Journal of Machine Learning Research 7 (2006) 1025 – 1044 Submitted 5/05; Revised 1/06; Published 6/06

One-Class Novelty Detection for Seizure Analysis from Intracranial EEG

Andrew B. Gardner AGARDNER@NEUROENG.ORG
Departments of Bioengineering and Neurology
University of Pennsylvania
301 Hayden Hall
Philadelphia, PA 19104, USA

Abba M. Krieger KRIEGER@WHARTON.UPENN.EDU
Department of Statistics, Wharton School
University of Pennsylvania
3733 Spruce Street
Philadelphia, PA 19104, USA

George Vachtsevanos GJV@ECE.GATECH.EDU
Department of Electrical and Computer Engineering
Georgia Institute of Technology
777 Atlantic Drive
Atlanta, GA 30332, USA

Brian Litt LITTB@MAIL.MED.UPENN.EDU
Departments of Bioengineering and Neurology
Hospital of the University of Pennsylvania
3 West Gates, 3400 Spruce Street
Philadelphia, PA 19104, USA

Editor: Leslie Pack Kaelbing

Abstract

This paper describes an application of one-class support vector machine (SVM) novelty detection for detecting seizures in humans. Our technique maps intracranial electroencephalogram (EEG) time series into corresponding novelty sequences by classifying short-time, energy-based statistics computed from one-second windows of data. We train a classifier on epochs of interictal (normal) EEG. During ictal (seizure) epochs of EEG, seizure activity induces distributional changes in feature space that increase the empirical outlier fraction. A hypothesis test determines when the parameter change differs significantly from its nominal value, signaling a seizure detection event. Outputs are gated in a "one-shot" manner using persistence to reduce the false alarm rate of the system. The detector was validated using leave-one-out cross-validation (LOO-CV) on a sample of 41 interictal and 29 ictal epochs, and achieved 97.1% sensitivity, a mean detection latency of

-7.58 seconds, and an asymptotic false positive rate (FPR) of 1.56 false positives per hour (Fp/hr). These results are better than those obtained from a novelty detection technique based on Mahalanobis distance outlier detection, and comparable to the performance of a supervised learning technique used in experimental implantable devices (Echauz et al., 2001). The novelty detection paradigm overcomes three significant limitations of competing methods: the need to collect seizure data, precisely mark seizure onset and offset times, and perform patient-specific parameter tuning for detector training.

Keywords: seizure detection, novelty detection, one-class SVM, epilepsy, unsupervised learning

1 Introduction

Epilepsy, a neurological disorder in which patients suffer from recurring seizures, affects approximately 1% of the world population. In the United States, 200,000 new cases are reported annually. There are more than 30 distinct classes of seizure. Their manifestations range from subtle, abnormal sensations to unpredictable changes in awareness, to immediate loss of consciousness and convulsions. In spite of available dietary, drug, and surgical treatment options, more than 25% of individuals with epilepsy have seizures that are uncontrollable (Kandel, Schwartz & Jessel, 1991). Daily life for these patients is greatly impaired—education, employment, and even transportation can become difficult endeavors. Many new therapies for medically resistant epilepsy are being investigated. Among the most promising are implantable devices that deliver local therapy, such as direct electrical stimulation or chemical infusions, to affected regions of the brain. These treatments rely on robust algorithms for seizure detection to perform effectively.

Over the past 30 years seizure detection technology has matured. Despite impressive advances, all reported approaches suffer from one or more of the following limitations:

- Accurate detection requires careful, patient-specific tuning
- Seizure detections do not occur "early enough" (i.e., interventions are more likely to be effective if therapy is administered with minimal delay following onset)
- *A priori* localization of the seizure focus is required
- Usefulness for poorly localized epilepsies is limited
- Seizure data (which is expensive to collect) is required for training

Techniques for overcoming some or all of these limitations hold promise for more precise and widely applicable methods to control or eliminate seizures. This paper presents one technique for improving the state of the art in seizure detection by reformulating the task as a time-series novelty detection problem. While seizure detection is traditionally considered a supervised learning problem (e.g., binary classification), an unsupervised approach allows for uniform treatment of seizure detection and prediction, and offers four key advantages for implementation. First, there is no need to perform supervised, patient-specific tuning during training. Second, the assumption that seizures are electrographically homogeneous—often required for classifier training due to very small data sets—is relaxed. Third, there is no need to collect seizure data for training. Such data collection is typically expensive (seizures occur infrequently, and patients must be continually monitored until an event is observed), and often invasive (e.g., craniotomy or burr hole for electrode implantation). Finally, there is no need to precisely mark seizure intervals. This practical issue is often overlooked, but is critical for training and validation: while expert markers usually agree on the presence of a seizure, there is considerable variability in marking its onset and offset.

Other researchers have investigated novelty detection for event detection from time series, for instance by directly extending the Incremental SVM algorithm (Tax & Laskov, 2003), or

modeling novelty region-of-support evolution to detect change-points (Desobry & Davy, 2003). In contrast to these approaches, we robustly detect empirical changes in the novelty parameter itself, and use these change-points to segment the (EEG) time series. For "properly chosen" features, novelties correspond well with the ictal events of interest, and our EEG time series are successfully segmented in a one-class-from-many manner.

2 Background

In this section, we present a brief review of seizure-related terminology, the seizure detection literature, and the one-class SVM.

2.1 Seizure-Related Terminology

Seizure analysis refers collectively to algorithms for seizure detection, seizure prediction, and automatic focus channel identification. These analyses are primarily performed on the EEG. In this study, analyses were carried out on the intracranial EEG (IEEG), which has considerably better spatial resolution, higher signal-to-noise ratio, and greater bandwidth than scalp EEG. When multiple channels are considered, the electrode location that exhibits the earliest evidence of seizure activity is labeled the *focus channel*. It is convenient to describe segments of the EEG signals by their temporal proximity to seizure activity. The *ictal* period refers to the time during which a seizure occurs. The *interictal* period is the time between successive seizures. The *unequivocal electrographic onset (UEO)* is defined as the earliest time that a seizure occurrence is evident to an epileptologist viewing an EEG without prior knowledge that a seizure follows; the *unequivocal clinical onset (UCO)* is the earliest time that a seizure occurrence is apparent by visually observing a patient. *Seizure onset* in this paper is synonymous with UEO. It is worth noting that the UEO almost always precedes the UCO by several seconds, and that many previously published papers defined "seizure onset" as the UCO.

2.2 Seizure Detection

Early attempts to detect seizures began in the 1970s (Viglione, Ordon & Risch, 1970; Liss, 1973) and primarily considered scalp EEG recordings to detect the clinical (and less frequently) electrographic onset of seizures. In 1990, Gotman reported a technique for automated seizure detection that achieved 76% detection accuracy at 1 Fp/hr for 293 seizures recorded from 49 patients (Gotman, 1990). In 1993, it was shown that the short-time mean Teager energy could be used to detect seizures from electrocorticograms (Zaveri, Williams & Sackellares, 1993). Their detector achieved 100% detection accuracy on an 11-seizure database. In 1995, Qu and Gotman presented an early seizure warning system trained on template EEG activity that achieved 100% detection accuracy at a mean detection latency of 9.35 seconds and false alarm rate of 0.2 Fp/hr (Qu & Gotman, 1995). Similar results were also reported using time- and frequency-domain features classified by a k-nearest neighbor classifier (Qu & Gotman, 1997). In 1998, Osorio et al. claimed 100% detection sensitivity with a mean detection latency of 2.1 seconds using a wavelet-based measure called seizure intensity. They analyzed a database of 125 patients, but the same data were used for training and validation (Osorio, Frei & Wilkinson, 1998). The algorithm was more extensively analyzed in 2002 using offline electrocorticogram recordings; again, 100% sensitivity was reported, with detection latencies ranging from 1.8 – 31.1 seconds (Osorio et al., 2002).

Several successful attempts at seizure detection using artificial neural network classifiers have been reported since 1996 (Khorasani & Weng, 1996; Webber et al., 1996; Gabor, 1998; Esteller, 2000). Evaluation of 31 distinct features (Esteller, 2000) showed that fractal dimension, wavelet packet energy, and mean Teager energy were especially promising for seizure detection. In 2001, Esteller reported a detector based on the line length feature that achieved a mean

detection latency of 4.1 seconds at a false alarm rate of 0.051 Fp/hr (Esteller et al., 2001). A total of 111 seizures (many subclinical) were used for validation. NeuroPace, Inc., subsequently reported a similar detector based upon this work that achieved 97% sensitivity at a mean detection latency of 5.01 seconds (Echauz et al., 2001). This detector was evaluated on 1265 hours of IEEG data, but was tuned heuristically in a patient-specific manner. The NeuroPace detector claims represent the state of the art in seizure detection performance. More complete reviews of the seizure detection and prediction literature are available elsewhere (Litt & Echauz, 2002; Gardner, 2004).

2.3 Novelty Detection

Traditional classification architectures rely on empirical risk minimization algorithms to specify "good" models for a classification decision function; as such, they are prone to over- or underfitting. In addition, their performance tends to be highly sensitive to parameter tuning and researcher skill. Statistical learning theory poses a structural risk minimization (SRM) criterion that balances the trade-off between good empirical performance (i.e., classification accuracy on training data) and good generalization ability (i.e., classification accuracy on unseen data). One popular application of SRM is the SVM, first presented in 1992 (Boser, Guyon & Vapnik, 1992). The basic idea behind the SVM is to find a hyperplane in a feature space that "optimally" separates two classes. Many other linear learning machines have been considered for this task, but the SVM yields a unique solution that can be shown to minimize the expected risk of misclassifying unseen examples (Vapnik, 1999). Training algorithms involve the solution of a well-known optimization problem, constrained quadratic programming, that is computationally efficient and yields global solutions. Several excellent tutorials provide historical context and details on the SVM (Burges, 1998; Bennett & Campbell, 2000; Müller et al., 2001).

In 1998, Schölkopf et al. introduced an extension to SVMs to estimate the support of a distribution (Schölkopf et al., 1999). Their motivation was to solve a simplified version of the density estimation problem, e.g., finding a minimum volume quantile estimator that is "simple." The solution they arrived at, the one-class SVM, was introduced for novelty detection.

Definition 1 (Novelty Detection). *Given a set of independent identically distributed (iid) training samples,* $x_1, \cdots, x_n \in X \subseteq \mathbf{R}^N$, *drawn from a probability distribution in feature space,* P, *the goal of novelty detection is to determine the "simplest" subset,* S, *of the feature space such that the probability that an unseen test point,* x', *drawn from* P *lies outside of* S *is bounded by an a priori specified value,* $\upsilon \in (0,1]$.

In the one-class formulation, data are first mapped into a feature space using an appropriate kernel function and then maximally separated from the origin using a hyperplane. The hyperplane parameters are determined by solving a quadratic programming problem, similar to the basic SVM case:

$$\min\left(\frac{1}{2}\|w\|^2 + \frac{1}{\nu l}\sum_{i=1}^{l}\xi_i - \rho\right) \tag{1}$$

subject to

$$(w \cdot \Phi(x_i)) \geq \rho - \xi_i \quad i = 1,2,\ldots,l \quad \xi_i \geq 0 \tag{2}$$

where w and ρ are hyperplane parameters, Φ is the map from input space to feature space, ν is the asymptotic fraction of outliers (novelties) allowed, l is the number of training instances, and ξ is a slack variable. For solutions to this problem, w and ρ, the decision function

$$f(x) = \text{sgn}(w \cdot \Phi(x) - \rho) \tag{3}$$

specifies labels for examples, e.g., -1 for novelty instances. Figure 1 shows the geometry of the one-class SVM in feature space.

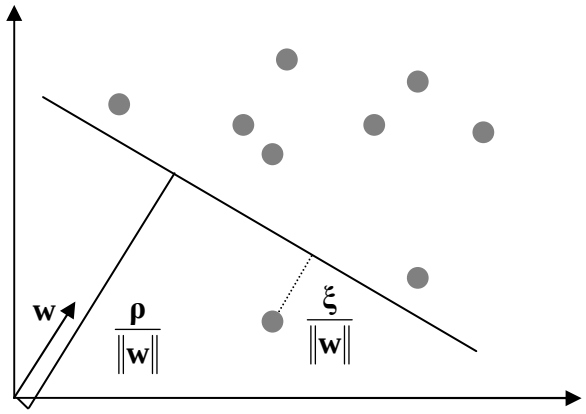

Figure 1: Geometry of the ν-SVM in feature space. Note the hyperplane and associated parameters, ρ and w, and the slack-variable, ξ, penalizing misclassifications.

Basic properties of the one-class SVM were proven in the initial paper (Scholkopf et al., 1999). The most important result is the interpretation of ν as both the asymptotic fraction of data labeled as outliers, and the fraction of support vectors returned by the algorithm. Implementation of the one-class SVM algorithm requires the following specifications: kernel function, kernel parameters, outlier fraction, and separating point in feature space. As with the basic SVM, there is no automatic method for specifying one-class SVM model parameters, but the interpretation of ν eases this task to some degree: the choice of outlier fraction should incorporate prior knowledge about the frequency of novelty occurrences (for example, a typical value for patient seizure frequency). Additionally, smaller values of ν increase the computational efficiency of the algorithm. The choice of origin as the separation point is arbitrary and affects the decision boundary returned by the algorithm. Other work (e.g., Hayton et al., 2001; Manevitz & Yousef, 2001) has addressed separation point selection given partial knowledge of outlier classes.

3 Methodology

In this section, we describe and discuss the experimental methods for detecting seizures under a novelty detection framework. A block diagram of this system is shown in Figure 2.

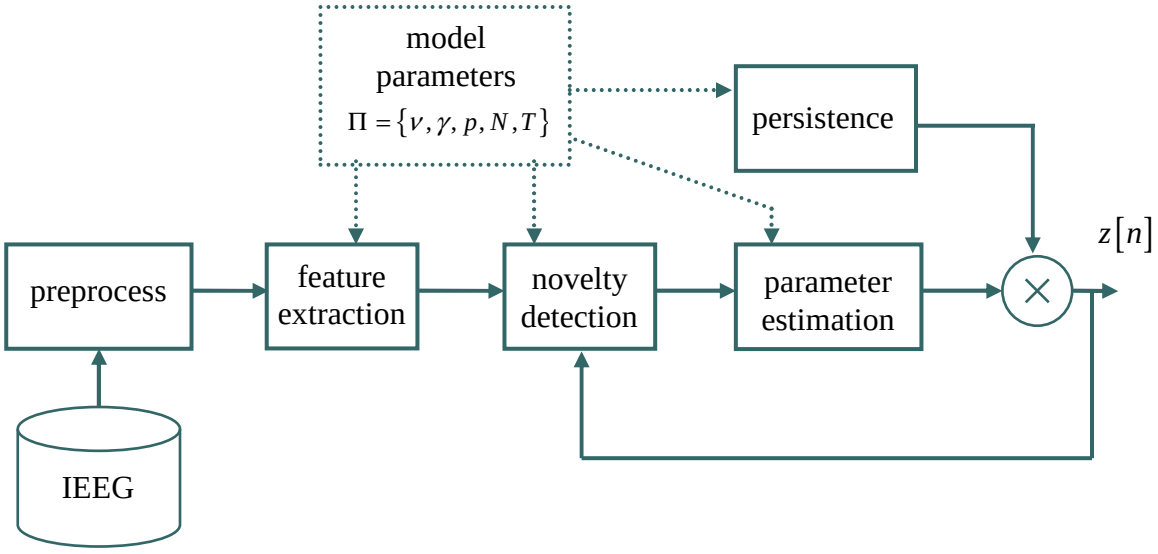

Figure 2: The seizure analysis architecture. IEEG time series data is block-processed in stages to produce the final output sequence, $z[n]$, indicating the presence/absence of ictal activity.

3.1 Human Data Preparation

The data analyzed were selected from intracranial EEG recordings of epilepsy patients implanted as part of standard evaluation for epilepsy surgery. Patients diagnosed with mesial temporal lobe epilepsy were observed in a hospital for 3 to 14 days. Between 20 and 36 electrodes were surgically placed either on the brain (grids or strips of electrodes), or in the brain substance (depth electrodes), and simultaneous IEEG and video were recorded. The IEEG data were amplified, bandpass-filtered (cutoffs at 0.1 Hz and 100 Hz), and digitized at 200 samples/second, 12 bits-per-sample resolution. Five consecutive patients with seizures arising from the temporal lobe(s) were selected for review, and the corresponding data were expertly and independently marked by two certified epileptologists to indicate UEO and UCO times. Collectively, these five patient records contain over 200 hours of data. Further details on this database are available elsewhere (D'Alessandro, 2001; Gardner, 2004).

Ictal epochs were selected from the focus channel for each temporal lobe seizure that a patient exhibited. Two patients exhibited some seizures with extra-temporal focal regions: those events were excluded from further analysis. Ictal epochs were extracted in a consistent manner such that the UEO occurred at a 10-minute offset within the epoch, allowing for analysis of both pre-ictal and post-ictal regimes. Interictal epochs from each patient were randomly selected. All epochs were expertly reviewed to ensure the absence of recording artifacts. The final data set consisted of 29 ictal- and 41 interictal epochs, each of 15-minute duration.

3.2 Feature Extraction

Many features have been proposed for seizure analysis (Esteller, 2000; D'Alessandro, 2001; Esteller et al., 2001). We selected a feature vector, q, composed of three energy-based statistics that have proven especially effective for seizure detection: mean curve length, CL; mean energy, E; and mean Teager energy, TE,

$$CL[n] = \log\left(\frac{1}{N} \sum_{m=n-N+2}^{n} \left| x[m] - x[m-1] \right| \right) \tag{4}$$

$$E[n] = \log\left(\frac{1}{N} \sum_{m=n-N+1}^{n} x[m]^2 \right) \tag{5}$$

$$TE[n] = \log\left(\frac{1}{N} \sum_{m=n-N+3}^{n} \left(x[m-1]^2 - x[m]\, x[m-2] \right) \right) \tag{6}$$

where $x[m]$ is an EEG time series, and N is the window length. We applied logarithmic scaling for feature normalization. Features were extracted using a block processing approach. In block processing, the data are windowed, a feature vector is computed, and the window is advanced in time. The selection of window length is an important issue (Esteller, 2000). Values typically range between 0.25 and 5 seconds; we used 1-second windows with 0.5-second overlap.

3.3 One-Class SVM

Feature extraction was performed on interictal epochs to generate feature vectors for training. A one-class SVM classifier was implemented using LIBSVM, a freely-available library of SVM tools available from http://www.csie.ntu.edu.tw/~cjlin/libsvm. A Gaussian radial basis function ($\gamma = 1.0$) was selected as the kernel function, and $\nu = 0.1$ was chosen consistent with the estimated fraction of ictal data. The resulting classifier model was stored for subsequent use in testing.

3.4 Parameter Estimation

For a stationary process, the one-class SVM novelty parameter, ν, asymptotically equals the outlier fraction. We exploit this property by training on features which strongly discriminate interictal from ictal EEG: features are stationary during interictal periods, but change markedly during periods of seizure activity, causing significant changes in the empirical outlier fraction.

We modeled classifier outputs, $y \in \{+1, -1\}$, as (iid) Bernoulli random variables where $P(novelty) = P(y = -1) = \nu$. We assumed that $\nu = \nu_0$ for interictal EEG, and $\nu = \nu_1 > \nu_0$ for ictal EEG. At each output sample, we computed the maximum likelihood estimate of the outlier fraction, $\hat{\nu}$, as

$$\hat{\nu} = \frac{n_{neg}}{n} = \frac{1}{2}\left(1 - \frac{1}{n}\sum_{i=1}^{n} y[i]\right) \tag{7}$$

where n_{neg} is the number of negative output occurrences in the n most recent samples of y. Note that the sequence length, n, affects the adaptation rate of the system. We then used this estimate to compute a seizure event indicator variable,

$$z[k] = \mathrm{sgn}\left(\hat{\nu} - C\right) \tag{8}$$

where $z = +1$ if a seizure is indicated or $z = -1$ otherwise, and $C \in [0,1]$ is a threshold parameter. Thresholding is equivalent to a standard hypothesis test of $H_0 : \nu = \nu_0$ vs. $H_1 : \nu = \nu_1$ where the null hypothesis is rejected if $\hat{\nu} > C$. For nominal values of $n = 20$ and $\nu_0 = 0.1$, we retained the null hypothesis (that is, we declared a frame to be interictal) if we observed fewer than five novelty outputs ($C = 0.8$). Under the iid assumption, this rule has a 4.33% chance of falsely rejecting the null hypothesis (i.e., producing a false positive). The chance of committing a Type II error (i.e., producing a false negative) depends on the unknown value ν_1. We calculated this error rate for several plausible values of ν_1 in Section 4.2.

3.5 Persistence (Detector Refractory Period)

During early experiments we observed that the detector tended to generate novelty events (i.e., "fire") in bursts, with increasing frequency near seizure onset. This behavior may indicate the presence of preictal states, periods of EEG activity that are likely to transition from interictal to ictal state. The bursty behavior can be problematic for performance assessment as multiple detections of a single seizure, or multiple false positive declarations may occur during a short interval of time. To address this problem, a refractory parameter, T_R, was introduced to the detection system. The refractory parameter specifies an interval during which the detector, if triggered, maintains its state and ignores subsequent triggers. In this sense it behaves like a "one-shot" device familiar from digital circuits. The use of this refractory rule is termed persistence.

Persistence offers an improvement to the basic system beyond false positive rate improvement: it allows for the characterization of the detector over a range of detection time horizons. As persistence decreases, one expects the false positive rate to increase and the detection latency to approach zero seconds. Conversely, as persistence increases, one expects the false positive rate to decrease, asymptotically approaching a value determined jointly by the novelty parameters of the system (some fraction of the data will always be novel) and the actual novelty rate due to epileptiform activity. Figure 3 illustrates the use of persistence. We heuristically set the detector persistence to $T_R = 180$ seconds for our experiments.

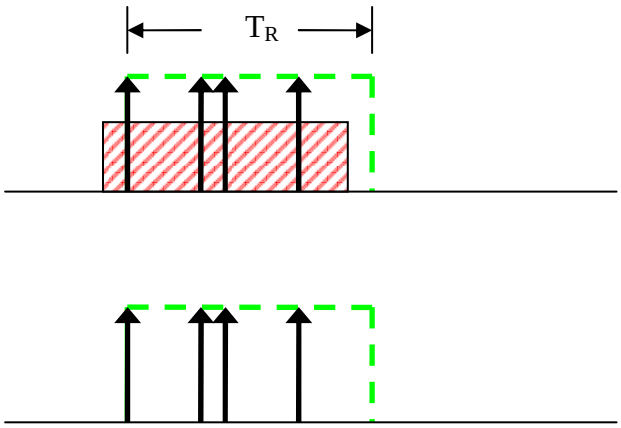

Figure 3: Examples of persistence for improving detector false alarm performance. (*Top*) Ictal epoch showing seizure activity (*red, diagonal hatching*). (*Bottom*) Interictal epoch. When persistence is applied, detections (*arrows*) are treated as a single event (*dashed green line*).

3.6 Performance Metrics

The detector was evaluated using LOO-CV and identical model parameters for each patient. Training was only performed using interictal epochs, however, testing was performed on each ictal segment, in addition to the withheld interictal epoch. This scheme yields $C\left(N_{BL},1\right)$ interictal- and $C\left(N_{BL},1\right)\times N_{SZ}$ ictal statistics per patient, where N_{BL}, N_{SZ} are the patient-specific number of interictal and ictal epochs, respectively. From these statistics we estimate three key performance metrics: sensitivity, false positive rate, and mean detection latency.

The detector's sensitivity (9) and false positive rate (10) measure its classification accuracy:

$$S = \frac{TP}{TP + FN} \tag{9}$$

$$FPR = \frac{FP}{T} \tag{10}$$

where *TP*, *FN*, and FP are the number of block true positives, block false negatives, and block false positives; and T is the duration (in hours) of the data analyzed. A block true positive occurs when the detector output, after applying persistence, correctly identifies an interval containing a seizure onset (c.f., Figure 4). Block false negatives and false positives occur when the detector incorrectly labels interictal and ictal intervals, respectively.

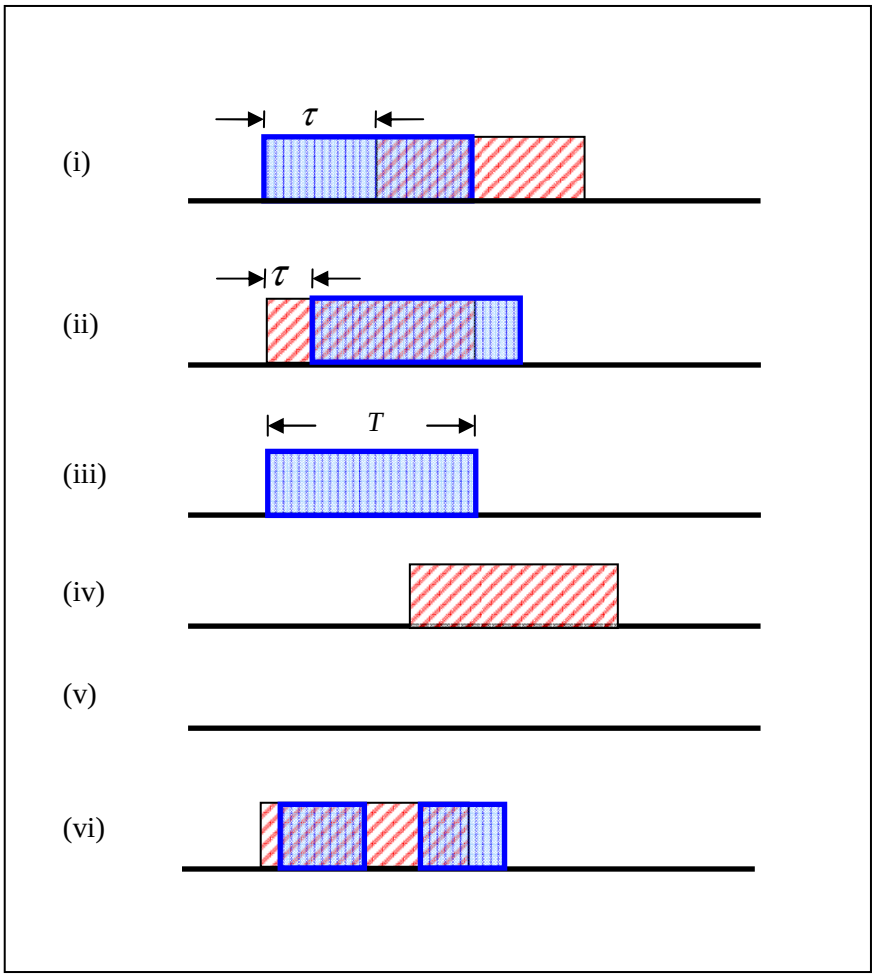

Figure 4: Temporal relationships considered in detector evaluation: intervals representing detected novelty (*blue, vertical hatching*) and ictal activity (*red, diagonal hatching*). *(i), (ii)* Two examples of block true positives (e.g., the novelty output interval overlaps ictal activity). The detection latency, τ, is also shown. An early-detection results in negative latency. *(iii)* A false positive detector error. *(iv)* A false negative detector error. *(v)* A true negative. *(vi)* An example of a degenerate case (multiple detection) producing both a true positive and a false positive event.

Mean detection latency (11) measures detector responsiveness:

$$\mu_\tau = \frac{1}{N}\sum_{i=1}^{N}\tau_i \tag{11}$$

where τ_i is the detection latency of each detected seizure. A negative latency indicates seizure event detection prior to the expert-labeled onset time.

3.7 Benchmark Novelty Detection

To provide a reference for the relative performance of our algorithm, and the general application of unsupervised learning to the seizure detection problem, we implemented a simple benchmark novelty detection algorithm.

During training, feature vectors, q, were extracted from IEEG time series (c.f. 3.2) and used to estimate the covariance matrix, Σ, and mean, μ, of the training data. We subsequently computed the Mahalanobis distances

$$D_M(q) = \sqrt{(q-\mu)^T \Sigma^{-1} (q-\mu)} \tag{12}$$

between each sample in the training data set and the centroid of the training set. An outlier threshold, K, was selected as the ν quantile of the Mahalanobis distances. As with the one-class SVM, we set $\nu = 0.1$.

During testing, feature vectors were thresholded to produce a frame-wise novelty sequence, y,

$$y[n] = \begin{cases} +1, & D_M(q_n) < K \\ -1, & D_M(q_n) \geq K \end{cases} \tag{13}$$

as a replacement for the SVM classifier output. This sequence was processed in the same manner as before (c.f. 3.4) to generate detections.

4 Results

In this section we present the results of both seizure detection approaches. Details on the effects of varying the one-class SVM model parameters—ν, γ, p, N, and T— and results from a genetic algorithm optimization are given in Gardner (2004).

4.1 Performance

Detection statistics from our LOO-CV analysis are presented in Table 1. Columns show patient id, the number of epochs analyzed for interictal (N_{BL}) and ictal data (N_{SZ}), the fraction of false positive detections on interictal (FP_{BL}) and ictal (FP_{SZ}) trials, the fraction of seizure epochs producing false positives (M_{FP}), the fraction of false negative detections (FN), and the mean detection latency ($\bar{\tau}$). The bottom row of the table shows aggregate statistics weighted by the number of seizures or number of baselines.

Patient	N_{BL}	N_{SZ}	*One-class Novelty*					*Benchmark Novelty*				
			FP_{BL}	FP_{SZ}	M_{FP}	FN	$\overline{\tau}$	FP_{BL}	FP_{SZ}	M_{FP}	FN	$\overline{\tau}$
1	6	5	0.00	0.10	0.20	0	2.07	0.17	0.13	0.20	0	4.60
2	9	7	0.22	0.43	0.57	0	-13.6	0.33	1.03[c]	0.71	0	-9.51
3	10	6	0.40[a]	0.13	0.17	0	6.57	0.40[a]	0.17	0.17	0	7.08
4	10	6	1.00	0.48	0.17	0.12	-6.57	1.00	0.48	1.00	0.17	-2.08
5	6	5	0.00	0.03	0.20	0	-27.0[b]	0.33	0.17	0.20	0	-24.9[b]
	41	**29**	**0.39**	**0.28**	**0.28**	**0.029**	**-7.58**	**0.49**	**0.47**	**0.45**	**0.041**	**-4.76**

[a]All false positives occurred on a single ictal epoch.
[b]Several seizure onsets were originally mislabeled by as much as 110 seconds. Results in this table are calculated from the corrected markings.
[c]Note that multiple false positive events per epoch can produce fractional values greater than one.

Table 1: Summary of detection statistics. Bottom row of table summarizes aggregate statistics.

We estimate the FPR over interictal EEG from the data in Table 1 by dividing FP_{BL} by the epoch duration (0.25 hours), yielding 1.56 Fp/hr for the one-class technique, and 1.96 Fp/hr for the benchmark technique. Since ictal events are rare, and the aggregate false positive rate on ictal segments is lower than the corresponding rate on interictal segments, we take the interictal FPR as an asymptotic measure of the overall FPR.

We reviewed the results for those patients (2, 4, and 5) with negative mean detection latencies. For each of these patients we found that the distribution of detection latencies was skewed, and a fraction (less than one-third) of the models detected seizures early. The median detection latencies for these patients, which might give a more balanced view of performance, ranged between 1.5 and 9.8 seconds for both models; the one-class delays were always less than the benchmarked values.

The SVM seizure detector achieved 97.1% sensitivity and a mean detection latency of -7.58 seconds at an estimated 1.56 Fp/hr. Representative IEEG time series, novelty sequences, and estimated outlier fractions for interictal- and ictal epochs are shown in Figures 5 and 6. As expected, the outlier fraction remained near its (small) nominal value except during periods of seizure activity. Onsets were detected quickly, and the entire seizure event—not just the onset—was correctly identified as novel. The near-zero false negative rate (FNR) of the detector was surprising because the data used for training originated from unknown states of consciousness (e.g., sleep or wake). Typically, seizure detection performance is drastically affected by patient state-of-consciousness; evaluation on larger data sets with concomitant sleep staging information will provide a better estimate of the true FNR.

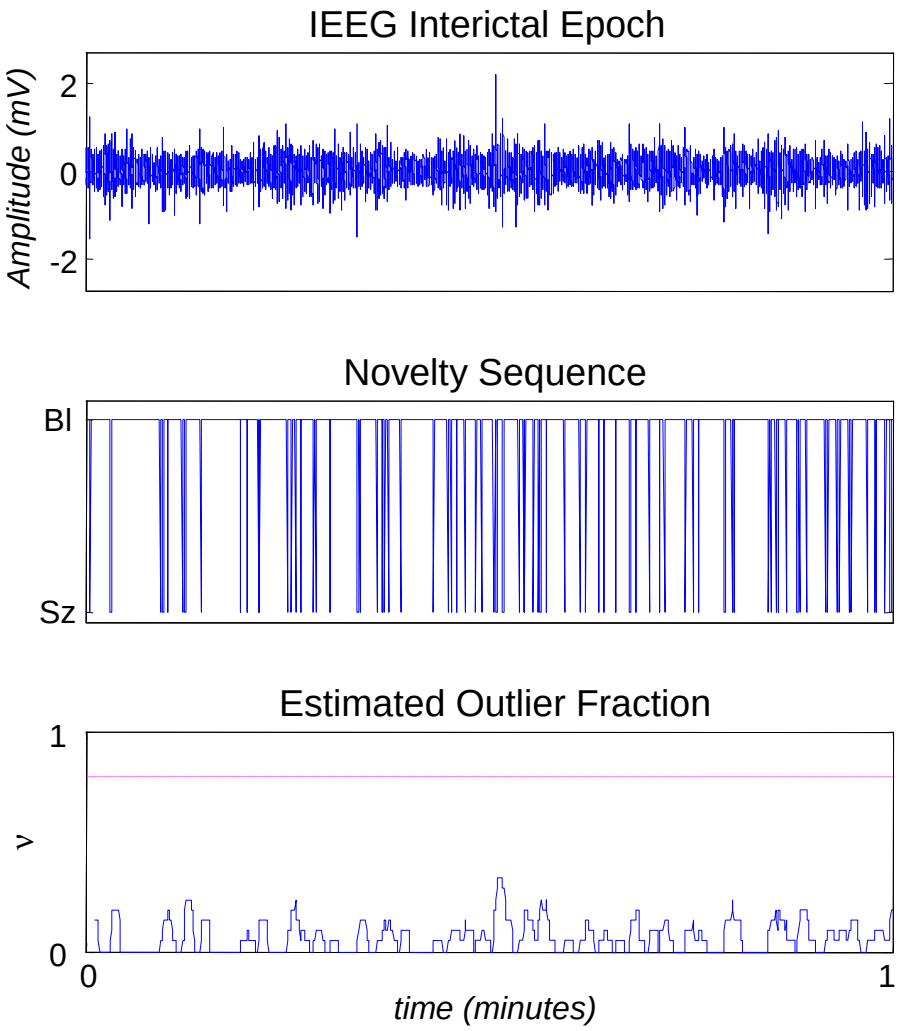

Figure 5: A typical interictal epoch. (*Top*) IEEG signal, (*Middle*) frame-wise output of the novelty detector, z, (*Bottom*) estimated outlier fraction (dashed line is 0.8). The mean of ν in this figure is 0.063.

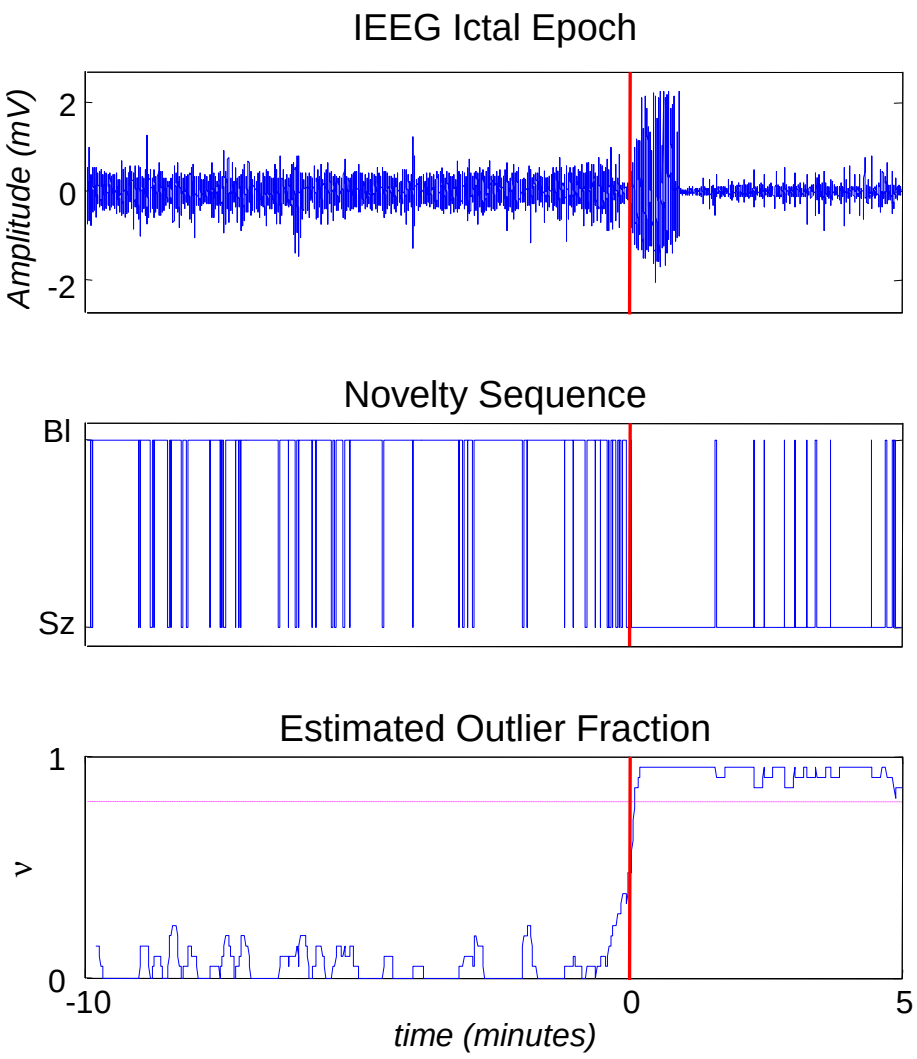

Figure 6: A typical ictal epoch. (*Top*) IEEG signal. The earliest electrographic change is visible as the beginning of the pinched region prior to the high-amplitude seizure onset. The UEO occurs at time zero, (*Middle*) frame-wise output of the novelty detector, z , (*Bottom*) estimated outlier fraction and 0.8 threshold. The detector has a latency of about 3 seconds in this example.

The SVM detector's mean detection latency outperformed all previously reported seizure detection algorithms. It should be noted, however, that this result is attributable to the large fraction of seizures (27%) that were detected early. This finding suggests the presence of two subclasses of seizures: those that are merely detectable, and those that may be predictable. These classes of seizures appear to be patient-dependent.

A direct comparison to other published detection algorithms is generally not meaningful due to the disparity of data sets that each research group operates on. However, NeuroPace (Echauz et al., 2001) previously evaluated their supervised algorithm on the same data set. While they did not perform cross-validation, and optimized in-sample for each patient, their reported results—a

mean detection latency of 5.01 seconds at 97% sensitivity and 0.013 Fp/hr—support the use of our approach.

The benchmark seizure detector achieved 95.8% sensitivity and a mean detection latency of -4.76 seconds at 1.96 Fp/hr. Both techniques are surprisingly effective at seizure detection, but the one-class SVM method performed consistently better, especially with respect to false positive rate. To gain insight into the relative performance difference between the two approaches, we examined the feature space for a single patient. Figure 7 shows the marginal distributions of features for both interictal and ictal data. It is clear from this figure that the feature distributions are highly skewed and possibly bimodal. An obvious explanation for the discrepancy in performance is that the normality assumption of the benchmark detector is severely violated, and the non-parametric estimation of the one-class SVM is better for modeling the data. The fact that the one-class SVM performs better, albeit on a limited number of patients, suggests that it tends to exclude vectors in feature space that appear more commonly when seizures occur as compared to the benchmark approach. Additionally, we examined the regions-of-support for this patient produced by each algorithm (Figure 8).

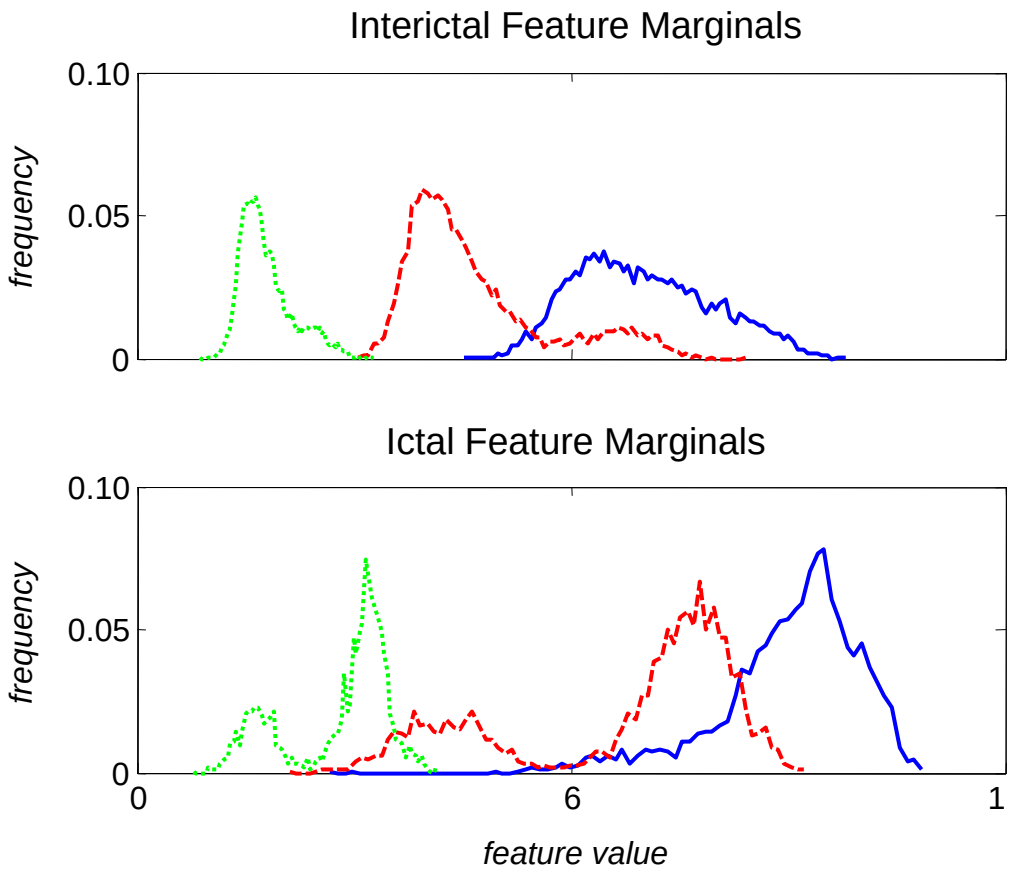

Figure 7: Representative marginals of the feature vector—E (*solid blue*), TE (*dashed red*), CL (*dotted green*)—for patient 5 corresponding to interictal (*top*) and ictal (*bottom*) frames.

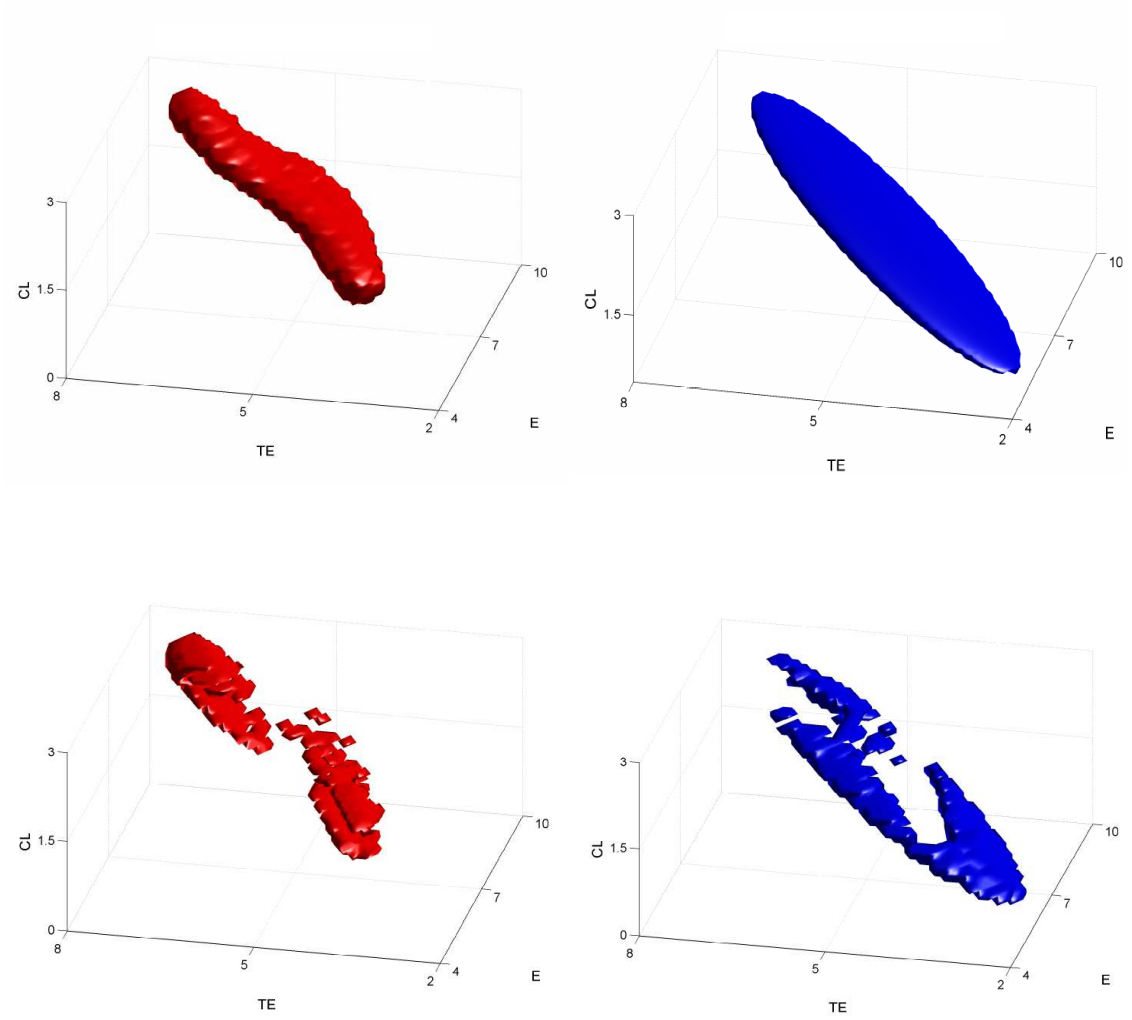

Figure 8: Representative isosurfaces in interictal feature space produced by each method. (*Top-left*) S_1, the $\nu = 0.1$ enclosing surface for the one-class SVM; (*Top-right*) S_2, the $\nu = 0.1$ enclosing surface for the benchmark method; (*Bottom-left*) $S_1 \setminus S_2$, the volume unique to the one-class SVM; (*Bottom-right*) $S_2 \setminus S_1$, the volume unique to the benchmark method.

Both approaches, SVM and Mahalanobis, find regions, S_1 and S_2, in feature space that include 90% of the observations from interictal data. It is interesting to note that, although the overlap of the regions, $S_1 \cap S_2$, must contain at least 80% of the training samples, 25.2% of the volume of S_1 and 37.5% of the volume of S_2 are non-intersecting. The minimum-volume property of the one-class SVM is also evident—S_1, is 84.4% of the volume of the benchmark technique—and may be a contributing factor to its increased performance over the Mahalanobis method.

4.2 Detector Output Analysis

We analyzed a sample of 850 interictal detector outputs and confirmed that the empirical outlier fraction equaled its nominal value, 0.10. We also investigated the performance of the detector output under the hypothesis $H_1 : \nu = \nu_1 > \nu_0$, assuming iid outputs. Illustratively, we considered $\nu_1 = 0.3$ and $\nu_1 = 0.5$ for the probabilities of a novelty occurrence during ictal epochs. Results in Table 1 show that the probability of falsely retaining the null hypothesis is small, and is of course smaller for $\nu_1 = 0.5$ than for $\nu_1 = 0.3$. This explains the superior FNR performance of the detector that we observed.

We performed logistic regressions between the outputs at times t, $t-1$, and $t-2$ to test our assumption that detector outputs are Bernoulli. We observed significant ($P < 0.001$) serial dependence. Empirically, the conditional probability of a novel detector output given a previous novelty output increases dramatically from 0.1 to 0.3. This analysis suggests that the detector output sequence obeys a Markov process where the probability at each point in time of a novelty is $P(z_t = -1) = \nu_0 = 0.1$, but the conditional probabilities for novelty outputs are $P(z_t = -1 \mid z_{t-1} = -1) = 0.3$ and $P(z_t = -1 \mid z_{t-1} = +1) = 0.078$.

We wrote a program to compute the probability of observing k novel outputs in N trials under the Markov process described above, and repeated our performance analysis. The results (Table 2) clearly show that the performance of the detector is worse under the serial dependence model.

	Binomial Output			Markovian Output		
	H_0	H_1	H_1	H_0	H_1	H_1
		$p_1 = 0.3$	$p_1 = 0.5$		$p_1 = 0.3$ $p_1^* = 0.5$	$p_1 = 0.5$ $p_1^* = 0.75$
Normal	0.9567	0.2374	0.0059	0.9175	0.3095	0.0736
Seizure	0.0433	0.7626	0.9941	0.0825	0.6905	0.8264

Table 2: An analysis of the hypothesis test for the detector output for both the binomial, and Markov cases for the rule where we declare an event if 5 or more out of 20 outputs are novelties. The estimated probability that an ictal frame is declared novel is p_1, and its corresponding conditional probability, $P(z_k = -1 \mid z_{k-1} = -1)$, is p_1^*.

5 Conclusions

Traditional approaches to seizure detection rely on binary classification. They require seizure data for training, which is difficult and invasive to collect, and do not address the class imbalance problem between interictal and ictal EEG, as less than 1% of EEG data from epileptic patients is seizure-related. These approaches assume that seizures develop in a consistent manner and seek to identify features and architectures that discriminate seizure EEG from "other" EEG. In contrast, we have presented a technique for seizure detection based on novelty detection that operates by modeling the dominant data class, interictal EEG, and declaring outliers to this class as seizure events. The success of our method relies on detecting change points in the empirical outlier fraction with respect to a feature space that strongly discriminates interictal from ictal EEG. If the feature space is well-chosen, the implication is that novelties are seizures.

In addition to achieving state-of-the-art performance, our technique overcomes three severe limitations of competing algorithms: (1) it does not need to be trained on seizures, (2) it does not require patient-specific tuning, and (3) it does not require knowledge of patient state-of-consciousness. While the false positive performance of the detector is not as good as other reported algorithms, this may be attributable to the presence of subclinical seizures, or other non-ictal anomalies in the data (e.g., normal periodic rhythms, artifacts, etc.). Furthermore, the acceptance by the research community of "hyperdetection strategies"—high false-positive rates and high-sensitivity detection—diminishes the emphasis placed on FPR metrics. For example, in early prototype reactive stimulation devices to treat seizures, the very brief and subthreshold stimulation involved in therapy appears to be well tolerated without any significant side-effects. In this setting, the need to prevent seizures (avoid false negative events), and the apparent relative harmlessness of false positive stimulations, encourage making the detector hypersensitive. As a final note, the entire algorithm is computationally efficient because of the use of the SVM and small novelty threshold.

Ongoing work includes methodological refinements for reducing FPR, and online implementations for validation on very large continuous, multichannel data sets.

Acknowledgements

This work was partly supported by funding from the Esther and Joseph Klingenstein Foundation, Whitaker Foundation, Epilepsy Foundation, American Epilepsy Society, Dana Foundation, Epilepsy Project, and the National Institutes of Health grant #RO1-NS041811-01.

References

K. P. Bennett and C. Campbell, "Support Vector Machines: Hype or Hallelujah?," *SIGKDD Explorations*, 2:1-13, 2000.

B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A Training Algorithm for Optimal Margin Classifiers." In *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, Pittsburgh, PA, USA, 1992.

C. J. Burges, "A Tutorial on Support Vector Machines for Pattern Recognition," *Data Mining and Knowledge Discovery*, 2:1-47, 1998.

M. D'Alessandro, "The Utility of Intracranial EEG Feature and Channel Synergy for Evaluating the Spatial and Temporal Behavior of Seizure Precursors." Ph.D. Dissertation, *Georgia Institute of Technology, Dept. of Electrical and Computer Engineering*. Atlanta, 2001.

F. Desobry and M. Davy, "Support Vector-Based Online Detection of Abrupt Changes." In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-03)*, part IV, pp. 872-875, 2003.

J. Echauz, R. Esteller, T. Tcheng, B. Pless, B. Gibb, E. Kishawi, and B. Litt, "Long-Term Validation of Detection Algorithms Suitable for an Implantable Device," *Epilepsia*, supplement 7, 42:35-36, Dec. 2001.

R. Esteller, "Detection of Seizure Onset in Epileptic Patients from Intracranial EEG Signals," Ph.D. Dissertation, *Georgia Institute of Technology, Dept. of Electrical and Computer Engineering*. Atlanta, 2000.

R. Esteller, J. Echauz, T. Tcheng, B. Litt, and B. Pless, "Line Length: An Efficient Feature for Seizure Onset Detection." In *Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2:1707-1710, 2001.

A. J. Gabor, "Seizure Detection Using a Self-organizing Neural Network: Validation and Comparison with other Detection Strategies," *Electroencephalography and Clinical Neurology*, 107(1):27 – 32, 1998.

A. Gardner, "A Novelty Detection Approach to Seizure Analysis from Intracranial EEG," Ph.D. Dissertation, *Georgia Institute of Technology, Dept. of Electrical and Computer Engineering*. Atlanta, 2004.

J. Gotman, "Automatic Seizure Detection: Improvements and Evaluation," *Electroencephalography and Clinical Neurophysiology*, 76:317-24, 1990.

P. Hayton, L. Tarrasenko, B. Schölkopf, and P. Anuzis, "Support Vector Novelty Detection Applied to Jet Engine Vibration Spectra." In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pp. 946-952. MIT Press, 2001.

E. R. Kandel, J. H. Schwartz, and T. M. Jessel, *Principles of Neural Science*. Prentice-Hall, New Jersey, 1991.

K. Khorasani and W. Weng, "An Adaptive Structure Neural Networks with Application to EEG Automatic Seizure Detection, " *Neural Networks*, 9(7):1223 – 1240, 1996.

S. Liss, "Apparatus for Monitoring and Counteracting Excess Brain Electrical Energy to Prevent Epileptic Seizures and the Like." US patent #3850161, 1973.

B. Litt and J. Echauz, "Prediction of Epileptic Seizures," *The Lancet Neurology*, 1:22-30, 2002.

L. Manevitz and M. Yousef, "One-Class SVMs for Document Classification," *Journal of Machine Learning Research*, 2:139-154, 2001.

K.-R. M. Müller; G. Ratsch, K. Tsuda, K.; B. Schölkopf, "An Introduction to Kernel-based Learning Algorithms," *IEEE Transactions on Neural Networks*, 2:181-201, 2001.

I. Osorio, M. G. Frei, and S. B. Wilkinson, "Real-time Automated Detection and Quantitative Analysis of Seizures and Short-term Prediction of Clinical Onset," *Epilepsia*, 39:615-27, 1998.

I. Osorio, M. G. Frei, J. Giftakis, T. Peters, J. Ingram, M. Turnbull, M. Herzog, M. T. Rise, S. Schaffner, R. A. Wennberg, T. S. Walczak, M. W. Risinger, and C. Ajmone-Marsan, "Performance Reassessment of a Real-time Seizure-detection Algorithm on Long ECoG Series," *Epilepsia*, 43:1522-1535, 2002.

H. Qu and J. Gotman, "A Seizure Warning System for Long-term Epilepsy Monitoring," *Neurology*, 45:2250-2254, 1995.

H. Qu and J. Gotman, "A Patient-specific Algorithm for the Detection of Seizure Onset in Long-term EEG Monitoring: Possible Use as a Warning Device," *IEEE Transactions on Biomedical Engineering*, 44:115-22, 1997.

B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, and R. Williamson, "Estimating the Support of a High-dimensional Distribution," Microsoft Research, Redmond, WA, USA MSR-TR-99-87, 1999.

D. Tax and P. Laskov, "Online SVM Learning: From Classification to Data Description and Back," *Proceedings of the 13th IEEE Workshop on Neural Network for Signal Processing*, pp. 499-508, 2003.

V. N. Vapnik, *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, 1999.

S. S. Viglione, V. A. Ordon, and F. Risch, "A Methodology for Detecting Ongoing Changes in the EEG Prior to Clinical Seizures." In *21st Western Institute on Epilepsy*, 1970.

W. R. Webber, R. P. Lesser, R. T. Richardson, and K. Wilson, "An Approach to Seizure Detection Using an Artificial Neural Network (ANN)," *Electroencephalography and Clinical Neurophysiology*, 98(4):250-2722, 1996.

H. P. Zaveri, W. J. Williams, and J. C. Sackellares, "Energy Based Detection of Seizures." In *15th Annual International Conference on Engineering and Medicine in Biology*, pp. 363-364, 1993.