Automatic Detection of Epileptic Seizure Onset and Termination using Intracranial EEG

by

Alaa Amin Kharbouch

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

Author .....
Department of Electrical Engineering and Computer Science
March 20, 2012

Certified by .....
John V. Guttag
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Chair, Department Committee on Graduate Students

Automatic Detection of Epileptic Seizure Onset and Termination using Intracranial EEG

by

Alaa Amin Kharbouch

Submitted to the Department of Electrical Engineering and Computer Science
on March 20, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

This thesis addresses the problem of real-time epileptic 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 out 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 tailored to individual patients.

In addition, we describe and evaluate an algorithm for the detection of the cessation of seizure activity within IEEG. Seizure end detection algorithms can enable important clinical applications such as the delivery of therapy to ameliorate post-ictal symptoms, the detection of status epilepticus, and the estimation of seizure duration. Our machine-learning-based approach is patient-specific. The algorithm is designed to search for the termination of electrographic seizure activity once a seizure has been discovered by a seizure onset detector. When tested on 65 seizures, 88% of all seizure ends were detected within 15 seconds of the time determined by a clinical expert to represent the electrographic end of a seizure.

We explore the effects of channel pre-selection on seizure onset detection. We evaluate and present the results from a seizure detector that has been restricted to use only a small subset of the channels available. These channels are manually

chosen to be those that show the earliest ictal activity. The results indicate that performance can suffer in many cases when the algorithm uses a small set of selected channels, often in the form of an increase in false alarm rate. This suggests that the inclusion of a full channel set 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 inter-ictal patterns. Finally, we present and evaluate an algorithm for patient-specific feature extraction, where the feature extraction process for a given patient leverages the training data available for that patient. The results from an evaluation of a detector that supplemented the original spectral energy features with features computed in a patient-specific manner show a significant improvement in 3 out of 5 patients. The results suggest that this is a promising avenue for further improvement in the performance of the seizure onset detector.

Thesis Supervisor: John V. Guttag
Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

I would like to thank my advisor, John Guttag, for giving me perspective and guidance when I needed it, and his generosity with his time. I would like to thank Dr. Sydney Cash for the many things he has taught me, and his hands-on approach. I would also like to thank Bill Freeman for his time and effort in his role as a committee member.

I would like to thank my friends, including Zahi, Demba and Paul.

I would like to thank Ali Shoeb for being a mentor and a friend.

I would like to thank my family, especially my parents. Many people think their parents are the best. It happens to be true in my case.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Epilepsy

Epilepsy is a neurological disorder that is characterized by the recurrence of seizures. The clinical symptoms of seizures can include convulsive movements, as well as alterations in behavior, sensation and consciousness. It is estimated that approximately 1% of the general population suffers from epilepsy, including more than 2.5 million people in the United States. For more than 20% of epilepsy patients, neither medications (which may be associated with negative side-effects) nor resective surgery (which requires intensive and invasive investigation), result in the elimination of seizures or constitute a sufficient treatment [32, 14, 21]. One of the most disruptive and disabling aspects of epilepsy is the uncertainty as to when the next seizure will strike.

Electrical activity in various areas of the brain can be measured through the placement of electrodes on the scalp, on the surface of the brain or within its depths. This neurophysiological data is broadly referred to as EEG (intracranial EEG is abbreviated as IEEG), and is often recorded for diagnostic purposes. The beginning of the measured electrical activity associated with the seizure is referred to as the *electrographic onset*. When the seizure is associated with clinical symptoms, the *clinical onset* is the point in time when clinical symptoms are first observed. The end of the electrographic activity associated with a seizure marks the seizure end, or seizure termination. While most seizures are self-terminating, and usually last less than 5 minutes, the serious state of persistent (and seemingly indefinite) seizure activity, known as *status epilepticus*, is associated with a significant risk of mortality. The

seizure itself is referred to as the *ictal* period (and seizure-related electrical activity as ictal activity), the period that immediately follows the end of a seizure is known as the post-ictal period. The post-ictal period itself is sometimes associated with clinical symptoms.

1.2 Seizure Onset Detection

Epilepsy is a disease characterized by recurrent episodes of dysfunctional brain activity. Yet, current approved 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 algorithm that can detect seizure activity early and accurately. In this thesis, we describe the architecture and performance of a real-time intracranial EEG (IEEG) seizure onset detector. Throughout this thesis, 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.

Two important considerations when assessing the performance of an automatic seizure onset detector are the delay between the seizure onset and the detection time, referred to as the detection *latency*, and the frequency of the occurrence of false detections, referred to as the *false alarm rate*. The reduction of latency and the lowering of false alarm rates are competing objectives, and the extent to which one is favored over the other is dictated by the application. Consider the example of a device that can detect a seizure's onset and deliver an electrical stimulation in response, which is designed to reduce the length or intensity of seizures. Low latencies are crucial for triggered neurostimulation applications, since the effectiveness of the therapy is delay-sensitive. In the case of an alerting application in a hospital setting, a low false alarm rate decreases the likelihood that alarms raised during seizures will be ignored by caregivers (or have the response to them delayed). If automated seizure detection is used to initiate the delivery of a drug (e.g. an anticonvulsant), then in most cases a very low false alarm rate would be crucial.

Rapid and reliable seizure onset detection from IEEG is challenging for a number of reasons. First, IEEG varies greatly across individuals with epilepsy [7]. The intracranial EEG associated with seizure onset in one patient can closely resemble a benign pattern within the IEEG of another patient. Furthermore, for individual patients there exists significant overlap in the IEEG associated with seizure and non-seizure states. In addition to inter-patient variability, there is also intra-patient 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 epilepsy patients transitions between regimes within both the seizure and non-seizure states, and is therefore a non-stationary process.

Previous work has introduced a wide variety of techniques for seizure detection using IEEG. Results from these studies show that many of the techniques struggle with high latencies or high false alarm rates or both. Furthermore, some aspects of these studies remain to be addressed. First, most IEEG detection studies use relatively short records; past studies have used no more than 30 hours per patient, and considerably less than 24 hours for many patients [23, 17, 12, 1]. 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 intracranial EEG seizure detection has focused on datasets with a small number of pre-selected electrodes[23, 1, 24, 28]. This may ignore useful information that is not obvious on visual examination, and in some cases necessitates additional assessment and culling of the dataset by an expert.

We treat seizure detection as a binary classification problem that involves separating seizure activity from non-seizure activity, and employ a machine-learning approach that is patient-specific, i.e., the classifier for each patient is trained using seizure and non-seizure examples (labeled as such) which have been extracted from the same patient's IEEG data. The division of a record of the brain's electrical activity by an expert into two encompassing classes, seizure and non-seizure, is consistent with standard clinical practice. We extract from each epoch features that correspond

to the log-energy in several frequency bands for every channel. A Support Vector Machine is trained using these feature vectors, and this constitutes the classifier which is tested against unseen continuous data from the same patient. This is similar to the approach to scalp EEG seizure detection of Shoeb *et al* [30, 29]. Our evaluation uses more than 875 hours of intracranial EEG in total, averaging more than 87 hours of continuous IEEG data per patient. We also present results from a study that 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 intracranial EEG across many channels and comparing time periods between seizures with seizure activity itself. Relative to previously published methods, ours exhibits high sensitivity, short latencies, and low false alarm rates.

We also explore the effects of channel pre-selection on seizure onset detection. We evaluate and present the results from a seizure detector that has been restricted to use only a small subset of the channels available. These channels are manually chosen to be those that show the earliest ictal activity. In this context, we also present and evaluate an algorithm for patient-specific feature extraction. With this approach, the process for computing features from each channel for a given patient leverages the training data available for that patient. This approach is designed to yield additional salient features that can be combined with the original feature set to lead to improved detection performance. The feature extraction process involves the representation of an epoch of IEEG on a channel as a linear combination of prototype components, which we refer to as dictionary elements. To learn a dictionary, the method generally seeks signal components that have a large presence in one class (e.g., seizure onsets), but are largely absent from another class. The discriminative value of a direction is quantified as the difference in the total power in that direction between the two classes. We compare the results using this approach for feature extraction against those obtained for the seizure detection with channel pre-selection using the generic spatio-spectral features only.

1.3 Seizure End Detection

Much effort has been dedicated to the detection of seizure events or seizure onsets. In contrast, little effort has been devoted to developing algorithms that can detect the termination of a seizure, even though such algorithms can enable important applications. They can facilitate the estimation of seizure duration, which could help physicians assess the efficacy of anti-epileptic drugs when combined with estimates of seizure frequency. A seizure end detector could also be used to control the delivery of therapies to control postictal symptoms, which can persist for anywhere from minutes to days. Finally, a seizure end detector can be used to detect the presence of status epilepticus. This can be accomplished, for example, by activating an alert when a seizure's end has not been detected within a certain period of time (e.g. 5 minutes) from its onset.

1.4 Intracranial EEG

Several types of neurophysiological data reflecting electrical activity are in use for diagnostics and research. Perhaps the most familiar is the scalp EEG (electroencephalogram), which is obtained by recording electrical activity measured by electrodes placed on the scalp. However, in this thesis we will focus on a dataset containing another form of EEG called intracranial EEG (IEEG), where the electrodes are placed inside the skull, and is thus categorized as an invasive form of EEG. The electrodes can be placed on the outer surface of the brain or cortex (these are known as grid or strip electrodes), or within brain matter (these are referred to as depth electrodes). In comparison to scalp EEG, IEEG signals reflect the activity of a smaller number of neurons, and therefore have a higher spatial resolution. IEEG recordings also have a higher bandwidth, and are thus sampled at higher rates. In addition, IEEG signals often show the signs of the start of seizure activity several seconds before scalp EEG recorded from the same patient [25]. Figures 1-1, 1-2, 1-3 show examples of IEEG from inter-ictal(non-seizure), ictal and post-ictal periods respectively.

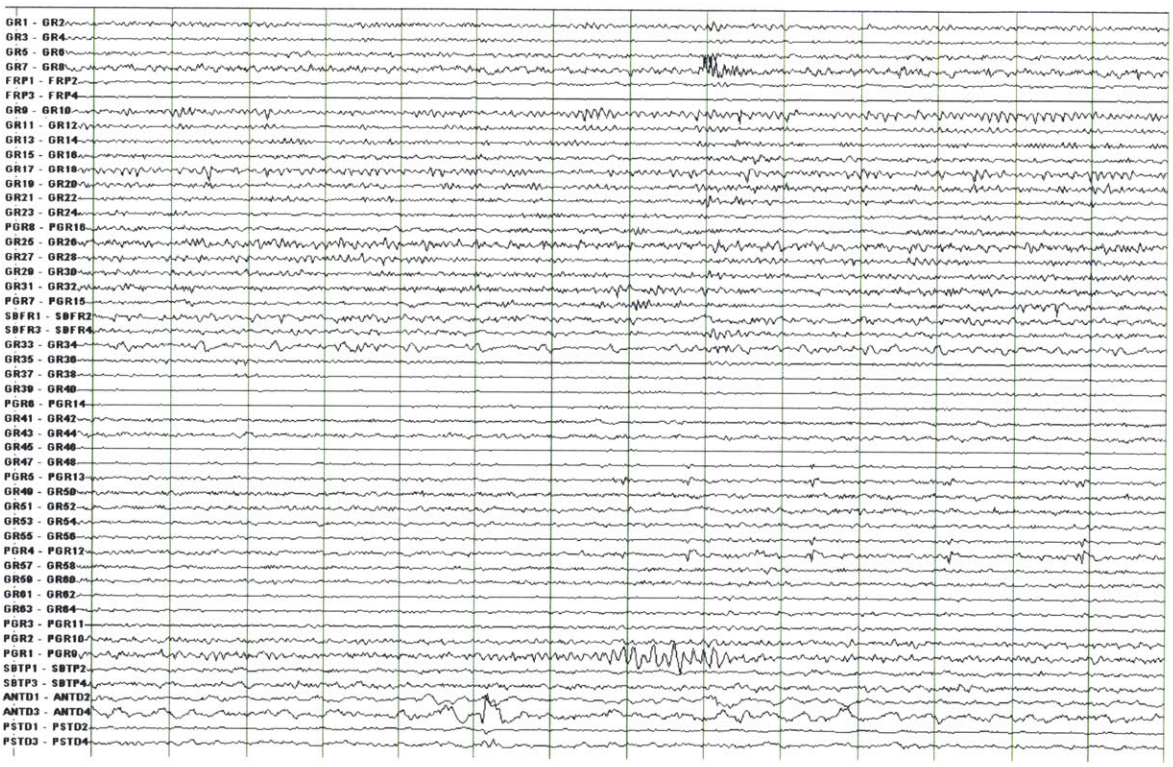

Figure 1-1. A 14-second window of intracranial EEG during an inter-ictal (i.e. non-seizure) period.

1.5 Previous work

This project builds on the work of Shoeb et al. [30, 29]. Several other methods have been published on seizure detection on noninvasive scalp EEG, including some which also incorporate a support vector machine as a component [11, 19, 33]. Our algorithm is designed to work with a different type of neurophysiological data (IEEG), and there are some differences in the methodology, which is described in Chapter 2.

The median patient false alarm rate (0.6/day) and the median/mean latency across all seizures (5 sec/6.9 sec) obtained using our algorithms are lower than those reported in previous work on seizure detection using IEEG, and the sensitivity was comparable (97%). It should be noted that the various studies mentioned in this section tested methods on different IEEG databases. There can also be some differences in evaluation methodology, and other aspects such as patient selection criteria.

Algorithms such as those described by Chan et al. [4] and Chua et al. [6] were designed for offline IEEG analysis, and cannot be implemented as part of a real-time warning system. Zhang et al. [35] report a sensitivity of 98.8%, a mean latency of 10.8 seconds, and a combined false alarm rate of 11.8/day for their patient-specific method when tested on IEEG from 21 patients. The authors report separate figures for false alarms dichotomized into "interesting" and "uninteresting" groups. Grewal and Gotman [12] 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. Aarabi et al. [1] report a sensitivity of 98.7%, an average latency of 11 seconds, and a false alarm rate of 6.5 per 24 hours for their method when tested on data from 6 contacts per patient in the Freiburg public database [18].

Gardner et al. [10] used a one-class support vector machine requiring only non-seizure data for training, and reported a false alarm rate of 37.4/day for their method when tested on 200 hours of data from 5 patients. That study allowed for negative latencies in cases where 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. [28] evaluated their two-channel patient-specific IEEG 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. [24] report impressive results. However, important differences distinguish their study from ours, including the automatic pre-selection 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 IEEG record had not been reviewed in its entirety to locate all instances of seizures. Thus, no clear comparison can be made.

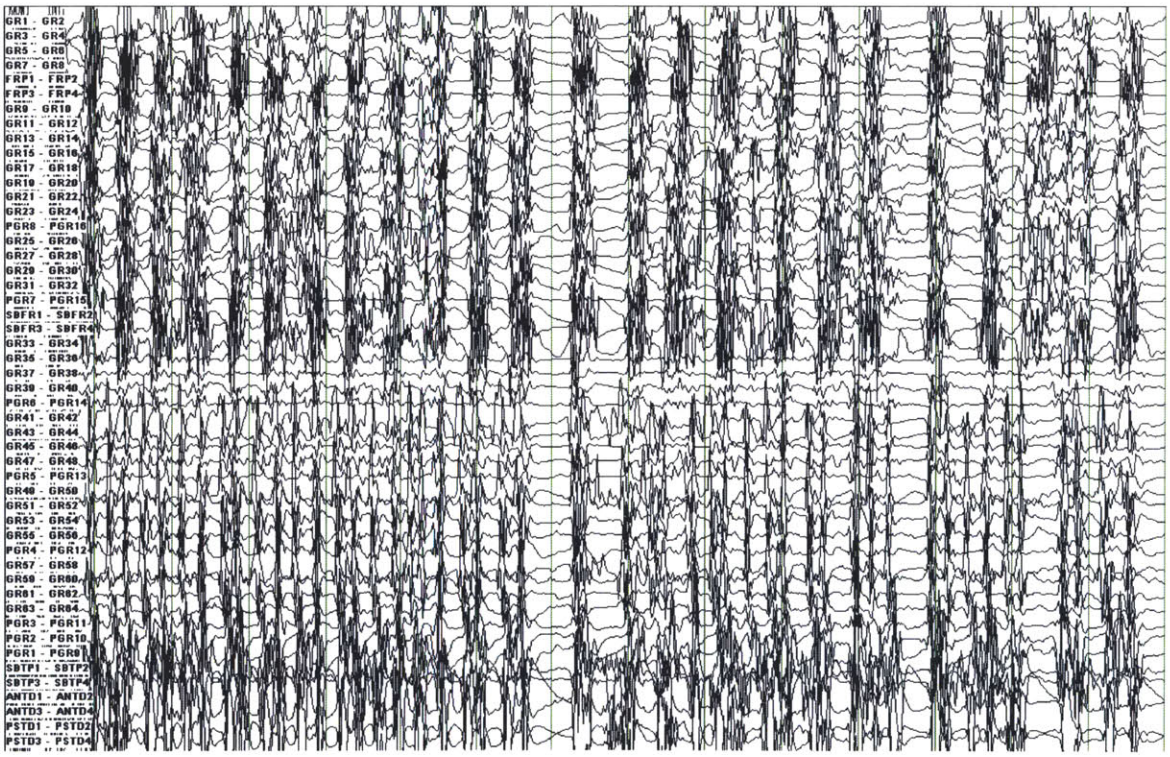

Figure 1-2. A 14-second window of intracranial EEG during a seizure

1.6 Contributions

Automated seizure onset detection enables multiple alerting and therapeutic applications, but existing approaches have been plagued with difficulties in either sensitivity, specificity, latency or some combination of these. In this thesis 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. The proposed algorithm performance compares favorably to existing methods, as determined by well-defined performance metrics. The median patient false alarm rate (0.6/day) and the median/mean latency across all seizures (5 sec/6.9 sec) obtained using our algorithm are lower than those reported in previous work on seizure detection using IEEG.

We analyzed more than 875 hours of continuous intracranial EEG recorded from 10 patients to evaluate our algorithm. The more than 87 hours of data per patient significantly exceeds the amount of data used in most previous studies. This allows for more realistic estimates of long-term performance. We provide some evidence for this with the finding that false alarms tend to temporally cluster, and are unevenly distributed in time (including across days).

Furthermore, the algorithm does not require the manual pre-selection of a small subset of channels by a clinician, as is the case with several previously published methods and studies. We present results from a study that uses the full set of intracranial electrodes of sufficient recording quality.

In addition, we describe and evaluate an algorithm for the detection of the cessation of seizure activity within IEEG. Seizure end detection algorithms can enable important clinical applications. Our machine-learning-based approach is patient-specific. When tested on 65 seizures, 88% of all seizure ends were detected within 15 seconds of the time determined by a clinical expert to represent the electrographic end of a seizure. For 5 out of 10 patients, 100% of seizure ends were detected within a 15-second margin of the expert-marked end.

We explore the effects of manual channel pre-selection on seizure onset detection for some patient datasets. We evaluate and present the results from a seizure detector

that has been restricted to use only a small subset of the channels available.

Finally, we propose an alternative feature extraction method that is tailored to a patient using the training data, and is designed to work with datasets limited to a few pre-selected channels. The feature extraction process involves the representation of an epoch of IEEG on a channel as a linear combination of vectors that we refer to as dictionary elements. To learn a dictionary, the method searches for directions in which the power in the two classes shows a large difference. The results from an evaluation of a detector that supplemented the original spectral energy features with features computed in a patient-specific manner show a significant improvement in 3 out of 5 patients. While not conclusive, the results suggest that this avenue is worthy of further exploration.

1.7 Thesis Outline

In Chapter 2 we describe and evaluate an algorithm for automatic seizure onset detection from intracranial EEG. We define some performance metrics, and evaluate the algorithm using a database containing continuous IEEG from ten patients. In Chapter 3 we describe and evaluate an algorithm for seizure end detection using IEEG. The algorithm is designed to search for the cessation of electrographic seizure activity once a seizure has been discovered by a seizure onset detector. We present the results from an evaluation of the algorithm using seizures from the same database used to test the seizure onset detector. In Chapter 4, we present the results when the seizure detector is restricted to the use of a small number of channels (2-4). The channels are manually selected based on visual examination for each patient, and correspond to the channels that show the earliest ictal activity. We also present and evaluate an algorithm for patient-specific feature extraction, an approach where the feature extraction process for a given patient leverages the training data available. A summary and some discussion and conclusions are included in Chapter 5.

Figure 1-3. A 14-second window of intracranial EEG during a post-ictal period, beginning less than 90 seconds after the expert-marked electrographic end of a seizure.

Chapter 2

Algorithm for Seizure Onset Detection

In this chapter we describe and evaluate an algorithm for automatic seizure onset detection from intracranial EEG. Our machine-learning approach is patient-specific. We define some performance metrics, and evaluate the algorithm using a database containing continuous IEEG from ten patients.

2.1 Feature Vector Extraction and Classification

We treat seizure detection as a binary classification problem that involves separating seizure activity from non-seizure 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. 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 [27, 30], with some essential changes also described in this section.

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 =+/-1 depending on whether X is representative of seizure or non-seizure IEEG. The function is derived using training sets of seizure and non-seizure 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 IEEG 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 [27]. 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 IEEG 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 non-seizure 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 filterbank 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. The beginnings of consecutive epochs are separated by 1 second. The filterbank is composed of M filters, and is illustrated in Figure 2-1.

The scalp EEG seizure detector in [27, 30] focuses on the frequency range 0.5-25Hz. However, the intracranial EEG is a signal of a higher bandwidth and carries relevant information at higher frequencies. Therefore, the IEEG detector filter bank considers a wider range of frequencies, although more emphasis is still placed on the lower frequency range. The 0.5-35Hz range is covered by more filters (12) each with a bandwidth of 3Hz, whereas the 35-105Hz range is covered by a lower density of filters (5 filters) each with a bandwidth of 15Hz, for a total of $M = 17$ filters. The 60 Hz

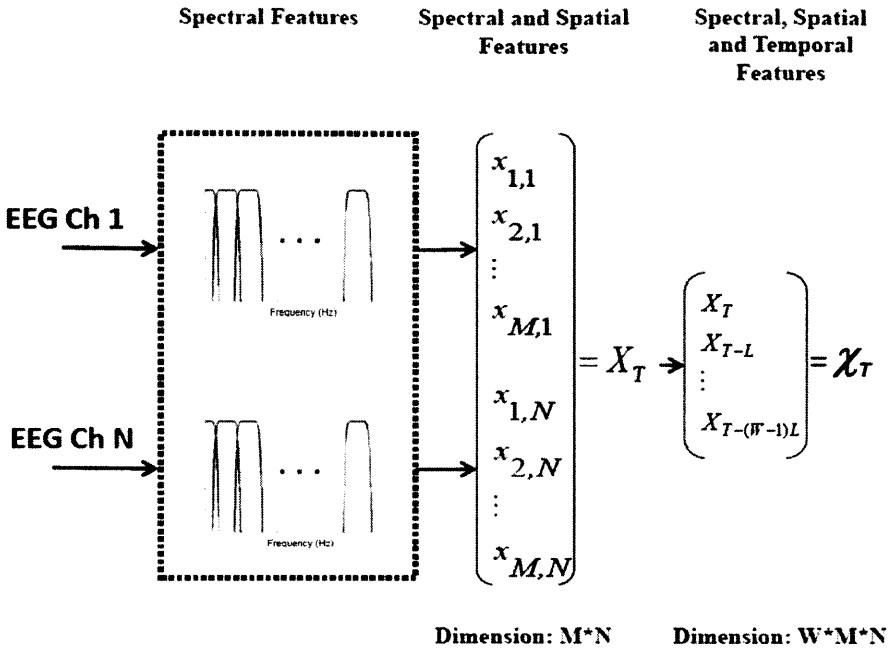

Figure 2-1. Feature vector formation steps.

region is neglected so as to remove the 60Hz electrical noise often contaminating the signals. A slightly different choice of M or the filter bandwidth may lead to similar (and possibly improved) results.

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 one-second EEG epoch at time t=T, we concatenate the $M = 17$ spectral energies extracted from each of N IEEG channels. This process forms a feature vector X_T with $M * N$ elements as shown in the middle portion of Figure 2-1. 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 mean-centered and then divided by the median absolute value.

The feature vector X_T does not capture how an epoch relates to those in the recent past. Consequently, 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 X_T, $X_{T-L}...X_{T-(W-1)L}$ from W contiguous, but non-overlapping $L = 1$ second epochs as shown on the right side of Figure 2-1. 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 X_T using a longer epoch length L. We set $W = 3$.

2.1.1 Training and Classification

The feature vector χ_T is classified as representative of seizure or non-seizure 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 non-seizure vectors computed from non-seizure IEEG. This results in a number of non-seizure training feature vectors that greatly outnumbers the number of seizure examples. In the training phase, the non-seizure 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 post-ictal period, which tends to be associated with electrographic qualities that distinguish it from the rest of the non-seizure activity. Within the SVMlight software package [16], the error cost parameter was set to $C = 1 \times 10^{-3}$ for all patients in the evaluation database. This parameter was not increased for the seizure class as that did not appear to lead to an overall improvement in results in preliminary analysis.

2.1.2 Artifact Rejection

An artifact rejection component works in conjunction with a trained classifier by checking for large differences between the minimum and maximum value of the signal for each channel within $W * L = 3$ 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 two minutes.

2.2 Data

The data used to evaluate our detector consists of more than 875 hours of continuous intracranial EEG sampled at 500 Hz. The data was recorded at Massachusetts General Hospital from 10 patients with focal epilepsy (5 female, mean age at onset of seizures 15 +/- 5 yrs S.D., mean age at surgery of 40 +/- 9 yrs S.D.). Etiologies included mesial temporal sclerosis (2), cortical dysplasia (2), post-traumatic epilepsy (1), and post-infectious epilepsy (1). In 4 patients the etiology was unknown.

The patients were all surgical candidates who required invasive monitoring, and therefore represent more complicated cases 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 paper was 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 [22]. For the hospitalized patients from whom the data for this study was recorded, the anti-epileptic drug levels were changed on a daily basis. The medications have a significant effect on both seizure and non-seizure IEEG. The lack of consistency in the magnitude or nature of these effects at different times complicates the evaluation of the detector.

While the recordings were being made, the patients experienced a total of 67 seizures (between 3 and 12 seizures for each 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 use 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 this data was performed retrospectively under the auspices of the local institutional review board in accordance with the Declaration of Helsinki.

2.3 Evaluation Methodology

Each patient's data was segmented into a number of records, where a record corresponds to a stretch of time in the IEEG recordings. Each record is up to 24 hours long, and contains at least one seizure (in a very small number of cases where two seizures are less than 15 minutes apart, they are included in the same record). If there is a separation of more than 24 hours between the seizure in the record and the preceding one, the beginning of the record is set to approximately 24 hours before the seizure, and its end is set to be approximately 20 minutes after (i.e., every seizure's post-ictal phase is included within the same record). If that distance 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 non-recording of electrodes were excluded from patient datasets. 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. The seizure detection algorithm is constrained to be causal, i.e., the decision at a given point in time as to whether a seizure is underway can only be made using past IEEG data leading up to that point. This is to ensure the algorithm can be compatible with a real-time application. In contrast, no causality constraint applies

as a human expert determines the seizure onset point that we use as a reference for evaluation. In fact, a clinician will often track backwards from more clearly ictal later-stage activity to determine a time point where the first change in activity leading up to it occurred. 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 anytime 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 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 testing scheme. We avoided an evaluation method based on leaving out epochs (as done in [20]) rather than records, since 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_R 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 $N_R - 1$ 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_R times so that each record is tested. In most cases $N_R - 1$ training seizures are used for each test record since most records contain a single seizure.

2.4 Results

Overall, 97% of the 67 test seizures were detected. As shown in Table 2.1, the sensitivity was 100% for most patients. The median latency with which the detector

declared the seizure onset (across all seizures) was 5 seconds. Some care is needed in the interpretation of some results (e.g., aggregates over all seizures in the database with an inconsistent number of seizures per patient) because of inter-patient differences or variations, which are discussed in later sections. The median latencies for each patient are shown in Table 2.1. The average false alarm rates for each patient are also shown in Table 2.1. The median false alarm rate was 0.6 false detections per 24 hour period. Figure 2-2 shows an example of one of many seizures that were detected with a short latency. The expert-marked onset is indicated by the dark vertical line, and the approximate time at which the algorithm detected the seizure is indicated by the red dotted line.

Patient	1	2	3	4	5
Total time tested(Hr)	52	127	71	29	96
Number of seizures	3	6	6	5	8
# of electrodes	98	56	114	34	68
Electrode types	Grid & Depth	Depth	Grid & Depth	Depth	Depth
Sensitivity	100%	100%	100%	100%	100%
Median Latency(sec)	6.5	3.25	3.25	5	5.75
False alarms/ 24Hr	0	0.4	1.7	0	0.8

Patient	6	7	8	9	10
Total time tested(Hr)	68	39	146	148	104
Number of seizures	3	7	12	11	6
# of electrodes	40	54	64	80	85
Electrode types	Depth	Depth	Depth	Depth	Grid & Depth
Sensitivity	66%	100%	92%	100%	100%
Median Latency(sec)	4.25	4.5	6	4.5	18.5
False alarms/ 24Hr	0	25.3	2.5	0.3	2.5

Table 2.1. 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.

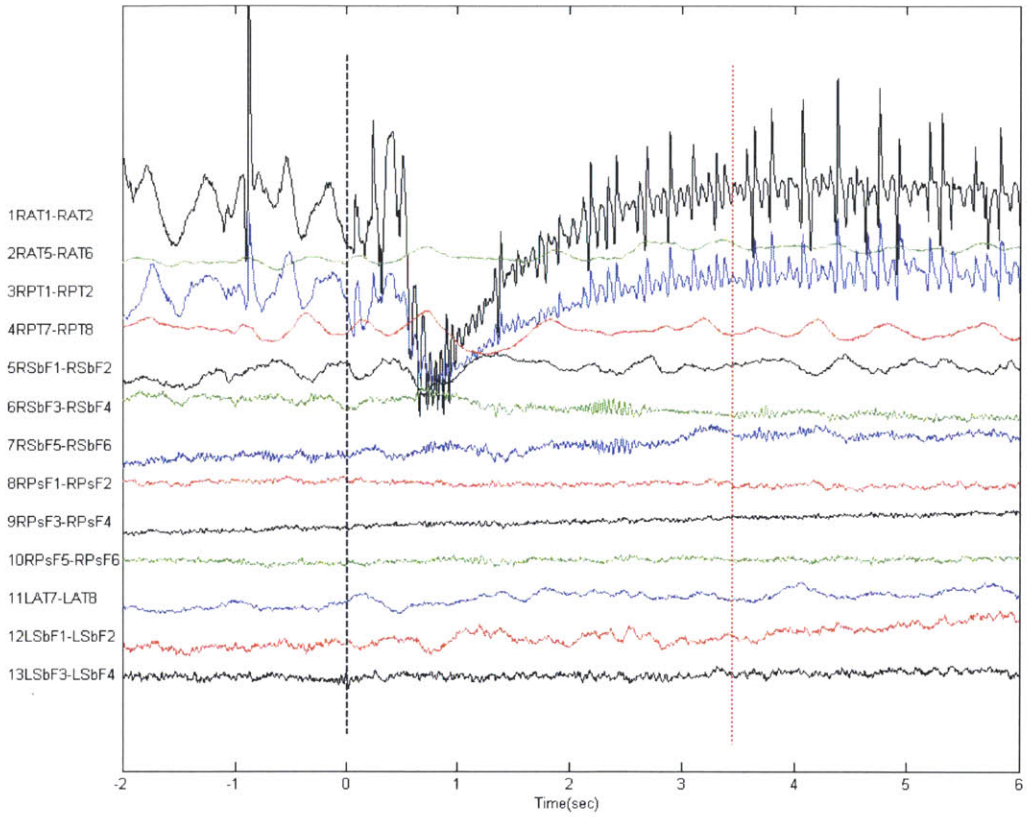

Figure 2-2. Intracranial EEG on a subset of channels in a 8-second epoch containing the onset of a Patient 5 seizure. The expert-marked onset is indicated by the dark vertical line, and the approximate time at which the algorithm detected the seizure is indicated by the red dotted line.

2.4.1 Sensitivity and Latency

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 dataset. Our detector misses or has a large latency when a test seizure differs greatly from all the training seizures. One example of this is the seizure from the Patient 6 dataset that was missed by the algorithm, shown in Figure 2-3. The two remaining seizures from this patient were grouped together, (by a clinician in the patient report) and one of them is shown in Figure 2-4. Clear differences between the seizure types include high amplitude rhythmic spiking activity on the first anterior temporal channel in the seizure shown in Figure 2-4 that is absent from the same channel in Figure 2-3. Rhythmic slowing in several anterior and posterior temporal contacts early in the onset of the seizure in Figure 2-4 also does not match the activity in the seizure in Figure 2-3.

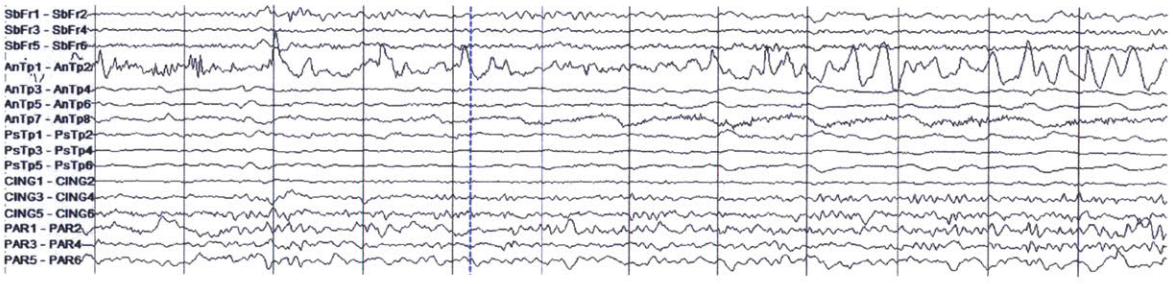

Figure 2-3. Intracranial EEG on a subset of channels in a 12-second epoch containing the onset of a Patient 6 seizure. The dashed vertical line indicates the expert-marked onset.

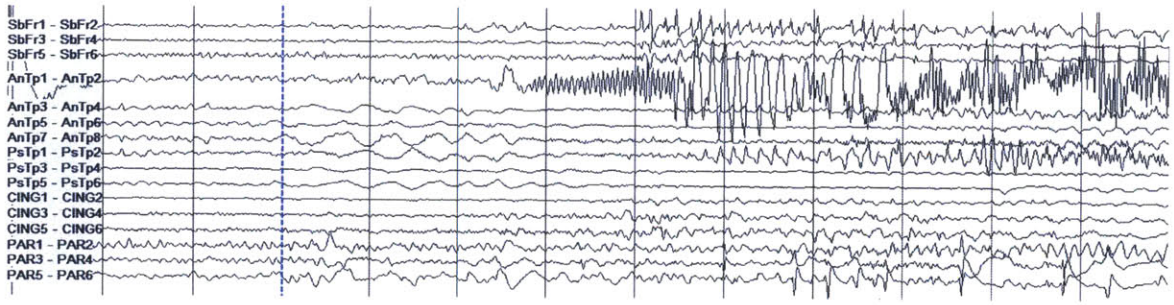

Figure 2-4. Intracranial EEG on a subset of channels in a 12-second epoch containing the onset of a Patient 6 seizure different from that of Figure 2-3. The dashed vertical line indicates the expert-marked onset.

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 were different from one another. For each 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 or lack of apparent clinical symptoms). This is possibly explained by the fact that the patient's dataset includes 12 seizures, and therefore, for most test seizures within this patient's dataset, the classifier has incorporated more examples of seizures of the same broad type, and learned some variation in IEEG patterns among them. This allows for a better performance, particularly in reduced latency in the case of this patient dataset.

2.4.2 False Alarms

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 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. Figures 2-5 and 2-6 show examples of activity that induced an alarm, and the seizure activity they resemble. The high amplitude rhythmic activity on the right subfrontal channel during the seizure in 2-5(A) also appears in the inter-ictal epoch where false detection occurred in testing, shown in 2-5(B). The type of activity highlighted in 2-5(B) lasted for approximately 8 seconds. Figure 2-6(A) shows the high-frequency activity on the right posterior temporal channel in a different seizure, and Figure 2-6(B) shows the similar activity on the same channels from an inter-ictal epoch where a false alarm was raised. Other false alarms appear to have been caused by non-physiological artifact that was not detected by the artifact rejection module.

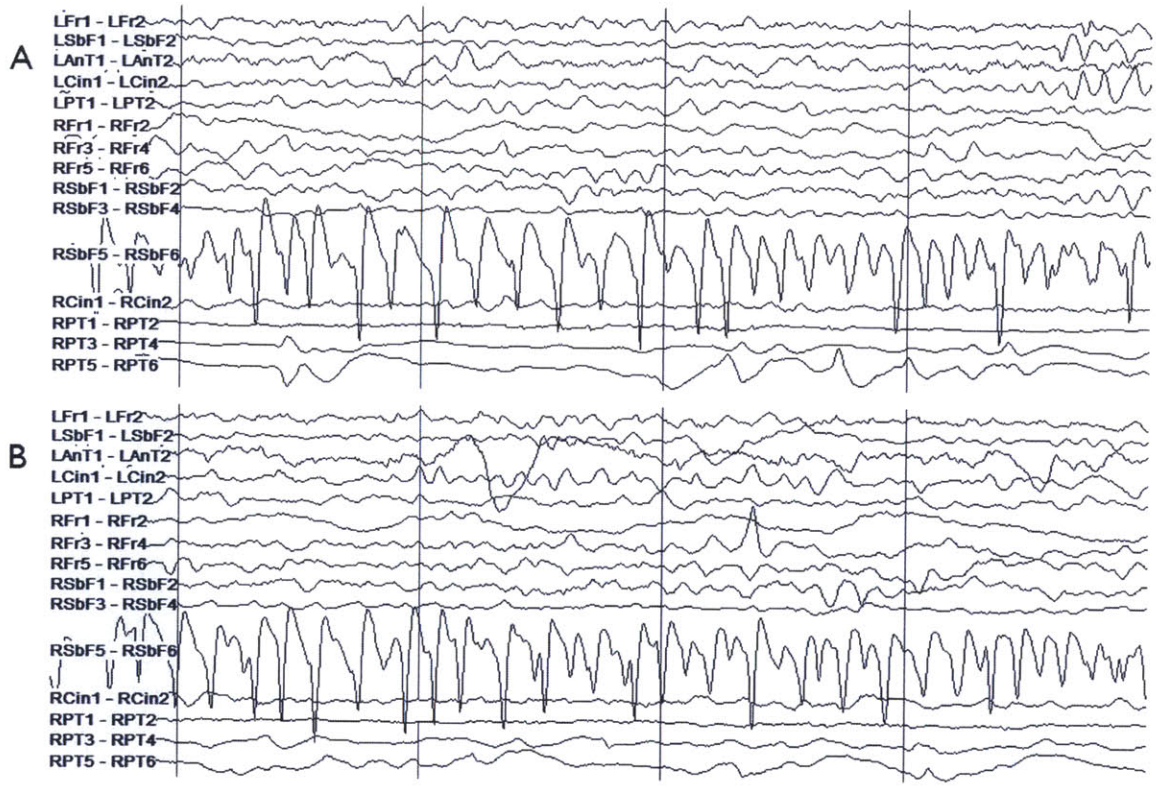

Figure 2-5. Patient 7 intracranial EEG on a subset of electrodes in two epochs showing similar activity, particularly on the right subfrontal channel 'RSbF 5 - RSbF 6'. (A) Seizure activity in a 4-second epoch beginning approximately 10 seconds after the expert-marked onset. (B) An interictal 4-second epoch during which false detection occurred. This activity was not judged to constitute a seizure by the clinicians.

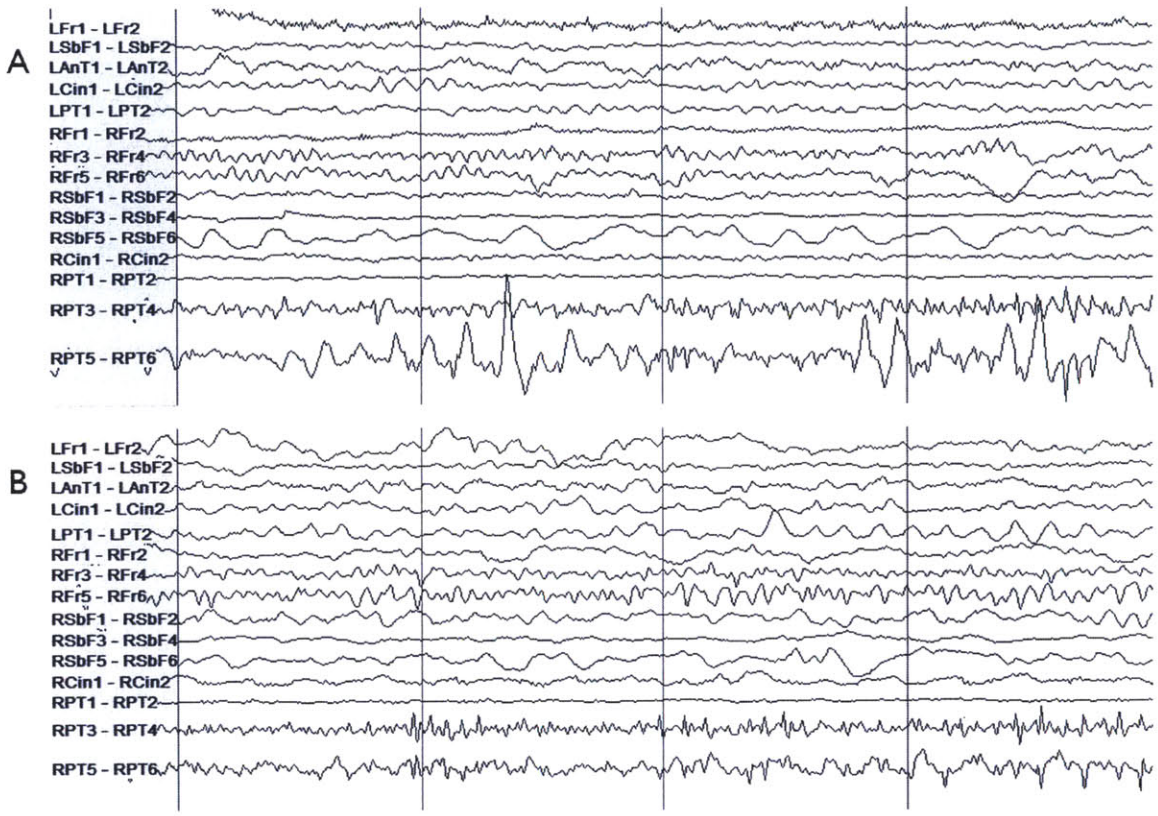

Figure 2-6. Patient 7 intracranial EEG on a subset of electrodes in two epochs showing similar activity, particularly on the right posterior temporal channels 'RPT 3- RPT 4' and 'RPT 5- RPT 6'. (A) Seizure activity in a 4-second epoch beginning approximately 7 seconds after the expert-marked onset. (B) An interictal 4-second epoch during which false detection occurred. As before, this activity was not judged to constitute ictal activity by clinical reviewers.

In many cases, the false alarm events are unevenly concentrated in different regions in time, i.e., they temporally cluster. In the case of Patient 3, 60% of false alarms (3 out of 5) 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. Figure 2-7 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 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.

2.4.3 Reduction of K parameter

We explored the possibility of reducing the latencies by adjusting the K parameter. The setting of $K = 2$ forces the algorithm to wait for 2 consecutive windows that the classifier has deemed to be part of a seizure onset before an alarm is raised. Setting $K = 1$ 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 an increase in the false alarm rate. The results obtained by setting $K = 1$ are shown in Table 2.2. 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/24hrs. 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.

Patient	1	2	3	4	5	6	7	8	9	10
Sensitivity	100%	100%	100%	100%	100%	66%	100%	92%	100%	100%
Median Latency(sec)	5.5	2.25	2.25	4	4.75	3.25	3.5	4	2.5	16.5
False alarms /24Hr	0	1.3	3.7	1.7	5	1.4	45.1	13.5	1	6.2

Table 2.2. Sensitivities, median latencies and false alarm rates for the particular case where K=1. The latencies decrease at the expense of higher false alarm rate.

2.5 Discussion

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

2.5.1 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 non-homogenously 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 was recorded, the medications levels were changed on a daily basis. This may have caused changes in the IEEG and unfamiliar electrical patterns that led to

false alarms (or diminished performance in other forms). 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.

2.5.2 Impact of Number of Electrodes

As opposed to several other investigations, the approach utilized 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 non-specific information that may increase the likelihood of incorrect classification. To better understand these tradeoffs 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 versus implantable devices, the number of channels in the latter is likely to increase.

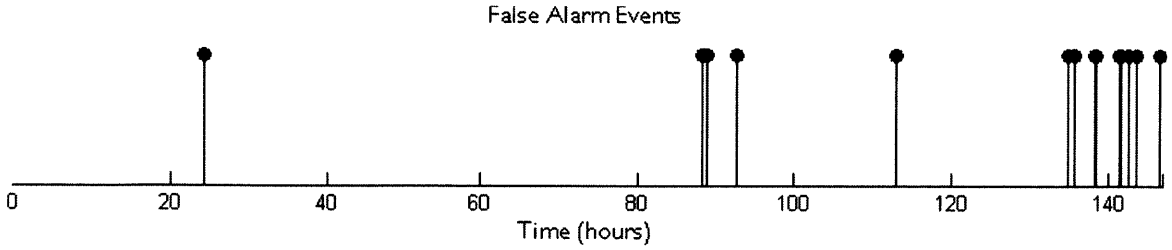

Figure 2-7. Approximate timing of false alarms that occurred in tests on Patient 8 data.

Chapter 3

Seizure End Detection in Intracranial EEG

In this chapter we describe and evaluate an algorithm for seizure end detection using IEEG. Whereas seizure onset detection was approached as a binary classification problem into inter-ictal(non-seizure) and seizure onset classes, the seizure end detector is tasked with discriminating between ictal and post-ictal activity. The algorithm searches for the cessation of electrographic seizure activity once a seizure has been discovered by a seizure onset detector. This approach was preferred over using a single detector to mark both the beginning and end of electrographic seizure activity for several reasons.

The characteristics of IEEG onset activity can differ greatly from those in the latter stages of a seizure. For example, seizures with focal activity (i.e., activity limited to a small number of nearby channels) during their onset can exhibit generalized activity towards their end. Other characteristics such as the fundamental frequency ictal activity can differ between the earliest and latest stages of a seizure. Finally, the post-ictal period tends to be associated with electrographic qualities (e.g. slowing and amplitude attenuation) that distinguish it from the rest of the inter-ictal activity.

As with seizure onset, the IEEG immediately following seizure end varies across patients and seizure types. Seizure end detection is further complicated in many cases by the gradual nature with which the IEEG changes at the end of seizures. When ictal-to-postictal IEEG transitions are gradual, electroencephalographers as well as algorithms may be challenged to identify the time of seizure termination.

Our machine-learning-based approach is patient-specific. We employ an approach where the seizure end detector is designed to discriminate between ictal and post-ictal activity, and is paired with a seizure onset detector that triggers it.

3.1 Seizure End Detection Algorithm and Evaluation

The seizure end detector performs a sliding window analysis of the IEEG following the beginning of a seizure. The analysis window used by the seizure end detector is 4 seconds long, and is advanced forward by 1-second increments. The detector extracts salient features from each analysis window, and then assembles those features into a feature vector. Next, the detector uses a classifier to determine whether a feature vector is representative of the ictal or post-ictal state. To prevent early seizure end declarations because due to the detector's failure to recognize the onset, at least three successive windows must be classified to belong to the seizure class before a seizure end can be declared. If $K_E = 5$ successive feature vectors are classified as belonging to the post-ictal state, then the detector declares that the seizure has ended.

Since ictal and post-ictal IEEG activity characteristically differ both in their spatial distribution and spectral structure, it is important that our feature vector capture these signal properties. The feature extraction process is very similar to those used for seizure onset detection described in Chapter 2. Each spectral feature in the feature vector X_T represents the logarithm of the total energy in a specific frequency band on a single channel. One difference in the case of seizure end detection is that we do not stack several of these vectors to form time-delay embedded feature vector. This is because later-stage seizure activity characteristics do not tend to change on a small timescale, or transition between different types of activity in a short period as is often the case with seizure onset activity.

The feature vector is classified as representative of ictal or post-ictal activity using a linear support vector machine (SVM). The SVM algorithm uses training ictal and post-ictal feature vectors to learn a decision boundary that separates these two classes of activity. Training ictal and post-ictal vectors are derived from multiple seizures of

a single patient. Once a decision boundary is learned, the SVM algorithm determines the class membership of a newly observed feature vectors based on which side of the boundary the vector falls. Within the SVMlight software package [16], the error cost parameter was set to C=1/50.

3.1.1 Data and Evaluation

We evaluated our methodology using the seizure data and associated post-ictal data in the database described in Chapter 2. This database contains 65 seizures recorded from 10 epilepsy patients admitted to Massachusetts General Hospital. For each seizure in the database, an expert electroencephalographer determined seizure onset and end times by examining the IEEG without knowledge of the determinations made by the algorithm. Table 3.1 contains some relevant information for all patients, such as the number of seizures recorded and the electrode types. Two seizures were omitted from the dataset because many channels were obscured by high-amplitude artifact for the majority of their duration. The IEEG during one of the seizures is shown in Figure 3-1.

Patient	1	2	3	4	5
Number of seizures	3	5	6	5	8
# of electrodes	98	56	114	34	68
Electrode types	Grid & Depth	Depth	Grid & Depth	Depth	Depth

Patient	6	7	8	9	10
Number of seizures	3	7	11	11	6
# of electrodes	40	54	64	80	85
Electrode types	Depth	Depth	Depth	Depth	Grid & Depth

Table 3.1. Patient data set information for the database used to evaluate the seizure end detector.

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 of seizure end and expert-marked seizure termination within the IEEG.

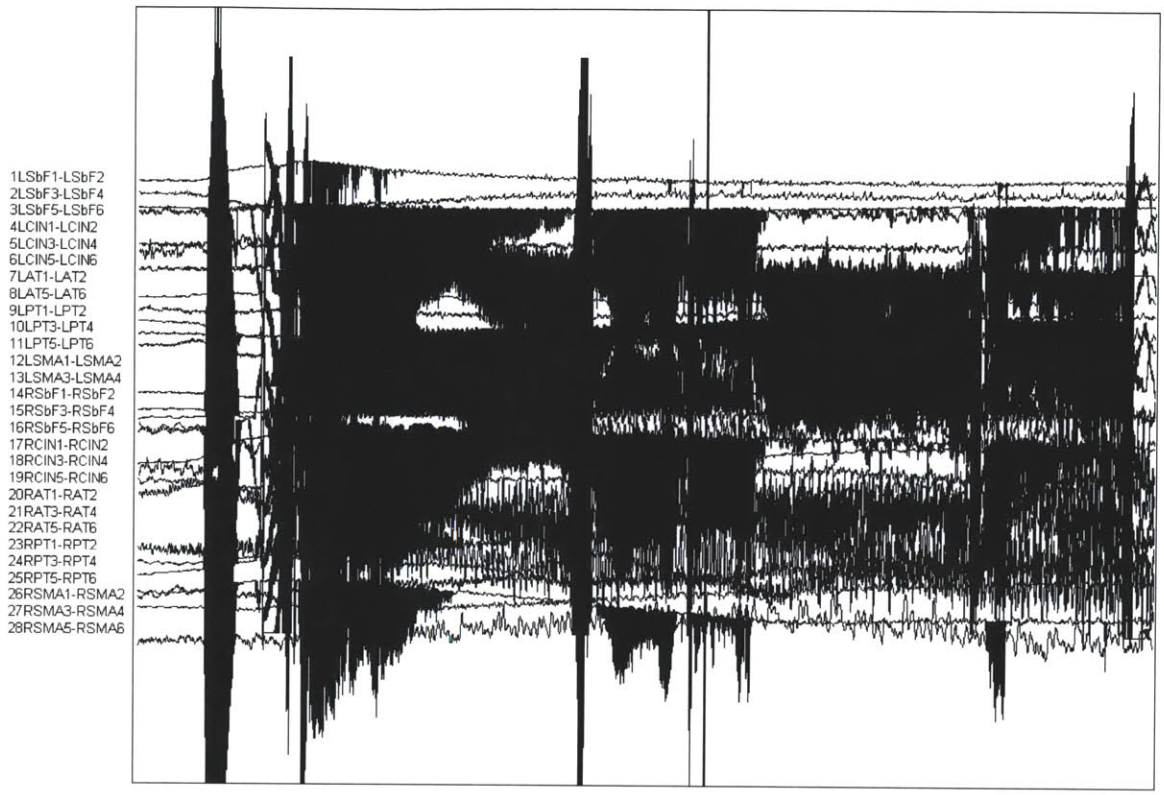

Figure 3-1. A 30-second window of intracranial EEG during a Patient 2 seizure omitted from the end detection study. Many channels are artifact-obscured for much of its duration.

The algorithm declares seizure end at the end of the last of $K_E = 5$ analysis windows that have been classified as representative of post-ictal activity. When we wish to ignore the sign of the error, we report the absolute end detection error, which is the absolute value of the end detection error. The sensitivity refers to the percentage of test seizures whose end the algorithm detected.

We use a leave-one-record-out testing procedure to evaluate the performance of our seizure end detector. To independently evaluate the seizure end detector, we assume a separate module will detect seizure onset, and test records extend from the electrographic onset of a seizure to $P_D = 8$ minutes following its end. The P_D parameter determines the length of the post-seizure period that the algorithm considers as post-ictal. However, this setting can result in the inclusion of data in the

set of post-ictal training examples that may be considered not to be post-ictal. The transition from post-ictal to background (or inter-ictal) IEEG is often very gradual, and there may not be a clear or unequivocal end to the post-ictal period. From each record we compute approximately 480 feature vectors labeled as post-ictal, and a number of seizure feature vectors proportional to the length of the seizure. A single exception was made for Patient 4, where seizures were less than 8 minutes apart, and the records ended 6 minutes after the seizure end in that case. Let N_i be the number of records in the dataset of patient i. A seizure end detector is trained on feature vectors derived from the ictal and post-ictal periods in N_{i-1} records. Next, the detector is tasked with detecting seizure end in the withheld record. These two steps are repeated N_i times so that each record from patient i is tested once. For each tested record, we note whether seizure end was detected, and if so, with what error.

3.2 Results and Discussion

For 5 out of 10 patients, 100% of seizure ends were detected within a 15-second margin of the expert-marked end. Overall, 88% of all seizure terminations were detected within 15 seconds of the marked seizure ends, and 86% were detected within a 10-second margin. This result can be compared to that obtained with the scalp EEG seizure termination detector in [31], where 81% of seizure ends were detected with an absolute detection error smaller than 15 seconds. A positive detection error (indicating late end detection) was obtained for a large majority of seizures, and this is mainly due to the algorithm requirement that $K_E = 5$ consecutive windows must be deemed by the classifier to correspond to post-ictal epochs before the end can be declared. Among those seizures for which an absolute detection error in excess of 15 seconds was obtained with this algorithm, the vast majority had a negative detection error (corresponding to premature seizure end detections). For a single seizure, a seizure end declaration that was induced by a short-lived artifact was ignored, and the algorithm was allowed to continue a search for the seizure end. We thus consider

artifact detection designed to work as part of a seizure end detector to be task that can be addressed separately.

As discussed in Chapter 2, the Patient 6 seizure of a distinct type relative to the two remaining seizures in the dataset was missed by the seizure onset detector in a leave-one-record-out test. The same phenomenon also lead to a large offset detection error for the seizure of a unique type. Another patient dataset with only 3 seizures also resulted in a low percentage of seizure ends detected within a 10-second margin (66% for Patient 1).

The algorithm declared a premature seizure end in a few cases due to a phenomenon where ictal activity persists on a small number of channels after an abrupt end of such activity on all the remaining channels. An example of this is illustrated in Figure 3-2, which shows a 20-second window of IEEG during a seizure from Patient 1. Whereas the high amplitude activity comes to a halt for most channels during this epoch, near-periodic activity (with decreased amplitude) continues on a few channels ('AnTD','PsTD' or 'SBTP' channels, all corresponding to temporal lobe depth electrodes). The end is declared during the period shown in that figure, but the expert-marked seizure end is approximately 45 seconds after the end of this window, resulting in a very high absolute detection error. The amplitude of this activity is increasingly and gradually attenuated on the five channels after it ends on all other channels. Combined with the cessation of ictal activity on a large majority of the channels, this may be why the classifier began prematurely classifying windows as belonging to the post-ictal class. In the remaining two seizures from this patient, all ictal activity ends abruptly and simultaneously on all channels. The IEEG around the marked seizure end for one of the two remaining seizures is shown in Figure 3-3. This is important to note because the occurrence of this phenomenon is not by itself sufficient to lead to erroneous algorithm end declarations. A premature end declaration resulted in cases where this phenomenon is unique to a single test seizure. If a similar pattern was shown in any of the other seizures from this patient's dataset, it is likely that this very premature end detection could have been avoided.

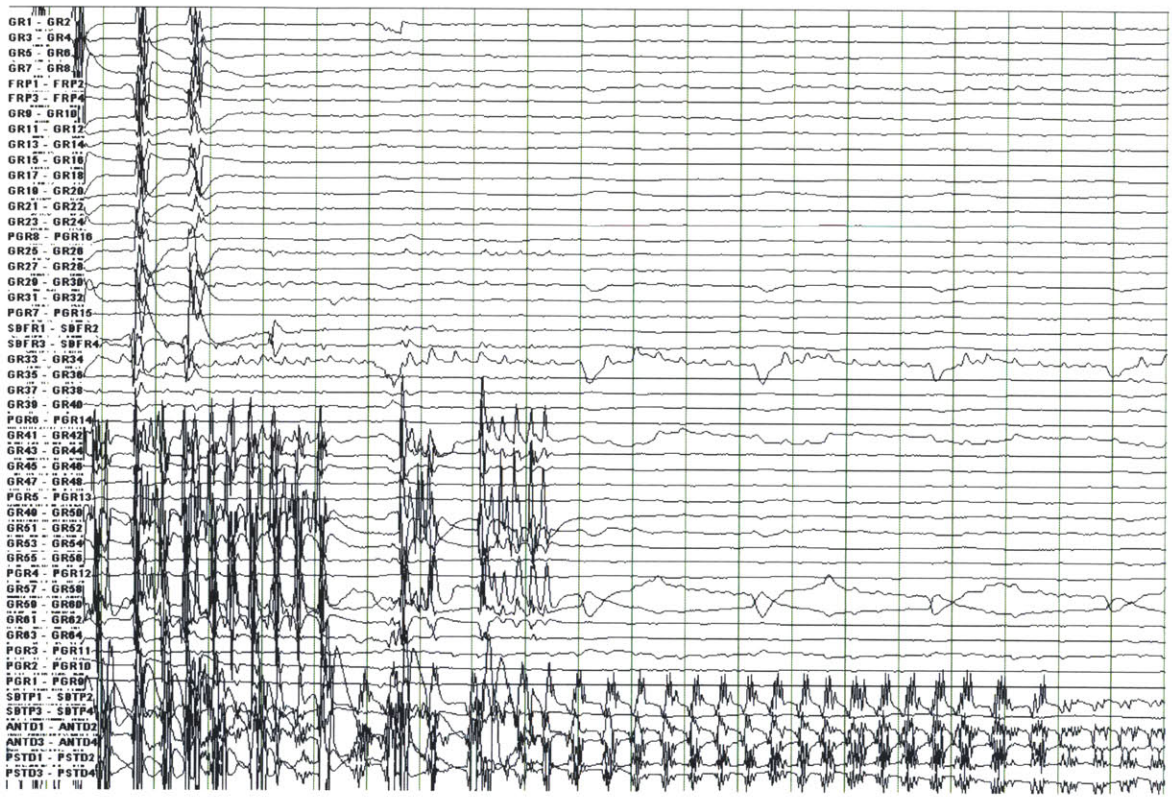

Figure 3-2. A 20-second window of intracranial EEG during a Patient 1 seizure. Whereas the high amplitude activity comes to a halt for most channels during this epoch, near-periodic activity (with decreased amplitude) continues on a few channels. A premature end detection was obtained for this seizure. The expert-marked seizure end is approximately 45 seconds after the end of this window

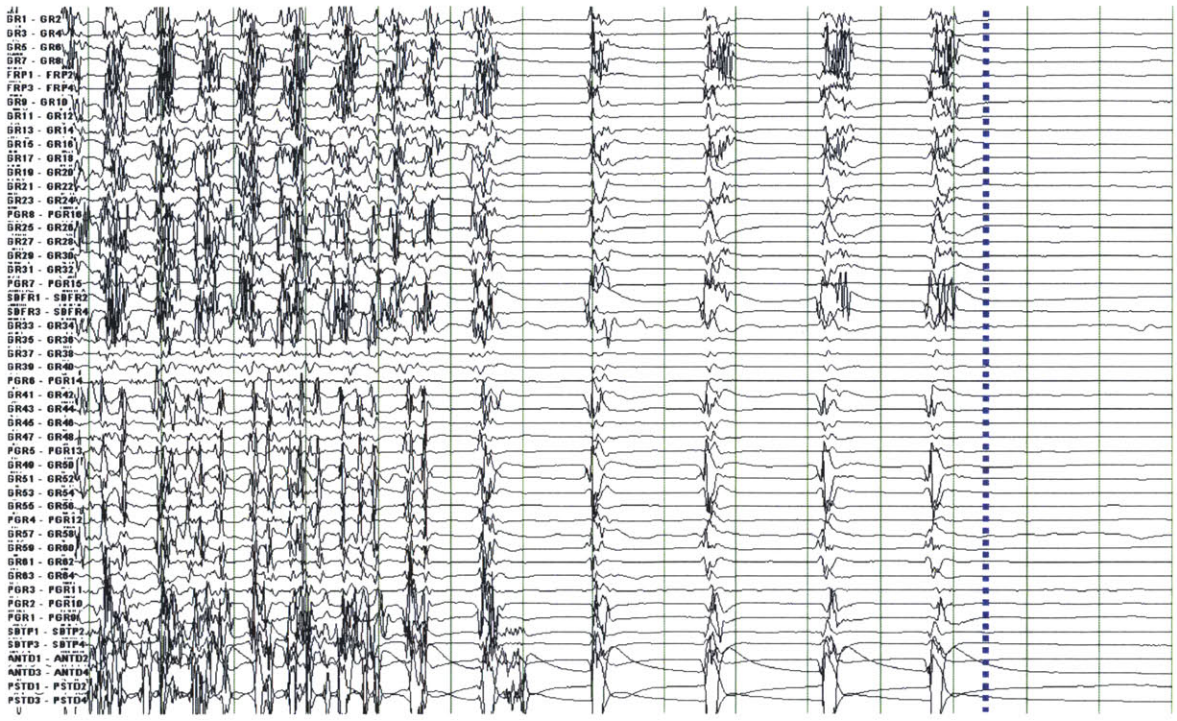

Figure 3-3. A 20-second window of intracranial EEG from Patient 1 showing the expert-marked end of a seizure (blue dotted line). Ictal activity ends abruptly on all channels nearly simultaneously, in contrast to the seizure end shown in Figure 3-2

We also had a premature end detection in a case where the highest-intensity ictal activity transitioned between different areas in the brain, with a lull in high amplitude activity in between. IEEG from a section of this patient's seizure are shown in Figures 3-4 and 3-5. A premature end detection resulted from an unusual seizure pattern where the seizure starts with ictal activity present in both the temporal lobe (as evidenced by the channels labeled 'RAT' and 'RPT'), and the frontal lobe ('RSbF' and 'RPsF' channels). The activity ceased in the temporal lobe, but continued (albeit somewhat attenuated and altered) in the frontal lobe. The frontal lobe activity evolves, as illustrated by Figure 3-5, which shows the IEEG during a 15-second window at a later stage in the seizure. In response primarily to the cessation of temporal lobe ictal activity the classifier begins to misclassify the epochs following that point as post-ictal, leading to an early seizure end declaration. However, as the frontal lobe activity progressed and evolved, the algorithm returned to (correctly) classifying the

later epochs as belonging to the ictal class

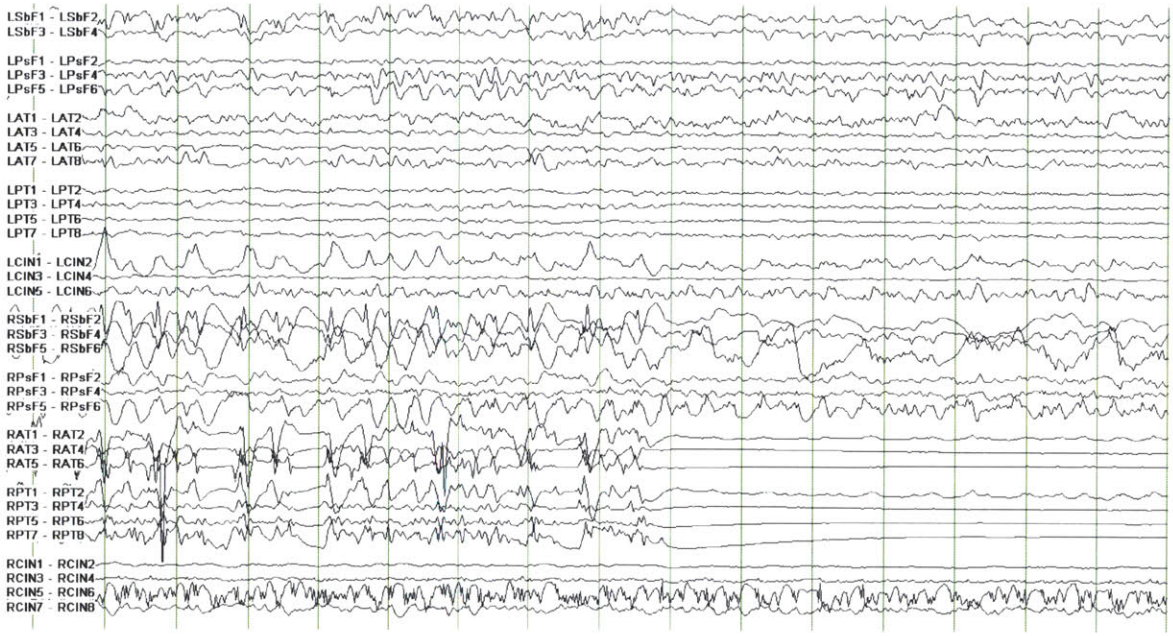

Figure 3-4. A 15-second window of intracranial EEG from a Patient 5 seizure, during which an early seizure end detection occurs. The seizure starts with ictal activity present in both the temporal lobe (as evidenced by the channels labeled 'RAT' and 'RPT'), and the frontal lobe ('RSbF' and 'RPsF' channels). The activity ceased in the temporal lobe, but continued (albeit somewhat attenuated and altered) in the frontal lobe.

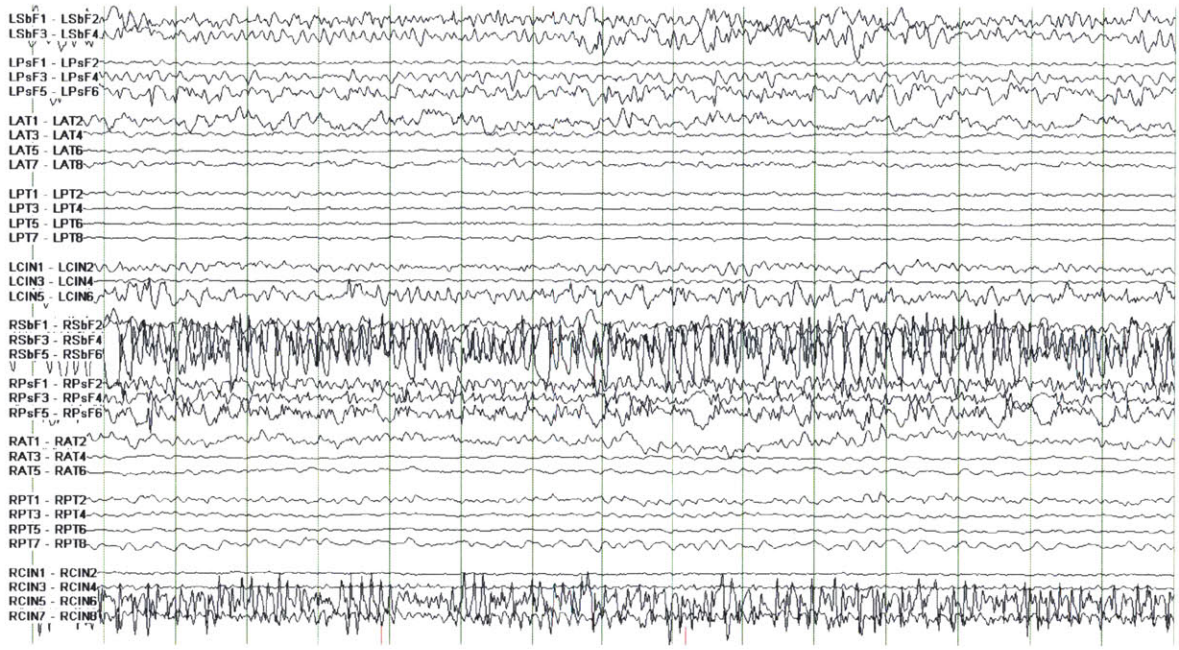

Figure 3-5. A 15-second window of intracranial EEG from Patient 5 seizure starting after the window shown in Figure 3-4. The frontal lobe activity ('RSbF' and 'RPsF' channels) has evolved leading up to this period, and the algorithm correctly classifies these epochs as ictal after a premature seizure end has been declared.

Only 3 out of 7 seizure ends were detected within a 10-second margin of the expert-labeled end for Patient 7. Despite some variation in the seizure patterns, the seizure onset detector demonstrated relatively low latency when tested on this patient data set. However, perhaps a greater degree of latter-stage ictal pattern diversity led to less impressive seizure end detection results. In addition, many of this patient's seizures featured a gradual ending, marked by an expert based on subtle changes in activity. In addition, it was very difficult to objectively establish a clear seizure end point, and subjectivity may have played a large role.

In a few cases an early seizure end detection resulted from the misclassification of clearly ictal activity. Although a seizure end was declared by the algorithm during periods with activity associated with the ictal state (such as high amplitude or rhythmic activity), this can occur when certain aspects of the spatio-spectral electrographic pattern has not been observed during seizures used for training. Two periods of ictal activity can differ in several ways, including the channels that show ictal activity, the amplitude of the activity, or the approximate dominant (fundamental) frequency of the activity. As an example Figure 3-6 shows a 7-second window of IEEG during which the algorithm declared a premature seizure end. This period features high amplitude activity on the channel labeled 'LAnT1-LAnT2', but the classifier fails to recognize a series of epochs as ictal, and declares an early seizure end. This may be because no similar period in any of the six remaining seizures used for training featured clear ictal activity on this channel.

Problems from cases like this can be remedied by an increase in the number of training seizures, to ensure that the classifier is exposed to as many of the patient's electrographic ictal patterns as possible. However, in absence of additional data, one way the likelihood of early termination detections in similar scenarios could be reduced is by supplementing the method with a patient non-specific component. This could be a knowledge-based approach, i.e. translating a clinician's characterization of indisputable ictal activity to some simple heuristics on a few measures (such as signal line-length or rough measures of signal periodicity). Alternatively, this component could be machine-learning-based, i.e., trained using the seizure and post-ictal data

from many patients. This component would likely have to be designed to be conservative in the sense that it errs on the side of declaring non-seizure activity, to avoid delaying the declaration of seizure termination. However, in cases like this with less subtle ictal activity this component can intervene to prevent an early declaration of seizure end. However, one issue that must be addressed before such is the lack of a consistent channel map across different patients for intracranial EEG. The channels correspond to a different set of locations on or within the brain for different patients. In addition to channels corresponding to different locations, the number of channels generally differs from patient to patient, and the feature extraction process described in this chapter will yield feature vectors of different sizes for different patient. Therefore, a way of transferring what is learned from electrodes in different configurations and locations from other patients must be devised before a similar patient-non-specific solution can be applied.

Figure 3-6. A 7-second window of intracranial EEG from a Patient 7 seizure during which an early seizure end is declared by the algorithm. This period features high-amplitude ictal activity on the channel labeled 'LAnT1-LAnT2'

Chapter 4

Channel Pre-selection and Patient-Specific Feature Extraction in Seizure Onset Detection

For the results in Chapter 2 we did not require the manual pre-selection of a small set of channels by an expert. In this chapter, we present the results when the seizure detector is restricted to the use of a small number of pre-selected channels (2-4). This restriction may, for example, be due to instrumentation power constraints. Another possible motivation for a reduction in the size of the channel set is an improvement in performance, since allowing all channels to be used may lead the algorithm to incorporate non-specific information (from channels not involved early in the onset) that may increase the likelihood of incorrect classification.

EEG channel selection could be accomplished in an automated manner, by means of algorithms that are designed to use training data to search for a reduced channel set that results in good performance [26]. However, in this chapter we explore the effect on automated seizure onset performance when the channels are manually selected based on visual examination, and correspond to the channels that display the earliest onset activity for each patient.

We also present and evaluate an algorithm for patient-specific feature extraction, an approach where the feature extraction process for a given patient leverages the training data available for that patient. The results from an evaluation of a detector that supplemented the original spectral energy features with features computed in a

patient-specific manner are also presented.

4.1 Channel Pre-selection Examples

In cases where a limited number of channels is used, it is still a goal of automatic seizure onset detection to declare a seizure with as low a latency as possible. Therefore, if the information from only a small number of channels will serve as input to an algorithm, they are usually selected to be those that display the earliest ictal activity (including the activity that was identified by an expert to determine the onset). This is often associated with an assumption that one or few seizures are representative of other or future seizures in terms of the identity of these channels. Without a similar assumption, one cannot ensure that cases where future seizures show the earliest activity on channels that have not been pre-selected for a particular patient will be avoided. This would predictably have undesired effects on detection latency. The feasibility of disposing of all but few of the IEEG channels is linked with spatial consistency in the onset activity. Therefore we focus only on patients whose seizures show the highest level of consistency in terms of the channels that show the earliest signs of seizure activity. In this section we discuss an example of a patient dataset that was included in the evaluation of the seizure onset detector with channel pre-selection, and a patient dataset that was not included in the channel-selection study.

We evaluated the onset detection algorithm with channel pre-selection using the IEEG database described in Chapter 2. One of the patients included in this channel selection study is Patient 2. Figure 4-1 shows a 6-second window of intracranial EEG from a Patient 2 seizure that includes its electrographic onset. The seizure begins with spikes on three channels labeled 'RAT1-RAT2', 'RAT3-RAT4' and 'RPT1-RPT2', followed by activity which evolves and becomes rhythmic. These channels all correspond to temporal lobe depth electrodes. Figure 4-2 shows the onset of another Patient 2 seizure. Similarly, the seizure also begins with spikes on the same three channels. Rhythmic activity also follows the seizure in Figure 4-2, although the activity that immediately follows the spike in Figure 4-1 is more irregular before it evolves. All six

of this patient's seizure onsets are signalled by similar spike activity (as determined by an expert), but most importantly for the purposes of this section, this occurs on the same set of three channels for all seizures. Given this high level of consistency in terms of the channels that exhibited the activity that was determined to mark the onset, this patient was included in the channel pre-selection study. The channels that were selected are the three aforementioned temporal lobe channels.

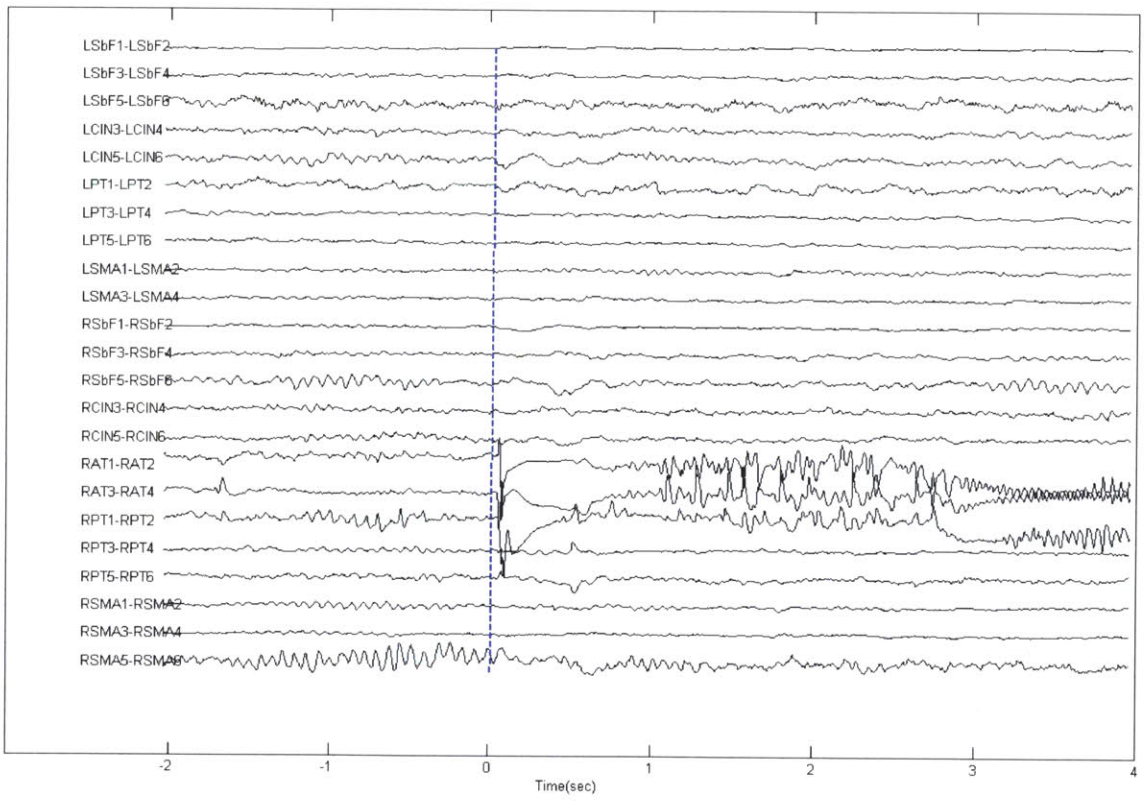

Figure 4-1. A 6-second window of intracranial EEG from a Patient 2 seizure showing the onset. The onset is indicated by the dotted blue line. The seizure begins with spikes on 3 channels labeled 'RAT1-RAT2', 'RAT3-RAT4' and 'RPT1-RPT2', followed by activity which evolves and becomes rhythmic.

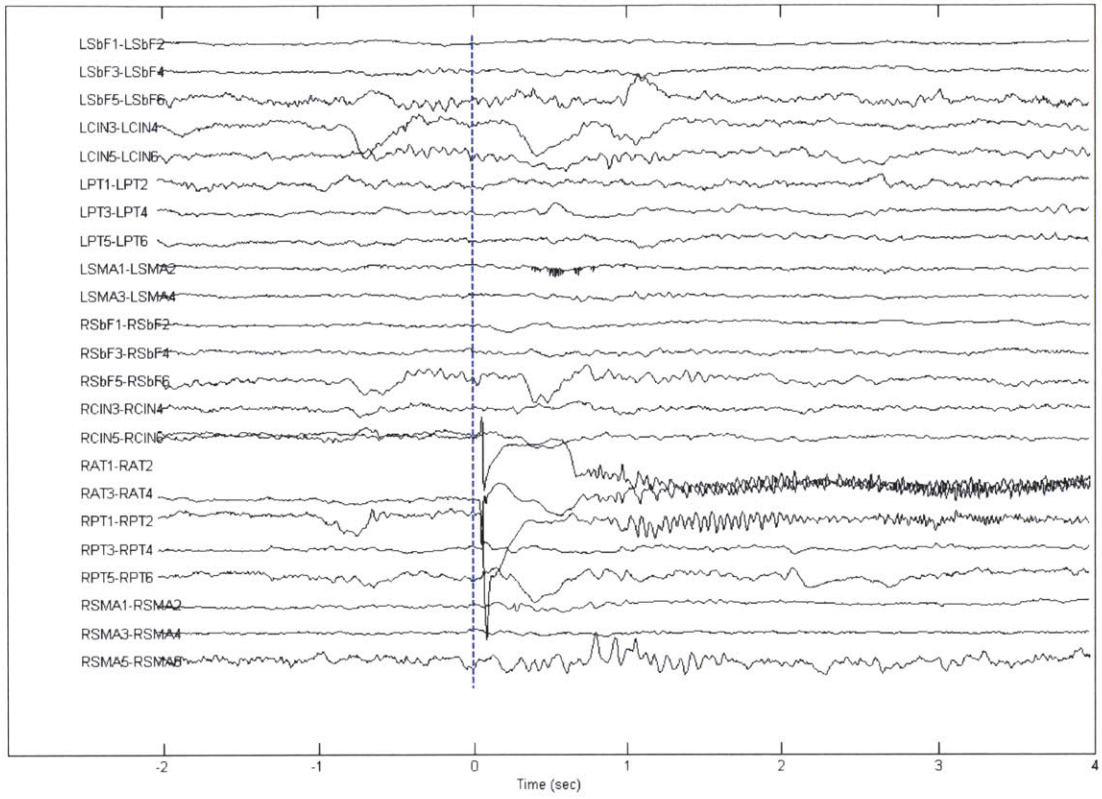

Figure 4-2. A 6-second window of intracranial EEG showing the onset of another Patient 2 seizure. The onset is indicated by the dotted blue line. The seizure also begins with spikes on 3 channels labeled 'RAT1-RAT2','RAT3-RAT4' and 'RPT1-RPT2', followed by rhythmic activity.

The three seizures of Patient 1 did not demonstrate a similar degree of spatial onset consistency. Figure 4-3 shows the onset of a Patient 1 seizure. The onset activity appears on the channels labeled 'ANTD' and 'PSTD', all corresponding to temporal lobe depth electrodes. Another seizure from this patient also features activity on the same channels at the expert-marked electrographic onset. Figure 4-4 shows the third (and last) seizure from the same patient. Unlike the other two recorded seizures in this dataset, this seizure's onset was determined by the small-amplitude rhythmic activity on the channel labeled 'PGR8-PGR13'. This seizure is unique in that it is the only seizure in the set where the earliest ictal activity appeared on this channel. Due to the inconsistency in the identity of the channels that show the earliest ictal activity, this patient data set was not used for the channel pre-selection study.

It should be noted that in a practical or clinical setting, it may be difficult to determine the number of seizures that must be observed before any conclusion regarding consistency in the identity of channels that show the earliest seizure activity can be reached. One aim of this chapter is to explore the effect of channel pre-selection on onset detection performance when the lack of consistency in the channels that display the earliest activity is not a (predictable) source of performance degradation, and the assumption regarding a high level of consistency holds. The patient data sets used in the channel pre-selection study include Patient 2, Patient 3, Patient 4, Patient 5 and Patient 9. Throughout this chapter they will instead be referred to as Patients A,B,C,D and E respectively. Table 4.1 contains some general information for all 5 patient data sets, including the number of pre-selected channels.

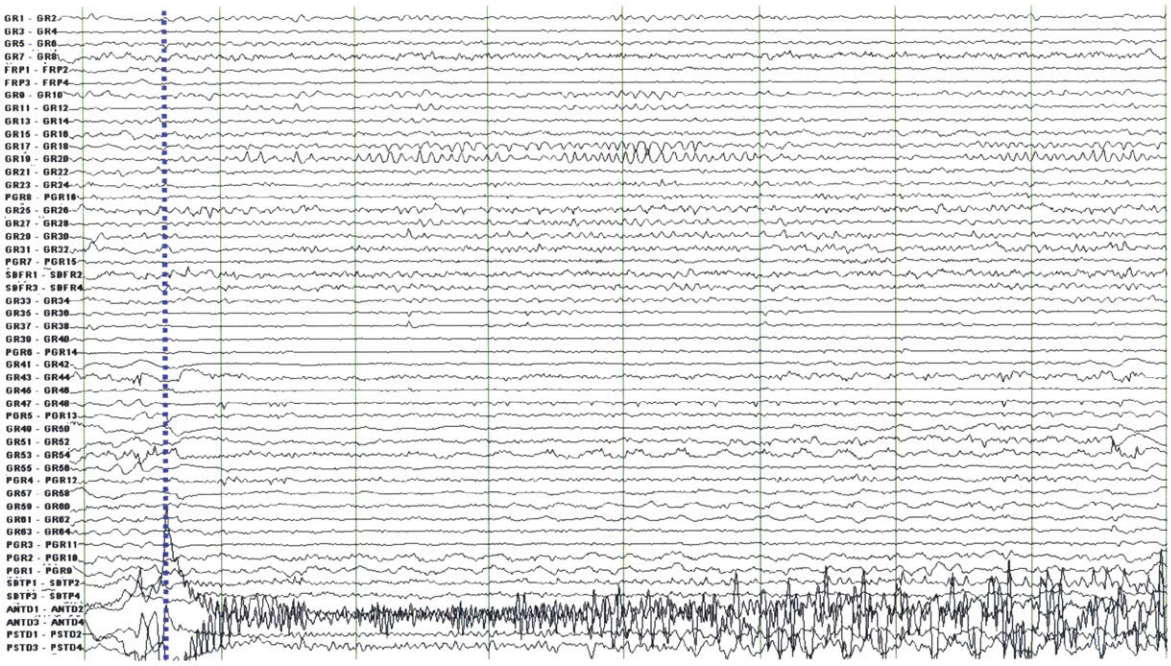

Figure 4-3. A 7-second window of intracranial EEG from a Patient 1 seizure showing the onset. The onset is indicated by the dotted blue line. The onset activity appears on the channels labeled 'ANTD' and 'PSTD', which correspond to temporal lobe depth electrodes

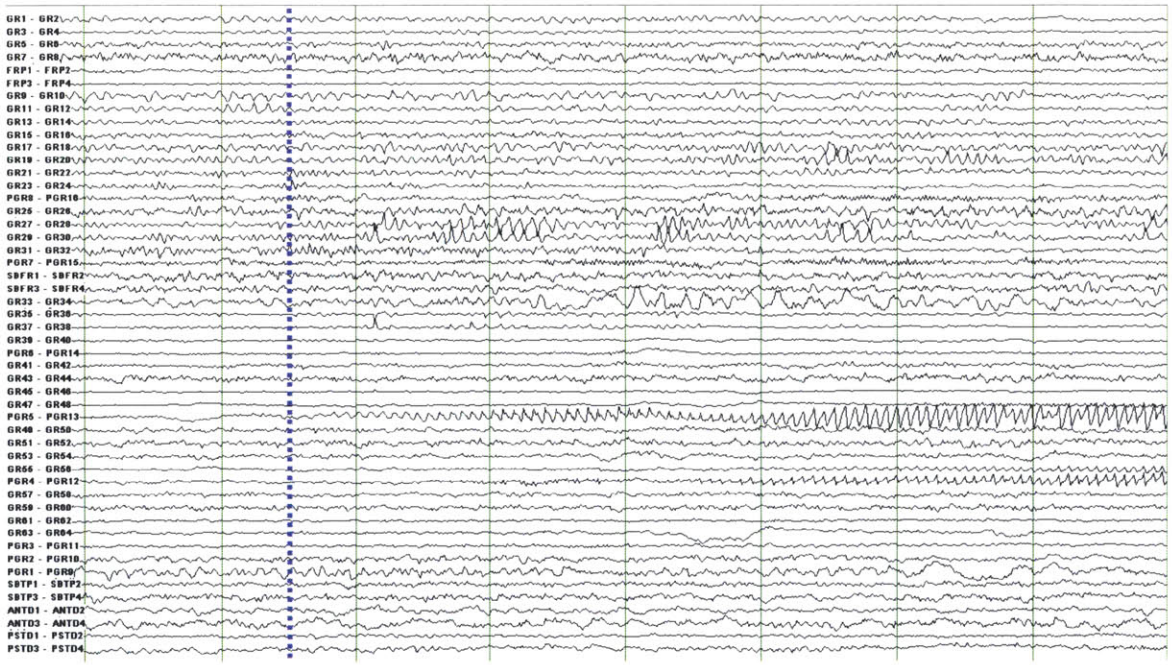

Figure 4-4. A 7-second window of intracranial EEG showing the onset of another Patient 1 seizure. The onset is indicated by the dotted blue line. Unlike the other two seizures recorded from this patient, the seizure begins with small-amplitude rhythmic activity on the channel labeled 'PGR8-PGR13'.

Patient	A	B	C	D	E
Total time tested(Hr)	127	71	29	96	148
Number of seizures	6	6	5	8	11
# of selected channels	3	4	2	2	4
Electrode types	Depth	Grid & Depth	Depth	Depth	Depth

Table 4.1. Patient data set information for the database used to evaluate the seizure end detector with channel pre-selection.

4.2 Performance of Onset Detection Algorithm With Channel Pre-selection

The results presented in this section are obtained by modifying the full-channel feature vectors (described in Chapter 2) so that only the features corresponding to pre-selected channels are preserved. The output of the artifact rejection component described in Section 2.1.2, which used the full channel set, was again used to prevent alarms during epochs where high-amplitude artifact was detected in this study. The SVM error cost parameter was set to $C = 1 \times 10^{-2}$ for all patients. Overall, 100% of the 36 test seizures were detected. The median latencies for each patient are shown in Table 4.2. The median false alarm rate was 1.5 false detections per 24 hour period. The average false alarm rates for each patient are also shown in Table 4.2. For comparison, the results from the full-channel study from Chapter 2 are repeated in Table 4.3 for convenience. Decreasing the cost parameter C led to a larger training error, particularly for the seizure class. We found that in many cases it was advantageous to allow for an increase in training error as this often resulted in a large decrease in false alarm rate at the expense of a relatively small increase in latency.

Patient	A	B	C	D	E
Sensitivity	100%	100%	100%	100%	100%
Median Latency(sec)	4	3.75	4	4.75	7
False alarms/ 24Hr	0.4	4.1	0.8	3.1	1.5

Table 4.2. Sensitivities, median latencies and false alarm rates obtained for each patient data set from evaluation of the seizure onset detector with channel pre-selection.

Patient	A	B	C	D	E
Sensitivity	100%	100%	100%	100%	100%
Median Latency(sec)	3.25	3.25	5	5.75	4.5
False alarms/ 24Hr	0.4	1.7	0	0.8	0.3

Table 4.3. Sensitivities, median latencies and false alarm rates obtained for each patient data set from evaluation of the seizure onset detector using a full channel set.

The median latency increased with channel pre-selection for 3 out of 5 patients, with the largest increase seen for Patient E (2.5 seconds). For the two patients where latency decreased (with channel pre-selection), the difference in median latency was exactly one second. In fact, the difference in median latency between the channel-selected and full-channel results did not exceed one second for 4 out of the 5 patients. The most common effect on results with channel pre-selection was an increase in false alarm rate, which was the case for 4 out of 5 patients. However, the effect of increase in the occurrence of false alarms may have been exaggerated in the case of Patient B, where more than 80% of false alarm occurred within a single 3-hour period (out of a total of 73 hours) due to the persistence of recording artifact in varying forms.

Overall, these results show that while a small reduction in latency may result from channel pre-selection in some cases, a more common effect on performance may be an increased susceptibility of the algorithm to false alarms. An example of a false alarm that occurred when the algorithm was restricted to using the selected channels is shown in Figure 4-5. A false alarm was raised during the 6-second window of IEEG in the figure, which shows the runs of irregular high-amplitude activity (at approximately 10Hz) during this epoch on the channel labeled 'RAT1-RAT2' (with matching lower amplitude activity on the 'RPT1-RPT2'). This activity is very similar to the activity that appears on the same channels during the onset of this patient's seizures. Useful information about the activity on channels in the full set that is not obvious on visual examination (and ignored in the channel pre-selection case) may have prevented this and similar false alarms when the algorithm used all available channels. Such examples are consistent with the conjecture that the inclusion of 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 inter-ictal patterns.

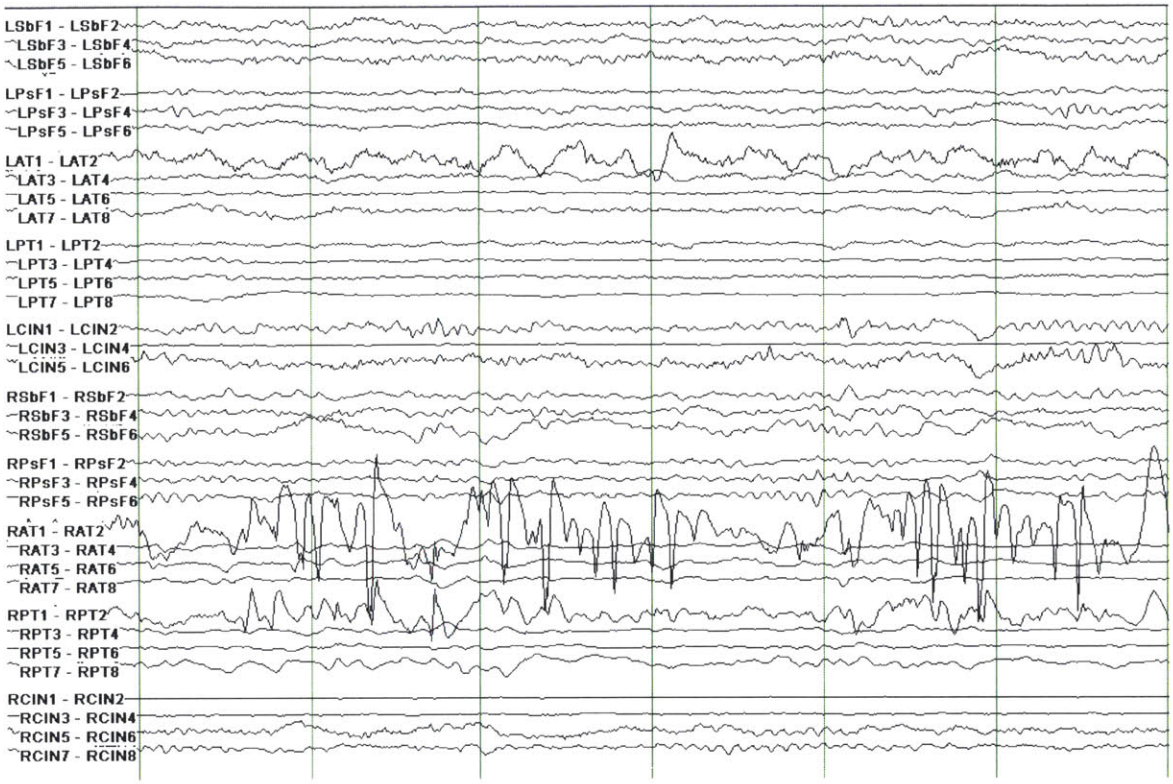

Figure 4-5. A 6-second window of intracranial EEG recorded from Patient D during which a false alarm occurred. The two channels that were pre-selected for this patient are labeled 'RAT1-RAT2' and 'RPT1-RPT2'. The onsets of a few of this patient's seizures feature similar activity

4.3 Patient-Specific Feature Extraction

Patient-Specific Feature Extraction (PSFE) refers to a method where the feature extraction stage is tailored to a patient using training data. The goal is to combine the original spectral features with salient features learned from labeled examples from the patient to improve detection performance. In general, the strategy is to learn certain directions in IEEG signal morphology or spectral signatures whose presence (or lack thereof) help indicate whether a given IEEG epoch belonged to the seizure or non-seizure class.

Figure 4-6 shows a high-level view of the algorithm with the original feature extraction strategy, as described in Chapter 2. A feature extraction component converts every window of multi-channel IEEG into a feature vector, and the feature vectors are either used to train an SVM (along with the respective labels) or to classify an epoch outside the training set as belonging to the seizure or non-seizure class. The feature extraction process uses a bandpass filter bank, and each filter is defined by the center frequency and the width of the pass-band. These properties are fixed, and do not vary from patient to patient. The respective energies of the output of the filterbank give us a coarse description of the spectral profile of the IEEG signal. However, it may be that the choice of filters can be better tailored to the patient, or that adjusting some of these parameters will lead to superior seizure detection performance. For example, for a certain patient a smaller bandwidth for filters at certain frequencies (and larger bandwidths for others) may be more effective. The above is just an example that introduces the concept that we can move away from the previous method towards a more patient-tailored mode, but the methods proposed here are not restricted or parameterized in precisely this fashion. However, they center around alternate signal descriptions, with some emphasis placed on discriminative value in their derivation.

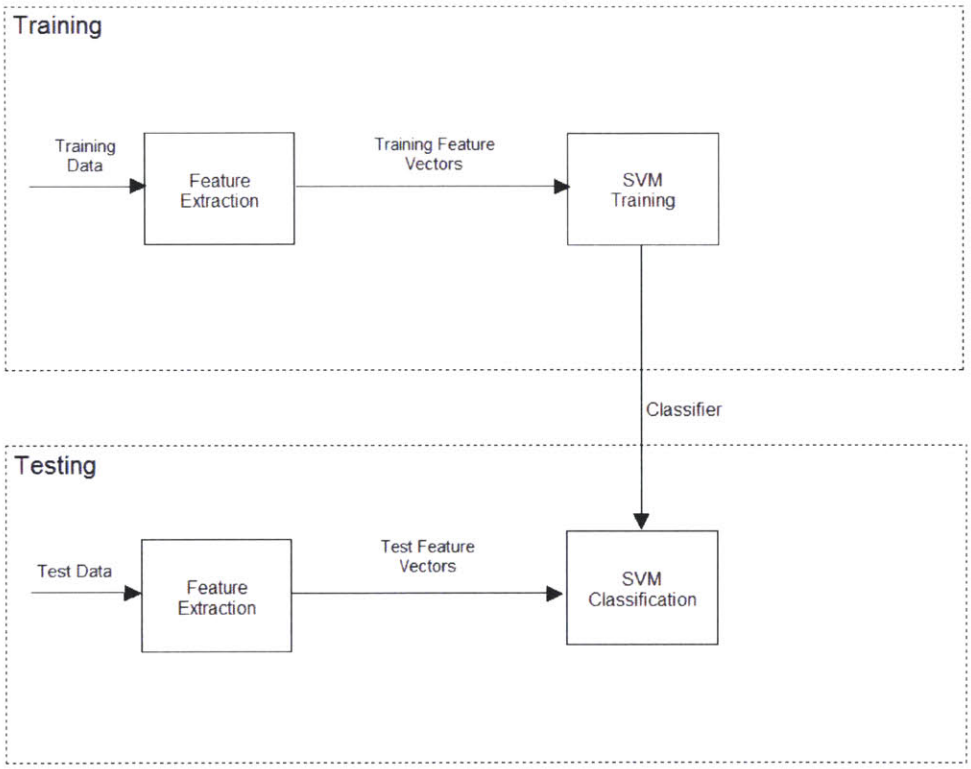

Figure 4-6. A high-level view of the algorithm with the original feature extraction strategy, as described in Chapter 2.

Figure 4-7 illustrates the role of the various components of the algorithm with the proposed patient-specific feature extraction method. Before training (or test) feature vectors can be computed, training data is used to determine a crucial aspect of the feature extraction process (during the stage represented in the figure as "feature extraction training"). The output of the Patient-specific feature extraction training component consists of a dictionary, which is a set of prototype sequences which can be combined to approximate an IEEG signal. The dictionary and the way it factors into feature extraction is further discussed in Section 4.3.1. The PSFE training (i.e. the dictionary learning) process is described in Section 4.3.2.

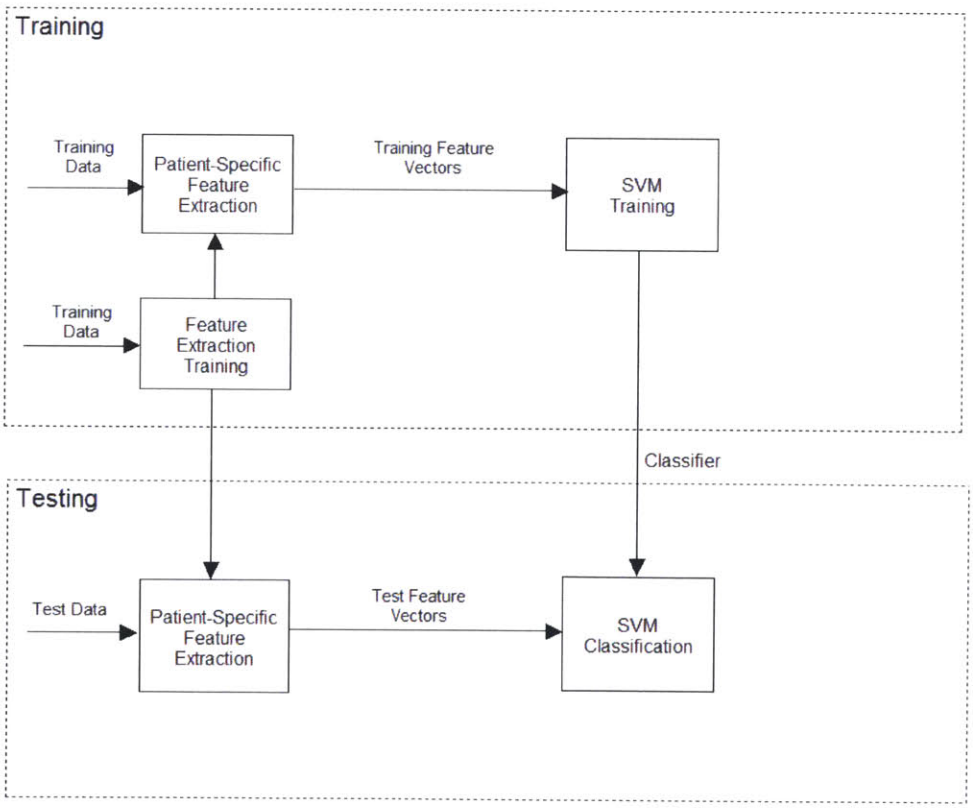

Figure 4-7. High-level illustration of the algorithm steps with the proposed patient-specific feature extraction method

The input to the patient-specific feature extraction training component comes in the form of seizure and non-seizure training examples, where each example corresponds to an epoch of IEEG on a single channel. Since low detection latency is an important goal in seizure detection, the examples from the seizure class are chosen to correspond to channels that show the earliest onset activity. The feature extraction recipe that was determined using only these channels may not be appropriate to represent the activity on the remaining channels, which can differ to a significant extent in many ways. Therefore, we evaluate the algorithm only on the patients for whom we have pre-selected the channels that show the earliest seizure activity, and focus on the case where the algorithm is restricted to using only these channels to compute feature vectors.

We limit the number of non-seizure training examples used for PSFE training, such that they do not greatly outnumber the number of training examples from the seizure class. This is done to reduce the computational requirement during this stage. In addition, this can prevent cases where a small minority of non-seizure examples exhibiting overlapping morphology with that seen during seizures is ignored in the PSFE training phase because this direction does not have significant presence in the non-seizure class on aggregate. The presence of such activity, even if relatively rare within the non-seizure class (i.e. as a ratio of the total number of examples), can still lead to a significant and undesirable number of false alarms. Selecting a smaller number of non-seizure training examples as an input to the PSFE training phase can avoid such a scenario. The non-seizure epochs for the PSFE training phase are selected to correspond to the support vectors that were determined through SVM training using the original (spectral energy) feature vectors described in Chapter 2.

In the following sections we provide some more detail about the feature extraction framework that will supplement the fixed bandpass filterbank of the previous algorithm, and introduce some mathematical notation.

4.3.1 Feature Extraction Steps

A signal $\mathbf{y}$ can be represented or approximated using a linear combination of other signals or vectors, which we refer to throughout this document as *dictionary elements*. We refer to a set of these vectors as a *dictionary*. In the context of PSFE, the modeled signals $\mathbf{y}$ represent a finite-length epoch of single-channel IEEG. If we denote an approximation of $\mathbf{y}$ as $\hat{\mathbf{y}}$, then

$$\hat{\mathbf{y}} = \sum_i x_i \mathbf{d_i}$$

where the $\mathbf{d_i}$ are the dictionary elements, and the coefficients x_i provide a recipe for synthesizing the approximation of the vector $\mathbf{y}$ from the dictionary. A more compact description of the previous equation can be written in matrix form

$$\hat{\mathbf{y}} = \mathbf{D}\mathbf{x}$$

where the $\mathbf{d_i}$ vectors form the columns of $\mathbf{D}$, and the ith entry of $\mathbf{x}$ is x_i.

When the dictionary elements are fixed or known, the coefficients in $\mathbf{x}$ form an alternative description of the signal. In the case where the representation is exact, it is a complete description, but we will not be concerned with this case throughout this chapter. We are interested in approximating the signal, i.e., allowing for some error. In such cases, a primary consideration is the reduction of the approximation error, which can be quantified as the mean square error

$$E(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2 = \|\mathbf{y} - \mathbf{D}\mathbf{x}\|_2^2 \tag{4.1}$$

where $\|\mathbf{z}\|_2^2 = z^{\mathrm{T}} z$ denotes the square of the Euclidean norm. The representation in a given dictionary (more precisely, the logarithm of the entries in the $\mathbf{x}$ vector) essentially constitutes the features we extract from an epoch from a single channel. We denote the number of dictionary elements as N_D. Since the number of coefficients is equal to the number of dictionary elements, the number of features extracted per channel for each epoch is equal to N_D. Given a dictionary, the features extracted for a given epoch correspond to coefficients that minimize the approxima-

tion/representation error. However, in the next section we describe how we derive a set of directions using a discriminative objective to obtain the dictionary.

The choice of the dictionary elements has a large influence on the resulting coefficients. Some properties or tendencies in the corresponding coefficient vectors may be desirable for a given application, such as sparsity for compression, or segregation of the effect of noise in a subset of the coefficients for denoising. The difference in the feature extraction process for each patient comes from the use of a different dictionary $\mathbf{D}$ for each patient. The dictionary elements $\mathbf{d_i}$ are learned using several examples $\mathbf{y_j}$, each corresponding to an epoch of IEEG on one channel.

The feature extraction methodology described here does not represent a complete departure from the original feature extraction phase, which computes spectral energy features. Previously, a given channel is passed through a filterbank consisting of simple bandpass filters, and the logarithm of the norms of the outputs are calculated as the features for that channel. If the raw IEEG is first pre-processed as described in Chapter 2, this approximately corresponds to setting the dictionary to the Fourier basis (each basis vector $\mathrm{d_i}$ is a sinusoid of a different frequency), and then summing the square magnitude of all the coefficients x_i corresponding to frequencies within a given band to compute the feature for that band-channel pair. The Fourier basis has a special property in that it is an orthogonal[1] basis (i.e. $\langle \mathrm{d_i}, \mathrm{d_j} \rangle = 0$ for $i \neq j$). A consequence of this suggests a way of computing features when we restrict ourselves to such bases/dictionaries. Specifically, when the basis set is orthonormal (i.e., an orthogonal set of vectors of unit norm), given a signal $\mathbf{y}$ each coefficient x_i can be obtained by taking an inner product between the signal and the basis vector. When perfect representation is not possible, the choice of coefficients that minimizes the mean square error is

$$x_i = \langle \mathbf{y}, \mathbf{d_i} \rangle$$

These coefficients can then be used to build the feature vector in cases where the dictionary elements are orthonormal.

[1]For the purposes of this document we focus on the inner product $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a^T b} = \sum_i a_i b_i$.

For the new PSFE features described in this chapter, the IEEG signals undergo some simple pre-processing to obtain a vector $\mathbf{y}$ for each window on a given channel. A derivative filter is applied, and this is followed by a computation of the square of the magnitude of samples of the discrete-time fourier transform.

The patient-specific feature extraction scheme is composed of a method to learn the dictionary, and a procedure to extract features once the dictionary is learned that we have described in this section. In a broad sense, we desire a choice of dictionary which yields coefficient vectors that exhibit different patterns depending on whether the signal to be modeled using the dictionary elements belongs to the seizure or non-seizure class. Said differently, we roughly wish to learn dictionaries that yield good separability between classes in the corresponding coefficient space. The dictionary learning (i.e., the feature extraction training) method we use is described in the next section.

4.3.2 Dictionary Learning Method

To learn a dictionary from a patient dataset, the general strategy is to find directions or morphologies that have a large presence in one class while having a minimal presence in the other class. More specifically, this methodology searches for directions in which the power in the two classes shows a large difference.

The dictionary elements $\mathbf{d_i}$ are learned using several examples, each corresponding to an epoch of IEEG on one channel. These can be compactly represented as the columns of a matrix. Each example belongs to either one of two classes (seizure and non-seizure), and these class labels are used in this learning process. It is convenient to separate the examples into a seizure matrix $\mathbf{Y_S}$ and non-seizure matrix $\mathbf{Y_{NS}}$.

We begin by specifying an objective function that quantifies a notion of discriminative value for a given dictionary element. We denote this function as $F(\mathbf{d})$, and we can determine a dictionary element by searching for maximizers of this function. We use the objective function

$$F(\mathbf{d}) = \frac{|\frac{1}{N_1}\sum_i \langle \mathbf{y_S^i}, \mathbf{d}\rangle^2 - \frac{1}{N_2}\sum_i \langle \mathbf{y_{NS}^i}, \mathbf{d}\rangle^2|}{\|\mathbf{d}\|_2^2} \tag{4.2}$$

where $\mathbf{y_S^i}$ denotes an example from the seizure class, $\mathbf{y_{NS}^i}$ denotes an example from the non-seizure class, and N_1 and N_2 represent the number of seizure and non-seizure examples respectively. We normalize by the square of the Euclidean norm of $\mathbf{d}$ (denoted by $||\mathbf{d}||_\mathbf{2}^\mathbf{2}$) to prevent arbitrary scaling of $\mathbf{d}$ from influencing the objective.

Our strategy is to successively search for dictionary elements. We can find the first dictionary element $\mathbf{d_1}$ by finding the vector that maximizes this function. However, after finding the first element, repeating this process to obtain the remaining dictionary members without any additional conditions will only yield an identical result. We therefore constrain the derived dictionary elements to be orthogonal. The orthogonality constraint on the set of dictionary elements provides a way of ensuring some diversity among the dictionary elements. As each dictionary element is sequentially derived using mathematical criteria that represents the separation between the two classes, the subsequent dictionary element is additionally constrained to be orthogonal to it (and others) to avoid redundant or extraneous features resulting from these solutions. A consequence of this is that the mathematical representation of the extracted features is simplified. As discussed in the previous section, when the dictionary is orthonormal, the least-square approximation coefficients can be obtained by simply taking the inner product between the dictionary elements and the IEEG window. Thus, a single matrix multiplication can yield all the features from several IEEG channels in a single epoch. In addition, the method described in this section searches for a dictionary element by maximizing the difference between the averages for the two classes of the square-magnitude of the corresponding feature. As an example that provides some intuition behind the interpretation of orthogonality, the frequency responses of the non-overlapping ideal bandpass filters used for the original spectral energy features are orthogonal (i.e. the inner product is zero). Increasing the overlap in frequency increases the inner product, so in that case orthogonality of the filter impulse responses indicates that they account for activity in different sections of the frequency axis.

To show how the maximizer of $F(\mathbf{d})$ can be found, we can rewrite the expression

in equation 4.2 as follows:

$$F(\mathbf{d}) = \frac{|\frac{1}{N_1}\sum_i \langle \mathbf{y_S^i}, \mathbf{d}\rangle^2 - \frac{1}{N_2}\sum_i \langle \mathbf{y_{NS}^i}, \mathbf{d}\rangle^2|}{\|\mathbf{d}\|_2^2} \tag{4.3}$$

$$= \frac{|\frac{1}{N_1}(\mathbf{d^T Y_S})(\mathbf{Y_S^T d}) - \frac{1}{N_2}(\mathbf{d^T Y_{NS}})(\mathbf{Y_{NS}^T d})|}{\|\mathbf{d}\|_2^2} \tag{4.4}$$

$$= \frac{|\mathbf{d^T}(\frac{1}{N_1}\mathbf{Y_S Y_S^T} - \frac{1}{N_2}\mathbf{Y_{NS} Y_{NS}^T})\mathbf{d}|}{\|\mathbf{d}\|_2^2} \tag{4.5}$$

$$= \frac{|\mathbf{d^T W d}|}{\|\mathbf{d}\|_2^2} \tag{4.6}$$

where

$$\mathbf{W} = \frac{1}{N_1}\mathbf{Y_S Y_S^T} - \frac{1}{N_2}\mathbf{Y_{NS} Y_{NS}^T} \tag{4.7}$$

is real-valued and symmetric. It follows from the Rayleigh-Ritz theorem [13] that the solution to this optimization problem is found by setting $\mathbf{d}$ to the eigenvector of the matrix $\mathbf{W}$ corresponding to the largest-magnitude eigenvalue. This solution corresponds to the first dictionary element, and we require that the rest of the dictionary elements be orthogonal to this vector. However, since the $\mathbf{W}$ matrix is symmetric, its eigenvectors are orthogonal. Therefore, all the dictionary elements can be determined by finding the eigenvectors corresponding the N_D largest eigenvalues, where N_D is the desired number of features per IEEG channel.

4.3.3 Seizure Onset Detection with PSFE Results

The results presented in this section are obtained by combining the channel-selected spectral energy features described in Section 4.2 with the PSFE features discussed in this chapter. The number of features per channel for the original feature set corresponded to the number of filters in the filterbank $M = 17$. However, in preliminary analysis it appeared that reducing this number for the PSFE case improved performance, and we set the number of PSFE features per channel to $N_D = 9$. The SVM error cost parameter was set to $C = 1 \times 10^{-3}$ for all patients. Overall, 100% of the 36 test seizures were detected. The median latencies for each patient are shown in Table 4.4. The median false alarm rate was 0.5 false detections per 24 hour period. The

average false alarm rates for each patient are also shown in Table 4.4. For comparison, the results from the channel pre-selection study with the original spectral energy features only in Section 4.2 are repeated in Table 4.5 for convenience.

Patient	A	B	C	D	E
Sensitivity	100%	100%	100%	100%	100%
Median Latency(sec)	3.75	5	4.5	4.25	6
False alarms/ 24Hr	0	4.4	0	1.5	0.5

Table 4.4. Sensitivities, median latencies and false alarm rates obtained for each patient data set from evaluation of the seizure onset detector using Patient-Specific Feature Extraction (PSFE).

Patient	A	B	C	D	E
Sensitivity	100%	100%	100%	100%	100%
Median Latency(sec)	4	3.75	4	4.75	7
False alarms/ 24Hr	0.4	4.1	0.8	3.1	1.5

Table 4.5. Sensitivities, median latencies and false alarm rates obtained for each patient data set from evaluation of the seizure onset detector with channel pre-selection using the patient non-specific spectral energy features only.

The estimated false alarm rate and median latency increased with the inclusion of the PSFE features for one patient (Patient B). The estimated false alarm rate decreased and the median latency increased for Patient C. However, the results indicate an improvement in performance for 3 out of the 5 patients (Patient A, Patient D and Patient E). Both the false alarm rate and the median latency fell when the PSFE features are used for these three patients. While the improvement in performance is not entirely consistent across all patients, these results are encouraging. The PSFE features appear to complement the patient non-specific features, and this suggests that this is a promising area of investigation. The methods we have discussed can be further refined, and further algorithm development may lead to more dramatic or consistent improvement in performance. Alternative dictionary learning methods could be considered, and more broadly, different approaches and implementations of patient-specific feature extraction. We hope this encourages the use of this and similar ideas (i.e., the introduction of learning in the feature extraction component) in

other EEG-related applications, or more broadly, in learning-based solutions to other biomedical problems.

4.3.4 Future Work

In this section we propose some future work in the form of an alternative approach to PSFE. The method outlined in this section involves a feature extraction procedure where coefficient vectors are computed from IEEG examples and a dictionary by means of sparse approximation. In sparse approximation, the coefficient vectors are restricted to have a certain number of zero entries. Therefore, this process seeks a least-squares optimal approximation using a linear combination of a limited number of the dictionary elements, and finding the subset of the dictionary that is most effective for this is part of the sparse approximation problem. The idea behind this form of feature extraction is to attempt to learn dictionaries such that certain coefficients are more likely to be zero for one class, and other coefficients tend to be equal to zero for the other class.

We seek a sparse coefficient vector $\mathbf{x}$ to approximate a signal given a (fixed) dictionary $\mathbf{D}$. More specifically, feature extraction through sparse approximation involves a search for a coefficient vector specified as

$$\arg\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{D}\mathbf{x}\|_2^2 \quad \text{such that } \|\mathbf{x}\|_0 \leq T \tag{4.8}$$

where $\|\mathbf{x}\|_0$ denotes the number of nonzero entries in $\mathbf{x}$. Good approximate solutions can be found by means of the Basis Pursuit algorithm [5], which instead solves the convex problem that minimizes the objective $\|\mathbf{y} - \mathbf{D}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_1$ where $\|\mathbf{x}\|_1 = \sum_i |x_i|$. The λ parameter trades off approximation error with sparsity of the resulting coefficient vector.

To learn a dictionary using training data, we can begin by leveraging existing work on dictionary learning for sparse approximation. The K-SVD algorithm finds both a dictionary and the corresponding coefficients that model a set of signals in $\mathbf{Y}$, by reducing the approximation error within strict sparsity constraints [2, 3]. In this way, the algorithm adapts a dictionary to a set of training signals. Specifically, it searches for solutions that reduce the objective $\sum_i \|\mathbf{y_i} - \mathbf{D}\mathbf{x_i}\|_2^2$ with the constraint that $\|\mathbf{x_i}\|_0 \leq T$ for all i.

The algorithm attempts to find a dictionary that minimizes the approximation error given sparsity constraints on the coefficient vectors $\mathbf{x_i}$. Alternatively, the problem can be formulated as the search for a dictionary that maximizes sparsity of the coefficient vectors given a limit to the approximation error on a set of signals in $\mathbf{Y}$.

The algorithm draws its inspiration from the K-means clustering algorithm, and can be viewed as a generalization of the K-means clustering algorithm. As is the case with the K-means algorithm, K-SVD only provides solutions at local minima. However, just as with K-means, multiple initializations of the optimization procedure can increase the likelihood that the algorithm still gives good or even optimal solutions.

The standard K-SVD algorithm derives a dictionary from a single set of given signals, but any consideration for class labels or discrimination between said classes is outside the context of the problem originally addressed by the algorithm. We therefore require a modification or extension of the K-SVD algorithm that is designed to leverage class labels to improve classification performance.

There are discriminant variants of the K-SVD algorithm that have been described in the literature [15, 34]. These methods alter the K-SVD algorithm by incorporating the error cost function of a linear classifier (using the representation coefficients associated with a dictionary) into the objective. The classification error function of a linear classifier on the coefficient vectors serves as an indicator of the separability of the training examples in the coefficient space for a given choice of dictionary. This may allow for the derivation of dictionaries that represent the data well (low approximation error), but also yield representations that a linear classifier can predict the labels from with some accuracy. These and other discriminative dictionary learning

methods can be explored to design an alternate method for patient-specific feature extraction.

Chapter 5

Summary and Discussion

This thesis addresses the computerized analysis of IEEG for the detection of the electrographic onset and end of epileptic seizures. We first described and demonstrated the utility of a learning-based algorithm for the early and accurate determination of seizure onsets from intracranial data.

We analyzed more than 875 hours of continuous intracranial EEG recorded from 10 patients to evaluate our algorithm. The study used more than 87 hours of data per patient, which significantly exceeds the amount of data used in most previous studies. This allows for more realistic estimates of long-term performance. We presented results from a study that uses the full set of intracranial electrodes of sufficient recording quality. The data used to evaluate our detector was recorded at Massachusetts General Hospital. The patients were all surgical candidates who required invasive monitoring, and therefore represent more complicated cases than the general population of patients with epilepsy. For the hospitalized patients from whom the data for this study was recorded, the anti-epileptic drug levels were changed on a daily basis. The medications have a significant effect on both seizure and non-seizure IEEG. The lack of consistency in the magnitude or nature of these effects at different times complicates the evaluation of the detector.

The results from this system provide evidence that patient-specific algorithms can outperform existing methods, especially when trying to balance early detection with a low false alarm rate. The proposed algorithm performance compares favorably to existing methods, as determined by well-defined performance metrics. Overall, 97% of the 67 test seizures were detected, and the sensitivity was 100% for most patients. The

median latency with which the detector declared the seizure onset (across all seizures) was 5 seconds. The median estimated false alarm rate was 0.6 false detections per 24 hour period.

The cases with relatively poor performance in terms of sensitivity or latency were cases with multiple classes of IEEG seizure activity within a patient dataset with a limited number of seizures. In one such case this was combined with considerable variation in morphology or the identity of the channels that displayed the earliest noted seizure activity within those classes. For each test seizure, the paucity or absence of sufficiently similar examples of onset activity from seizures in the training set contributed to the decline in performance. A lower latency was obtained for a dataset that included a large number of seizures recorded from a patient for whom the clinician enumerated several seizure types.

The false alarm events are unevenly concentrated in different regions in time, i.e., they temporally cluster. The temporal clustering of false alarm times has application-specific implications for the potential utility of this system. The IEEG of patients with epilepsy transitions between regimes within both the seizure and the non-seizure states and is, therefore, a nonstationary process. Different periods feature different types of activity. Some activity may be common during some periods, and absent or suppressed in others. The characteristics of IEEG are subject to influence by more global processes or physiological states, such as as circadian rhythms and catamenial cycles. Also, the daily change in medication levels may have played a role in the number (and perhaps the distribution) of false alarms.

In addition, we described and evaluated an algorithm for the detection of the end of seizure activity within IEEG. The algorithm is designed to be paired with a seizure onset detector that triggers its search for the seizure end point. A seizure end detection algorithm can enable important applications, including the delivery of therapies to control postictal symptoms, and the detection of status epilepticus. For 5 out of 10 patients, 100% of seizure ends were detected within a 15-second margin of the expert-marked end. Overall, 88% of all seizure terminations were detected within 15 seconds of the expert-marked seizure ends, and 86% were detected within a

10-second margin. This result can be compared to that obtained with the scalp EEG seizure termination detector in [31], where 81% of seizure ends were detected with an absolute detection error smaller than 15 seconds.

Large variations in late-stage IEEG patterns or patterns that are unique to a single seizure in a patient dataset lead to poor seizure end detection performance in the case of a few seizures. Problems from cases like this might be remedied by an increase in the number of training seizures, to ensure that the classifier is exposed to as many of the patient's electrographic ictal patterns as possible. However, in absence of additional data, one way the likelihood of early termination detections in similar scenarios could be reduced is by supplementing the method with a patient non-specific component. This could be a knowledge-based approach, i.e., translating a clinician's characterization of indisputable ictal activity to some simple heuristics on a few measures (such as signal line-length or rough measures of signal periodicity). Alternatively, this component could be machine-learning-based, i.e., trained using the seizure and post-ictal data from many patients. However, a way of transferring what is learned from electrodes in different configurations and locations from other patients must be devised before such a patient-non-specific solution can be applied.

We explored the effects of manual channel pre-selection on seizure onset detection for some patient data sets. We evaluated and presented the results from a seizure detector that has been restricted to use only a small subset of the channels available. The channels that are selected were those that displayed the earliest onset activity for each patient. The feasibility of disposing of all but few of the IEEG channels is linked with spatial consistency in the onset activity. Therefore we focused only on patients whose seizures show the highest level of consistency in terms of the channels that show the earliest signs of seizure activity. The results indicate that performance can suffer in many cases when the algorithm uses a small set of selected channels, often in the form of an increase in false alarm rate. This suggests that the inclusion of a full channel set 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 inter-ictal patterns.

These results may shed some light on the physiological aspects of seizure generation and propagation. Information from channels that do not display clear ictal activity during the onset appear to be useful for discriminating seizure from non-seizure activity. Therefore, one can speculate that the activity from the regions of the brain to which they correspond is related to the seizure activity, even if clear ictal activity only appears in a different region. During a seizure where ictal activity is limited to a certain region the brain, there can be physiological and subtle electrical changes in other brain regions, as well as clinical symptoms associated with impairment in the regions outside the focus [9, 8]. The small region where clear ictal activity first appears may not be the only significant participant in the genesis of a seizure, and seizure generation or propagation may be contingent on the state of activity in other regions of the brain.

Finally, we proposed an alternative feature extraction method that is tailored to a patient using the training data, and is designed to work with datasets limited to a few pre-selected channels. The feature extraction process involves the representation of an epoch of IEEG on a channel as a linear combination of vectors that we refer to as dictionary elements. To learn a dictionary, the method searches for directions in which the power in the two classes shows a large difference. The results from an evaluation of a detector that supplemented the original spectral energy features with features computed in a patient-specific manner show a significant improvement in 3 out of 5 patients. The features computed in a patient-specific way appear to complement the patient non-specific features, and this suggests that this is a promising area of investigation.

The patient-specific feature extraction methods we have discussed can be further refined, and further algorithm development may lead to more dramatic or consistent improvement in performance. Alternative dictionary learning methods could be considered, and more broadly, different approaches and implementations of patient-specific feature extraction. We hope this encourages the use of this and similar ideas (i.e., the introduction of learning in the feature extraction component) in other EEG-related applications, or more broadly, in learning-based solutions to other biomedical

problems.

Bibliography

[1] A. Aarabi, R. Fazel-Rezai, and Y. Aghakhani. A fuzzy rule-based system for epileptic seizure in intracranial eeg. *Clinical Neurophysiology*, 120:1648–1657, 2009.

[2] M. Aharon, M. Elad, and A. Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2005.

[3] M. Aharon, M. Elad, and A. Bruckstein. K-svd and its non-negative variant for dictionary design. In *Proceedings of SPIE*, volume 5914, pages 327–339, August 2005.

[4] A. M. Chan, F. T. Sun, E. H. Boto, and B. M. Wingeier. Automated seizure onset detection for accurate onset time determination in intracranial eeg. *Clinical Neurophysiology*, 119:2687–1696, 2008.

[5] S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *J. Sci Comp.*, 20(1):33–61, 1999.

[6] E. C. Chua, K. Patel, M. Fitzsimons, and C. J. Bleakley. Improved patient specific seizure detection during pre-surgical evaluation. *Clinical Neurophysiology*, 122:672–679, 2011.

[7] J. Echauz, S. Wong, O. Smart, A. Gardner, G. Worrell, and B. Litt. Computation applied to clinical epilepsy and antiepileptic devices. In I. Soltesz and K. Staley, editors, *Computational neuroscience in epilepsy*, pages 544–558. Academic Press, London, 2008.

[8] D. Englot, A. Mishra, P. Mansuripur, P. Herman, F. Hyder, and H. Blumenfeld. Remote effects of focal hippocampal seizures on the rat neocortex. *The Journal of Neuroscience*, 28(36):9066–9081, 2008.

[9] D. Englot, L. Yang, H. Hamid, N. Danielson, X. Bai, A. Marfeo, L. Yu, A. Gordon, M. Purcaro, J. Motelow, R. Agarwal, D. Ellens, J. Golomb, M. Shamy, H. Zhang, C. Carlson, W. Doyle, O. Devinsky, K. Vives, D. Spencer, S. Spencer, C. Schevon, H. Zaveri, and H. Blumenfeld. Impaired consciousness in temporal lobe seizures: role of cortical slow activity. *Brain*, 133:3764–3777, 2010.

[10] A. Gardner, A. Krieger, G. Vachtsevanos, and B. Litt. One-class novelty detection for seizure analysis from intracranial eeg. *Journal of Machine Learning Research*, 7:1025–1044, 2006.

[11] B. Gonzalez-Vellon, S. Sanei, and J. A. Chambers. Support vector machines for seizure detection. In *Proceedings of the 3rd IEEE International Symposium on Signal Processing and Information Technology*, pages 126–129, December 2003.

[12] S. Grewal and J. Gotman. An automatic warning system for epileptic seizures recorded on intracerebral eegs. *Clinical Neurophysiology*, 116:2460–2472, 2005.

[13] R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge Press, Cambridge, England, 1990.

[14] Leon Iasemidis. Epileptic seizure prediction and control. *IEEE Transactions on Biomedical Engineering*, 50(5):549–558, 2003.

[15] Z. Jiang, Z Lin, and L.S. Davis. Learning a discriminative dictionary for sparse coding via label consistent k-svd. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1697–1704.

[16] T. Joachims. Making large-scale svm learning practical. In B. Schlkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods: Support Vector Learning*, part 11, pages 169–184. MIT Press, Cambridge, MA, 1999.

[17] Y.U. Khan and J. Gotman. Wavelet based automatic seizure detection in intracerebral electroencephalogram. *Clinical Neurophysiology*, 114:898–908, 2003.

[18] T. Maiwald, M. Winterhalder, R. Aschenbrenner-Scheibe, H. U. Voss, A. Schulze-Bonhage, and J. Timmer. Comparison of three nonlinear seizure prediction methods by means of the seizure prediction characteristic. *Physica D*, 194:357–368, 2004.

[19] 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*, 25:119–131, 2008.

[20] P. Mirowski, D. Madhavan, Y. Lecun, and R. Kuzniecky. Classification of patterns of eeg synchronization for seizure prediction. *Clinical Neurophysiology*, 120:1927–1940, 2009.

[21] F. Mormann, R. Andrzejak, C. Elger, and K. Lehnertz. Seizure prediction: the long and winding road. *Brain*, 130:314–333, 2007.

[22] M. Morrell. Brain stimulation for epilepsy: can scheduled or responsive neurostimulation stop seizures? *Current Opinion in Neurology*, 19(2):164–168, April 2006.

[23] I. Osorio, M. Frei, and S. Wilkinson. Real-time automated detection and quantitative analysis of seizures and short-termprediction of clinical onset. *Epilepsia*, 39(6):615–627, 1998.

[24] I. Osorio, M. G. Frei, and J. Giftakis et al. Performance reassessment of a real-time seizure-detection algorithm on long ecog series. *Epilepsia*, 43(1):522–535, 2002.

[25] S. Pacia and J. Ebersole. Intracranial eeg substrates of scalp ictal patterns from temporal lobe foci. *Epilepsia*, 38(6):642–654, 1997.

[26] E. Shih and A. Shoeb snd J. Guttag. Sensor selection for energy-efficient ambulatory medical monitoring. In *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services*, 2009.

[27] A. Shoeb. *Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment.* PhD dissertation, Massachusetts Institute of Technology, 2009.

[28] A. Shoeb, D. Carlson, E. Panken, and T. Denison. A micropower support vector machine based seizure detection architecture for embedded medical devices. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4202–4205, Minneapolis, September 2009.

[29] A. Shoeb, H. Edwards, J. Connolly, B. Bourgeois, S. T. Treves, and J. Guttag. Patient-specific seizure onset detection. *Epilepsy Behavior*, 5:483–498, 2004.

[30] A. Shoeb and J. Guttag. Application of machine learning to epileptic seizure detection. In J. Furnkranz and T. Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 975–982, Haifa, Israel, June 2010.

[31] A. Shoeb, A. Kharbouch, J. Soegaard, S. Scachter, and J. Guttag. A machine-learning algorithm for detecting seizure termination in scalp eeg. *Epilepsy and Behavior*, 22, 2011.

[32] Ali Shoeb, H. Edwards, J. Connolly, B. Bourgeois, S.T. Treves, and J. Guttag. Patient-specific seizure onset detection. *Epilepsy and Behavior*, 5(4):483–498, 2004.

[33] A. Temko, E. Thomas, W. Marnane, G. Lightbody, and G. Boylan. Eeg-based neonatal seizure detection with support vector machines. *Clinical Neurophysiology*, 122:464–473, 2011.

[34] Q. Zhang and B Li. Discriminative k-svd for dictionary learning in face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2691–2698.

[35] Y. Zhang, G. Xu, J. Wang, and L. Liang. An automatic patient-specific seizure onset detection method in intracranial eeg based on incremental nonlinear dimensionality reduction. *Computers in Biology and Medicine*, 40:889–899, 2010.