which accounts for the absence of a bar.

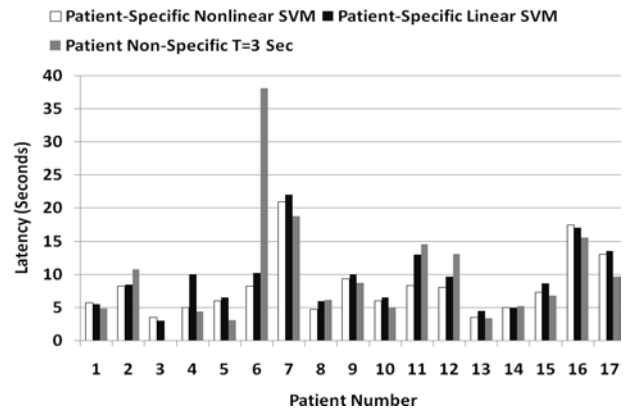


Figure 8: Patient specific and non-specific detector latencies

Figure 9 illustrates the number of false detections declared by the detectors for each of the 17 test subjects (40 hours of test data). The patient-specific detector employing a nonlinear SVM committed 19 false detections (11 false detections per day); the same detector using a linear SVM committed 28 false detections (17 false detections per day). The patient non-specific detector committed 126 false detections (75 false detections per day).

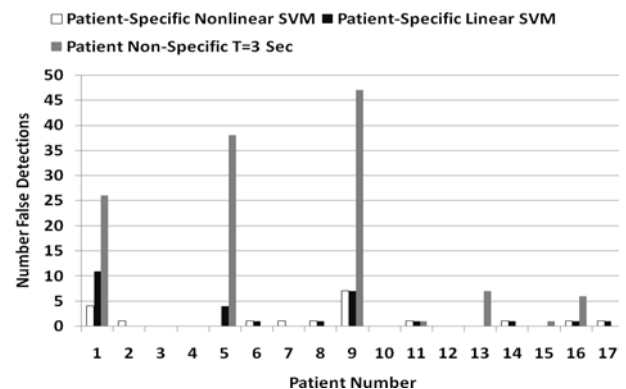


Figure 9: Patient specific and non-specific detector false alarm counts

When the patient non-specific detector was set to declare a seizure following $T=10$ seconds, it committed 17 false detections (10 false detections per day) as illustrated in Figure 10 alongside the patient-specific method's false detections.




Figure 10: Patient specific and non-specific detector false alarm counts

Furthermore, with T=10 seconds the patient non-specific algorithm detected 41/60 seizures within 18.7+/-10.8 seconds as shown in Figure 11 alongside the patient-specific detector latencies.




Figure 11: Patient specific and non-specific detector latencies


This comparison demonstrates that the patient-specific approach detects a larger percentage of seizures with both a lower false-detection rate and smaller detection latency when compared to the patient non-specific approach.

IV. IMPLEMENTATION

Both the patient-specific and non-specific algorithms were implemented on an ultra low-power neural IC [4] containing an analog front-end for IEEG signal acquisition and extraction of power within a configurable frequency band. The analog front-end is followed by a digital back-end that samples the analog power profile; assembles these samples into a feature vector; and classifies the feature vector using the algorithms discussed in Section II. This signal chain is illustrated the Figure 12.

By delegating the feature extraction to the analog front-end, the digital back-end remains idle until it is necessary to sample the power profile and test for the presence of seizure activity; this sampling can occur at a rate as low as 1Hz. Had feature extraction been delegated to the digital back-end, then sampling would have had to occur at rate suitable for IEEG signals, which is typically 200-250 Hz.

Analog :
Senses IEEG and Extracts Spectral Power



Digital:
Samples Spectral Power, Creates Feature Vector, and Classifies




Figure 12

Our measurements indicate that the total current drawn by the patient non-specific detector was 32 μ A. The patient-specific detector drew 12 μ A when using a linear SVM and 56 μ A when using a nonlinear SVM. These measurements suggest that the patient-specific detector employing a linear SVM exhibits the best combination of accuracy, latency, and power consumption.

V. DISCUSSION AND CONCLUSION

Patient-specific detectors synthesized using machine learning techniques can improve the performance of algorithm-driven neurostimulators. Moreover, embedding these algorithms within fully implantable neurostimulators is possible when simple decision boundaries or model-order reduction techniques are applied. Finally, it is important to note that a patient-specific, machine learning based detector requires an expert to designate non-seizure activity and target seizure activity in a training set; software that facilitates this labeling process is essential for successful deployment.

REFERENCES

- [1] Schuele SU, Lüders HO. "Intractable epilepsy: management and therapeutic alternatives". *Lancet Neurol* 2008 Jun; 7(6):514-24.
- [2] Theodore WH, Fisher RS. "Brain stimulation for epilepsy". *Lancet Neurol*. 2004 Feb;3(2):111-8.
- [3] Soltesz I, Staley K. *Computational Neuroscience in Epilepsy*. Elsevier Inc, 2008, chapters 28, 32.
- [4] Avestruz A, et al. "5 μ W/Channel spectral analysis IC for chronic bidirectional brain-machine interfaces" *IEEE J Solid-State Circuits*, Vol 43, No. 12, Dec 2008.
- [5] Shawe-Taylor J, Cristianini N. *Support vector machines and other kernel-based learning methods*. Cambridge University Press, 2000.
- [6] Osorio I, Frei MG, Wilkinson SB. "Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset." *Epilepsia* 1998 Jun;39(6):615-27.
- [7] Grewal S, Gotman J. "An automatic warning system for epileptic seizures recorded on intracerebral eegs." *Clin Neurophysiol*. 2005 Oct; 116(10):2460-72.
- [8] Statistical Pattern Recognition Toolbox For Matlab (STPRTool): <http://cmp.felk.cvut.cz/cmp/software/stprtool/index.html>