

MIT Open Access Articles

An algorithm for seizure onset detection using intracranial EEG

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Kharbouch, Alaa, Ali Shoeb, John Guttag, and Sydney S. Cash. "An Algorithm for Seizure Onset Detection Using Intracranial EEG." *Epilepsy & Behavior* 22 (December 2011): S29–S35.

As Published: <http://dx.doi.org/10.1016/j.yebeh.2011.08.031>

Publisher: Elsevier

Persistent URL: <http://hdl.handle.net/1721.1/100243>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



Published in final edited form as:

Epilepsy Behav. 2011 December ; 22(0 1): S29–S35. doi:10.1016/j.yebeh.2011.08.031.

An algorithm for seizure onset detection using intracranial EEG

Alaa Kharbouch^{a,*}, Ali Shoeb^b, John Guttag^a, and Sydney S. Cash^{b,c}

^aDepartment of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA

^bDepartment of Neurology, Massachusetts General Hospital, Boston, MA, USA

^cHarvard Medical School, Boston, MA, USA

Abstract

This article addresses the problem of real-time seizure detection from intracranial EEG (IEEG). One difficulty in creating an approach that can be used for many patients is the heterogeneity of seizure IEEG patterns across different patients and even within a patient. In addition, simultaneously maximizing sensitivity and minimizing latency and false detection rates has been challenging as these are competing objectives. Automated machine learning systems provide a mechanism for dealing with these hurdles. Here we present and evaluate an algorithm for real-time seizure onset detection from IEEG using a machine-learning approach that permits a patient-specific solution. We extract temporal and spectral features across all intracranial EEG channels. A pattern recognition component is trained using these feature vectors and tested against unseen continuous data from the same patient. When tested on more than 875 hours of IEEG data from 10 patients, the algorithm detected 97% of 67 test seizures of several types with a median detection delay of 5 seconds and a median false alarm rate of 0.6 false alarms per 24-hour period. The sensitivity was 100% for 8 of 10 patients. These results indicate that a sensitive, specific, and relatively short-latency detection system based on machine learning can be employed for seizure detection from EEG using a full set of intracranial electrodes to individual patients.

Keywords

Seizure detection; Intracranial electroencephalography; Patient specific

1. Introduction

Epilepsy is a disease characterized by recurrent episodes of dysfunctional brain activity associated in time with changes in behavior. Yet, current therapies do not take into account the episodic nature of epileptic seizures. Therefore, a goal of current research is to develop seizure-triggered diagnostic, therapeutic, and alerting systems. Central to these systems is an

© 2011 Elsevier Inc. All rights reserved.

*Corresponding author at: Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 32 Vassar Street, Room 32-G915, Cambridge, MA 02139, USA. aak@mit.edu (A. Kharbouch).

Ethical approval

We confirm that we have read the Journal's position on issues involved in ethical publication and affirm that this report is consistent with those guidelines.

Conflict of interest statement

The Massachusetts Institute of Technology has licensed a seizure onset detection algorithm to Cyberonics, Inc., and J Guttag has served as a consultant to Cyberonics, Inc. None of the remaining authors has any conflict of interest to disclose.

algorithm that can detect seizure activity early and accurately. In this article, we describe the architecture and performance of a real-time intracranial EEG (IEEG) seizure onset detector.

Rapid and reliable seizure onset detection from IEEG is challenging for a number of reasons. First, IEEG varies greatly across individuals with epilepsy [1]. In fact, the intracranial EEG associated with seizure onset in one patient can closely resemble a benign pattern within the IEEG of another patient. Furthermore, there exists significant overlap in the IEEG associated with seizure and nonseizure states. In addition to interpatient variability, there is also intrapatient variability. The identity of the IEEG channels involved and the evolution of the earliest seizure activity can differ within an individual, particularly when seizures arise from different brain regions. Moreover, the IEEG of patients with epilepsy transitions between electrographic stages within both the seizure and the nonseizure states and is, therefore, a nonstationary process.

Previous work has introduced a wide variety of techniques. Results from these studies show that many of the techniques struggle with high latencies, false alarm rates, or both. Furthermore, some aspects of these studies remain to be addressed. First, most IEEG detection studies use only relatively short records; past studies have used no more than 30 hours per patient and considerably less than 24 hours for many patients [2–5]. This may lead to an inaccurate estimate of the false alarm rate and may not faithfully represent long-term performance in a clinical application.

In addition, previous work in the area of IEEG seizure detection has focused on data sets with a small number of preselected electrodes [2,5–7]. This may ignore useful information that is not obvious on visual examination and, in some cases, necessitates additional assessment and culling of the data set by an expert.

We employ a machine-learning approach that is patient specific (i.e., the classifier for each patient is trained using the same patient's previously recorded IEEG data). Our evaluation uses more than 875 hours of IEEG in total, averaging more than 87 hours of continuous IEEG data per patient. This study also uses the full set of intracranial electrodes of sufficient recording quality. The algorithm detects seizures by examining the short-term evolution of spectral properties of the IEEG across many channels and comparing periods between seizures with seizure activity itself. Relative to previously published methods, ours exhibits high sensitivity, short latencies, and low false alarm rates.

2. Methods

We treat seizure detection as a binary classification problem that involves separating seizure activity from nonseizure activity. We adopt a patient-specific approach to seizure detection to overcome the cross-patient variability in ictal and interictal IEEG patterns and to exploit the consistency within ictal patterns emerging from the same brain region within the same patient. The key to our detector's high accuracy is a feature vector that unifies in a single feature space the time evolution of spectral properties of the brain's electrical activity as recorded by several IEEG electrodes. The algorithm presented is based in part on the algorithm in [8,9], with some essential changes also described in this section.

2.1. Detection algorithm

Our goal is to construct a function $f(X)$ that maps a feature vector X derived from an epoch of IEEG onto the labels $Y = \pm 1$ depending on whether X is representative of seizure or nonseizure IEEG. The function is derived using training sets of seizure and nonseizure feature vectors specific to an individual patient. In this section we discuss how we construct the feature vector X , the discriminant function $f(X)$, and the training sets.

Features important for characterizing IIEEG activity include its spectral distribution, the channels on which it manifests, and its short-term temporal evolution. The following subsections illustrate how these features are extracted and encoded. We use spectral energy features similar to those that have been shown to be effective in the seizure detection scheme of Shoeb [8]. Each spectral feature represents the logarithm of the total energy in a specific frequency band on a single channel.

EEG signals generally have a spectral amplitude profile that is inversely proportional to frequency. To remove this frequency-domain trend, a derivative filter is applied to all channels as an added first step in the feature extraction phase for the IIEEG detector. This introduces more parity in the scaling of spectral content at different frequencies.

Considering the multiple frequency components that compose the activity associated with seizure onset is essential to detecting seizures with high accuracy. The dominant spectral content of a seizure epoch may overlap the dominant frequency of an epoch of nonseizure activity, but they can still be distinguished by the presence or absence of other spectral components. We extract the spectral structure of a sliding window of length $L = 1$ second by passing it through a filter bank and then measuring the energy falling within the passband of each filter. This choice of epoch length provides sufficient time resolution to capture discrete electrographic events, and also provides sufficient frequency resolution when compared against the bandwidth of the bandpass filters. Consecutive epochs are separated by 1 second. The filter bank is composed of $M = 17$ filters and is illustrated in Fig. 1.

The scalp EEG seizure detector in [8,9] focuses on the frequency range 0.5–25 Hz. However, the IIEEG is a signal of a higher bandwidth and carries relevant information at higher frequencies. Therefore, the IIEEG detector filter bank considers a wider range of frequencies, although more emphasis is still placed on the lower-frequency range. The range 0.5–35 Hz range is covered by more filters (12), each with a bandwidth of 3 Hz, whereas the range 35–105 Hz range is covered by a lower density of filters (5), each with a bandwidth of 15 Hz. The 60-Hz region is neglected so as to remove the 60-Hz electrical noise often contaminating the signals from consideration.

For channel k , the energy measured by filter i is denoted by the feature $x_{i,k}$. To capture the spectral and spatial information contained within each 1-second EEG epoch at time $t = T$, we concatenate the $M = 17$ spectral energies extracted from each of N EEG channels. This process forms a feature vector \mathbf{X}_T with $M \times N$ elements as shown in the middle portion of Fig. 1. Each feature is multiplied by a coefficient to normalize the approximate scaling of each feature in the set. The scaling and normalization coefficients are determined using the training points only. This is to avoid situations where some features do not play a significant role in the classifier structure because of the dominance of features with much larger scaling. The features are median centered and then divided by the median absolute value.

The feature vector \mathbf{X}_T does not capture how an epoch relates to those in the recent past. Consequently, \mathbf{X}_T cannot represent how a seizure emerges from the background or how it evolves. To capture this information, we form the time-delay embedded feature vector χ_T by stacking the vectors $\mathbf{X}_T, \mathbf{X}_{T-L}, \dots, \mathbf{X}_{T-(W-1)L}$ from W contiguous, but nonoverlapping $L = 1$ second epochs as shown on the right side of Fig. 1. This approach allows the timing and order of discrete events to be encoded to some extent, and it is not equivalent to forming a feature vector \mathbf{X}_T using a longer epoch length L . We set $W = 3$.

The feature vector χ_T is classified as representative of seizure or nonseizure activity using a linear support vector machine (SVM). We train the SVM on seizure vectors computed from the first $S = 20$ seconds of each training seizure, and on nonseizure vectors computed from

nonseizure IEEG. In the training phase, the nonseizure feature vectors were subsampled such that only every sixth epoch in the training set was used due to memory limitations. An exception to this is made for the 20-minute period following any seizure, to ensure there are enough training examples to describe the postictal period, which tends to be associated with electrographic qualities that distinguish it from the rest of the nonseizure activity. Within the SVM^{light} software package [10], the error cost parameter was set to $C = 1/1000$.

An artifact rejection component works in conjunction with a trained classifier by checking for large differences between the minimum and maximum values of the signal for each channel within $W \times L = 3$ seconds to reject high-amplitude artifacts. If more than 20% of the channels are deemed to contain artifact then a seizure detection alarm is prevented at the corresponding time. More specifically, a seizure detection alarm is raised when feature vectors corresponding to $K = 2$ consecutive epochs are classified as belonging to the seizure class, and no artifact is detected (on enough channels) within them by the artifact rejection module. The alarm is turned off if the classifier does not detect a seizure in any epochs for more than 2 minutes.

2.2. Data and evaluation methodology

The data used to evaluate our detector consist of more than 875 hours of continuous IEEG sampled at 500 Hz. The data were recorded at Massachusetts General Hospital from 10 patients with focal epilepsy (5 females, age at onset = 15 ± 5 years [mean \pm SD], age at surgery = 40 ± 9 years [mean \pm SD]). Etiologies included mesial temporal sclerosis (2), cortical dysplasia (2), posttraumatic epilepsy (1), and postinfectious epilepsy (1); in 4 patients the etiology was unknown.

The patients were all surgical candidates who required invasive monitoring and, therefore, represent cases more complicated than the general population of patients with epilepsy. They are also not necessarily representative of the population most likely to benefit from a seizure detection system. For example, patients for whom resection would incur too much risk (e.g., because the seizures arise from eloquent cortex) would be prime candidates for an implanted seizure detection and control system. The data used for this article were collected for clinical purposes, and once enough seizures were observed data collection was halted. If a similar system were to be deployed, sufficient data would be collected from a patient specifically for the purpose of training the system, as is the case (albeit with manual tuning) with existing devices [11].

While the recordings were being made, the patients experienced a total of 67 seizures (between 3 and 12 seizures per patient). For each seizure, an expert indicated the earliest IEEG change associated with the seizure. The patients were consecutively chosen, with two patients omitted because the clinicians were unable to reliably determine seizure onsets. We used a bipolar montage, consistent with the montage used by clinicians who reviewed each patient's data and authored an accompanying IEEG report. In most cases the recording from each electrode is used for one channel. Analysis of these data was performed retrospectively under the auspices of the local institutional review board in accordance with the Declaration of Helsinki.

Data for each patient were segmented into a number of records, where a record corresponds to a stretch of time in the IEEG recordings. Every record is up to 24 hours long, and contains at least one seizure (in the very small number of cases in which two seizures occurred less than 15 minutes apart, both are included in the same record). If the separation between the seizure in the record and the preceding one is greater than 24 hours, the beginning of the record is set to approximately 24 hours before the seizure, and its end, to approximately 20 minutes after (i.e., every seizure's postictal phase is included within the same record). If that

period is less than 24 hours, then the start of the record is set to the end of the previous record, and the end is set to be no earlier than 20 minutes after the end of the seizure. In that case the endpoint of the record may be moved to a later point to maximize record length, as long as it does not cross the beginning of the next record and the 24-hour record duration limit is maintained. Some stretches of time corresponding to disconnects or nonrecording of electrodes were excluded from patient data sets. One artifact-obscured seizure was omitted from any training sets, but was still used for testing.

A seizure onset is considered to be the earliest point of unequivocal change in the IEEG waveforms leading up to a seizure, as judged by an expert electroencephalographer. We characterized our detector's performance in terms of sensitivity, specificity, and latency. Sensitivity refers to the percentage of test seizures detected. A seizure is considered successfully detected if an alarm is raised between its (expert-marked) onset and its end. The false alarm rate refers to the average number of times, per 24 hours, that the detector incorrectly declared the onset of seizure. Alarms that begin outside intervals between a seizure onset and the end of the same seizure are considered *false alarms*. The delay between the expert-identified electrographic onset and the time a detection algorithm indicates that a seizure has been detected is referred to as the *detection latency*.

To estimate our detector's performance on data from a given patient, we used a leave-one-record-out cross-validation scheme. We avoided an evaluation method based on leaving out epochs [12] rather than records, as that approach leads to unrepresentative and misleadingly good results by including in the training set feature vectors in close temporal proximity to those in the test data.

Let N denote the number of records for a given patient. To estimate the detector's latency, sensitivity, and false alarm rate we train the detector on $N - 1$ records from the patient. The detector is then tasked with detecting the seizure in the withheld record. For each round we record whether the test seizure was detected and, if so, with what latency. Any alarms beginning outside the seizure are also counted as false alarms. This process is repeated N times so that each record is tested. In most cases, $N - 1$ training seizures are used for each test record, as each record contains a single seizure.

3. Results

Overall, 97% of the 67 test seizures were detected. The median latency with which the detector declared seizure onset (across all seizures) was 5 seconds. The median latencies for each patient are listed in Table 1. The average false alarm rates for each patient are also listed in Table 1. The median false alarm rate was 0.6 false detections per 24-hour period.

A notable phenomenon is that of single seizures that were deemed by a clinician to be of a distinct seizure type relative to all the other remaining seizures in the patient data set. Our detector misses or has a long latency when a test seizure largely differs from all the training seizures. One example of this is the seizure from Patient 6's data set that was missed by the algorithm, shown in Fig. 2. The two remaining seizures from this patient were grouped together (by a clinician in the patient report) and one of them is shown in Fig. 3. Clear differences between the seizure types include high-amplitude rhythmic spiking activity on the first anterior temporal channel in the seizure shown in Fig. 3 that is absent from the same channel in Fig. 2. Rhythmic slowing in several anterior and posterior temporal contacts early in the onset of the seizure in Fig. 3 also does not match the activity in the seizure in Fig. 2.

The algorithm returned a notably long median detection latency for Patient 10. The clinician report for Patient 10 indicates at least three classes of IEEG onset activity patterns across six total seizures, and notes some difference in morphology on the channels that displayed the

earliest noted seizure activity within those classes. The clinical conclusion from this IEEG was that the patient had multifocal epilepsy, although the possibility of a single focus with heterogeneous propagation cannot be entirely excluded. As a result, the onset of nearly every seizure was unique in some way. The seizures in the training data, therefore, differed from one another. For every given test seizure, the paucity or absence of sufficiently similar examples of onset activity from seizures in the training set contributed to the relatively poor detection latency.

Not all cases that include multiple distinct seizure types yield poor results. More favorable results were obtained for Patient 8, for whom the clinician enumerated several seizure types (divided between left and right hemisphere onsets and presence and lack of apparent clinical symptoms). This is possibly explained by the fact that the patient's data set includes 12 seizures, and therefore, for most test seizures within this patient's data set, the classifier has incorporated more examples of seizures of the same broad type, and observed some variation in IEEG patterns among them. This allows for a better performance, particularly in reduced latency in the case of this patient's data set.

The largest false alarm rate estimate was obtained for Patient 7. Some false alarms were due to mimics of seizure onset activity; many had a morphological appearance similar to that of ictal manifestations seen in other seizures. These mimics were shorter than typical clinical seizures or were too variable in length or appearance to have been judged a seizure by clinicians. Figs. 4 and 5 illustrate examples of activity that induced alarm and the seizure activity they resemble. The high-amplitude rhythmic activity on the right subfrontal channel during the seizure in Fig. 4A also appears in the interictal epoch where false detection occurred in testing, shown in Fig. 4B. The type of activity highlighted in Fig. 4B lasted approximately 8 seconds. Fig. 5A shows the high-frequency activity on the right posterior temporal channel in a different seizure, and Fig. 5B shows similar activity on the same channels from an interictal epoch in which a false alarm was raised. Other false alarms appear to have been caused by nonphysiological artifact that was not detected by the artifact rejection module.

In many cases, the false alarm events are unevenly concentrated in different regions in time; that is, they temporally cluster. In the case of Patient 3, 60% of false alarms (three of five) occur within a 15-minute block of time (out of a total of approximately 71 hours). Another contiguous 2-hour block contains the two remaining false alarms. Fig. 6 contains a false alarm event plot for Patient 8. It is formed using a timeline that combines all the records in correct temporal order. An event shown on the patient's false alarm plot timeline indicates the occurrence of a false alarm in a leave-one-out test on the record in which the time point fell. For this patient, 66% of false alarms fall within a single contiguous block less than 12 hours long (out of a total of 146 hours). These false alarms seem to have been triggered by types of interictal electrical activity that do not appear with significant frequency outside the record that contains this period. The lack of prior examples of similar activity in interictal periods in the training set may explain the large number of false alarms for this record. Interestingly, this cluster of false alarms begins in close proximity to the point of lowest drug concentrations during the admission period for this patient.

We explored the possibility of reducing the latencies by adjusting the K parameter. Setting K equal to 2 forces the algorithm to wait for two consecutive windows that the classifier has deemed to be part of a seizure onset before an alarm is raised. Setting K equal to 1 has the effect of reducing latencies (by a minimum of $L = 1$ second for every seizure for this choice of epoch length) at the expense of a false alarm rate that likely increases. The results obtained by setting K equal to 1 are summarized in Table 2. The median latency across all detected seizures falls to 3.5 seconds. Although the latencies decrease, the median patient

false alarm rate rises to 2.7 false alarms/24 hours. This alternative trade-off point may be useful for certain applications where the consequences of false alarms are less severe or lower latencies are more crucial.

4. Discussion

We have described a patient-specific algorithm that detects seizures by examining the short-term evolution of spectral properties across several IIEEG channels. The results of this evaluation, using an average of more than 3 days of data per patient and a full IIEEG electrode set, show the efficacy of using a patient-specific algorithm for automatic seizure detection as determined by clinically relevant performance metrics.

4.1. Comparison with other approaches

Although this project builds on the work of Shoeb et al. [9, 13], several methods have been published on seizure detection on noninvasive EEG, including some which also incorporate a support vector machine as a component [14–16].

The median patient false alarm rate (0.6/day) and the median/mean latency across all seizures (5 seconds/6.9 seconds) obtained using our algorithms are lower than those reported in previous work on seizure detection using IIEEG, and the sensitivity was comparable (97%). Algorithms such as those described by Chan et al. [17] and Chua et al. [18] were designed for offline IIEEG analysis, and cannot be implemented as part of a real-time warning system. The patient-specific method of Zhang et al. [19] yielded a sensitivity of 98.8%, a mean latency of 10.8 seconds, and a combined false alarm rate of 11.8/day when tested on IIEEG from 21 patients. The authors report separate figures for false alarms dichotomized into “interesting” and “uninteresting” groups. Grewal and Gotman [4] describe an algorithm with tunable parameters that can be set for a given patient using data from that patient. Using tuned parameters they report a sensitivity of 89.7%, a median latency of 17.1 seconds, and a false alarm rate of 5.3/day. The method of Aarabi et al. [5] yielded a sensitivity of 98.7%, an average latency of 11 seconds, and a false alarm rate of 6.5 per 24 hours when tested on data from six contacts per patient in the Freiburg public database [20]. The method of Gardner et al. [21] used a one-class support vector machine requiring only nonseizure data for training, and yielded a false alarm rate of 37.4/day when tested on 200 hours of data from five patients. That study allowed for negative latencies in cases in which an alarm began up to 3 minutes before a seizure, and negative average latencies were reported for several patients (although the median latencies for these patients were positive). Shoeb et al. [7] evaluated their two-channel patient-specific IIEEG detector, which was constrained to compute the energy in two bands, on 81 hours of data from 17 subjects. A mean latency of 9.3 seconds and a false alarm rate of 11/day were obtained in that study. Osorio et al. [6] report impressive results. However, important differences distinguish their study from ours, including the automatic preselection of a subset of channels using patient data in their study. They also use a different evaluation methodology that includes the designation of some events that do not fall under the category of true positives as epileptiform discharges rather than false alarms. Visual review of all automated detections and a sample of interictal segments were used to determine sensitivity and specificity, and the IIEEG record had not been reviewed in its entirety to locate all instances of seizures. Thus, no clear comparison can be made.

4.2. False alarm generation and evaluation

In addition to recording the false alarm rate, we examined the nature and distribution of false alarms. False alarms were often caused by electrical artifacts as well as events that had features of seizure activity. Moreover, false alarms were nonhomogeneously spaced.

The temporal clustering of false alarm times has application-specific implications for the potential utility of this system. For example, consider the false alarms generated by a seizure detector for an ambulatory alerting system in a given week. The overall disruption caused by these alarms occurring a few minutes apart from one another may be less objectionable than the case where each occurs on a different day. However, the extent to which some factors may influence the number and distribution of events is not clear. For the hospitalized patients from whom the data for this study were recorded, the medication levels were changed on a daily basis. This may have caused changes in the IEEG and unfamiliar electrical patterns that led to false alarms. Nonetheless, the finding that false alarms are clustered temporally implies that long recordings are necessary in the evaluation of a seizure detector. Long records provide an ability to obtain far more representative estimates of the false alarm rate.

4.3. Impact of number of electrodes

As opposed to several other investigations, the approach used here incorporated a full set of intracranial electrodes. Overall, it is not clear what the effect of using a small subset of the channels has on the performance of classifiers. There are, however, some factors to consider. Using a full set of electrodes makes the clinical deployment of the algorithm simpler from the user's standpoint. For some patients it can be difficult to select only a few channels because ictal onset changes may involve many channels and be subtly different for each seizure. Furthermore, including all channels allows the system to leverage information that is not readily apparent to a clinical reader (from regions seemingly not involved in the onset) to better differentiate ictal and interictal patterns. On the other hand, allowing all channels to be used may lead the learning algorithm to incorporate nonspecific information that may increase the likelihood of incorrect classification. To better understand these trade-offs, a subsequent investigation of different channel counts using a given detection algorithm is necessary. Although there is currently a considerable difference between the number of channels in monitoring units and that in implantable devices, the number of channels in the latter is likely to increase.

In summary, we have demonstrated the utility of a learning-based algorithm for the early and accurate determination of seizure onsets from intracranial data. We have focused on long records from clinically complicated patients to evaluate this approach. The results from this system provide additional evidence of the potential advantages patient-specific algorithms compared with existing methods, especially when trying to balance early detection with a low false alarm rate.

Acknowledgments

A.K. was supported by grants from the Center for Integration of Medicine & Innovative Technology as well as funds from Quanta Computer, Inc and Cyberonics, Inc. SSC was supported by the Center for Integration of Medicine & Innovative Technology and the National Institute of Neurological Disorders and Stroke (NINDS-NS062092). The authors thank the reviewers for their comments.

References

1. Echaux, J.; Wong, S.; Smart, O.; Gardner, A.; Worrell, G.; Litt, B. Computation applied to clinical epilepsy and antiepileptic devices. In: Soltesz, I.; Staley, K., editors. Computational neuroscience in epilepsy. London: Academic Press; 2008. p. 544-558.
2. Osorio I, FreiM, Wilkinson S. Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset. *Epilepsia*. 1998; 39:615-627. [PubMed: 9637604]
3. Khan YU, Gotman J. Wavelet based automatic seizure detection in intracerebral electroencephalogram. *Clin Neurophysiol*. 2003; 114:898-908. [PubMed: 12738437]

4. Grewal S, Gotman J. An automatic warning system for epileptic seizures recorded on intracerebral EEGs. *Clin Neurophysiol.* 2005; 116:2460–2472. [PubMed: 16125459]
5. Aarabi A, Fazel-Rezai R, Aghakhani Y. A fuzzy rule-based system for epileptic seizure in intracranial EEG. *Clin Neurophysiol.* 2009; 120:1648–1657. [PubMed: 19632891]
6. Osorio I, Frei MG, Giftakis J, et al. Performance reassessment of a real-time seizure-detection algorithm on long ECoG series. *Epilepsia.* 2002; 43:1522–1535. [PubMed: 12460255]
7. Shoeb, A.; Carlson, D.; Panken, E.; Denison, T. A micropower support vector machine based seizure detection architecture for embedded medical devices. Presented at the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC); September 2009; Minneapolis, Minnesota. p. 4202-4205.
8. Shoeb, A. Ph.D. thesis. Massachusetts Institute of Technology; 2009. Application of machine learning to epileptic seizure onset detection and treatment.
9. Shoeb, A.; Gutttag, J. Application of machine learning to epileptic seizure detection. In: Fürnkranz, J.; Joachims, T., editors. *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*; June 2010; Haifa, Israel. p. 975-982.
10. Joachims, T. *Making large-scale SVM learning practical.* Cambridge: MIT Press; 1999.
11. Morrell M. Brain stimulation for epilepsy: can scheduled or responsive neurostimulation stop seizures? *Curr Opin Neurol.* 2006; 19:164–168. [PubMed: 16538091]
12. Mirowski P, Madhavan D, Lecun Y, Kuzniecky R. Classification of patterns of EEG synchronization for seizure prediction. *Clin Neurophysiol.* 2009; 120:1927–1940. [PubMed: 19837629]
13. Shoeb A, Edwards H, Connolly J, Bourgeois B, Treves ST, Gutttag J. Patient-specific seizure onset detection. *Epilepsy Behav.* 2004; 5:483–498. [PubMed: 15256184]
14. Gonzalez-Vellon, B.; Sanei, S.; Chambers, JA. Support vector machines for seizure detection. *Proceedings of the 3rd IEEE International Symposium on Signal Processing and Information Technology*; December 2003; p. 126-129.
15. Meier R, Dittrich H, Schulze-Bonhage A, Aertsen A. Detecting epileptic seizures in long-term human EEG: a new approach to automatic online and real-time detection and classification of polymorphic seizure patterns. *J Clin Neurophysiol.* 2008; 25:119–131. [PubMed: 18469727]
16. Temko A, Thomas E, Marnane W, Lightbody G, Boylan G. EEG-based neonatal seizure detection with support vector machines. *Clin Neurophysiol.* 2011; 122:464–473. [PubMed: 20713314]
17. Chan AM, Sun FT, Boto EH, Wingeier BM. Automated seizure onset detection for accurate onset time determination in intracranial EEG. *Clin Neurophysiol.* 2008; 119:2687–2696. [PubMed: 18993113]
18. Chua EC, Patel K, Fitzsimons M, Bleakley CJ. Improved patient specific seizure detection during pre-surgical evaluation. *Clin Neurophysiol.* 2011; 122:672–679. [PubMed: 21130030]
19. Zhang Y, Xu G, Wang J, Liang L. An automatic patient-specific seizure onset detection method in intracranial EEG based on incremental nonlinear dimensionality reduction. *Comput Biol Med.* 2010; 40:889–899. [PubMed: 20951372]
20. Maiwald T, Winterhalder M, Aschenbrenner-Scheibe R, Voss HU, Schulze-Bonhage A, Timmer J. Comparison of three nonlinear seizure prediction methods by means of the seizure prediction characteristic. *Phys D.* 2004; 194:357–368.
21. Gardner A, Krieger A, Vachtsevanos G, Litt B. One-class novelty detection for seizure analysis from intracranial EEG. *J Mach Learn Res.* 2006; 7:1025–1044.




Fig. 1. Feature vector formation steps.




Fig. 2. Intracranial EEG on a subset of channels in a 12-second epoch containing the onset of a seizure in Patient 6. The dashed vertical line indicates the expert-marked onset.



Fig. 3. Intracranial EEG on a subset of channels in a 12-second epoch containing the onset of a seizure in Patient 6 different from the seizure in Fig. 2. The dashed vertical line indicates the expert-marked onset.




Fig. 4. Patient 7's intracranial EEG on a subset of electrodes in two epochs showing similar activity, particularly on the right subfrontal channel "RSbF 5–RSbF 6." (A) Seizure activity in a 4-second epoch beginning approximately 10 seconds after the expert-marked onset. (B) An interictal 4-second epoch during which false detection occurred. This activity was not judged to constitute a seizure by the clinicians.



Fig. 5. Patient 7's intracranial EEG on a subset of electrodes in two epochs showing similar activity, particularly on the right posterior temporal channels "RPT 3–RPT 4" and "RPT 5–RPT 6." (A) Seizure activity in a 4-second epoch beginning approximately 7 seconds after the expert-marked onset. (B) An interictal 4-second epoch during which false detection occurred. As before, this activity was not judged to constitute ictal activity by clinical reviewers.



Fig. 6. Approximate timing of false alarms that occurred in tests on data from Patient 8.

Patient data set information and sensitivity, median latency, estimated false alarm rate obtained for each patient data set from evaluation of the seizure onset detector.

Table 1

	Patient									
	1	2	3	4	5	6	7	8	9	10
Total time tested (h)	52	127	71	29	96	68	39	146	148	104
Number of seizures	3	6	6	5	8	3	7	12	11	6
Number of electrodes	98	56	114	34	68	40	54	64	80	85
Electrode type	Grid and depth	Depth	Grid and depth	Depth	Depth	Depth	Depth	Depth	Depth	Grid and depth
Sensitivity	100%	100%	100%	100%	100%	66%	100%	100%	100%	100%
Median latency (s)	6.5	3.25	3.25	5	5.75	4.25	4.5	6	4.5	18.5
False alarms/24 h	0	0.4	1.7	0	0.8	0	25.3	2.5	0.3	2.5

Table 2

Sensitivities, median latencies, and false alarm rates for the particular case where $K = 1^a$.

	Patient									
	1	2	3	4	5	6	7	8	9	10
Sensitivity	100%	100%	100%	100%	100%	66%	100%	91%	100%	100%
Median latency (s)	5.5	2.25	2.25	4	4.75	3.25	3.5	4	2.5	16.5
False alarm rate (per 24 h)	0	1.3	3.7	1.7	5	1.4	45.1	13.5	1	6.2

^aThe latencies decrease at the expense of higher false alarm rate.