33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

An Algorithm For Detecting Seizure Termination in Scalp EEG

Ali Shoeb, Alaa Kharbouch, Jacqueline Soegaard, Steven Schachter, John Guttag

Abstract— **Little effort has been devoted to developing algorithms that can detect the cessation of seizure activity in scalp EEG. Such algorithms could facilitate clinical applications such as the estimation of seizure duration or the delivery of therapies designed to mitigate postictal period symptoms. In this paper, we present a method for detecting the termination of seizure activity. When tested on 133 seizures from a public database, our method detected the end of 132 seizures with a mean absolute error of 10.3 $\pm$ 5.5 seconds of the time marked by an electroencephalographer. Furthermore, by pairing our seizure end detector with a previously published seizure onset detector, we could automatically estimate the duration of 85% of test seizures within a 15 second error margin.**

I. INTRODUCTION

Much effort has been devoted to developing computerized algorithms for detecting the presence of seizure activity in recordings of the scalp electroencephalogram (EEG). Published algorithms include patient-specific [1] and non-specific [2] algorithms capable of detecting seizures in either an online or offline fashion.

While some effort has been devoted to developing methods for detecting the end of the non-epileptic seizures that accompany electroconvulsive therapy [3], little effort has been devoted to investigating how to reliably detect the cessation of epileptic seizures. Such algorithms could facilitate novel clinical applications. For instance, detecting the end of a seizure could prompt the delivery of a therapy that decreases the side effects of the postictal period [4].

Implicit within the literature on seizure detection is the assumption that the same decision rule used to detect a seizure's onset can be used to detect its end. However, this approach yields poor results because a seizure's onset and end can be very different. Seizures with focal activity following their onset may exhibit generalized activity following their end, and generalized seizures may exhibit focal activity following their end [5]. Furthermore, the rhythmic activity accompanying the onset of a seizure typically has a fundamental frequency within the alpha, beta, theta, or delta bands [2], while EEG following a seizure, which is known as postictal EEG, often exhibits delta wave slowing, amplitude attenuation, or a combination of these phenomena [5].

Similar to seizure onset, postictal EEG also varies across patients. For example, the postictal EEG of patient A, shown in the left panel of Figure 1, exhibits generalized amplitude attenuation following seizure termination at 3380 seconds. In contrast, the postictal EEG of patient B, shown in the right panel Figure 1, exhibits high-amplitude, slow activity following seizure termination at 2925 seconds.

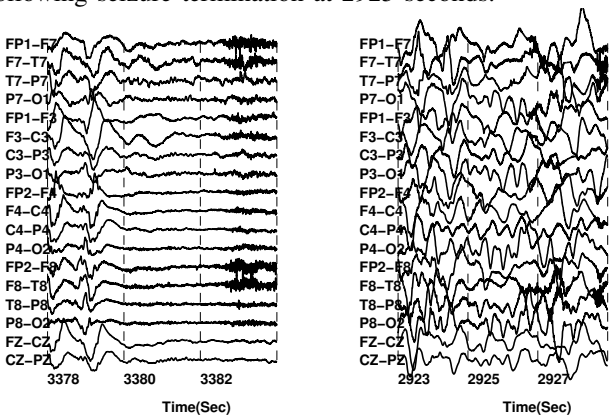

Fig. 1: Postictal EEG of patient A (left) and B (right).

In this paper we present a method for detecting the end of seizure activity. Our method, which can be patient-specific or non-specific, involves initiating a search for the end of a seizure once its beginning has been identified. By separating the tasks of seizure onset and end detection, we transform seizure end detection to a binary classification problem that involves separating ictal and postictal EEG.

II. DETECTING SEIZURE TERMINATION

The seizure end detector performs a sliding window analysis of the EEG once seizure onset has been detected. The analysis window is 5 seconds long, and is advanced by 1-second increments. The detector extracts features from each analysis window, and then assembles a feature vector X. Next, the detector uses a classifier to determine whether X is representative of the ictal or postictal state. If $L = 5$ successive vectors are classified as postictal, then the detector declares that the seizure has ended.

A. Feature Vector Design

We capture the spatial and spectral properties of each EEG epoch in a $M = 25$ dimensional vector X. Each element in X corresponds to the average, across N channels, signal energy within one of M contiguous frequency bands. The bands are 1 Hz wide and span the frequency range 0-M Hz (e.g. 0-1,1-2,...,24-25 Hz). We assemble X as described next.

First, we compute x_{ij}, the energy within band $i = 1,2,\ldots,M$ on channel $j = 1,2,\ldots,N$. We do so by computing the periodogram of channel j, and then summing the values of that function over the indices corresponding to the frequency band i. We estimate the periodogram of an epoch from channel j by applying the Welch algorithm

Ali Shoeb (ashoeb@partners.org) is with the Massachusetts General Hospital, Boston, USA. Steven Schachter is with the Center for the Integration of Medicine and Innovative Technology, Boston, USA. Alaa Kharbouch, Jacqueline Soegaard, and John Guttag are with the Massachusetts Institute of Technology, Boston, USA.

with 1 second window, 50% overlap, and $N_{FFT} = 4096$. Next, we compute the elements of X. Each element X_i corresponds to a weighted average of x_{ij} across N channels. Thus, $X_i = (\sum_{j=1}^{N} w_j x_{ij})(\sum_{j=1}^{N} w_j)^{-1}$ for $i = 1,2,\ldots,M$.

In the patient non-specific setting, no information about the test patient is available to guide the weighting of x_{ij}, so $w_j = 1$ for $j = 1,2,\ldots,N$. In the patient-specific setting, channels are weighted based on how differently they appear in the ictal and postictal states. If $\bar{x}_{ij,ictal}$ and $\bar{x}_{ij,postictal}$ are the mean x_{ij} extracted from a patient's training ictal and postictal epochs, then $w_j = \sum_{i=1}^{M} |\bar{x}_{ij,ictal} - \bar{x}_{ij,postictal}|$.

B. Feature Vector Classification

To classify X as ictal or postictal, we use a support-vector machine (SVM). Specifically, the implementation in the Statistical Pattern Recognition Toolbox (STPRTool) [6]. A SVM assigns X to one of these classes based on which side of a decision boundary it falls. The decision boundary, which is determined using training data, can be chosen to be linear or nonlinear. Which boundary type to choose is influenced by the amount of training data available, and its distribution within the feature space.

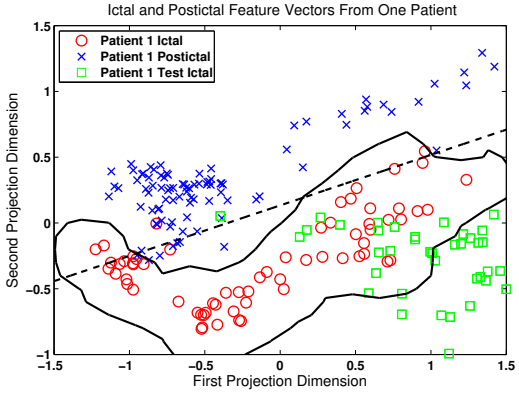

Fig. 2: Separation of ictal and postictal feature vectors from a patient using linear and nonlinear decision boundaries.

In the patient-specific setting, training data is limited. Consequently, nonlinear boundaries that tightly fit the distribution of training samples may result in poor performance on seizures not well-represented by those samples. As an example, Figure 2 shows the linear and nonlinear boundaries learned using two-dimensional projections of training ictal (red) and postictal (blue) vectors from a patient. The figure also shows how both boundaries perform when used to classify test ictal vectors (green) extracted from the same patient during an unusual seizure. The nonlinear boundary encloses a few test vectors. In contrast, the linear boundary correctly classifies all but one of the test vectors as ictal. In this work, we use a L_2 soft-margin SVM (*svm2* in STPRTool) with a linear kernel and regularization constants $C = 5$ for ictal and $C = 1$ for postictal when operating the detector in the patient-specific mode.

In the patient non-specific setting, training datasets exhibit greater diversity and cannot be well-separated by a linear boundary. As an example, Figure 3 shows training ictal (circles) and postictal (crosses) vectors from two patients, and a nonlinear boundary that can separate those vectors better than any other linear boundary. In this work, we use a L_2 soft-margin SVM with radial basis function kernel, kernel parameter 0.25, and default regularization constants when operating the detector in the patient non-specific mode.

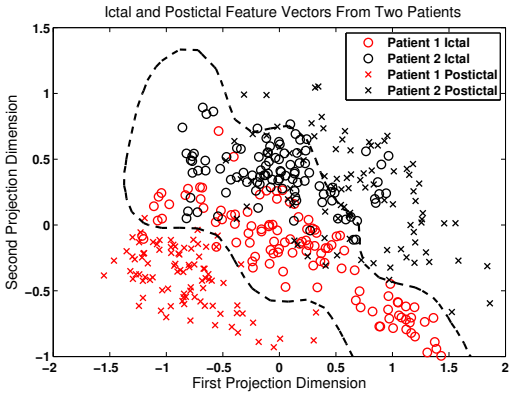

Fig. 3: Separation of ictal and postictal feature vectors from two patients using a nonlinear decision boundary.

III. Testing Methodology

A. Dataset

We evaluated our detector on the CHB-MIT database[1], which contains 133 seizures from 22 pediatric patients; 7 patients had partial-onset seizures, 7 had secondarily-generalized seizures, and 8 had generalized seizures. Each seizure's onset and end was determined by an expert. Figure 4 shows each patient's minimum, maximum, and mean seizure length. The numeral located over each bar indicates the number of seizures available for each patient.

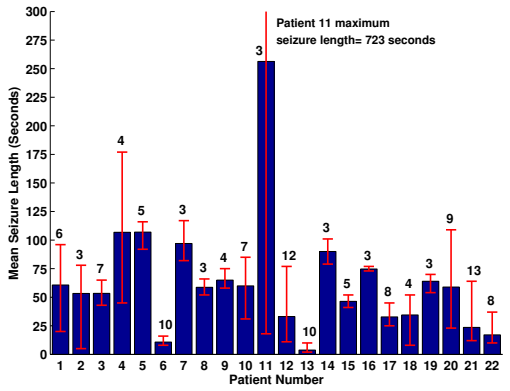

Fig. 4: Number of seizures and minimum, maximum, and mean seizure length for each database patient. EEG sampled at 256 Hz.

B. Evaluation Metrics

We use multiple metrics to assess the performance of the seizure end detector. The first metric, *end detection error*, measures the time difference between algorithm declaration

[1] http://physionet.org/physiobank/database/chbmit/

of seizure end and expert-marked seizure termination. When we wish to ignore the sign of the error, we report the absolute value of the end detection error. The next metric, *sensitivity*, measures the percentage of test seizures whose end the algorithm detected. The third metric, *classification accuracy*, refers to the percentage of ictal and postictal feature vectors that the end detector correctly classifies. Finally, *seizure length estimation error*, reflects how well the combination of a patient-specific seizure onset and end detector measure seizure duration.

C. Evaluation Scheme

We use a leave-one-record-out procedure to evaluate the performance of our seizure end detector. Since we assume a separate module will detect seizure onset, test records extend from the electrographic onset of a seizure to 90 seconds following its end. Records contain, on average, 64 $\pm$ 51 ictal vectors and 90 postictal vectors. Let N_i be the number of records that belong to patient i.

In the patient-specific context, the end detector is trained on vectors derived from the ictal and postictal periods in $N_i - 1$ records. Next, the detector is tasked with detecting seizure termination in the withheld record. These two steps are repeated N_i times so that each record from patient i is tested once.

In the patient non-specific context, the end detector is trained on feature vectors extracted from the records of all database patients other than patient i. The detector is then tasked with detecting seizure terminations in all N_i records belonging to patient i.

IV. Results

A. End detection error, Sensitivity, Classification accuracy

The patient-specific detector recognized 132/133 seizure ends with a classification accuracy of 90% and an average, absolute end detection error of 10.3 ± 5.5 seconds. The patient non-specific detector recognized all seizure ends with a classification accuracy of 84% and an average, absolute error of 8.9 ± 2.3 second.

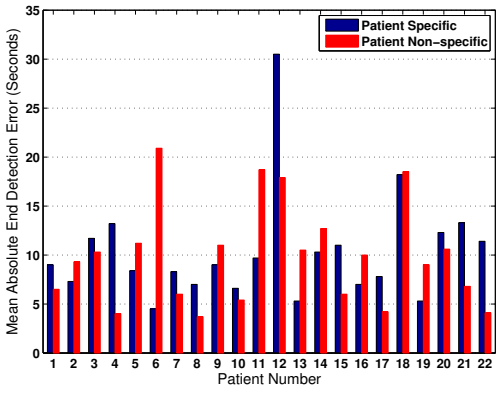

Fig. 5: Patient-specific and non-specific mean, absolute end detection error.

Figure 5 shows the mean, absolute end detection error for each detection scheme/patient pair. For 16 of 22 patients (e.g. patients 1, 2, 3, 5) the patient-specific and non-specific detectors exhibited absolute errors that differed by less than 5 seconds.

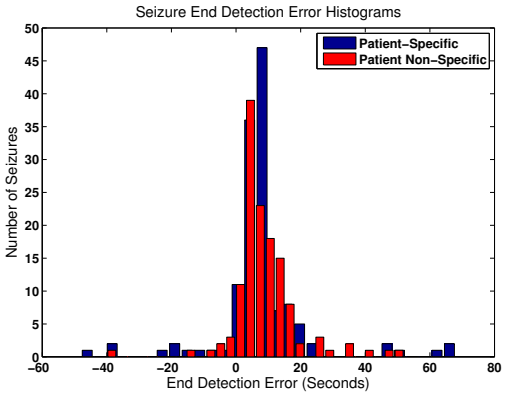

Fig. 6: Histogram of patient-specific and non-specific end detection errors.

Figure 6 shows histograms of the signed end detection error for the patient-specific and non-specific detectors. For both detectors, the majority of seizure ends were detected following the end marked by the expert, which results in a positive error. Each detector noted the end of 81% of the test seizures within 15 seconds of the end marked by the expert.

B. Seizure Length Estimation

Figure 7 shows the average, signed seizure length estimation error obtained using the combination of a published patient-specific seizure onset detector [7], and the end detector described in this paper. The combined onset and end detector was able to accurately estimate the length of both short seizures (patients 6, 13) and long seizures (patients 5, 7). The patients for whom the mean seizure length estimation error was close to or greater than 10 seconds (patients 3, 12, and 18) did not have unusually long seizures or a small number of training seizures. Instead, the EEG of these patients exhibited ictal or postictal activity that caused the end detector to declare seizure termination too early or too late as illustrated in the following section.

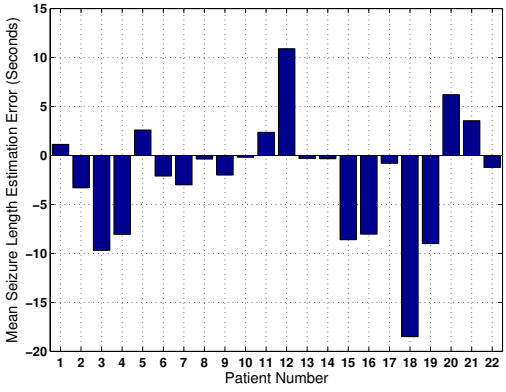

Fig. 7: Seizure length estimation error for a combined patient-specific onset and end detector.

C. Case Studies

The patient-specific seizure end detector exhibited robust performance on the data belonging to patient 11. The database contains three seizures for patient 11 of length 18, 28, and 723 seconds. As shown in the left panel of Figure 8, short seizures exhibited low-amplitude, 4-5 Hz activity that was localized to the channels FP_1-F_7 through P_7-O_1. A portion of the long seizure is shown in the right panel of Figure 8. This seizure exhibited high amplitude, 2-3 Hz activity on the channels involved in the shorter seizures, and the channels P_4-O_2, T_8-P_8. Although the long seizure is different, the detector was able to correctly identify its endpoint with a 11 second end detection error when trained only on the two shorter seizures.

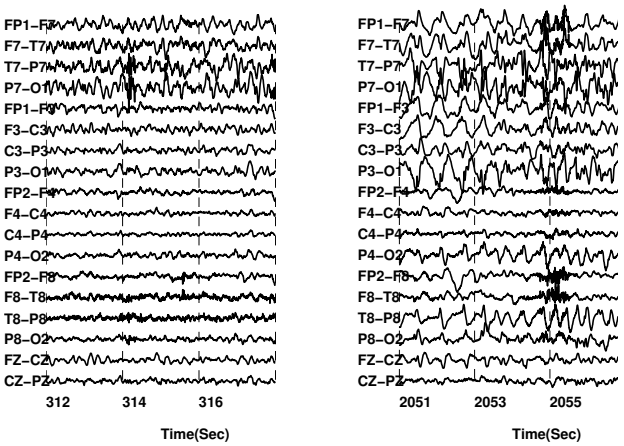

Fig. 8: The patient-specific seizure end detector continues to recognize a patient's ictal activity even after significant change in its spectral and spatial chracter.

Seizures that gradually transitioned from the ictal to the postictal state caused the detector to exhibit large seizure end detection errors. As an example, Figure 9 shows a seizure belonging to patient 1. Although an expert determined that this seizure ends at 1491 seconds, the rhythmic activity on many EEG channels (e.g. FP_1-F_7, F_3-C_3) changes only slightly. This caused the algorithm to postpone declaring seizure termination until 1510 seconds.

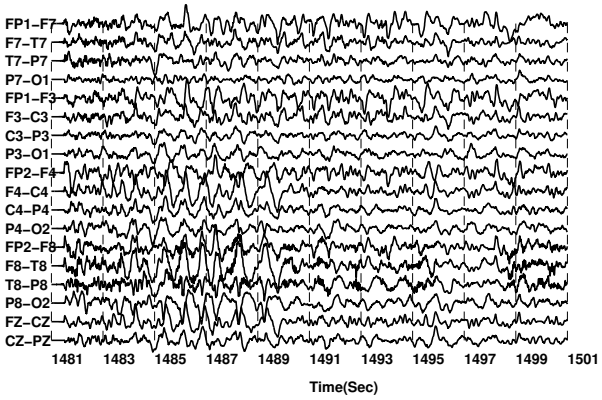

Fig. 9: Similarity in the spectral character of ictal and postictal activity causes the patient-specific seizure end detector to delay declaring seizure termination.

In contrast, the detector declared seizure termination prematurely whenever significant artifacts overlapped seizure activity. As an example, Figure 10 shows a stretch of artifact-corrupted, ictal EEG belonging to patient 3. The rhythmic ictal activity appears most prominently on the channel F_7-T_7 between 1746-1750 seconds. In this case, averaging across channels obscured the ictal activity, and caused the detector to incorrectly conclude that the seizure ends 40 seconds before the actual endpoint.

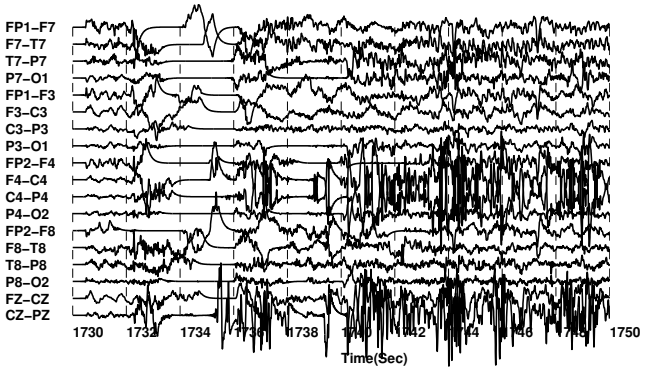

Fig. 10: Recording artifacts cause the seizure end detector to declare seizure termination prematurely.

V. CONCLUSIONS

We presented a machine learning based method for detecting the end of seizures. Our method exhibits excellent performance in both the patient-specific and non-specific settings thanks to careful feature vector and classifier design. Furthermore, our detector can be easily paired with previously published seizure onset detectors to facilitate applications such as epileptic seizure duration estimation.

REFERENCES

[1] G. R. Minasyan, J. B. Chatten, M. J. Chatten, and R. N. Harner, "Patient-specific early seizure detection from scalp electroencephalogram." *Journal of clinical neurophysiology*, vol. 27, no. 3, pp. 163–178, jun 2010.

[2] R. Meier, H. Dittrich, A. Schulze-Bonhage, and A. Aertsen, "Detecting Epileptic Seizures in Long-term Human EEG: A New Approach to Automatic Online and Real-Time Detection and Classification of Polymorphic Seizure Patterns," *Journal of Clinical Neurophysiology*, vol. 25, no. 3, pp. 119–131, jun 2008.

[3] B. Gangadhar and D. Dutt, "Automation of seizure duration estimation during ECT; use of fractal dimension," *IEEE Engineering in Medicine and Biology Society Conference*, pp. 3.37–3.38, 1995.

[4] K. Vonck, R. Raedt, and P. Boon, "Vagus nerve stimulation and the postictal state," *Epilepsy & Behavior*, vol. 19, no. 2, pp. 182–185, oct 2010.

[5] N. K. So and W. T. Blume, "The postictal EEG," *Epilepsy & Behavior*, vol. 19, no. 2, pp. 121–126, oct 2010.

[6] V. Franc and V. Hlaváč, "A new feature of the statistical pattern recognition toolbox," in *Computer Vision, Computer Graphics and Photogrammetry – a Common Viewpoint*, S. Scherer, Ed. Österreichische Computer Gesellschaft, June 2001, pp. 143–150. [Online]. Available: http://cmp.felk.cvut.cz/cmp/software/stprtool/

[7] A. Shoeb and J. Guttag, "Application of machine learning to epileptic seizure detection," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, J. Fürnkranz and T. Joachims, Eds. Haifa, Israel: Omnipress, June 2010, pp. 975–982. [Online]. Available: http://www.icml2010.org/papers/493.pdf