Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment

by

Ali Hossam Shoeb

Bachelor of Science in Electrical Engineering and Computer Science, Massachusetts Institute of Technology (2003)
Master of Engineering in Electrical Engineering and Computer Science, Massachusetts Institute of Technology (2003)

Submitted to the Harvard-MIT Division of Health Sciences and Technology
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical and Medical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2009

© Massachusetts Institute of Technology 2009. All rights reserved.

Author .....
Harvard-MIT Division of Health Sciences and Technology
September 10, 2009

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

Accepted by .....
Ram Sasisekharan
Director of Harvard-MIT Division of Health Sciences and Technology

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
OCT 0 2 2009
LIBRARIES

ARCHIVES

Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment

by

Ali Hossam Shoeb

Submitted to the Harvard-MIT Division of Health Sciences and Technology
on September 10, 2009, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical and Medical Engineering

Abstract

Epilepsy is a chronic disorder of the central nervous system that predisposes individuals to experiencing recurrent seizures. It affects 3 million Americans and 50 million people world-wide.

A seizure is a transient aberration in the brain's electrical activity that produces disruptive physical symptoms such as a lapse in attention and memory, a sensory hallucination, or a whole-body convulsion. Approximately 1 out of every 3 individuals with epilepsy continues to experience frequent seizures despite treatment with multiple anti-epileptic drugs. These intractable seizures pose a serious risk of injury, limit the independence and mobility of an individual, and result in both social isolation and economic hardship.

This thesis presents novel technology intended to ease the burden of intractable seizures. At its heart is a method for computerized detection of seizure onset. The method uses machine learning to construct patient-specific classifiers that are capable of rapid, sensitive, and specific detection of seizure onset. The algorithm detects the onset of a seizure through analysis of the brain's electrical activity alone or in concert with other physiologic signals. When trained on 2 or more seizures and tested on 844 hours of continuous scalp EEG from 23 pediatric epilepsy patients, our algorithm detected 96% of 163 test seizures with a median detection delay of 3 seconds and a median false detection rate of 2 false detections per 24 hour period.

In this thesis we also discuss how our detector can be embedded within a low-power, implantable medical device to enable the delivery of just-in-time therapy that has the potential to either eliminate or attenuate the clinical symptoms associated with seizures.

Finally, we report on the in-hospital use of our detector to enable delay-sensitive therapeutic and diagnostic applications. We demonstrate the feasibility of using the algorithm to control the Vagus Nerve Stimulator (an implantable neurostimulator for the treatment of intractable seizures), and to initiate ictal SPECT (a functional neuroimaging modality useful for localizing the cerebral site of origin of a seizure).

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

Acknowledgments

I am grateful to the individuals that participated in the clinical studies described in this thesis. I would like to thank them for their patience, generosity, and courage. I would like to thank them for teaching me about living with epilepsy, for redefining my understanding of strength, and for encouraging me to continue my research.

Thanks to Drs. Blaise Bourgeois and S. Ted Treves for their mentorship while I conducted ictal SPECT studies at Children's Hospital Boston. Thanks to Jack Connolly and Herman Edwards for teaching me much of what I know about EEG.

Thanks to Dr. Steven Schachter for teaching me how to design clinical studies, and for exposing me to the clinical management of epilepsy. Thanks to Dr. Steven Schachter, Dr. Trudy Pang, and the nurses of the Beth Israel Deaconess Medical Center for making the vagus nerve stimulation studies possible. The prototype system used in these studies could not have been built without the help of Wayne Ryan and Ronald Wiken, and the input of professors Steven Leeb, Jeff Lang, Markus Zahn, and David Perreault.

Thanks to my thesis committee chair, Prof. Tomas Lozano-Perez, for technical insights that improved the performance of the seizure detection algorithm. Thanks also to my thesis reader, Dr. Sydney Cash, for introducing me to the fascinating world of intracranial EEG.

I am grateful to my thesis advisor Prof. John Guttag. John has helped me grow as an engineer, researcher, teacher, and person over the course of my undergraduate and graduate careers at MIT.

Thanks to Dorothy Curtis for helping me surpass the hurdles that often get in the way of a graduate student. Thanks for entertaining endless discussion of drug pumps, sewing machines, 9V batteries, and electromagnets.

Thanks to Zeeshan Syed and Asfandyar Qureshi for being there throughout both my undergraduate and graduate years at MIT. Thanks to Eugene Shih for always encouraging me to do better, and for finding ways to make long days in the lab fun. Thanks to Jenna Wiens and Anima Singh for cheering me on during the last days

of thesis writing, and for saying "you are going to be here forever". Thanks to Alaa Kharbouch, Zahi Karam, Georges Aoude, and Ram Srinivasan for always checking up on me, and for things too many to list. Thanks to James Geraci for being the best TA. Thanks to Thomas Heldt for being an inspiration. Thanks to Naveen Verma for an exciting technical collaboration that taught me so much. Thanks to Deema Arafah for her friendship, and for a memorable and comical discussion about how not to introduce this work.

Marwa, Fawkia, and Hossam, thanks for always believing in me.

Contents

1 Introduction **27**

1.1 Epilepsy 27

1.2 Seizure Detection Algorithms and Applications 29

1.2.1 Applications of Seizure Onset Detection 30

1.2.2 Applications of Seizure Event Detection 30

1.2.3 Application Dependant Performance 31

1.3 Why is Seizure Detection Challenging? 31

1.4 Related Work 32

1.4.1 Scalp EEG Seizure Event Detectors 33

1.4.2 Scalp EEG Seizure Onset Detectors 34

1.5 Thesis Contributions 35

1.6 Thesis Outline 38

2 Seizures and the Electroencephalogram **39**

2.1 Epileptic Seizures 39

2.2 Scalp Electroencephalogram 40

2.2.1 Seizures Within the Scalp Electroencephalogram 44

2.3 Intracranial Electroencephalogram 48

2.3.1 Seizures Within the Intracranial Electroencephalogram 50

2.4 Summary 53

3 Patient-Specific Seizure Onset Detection using the Scalp EEG **55**

3.1 Overview of Binary Classification 55

3.2 Feature Vector Design 56
3.2.1 Spectral Features 56
3.2.2 Spatial Features 60
3.2.3 Time Evolution 63
3.3 Feature Vector Classification 66
3.4 Patient-Specific Detector Architecture 70
3.5 Instantiating Detector Parameters 73
3.5.1 EEG Epoch Length: L 73
3.5.2 Number of EEG Channels: N 74
3.5.3 Number of filters: M 74
3.5.4 Number of feature vectors in $\mathcal{X}_{\mathcal{T}}$: W 74
3.5.5 SVM Parameters: γ, J, and C 75
3.5.6 Training Parameters: H, K, and S 75

4 Scalp EEG Data and Testing Methodology **77**
4.1 Scalp EEG Data Set 77
4.2 Performance Metrics 78
4.3 Performance Metric Measurement 78

5 Performance **81**
5.1 Patient-Specific Detector Performance 81
5.1.1 Latency 81
5.1.2 Sensitivity 82
5.1.3 Specificity 82
5.2 Varying Patient-Specific Detector Parameters 87
5.2.1 Varying the Number of Filters: M 87
5.2.2 Varying the Number of Feature Vectors in $\mathcal{X}_{\mathcal{T}}$: W 88
5.2.3 Varying the Number of Training Seizures: K 89
5.2.4 Varying Training Time into a Seizure: S 90
5.3 Patient-Specific and Non-specific Seizure Detection 92
5.3.1 Performance Comparison 92

5.4 Case Studies 95
5.4.1 Latency 95
5.4.2 Sensitivity 97
5.4.3 Specificity 98

6 Seizure Onset Detection Using Physiologic Signal Fusion 101
6.1 Why Use a Second Physiologic Signal? 101
6.2 The Electrocardiogram as a Second Signal 102
6.3 Patient-Specific, EEG-ECG-based Seizure Detection 102
6.4 Case Studies 104
6.4.1 Data 104
6.4.2 Case 1 104
6.4.3 Case 2 106
6.5 Importance of Patient-Specificity 111

7 Seizure-Triggered Vagus Nerve Stimulation 113
7.1 Vagus Nerve Stimulation 113
7.2 Methods 115
7.2.1 System Overview 115
7.2.2 Study Protocol 116
7.3 Case Studies 117
7.3.1 Patient A 117
7.3.2 Patient B 121
7.3.3 Patient C 125

8 Seizure-Triggered Single Photon Emission Computed Tomography 129
8.1 SPECT in Epilepsy 129
8.2 Methods 132
8.2.1 System Overview 132
8.2.2 Study Procotol 132
8.3 Results of Clinical Evaluation 134

9 Patient-Specific Seizure Onset Detection using iEEG **137**

9.1 Why Detect Seizure Onset Using iEEG? 137

9.2 Patient-Specific Detector Architecture 138

9.2.1 Cost of Support-Vector Machine Classification 139

9.3 Patient Non-specific Detector Architecture 142

9.4 iEEG Data Set 143

9.5 Performance Comparison 143

9.6 Reduced and Non-Reduced Support-Vector Machines 145

9.7 Case Studies 148

9.7.1 Latency 148

9.7.2 False Detections 149

9.8 Predicting Clinical Seizure Onset 149

9.9 Implementation 150

10 Conclusion and Future Work **153**

10.1 Goals and Contributions 153

10.2 Future Work 155

10.2.1 Feature Vector Enhancement 155

10.2.2 Detecting Seizure Cessation 155

10.2.3 Closed-Loop, Non-invasive Brain Stimulation 156

List of Figures

2-1 EEG electrodes arrayed symmetrically across the scalp provide a temporal and spatial summary of the synchronous firing of tens of millions of neurons within the brain. 41

2-2 Example of 10 seconds of awake EEG interrupted by an eye-blink at 37 seconds. The eye-blink results in a downward deflection on the EEG channels $\{FP1 - F7, FP1 - F3, FP2 - F4, FP2 - F8\}$. 42

2-3 Example of 10 seconds of sleep EEG interrupted by 11 Hz oscillations known as sleep spindles. The sleep spindles are most visible on the channel $FP2 - F4$ between 8-10 and 12-14 seconds. 43

2-4 Example of 10 seconds of awake EEG interrupted by the rhythmic, high-frequency activity associated with chewing. The high-frequency activity caused by chewing can be seen on the EEG channels $F7 - T7$ and $F8 - T8$. 43

2-5 Example of a seizure within the scalp EEG of Patient A. The seizure, which begins at 1723 seconds, involves flattening of the EEG signal across all channels followed by the appearance of a beta band rhythm on the channels $\{F3 - C3, C3 - P3\}$. 44

2-6 Example of a seizure within the scalp EEG of Patient B. The seizure, which begins at 6313 seconds, involves the appearance of a theta band rhythm on the channels $\{F7 - T7, T7 - P7\}$. 45

2-7 Example of an abnormal discharge within the scalp EEG of Patient A. The discharge, which occurs between 2884-2892 seconds, involves a repeating pattern of high-amplitude spikes followed by broad waves that can be seen on most EEG channels. 46

2-8 Example of abnormal rhythmic activity within the scalp EEG of Patient B. The rhythmic activity, which occurs between 6126-6130 seconds, involves a theta band rhythm on the frontal channels $\{FP1 - F3, F3 - C3\}$. 47

2-9 A second seizure within the scalp EEG of Patient A. The second seizure, which begins at 6210 seconds, resembles the seizure shown in Figure 2-5. 47

2-10 A second seizure within the scalp EEG of Patient B. The second seizure, which begins at 2381 seconds, resembles the seizure shown in Figure 2-6. 48

2-11 Example of a seizure within the scalp EEG of Patient C. The seizure, which begins at 1486 seconds, involves rapid eye-blinking followed by the appearance of a 3-4 Hz theta wave on the channel $T4 - T6$. . . . 49

2-12 Example of a seizure recorded from Patient C approximately 10 months after recording the seizure in Figure 2-11. 49

2-13 Example of seizures within the iEEG of patients D (top panel) and C (bottom panel). The seizure of Patient D begins at 5 seconds and consists of a few spikes that evolve into a high-amplitude spike train. The seizure of Patient E begins at 5 seconds and consists of a few spikes followed by low-amplitude, high-frequency activity. 51

2-14 The iEEG of Patient F contains rhythmic activity that is associated with a seizure (top panel) as well as rhythmic activity that is not associated with any clinical symptoms (bottom panel). 52

2-15 Example of two seizures within the iEEG of Patient D. Both seizures begin at 5 seconds and consist of a few spikes that evolve into a high-amplitude spike train. 52

3-1 Example of a seizure whose onset is associated with rhythmic activity that contains a mixture of prominent frequency components. Seizure onset follows 2994 seconds and mostly involves channels on the right side of the head (channels with even numerals). 57

3-2 Frequency spectrum of the channel $FP2-F4$ following the onset of the seizure illustrated in Figure 3-1. The spectrum contains large spectral components at 2, 5, and 11 Hz. 57

3-3 Superposition of the frequency spectra of an eye-blink and the rhythmic activity observed on the channel $FP2 - F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains an 11 Hz component that is absent from the spectrum of the eye-blink 58

3-4 Superposition of the frequency spectra of a sleep spindle and the rhythmic activity observed on the channel $FP2 - F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains low-frequency spectral components that are larger than those in the spindle spectrum. 59

3-5 Superposition of the frequency spectra of chewing and the rhythmic activity observed on the channel $FP2 - F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains less high-frequency content relative to the chewing spectrum. 59

3-6 M-band filterbank that measures the energy within the spectral components of an L second epoch taken from a single EEG channel. . . . 60

3-7 Energy within the frequency bands defined by the rightmost three filters (thick line) of the filterbank differentiates between the seizure spectrum (red) and chewing spectrum (blue) 61

3-8 The spatial distribution of EEG activity can be used to differentiate seizure from the non-seizure activity. The sleep spindles between 2989-2992 seconds do no involve activity on the channels $C4 - P4$ and $T8 - P8$. In contrast, the seizure activity following 2994 seconds involves both of these channels. 61

3-9 Formation of an intermediate feature vector X_T that captures the spectral and spatial properties of an epoch at time $t = T$ 62

3-10 Scalp EEG seizure involving a sequence of discrete spectral events. The first event, at 1723 seconds, is a spike involving all EEG channels. Next, all EEG channels exhibit a period of low-amplitude EEG. Finally, at 1725 seconds, a beta band rhythm that increases in amplitude and decreases in frequency appears on the channel $F3 - C3$. 64

3-11 Spectrogram of Channel $F3 - C3$ illustrating sequence of events that compose the onset and evolution of the seizure in Figure 3-10. 65

3-12 Spectrogram of channel $F3 - C3$ illustrating how the feature vector $\mathcal{X}_T$ models the transition from background to seizure onset. The feature vector X_T captures the period of low amplitude activity, X_{T-2} captures the spike, and X_{T-4} captures the background EEG activity preceding the spike. 65

3-13 Spectrogram of channel $F3 - C3$ illustrating how the feature vector $\mathcal{X}_T$ models the evolution of a seizure. The feature vector X_T captures the beta band rhythm, X_{T-2} captures the period of low amplitude activity, and X_{T-4} captures the spike. 66

3-14 Two-dimensional seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The first dimension of this feature space is defined by the energy in the spectral band 0-16 Hz and the second dimension corresponds to the energy in the 25 ± 11 Hz band. The seizure vectors extracted from the first 7 seconds of the seizure are numbered 1-5. 67

3-15 Linear decision boundary separating seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The decision boundary was determined using the SVM learning algorithm. 68

3-16 Nonlinear decision boundary separating seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The decision boundary was determined using a radial basis kernel and the SVM learning algorithm. 69

3-17 Patient-specific seizure onset detector architecture. 71

3-18 Generating training non-seizure feature vectors. 72

3-19 Generating training seizure feature vectors. 73

5-1 Percentage of 163 test seizures detected within a specified latency. The detector notes the onset of 50% of all test seizures within 3 seconds. . 82

5-2 Latency with which the detector notes the onset of seizures for each of the 23 test subjects. Each dot represents a seizure. A numeral is placed next to dots that represent more than one seizure. For most patients, the majority of seizures are detected within 5 seconds. . . . 83

5-3 Example of a seizure within the scalp EEG of Patient 15. The seizure, which begins at 272 seconds, consists of a theta band rhythm that is most prominently seen on the channel $T7 - P7$. 83

5-4 Example of another type of seizure observed within the scalp EEG of Patient 15. This seizure, which begins at 876 seconds, consists of a train of spikes on the channel $P7 - O1$. The detector fails to detect the onset of this seizure since its spectral and spatial characteristics differ from those of training seizures. Training seizures resemble the seizure shown in Figure 5-3. 84

5-5 Sensitivity of the patient-specific seizure detector. Red bars show the number of test seizures available for each subject, and black bars show the number of test seizures recognized by the detector. Overall, the detector recognized 96% of 163 test seizures. 84

5-6 Specificity of the patient-specific seizure detector. For most patients (18 of 23), the detector declares between 0 and 5 false detections per 24 hours. 85

5-7 Example of a seizure within the scalp EEG of Patient 13. The seizure, which begins at 640 seconds, consists of a delta band rhythm that is most prominent on the channel $FP1 - F7$. 86

5-8 Example of a short burst of rhythmic activity within the scalp EEG of Patient 13. The burst, which occurs between 233-235 seconds, consists of a delta band rhythm that is most prominent on the channel $FP1 - F7$. Since the burst resembles the seizure in Figure 5-7, the detector declares a false detection following the onset of the burst. 86

5-9 False detections declared by the detector are not uniformly distributed across test records. For Patient 13, 75% of the test records resulted in 0 false detections. For Patient 12, 68% of the test records resulted in 0 false detections. 87

5-10 The false detection rate of the detector decreases as number of filters used to construct X_T is increased from 2 to 8. A two-filter filterbank yields a false detection rate of 8 false detections per 24 hours. An eight-filter filterbank results in a false detection rate smaller than 4 false detections per 24 hours. 88

5-11 Increasing the number of feature vectors within $\mathcal{X}_T$ decreases the detector's false detection rate and increases its detection delay. To detect seizures with an average latency shorter than 5 seconds one should set $W = 3$. 89

5-12 Increasing the number of training seizures decreases the detector's miss rate and its detection delay. Including more than three training seizures results in a marginal improvement in detector's miss rate. 90

5-13 Including seizure vectors derived from the first $S > 12$ seconds of each training seizure does not improve seizure detection delay significantly. For $S > 12$ seconds the detector's mean latency is less than 3 seconds. 91

5-14 Including seizure vectors derived from the first $S > 18$ seconds of each training seizure does not improve seizure detection rate significantly. For $S > 18$ seconds the detector recognizes more than 95% of all test seizures. 91

5-15 Increasing S increases the detector's false detection rate. 92

5-16 Comparison of the number of test seizures recognized by the Reveal (sensitive setting) and patient-specific algorithms. Black bars show the number of test seizures available for each subject, red bars show the number of test seizures recognized by the patient-specific algorithm, and blue bars show the number of seizures recognized by the Reveal algorithm. The Reveal algorithm detected 74% of 152 test seizures and our patient-specific method detected 96% of all test seizures. 93

5-17 Comparison of the number of false detections declared by the Reveal (sensitive setting) and patient-specific algorithms. For some patients the Reveal algorithm declared an excess of 100 false detections per 24 hour period. 94

5-18 Comparison of the number of false detections declared by the Reveal (sensitive setting) and patient-specific algorithms. Expanded y-axis. . 94

5-19 Comparison of the number of test seizures recognized by the Reveal algorithm (specific setting) and our patient-specific algorithm. Black bars show the number of test seizures available for each subject, red bars show the number of test seizures recognized by the patient-specific algorithm, and blue bars show the number of seizures recognized by the Reveal algorithm. The Reveal algorithm detected 61% of 152 test seizures and our patient-specific method detected 96% of all test seizures. 95

5-20 Comparison of the number of false detections declared by the Reveal (specific setting) and patient-specific algorithms. For some patients the Reveal algorithm declared an excess of 100 false detections per 24 hour period. 96

5-21 Comparison of the number of false detections declared by the Reveal (specific setting) and patient-specific algorithms. Expanded y-axis. . . 96

5-22 Seizure recorded within the scalp EEG of Patient G. The seizure, which begins at 701 seconds, consists of beta band activity that is most prominent on the channels $\{F4 - C4, C4 - P4, F8 - T8, T8 - P8\}$. 97

5-23 Large amplitude theta band activity can be seen on all EEG channels twenty-nine seconds after the onset of the seizure shown in Figure 5-22. 98

5-24 Seizure recorded within the scalp EEG of Patient H. The seizure, which begins at 977 seconds, consists of rapid eye-blinking that is later accompanied by theta band activity on the channels $\{T4 - T6, T6 - O2\}$ at 983 seconds. 99

5-25 Seizure recorded within the scalp EEG of Patient J. The onset of this seizure invovles a generalized spike at 12231 seconds that is followed by generalized rhythmic activity. 100

5-26 Rhythmic activity commonly observed within the awake scalp EEG of Patient J. 100

6-1 Block diagram of patient-specific seizure onset detector that combines features extracted from the EEG and ECG signals. 103

6-2 Example of a seizure within the scalp EEG of Case 1. The seizure, which begins at 1486 seconds, involves rapid eye-blinking that results in downward deflections on frontal EEG channels (e.g. $\{FP1 - F3, FP2 - F4\}$). Coincident with the onset of rapid eye-blinking, the patient's heart rate accelerates as shown in Figure 6-3. Later, at 1492 seconds, a 3-4 Hz theta wave appears on the EEG channel $T4 - T6$. 105

6-3 Seizure onset, at 1486 seconds, is associated with an acceleration of the patient's heart rate. 106

6-4 Comparison of the minimum, maximum, and mean detection delays of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a shorter mean detection delay. . 107

6-5 Comparison of the false detection rates of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a smaller false detection rate. 107

6-6 Example of a seizure within the scalp EEG of Case 2. The seizure, which begins at 56 seconds, involves a 12 second period of low-amplitude EEG activity across most EEG channels. At the same time, the patient's heart rate accelerates as shown in Figure 6-7. Later, at 68 seconds, 1-2 Hz generalized, rhythmic activity develops. 108

6-7 Seizure onset, at 56 seconds, is associated with an acceleration of the patient's heart rate. 109

6-8 Comparison of the minimum, maximum, and mean detection delays of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The two detectors have comparable seizure detection delays. 110

6-9 Comparison of the false detection rates of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a smaller false detection rate. 110

7-1 Block diagram of closed-loop vagus nerve stimulation system. 116

7-2 Example of a seizure within the scalp EEG of Patient A. The seizure, which begins at 1486 seconds, involves rapid eye-blinking that results in downward deflections on frontal EEG channels (e.g. $\{FP1-F3, FP2-F4\}$). Coincident with the onset of rapid eye-blinking, the patient's heart rate accelerates as shown in Figure 7-3. Later, at 1492 seconds, a 3-4 Hz theta wave appears on the EEG channel $T4-T6$. 118

7-3 Seizure onset, at 1486 seconds, is associated with an acceleration of the patient's heart rate. 119

7-4 Block diagram of patient-specific seizure onset detector that combines features extracted from the EEG and ECG signals. 120

7-5 Initiation of VNS following computerized detection of the onset of a seizure from Patient A. Seizure onset begins with rapid eye blinking at 992 seconds. The detector declared seizure onset at 996 seconds, and initiated VNS in response. Beginning at 1002 seconds, a spike-train appears on the "VNS" channel confirming the initiation of vagus nerve stimulation. 120

7-6 Seizure onset, at 992 seconds, is associated with an acceleration of the patient's heart rate. 121

7-7 Typical EEG change associated with onset of Patient B's seizure. The seizure, which begins at 1340 seconds, involves an electrodecrement most prominent on the occipital channels $P3-O1$ and $P4-O2$. Next, at the 1348 seconds, a 3-5 Hz occipital rhythm emerges from the electrodecrement and rapidly generalizes. By 1360 seconds, muscle activity associated with the tonic-clonic phase of the seizure becomes visible. . 123

7-8 Block diagram of patient-specific seizure onset detector that uses features extracted from the EEG signal. 123

7-9 Initiation of VNS following computerized detection of the onset of a seizure from Patient B. The seizure begins with an electrodecrement at 1485 seconds. At 1501 seconds, the computerized system detected the muscle activity associated with the tonic-clonic phase of the seizure and initiated VNS in response. Evidence of VNS generator activity can be seen on the VNS channel at the conclusion of the seizure as shown in Figure 7-10. 124

7-10 The spike-train seen on the VNS channel between 1555-1565 seconds confirms the automatic initiation of vagus nerve stimulation following the onset of the seizure illustrated in Figure 7-9. 124

7-11 Example of a seizure within the scalp EEG of Patient C. The seizure, which begins at 248 seconds, involves a 14 second period of low-amplitude EEG activity across most EEG channels. Later, at 262 seconds, generalized, rhythmic activity develops. 126

7-12 Initiation of VNS following computerized detection of the onset of a seizure from Patient C. The seizure begins with an electrodecrement at 56 seconds. The computerized system noted seizure activity at 59 seconds and initiated vagus nerve stimulation in response. Evidence of VNS generator activity can be seen on the VNS channel starting at 69 seconds. 127

7-13 The spike-train seen on the VNS channel (first channel from the bottom), between 68-84 seconds, confirms the automatic initiation of vagus nerve stimulation following the onset of the seizure illustrated in Figure 7-12. 127

8-1 SPECT Image taken outside the seizure state (Interictal) and during the seizure state (Ictal I). During a seizure the seizure focus is hyperperfused and appears as a bright spot in a SPECT Image. The third image (Ictal II) is the result of subtracting the first two images. . . . 130

8-2 Delay associated with injecting the ictal SPECT radiotracer ranges between 10-130 seconds for 18 pediatric patients undergoing ictal SPECTs at Children's Hospital Boston. 131

8-3 Block diagram of system for automatic infusion of ictal SPECT radiotracer. 132

8-4 Protocol for evaluating automatic system for infusion of ictal SPECT radiotracer. See section 8.2.2 for details. 133

8-5 Comparison of delay between seizure onset and completion of radiotracer infusion by the clinical team (Group 1) and the computerized system (Group 2). 135

8-6 Latency and number of false detections declared by the computerized system during eight Ictal SPECT studies. 136

9-1 Block diagram of a patient-specific seizure detection algorithm that uses spectral and spatial features extracted from two iEEG channels. 139

9-2 Superposition of a nonlinear SVM boundary requiring $N_{SV} = 50$ terms (solid line) and an approximation of that boundary using $M_{SV} = 8$ terms (dashed line). 142

9-3 Block diagram of patient non-specific seizure detection algorithm. . . 143

9-4 Comparison of the detection latencies of the patient-specific and patient non-specific seizure detectors. The patient-specific detector that used a nonlinear SVM detected seizures within 8.5 ± 5.1 seconds. The same detector using a linear SVM detected seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm (with T=3) detected seizures within 10.5 ± 8.7 seconds. 144

9-5 Comparison of the number of false detections declared by the patient-specific and patient non-specific seizure detectors during 40 hours of non-seizure data. The patient-specific detector with a nonlinear SVM declared 19 false detections. The same detector using a linear SVM declared 28 false detections. The patient non-specific detector (with T=3) declared 126 false detections. 145

9-6 Comparison of the number of false detections declared by the patient-specific and patient non-specific seizure detectors during 40 hours of non-seizure data. The patient-specific detector with a nonlinear SVM declared 19 false detections. The same detector using a linear SVM declared 28 false detections. The patient non-specific detector (with T=10) declared 17 false detections. 146

9-7 Comparison of the detection latencies of the patient-specific and patient non-specific seizure detectors. The patient-specific detector that used a nonlinear SVM detected seizures within 8.5 ± 5.1 seconds. The same detector using a linear SVM detected seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm (with T=10) detected seizures within 18.7 ± 10.8 seconds. 146

9-8 Comparison of latency of detectors that use support-vector machines with different number of support-vectors 147

9-9 Comparison of number of false detections declared by detectors that use support-vector machines with different number of support-vectors 147

9-10 Example of a seizure within the iEEG of Patient 6. The seizure, which begins at 2 seconds, consists of a few spikes that evolve into a high-amplitude spike train. The patient-specific detector recognized the seizure at 16 seconds, while the patient non-specific detector did so at 40 seconds. 148

9-11 Example of a seizure within the iEEG of Patient 5. The seizure, which begins at 6 seconds, consists of a high-frequency rhythm that increases in amplitude and decreases in frequency as the seizure progresses. . . 149

9-12 Example of a burst of rhythmic non-seizure activity within the iEEG of Patient 5. The patient non-specific detector declared this burst as a seizure event even though its spatial and spectral character differs significantly from that of the seizure shown in Figure 9-11. 150

9-13 Detection delay relative to the clinical onset of a seizure. An alarm based on detecting the electrical onset of a seizure using our patient-specific method could provide some patients, such as patients 3 and 6, enough time to prepare for the clinical onset of a seizure. 151

9-14 Implementation of our machine-learning based, patient-specific detector on the hardware described in [2]. An analog front-end processes two iEEG channels, and for each channel, extracts the spectral power within two configurable frequency bands. A digital back-end samples the analog power profile, assembles these samples into a feature vector, and classifies the feature vector using a support-vector machine. . . . 151

List of Tables

3.1 Patient-specific Detector Parameters 74

Chapter 1

Introduction

1.1 Epilepsy

Epilepsy is a chronic disorder of the central nervous system that predisposes individuals to experiencing recurrent *seizures*. A seizure is a sudden, transient aberration in the brain's electrical activity that produces disruptive symptoms. These symptoms range between a lapse in attention, a sensory hallucination, or a whole-body convulsion. In the book "Brainstorms: Epilepsy in Our Words" [49], individuals with epilepsy describe what it is like to have a seizure:

> *I only experience my seizures as I am falling off to sleep...my symptoms include a shock-like feeling inside my head, a twitch of one or more limbs, a shock that makes the trunk of my body jump.*
>
> *I experience a combination of deja vu with extreme fear. Nothing I do takes me out of the deja vu...the general feeling is of being in front of an oncoming train with no way to escape.*

Epilepsy is not a single disease, but a family of syndromes that share the feature of recurrent seizures. Epilepsy may develop as a result of inheriting a mutation in a molecular mechanism that regulates neuron behavior, migration, or organization. Alternatively, it may develop as a result of brain trauma such as a severe blow to the head, a stroke, a cerebral infection, or a brain malignancy [8].

Fifty million people world-wide are diagnosed with epilepsy. In the United States epilepsy affects 3 million people and is the third most common neurologic disorder after Alzheimer's disease and stroke. In an unfortunate subset of 1.2 million individuals, frequent, unpredictable seizures persist despite treatment by one or multiple anti-epileptic drugs. These types of seizures are known as medically intractable seizures.

> *The worst part of having seizures is knowing that they can happen anytime and even though drugs control mine most of the time they occasionally break through.*

Medically intractable seizures severely limit the independence and mobility of an individual and, as a consequence, can result in social isolation and economic hardship. Most concerning is that refractory seizures significantly increase an individual's chance of experiencing burns, lacerations, skull fractures, and even sudden unexpected death [16, 17].

> *I knew, also instantaneously, that I must keep this affliction secret, and that this condition and the dark secret of it set me apart from others. I was to live my life feeling this isolation and separation.*

The negative influence of uncontrolled seizures extends beyond the individual to affect their family members, friends, and the whole of society. The families and friends of people with epilepsy experience chronic anxiety and rearrange their lives to ensure the safety of their loved one. Society incurs an annual loss of 12.5 billion dollars in health care costs and losses in productivity [6]. There is a need for novel therapies that better control seizures as well as technology that helps both the individual and their family to cope with the consequences of seizures.

Computerized seizure onset detection will enable the engineering of novel therapeutic and alerting systems that may ease the burden of intractable seizures. A therapeutic system capable of detecting and reacting to the onset of a seizure could administer a local electrical [60], thermal [47], or neurochemical [69] stimulus that

halts the progression of a seizure prior to the development of clinical symptoms. Moreover, just-in-time, local therapy could relieve patients of the toxic side-effects that accompany systemic administration of multiple anti-epileptic drugs. An alerting system equipped with seizure onset detection could warn the patient of the seizure prior to the development of debilitating symptoms, or could notify a family member so that the consequences of a seizure are limited. Knowledge that a reliable warning will be issued rapidly following seizure onset may restore within individuals the confidence to overcome the limits on life that accompany seizures.

In this thesis, we describe a patient-specific algorithm capable of rapidly detecting seizure onset through analysis of the brain's electrical activity alone or in concert with other physiologic signals. Moreover, we demonstrate the feasibility of using our detector to control the Vagus Nerve Stimulator (an implantable neurostimulator for the treatment of intractable seizures), and to initiate ictal SPECT (a functional neuroimaging modality useful for localizing the cerebral site of origin of a seizure).

In this chapter, we describe different types of seizure detectors and their potential role in diagnosistic, therapeutic, and alerting applications. We also highlight the challenges associated with the computerized detection of seizures using the brain's electrical activity. Finally, we review the methodology and performance of previously published algorithms and highlight how our approach contributes to the field of seizure detection.

1.2 Seizure Detection Algorithms and Applications

A seizure detector can be classified as either a *seizure onset detector* or as a *seizure event detector*. The purpose of a seizure onset detector is to recognize that a seizure has started with the shortest possible delay, but not necessarily with the highest possible accuracy. In contrast, the purpose of a seizure event detector is to identify seizures with the greatest possible accuracy, but not necessarily with the shortest delay. Seizure onset detectors are suited for applications requiring a rapid response to a seizure, while seizure event detectors are suited for applications requiring an accu-

rate account of seizure activity over a period of time. The following sections present applications of both detector types and discuss how the specifics of an application dictate how to balance detection delay against accuracy of detection.

1.2.1 Applications of Seizure Onset Detection

Computerized seizure onset detection can facilitate the initiation of delay-sensitive diagnostic, therapeutic, and alerting procedures. Within the realm of diagnosis, seizure onset detection could be used to quickly initiate functional neuroimaging studies designed to localize the cerebral origin of a seizure. In this case, rapid initiation is important since the accuracy of such imaging studies diminishes the greater the delay between seizure onset and infusion of the imaging radiotracer [10]. Within the realm of therapy, seizure onset detection could be used to trigger neurostimulators designed to affect the progression of a seizure [60]. In this application, rapid initiation is important since the likelihood of affecting a seizure seems to decrease the longer the delay between the onset of a seizure and the start of stimulation [22]. Finally, within the realm of alerting, seizure onset detection could prompt a patient or care provider to ensure safety or administer a fast-acting anticonvulsant. In this scenario, waiting too long to alert the patient or care provider increases the chances that symptoms of the seizure will leave them unable to respond.

1.2.2 Applications of Seizure Event Detection

Computerized seizure event detectors can enable physicians to better titrate therapy (pharmacologic or otherwise) over time. Currently, anti-seizure therapy is dispensed based on a individual's account of the number of and severity of the seizures they experienced between clinic visits. Unfortunately, too many patients produce an inaccurate tally [24], which may lead a physician to prescribe too much, too little, or the incorrect medication. Prescribing too much medication results in toxic side-effects while prescribing too little means a continuation of frequent seizures. A wearable device capable of computerized seizure event detection within the ambulatory setting

could provide physicians with a summary of the number, frequency, duration and time of day an individual experiences seizures. By correlating this information with different medication regimens a physician could more quickly converge on a treatment plan that maximally benefits the individual.

1.2.3 Application Dependant Performance

For a seizure onset detector the speed with which a seizure is recognized can be increased at the expense of seizure detection accuracy. The degree to which one favors detection speed or accuracy is dictated by the problem to be solved. For instance, an application such as automatically infusing a radiotracer for the purpose of a neuroimaging study requires excellent detection accuracy even if it is at the expense of latency. In contrast, applications such as automatically initiating neurostimulation call for emphasizing detection speed over accuracy because of the benign effect of delivering many types of neurostimulation outside of the seizure state [60].

1.3 Why is Seizure Detection Challenging?

Seizure onset and event detection is most often accomplished through analysis of the *Electroencephalogram* (EEG). The EEG is a multichannel recording of the electrical activity generated by collections of neurons within the brain; different channels reflect the activity within different brain regions. When the EEG is measured using non-invasive electrodes arrayed on an individual's scalp it is referred to as *scalp EEG*; and when it is measured using electrodes placed on the surface of the brain or within its depths it is referred to as intracranial EEG.

The property of scalp and intracranial EEG that most complicates the seizure detection task is its *variability* across individuals with epilepsy [18], both in the seizure and non-seizure states. Typically, following the onset of a seizure, a set of EEG channels develops rhythmic activity that reflects underlying neuronal hypersynchrony. Both the location of the involved EEG channels as well as the spectral content of the rhythmic activity varies across individuals. Furthermore, the EEG signature of one

patient's seizure may closely resemble the signature of abnormal, non-seizure EEG gathered from the same patient or from a different patient [46].

Within the scalp EEG the seizure detection task is further complicated by the physical properties of the signal. The scalp EEG is most sensitive to the activity of neurons on the brain surface; consequently, the activity of neurons within deep brain structures has almost no influence on the scalp EEG. When the epileptic neural network is deep within the brain, the scalp EEG may reflect physical sequelae of the seizure, such as repetitive eye-blinks (eye flutter) or muscle contractions, before reflecting hypersynchronous neural activity. Seizures of this type are difficult to detect with high specificity and low latency since activity such as eye flutter and muscle contractions are routinely observed as an individual partakes in the activities of daily life.

Another property of scalp EEG that makes seizure detection challenging is its susceptibility to contamination by non-physiologic sources. The sway of EEG electrode cables, alterations in the electrode-skin interface, and the coupling of AC harmonics from electric machinery can all produce spectral changes that affect the performance of a seizure detector.

1.4 Related Work

Research into seizure detection methods began with the development of seizure event detectors [18]. The detectors developed were meant to detect the seizures of any individual with epilepsy, i.e., they were *patient non-specific*. The variability within EEG, along with the challenges discussed in Section 1.3, severely limited the detection accuracy of these patient non-specific detectors. To improve performance, investigators developed *patient-specific* event detectors, i.e., detectors that could be tailored to the EEG of an individual [44]. These detectors exhibited improved performance because seizure and non-seizure EEG recorded from an individual exhibits less variability as shown in Chapter 3. Years later, the development of diagnostic and therapeutic applications that require initiation following seizure onset motivated the development

of seizure onset detection algorithms.

1.4.1 Scalp EEG Seizure Event Detectors

One of the earliest patient non-specific seizure event detectors was the one developed by Gotman [18] in 1982. The Gotman algorithm searches for the hallmark sign of seizures: sustained rhythmic activity. The algorithm sequentially searches a number of EEG channels for the presence of rhythmic activity with a dominant frequency between 3-20 Hz and an amplitude at least 3 times greater than that of a background window; whenever the degree of rhythmicity exceeds a threshold on at least two channels and persists for 4 seconds a seizure is declared.

The Gotman algorithm successfully detects seizures whose evolution includes sustained rhythmic activity with a fundamental below 20 Hz; it is not successful in detecting seizures consisting of EEG containing a mixture of frequencies or those with low amplitude high frequency activity. Since the scalp EEG of individuals with epilepsy contains pathologic, normal, and artifact-induced bursts of rhythmic activity, a significant fraction of detections produced by the Gotman algorithm are not associated with seizures [18]. A recent evaluation of the Gotman algorithm on 652 hours of scalp EEG that included 126 seizures from 28 patients [48] demonstrated that this approach detects 50% of test seizures and declares 0.5 false detections per hour.

Since Gotman's work, investigators have developed seizure event detectors that utilize more sophisticated signal processing to characterize the rhythmicity associated with seizures as well as more sophisticated schemes for determining whether that activity is representative of an ongoing seizure. An example of such efforts is the Reveal seizure detector developed by Wilson [64]. The Reveal algorithm decomposes 2 second EEG epochs from each input channel into time-frequency atoms using the Matching Pursuit algorithm. Reveal then employs hand-coded and neural network rules to determine whether features derived from the atoms of a channel are consistent with a seizure taking place on that channel. The thresholds for some of the neural network rules are determined using both archetypal seizures from individuals with

epilepsy and background EEG from individuals without epilepsy.

In [64] Wilson reported that the Reveal algorithm detected 76% of 672 seizures gathered from 426 individuals with epilepsy, and that it declared false detections at a rate of 0.11 false detections per hour when tested on data from individuals without epilepsy. When the Reveal algorithm is made patient-specific [65, 66], they reported that it improved the specificity of the original algorithm from 0.62 false-detection per hour to 0.34 false detections per hour while improving the sensitivity to 78%. The Reveal algorithm has a sensitivity superior to that of the classic Gotman algorithm, but as we show in Chapter 3, it has poor specificity when processing the scalp EEG of patients with abnormal, non-seizure rhythmic activity.

1.4.2 Scalp EEG Seizure Onset Detectors

Saab developed a patient non-specific seizure onset detector [48]. Saab's algorithm uses features derived from a wavelet decomposition of each EEG channel to estimate the probability of a seizure event. Whenever the probability exceeds a user defined threshold for a given period of time, the algorithm declares the onset of a seizure. When evaluated on 652 hours of scalp EEG that included 126 seizures from 28 patients, Saab's algorithm detected 78% of seizures with a median detection latency of 9.8 seconds and a false detection rate of 0.86 false detections per hour. Saab reported that missed seizures included those with onsets characterized by focal activity, mixed frequencies, or short duration; and that false detections were mainly caused by short bursts of rhythmic activity, rapid eye blinking, and chewing.

Qu developed the first patient-specific seizure onset detection algorithm [43, 44, 45, 46]. Qu's patient-specific algorithm relies on a nearest-neighbor classifier to assign a list of features, or feature vector, to the seizure or non-seizure class. The classifier is trained on seizure and non-seizure feature vectors from a single individual. The feature vector consists of measures of the EEG's average amplitude, dominant frequency, and rhythmicity. The classifier sequentially classifies feature vectors derived from the available EEG channels, and declares a seizure if the set of positively classified channels matches half of those chosen by an expert. When tested by Qu on 29.7 hours and

47 seizures from 12 patients, the method detected 100% of seizures with an average delay of 9.35 seconds and a false alarm rate of 0.03 false detections per hour. The non-seizure EEG that Qu used to calculate the false detection rate of his algorithm was formed by concatenating segments of EEG extracted at regular intervals from several days of data. When compared to Saab's work, Qu's work illustrates that a patient-specific approach can result in improved sensitivity and specificity, but not necessarily an improvement in detection latency.

Meier developed a seizure onset detection system that is patient non-specific but is seizure-specific [35]. Meier grouped seizures in a database into 6 categories based on the frequency of the dominant rhythm that appears following seizure onset. He then trained a set of support-vector machines, one for each seizure type, to determine whether a feature vector extracted from an EEG epoch is consistent with one of the seizure types. Rather than extract and then classify feature vectors from one channel after another, Meier extracts a single feature vector that includes the average, across channels, of signal properties such as the number of zero-crossings, wavelet coefficient power, and cross-correlation. When evaluated on 91 seizures and 1,360 hours of non-seizure EEG from 57 patients, Meier's algorithm detected 96% of the test seizures with an average detection delay of 1.6 seconds and false alarm rate of 0.45 false detections per hour. Meier's approach depends on the test seizure being a member of one of the 6 defined categories as well as it being recorded using the same number and position of channels used to record the training seizures. Seizures whose onsets lack the development of rhythmic activity and instead reflect physical sequelae of the seizure, such as eye-flutter, do not fall within the defined categories; consequently, such seizures will be detected later or not at all.

1.5 Thesis Contributions

In this thesis we develop a patient-specific algorithm for detecting seizure onset that improves technically upon existing methods in the following ways:

- **Enhanced Performance:** When trained on 2 or more seizures from the same

patient and tested on 844 hours of continuous scalp EEG from 23 pediatric subjects, our algorithm detected 96% of 163 test seizures with a median detection delay of 3 seconds (average 4.6 seconds) and a median false detection rate of 0.07 false detections per hour (average 0.13 false detection per hour). When evaluated on the same data set, the Reveal algorithm [64] detected 61% of seizures with a false detection rate of 1.6 false detections per hour.

Relative to Saab's algorithm [48], our algorithm exhibits a shorter detection delay, a higher seizure detection rate, and a lower false detection rate.

Relative to Qu [46], our detector exhibits a shorter detection delay, comparable seizure detection rate, and a higher false detection rate. Unlike Qu's algorithm, ours does not require an expert to identify relevant EEG channels. Furthermore, to estimate the seizure detection rate of our method, we used a larger data set than that used by Qu (163 seizures from 23 patients vs 47 seizures from 12 patients). We also used all 844 hours of non-seizure EEG in order to reliably estimate the false detection rate of our algorithm. In contrast, Qu tested his algorithm on 29.7 hours of non-seizure EEG formed by concatenating EEG segments extracted at regular intervals from several days of data.

Relative to Meier [35], our detector has a longer detection delay, equal seizure detection rate, and a lower false detection rate. However, unlike Meier's algorithm, ours is not restricted to detecting a defined set of seizure types from a particular population of individuals with epilepsy.

It is important to note that each of Saab, Qu, and Meier used different data sets from our own when evaluating the performance of their algorithms.

- **Uses Multiple Physiologic Signals:** Our method can automatically learn how to use multiple sources of physiologic information to detect seizures whenever the scalp EEG alone is unreliable. This capability is important for the detection of seizures whose onsets lack the development of rhythmic activity and instead reflect physical sequelae of the seizure such as eye-flutter, muscle contractions, or changes in heart rate. Such seizures are not easily grouped into

the categories defined by Meier [35] and therefore will be detected later or not at all.

- **Minimal User Intervention:** Our algorithm does not require a user to define the values of key algorithm parameters. In our approach, a user only needs to define the onset of activity associated with a seizure in a set of physiologic signals. The relationships between these signals that distinguish the seizure and non-seizure periods are automatically learned.

- **Suitable for Implantable Medical Devices:** Our algorithm can be adapted for the detection of seizure onset within intracranial EEG, and can be implemented on the low-power hardware of an implantable neurostimulator. When evaluated on 81 hours of intracranial EEG containing 61 seizures and gathered from 17 adult subjects, our algorithm detected 60/61 seizures within 9.3 seconds, declared a total of 28 false detections, and consumed 12μA when implemented on the hardware in [2]. When evaluated on the same data set and hardware, a patient non-specific algorithm based on [38] detected 41/61 seizures within 18.7 seconds, declared a total of 17 false detections, and consumed 32μA.

In this thesis we also discuss the application of our patient-specific algorithm to the following delay-sensitive therapeutic and diagnostic applications:

- **Non-invasive Closed-Loop Control of the Vagus Nerve Stimulator:** Using our seizure onset detector we designed and clinically evaluated the first non-invasive system that initiates Vagus Nerve Stimulation (VNS) in response to detecting the onset of a seizure using multiple physiologic signals [57]. As an example of the system's capabilities, during an 81 hour clinical test of the system on a patient, the computerized system detected 5/5 seizures and initiated VNS within 5 seconds of the appearance of ictal discharges in the EEG.

- **Computerized Initiation of Ictal SPECT Studies:** Using our seizure onset detector we designed and clinically evaluated a system for initiating a functional neuroimaging study following seizure onset. The neuroimaging modality, ictal

SPECT, is used to radiographically localize the cerebral origin of a seizure. Our system could initiate injection of the radiotracer used for ictal SPECT within 19.3 ± 2.3 seconds in 8/8 prospective trials, while the clinical team required 27.7 ± 8.5 seconds, and failed to initiate ictal SPECT in one of the trials.

1.6 Thesis Outline

This thesis is organized as follows: In Chapter 2 we provide background material on seizures and both the scalp and intracranial electroencephalograms. Next, in Chapters 3-5, we develop and analyze the performance of a machine-learning based, patient-specific algorithm for the detection of seizure onsets within scalp EEG. In Chapter 6, we illustrate how the detector in Chapter 3 can be extended with information from other physiological signals in order to detect seizures whose onset does not immediately involve the development of rhythmic activity within the scalp EEG.

We then present applications of our seizure detection methodology. In Chapter 7 we illustrate how the methods developed in Chapters 3 and 6 were integrated into a real-time system that initiates vagus nerve stimulation in response to detecting the onset of a seizure. Chapter 8 illustrates how the method developed in Chapter 3 was used in a real-time system that infuses the radiotracer used in ictal SPECT following seizure onset detection. Finally, in Chapter 9, we adapt the feature extraction and classification stages of our detector so that it may be embedded within a low-power, implantable medical device.

Chapter 2

Seizures and the Electroencephalogram

This chapter reviews the pathophysiology underlying seizures as well as their clinical manifestation and categorization. This chapter also reviews properties of the electroencephalogram and the electrographic characteristics of seizures.

2.1 Epileptic Seizures

Neurons are cells within the brain capable of generating, propagating, and processing electric signals. Neurons connect to other neurons in order to form functional networks, and the brain can be viewed as a collection of interacting neural networks. The inputs to a neural network can be excitatory or inhibitory. Excitatory inputs promote activity among neurons within a network and inhibitory inputs suppress it [30].

Epileptic seizures are transient periods involving the hyperactivity and hypersynchronization of a large number of neurons within one or more neural networks. These transient states arise because of a perturbation that creates an imbalance favoring the excitation of a neural network over its inhibition. The imbalance may arise because of defects within a neuron, such as an ion channel dysfunction; defects in connections between neurons, such as deficient inhibitory neurotransmitter synthesis; or defects in

neural network organization, such as the formation of aberrant excitatory connections between neurons. Defects within neurons, neuronal connections, or neural network orginization may result from a genetic disorder or from trauma to the central nervous system during life.

Epileptic seizures are broadly classified according to their cerebral site of origin and spread. *Focal* seizures arise from a localized region of the brain's cortex and have clinical manifestations that reflect that region of the brain. As an example, a focal seizure originating in the temporal lobe, the part of the brain that processes emotions and short-term memory, may result in feelings such as euphoria, fear, and deja vu or hallucinations of taste or smell. Focal seizures may spread to involve other regions of the brain or the entire brain. As an example, a seizure originating in the left motor cortex may result in jerking movements of the right upper extremity. When the seizure spreads to adjacent areas and then the entire brain, whole-body convulsions ensue.

Generalized seizures begin with abnormal electrical activity that appears to encompass the entire cerebral cortex. The manifestations of such widespread abnormal electrical activity often includes the loss of consciousness. Motor manifestation of these seizures may include whole-body rigidity and jerking (generalized tonic-clonic seizure) or whole-body loss of muscle tone (atonic seizure). A seizure that begins focally and then generalizes is referred to as a *secondarily generalized* seizure.

2.2 Scalp Electroencephalogram

The scalp electroencephalogram (scalp EEG) is a non-invasive measure of the electrical potentials generated by the activity of tens of millions of neurons within the brain. The scalp EEG is usually measured through electrodes that are symmetrically arrayed on the scalp as shown in Figure 2-1. An EEG signal, or channel, is formed by taking the difference between potentials measured at two electrodes. For example, the channel $FP1 - F7$ is formed by taking the difference between the potentials measured at the electrodes $FP1$ and $F7$. Each EEG channel summarizes activity

localized within a region of the brain; for instance, the channel $FP1 - F7$ reflects neural activity originating within the frontal lobe of the left hemisphere. The onset of a focal seizure involves a change in activity on the few scalp EEG channels that lie above or near the site of the brain giving rise to a seizure; on the other hand, the onset of a generalized seizure involves activity on all scalp EEG channels.

The physics of EEG generation constrains both the origin and characteristics of neural activity visible within the scalp EEG. In particular, the neurons that contribute the most to the scalp EEG are those closest to the scalp surface; in contrast, the activity of neurons buried within deep brain structures is not observable. Furthermore, the cerebrospinal fluid and skull surrounding the brain act as attenuators that greatly diminish the amplitude of higher frequency neural oscillations. An important consequence of these physical limitations is that certain types of seizures, namely those involving a small, deep region within the brain cannot be observed using the scalp EEG.

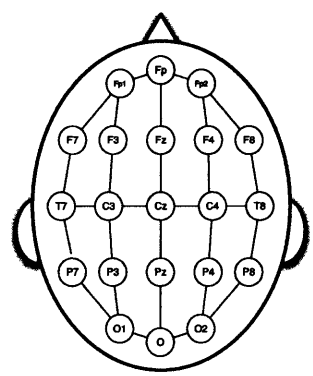

Figure 2-1: EEG electrodes arrayed symmetrically across the scalp provide a temporal and spatial summary of the synchronous firing of tens of millions of neurons within the brain.

Electroencephalographers describe EEG activity in terms of its spatial distribution on the scalp (frontal, posterior, lateral, and bilateral) as well as its dominant frequency component. An EEG wave is classified as having a *delta* component if its dominant frequency component f is ≤ 4 Hz, a *theta* component if $4 < f < 8$ Hz, an *alpha* component when $8 \leq f < 12$ Hz, a *beta* component when $12 \leq f < 30$ Hz, or a *gamma* component when $f \geq 30$ Hz. For example, the *alpha wave* in EEG parlance

refers to a 10 Hz rhythm that appears most prominently on posterior channels when a subject closes their eyes and relaxes.

Scalp EEG activity is modulated by the state of vigilance of an individual. In particular, the dominant frequency and spatial distribution of EEG activity during the awake state is different than that during sleep. As an example, Figure 2-2 illustrates awake EEG activity interrupted by an eye-blink at 37 seconds. The awake EEG background is primarily composed of low frequency activity. The eye-blink results in a downward deflection of the EEG signal on the channels $\{FP1-F7, FP1-F3, FP2-F4, FP2-F8\}$. Figure 2-3 illustrates EEG activity recorded during sleep. The 11 Hz oscillation, observed most prominently on the channel $FP2-F4$ between 12-14 seconds, is known as a *sleep spindle.*

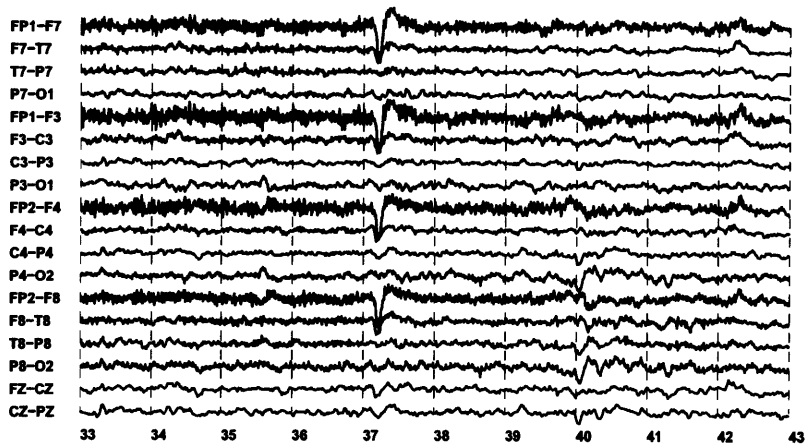

Figure 2-2: Example of 10 seconds of awake EEG interrupted by an eye-blink at 37 seconds. The eye-blink results in a downward deflection on the EEG channels $\{FP1-F7, FP1-F3, FP2-F4, FP2-F8\}$.

Scalp EEG is also easily corrupted by both physiological and non-physiological artifacts. Physiological artifacts include sweat, chewing, eye-blinks, and scalp muscle contractions. Non-physiological artifacts include power-line noise and EEG cable motion. As an example, Figure 2-4 illustrates how chewing affects the scalp EEG. Chewing results in rhythmic, high-frequency activity that is most prominently observed on channels on either side of the scalp (e.g. $F7-T7$ and $F8-T8$).

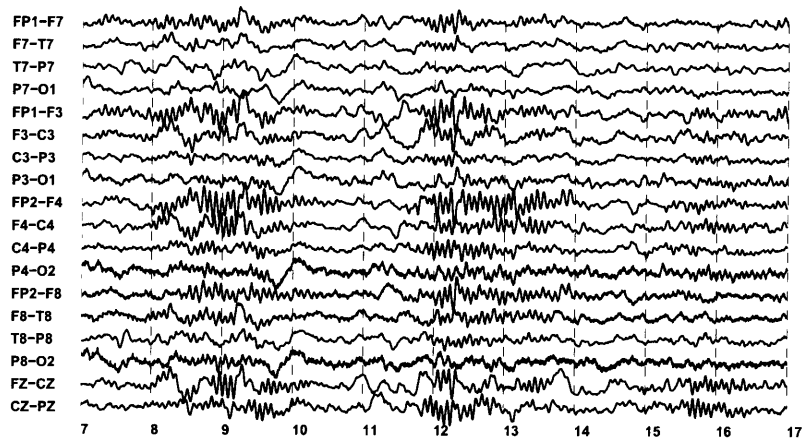

Figure 2-3: Example of 10 seconds of sleep EEG interrupted by 11 Hz oscillations known as sleep spindles. The sleep spindles are most visible on the channel $FP2 - F4$ between 8-10 and 12-14 seconds.

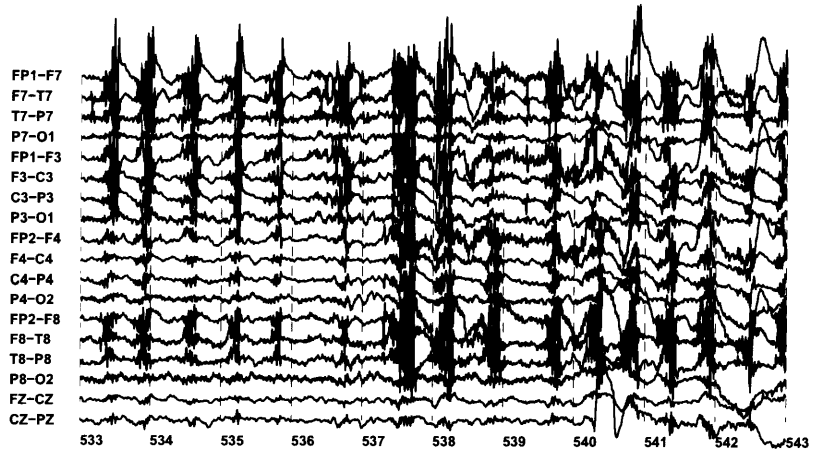

Figure 2-4: Example of 10 seconds of awake EEG interrupted by the rhythmic, high-frequency activity associated with chewing. The high-frequency activity caused by chewing can be seen on the EEG channels $F7 - T7$ and $F8 - T8$.

2.2.1 Seizures Within the Scalp Electroencephalogram

Within the scalp EEG, seizures manifest as a sudden redistribution of spectral energy on a set of EEG channels. The spectral energy redistribution is caused by hypersynchrony of neurons within an epileptic neural network, and consists of an appearance or disappearance of frequency components within the 0-25 Hz band [19]. However, which spectral components vanish or rise to prominence varies across patients. Furthermore, the EEG channels demonstrating the spectral energy change also varies across patients since it is a function of the cerebral site of origin of a seizure.

As an example, Figure 2-5 and Figure 2-6 illustrate seizures from patients A and B respectively. The seizure in Figure 2-5 begins at 1723 seconds and consists of flattening of the EEG signal across all channels followed by the appearance of a beta band rhythm on the channels $\{F3 - C3, C3 - P3\}$. Then, over the course of a few seconds, the amplitude of this rhythm increases as its frequency decreases and settles within the theta band.

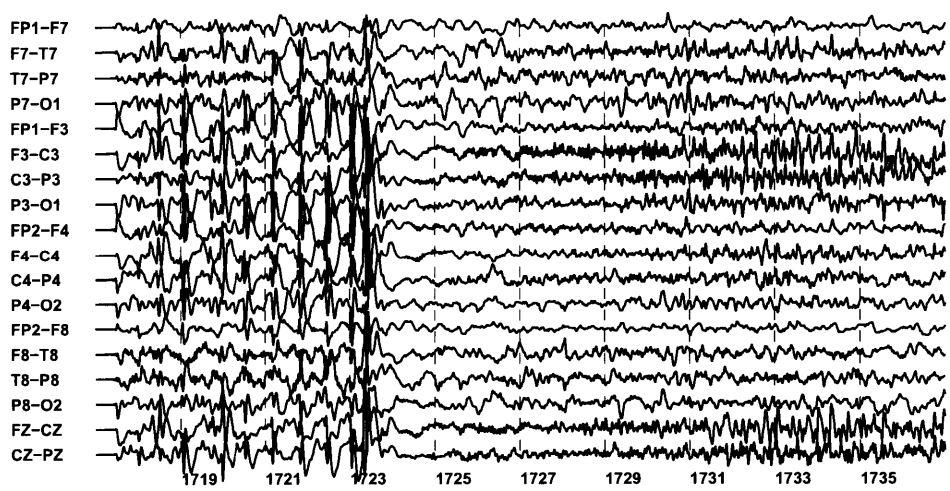

Figure 2-5: Example of a seizure within the scalp EEG of Patient A. The seizure, which begins at 1723 seconds, involves flattening of the EEG signal across all channels followed by the appearance of a beta band rhythm on the channels $\{F3-C3, C3-P3\}$.

The seizure in Figure 2-6 begins at 6313 seconds with the onset of a theta band

rhythm that is most prominent on the channels $\{F7 - T7, T7 - P7\}$. Other EEG channels also exhibit a change following seizure onset. The channel $\{C3 - P3\}$ develops a theta band rhythm while the channel $\{FP2 - F8\}$ develops a delta band rhythm.

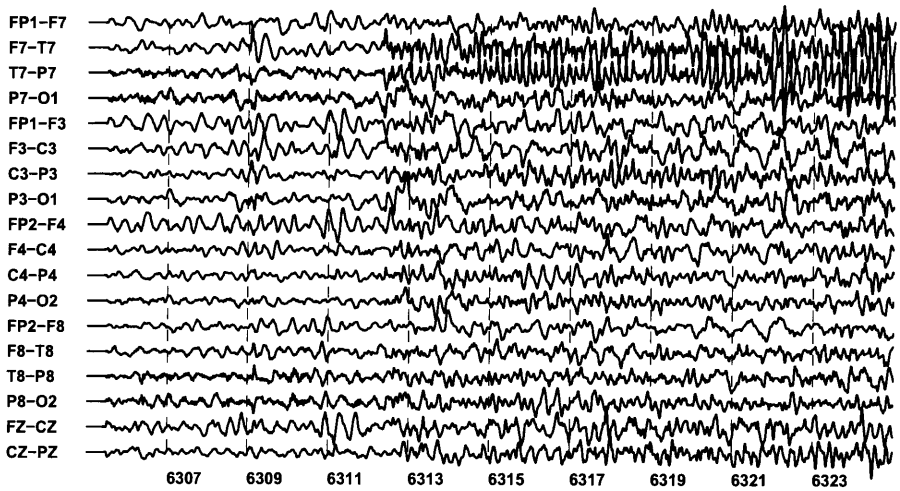

Figure 2-6: Example of a seizure within the scalp EEG of Patient B. The seizure, which begins at 6313 seconds, involves the appearance of a theta band rhythm on the channels $\{F7 - T7, T7 - P7\}$.

Between seizures, the EEG of an individual with epilepsy may exhibit abnormal rhythmic activity or discharges. The spatial and spectral characteristics of these discharges varies across patients. Figure 2-7 illustrates an abnormal discharge within the EEG of Patient A. The discharge, which involves most EEG channels, falls between 2884-2892 seconds, and is characterized by a repeating pattern of high-amplitude spikes followed by broad waves. While these discharges are seen frequently in the awake EEG of Patient A, they are not accompanied by the physical symptoms associated with Patient A's seizure (Figure 2-5). Consequently, a detector designed to react to the seizures of Patient A should not produce an alarm upon the onset of one of these discharges. In another patient, this type of activity may be associated with physical symptoms, as is visible in the seizure illustrated in Figure 5-25.

Figure 2-8 illustrates abnormal rhythmic activity observed within the EEG of

Figure 2-7: Example of an abnormal discharge within the scalp EEG of Patient A. The discharge, which occurs between 2884-2892 seconds, involves a repeating pattern of high-amplitude spikes followed by broad waves that can be seen on most EEG channels.

Patient B. In this case, a theta band rhythm is observed on the frontal channels $\{FP1 - F3, F3 - C3\}$ during the awake state between 6126-6130 seconds. This activity is both frequent and abnormal, but it is not accompanied by the physical symptoms associated with Patient B's seizure (Figure 2-6). Once again, a detector designed to react to the seizures of Patient B should not produce an alarm upon the onset of this rhythmic activity.

The previous examples illustrate the variability in the spectral and spatial signature of seizure and non-seizure activity across patients. This variability is the primary reason why patient non-specific detectors exhibit poor sensitivity and specificity for seizure events [18]. In contrast, for a given individual, seizures that emerge from the same cerebral site exhibit similar clinical symptomatology and electrographic spatial and spectral characteristics. As an example, Figures 2-9 and 2-10 illustrate more seizures recorded from patients A and B respectively. Note the similarity in the spatial and spectral character of these seizures and those illustrated in Figures 2-5 and 2-6.

The long-term stability of an individual's seizure signature has not been studied

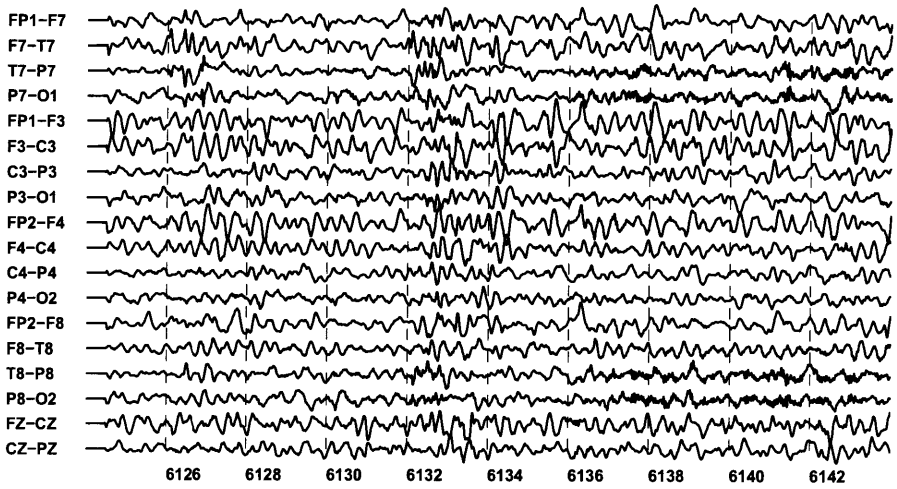

Figure 2-8: Example of abnormal rhythmic activity within the scalp EEG of Patient B. The rhythmic activity, which occurs between 6126-6130 seconds, involves a theta band rhythm on the frontal channels $\{FP1 - F3, F3 - C3\}$.

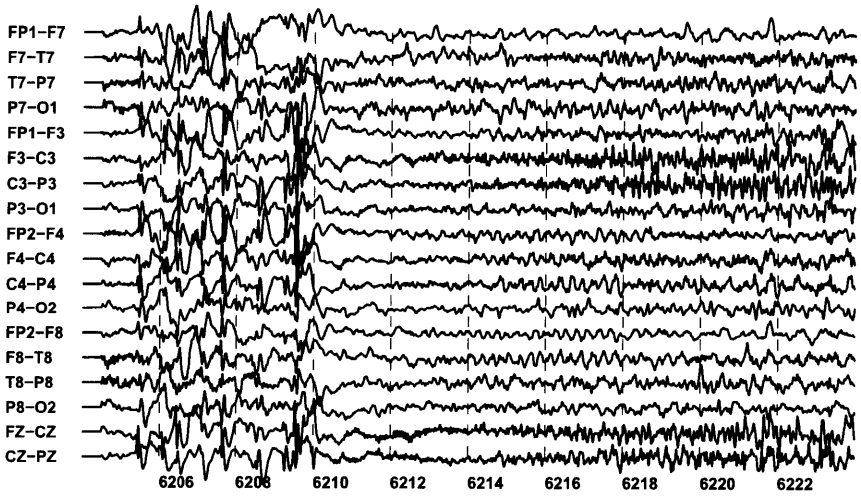

Figure 2-9: A second seizure within the scalp EEG of Patient A. The second seizure, which begins at 6210 seconds, resembles the seizure shown in Figure 2-5.

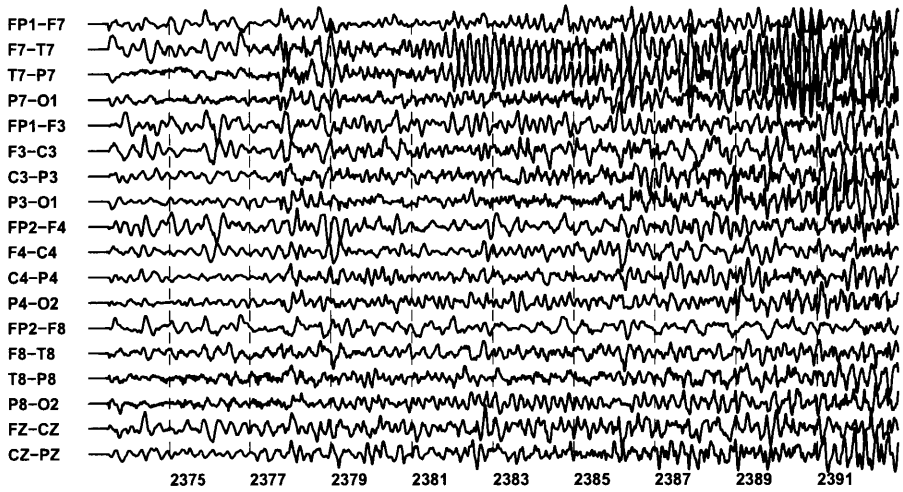

Figure 2-10: A second seizure within the scalp EEG of Patient B. The second seizure, which begins at 2381 seconds, resembles the seizure shown in Figure 2-6.

rigorously. However, in our experience, seizures recorded months apart from patients without a progressive brain disorder exhibit very similar spectral and spatial characteristics. As an example, Figures 2-11 and 2-12 illustrate two seizures that were recorded approximately 10 months apart from Patient C. Both seizures begin with rapid eye-blinking, at 1488 seconds and 992 seconds respectively, which manifests in the EEG as high-amplitude, downward deflections on the channels $\{FP1-F3, FP2-F4, FP1-F7, FP2-F8\}$. Later, both seizures exhibit a 3-4 Hz theta wave most prominently on the EEG channel $T4-T6$ at 1492 seconds and 998 seconds respectively.

2.3 Intracranial Electroencephalogram

Similar to scalp EEG, the intracranial Electroencephalogram (iEEG) provides a spatial and temporal summary of the electrical activity of a population of neurons. Since intracranial EEG electrodes are placed on the brain cortex or deep within brain struc-

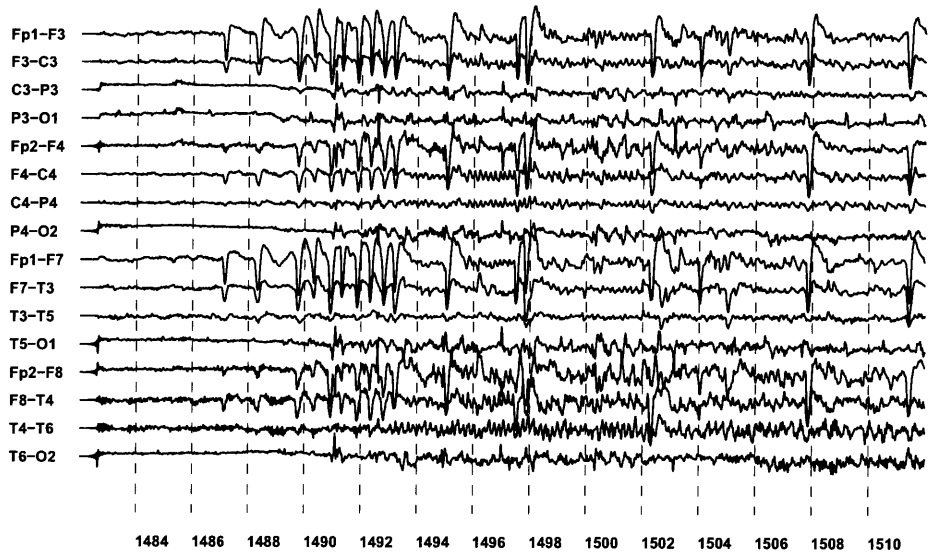

Figure 2-11: Example of a seizure within the scalp EEG of Patient C. The seizure, which begins at 1486 seconds, involves rapid eye-blinking followed by the appearance of a 3-4 Hz theta wave on the channel $T4 - T6$.

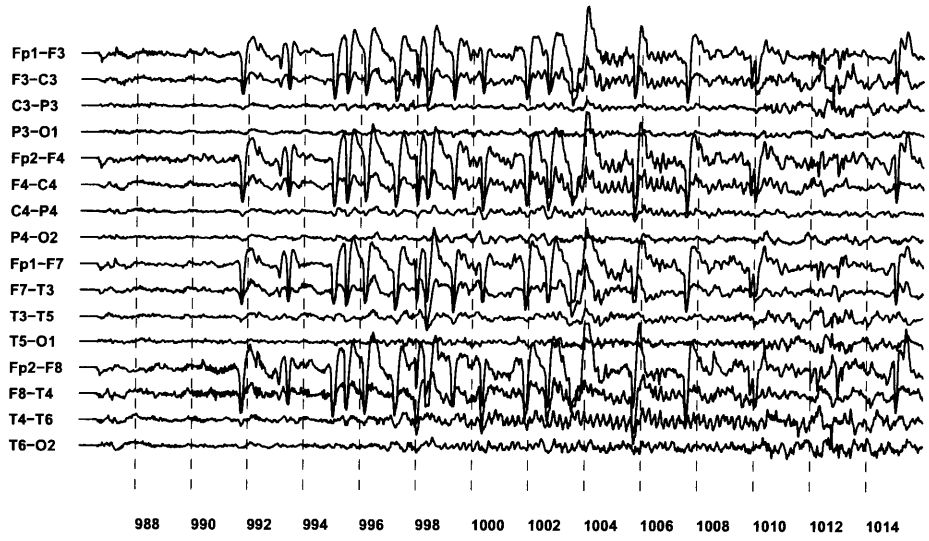

Figure 2-12: Example of a seizure recorded from Patient C approximately 10 months after recording the seizure in Figure 2-11.

tures, the intracranial EEG signal summarizes the activity of a smaller population of neurons and is immune to corruption by many of the physiologic and environmental artifacts that plague scalp EEG. Because they are highly invasive, iEEG electrodes can only be implanted in a limited number of brain sites at any given time; so intracranial EEG offers higher spatial resolution but spatial coverage that is worse than that of scalp EEG.

The higher spatial resolution of iEEG allows one to notice the neuronal hypersynchrony associated with a seizure tens of seconds before the same phenomenon is noticeable within the scalp EEG [39]. At the same time, the higher spatial resolution of iEEG permits the recording of a wider gamut of abnormal, non-seizure activity that is not visible within the scalp EEG [26, 59, 62]. This activity can confound the seizure detection process.

2.3.1 Seizures Within the Intracranial Electroencephalogram

The manifestation of seizures within the intracranial EEG involves a sudden redistribution of spectral energy on a set of iEEG channels. The spectral energy change typically consist of an appearance or disappearance of frequency components within 0-65 Hz band [20]. This frequency range is wider than that associated with scalp EEG because of the absence of the attenuating effect of the cerebrospinal fluid and skull. Prior to the spectral energy change, discrete events such as a spike or train of spikes may appear.

As was the case with scalp EEG, the manifestation of seizures within the iEEG varies significantly across patients. To observe this variability consider Figure 2-13. The top panel corresponds to a seizure from Patient D while the bottom panel corresponds to a seizure from Patient E. The seizure of Patient D begins at 5 seconds and consists of a few spikes that transition into a high-amplitude spike train with a period of 1 Hz. The seizure of Patient E also begins at 5 seconds and consists of a few spikes followed by low-amplitude, high-frequency activity.

Outside the context of a seizure the intracranial EEG of an individual with epilepsy may contain a wide variety of rhythmic activity with characteristics that resemble

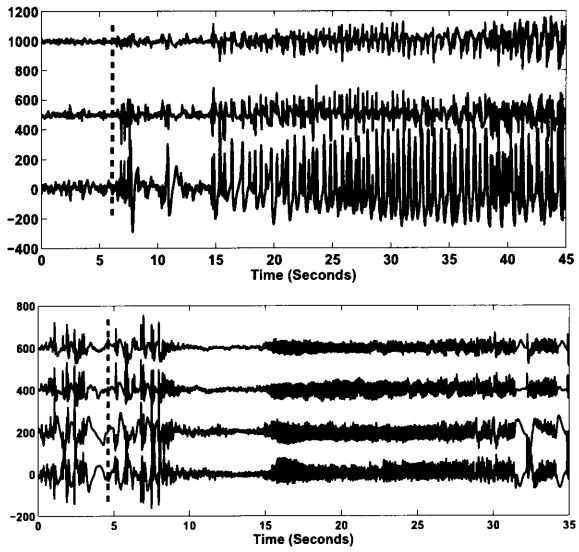

Figure 2-13: Example of seizures within the iEEG of patients D (top panel) and C (bottom panel). The seizure of Patient D begins at 5 seconds and consists of a few spikes that evolve into a high-amplitude spike train. The seizure of Patient E begins at 5 seconds and consists of a few spikes followed by low-amplitude, high-frequency activity.

those of a seizure. As an example, consider Figure 2-14. The top panel of Figure 2-14 illustrates a seizure recorded from Patient F while the bottom panel illustrates a burst of rhythmic activity recorded from the same patient. An algorithm sensitive only to changes in spectral energy without regard to the frequency range or channels on which those changes occur [38] may falsely declare the activity in the bottom panel as a seizure.

For a given patient, the spectral character of seizures within the iEEG is similar from one seizure to the next provided that the seizures originate from the same cerebral site [14, 29]. As an example, consider the two seizures from Patient D illustrated in Figure 2-15. Both seizures begin with a few spikes that transition into a high-amplitude spike train with a period of 1 Hz. However, the spatial character of seizures within the iEEG, that is which iEEG channels are involved, tends to vary more than in the setting of scalp EEG. More specifically, for a given patient, neighboring iEEG channels may alternate in exhibiting the first signs of iEEG activity associated with the onset of a seizure from the same brain region.

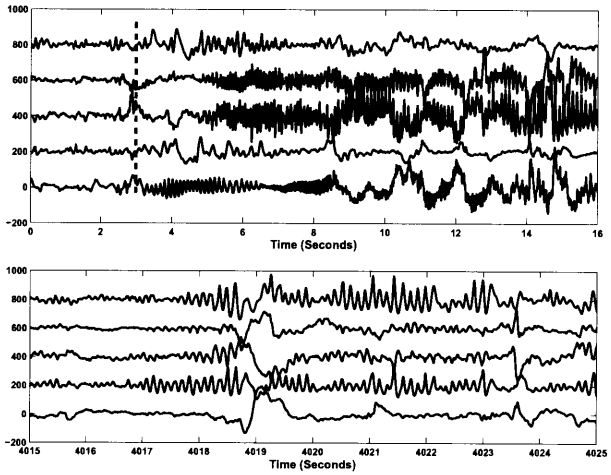

Figure 2-14: The iEEG of Patient F contains rhythmic activity that is associated with a seizure (top panel) as well as rhythmic activity that is not associated with any clinical symptoms (bottom panel).

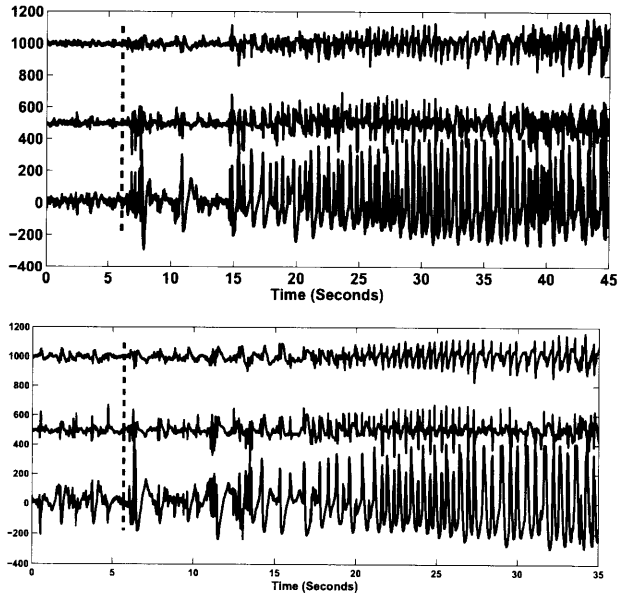

Figure 2-15: Example of two seizures within the iEEG of Patient D. Both seizures begin at 5 seconds and consist of a few spikes that evolve into a high-amplitude spike train.

2.4 Summary

Seizures are transient periods of neural network hyperactivity and hypersynchrony. The clinical symptoms that accompany a seizure depend both on its cerebral site of origin as well as its pattern of spread to surrounding brain regions. The scalp and intracranial electroencephalograms, which measure the aggregate electrical activity of populations of neurons, can be used to detect seizures. The scalp EEG offers poor spatial resolution but high spatial coverage while the intracranial EEG offers high spatial resolution but less spatial coverage.

Within both the scalp and intracranial electroencephalograms seizures manifest as a redistribution of spectral energy on a set of channels. The manner with which spectral energy is redistributed varies significantly across patients. For a given individual, seizures from the same cerebral site generally exhibit stereotyped onset and evolution.

Chapter 3

Patient-Specific Seizure Onset Detection using the Scalp EEG

In this chapter we present the architecture of a novel algorithm for detecting the onset of a seizure within the scalp EEG. The algorithm treats the seizure onset detection problem as a binary classification task that is to be solved within a patient-specific context. The feature vector used by the algorithm encodes the time evolution of spectral and spatial properties of the scalp EEG.

3.1 Overview of Binary Classification

Binary classification is the process of assigning an observation to one of two classes or categories in a manner that optimizes a chosen performance objective; for instance, minimizing the probability of an erroneous classification. As an example of binary classification, consider the problem of determining the absence or presence of an aircraft using returned radar pulses. In this scenario the observation is the radar pulse, which represents the presence (Class C_1) or absence (Class C_2) of an aircraft. In the seizure onset detection problem the observation is an EEG epoch and the two classes are seizure activity (Class C_1) and non-seizure activity (Class C_2).

Determining the class membership of an observation involves two steps. First, salient properties or features that most discriminate between instances of each class

are extracted from the observation and assembled into a vector of features. In the case of radar signal processing, the features may be the energy within a series of returned radar pulses. Next, a classifier that is trained to recognize the difference between feature vectors extracted from instances of each class determines the class membership of the observation based on its associated feature vector.

The success of a binary classification task depends strongly on which features are extracted from an observation as well as on the classifier used to determine class membership. Features that have a distribution of values for observations belonging to class C_1 and another, very different distribution for observations belonging to class C_2 result in good performance. A classifier capable of deducing from training data a simple rule that accurately distinguishes between feature vectors extracted from instances of each class is also important to good performance.

3.2 Feature Vector Design

Seizures are dynamic processes. The onset of a seizure induces within the scalp EEG a sudden change in the spectral energy distribution of a set of EEG channels. As the seizure progresses, the spectral and spatial character of the EEG continues to evolve. The following sections detail how the feature vector used in our detector is designed to encode the time evolution of spectral and spatial properties within the scalp EEG.

3.2.1 Spectral Features

Following the onset of most seizures, scalp EEG channels that record brain activity within regions involved in the seizure exhibit rhythmic activity. The spectral structure of this rhythmic activity may be composed of multiple frequency components. As an example, Figure 3-1 illustrates a seizure that begins following 2994 seconds, and Figure 3-2 illustrates the frequency spectrum of the channel $FP2 - F4$ between 2994-2999 seconds. For this seizure, the rhythmic activity on the channel $FP2 - F4$ is composed primarily of frequency components at 2, 5, and 11 Hz.

Considering several spectral components that constitute an EEG epoch, and not

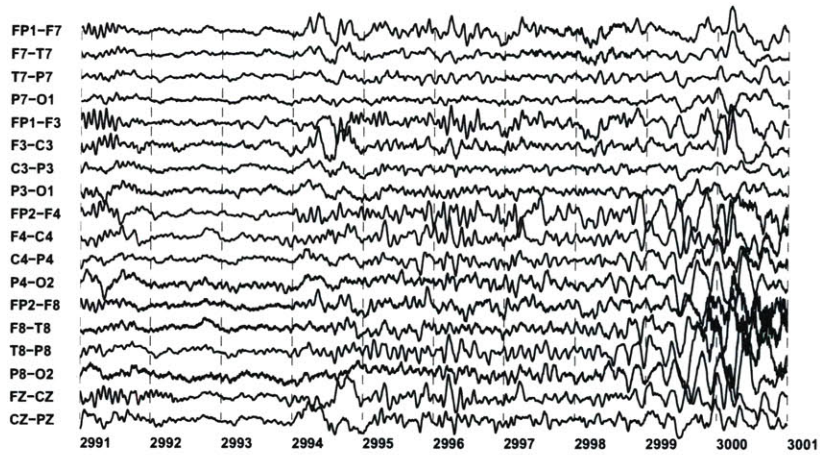

Figure 3-1: Example of a seizure whose onset is associated with rhythmic activity that contains a mixture of prominent frequency components. Seizure onset follows 2994 seconds and mostly involves channels on the right side of the head (channels with even numerals).

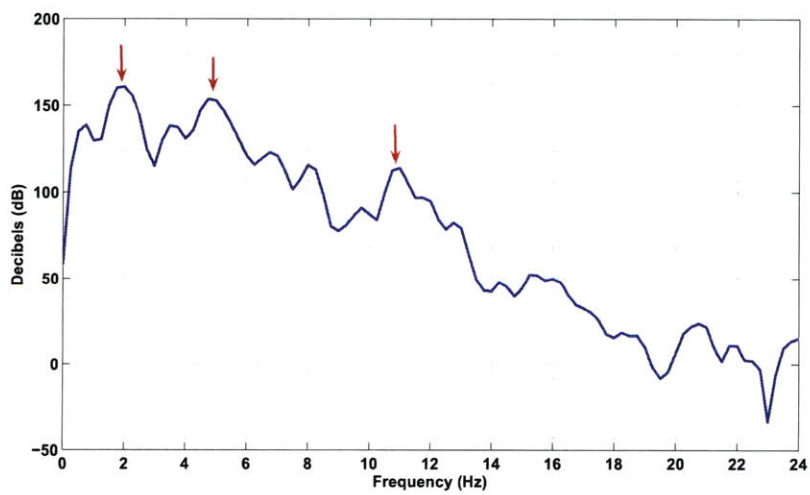

Figure 3-2: Frequency spectrum of the channel $FP2 - F4$ following the onset of the seizure illustrated in Figure 3-1. The spectrum contains large spectral components at 2, 5, and 11 Hz.

simply the dominant spectral component as was done in early seizure detection algorithms [18], is important in detecting seizure onsets with high specificity. The dominant frequency of a seizure onset epoch may overlap the dominant frequency of an epoch of non-seizure activity, but what differentiates the two is the presence or absence of less dominant spectral components. The following examples illustrate this point.

Figure 3-3 illustrates the spectrum of an eye-blink superimposed on the seizure spectrum (Figure 3-2). The two spectra overlap in the frequency range 0-4 Hz, but the spectrum of the seizure activity contains an 11 Hz component that is absent from the spectrum of the eye-blink.

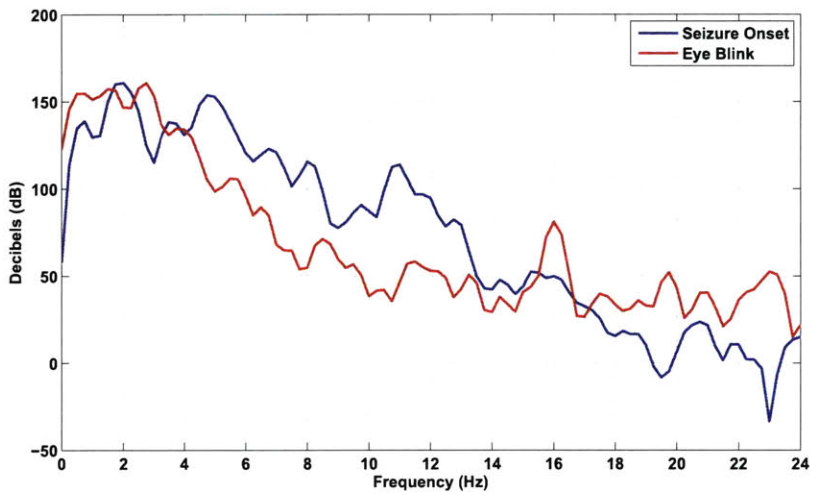

Figure 3-3: Superposition of the frequency spectra of an eye-blink and the rhythmic activity observed on the channel $FP2 - F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains an 11 Hz component that is absent from the spectrum of the eye-blink

Figure 3-4 illustrates the spectrum of a sleep spindle superimposed on the seizure spectrum. In this case, the two spectra overlap in the frequency range 10-12 Hz, but the seizure spectrum contains larger low-frequency spectral components.

Finally, Figure 3-5 illustrates the spectrum of chewing superimposed on the seizure spectrum. The two spectra overlap within the frequency range 0-12 Hz, but the seizure spectrum contains less high-frequency content relative to the chewing spectrum.

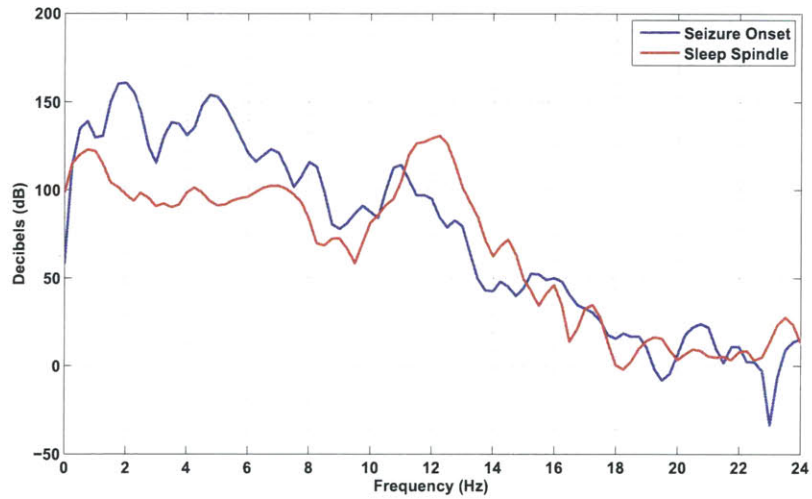

Figure 3-4: Superposition of the frequency spectra of a sleep spindle and the rhythmic activity observed on the channel $FP2-F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains low-frequency spectral components that are larger than those in the spindle spectrum.

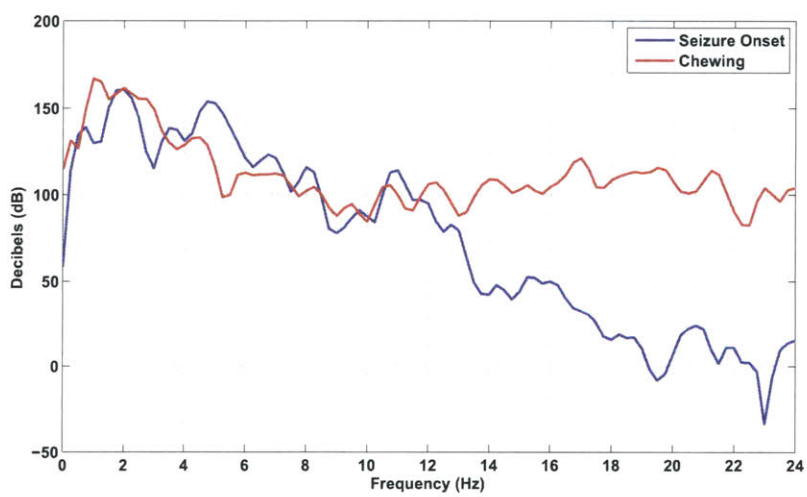

Figure 3-5: Superposition of the frequency spectra of chewing and the rhythmic activity observed on the channel $FP2 - F4$ following the onset of the seizure in Figure 3-1. The seizure spectrum contains less high-frequency content relative to the chewing spectrum.

Our feature vector encodes the spectral structure of an L second epoch from EEG channel k by passing the epoch through the M-band filterbank shown in Figure 3-6, and then measuring the energy in the subband signals S_i, $i = 1, \ldots, M$. The energy in the subband signal S_i from channel k is represented by the feature $x_{i,k}$. The filterbank spans the frequency range 0.5-24 Hz since within this range one observes most physiologic and pathophysiologic scalp EEG activity [19].

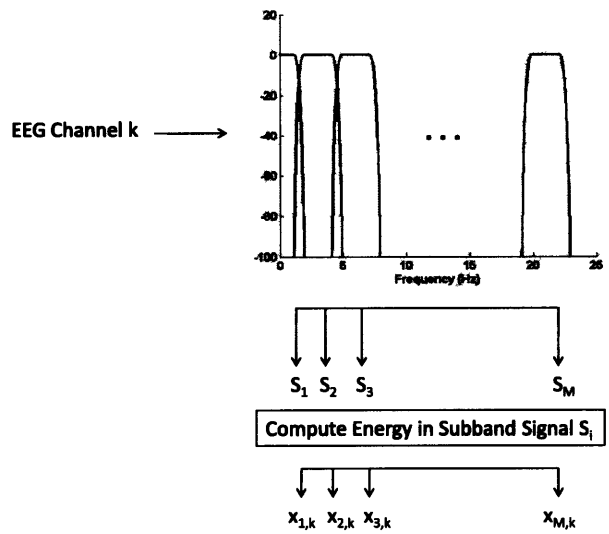

Figure 3-6: M-band filterbank that measures the energy within the spectral components of an L second epoch taken from a single EEG channel.

As an example, Figure 3-7 illustrates how energies within the frequency bands defined by an $M = 8$ filter filterbank can differentiate between the spectra of seizure and non-seizure activity. The energy within the frequency bands defined by the rightmost three filters (thick line) differentiate between the seizure spectrum (red) and the spectrum associated with chewing (blue).

3.2.2 Spatial Features

The set of EEG channels on which EEG activity is observed, which is referred to as its spatial distribution, can be important in differentiating between seizure onset and non-seizure activity.

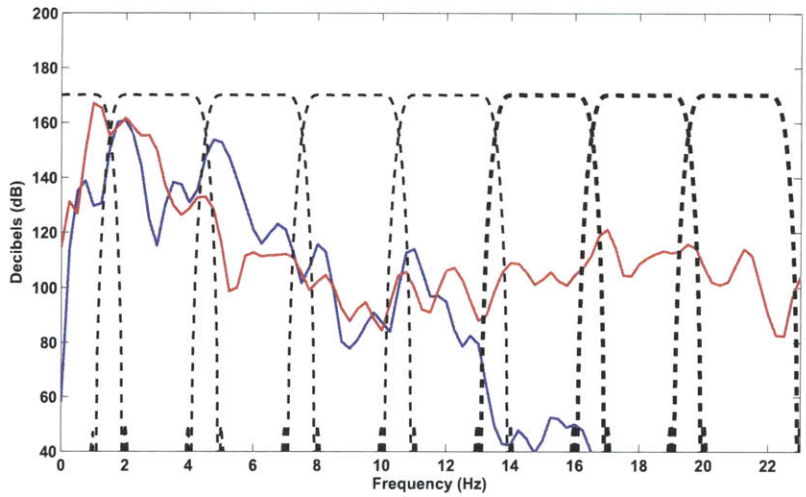

Figure 3-7: Energy within the frequency bands defined by the rightmost three filters (thick line) of the filterbank differentiates between the seizure spectrum (red) and chewing spectrum (blue)

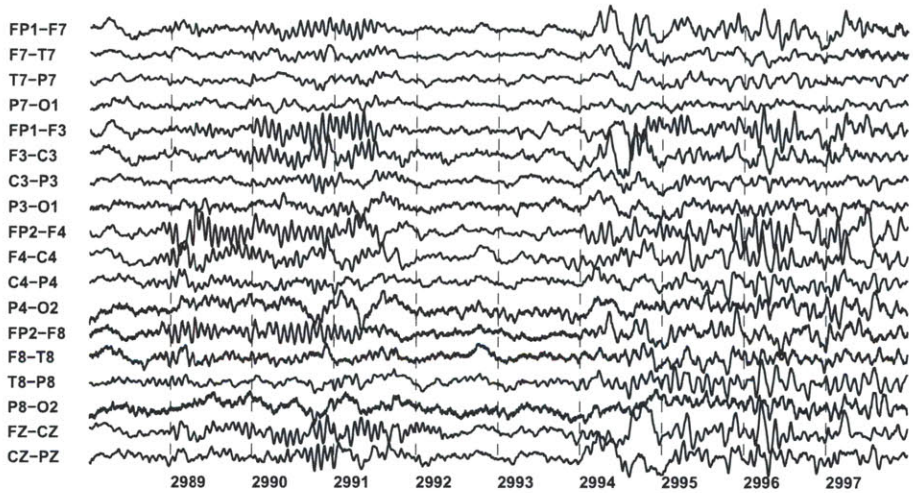

Figure 3-8: The spatial distribution of EEG activity can be used to differentiate seizure from the non-seizure activity. The sleep spindles between 2989-2992 seconds do no involve activity on the channels $C4 - P4$ and $T8 - P8$. In contrast, the seizure activity following 2994 seconds involves both of these channels.

As an example, Figure 3-8 illustrates sleep spindles between 2989-2992 seconds, and seizure activity following 2994 seconds. The sleep spindle and seizure activity are both rhythmic, but the seizure involves EEG channels ($C4 - P4$ and $T8 - P8$) that are not involved in the sleep spindle event. Consequently, the absence of activity on a set of channels during a non-seizure event and its presence during seizure onset can serve to differentiate EEG activity beyond what is possible using spectral information alone.

To simultaneously capture the spectral and spatial information contained within an L second EEG epoch at time $t = T$, we concatenate the M spectral energies extracted from each of N EEG channels to form a composite feature vector X_T with $M \times N$ elements. The process of generating X_T is illustrated in Figure 3-9. It is important to note that the feature vector X_T *automatically* captures relationships between the spectral structure of different EEG channels. Previous algorithms relied on experts to define those relationships [44, 46].

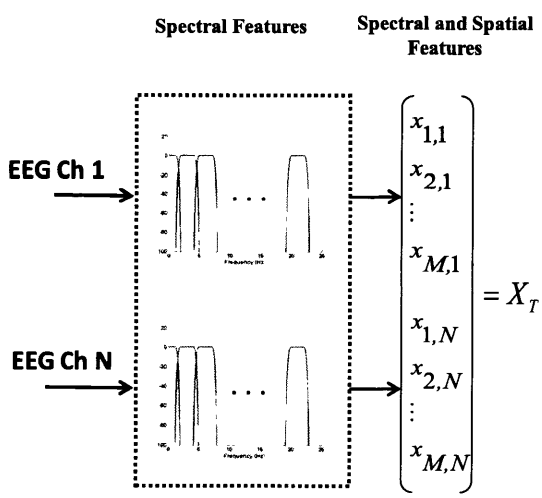

Figure 3-9: Formation of an intermediate feature vector X_T that captures the spectral and spatial properties of an epoch at time $t = T$

3.2.3 Time Evolution

The feature vector X_T developed in section 3.2.2 uses all EEG channels to capture the spatial and spectral content of an EEG epoch at time $t = T$, but it does not capture how the current epoch relates to those in the recent past. Consequently, the feature vector X_T cannot represent how a seizure emerges from the background EEG nor how it evolves during the onset period. In order to capture the relation between feature vectors across time we form a composite feature vector $\mathcal{X}_T$ that is the result of concatenating feature vectors from W non-overlapping epochs. This is similar to what was done in [7] for off-line intracranial seizure onset detection. The feature vector $\mathcal{X}_T$ is constructed as shown in equation 3.1, where X_T is a spectral and spatial feature vector extracted from an epoch of length L seconds at time $t = T$.

$$\mathcal{X}_T = [X_{T-(W-1)L} \ \ \ldots \ \ X_{T-2L} \ \ X_{T-L} \ \ X_T] \tag{3.1}$$

Encoding the temporal evolution of EEG through the formation of $\mathcal{X}_T$ is not analogous to forming a single feature vector X_T using a longer epoch length L. In the former approach, discrete events are preserved and in the latter the spectral signature of discrete events within the longer epoch is smeared.

To appreciate the value of capturing the temporal evolution of the feature vectors X_T consider deriving the composite feature vector $\mathcal{X}_T$ for the seizure shown in Figure 3-10 using only the channel $F3 - C3$ (recall that in the actual detector all channels are used). In this example the epoch length used to from X_T is $L = 2$ seconds long and the feature vector $\mathcal{X}_T$ is formed by concatenating $W = 3$ feature vectors: X_{T-4}, X_{T-2}, and X_T.

The seizure in Figure 3-10 involves a sequence of events that can be modeled through the feature vector $\mathcal{X}_T$. The seizure begins at 1723 seconds with a spike followed by a period of low amplitude EEG known as an *electrodecrement*. Next, a beta band rhythm appears most prominently on the channel $F3 - C3$, and, over the period of a few seconds, increases in amplitude and decreases in frequency towards

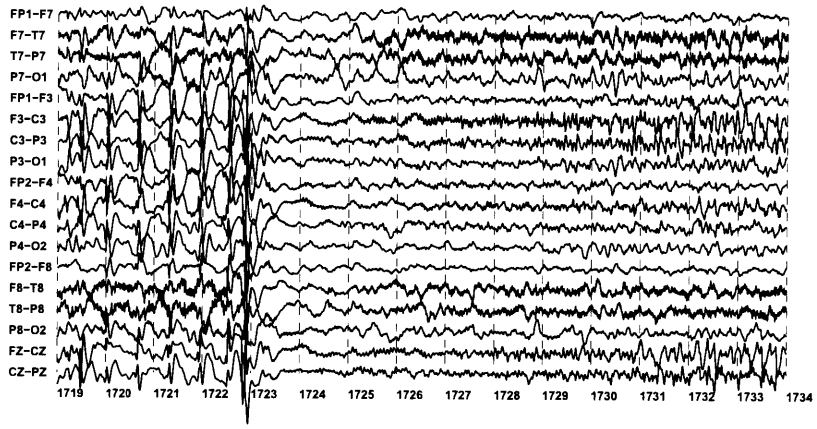

Figure 3-10: Scalp EEG seizure involving a sequence of discrete spectral events. The first event, at 1723 seconds, is a spike involving all EEG channels. Next, all EEG channels exhibit a period of low-amplitude EEG. Finally, at 1725 seconds, a beta band rhythm that increases in amplitude and decreases in frequency appears on the channel $F3 - C3$.

the theta band.

To more clearly demonstrate the series of spectral events that constitute the seizure in Figure 3-10, we show the *spectrogram* of channel $F3 - C3$ in Figure 3-11. A spectrogram is a plot of the spectral content of a signal as function of time. In a spectrogram spectral components with greater energy are shown in darker colors and those with less energy in brighter colors. In Figure 3-11 one can see a high energy event with spectral content up to 30 Hz between 3.5-4.5 seconds, this corresponds to the spike at 1723 seconds in Figure 3-10. Between 4.5-6.5 seconds a decrease in energy (brighter colors) in the frequency band 0-10 Hz is observed, this corresponds to the electrodecrement between 1723-1725 seconds in Figure 3-10. Finally, between 6.5-13.5 seconds one observes a negatively sloped, dark colored line representing a spectral component that begins with a frequency of 20 Hz and decreases towards a frequency of 5 Hz. This corresponds to the onset and evolution of the beta band rhythm at 1725 seconds in Figure 3-10.

Figures 3-12 and 3-13 illustrate how the composite feature vector $\mathcal{X}_T$ is able to model the sequence of events associated with both seizure onset and evolution. In Figure 3-12 note how $\mathcal{X}_T$, which is composed of the triplet of feature vectors X_{T-4},

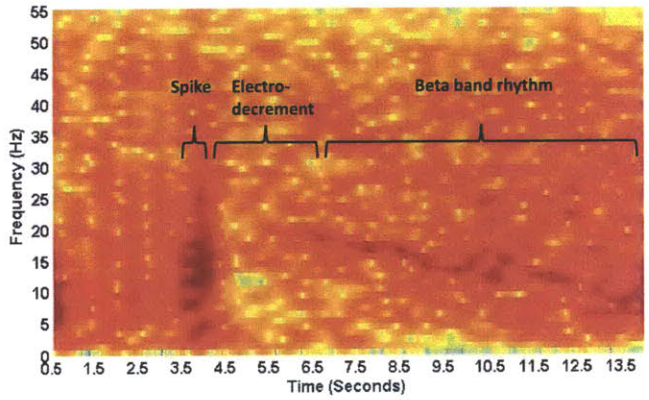

Figure 3-11: Spectrogram of Channel $F3 - C3$ illustrating sequence of events that compose the onset and evolution of the seizure in Figure 3-10.

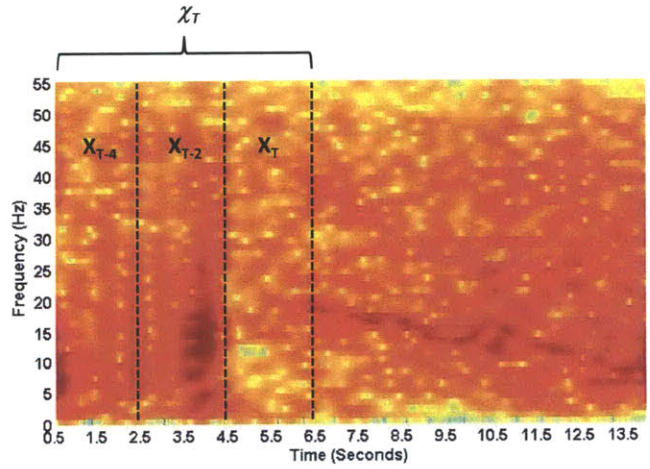

Figure 3-12: Spectrogram of channel $F3 - C3$ illustrating how the feature vector $\mathcal{X}_T$ models the transition from background to seizure onset. The feature vector X_T captures the period of low amplitude activity, X_{T-2} captures the spike, and X_{T-4} captures the background EEG activity preceding the spike.

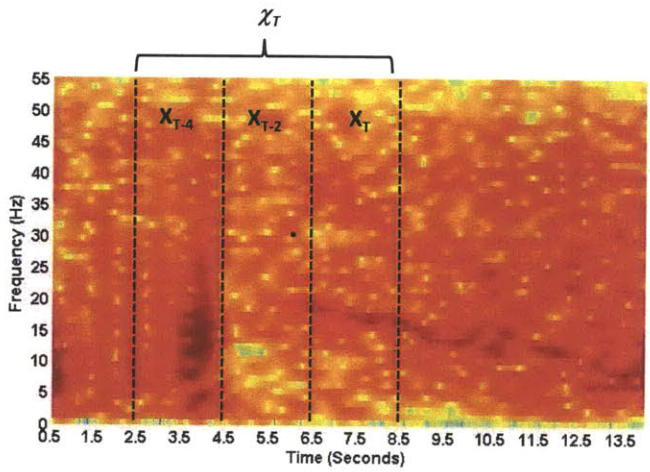

Figure 3-13: Spectrogram of channel $F3 - C3$ illustrating how the feature vector $\mathcal{X}_T$ models the evolution of a seizure. The feature vector X_T captures the beta band rhythm, X_{T-2} captures the period of low amplitude activity, and X_{T-4} captures the spike.

X_{T-2}, and X_T spans the transition from background to seizure activity. The feature vector X_T captures the period of low amplitude activity, X_{T-2} captures the spike, and X_{T-4} captures the background EEG activity preceding the spike. When shifted forward in time as shown in Figure 3-13, the composite feature vector $\mathcal{X}_T$ models the seizure's evolution. The feature vector X_T captures the beta band rhythm, X_{T-2} captures the period of low amplitude activity, and X_{T-4} captures the spike.

3.3 Feature Vector Classification

To classify the feature vector $\mathcal{X}_T$ as representative of seizure or non-seizure activity we used the Support-Vector Machine (SVM) algorithm [9]. The SVM algorithm uses training seizure and non-seizure feature vectors to learn a *decision boundary* that separates these two classes of activity. Once the decision boundary is learned, the SVM algorithm determines the class membership of a newly observed feature vector $\mathcal{X}_T$ based on which side of the boundary the vector falls.

In its simplest form, the SVM algorithm determines a decision boundary in the

form of a high-dimensional plane, or hyperplane, in the space of features. The hyperplane is chosen to maximize the classification *margin*, which is the geometric distance between the hyperplane and the boundary cases of each class [61]. The boundary cases are known as *support-vectors*. If the classes cannot be well separated by a hyperplane, as is the case in the seizure onset detection problem, the SVM can be used to determine a quadratic or circumferential decision boundary. Conceptually, the SVM determines a nonlinear boundary in the input feature space by solving for a linear boundary in a higher-dimensional feature space induced by a *kernel*.

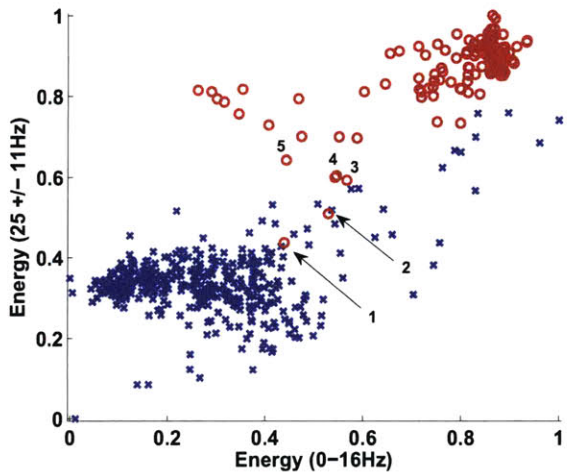

Figure 3-14: Two-dimensional seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The first dimension of this feature space is defined by the energy in the spectral band 0-16 Hz and the second dimension corresponds to the energy in the 25 ± 11 Hz band. The seizure vectors extracted from the first 7 seconds of the seizure are numbered 1-5.

As an example, Figure 3-14 plots, in a two-dimensional feature space, iEEG seizure feature vectors (red circles) and non-seizure feature vectors (blue crosses) extracted from a single iEEG channel. The first dimension of this feature space is defined by the energy in the spectral band 0-16 Hz and the second is the energy in the 25 ± 11 Hz band. The seizure feature vectors (red circles) form two clusters: one cluster with fewer feature vectors that lies in close proximity to the non-seizure feature vectors, and a second cluster with a much greater number of feature vectors that lies far away from the non-seizure vectors. The seizure vectors closest to the non-seizure vectors

are associated with the seizure onset period and those further away are associated with later stages of the seizure. In Figure 3-14, the seizure vectors numbered 1-5 are extracted from the first 7 seconds of a seizure.

Figure 3-15 shows the linear decision boundary separating the iEEG seizure feature vectors from the non-seizure vectors. As can be seen the two classes are not well separated by a linear decision boundary as there are numerous crosses on the side of the hyperplane designated for seizure feature vectors.

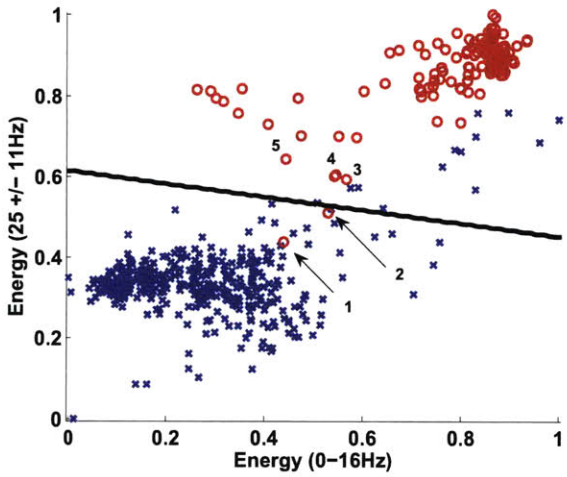

Figure 3-15: Linear decision boundary separating seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The decision boundary was determined using the SVM learning algorithm.

Figure 3-16 shows a nonlinear decision boundary separating the same iEEG seizure and non-seizure feature vectors. The nonlinear boundary, which is computed using a *radial basis* kernel, is able to better separate the seizure and non-seizure feature vectors.

The earliest seizure vector that is correctly classified by either the linear or non-linear decision boundary is the third seizure vector. The SVM learning algorithm allows the user to specify the value of a parameter J that controls the relative cost of misclassifying seizure and non-seizure vectors. If the value of J is set so that the cost of misclassifying a seizure vector far exceeds the cost of misclassifying a non-seizure vector, then the SVM learning algorithm could produce a boundary that can cor-

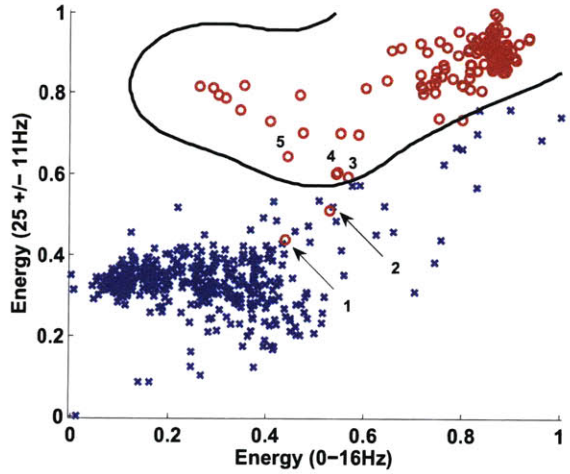

Figure 3-16: Nonlinear decision boundary separating seizure (red) and non-seizure (blue) feature vectors extracted from a single iEEG channel. The decision boundary was determined using a radial basis kernel and the SVM learning algorithm.

rectly classify the first and second seizure vectors at the expense of classifying more non-seizure vectors incorrectly.

The mapping function $f(\cdot)$ that establishes whether a feature vector falls within the region of the feature space defined by seizure activity is expressed as shown in equation 3.2 in the case of a linear boundary

$$\begin{aligned} f(X) &= \text{Seizure} \;\; if \;\; W^T X + \beta > 0 \\ f(X) &= \text{Non-seizure} \;\; if \;\; W^T X + \beta \leq 0 \end{aligned} \tag{3.2}$$

Where the vector W, which is the normal to the separating hyperplane, and the bias term β are parameters determined by the SVM learning algorithm [9]. In the case of a nonlinear boundary determined using a radial basis kernel, the discriminant function takes the form shown in equation 3.3. The coefficients α_i, support-vectors X_i, and bias term β are again determined by the SVM learning algorithm. The parameter γ, which controls how tightly the nonlinear boundary circumscribes a class, is user-defined.

$$f(X) = \text{Seizure } if \ \{\sum_{i=1}^{N_{SV}} \alpha_i \exp(-\gamma \|X - X_i\|^2)\} + \beta > 0 \tag{3.3}$$

$$f(X) = \text{Non-seizure } if \ \{\sum_{i=1}^{N_{SV}} \alpha_i \exp(-\gamma \|X - X_i\|^2)\} + \beta \leq 0$$

The number of support-vectors N_{SV} is partly governed by the complexity of the classification task. As the similarity between an individual's seizure and non-seizure activity increases more support-vectors are needed in order to define a more complex decision boundary, and as a result the computational cost of equation 3.3 increases.

The support-vector machine algorithm is well-suited for the task of seizure onset detection because its learning mechanism focuses on determining a decision boundary that separates the boundary cases of each class; recall that the boundary cases of the seizure class are associated with the onset of the seizure. Support-vector machines have also been shown to perform well in high-dimensional classification tasks [27] and in the setting where one class has a smaller number of training examples [1].

3.4 Patient-Specific Detector Architecture

The block diagram in Figure 3-17 illustrates the processing stages of the patient-specific detector. The detector passes L-second epochs from each of N EEG channels through a filterbank. In turn, the filterbank computes for each channel M features that correspond to the energies within M frequency bands. The M features extracted from each of the N channels are then concatenated to form a $M \times N$ element vector X_T that automatically captures the spectral and spatial relations between channels. Next, the W feature vectors X_T, X_{T-L},..., and $X_{T-(W-1)L}$ are concatenated to form the feature vector $\mathcal{X}_T$. This vector, which contains $W \times M \times N$ elements, automatically captures the time evolution of spectral and spatial relations among the input EEG channels.

Finally, the feature vector $\mathcal{X}_T$ is assigned to the seizure or non-seizure class using a

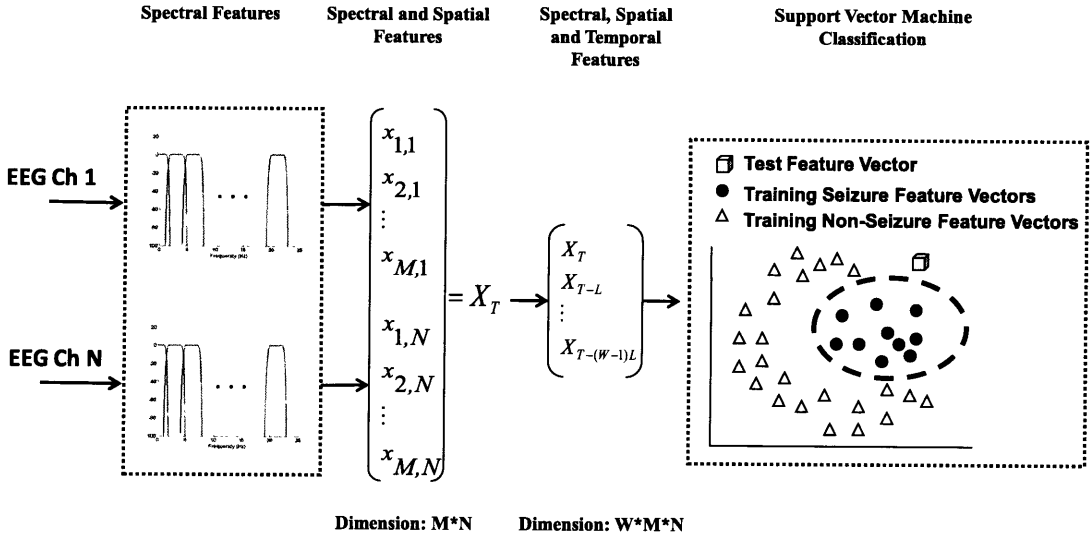

Figure 3-17: Patient-specific seizure onset detector architecture.

two-class support-vector machine classifier with cost factor J, radial basis kernel with parameter γ, and trade-off between classification margin and error C. The detector declares a seizure after one feature vector $\mathcal{X}_T$ is assigned to the seizure class.

Since seizure and non-seizure activity are generally stereotypical for a patient and highly variable across patients (Chapter 2), the support-vector machine is trained on training seizure and non-seizure vectors from a single patient. The training non-seizure vectors are extracted from H hours of continuously recorded scalp EEG. The seizure vectors are derived from the first S seconds following the onset of K training seizures.

Figure 3-18 illustrates how successive non-seizure training vectors are extracted from a non-seizure training record when L=2 and W=3. For these parameter values, a six second window is shifted, one second at a time, across a non-seizure record. For each positioning of the six second window within the non-seizure record, a training non-seizure feature vector $\mathcal{X}$ is generated. More specifically, in Figure 3-18, the time interval 0-6 seconds is broken up into three, non-overlapping 2 second epochs that are used to generate the feature vectors X_2, X_4, and X_6. These vectors are then grouped

to form the first training non-seizure feature vector $\mathcal{X}_6$. Next, the time interval 1-7 seconds is broken up into three, non-overlapping 2 second epochs that are used to generate the feature vectors X_3, X_5, and X_7. These vectors are then grouped to form the second training non-seizure feature vector $\mathcal{X}_7$. The third training non-seizure feature vector $\mathcal{X}_8$ is formed by grouping the feature vectors X_4, X_6, and X_8 extracted from the time interval 2-8 seconds. The training non-seizure vector $\mathcal{X}_8$ overlaps the training vector $\mathcal{X}_6$ in composition.

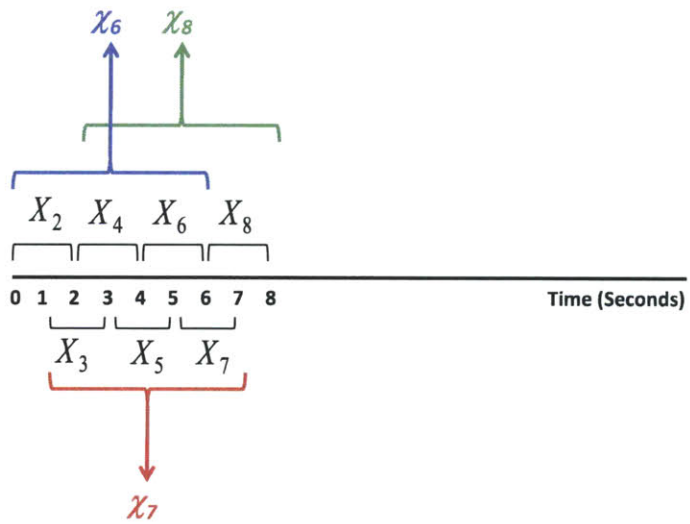

Figure 3-18: Generating training non-seizure feature vectors.

The process used to generate training seizure vectors is similar to that used to generate training non-seizure vectors. Figure 3-19 illustrates how successive seizure training vectors are extracted from a training seizure when L=2 and W=3. In Figure 3-19, seizure onset is at time 0 seconds. The first training seizure vector $\mathcal{X}_2$ is formed by grouping the feature vectors extracted from the six second window spanning the time interval between -4 and +2 seconds. The second training seizure vector $\mathcal{X}_3$ is formed by grouping the feature vectors extracted from the six second window spanning the time interval between -3 and +3 seconds. This process is repeated until one has

advanced S seconds into the seizure.

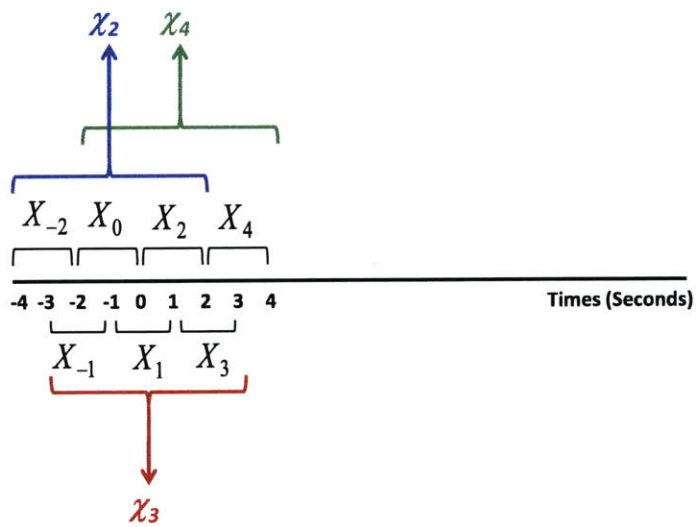

Figure 3-19: Generating training seizure feature vectors.

Table 3.1 summarizes the parameters that control our patient-specific detection algorithm.

3.5 Instantiating Detector Parameters

In this section we assign values to the detector parameters listed in Table 3.1. In Chapter 5 we illustrate how perturbing the assigned value of some of these parameters impacts seizure onset detection performance.

3.5.1 EEG Epoch Length: L

Most of the spectral energy of physiologic and pathophysiologic scalp EEG activity falls within the 0.5-24 Hz frequency band [19]. In order to resolve the lowest end of the this frequency range, the detector extracts spectral features from a $L = 2$ second

Table 3.1: Patient-specific Detector Parameters

Parameter	Description
L	EEG epoch length
N	Number of EEG channels
M	Number of filters in filterbank
W	Number of feature vectors that form $\mathcal{X}_T$
γ	SVM radial basis kernel parameter
J	Relative cost of misclassifications
C	Trade-off between classification margin and error
H	$\sharp$ of hours used to extract training non-seizure vectors
S	$\sharp$ of seconds into seizure used to extract training seizure vectors
K	$\sharp$ of seizures used to extract training seizure vectors

epoch. Other seizure detection algorithms have used 2.5 second [46] and 6 second [37] EEG epochs.

3.5.2 Number of EEG Channels: N

All available EEG channels are provided as inputs to the seizure detection algorithm. The scalp EEG data set used to evaluate the algorithm, which is described in Chapter 4, contained $N = 18$ EEG channels. However, in [54], Shih shows that limiting the detector presented in this chapter to using $N = 5$ channels does not significantly impact seizure detection performance.

3.5.3 Number of filters: M

The detector uses an $M = 8$ uniform bandwidth filterbank. The fiterbank spans the 0-25 Hz frequency range using filters that are 3 Hz wide. In Chapter 5 we demonstrate the impact of using a filterbank with $M < 8$ filters.

3.5.4 Number of feature vectors in $\mathcal{X}_T$: W

The detector groups $W = 3$ feature vectors to form the feature vector $\mathcal{X}_T$. With $W = 3$, the detector examines the evolution of the EEG over a period of 6 seconds

when determining whether or not to declare a seizure. As will be shown in Chapter 5, using fewer vector to form $\mathcal{X}_T$ increases false detections, and using more vectors increases detector latency.

3.5.5 SVM Parameters: γ, J, and C

The SVM used by the detector is trained using the *SVMLight* software package [28] with a cost factor $J = 1$, RBF kernel parameter of $\gamma = 0.1$, and trade-off between classification margin and error $C = 1$. We showed in [55, 56] that these values, after adjusting for SVM software package differences, work well on an independent data set consisting of 139 seizures and 49 hours of non-seizure EEG recorded from 36 pediatric epilepsy patients.

3.5.6 Training Parameters: H, K, and S

The detector should be trained on non-seizure vectors extracted from $H \geq 24$ hours of non-seizure EEG so that it is exposed to awake, sleep, abnormal, and artifact-contaminated EEG. The detector should be trained on as many seizures originating from the same seizure focus; however, as shown in Chapter 5, performance gains are marginal for $K > 3$ seizures. For any given seizure, the detector is trained on seizure vectors extracted from the first $S = 20$ seconds of a seizure. Seizure onset is marked by an electroencephalographer and corresponds to the onset of rhythmic activity associated with a clinical seizure. In Chapter 5 we demonstrate the impact of training on $S < 20$ seconds.

Chapter 4

Scalp EEG Data and Testing Methodology

In this chapter we describe the scalp EEG data and testing methodology used to estimate the latency, sensitivity, and specificity of the patient-specific seizure onset detector described in Chapter 3.

4.1 Scalp EEG Data Set

The data set used to evaluate the performance of our patient specific detector consists of continuous Scalp EEG recordings from 23 pediatric patients (age < 18) undergoing medication withdrawal for epilepsy surgery evaluation at Children's Hospital Boston. The EEG was sampled at 256 Hz and recorded using an 18-channel, 10-20 bipolar montage. According to the 10-20 bipolar montage, EEG electrodes are arrayed on the scalp as shown in Figure 2-1, and the EEG channels formed are those illustrated in the EEG tracings discussed in Chapter 2. Overall, this 23 patient data set contained 844 hours of continuously recorded EEG and 163 seizures.

The scalp EEG data set is segmented into records. A record is typically one hour long. Records that do not contain a seizure are called *non-seizure records* and those that contain one or more seizures are called *seizure records.*

4.2 Performance Metrics

Three metrics will be used to characterize the performance of our seizure onset detector.

1. **Electrographic Seizure Onset Detection Latency** $EO_{Latency}$. The electrographic onset of a seizure refers to the onset of scalp EEG changes associated with a seizure. The clinical onset of a seizure refers to the onset of its physical or cognitive symptoms. In scalp EEG, the electrographic onset of a seizure may or may not precede its clinical onset; in contrast, in iEEG, the electrographic onset of a seizure tends to precede its clinical onset. $EO_{Latency}$ refers to the delay between electrographic onset and detector recognition of seizure activity. By definition $EO_{Latency} \geq 0$.

2. **Sensitivity** S refers to the percentage of test seizures identified by a detector.

3. **False Alarms Per Hour** FA refers to the number of times, over the course of an hour, that a detector declares the onset of seizure activity in absence of an actual seizure.

4.3 Performance Metric Measurement

To estimate our detector's performance on data recorded from a patient, we use a *leave-one-record-out* cross-validation scheme. For the purposes of this discussion, Let N_{NS} denote the number of non-seizure records and let N_S denote the number of seizure records.

Measuring $EO_{Latency}$ and S

The detector is trained on the N_{NS} non-seizure records of a patient (median $N_{NS} = 33$) as well as $N_S - 1$ records containing a seizure (median $N_S = 5$). The detector is then tasked with detecting seizures in the mth seizure record; the record that was withheld from the training set. This process is repeated N_S times so that each of

the N_S seizure records is tested once; a seizure record is never simultaneously in the training and testing sets.

Let $S_m \in \{0, 1\}$ be a binary variable that denotes whether the detector noted the mth seizure, and let $EO_{Latency,m}$ denote the latency with which the detector notes that seizure's electrographic onset. Let $FA_{S,m}$ be the number of false alarms declared while processing the mth seizure record. Finally, let K denote the total number of detected seizures. Equation (4.1) shows how these quantities are combined into an estimate of the detector's performance for a given patient.

$$S = \frac{1}{N_S} \sum_{m=1}^{N_S} S_m \tag{4.1}$$

$$EO_{Latency} = \frac{1}{K} \sum_{m=1}^{N_S} S_m * EO_{Latency,m}$$

Measuring FA

A detector is trained on the N_S seizure records of a patient as well as $N_{NS} - 1$ non-seizure records. The detector is then tasked with processing the nth non-seizure record; the record that was withheld from the training set. This process is repeated N_{NS} times so that each of the N_{NS} non-seizure records is tested once; a non-seizure record is never simultaneously in the training and testing sets. Let $FA_{NS,n}$ be the number of false alarms declared while processing the nth non-seizure record. Equation (4.2) is used to estimate the detector's false alarm rate.

$$FA = \frac{1}{N_{NS} + N_S} * (\sum_{n=1}^{N_{NS}} FA_{NS,n} + \sum_{m=1}^{N_S} FA_{S,m}) \tag{4.2}$$

Chapter 5

Performance

In this chapter we evaluate the performance of the detector described in Chapter 3 using the scalp EEG data and testing methodology described in Chapter 4. We also analyze how varying certain detector parameters impacts detector performance. Finally, we compare the performance of our algorithm to that of a state-of-the-art patient non-specific seizure detector.

5.1 Patient-Specific Detector Performance

5.1.1 Latency

Figure 5-1 illustrates the percentage of test seizures detected within a specified latency. As can be seen, 50% of the 163 test seizures are detected within 3 seconds, 71% within 5 seconds, and 91% within 10 seconds. The mean latency with which the detector declared the onset of a test seizure is 4.6 seconds.

Figure 5-2 illustrates the latency with which the detector declared the onset of each seizure for each of the 23 subjects. A seizure is represented by a blue dot, and the height of a dot represents the latency associated with a seizure detection. If more than one seizure was detected with the same latency, then the numeral next to a dot indicates the number of seizures it represents. For most patients, the majority of seizures are detected within 5 seconds.

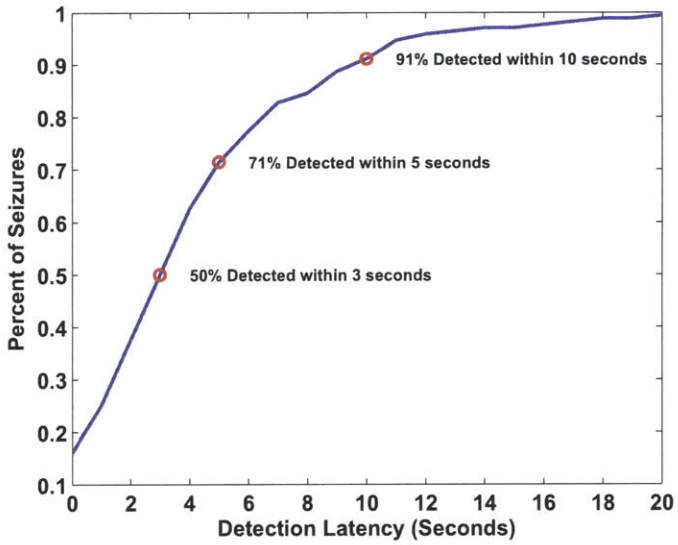

Figure 5-1: Percentage of 163 test seizures detected within a specified latency. The detector notes the onset of 50% of all test seizures within 3 seconds.

For Patient 15 one of the seizures is detected with a 55 second delay. The reason for the long detection latency is that the EEG at the onset of this seizure differed in spectral and spatial characteristics from training seizures. Figure 5-3 illustrates a typical training seizure. The seizure begins following 272 seconds with a theta band rhythm that is most prominently seen on the channel $T7 - P7$. Figure 5-4 illustrates the seizure that the algorithm failed to detect promptly. This seizure, which begins at 876 seconds, involves a train of spikes on the channel $P7 - O1$.

5.1.2 Sensitivity

Figure 5-5 shows the sensitivity with which the patient-specific algorithm detects the test seizures of each of the 23 subjects. Overall, 96% of 163 test seizures were detected.

5.1.3 Specificity

Figure 5-6 shows the number of false detections declared by the patient-specific detector in a period of 24 hours for each of the 23 subjects. For most patients (18 of

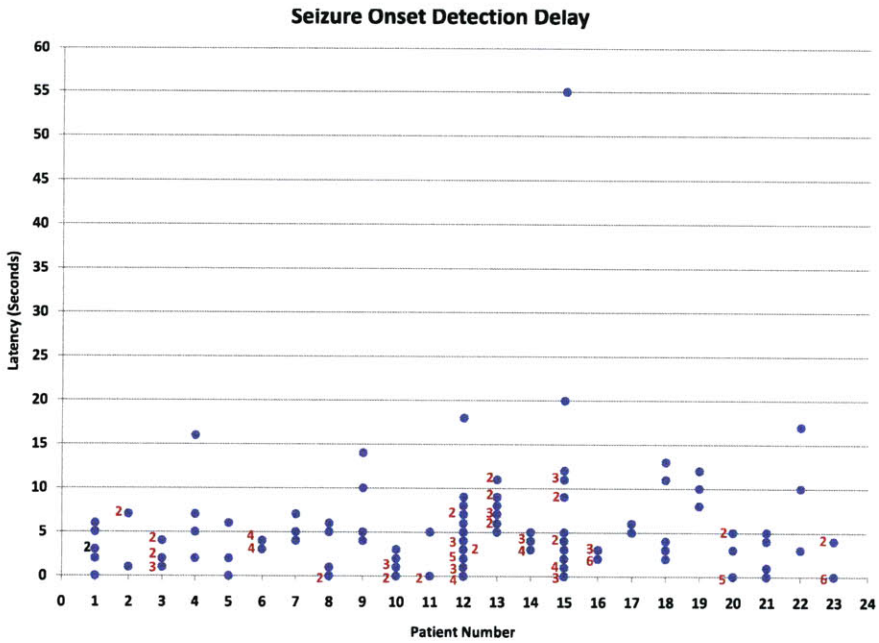

Figure 5-2: Latency with which the detector notes the onset of seizures for each of the 23 test subjects. Each dot represents a seizure. A numeral is placed next to dots that represent more than one seizure. For most patients, the majority of seizures are detected within 5 seconds.

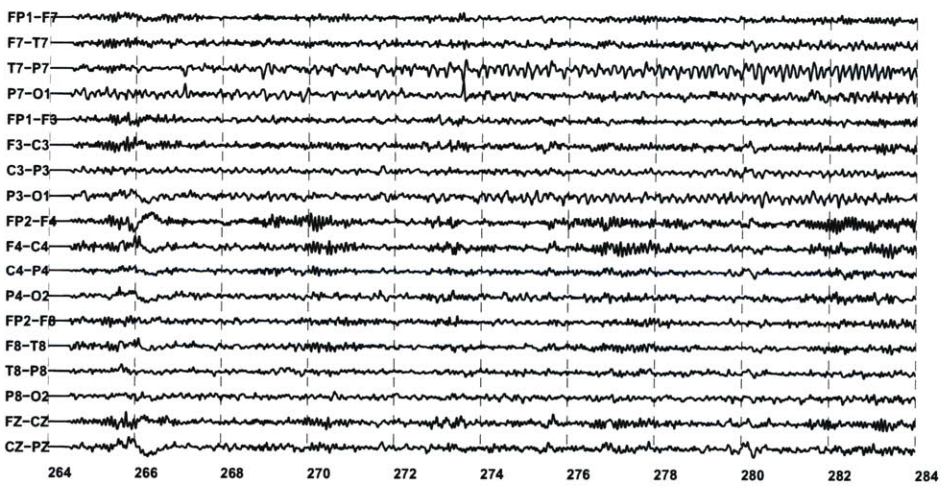

Figure 5-3: Example of a seizure within the scalp EEG of Patient 15. The seizure, which begins at 272 seconds, consists of a theta band rhythm that is most prominently seen on the channel $T7 - P7$.

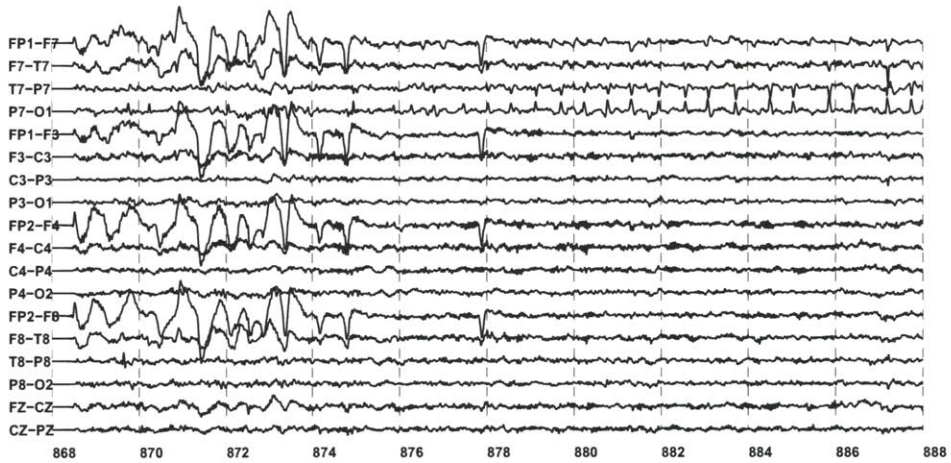

Figure 5-4: Example of another type of seizure observed within the scalp EEG of Patient 15. This seizure, which begins at 876 seconds, consists of a train of spikes on the channel $P7 - O1$. The detector fails to detect the onset of this seizure since its spectral and spatial characteristics differ from those of training seizures. Training seizures resemble the seizure shown in Figure 5-3.

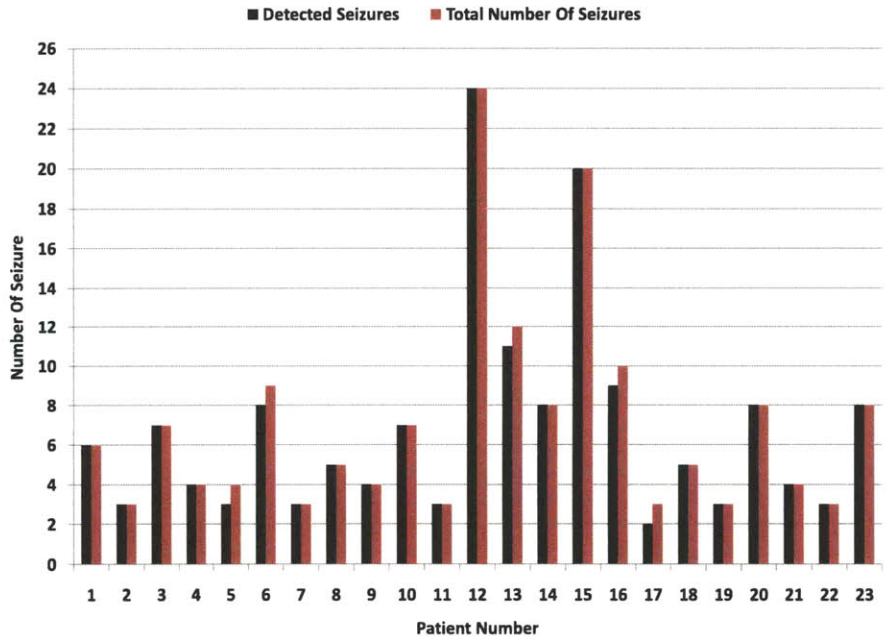

Figure 5-5: Sensitivity of the patient-specific seizure detector. Red bars show the number of test seizures available for each subject, and black bars show the number of test seizures recognized by the detector. Overall, the detector recognized 96% of 163 test seizures.

23), the detector declares between 0 and 5 False detections per 24 hours.

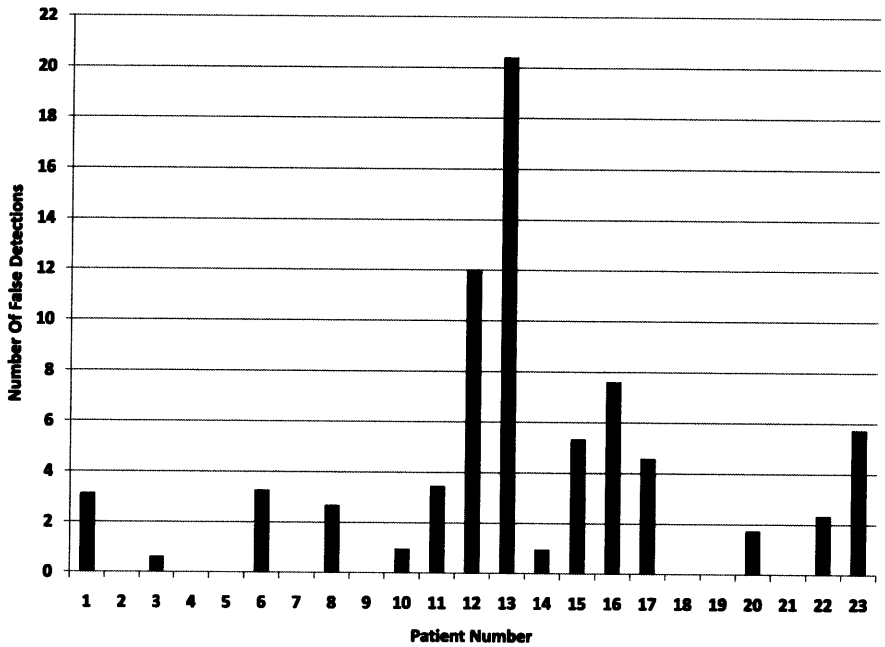

Figure 5-6: Specificity of the patient-specific seizure detector. For most patients (18 of 23), the detector declares between 0 and 5 false detections per 24 hours.

For Patient 13 the algorithm declares 20 false detections per 24 hour period. The reason for the high false alarm rate is the presence of short bursts of rhythmic activity that match the seizure onset signature both in spectral and spatial character. Figure 5-7 illustrates the onset of a typical seizure recorded from Patient 13. The seizure begins following 640 seconds with a delta band rhythm that is most prominent on the channel $FP1 - F7$. Figure 5-8 illustrates the burst of rhythmic activity that resulted in a false detection. The burst, which occurs between 233-235 seconds, also consists of a delta band rhythm that is most prominent on the channel $FP1 - F7$. The main feature that distinguishes the seizure from the burst is the seizure's temporal extent. To improve the specificity of the detection algorithm for Patient 13, one could declare a seizure only when activity suspected of being a seizure persists for a duration of time that exceeds the average length of the bursts of rhythmic activity. The cost of such a modification will be an increase in detection latency.

For the two patients with the highest number of false detections, patients 12 and

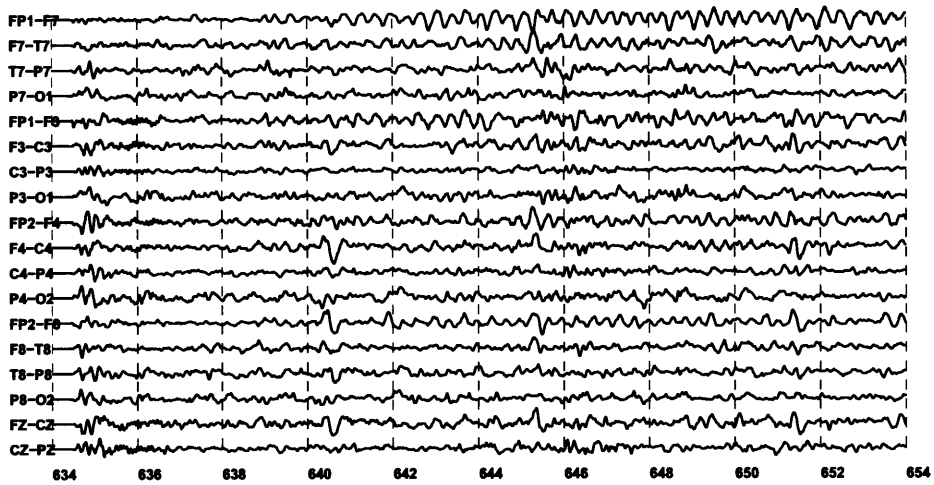

Figure 5-7: Example of a seizure within the scalp EEG of Patient 13. The seizure, which begins at 640 seconds, consists of a delta band rhythm that is most prominent on the channel $FP1 - F7$.

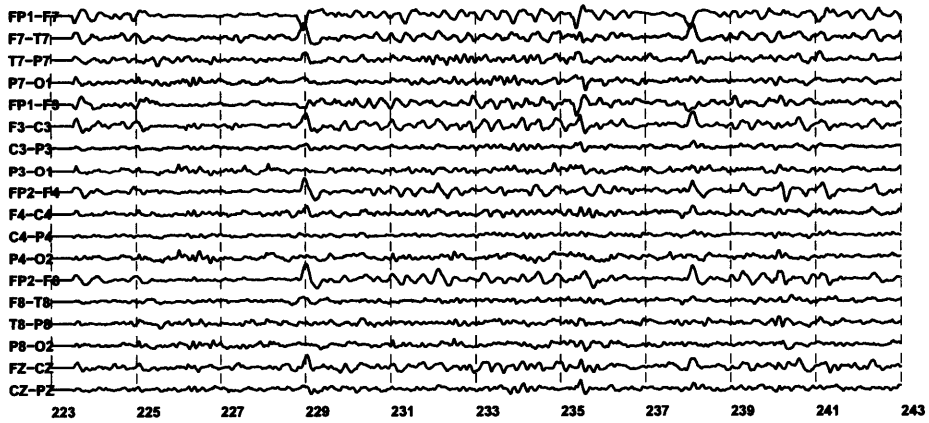

Figure 5-8: Example of a short burst of rhythmic activity within the scalp EEG of Patient 13. The burst, which occurs between 233-235 seconds, consists of a delta band rhythm that is most prominent on the channel $FP1 - F7$. Since the burst resembles the seizure in Figure 5-7, the detector declares a false detection following the onset of the burst.

13, we note that false detections tend to occur in a small fraction of the 1 hour records processed by the detector. Figure 5-9 illustrates that 75% of the test records (24 of 33 hours) resulted in 0 false detections in the case of Patient 13. For Patient 12, 68% of the test records (11 of 16 hours) resulted in 0 false detections.

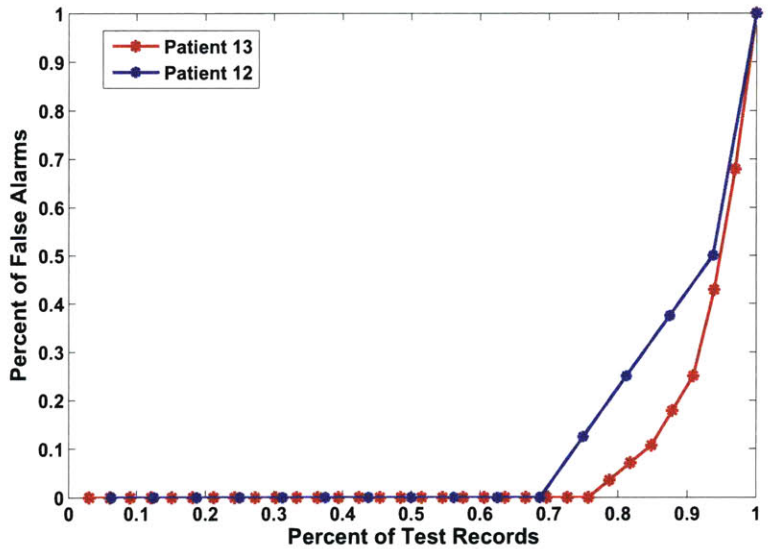

Figure 5-9: False detections declared by the detector are not uniformly distributed across test records. For Patient 13, 75% of the test records resulted in 0 false detections. For Patient 12, 68% of the test records resulted in 0 false detections.

5.2 Varying Patient-Specific Detector Parameters

5.2.1 Varying the Number of Filters: M

On the scalp EEG data set discussed in section 4.1 we profiled the performance of our patient-specific detector as a function of uniformly splitting the bandwidth 0-25 Hz using 2, 4, and 8 equal bandwidth filters. In these experiments the number of filters M used to generate the feature vector X_T varies, but the number of feature vectors that are grouped to form $\mathcal{X}_T$ is fixed at W=3.

Increasing the number of filters within the filterbank does not greatly impact the detection latency or the sensitivity of the detector. The average, across 23 subjects,

latency of detectors with 2, 4, and 8 filter filterbanks was 4.4, 4.7, and 4.8 seconds respectively, and all detectors had a sensitivity of 96%. Increasing the number of filters appears to improve the specificity of a detector as shown in Figure 5-10. The average specificity of a detector that uses a two-filter filterbank is greater than 8 false detections per 24 hours, and that of a detector with an eight-filter filterbank is less than 4 false detections per 24 hours.

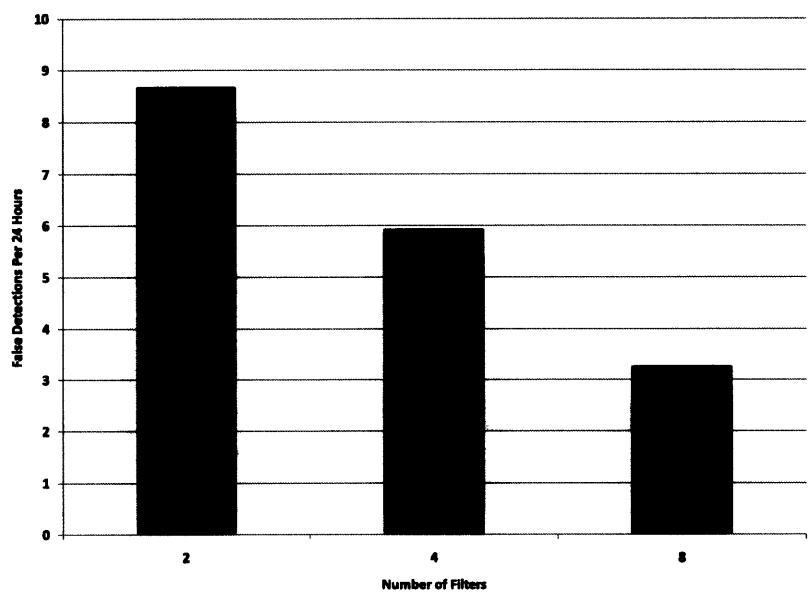

Figure 5-10: The false detection rate of the detector decreases as number of filters used to construct X_T is increased from 2 to 8. A two-filter filterbank yields a false detection rate of 8 false detections per 24 hours. An eight-filter filterbank results in a false detection rate smaller than 4 false detections per 24 hours.

5.2.2 Varying the Number of Feature Vectors in $\mathcal{X}_T$: W

On the scalp EEG data set discussed in section 4.1 we profiled the performance of our patient-specific detector as the number feature vectors W used to form $\mathcal{X}_T$ is changed. In these experiments a $M = 8$ filter filterbank was used to generate the feature vectors $X_T, X_{T-L}, \ldots, X_{T-(W-1)L}$ that are embedded within $\mathcal{X}_T$.

The sensitivity of the patient-specific detector remained 96% as the number of

feature vectors within $\mathcal{X}_T$ was increased from 2 through 4. However, as W is increased the false detection rate of the detector decreases while its latency increases as shown in Figure 5-11. In order to detect seizures with an average latency shorter than 5 seconds one should set $W = 3$.

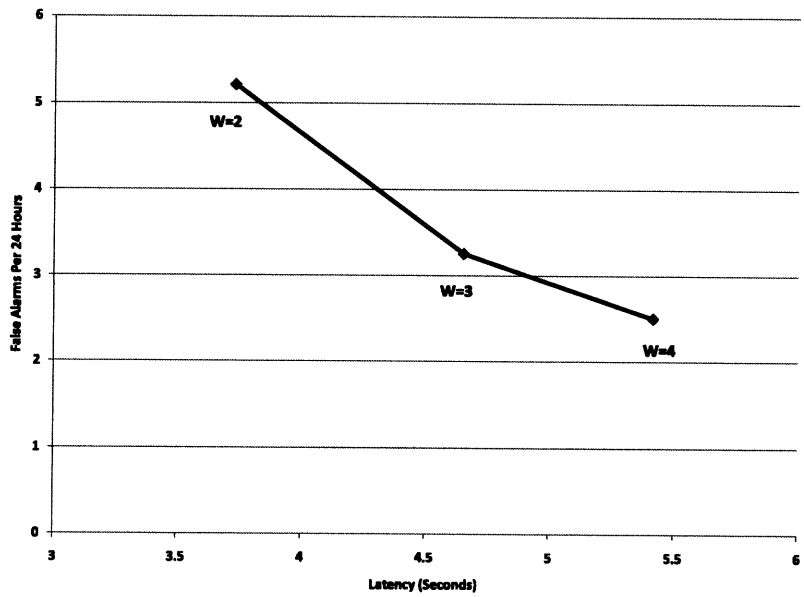

Figure 5-11: Increasing the number of feature vectors within $\mathcal{X}_T$ decreases the detector's false detection rate and increases its detection delay. To detect seizures with an average latency shorter than 5 seconds one should set $W = 3$.

5.2.3 Varying the Number of Training Seizures: K

The performance of the detector improves as more seizures are included in the training set. Figure 5-12 illustrates how the average detection latency and miss rate decrease with an increasing number of training seizures for five randomly selected patients. With a single training seizure, the detector has a latency greater than 7 seconds and misses more than 45% of the test seizures. With three training seizures, the detector has a latency close to 4 seconds and misses less than 5% of test seizures. The addition of a fourth seizure improves the detector's latency to 3.5 seconds but does have a great impact on the detector's miss rate.

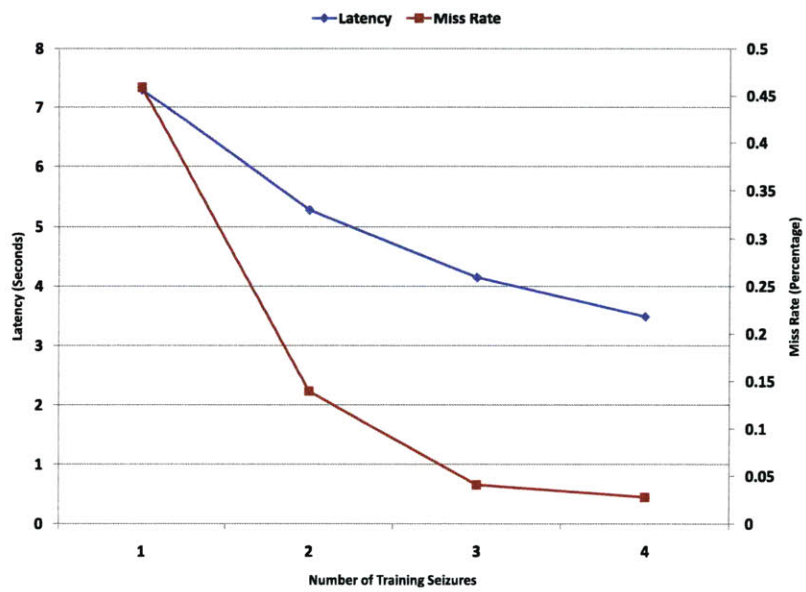

Figure 5-12: Increasing the number of training seizures decreases the detector's miss rate and its detection delay. Including more than three training seizures results in a marginal improvement in detector's miss rate.

5.2.4 Varying Training Time into a Seizure: S

In this section we examine how the performance of the detector changes when training on seizure feature vectors derived from the first $S = 6$, 12, 18, or 24 seconds of training seizures. Figure 5-13 shows that the detector's average latency does not improve significantly for $S > 12$ seconds. Moreover, the detector's sensitivity exceeds 90% for $S > 18$ seconds as shown in Figure 5-14. Finally, Figure 5-15 shows that the detector's false detection rate increases as the value of S is increased.

One explanation for the increase in false detections with increasing S assumes that the decision boundary developed by the SVM learning algorithm encloses a single, contigous region of the feature space as illustrated in Figure 3-16. With this assumption, increasing S forces the SVM to develop a wider decision region in order to enclose feature vectors from the early and late portions of a seizure. Such a boundary is more likely to result in false detections since it does not tightly circumbscribe the training seizure vectors.

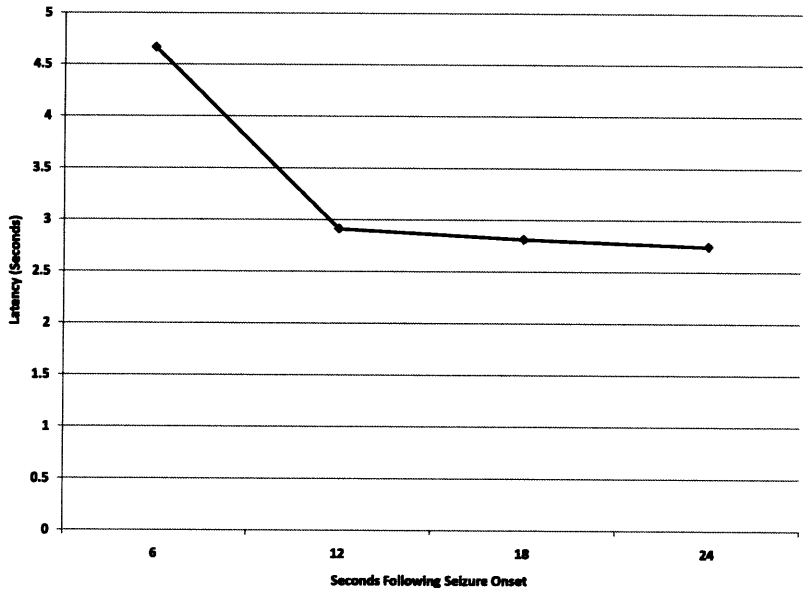

Figure 5-13: Including seizure vectors derived from the first $S > 12$ seconds of each training seizure does not improve seizure detection delay significantly. For $S > 12$ seconds the detector's mean latency is less than 3 seconds.

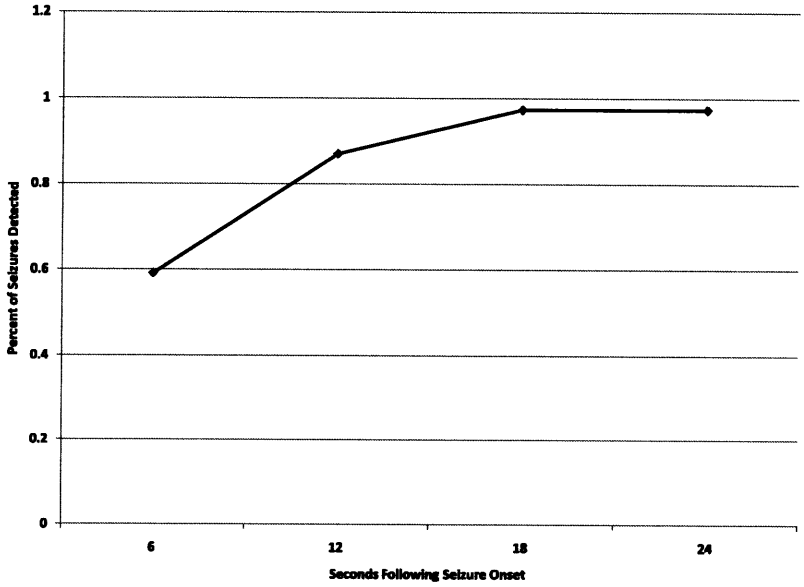

Figure 5-14: Including seizure vectors derived from the first $S > 18$ seconds of each training seizure does not improve seizure detection rate significantly. For $S > 18$ seconds the detector recognizes more than 95% of all test seizures.

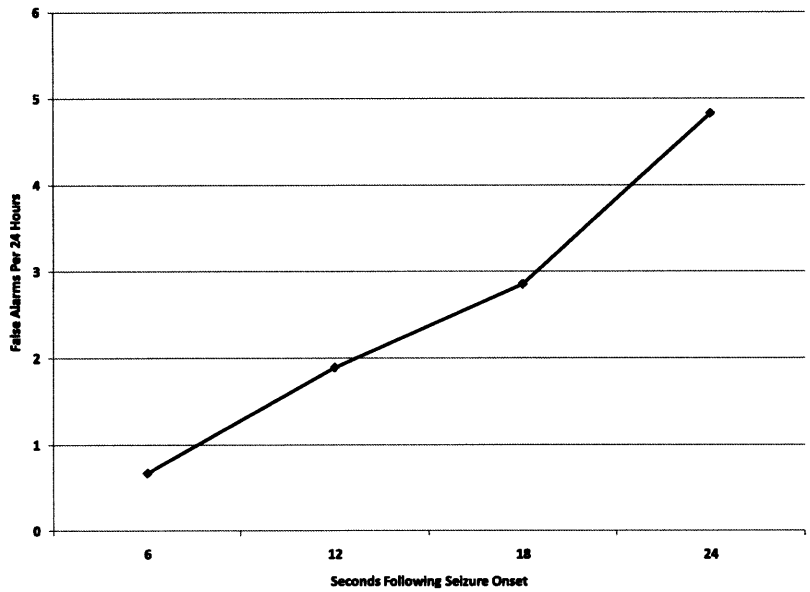

Figure 5-15: Increasing S increases the detector's false detection rate.

5.3 Patient-Specific and Non-specific Seizure Detection

In this section we compare the performance of our patient-specific detector to the performance of the Reveal algorithm, a patient non-specific detector. Recall that Reveal examines one EEG channel at a time and then employs hand-coded and neural network rules to determine whether features derived from a time-frequency decomposition of a channel are consistent with a seizure event. The thresholds for some of the neural network rules are determined using both archetypal seizures from individuals with epilepsy and background EEG from individuals without epilepsy.

5.3.1 Performance Comparison

We evaluated the performance of the Reveal algorithm under a *sensitive* and a *specific* setting. Under the sensitive setting, the Reveal algorithm declares a seizure whenever 5 seconds of EEG resembles training seizures with a score of 0.9 out of 1. Under the specific setting, a seizure is declared whenever 20 seconds of EEG resembles training

seizures with a score of 0.9 out of 1.

Under the sensitive setting, the Reveal algorithm detected 74% of 152 test seizures while our patient-specific method detected 96% of the test seizures. Figure 5-16 shows, for each patient, the number of seizures detected by our method and the Reveal algorithm.

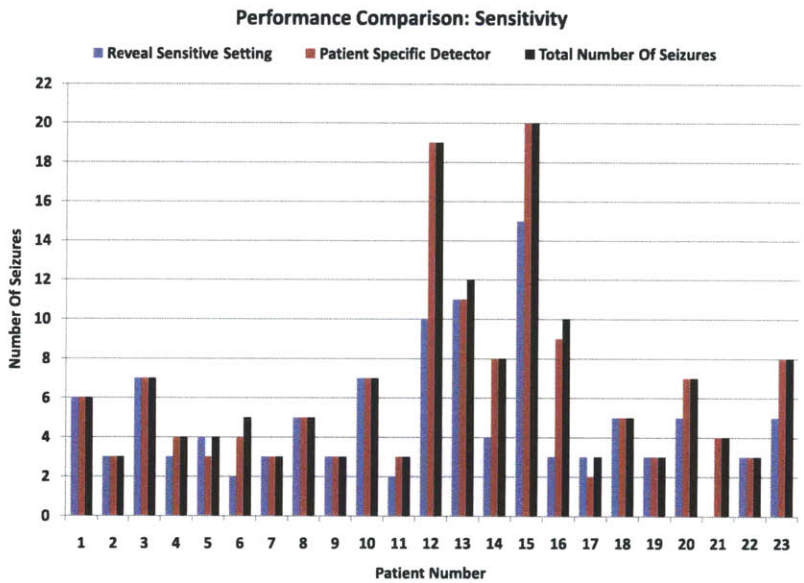

Figure 5-16: Comparison of the number of test seizures recognized by the Reveal (sensitive setting) and patient-specific algorithms. Black bars show the number of test seizures available for each subject, red bars show the number of test seizures recognized by the patient-specific algorithm, and blue bars show the number of seizures recognized by the Reveal algorithm. The Reveal algorithm detected 74% of 152 test seizures and our patient-specific method detected 96% of all test seizures.

The specificity of the Reveal algorithm was poor as illustrated by the number of false detections declared in a 24 hour period (Figure 5-17). For some patients the Reveal algorithm declared an excess of 100 false detections per 24 hour period. Figure 5-18 expands the y-axis of Figure 5-17 so that the false detection rates of the patient-specific detector are visible.

Under the specific setting, the sensitivity of the Reveal algorithm decreased to 61%. A comparison, for each patient, of the number of seizures detected by our

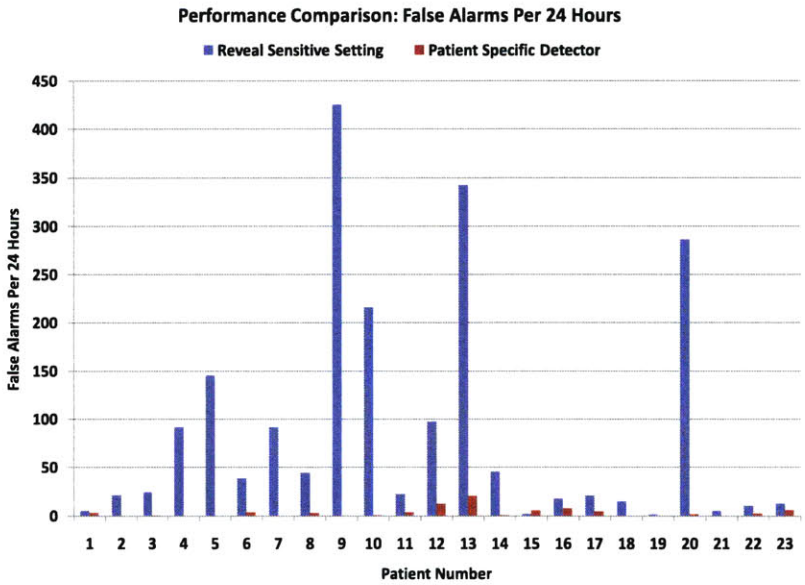

Figure 5-17: Comparison of the number of false detections declared by the Reveal (sensitive setting) and patient-specific algorithms. For some patients the Reveal algorithm declared an excess of 100 false detections per 24 hour period.

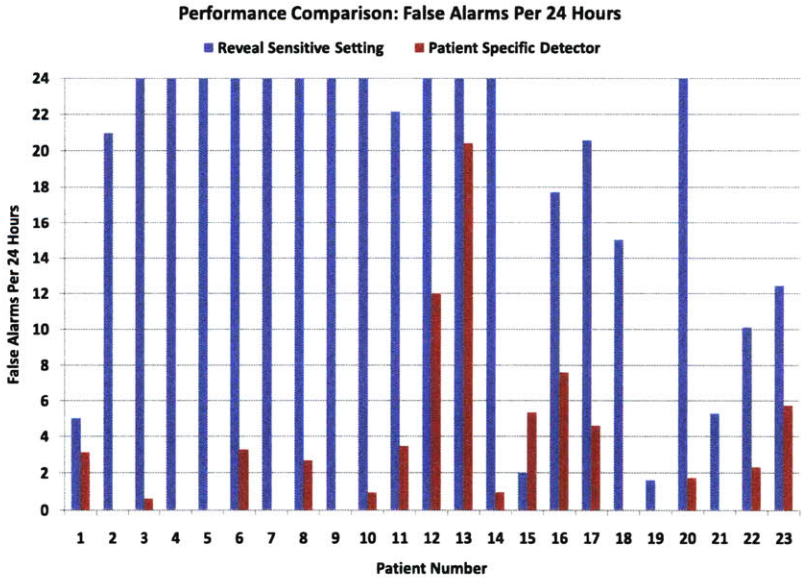

Figure 5-18: Comparison of the number of false detections declared by the Reveal (sensitive setting) and patient-specific algorithms. Expanded y-axis.

method and the Reveal algorithm is shown in Figure 5-19. As expected, the specificity of Reveal improves under this setting but it remains worse than that of the patient-specific detector as shown in Figures 5-20 and 5-21.

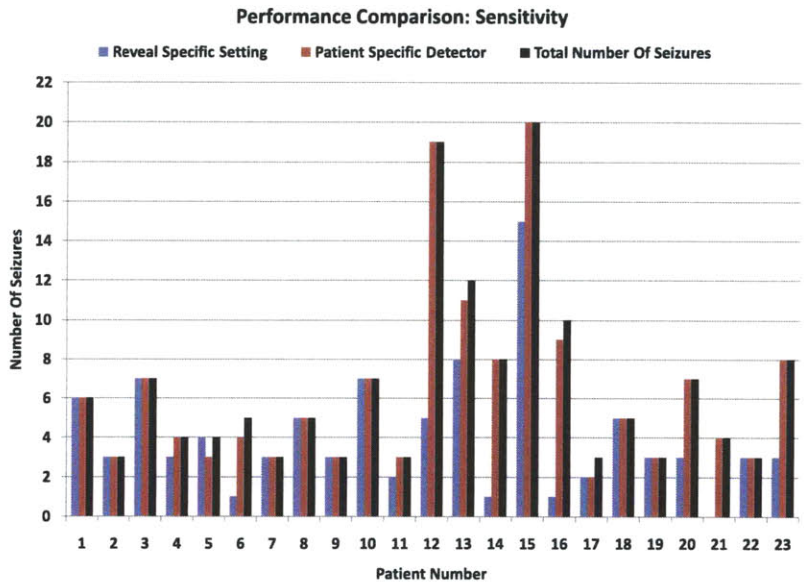

Figure 5-19: Comparison of the number of test seizures recognized by the Reveal algorithm (specific setting) and our patient-specific algorithm. Black bars show the number of test seizures available for each subject, red bars show the number of test seizures recognized by the patient-specific algorithm, and blue bars show the number of seizures recognized by the Reveal algorithm. The Reveal algorithm detected 61% of 152 test seizures and our patient-specific method detected 96% of all test seizures.

5.4 Case Studies

The following examples illustrate why a patient-specific approach can yield better seizure detection performance.

5.4.1 Latency

Figure 5-22 illustrates a seizure recorded within the scalp EEG of Patient G. The onset of the seizure occurs at 701 seconds and consists of beta band activity that is

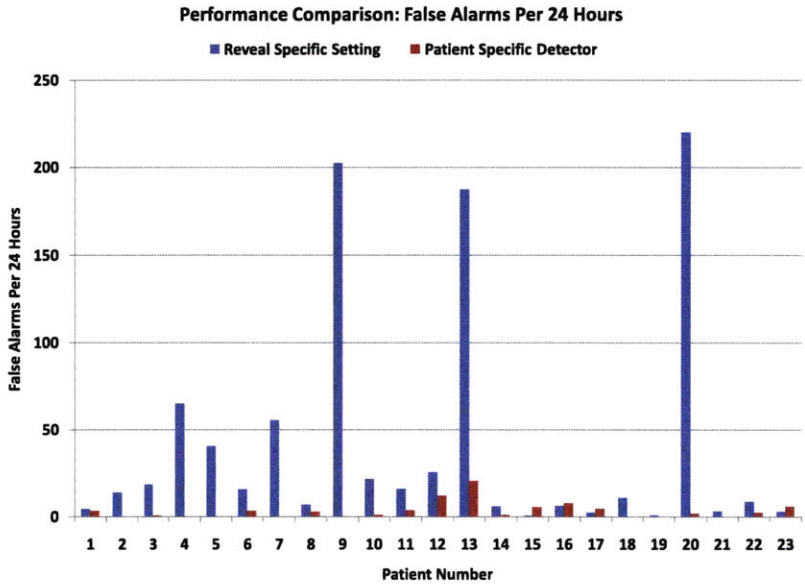

Figure 5-20: Comparison of the number of false detections declared by the Reveal (specific setting) and patient-specific algorithms. For some patients the Reveal algorithm declared an excess of 100 false detections per 24 hour period.

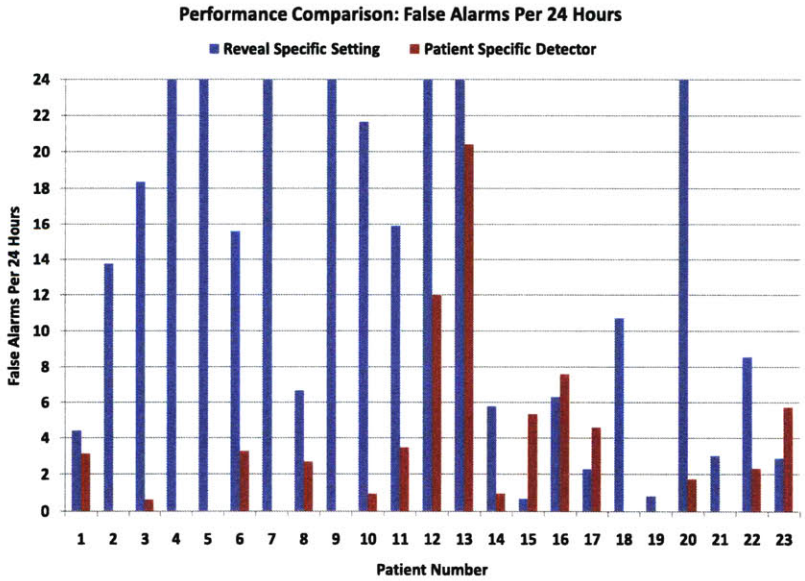

Figure 5-21: Comparison of the number of false detections declared by the Reveal (specific setting) and patient-specific algorithms. Expanded y-axis.

most prominent on the channels $\{F4-C4, C4-P4, F8-T8, T8-P8\}$. Twenty-nine seconds into the seizure, large amplitude theta band activity can be seen on all EEG channels as shown in Figure 5-23.

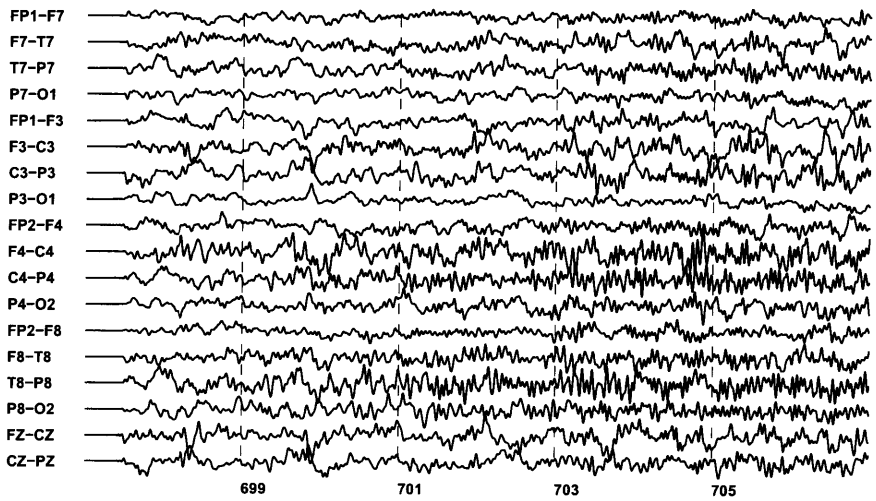

Figure 5-22: Seizure recorded within the scalp EEG of Patient G. The seizure, which begins at 701 seconds, consists of beta band activity that is most prominent on the channels $\{F4-C4, C4-P4, F8-T8, T8-P8\}$.

In this case, our patient-specific detector demonstrated improved detection latency relative to the Reveal algorithm. The patient-specific detector, on average, recognized that a seizure was ongoing within 7 seconds of its electrograpihc onset. On the other hand, the Reveal algorithm detected the presence of seizure activity 21.5 seconds following seizure onset. The Reveal algorithm could not detect the subtle EEG change associated with seizure onset, but could detect the high-amplitude, generalized activity that emerged towards the end of the seizure.

5.4.2 Sensitivity

Figure 5-24 illustrates a seizure recorded within the scalp EEG of Patient H. The seizure begins at 977 seconds with rapid eye-blinking that manifests as large, downward deflections on the channels $FP1-F3, FP2-F4, FP1-F7, FP2-F8$. Later,

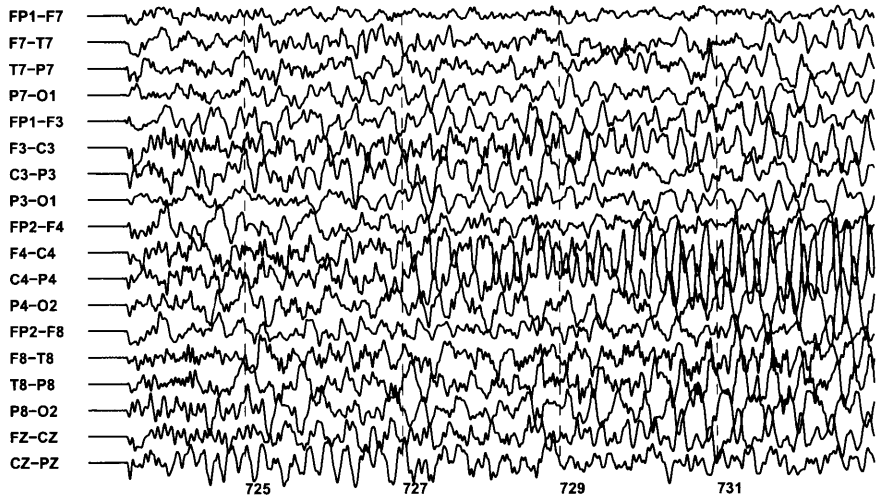

Figure 5-23: Large amplitude theta band activity can be seen on all EEG channels twenty-nine seconds after the onset of the seizure shown in Figure 5-22.

at 983 seconds, the eye-blinks are accompanied by theta band activity on the channels $\{T4 - T6, T6 - O2\}$.

In this case, the Reveal algorithm failed to detect any of the test seizures. A possible explanation for Reveal's poor sensitivity is the overlap of eye-blinks on the frontal channels (channels with the letter F) with seizure discharges on the temporal channels (channels with the letter T). Patient non-specific algorithms, such as the Reveal algorithm, are generally designed to classify EEG epochs containing rapid eye-blinks as corrupt and unlikely to contain a seizure. In contrast, our patient-specific detector could detect these seizures because it learned that the key feature that differentiates seizure-associated eye-blinks from others is the presence of seizure discharges on the temporal channels.

5.4.3 Specificity

Figure 5-26 illustrates a seizure recorded within the scalp EEG of Patient J. The onset of this seizure invovles a generalized spike at 12231 seconds that is followed by generalized rhythmic activity. Both the Reveal and patient-specific algorithms

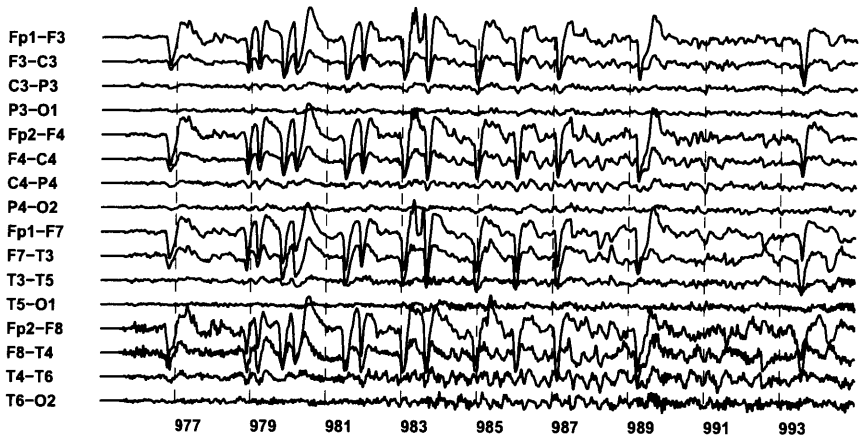

Figure 5-24: Seizure recorded within the scalp EEG of Patient H. The seizure, which begins at 977 seconds, consists of rapid eye-blinking that is later accompanied by theta band activity on the channels $\{T4 - T6, T6 - O2\}$ at 983 seconds.

were successful in detecting the onset of seizures recorded from Patient J. However, the Reveal algorithm exhibited poor specificity because it declared false detections whenever the rhythmic activity illustrated in Figure 5-26 occurred. This rhythmic activity is commonly observed within the awake EEG of Patient J. The patient-specific detector correctly classified the rhythmic activity as belonging to the non-seizure class because it had been labeled as such in the training set.

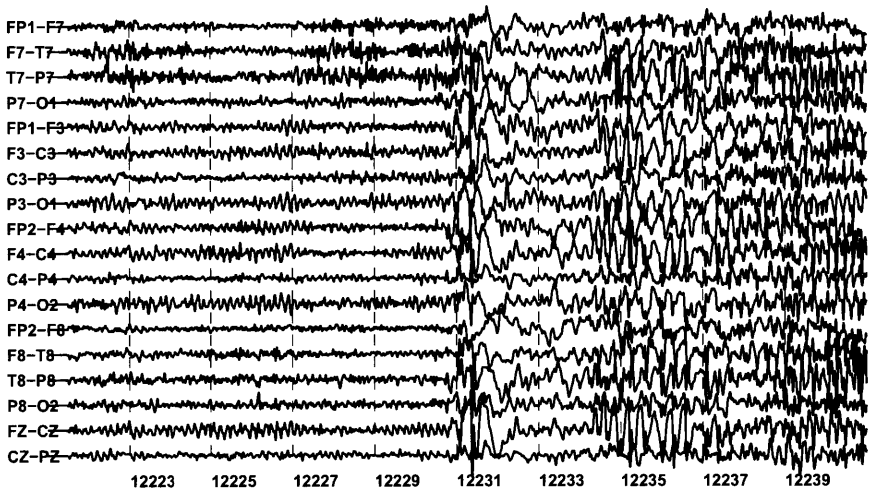

Figure 5-25: Seizure recorded within the scalp EEG of Patient J. The onset of this seizure invovles a generalized spike at 12231 seconds that is followed by generalized rhythmic activity.

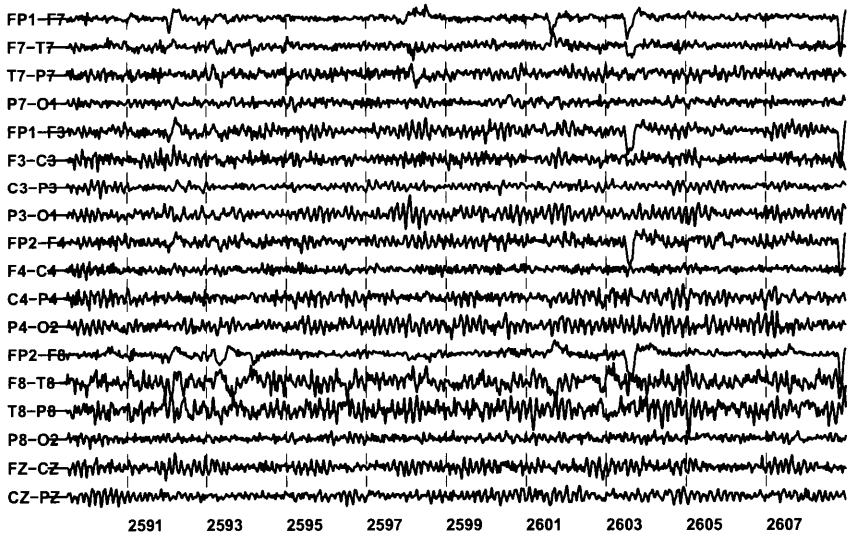

Figure 5-26: Rhythmic activity commonly observed within the awake scalp EEG of Patient J.

Chapter 6

Seizure Onset Detection Using Physiologic Signal Fusion

Using the scalp electroencephalogram (scalp EEG) to detect seizure onsets that are not associated with rhythmic EEG activity is challenging. In this chapter we illustrate how supplementing information extracted from the scalp EEG with information extracted from a second physiologic signal can improve the detection of these types of seizures.

6.1 Why Use a Second Physiologic Signal?

As discussed in Chapter 2, seizures involve the hyperactivity and hypersynchrony of a population of neurons. At the level of the scalp EEG, this coherent neural firing gives rise to rhythmic activity with a dominant frequency between 0-25 Hz. However, when the underlying neural hypersynchrony involves a neural network deep within the brain, the earliest scalp EEG changes may not reflect the underlying neuronal hypersynchrony, but physical consequences of the seizure such as rapid eye-blinks or muscle contractions. Once the entrained nerual mass is large enough, rhythmic activity reflective of neuronal hypersynchrony becomes manifest within the scalp EEG.

Detecting seizures of this kind using the scalp EEG is challenging since the EEG changes observed following seizure onset are also seen routinely during non-seizure

states. In fact, classic seizure detection methods insert processing stages to remove such activity [34, 48] because detecting it would result in poor specificity.

In order to detect these types of seizures a detector requires information beyond that within the scalp EEG to ascertain whether or not a seizure is taking place. The additional information can be derived using a second physiologic signal whose dynamics are influenced by the seizure. The second physiologic signal and the scalp EEG will complement each other and improve seizure onset detection if the changes, in each of these signals, that suggest the onset of a seizure rarely coincide during non-seizure states and often coincide at the time of an actual seizure. Patient-specificity remains essential to the success of this approach since the manner with which the scalp EEG and the secondary signal change during seizure and non-seizures states varies across patients.

6.2 The Electrocardiogram as a Second Signal

A variety of physiologic signals may be used to supplement the scalp EEG in order to improve seizure detection performance. For example, seizures resulting in repetitive motor activity may become readily detectable if scalp EEG data is supplemented with accelerometer sensor data [40]. For other types of seizures, especially those originating within or spreading to the temporal lobes, seizures are associated with electrocardiographic changes. The most common ECG change associated with seizures is a heart rate acceleration (tachycardia) [33].

6.3 Patient-Specific, EEG-ECG-based Seizure Detection

As was the case in Chapter 3, a classifier will be used to determine whether an observed feature vector is representative of an individual's seizure or a non-seizure state; however, in this chapter, the feature vector will contain information extracted from both the EEG and ECG.

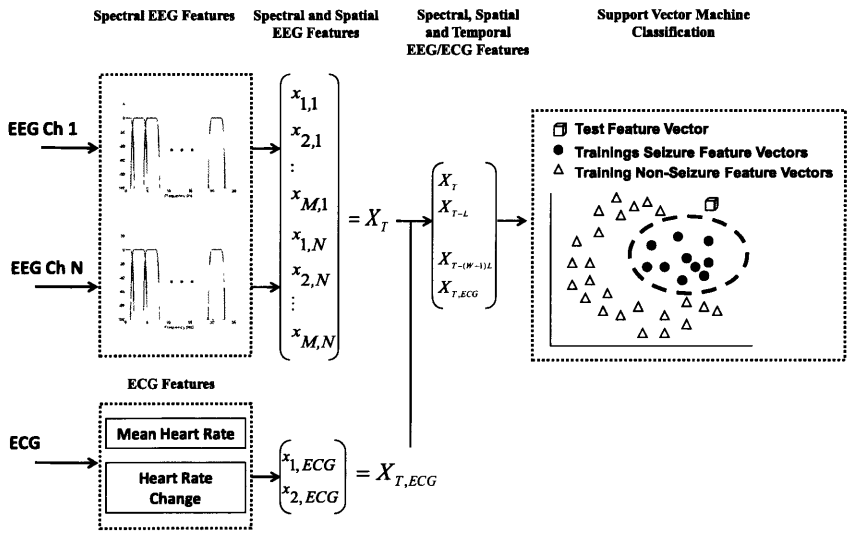

Figure 6-1: Block diagram of patient-specific seizure onset detector that combines features extracted from the EEG and ECG signals.

The block diagram in Figure 6-1 illustrates the architecture of the combined EEG and ECG patient-specific seizure onset detector. The detector passes a L second epoch from $N = 16$ EEG channels and a lead-II ECG channel through a feature extractor. The feature extractor uses a filterbank with $M = 8$ filters to compute for each EEG channel the energy in eight overlapping frequency bands; the frequency bands are each 3 Hz wide and span 0-25Hz. The 8 features extracted from each of the 16 channels are then concatenated to form a 128 dimensional feature vector X_T that automatically captures the spectral and spatial correlations between channels.

The feature extractor uses the ECG channel to extract a feature vector $X_{T,ECG}$ that is composed of two features. The first feature, $X_{1,ECG}$, corresponds to the mean heart rate within the L second epoch. The second feature, $X_{2,ECG}$, corresponds to the difference between the instantaneous heart rate measured at the start and end of the L second epoch. Outside the context of a seizure, $X_{2,ECG} \approx 0$ as successive, instantaneous heart rate measurements within the L second epoch will oscillate about a mean heart rate. In contrast, for patients experiencing seizure-associated tachycar-

dia, $X_{2,ECG} > 0$ since the instantaneous heart rate at the end of the epoch will exceed that at the start of the epoch.

Next, $W = 3$ EEG feature vectors X_T, X_{T-L}, and X_{T-2L} as well as the ECG feature vector $X_{T,ECG}$ are concatenated to form a single feature vector $\mathcal{X}_T$. This feature vector encodes the time evolution of spectral and spatial properties of the scalp EEG as well as heart rate and heart rate change information extracted from the ECG. This encoding is *automatically* generated without user intervention. Finally, the feature vector $\mathcal{X}_T$ is assigned to the seizure or non-seizure state using a two-class support vector machine (Chapter 3).

6.4 Case Studies

In this section we present two case studies that illustrate how a detector that combines information extracted from the scalp EEG and surface ECG can perform better than a detector that only uses information extracted from the scalp EEG.

6.4.1 Data

The data set used to assess our detection methodology was gathered from two patients. From the first patient, an adult woman with refractory complex partial seizures, 66 hours of continuous EEG-ECG that includes 10 seizures were collected. From the second patient, an adult male with a childhood-onset epilepsy syndrome known Lennox-Gastaut, 34 hours of continuous EEG-ECG that includes 3 seizures were collected.

6.4.2 Case 1

Figure 6-2 shows an EEG tracing of a typical seizure recorded from the first subject. The onset of the seizure, at 1486 seconds, involves rapid eye-blinking (eye flutter), which manifests in the EEG as high-amplitude deflections on the channels $\{FP1 - F3, FP2 - F4, FP1 - F7, FP2 - F8\}$.

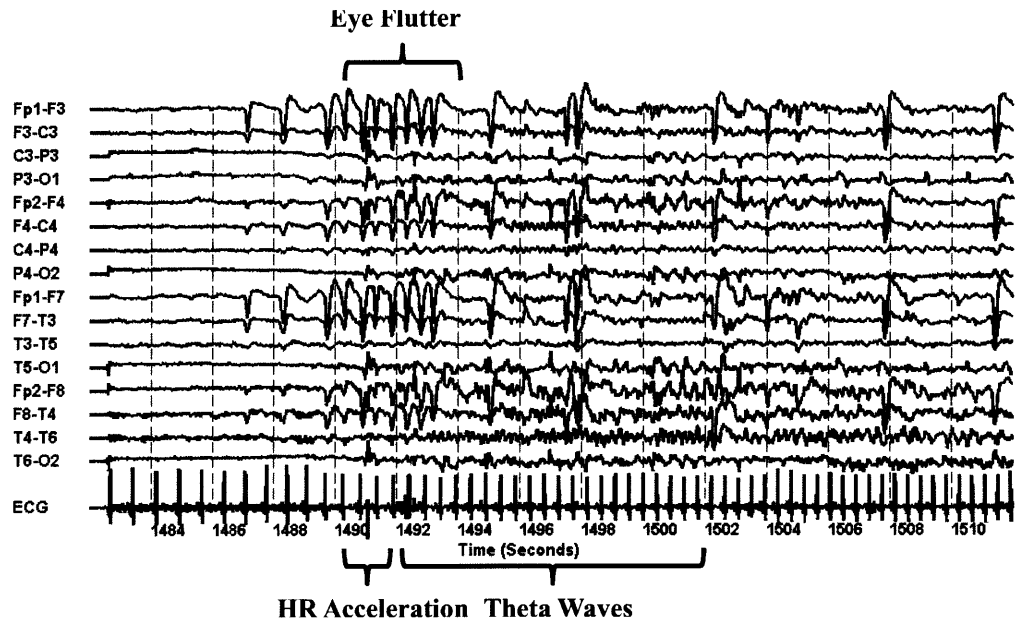

Figure 6-2: Example of a seizure within the scalp EEG of Case 1. The seizure, which begins at 1486 seconds, involves rapid eye-blinking that results in downward deflections on frontal EEG channels (e.g. $\{FP1 - F3, FP2 - F4\}$). Coincident with the onset of rapid eye-blinking, the patient's heart rate accelerates as shown in Figure 6-3. Later, at 1492 seconds, a 3-4 Hz theta wave appears on the EEG channel $T4 - T6$.

With the onset of eye-flutter, the patient's heart rate accelerates as can be seen on the bottom channel in Figure 6-2 and in the heart rate profile in Figure 6-3. Finally, at 1492 seconds, 3-4 Hz theta waves appear on the EEG channel $T4 - T6$, while the patient's heart rate remains elevated. Electrographic signs of the seizure end around 1510 seconds, but clinical symptoms of the seizure persist for another 2-3 minutes.

Performance Comparison

To evaluate the utility of combining EEG and ECG information in this case, we compared the performance of two detectors. One detector classifies a feature vector synthesized using EEG and ECG as in Figure 6-1, and the other classifies a feature vector synthesized using EEG features as in Figure 3-17 of Chapter 3. Both detectors are trained on the $S = 20$ seconds following the onset of eye-flutter within training seizures, and both detectors process $L = 5$ second epochs. The SVM used within the detector that combined EEG and ECG features was trained with $\gamma = 0.007$ and

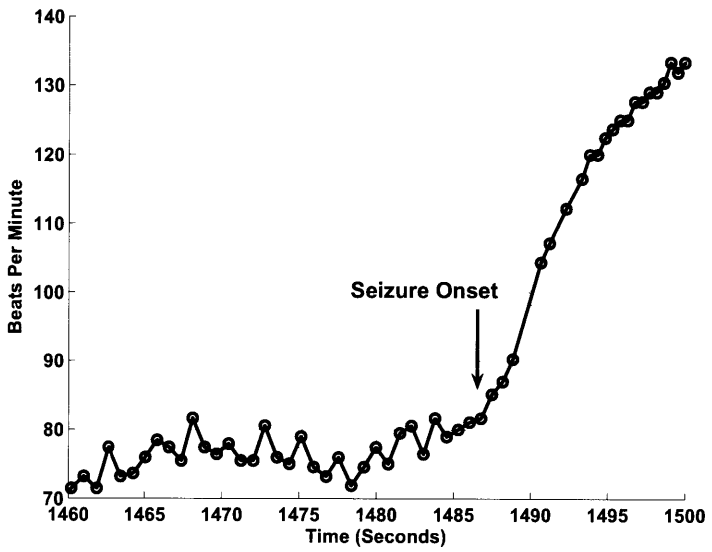

Figure 6-3: Seizure onset, at 1486 seconds, is associated with an acceleration of the patient's heart rate.

$J = 5$, while the SVM used within the detector that only used EEG was trained with $\gamma = 0.003$ and $J = 5$. The SVM parameters were chosen to bias each detector towards recognizing seizure onset with a short delay.

Figure 6-4 illustrates the mean, minimum, and maximum latency with which each detector declares the ten test seizures, and Figure 6-5 shows the false alarm rate exhibited by each detector while analyzing the 66 hours of non-seizure test data.

In this case, the detector that combined EEG and ECG information detected seizures with a mean detection latency (2.7 seconds) that is shorter than that of the detector which relied solely on EEG information (4.2 seconds). Furthermore, the detector that fused EEG and ECG information had a false detection rate (less than 5 false detections per 24 hours) that is lower than that of the detector which relied on EEG information alone (more than 10 false detections per 24 hours).

6.4.3 Case 2

Figure 6-6 shows an EEG tracing of a typical seizure recorded from the second subject. The onset of the seizure, at 56 seconds, consists of an electrodecrement involving all

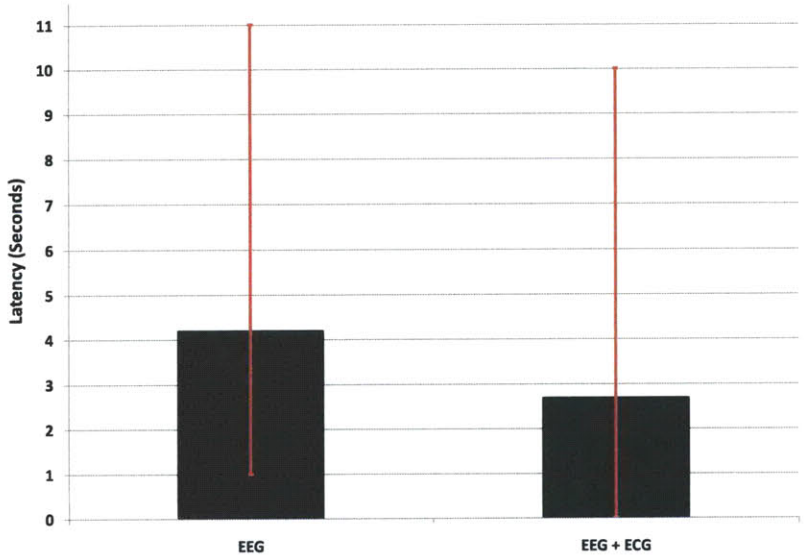

Figure 6-4: Comparison of the minimum, maximum, and mean detection delays of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a shorter mean detection delay.

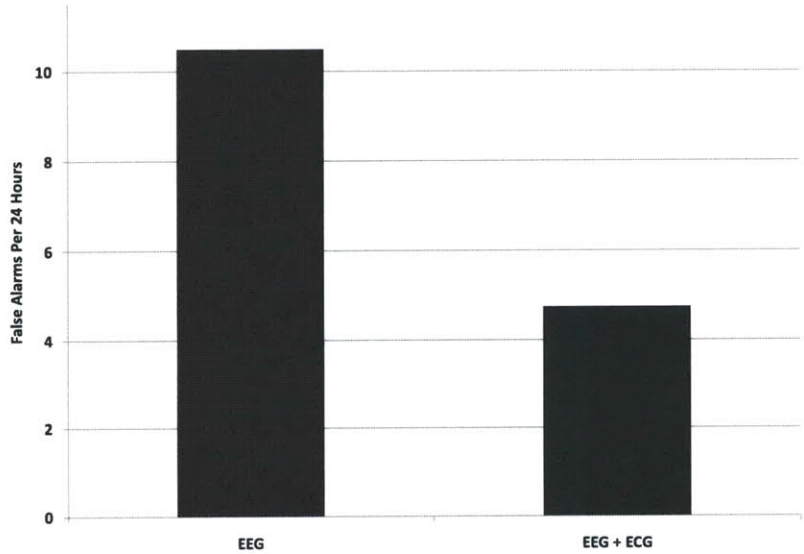

Figure 6-5: Comparison of the false detection rates of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a smaller false detection rate.

EEG channels and lasting for 12 seconds. Next, at 68 seconds, 1-2 Hz generalized, rhythmic activity develops and persists for another 30 seconds. Clinically, throughout the period of the electrodecrement the patient experiences tonic contraction of major muscle groups.

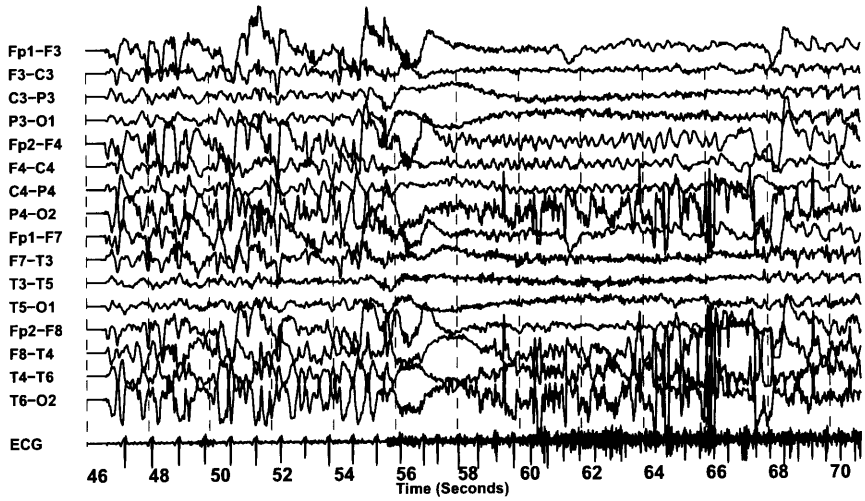

Figure 6-6: Example of a seizure within the scalp EEG of Case 2. The seizure, which begins at 56 seconds, involves a 12 second period of low-amplitude EEG activity across most EEG channels. At the same time, the patient's heart rate accelerates as shown in Figure 6-7. Later, at 68 seconds, 1-2 Hz generalized, rhythmic activity develops.

With the onset of the electrodecrement, the patient's heart rate accelerates as can be seen on the bottom channel in Figure 6-6 and in the heart rate profile in Figure 6-7. Throughout most of the period of generalized, rhythmic activity the patient's heart rate remains elevated.

Performance Comparison

To evaluate the utility of combining EEG and ECG information in this case, we compared the performance of two detectors. One detector classifies a feature vector synthesized using EEG and ECG as in Figure 6-1, and the other classifies a feature

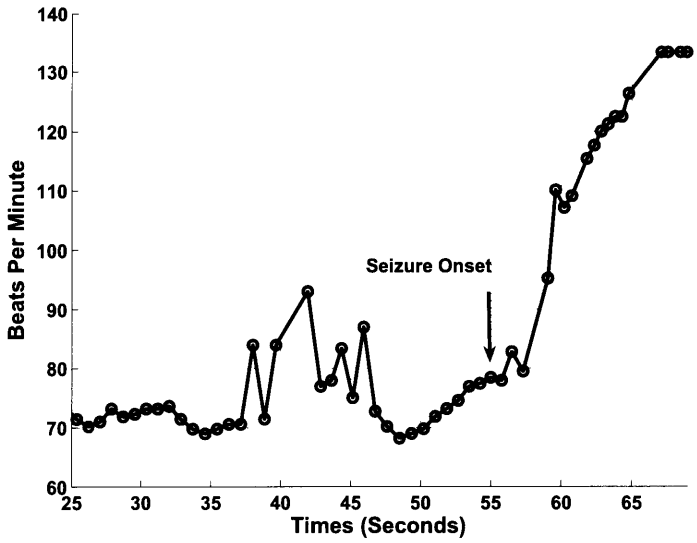

Figure 6-7: Seizure onset, at 56 seconds, is associated with an acceleration of the patient's heart rate.

vector synthesized using EEG features as in Figure 3-17 of Chapter 3. Both detectors are trained on the $S = 20$ seconds following the onset of the electrodecrement within training seizures, and both detectors process $L = 2$ second epochs. Similar to Case 1, the SVM used within the detector that combined EEG and ECG features was trained with $\gamma = 0.007$ and $J = 5$, while the SVM used within the detector that only used EEG was trained with $\gamma = 0.003$ and $J = 5$.

Figure 6-8 illustrates the mean, minimum, and maximum latency with which each detector declares the 3 test seizures. Figure 6-9 shows the false alarm rate exhibited by each detector while analyzing the 34 hours of non-seizure test data.

In this case, the detector that combined EEG and ECG information detected seizures with a mean detection latency (6 seconds) that is comparable to that of the detector which relied solely on EEG information (5 seconds). However, the detector that fused EEG and ECG information had a false detection rate (4 false detections per 24 hours) that is lower than that of the detector which relied on EEG information alone (more than 10 false detections per 24 hours).

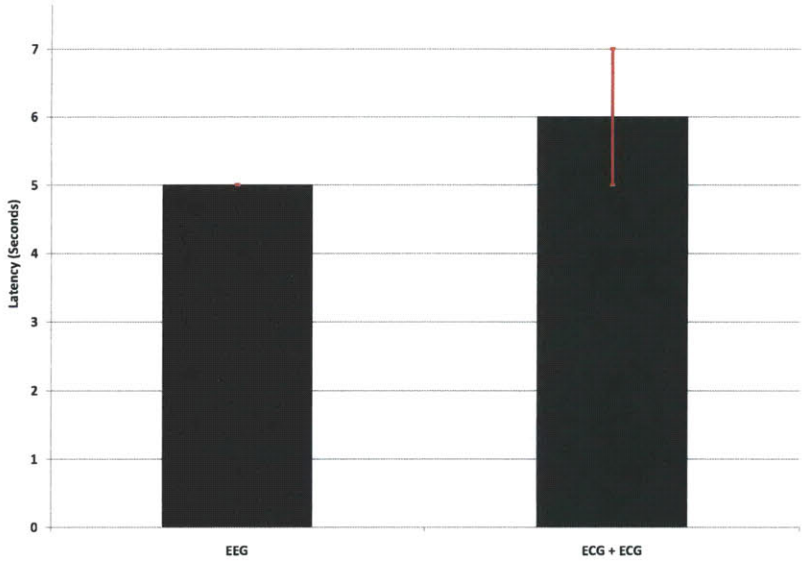

Figure 6-8: Comparison of the minimum, maximum, and mean detection delays of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The two detectors have comparable seizure detection delays.

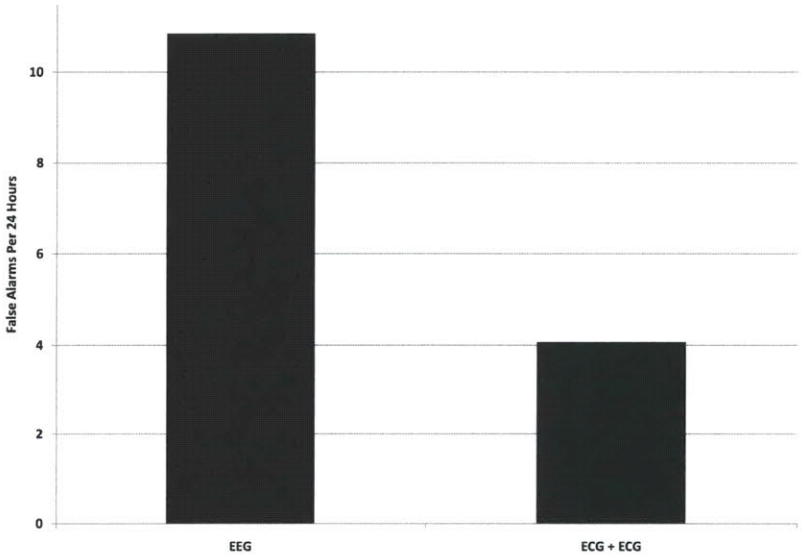

Figure 6-9: Comparison of the false detection rates of two detectors. One detector classifies a feature vector composed solely of EEG features, while the other uses a feature vector that combines EEG and ECG features. The detector that fuses features extracted from the EEG and ECG signals has a smaller false detection rate.

6.5 Importance of Patient-Specificity

We evaluated the performance of the Reveal algorithm (section 1.4.1) on Cases 1 and 2 to demonstrate the importance of both physiologic signal fusion and patient-specificity to detecting the class of seizures discussed in this chapter. The Reveal algorithm is only capable of processing EEG signals and is patient non-specific. When tested on Case 1, the Reveal algorithm detected 1 of 10 seizures. When tested on Case 2, the Reveal algorithm missed the electrodecrement in 2 out of 3 test seizures, and instead detected the generalized 1-2 Hz activity that develops near the end of the test seizures.

Chapter 7

Seizure-Triggered Vagus Nerve Stimulation

In this chapter we illustrate the feasibility of applying the seizure detection methodology developed in Chapters 3 and 6 to the initiation of a delay-sensitive therapeutic application. More specifically, the delivery of vagus nerve stimulation following seizure onset.

7.1 Vagus Nerve Stimulation

For most of the history of epilepsy, pharmacologic, surgical, and diet-based treatments were the only therapeutic options available to patients with refractory epilepsy. In 1997, the United States Food and Drug Administration approved the use of Vagus Nerve Stimulation (VNS) as an adjunctive therapy for the treatment of refractory seizures; the Vagus Nerve Stimulator became the first implantable device for the treatment of epilepsy [3, 50, 51]. The VNS pulse generator is implanted in an infraclavicular, subcutaneous pocket and delivers stimuli to the midcervical portion of the left vagus nerve.

Vagus Nerve Stimulation therapy is delivered in two modes. In automatic mode, the implanted pulse-generator automatically delivers vagus nerve stimuli at programmed intervals. In on-demand mode, the patient or their caregiver initiates vagus

nerve stimulation in response to symptoms of a seizure. Initiating on-demand stimulation requires holding a permanent magnet for 1-2 seconds over the implanted pulse generator.

Several clinical trials have been conducted to assess the therapeutic efficacy of automatic mode vagus nerve stimulation. The E03 and E05 trials demonstrated that automatic mode VNS reduces seizure frequency by more than 50% in more than 20% of patients within 3 months of device implantation [21, 23]. The XE5 trial demonstrated that the efficacy of automatic vagus nerve stimulation is both long-lasting and improves significantly with time [11]. Based on these and other studies, vagus nerve stimulation was deemed a safe and effective adjunctive therapy for intractable epilepsy.

Pre-clinical studies involving rat [67] and canine [68] models of seizures demonstrated that initiating vagus nerve stimulation during a seizure could terminate it or lessen its severity. Clinical studies designed to quantify the therapeutic impact of on-demand mode vagus nerve stimulation have also been conducted. Hammond [22] recorded an EEG tracing illustrating the abrupt termination of an electrographic and behavioral seizure following the initiation of on-demand VNS. Morris [25] retrospectively analyzed seizure diaries from the E04 trial [32] and noted that 53% of patients capable of receiving on-demand stimulation reported experiencing seizure termination or diminution. Similarly, Boon [4] noted that two-thirds of patients receiving on-demand stimulation reported being able to interrupt seizures. While the results in [25] and [4] are encouraging, it is important to note that they were determined using patient and caregiver accounts of seizures as opposed to EEG tracings illustrating seizure termination or diminution.

A significant proportion of VNS therapy patients depend on others to initiate on-demand stimulation because the physical or cognitive symptoms of a seizure leave them unable to do so themselves[4]. Depending on a caregiver to initiate on-demand stimulation has two consequences: 1) it denies patients the therapeutic benefit of on-demand stimulation in the absence of caregivers, and 2) it leads to an inconsistent ability to terminate or attenuate seizures since caregivers may not always be able

to initiate on-demand stimulation immediately upon the clinical (let alone the electrographic) onset of a seizure. Since there is anecdotal evidence suggesting that the likelihood of affecting seizure progression decreases as the time between seizure onset and the start of stimulation increases [22], this delay in detection may have significant consequences.

A system that automatically initiates on-demand mode VNS following computerized seizure onset detection could relieve caregivers, increase a patient's sense of independence and security, and more frequently terminate or diminish seizure symptoms in individuals with seizures that respond to on-demand stimulation. The system may also benefit patients who never realized they were on-demand responders since neither they nor their caregivers were able to initiate stimulation following the onset of a seizure.

In this chapter, we describe the design and clinical evaluation of a non-invasive computerized system that automatically initiates on-demand VNS following the detection of a prespecified seizure. The computerized system detects these events through the patient-specific seizure detection methodology developed in Chapters 3 and 6.

7.2 Methods

7.2.1 System Overview

Figure 7-1 shows a block diagram of the computerized system; all the components of the computerized system are external to the patient. The computerized system is composed of a commercial acquisition system (Digitrace 1800 Plus from SleepMed Inc.) that collects the EEG and ECG of a patient, a computer that analyzes the EEG and ECG in real-time using the algorithms presented in Chapters 3 and 6, and an electromagnet that is worn by the patient and positioned so that it rests over the implanted VNS pulse generator.

When the computerized system detects the onset of a seizure, it energizes the

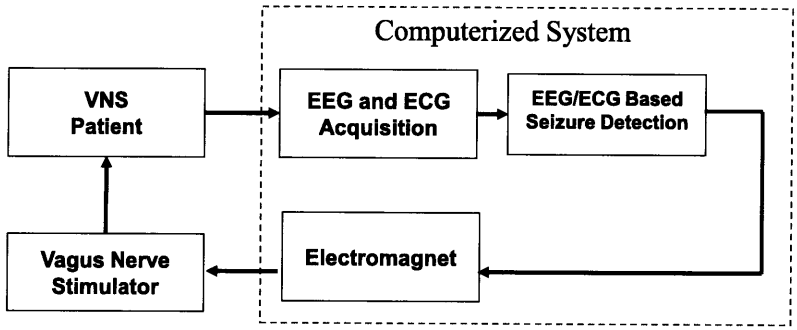

Figure 7-1: Block diagram of closed-loop vagus nerve stimulation system.

electromagnet worn by the patient. The magnetic field produced by the electromagnet triggers the implanted generator to initiate on-demand stimulation of the vagus nerve. The electromagnet initiates on-demand stimulation through the same mechanism that is triggered when a permanent magnet is held over the implanted generator.

7.2.2 Study Protocol

The clinical evaluation of the computerized system followed a protocol approved by the Committee on Clinical Investigations at the Beth Israel Deaconess Medical Center (BIDMC), Boston, Massachusetts, USA, and the Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts, USA. Study participants were adult, long-term users of VNS who reported no adverse effects from on-demand VNS. Participants were admitted to the BIDMC General Clinical Research Center (GCRC) for a period lasting up to five days. Admission to the GCRC took place after study staff obtained informed consent, and confirmed that energizing the computerized system's electromagnet reliably initiates on-demand stimulation. No changes to antiepileptic drug (AED) regimens or VNS stimulation parameters were made during the study period.

During the admission period Video, EEG, and ECG signals were recorded. After

the first 24 hours of the study, an electroencephalographer reviewed the collected EEG and highlighted seizure events. Next, the computerized system was trained to detect the identified seizure events using the EEG gathered in the first 24 hours of the study. On subsequent days of the study, the computerized system was set to automatically activate the participant's VNS generator whenever real-time detection of the target seizure occurred.

7.3 Case Studies

In this section we examine three cases that involved the initiation of on-demand vagus nerve stimulation in response to computerized seizure onset detection.

7.3.1 Patient A

Patient A Medical History

Patient A and Case 1 of Chapter 6 correspond to the same individual. Patient A is a 39-year-old woman with a long history of refractory complex partial seizures. At the time of admission to our study, she was experiencing 30-40 seizures per month. Patient A's seizures last for 1-2 minutes and consist of repeatedly asking questions and blank stares; her seizures are not accompanied by automatisms or tonic-clonic movements. Following a seizure, patient A is confused for a few minutes and requires reorientation by friends or family. Since Patient A does not experience a warning at the onset of a seizure, she does not use the on-demand mode VNS.

Patient A Seizures

Figure 7-2 illustrates an EEG tracing of a typical seizure recorded from patient A. The onset of the seizure, at 1486 seconds, involves rapid eye-blinking (eye flutter) which manifests in the EEG as high-amplitude deflections on the channels $\{FP1 - F7, FP1 - F3, FP2 - F8, FP2 - F7\}$. Coincident with the onset of eye-flutter is an acceleration of the patient's heart rate, as seen on the bottom channel in Figure

7-2, and plotted in Figure 7-3. Finally, at 1492 seconds, 3-4 Hz theta waves appear in the EEG on the temporal channel $T4 - T6$. Electrographic signs of the seizure end around 1510 seconds, but clinical symptoms of the seizure persist for another 2-3 minutes.

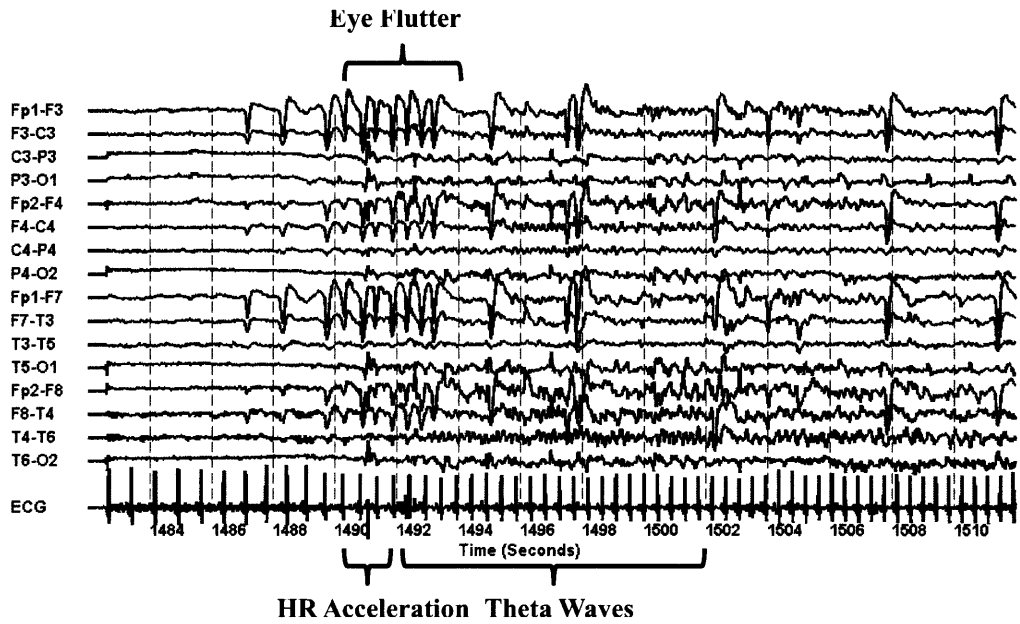

Figure 7-2: Example of a seizure within the scalp EEG of Patient A. The seizure, which begins at 1486 seconds, involves rapid eye-blinking that results in downward deflections on frontal EEG channels (e.g. $\{FP1 - F3, FP2 - F4\}$). Coincident with the onset of rapid eye-blinking, the patient's heart rate accelerates as shown in Figure 7-3. Later, at 1492 seconds, a 3-4 Hz theta wave appears on the EEG channel $T4 - T6$.

A total of 5 prerecorded seizures as well as two days of prerecorded awake and sleep EEG were available to train a patient-specific detector to differentiate between the seizure and non-seizure EEG of Patient A. The training seizure and non-seizure data were recorded six months prior to patient A's admission to the study.

Patient-Specific, EEG-ECG-based Initiation of VNS

The earliest events that define Patient A's seizures are eye-flutter within the EEG and a heart rate acceleration within the ECG. In order to reliably initiate VNS following the coincidence of these events we used the detector architecture shown in Figure 7-4. This architecture is similar to that developed in Chapter 6 except that the feature

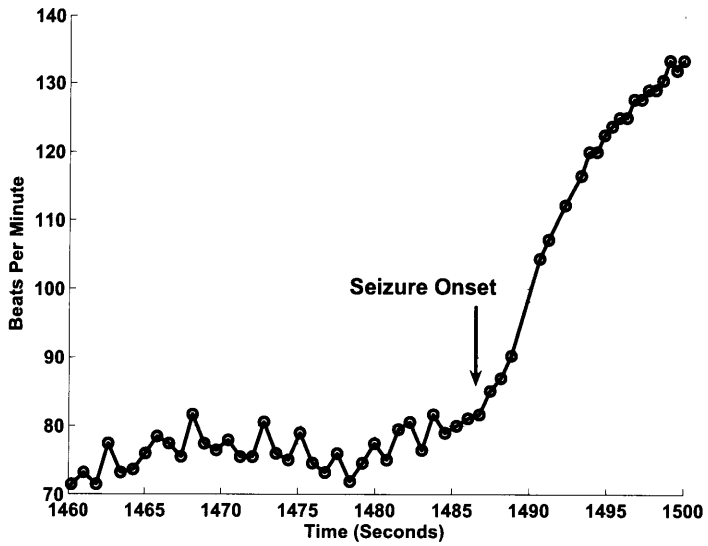

Figure 7-3: Seizure onset, at 1486 seconds, is associated with an acceleration of the patient's heart rate.

vector used by the detector involves only spectral and spatial information extracted from the EEG (no time evolution) and heart rate and heart rate change information extracted from the ECG.

Figure 7-5 illustrates a seizure detected using the architecture illustrated in Figure 7-4. Seizure onset begins with rapid eye blinking at 992 seconds. The detector declared seizure onset at 996 seconds, which is 4 seconds after the onset of eye-flutter and 2 seconds before the onset of theta wave activity. Beginning at 1002 seconds a spike-train appears on the "VNS" channel confirming the initiation of vagus nerve stimulation. In this case, VNS was initiated 1 second after the onset of theta wave activity. The heart rate acceleration that triggered the detector is shown in Figure 7-6; note its similarity to the training heart rate profile illustrated in Figure 7-3. The computerized system automatically detected and initiated VNS on four other seizures as described in [57].

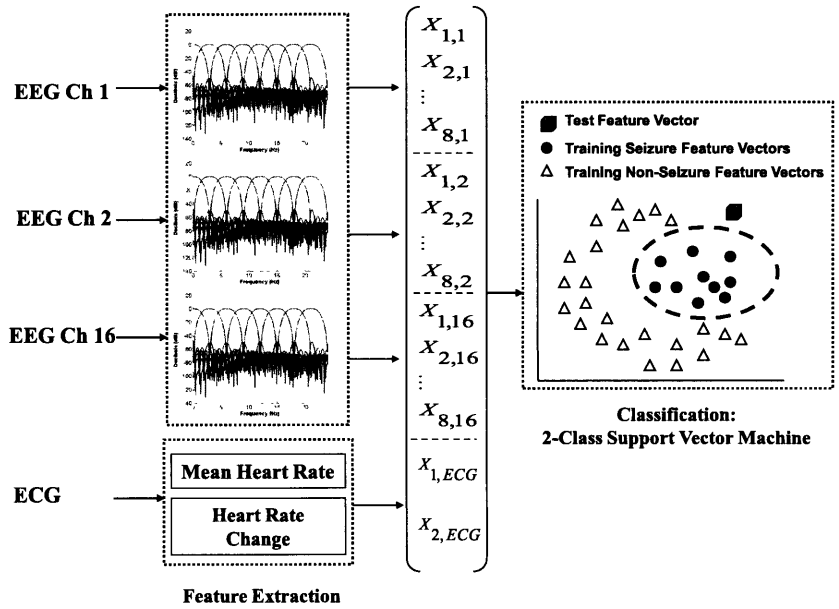

Figure 7-4: Block diagram of patient-specific seizure onset detector that combines features extracted from the EEG and ECG signals.

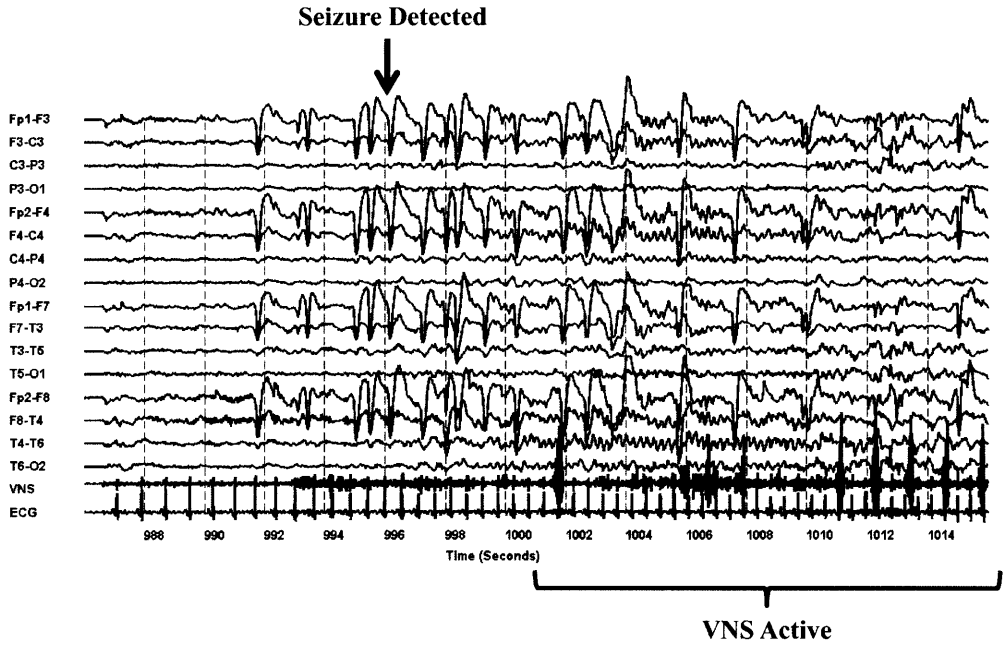

Figure 7-5: Initiation of VNS following computerized detection of the onset of a seizure from Patient A. Seizure onset begins with rapid eye blinking at 992 seconds. The detector declared seizure onset at 996 seconds, and initiated VNS in response. Beginning at 1002 seconds, a spike-train appears on the "VNS" channel confirming the initiation of vagus nerve stimulation.

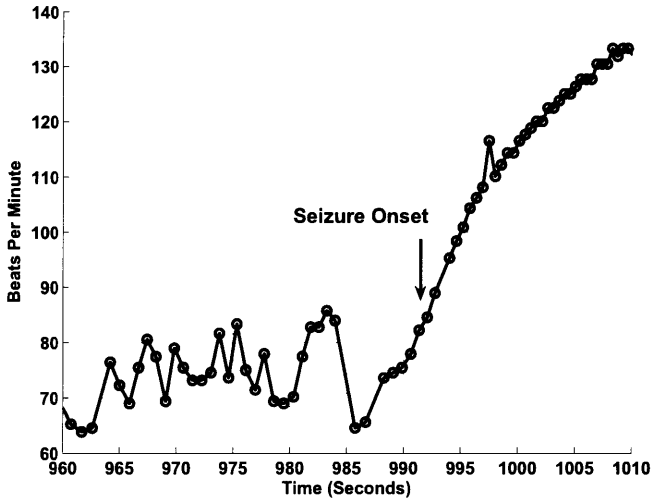

Figure 7-6: Seizure onset, at 992 seconds, is associated with an acceleration of the patient's heart rate.

Comparing Stimulated and Non-stimulated Seizures

The computerized system did not initiate on-demand VNS following the onset of one of patient A's seizures. This was done intentionally so that we could compare the electrographic and behavioral character associated with stimulated seizures and a non-stimulated seizure.

We did not detect a difference in the electrographic or behavioral character of stimulated and non-stimulated seizures. Three possible explanations for this include: 1) the on-demand stimulus was not initiated early enough in the course of the seizure, 2) the on-demand VNS stimulus parameters (pulse current, frequency, and duration) were not set appropriately 3) Patient A is not a responder to on-demand mode VNS.

7.3.2 Patient B

Patient B Medical History

Patient B is a 41-year-old woman with a long history of simple partial, complex partial and secondarily generalized seizures. At the time of her admission to our study, she was experiencing 5-6 seizures per month. Patient B's secondarily generalized seizures

consist of a tonic phase followed by a clonic phase, and finally a post-ictal phase. The tonic-clonic phase lasts for 1 minute. The post-ictal phase, which is characterized by generalized EEG slowing, lasts for 30 minutes. According to patient B's caregiver, initiating VNS early in the ictal phase results in reduced anxiety and confusion during the post-ictal phase.

Patient B Seizures

During the study period, patient B experienced two secondarily generalized seizures. The first was used to train the computerized system, and the second to investigate the impact of initiating on-demand VNS early in the course of patient B's seizure.

Figure 7-7 illustrates an EEG trace of the first seizure. The seizure begins following the 1340 seconds with an electrodecrement most prominent on the occipital channels $\{P3 - O1, P4 - O2\}$. Next, at 1348 seconds, a 3-5 Hz occipital rhythm emerges from the electrodecrement and rapidly generalizes. By 1360 seconds, muscle activity associated with the tonic-clonic phase of the seizure becomes visible. The tonic-clonic phase of the first seizure lasted for 55 seconds and the post-ictal phase lasted for 29 minutes. During the post-ictal phase, the patient was unaware of her surroundings and very anxious. The computerized system, using a detector of the form illustrated in Figure 7-8, was trained to recognize the occipital rhythm and muscle activity associated with the tonic-clonic phase of the seizure.

Patient-Specific, EEG-based Initiation of VNS

Figure 7-9 shows an EEG trace of the seizure that triggered the system to initiate on-demand VNS. The seizure begins with an electrodecrement at 1485 seconds. The computerized system failed to detect the onset of the focal occipital rhythm at 1490 seconds. Instead, at 1501 seconds, it detected the muscle activity associated with the tonic-clonic phase of the seizure. Evidence of initiating on-demand VNS cannot be seen near the time of seizure detection because the VNS channel is obscured by muscle artifact. However, at the conclusion of the tonic-clonic phase, evidence of VNS generator activity can be seen on the VNS channel as shown in Figure 7-10.

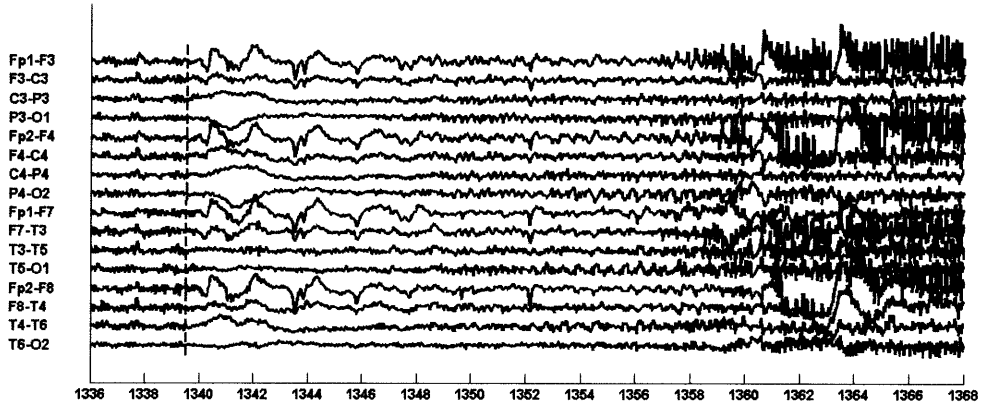

Figure 7-7: Typical EEG change associated with onset of Patient B's seizure. The seizure, which begins at 1340 seconds, involves an electrodecrement most prominent on the occipital channels $P3 - O1$ and $P4 - O2$. Next, at the 1348 seconds, a 3-5 Hz occipital rhythm emerges from the electrodecrement and rapidly generalizes. By 1360 seconds, muscle activity associated with the tonic-clonic phase of the seizure becomes visible.

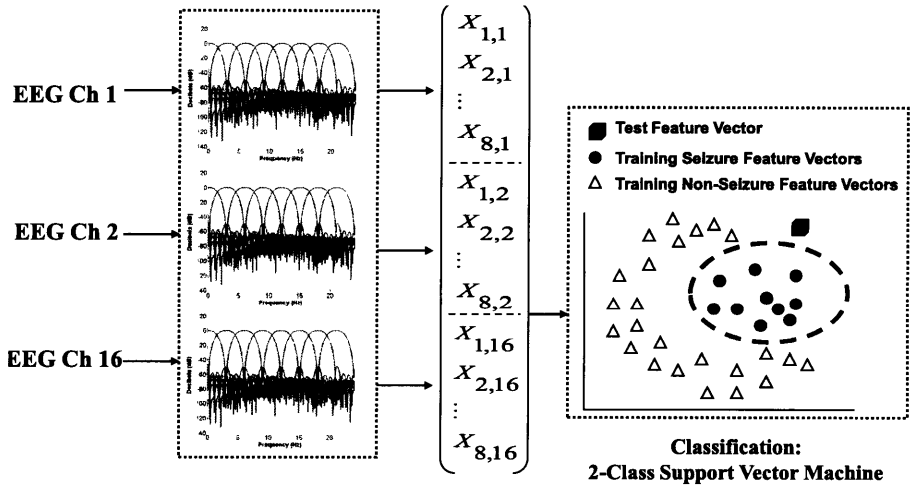

Figure 7-8: Block diagram of patient-specific seizure onset detector that uses features extracted from the EEG signal.

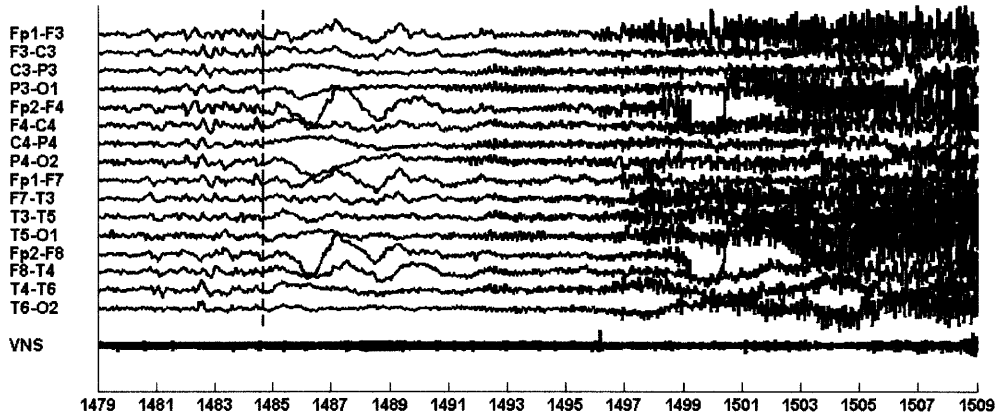

Figure 7-9: Initiation of VNS following computerized detection of the onset of a seizure from Patient B. The seizure begins with an electrodecrement at 1485 seconds. At 1501 seconds, the computerized system detected the muscle activity associated with the tonic-clonic phase of the seizure and initiated VNS in response. Evidence of VNS generator activity can be seen on the VNS channel at the conclusion of the seizure as shown in Figure 7-10.

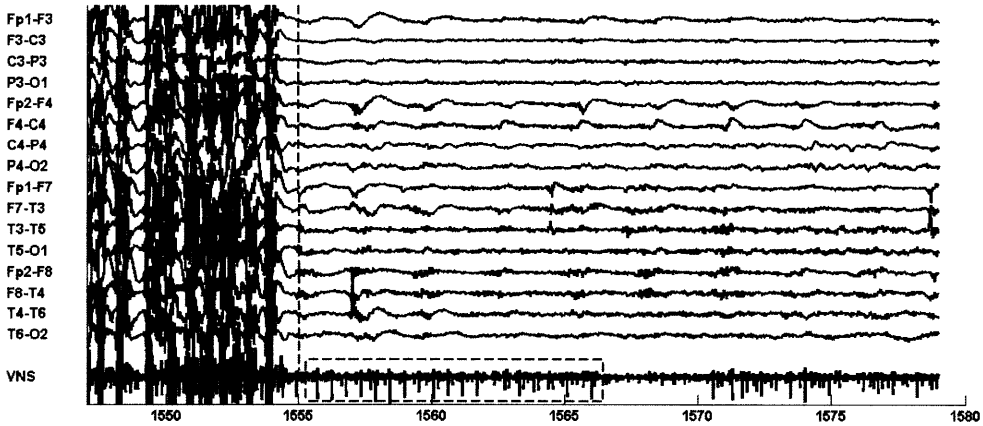

Figure 7-10: The spike-train seen on the VNS channel between 1555-1565 seconds confirms the automatic initiation of vagus nerve stimulation following the onset of the seizure illustrated in Figure 7-9.

Comparing Stimulated and Non-stimulated Seizures

Similar to the first seizure, the tonic-clonic phase of the second seizure lasted for 65 seconds and the post-ictal phase for approximatly 30 minutes. We believe that the length and character of these phases was unaltered because the computerized system initiated on-demand VNS during the generalized (tonic-clonic) portion of the seizure rather than at the focal stage of the seizure. However, behaviorally, relative to the first seizure, patient B was significantly less anxious during the post-ictal phase. Our observation that applying on-demand VNS during the seizure improves post-ictal recovery from the seizure is consistent with the experience of patient B's caregiver.

7.3.3 Patient C

Patient C Medical History

Patient C and Case 2 of Chapter 6 correspond to the same individual. Patient C is a 24 year-old man with Lennox-Gastaut syndrome. At the time of his admission to our study, he was experiencing daily, generalized tonic seizures. The tonic phase of the seizure lasts between 10-20 seconds and the post-ictal phase for approximately an hour. Since Patient C does not experience a warning prior to the onset of a seizure he does not use the on-demand mode of VNS.

Patient C Seizures

Figure 7-11 illustrates an EEG trace of a typical seizure recorded from Patient C. The onset of the seizure, at 248 seconds, involves an electrodecrement that includes all EEG channels and lasts for 14 seconds. During the electrodecrement the patient experiences generalized, tonic muscle contractions. At 262 seconds, both the electrodecrement and tonic contractions end and the EEG exhibits generalized rhythmic activity. The computerized system, using a detector of the form illustrated in Figure 7-8, was trained to recognize the electrodecrement associated with the tonic phase of the seizure.

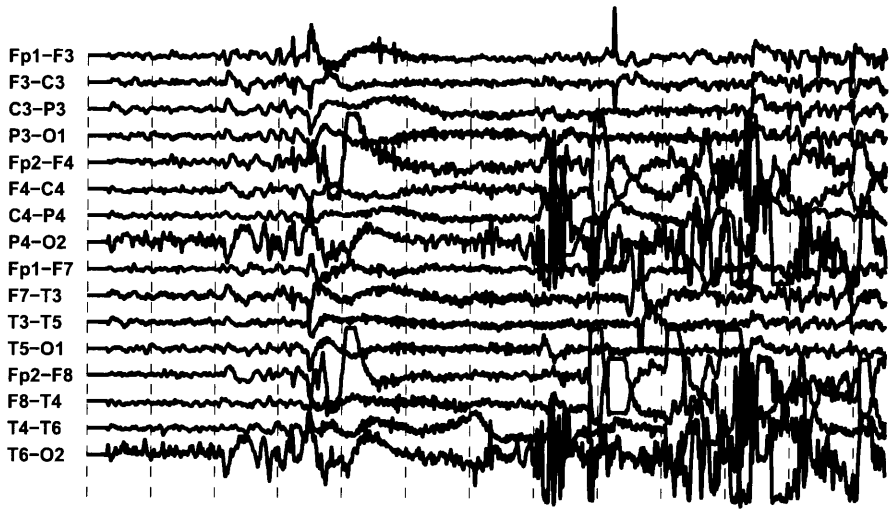

Figure 7-11: Example of a seizure within the scalp EEG of Patient C. The seizure, which begins at 248 seconds, involves a 14 second period of low-amplitude EEG activity across most EEG channels. Later, at 262 seconds, generalized, rhythmic activity develops.

EEG-based Initiation of VNS

Figure 7-12 shows a seizure that triggered the system to initiate on-demand VNS. The system noted seizure activity at 59 seconds (3 second latency) and initiated vagus nerve stimulation in response. Beginning at 69 seconds, one can see a spike-train on the VNS channel indicating successful initiation of vagus nerve stimulation; the spike-train can be seen to continue in Figure 7-13.

Comparing Stimulated and Unstimulated Seizures

We did not detect a change in the electrographic or behavioral character of Patient C's seizure following the initiation of on-demand VNS. Three possible explanations for this include: 1) the on-demand stimulus was not initiated early enough in the course of the seizure 2) the on-demand VNS stimulus parameters (pulse current, frequency, and duration) were not set appropriately 3) Patient C is not a responder to on-demand mode VNS.

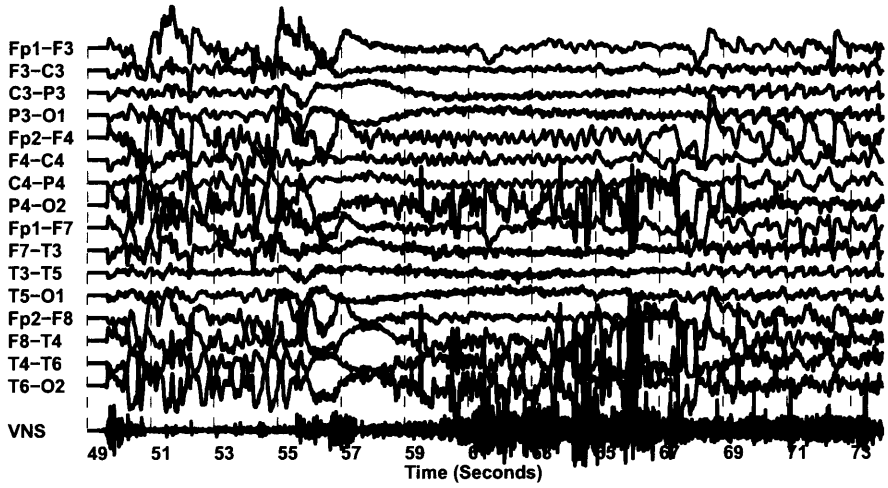

Figure 7-12: Initiation of VNS following computerized detection of the onset of a seizure from Patient C. The seizure begins with an electrodecrement at 56 seconds. The computerized system noted seizure activity at 59 seconds and initiated vagus nerve stimulation in response. Evidence of VNS generator activity can be seen on the VNS channel starting at 69 seconds.

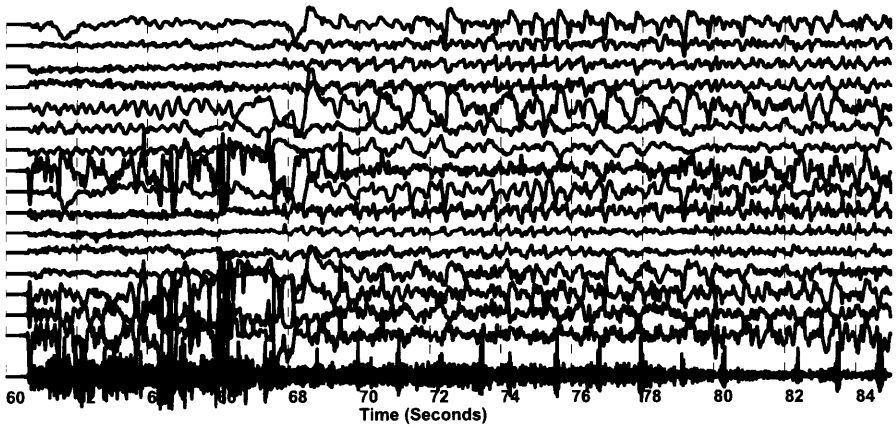

Figure 7-13: The spike-train seen on the VNS channel (first channel from the bottom), between 68-84 seconds, confirms the automatic initiation of vagus nerve stimulation following the onset of the seizure illustrated in Figure 7-12.

Chapter 8

Seizure-Triggered Single Photon Emission Computed Tomography

In this chapter we illustrate the feasibility of applying the seizure detection methodology developed in Chapter 3 to the initiation of a delay-sensitive diagnostic application. More specifically, the infusion of a neuroimaging radiotracer following seizure onset.

8.1 SPECT in Epilepsy

Surgical management of medically intractable seizures requires the identification and resection of an individual's *seizure focus*, the cerebral site that gives rise to a seizure. Defining the seizure focus usually involves electrical characterization of seizures through EEG, anatomical imaging of the brain through magnetic resonance imaging (MRI), and functional imaging of the brain during the seizure and non-seizure states using Single Photon Emission Computed Tomography (SPECT) [5].

A SPECT image reflects differences in blood flow, or perfusion, to various regions within the brain. Since the seizure focus tends to be hypermetabolic during a seizure, demanding a greater share of cerebral blood flow relative to surrounding regions, a SPECT image taken at the time of a seizure (an ictal SPECT) reveals the seizure focus as a region of hyperperfusion as shown in Figure 8-1. In order to obtain an ictal SPECT image that highlights the seizure focus well, a radiotracer must be injected

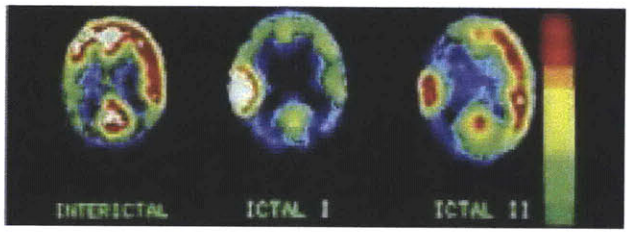

Figure 8-1: SPECT Image taken outside the seizure state (Interictal) and during the seizure state (Ictal I). During a seizure the seizure focus is hyperperfused and appears as a bright spot in a SPECT Image. The third image (Ictal II) is the result of subtracting the first two images.

into a patient soon after the electrical onset of a seizure [63]. To do this one could have: 1) an EEG technician continuously monitor a patient's EEG in real-time in order to rapidly detect and declare the electrical onset of a seizure 2) a trained nurse remain near the patient so that the correct radiotracer dose is safely and promptly administered following declaration of seizure onset by the EEG technician.

Such an approach is costly and nearly impossible to do well when multiple patients need an ictal SPECT within a hospital's Epilepsy Monitoring Unit. In practice, a trained nurse uses a syringe to inject the radiotracer dose after a caregiver near the patient observes the clinical manifestations of a seizure. This results in appreciable injection delays because the seizure's clinical onset tends to lag its electrographic onset; early signs of the seizure's clinical onset are subtle; and the trained nurse is far away from the patient. In our experience injections are started 10-130 seconds after seizure onset as shown in Figure 8-2. Large delays lead to poor localization of the seizure focus due to the visualization of secondarily activated foci in addition to the primary seizure focus.

A system capable of automatically detecting the electrical onset of a seizure and dispensing the appropriate dose of radiotracer using a drug infusion pump could result in shorter and less varied radiotracer injection delays. Overcoming the technical and logistical challenges inherent to the practice of ictal SPECT may increase both the utility and efficacy of ictal SPECT as a tool for seizure focus localization. This will

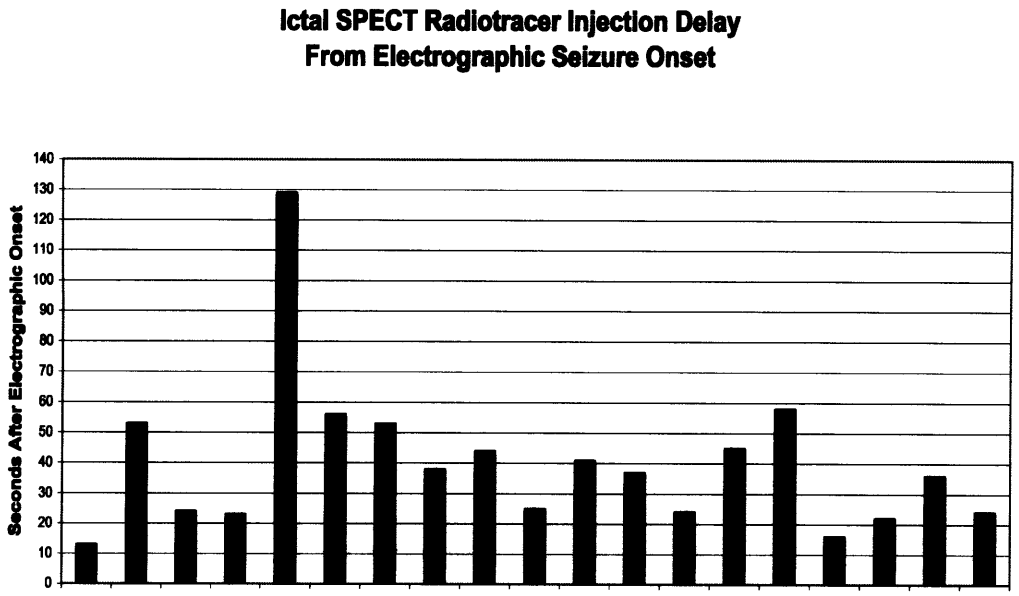

Figure 8-2: Delay associated with injecting the ictal SPECT radiotracer ranges between 10-130 seconds for 18 pediatric patients undergoing ictal SPECTs at Children's Hospital Boston.

particularly benefit epilepsy surgery candidates dependent on functional neuroimaging for the definition of seizure foci that are difficult to visualize using anatomical imaging modalities such as MRI.

In this chapter we describe the design and clinical evaluation of a computerized system that automatically initiates radiotracer infusion following the detection of the electrical onset of a seizure. The computerized system detects seizure onset through the patient-specific methodology developed in Chapter 3. Other automated ictal SPECT systems [15, 53] did not involve computerized detection of the electrical onset of a seizure. Instead, the clinical onset of a seizure prompted a health care provider to remotely initiate radiotracer infusion using a computer-controlled drug pump. When the system described in [15] was evaluated on 26 patients, radiotracer injection delays ranged between 3-48 seconds.

8.2 Methods

8.2.1 System Overview

Figure 8-3 shows a block diagram of the computerized system. The computerized system is composed of a commercial acquisition system (BioLogic Inc.) that collects the EEG of a patient, a computer that analyzes the EEG in real-time using the patient-specific algorithm developed in Chapter 3, and a computer-controlled drug pump that dispenses the ictal SPECT radiotracer.

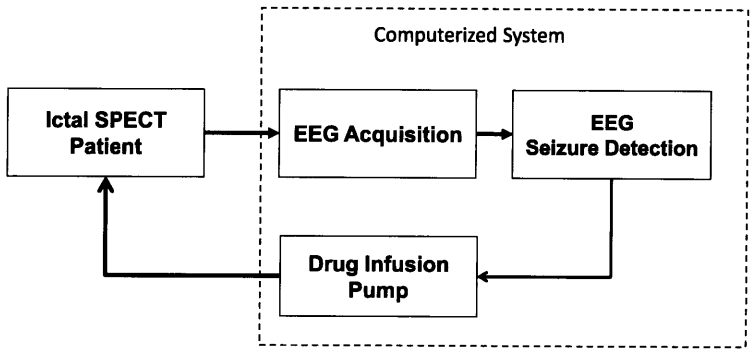

Figure 8-3: Block diagram of system for automatic infusion of ictal SPECT radiotracer.

When the computer detects the onset of a seizure through the analysis of the EEG signal stream, it computes the appropriate dose of radiotracer to infuse. The dose is a function of how long after radiotracer preparation the seizure occurs. Finally, the computer issues a command to the drug pump to dispense the radiotracer dose as a bolus into the patient.

8.2.2 Study Procotol

The clinical evaluation of the computerized system followed a protocol approved by the Institutional Review Board at Children's Hospital Boston (CHB), Boston, Massachusetts, USA. The protocol involved two groups that were simultaneously active,

but which were blinded to each other's activities. The first group (Group 1) included a patient scheduled for an Ictal SPECT as part of his/her hospital admission, and the medical staff responsible for carrying out that ictal SPECT according to CHB protocol. The second group (Group 2) included the computerized system. The two groups are shown in Figure 8-4.

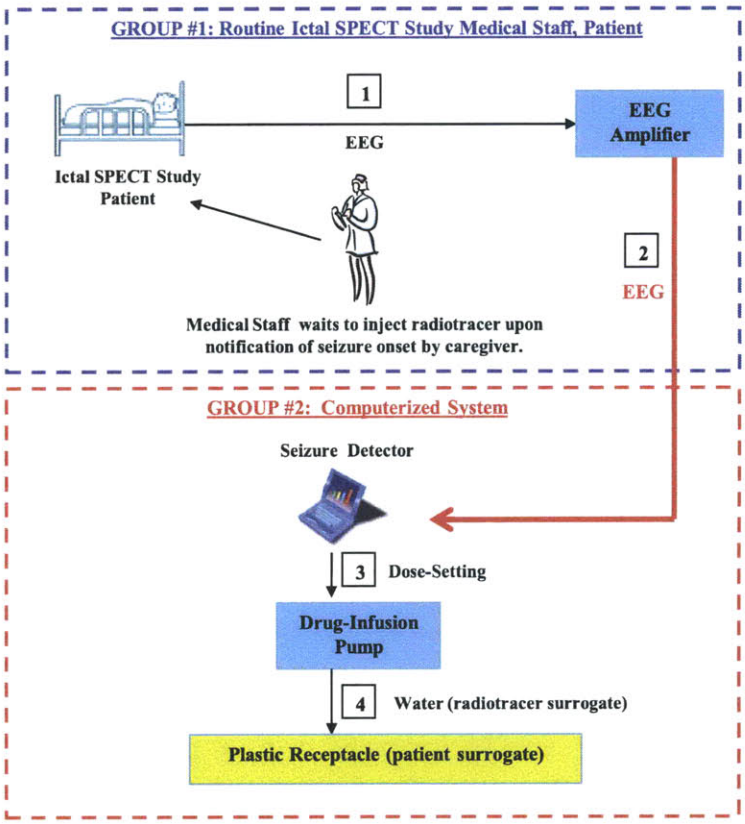

Figure 8-4: Protocol for evaluating automatic system for infusion of ictal SPECT radiotracer. See section 8.2.2 for details.

Group 1 carried out a routine ictal SPECT procedure in accordance with the CHB ictal SPECT protocol. That protocol involved inserting an intravenous line into the patient for radiotracer delivery; recording live scalp EEG from the patient as shown in link 1 of Figure 8-4; awaiting notification of the clinical onset of the patient's seizure by an attending caregiver; confirming seizure activity on the recorded EEG;

and finally mobilizing a trained nurse to the patient's room to deliver the appropriate radiotracer dose as well as a volume of saline to flush the intravenous line connected to the patient.

The scalp EEG captured from the patient in Group 1 was forwarded in real-time to the computerized system in Group 2 as shown by links 1 and 2 of Figure 8-4. The system processed the EEG using the patient-specific detector discussed in Chapter 3. The detector was trained using seizure and non-seizure data recorded between the time of admission and the ictal SPECT study. Upon detecting the electrographic onset of a seizure, the system determined the appropriate dose to be delivered and instructed the drug pump to initiate radiotracer delivery as shown in links 3 and 4 of Figure 8-4. The pump delivered a volume of water (a surrogate for the radiotracer) into a plastic receptacle (a surrogate for the patient). The volume of water equaled the sum of the appropriate radiotracer dose and the volume of saline injected by the nurse to flush the patient's intravenous line.

8.3 Results of Clinical Evaluation

We evaluated the performance of the computerized system during eight ictal SPECT trials. Figure 8-5 shows the time elapsed between electrographic seizure onset and dispensing of the radiotracer dose for both the computerized system (column 2) and the ictal SPECT clinical team (column 3); also shown in Figure 8-5 is the delay with which the computerized system first recognizes the electrographic onset of a seizure (column 1).

We find that our computerized system recognized the electrographic onset of a seizure within 5.1 ± 2.8 seconds and that it completed dispensing the appropriate radiotracer volume within 19.3 ± 2.3 seconds. We attribute the large delay between recognition of seizure onset and pump infusion of the requisite volumes to the slow rate with which the pump selected for the study dispenses fluid. A faster pump would have resulted in a shorter delay. The ictal SPECT clinical team completed infusion of the radiotracer within 27.7 ± 8.5 seconds. For patient number 7, the first seizure went

unnoticed by the clinical team (absent column 3 of Figure 8-5), but was detected by the computerized system. Missing the first seizure of patient 7 prompted the need for a second ictal SPECT study, which both the clinical team and system detected.

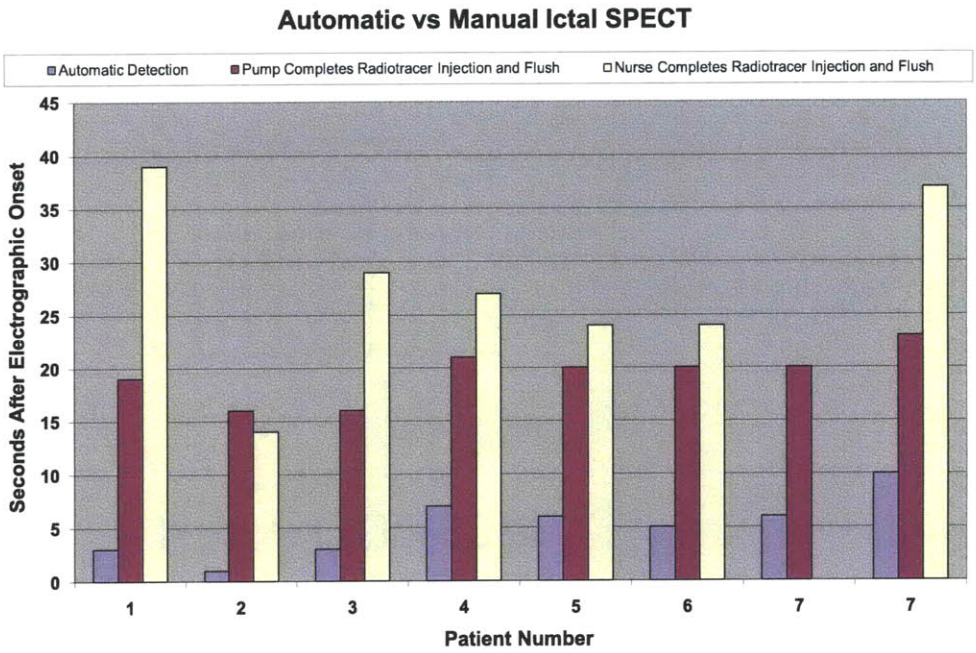

Figure 8-5: Comparison of delay between seizure onset and completion of radiotracer infusion by the clinical team (Group 1) and the computerized system (Group 2).

The computerized system incorrectly dispensed the surrogate radiotracer on 4 occasions over the course of 33 hours. Figure 8-6 shows the latency with which each of the eight seizures were detected (blue bars) as well as the number of false positives (FP) and total number of hours processed per ictal SPECT study (text above blue bars). Two of the eight ictal SPECT trials accounted for all the false detections generated by the computerized system. This result suggests that with careful patient selection the system can be used to initiate radiotracer infusion without concern that injections stem from false detections.

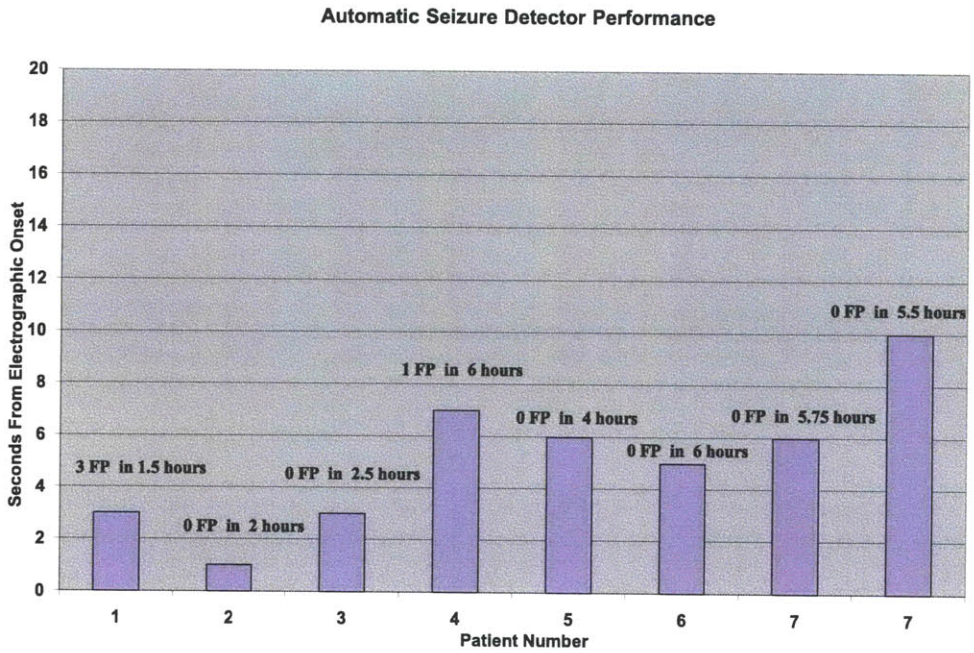

Figure 8-6: Latency and number of false detections declared by the computerized system during eight Ictal SPECT studies.

Chapter 9

Patient-Specific Seizure Onset Detection using iEEG

In this chapter we adapt the patient-specific detector developed in Chapter 3 for the purpose of detecting the onset of a seizure using the intracranial Electroencephalogram (iEEG). Since the preferred embodiment of an intracranial seizure onset detector involves implementation on a low-power, computationally-constrained device, we formulate the algorithm with this constraint in mind. In this chapter we also compare the performance of our algorithm to that of a state-of-the-art, patient non-specific intracranial seizure onset detector.

9.1 Why Detect Seizure Onset Using iEEG?

As discussed in Chapter 2, the intracranial Electroencephalogram provides a spatial and temporal summary of the electrical activity of a population of neurons. However, unlike the scalp EEG, intracranial EEG provides a summary with higher spatial resolution since each iEEG electrode samples the activity of a smaller population of neurons. The higher spatial resolution of iEEG allows one to notice the neuronal hypersynchrony associated with a seizure tens of seconds before the same phenomenon is noticeable within the scalp EEG [39]. At the same time, the higher spatial resolution of iEEG permits the recording of a wider gamut of abnormal, non-seizure activity that

is not visible within the scalp EEG [26, 59, 62].

An implantable medical device capable of detecting and reacting to the onset of a seizure within the iEEG can facilitate applications that are not possible using devices lacking seizure detection capabilities. The device could 1) alert patients of the electrical onset of a seizure before the development of clinical symptoms 2) align a stimulus with the onset of a seizure to suppress the seizure 3) maintain an account of seizure activity so that physicians can objectively determine the efficacy of a therapeutic device.

Performing rapid and reliable seizure onset detection within an implantable medical device is challenging. Algorithms for detecting seizure onset within iEEG signals must be capable of detecting seizures given substantial variability of seizure characteristics across patients, and overlap between the features of seizure and non-seizure activity within the iEEG of a patient. At the same time, such algorithms need to consume a small amount ($\sim 50\mu W$) of the total power budget of the medical device.

9.2 Patient-Specific Detector Architecture

The implantable neurostimulator on which we sought to implement a seizure onset detector is described fully in [2]. This neurostimulator is composed of an analog front-end and a digital back-end. The analog front-end is capable of processing two iEEG channels. For each channel, the analog front-end can extract the spectral power within two bands with configurable center frequency and bandwidth. The digital back-end, which has limited computational capabilities, can be used to implement decision logic that ascertains the presence or absence of seizure activity.

The architecture of our patient-specific detector for the hardware in [2] is illustrated in Figure 9-1. The detector configures the neurostimulator's analog front-end to extract, from each of two iEEG channels, the power within the frequency bands 0-16 Hz and 15-37 Hz. Next, within the neurostimulator's digital back-end, the detector samples the spectral power produced by the front-end and forms a 4-dimensional feature vector that encodes spectral and spatial dependencies between the input chan-

nels. Finally, the detector classifies the observed feature vector using a support-vector machine (SVM) trained to differentiate between a patient's seizure and non-seizure activity using parameters $\gamma = 0.1$, $J = 17$, and a value of C left at the default value of the *SVMLight* software package [28]. Implementation of a support-vector machine with a nonlinear decision boundary required simplification of the SVM discriminant function using the methodology discussed in the following section.

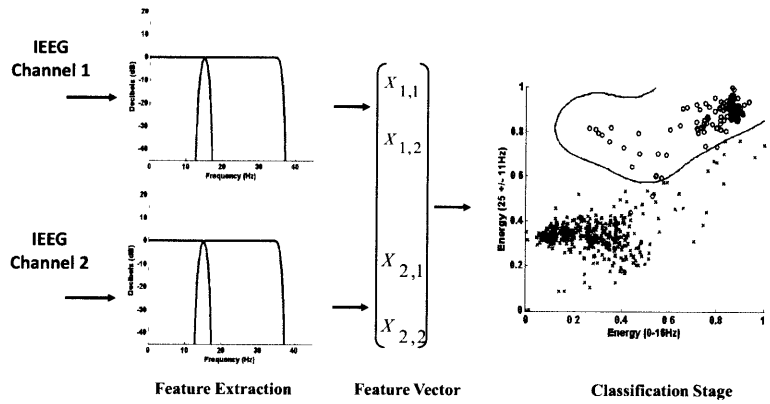

Figure 9-1: Block diagram of a patient-specific seizure detection algorithm that uses spectral and spatial features extracted from two iEEG channels.

9.2.1 Cost of Support-Vector Machine Classification

When the support-vector machine learning algorithm is used to determine a linear decision boundary, the resulting discriminant function (equation 9.1) requires little computation to implement. In the case of the detection architecture shown in Figure 9-1, this discriminant function would involve 4 addition and 4 multiplication operations.

$$\begin{aligned} f(X) &= \text{Seizure} \quad if \;\; W^T X + \beta > 0 \\ f(X) &= \text{Non-seizure} \quad if \;\; W^T X + \beta \leq 0 \end{aligned} \tag{9.1}$$

In the case of a nonlinear boundary, determined using a radial basis kernel, the resulting discriminant function (equation 9.2) is more complex. The discriminant

function contains a number of terms, N_{SV}, equal to the number of support-vectors determined by the support-vector machine learning algorithm. The number of support-vectors is partly governed by the complexity of the classification task. As the similarity between an individual's seizure and non-seizure activity increases, more support-vectors are needed in order to define a more complex decision boundary, and as a result the computational cost of equation 9.2 increases. In our experience, for a detector with the architecture illustrated in Figure 9-1, N_{SV} ranges between 50 and 300. This precludes implementation of the discriminant function in equation 9.2 on the hardware described in [2].

$$f(X) = \text{Seizure} \;\; if \;\; \{\sum_{i=1}^{N_{SV}} \alpha_i \exp(-\gamma \|X - X_i\|^2)\} + \beta > 0 \tag{9.2}$$
$$f(X) = \text{Non-Seizure} \;\; if \;\; \{\sum_{i=1}^{N_{SV}} \alpha_i \exp(-\gamma \|X - X_i\|^2)\} + \beta \leq 0$$

In the following section, we broadly outline the steps involved in a model order reduction technique developed by Schölkopf [52]. This technique allows the nonlinear discriminant function in equation 9.2 to be expressed using $M_{SV} < N_{SV}$ terms.

Reducing The Number of Support-Vectors

Conceptually, the support-vector machine determines a nonlinear boundary in the space of feature vectors by determining a maximum margin, linear boundary in the higher-dimensional space induced by the kernel. Denote a feature vector in the feature space as X_i and its image in the higher-dimensional space as $\phi(X_i)$ where $\phi(\cdot)$ is a map from the feature space to the higher-dimensional space. In the higher-dimensional space the linear boundary is defined by its normal vector W, which can be expressed as a linear combination of the higher-dimensional image of N_{SV} support-vectors as shown in equation 9.3.

$$W = \sum_{i=1}^{N_{SV}} \alpha_i \phi(X_i) \tag{9.3}$$

In order to obtain a nonlinear decision boundary in the feature space that uses $M_{SV} < N_{SV}$ terms, we seek a normal vector $\widetilde{W}$ that is close to W in the sense that $\| \widetilde{W} - W \|^2$ is small. Moreover, $\widetilde{W}$ should be expressed as a combination of the higher-dimensional image of M_{SV} new support-vectors S_i as shown in equation 9.4. Since the L_2-norm of the difference between $\widetilde{W}$ and W is being minimized, the higher dimensional images of feature vectors need not be computed because expansion of the expression being minimized yields inner-products of higher-dimensional vectors that can be evaluated using a kernel.

$$\widetilde{W} = \sum_{i=1}^{M_{SV}} \eta_i \phi(S_i) \tag{9.4}$$

The M_{SV} new support vectors S_i are obtained through a sequential process. The first, scaled vector $\eta_1 S_1$ is chosen as the best single vector approximation of W through solving the following minimization problem

$$\eta_1 S_1 = \min_{\eta, S} \| \eta\phi(S) - W \|^2 \tag{9.5}$$

The next scaled vector $\eta_2 S_2$ is chosen as the best single vector approximation of the residual vector $Y = W - \eta_1 S_1$ through solving the following optimization problem

$$\eta_2 S_2 = \min_{\eta, S} \| \eta\phi(S) - Y \|^2 \tag{9.6}$$

This process is repeated until the set of new support vectors $\eta_i S_i,\ i = 1 \ldots M_{SV}$ is complete. With the new set of support vectors we can reexpress the discriminant

function in equation 9.2 using M_{SV} instead of N_{SV} terms as shown in equation 9.7.

$$f(X) = \text{Seizure } if \ \{\sum_{i=1}^{M_{SV}} \eta_i \exp(-\gamma\|X - S_i\|^2)\} + \beta > 0 \tag{9.7}$$
$$f(X) = \text{Non-Seizure } if \ \{\sum_{i=1}^{M_{SV}} \eta_i \exp(-\gamma\|X - S_i\|^2)\} + \beta \leq 0$$

As an example, Figure 9-2 shows the superposition of a nonlinear boundary requiring $N_{SV} = 50$ terms (solid line) and an approximation of that boundary using $M_{SV} = 8$ terms (dashed line). In our evaluation of the patient-specific detector described in section 9.2 we restricted equation 9.7 to M=8 terms.

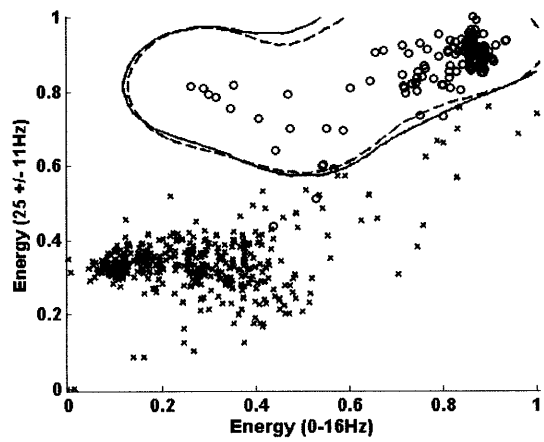

Figure 9-2: Superposition of a nonlinear SVM boundary requiring $N_{SV} = 50$ terms (solid line) and an approximation of that boundary using $M_{SV} = 8$ terms (dashed line).

9.3 Patient Non-specific Detector Architecture

The patient non-specific detector evaluated in this chapter is a simplification of the Osorio-Frei seizure detector [38]. The architecture of this detector is shown in Figure 9-3. The detector monitors, for each channel, the ratio of current and background energy in the 15-37 Hz frequency band; the current energy is computed using a window of 2 seconds while the background energy is computed using a window of 30 minutes.

Whenever the ratio of foreground to background energy on any channel exceeds a threshold R for T seconds a seizure is declared.

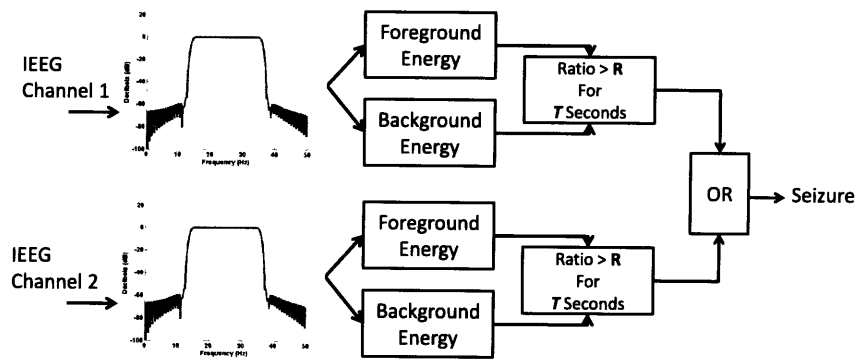

Figure 9-3: Block diagram of patient non-specific seizure detection algorithm.

9.4 iEEG Data Set

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

9.5 Performance Comparison

The performance of the patient-specific detector using both linear and nonlinear SVMs was compared to the performance of the patient non-specific detector set to declare a seizure following $T = 3$ seconds and $T = 10$ seconds.

Figure 9-4 illustrates the average latency with which the patient-specific and non-specific detectors declare the onset of a seizure for each of the 17 subjects. The

patient-specific detector that used a nonlinear SVM (reduced to use M_{NSV}=8 terms) detected 60/61 seizures within 8.5 ± 5.1 seconds. The same detector using a linear SVM detected 60/61 seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm detected 55/61 seizures within 10.5 ± 8.7 seconds. Note that the patient non-specific detector failed to detect the seizures of Patient 3, which accounts for the absence of a bar.

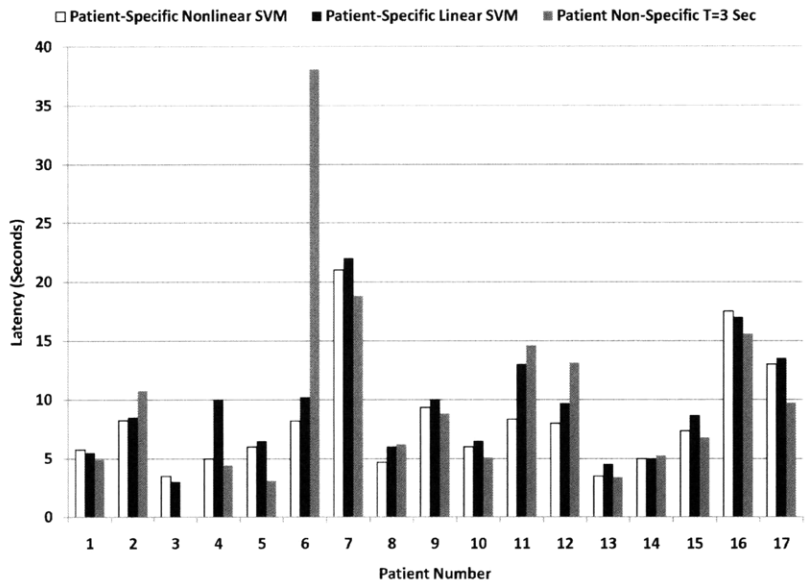

Figure 9-4: Comparison of the detection latencies of the patient-specific and patient non-specific seizure detectors. The patient-specific detector that used a nonlinear SVM detected seizures within 8.5 ± 5.1 seconds. The same detector using a linear SVM detected seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm (with T=3) detected seizures within 10.5 ± 8.7 seconds.

Figure 9-5 illustrates the number of false detections declared by the detectors for each of the 17 test subjects (40 hours of test data). The patient-specific detector employing a nonlinear SVM made 19 false detections; the same detector using a linear SVM made 28 false detections. The patient non-specific detector made 126 false detections.

When the patient non-specific detector was set to declare a seizure following T=10 seconds, it made 17 false detections as illustrated in Figure 9-6 alongside the patient-specific method's false detections. Furthermore, with T=10 seconds the patient non-specific algorithm detected 41/60 seizures within 18.7 ± 10.8 seconds as shown in

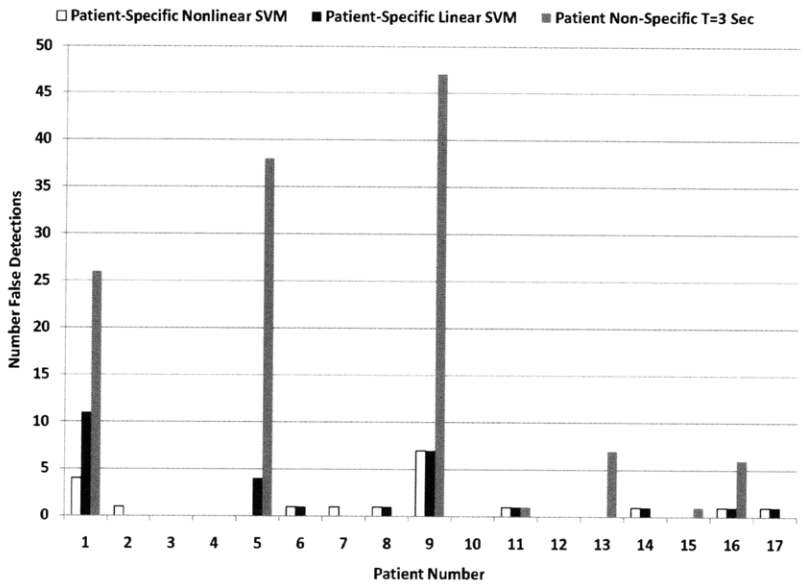

Figure 9-5: Comparison of the number of false detections declared by the patient-specific and patient non-specific seizure detectors during 40 hours of non-seizure data. The patient-specific detector with a nonlinear SVM declared 19 false detections. The same detector using a linear SVM declared 28 false detections. The patient non-specific detector (with T=3) declared 126 false detections.

Figure 9-7 alongside the patient-specific detector latencies.

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

9.6 Reduced and Non-Reduced Support-Vector Machines

In this chapter the number of support-vectors used in equation 9.7 was reduced to $M_{SV} = 8$ using the method developed by Schölkopf [52]. In this section we compare the performance of discriminant functions that use different number of support-vectors in order to show that little performance is lost as a result of reducing the number of support-vectors used to classify an observed feature vector.

Figures 9-8 and 9-9 compare the performance of support-vector machines that

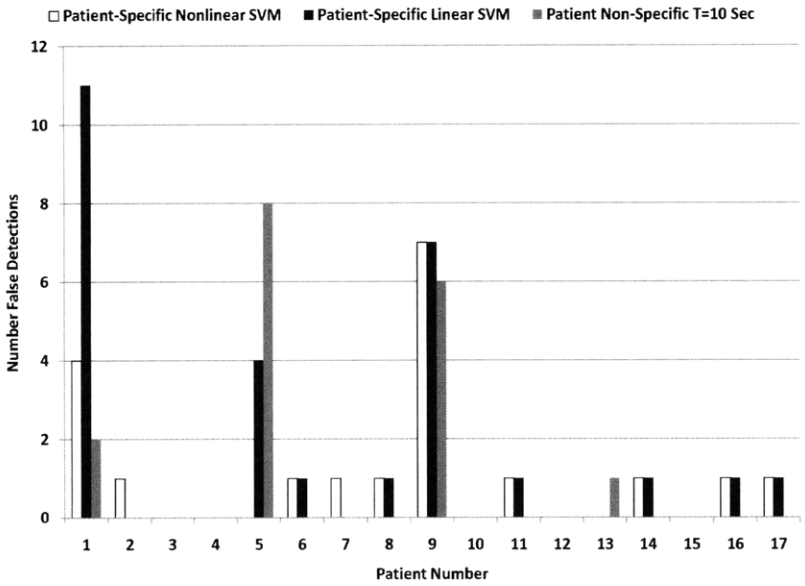

Figure 9-6: Comparison of the number of false detections declared by the patient-specific and patient non-specific seizure detectors during 40 hours of non-seizure data. The patient-specific detector with a nonlinear SVM declared 19 false detections. The same detector using a linear SVM declared 28 false detections. The patient non-specific detector (with T=10) declared 17 false detections.

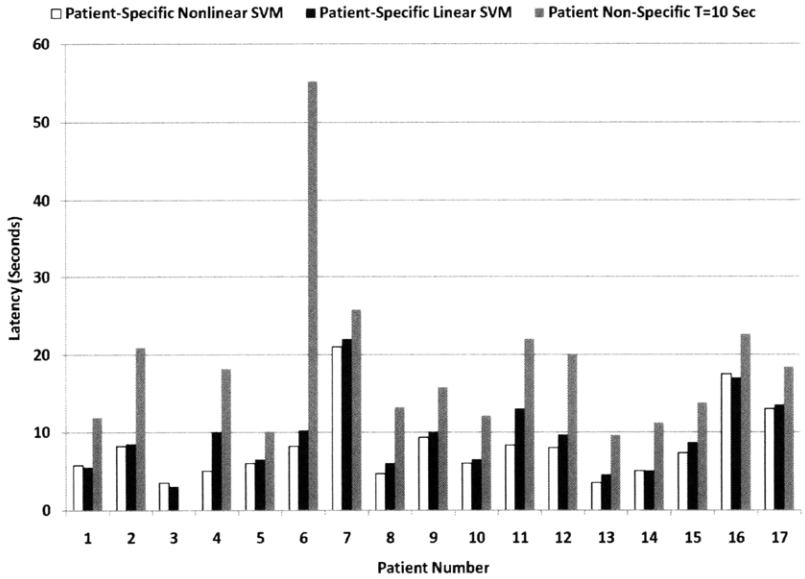

Figure 9-7: Comparison of the detection latencies of the patient-specific and patient non-specific seizure detectors. The patient-specific detector that used a nonlinear SVM detected seizures within 8.5 ± 5.1 seconds. The same detector using a linear SVM detected seizures within 9.3 ± 4.8 seconds. The patient non-specific algorithm (with T=10) detected seizures within 18.7 ± 10.8 seconds.

use $M_{SV} = 8$ and $M_{SV} = 16$ support-vectors with a support-vector machine that is allowed to use all support-vectors. On average, the support-vector machine that used all support-vectors required 321 support-vectors. Figure 9-8 illustrates that significantly reducing the number of support-vectors does not impact a detector's latency, and Figure 9-9 shows that the support-vector machines with fewer support-vectors declared only a few extra false detections.

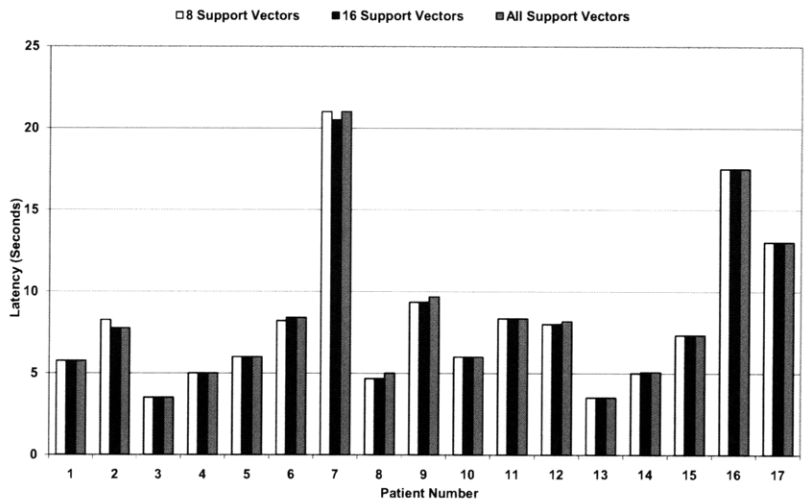

Figure 9-8: Comparison of latency of detectors that use support-vector machines with different number of support-vectors

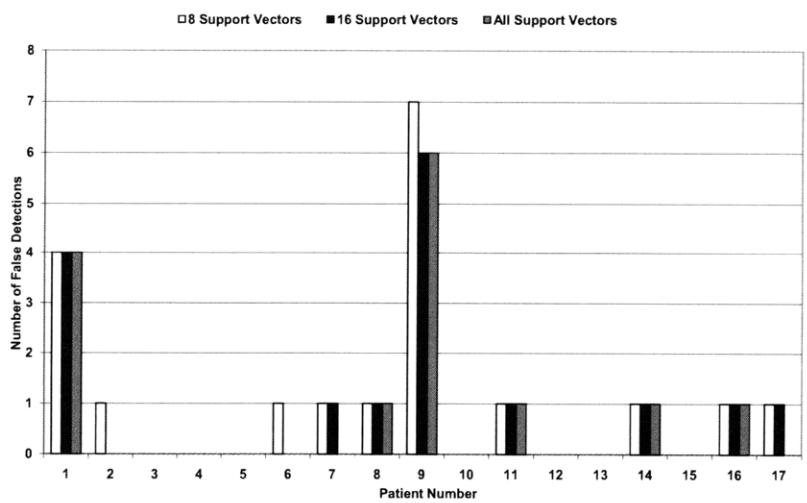

Figure 9-9: Comparison of number of false detections declared by detectors that use support-vector machines with different number of support-vectors

9.7 Case Studies

The following examples illustrate why a patient-specific approach can yield better seizure onset detection performance.

9.7.1 Latency

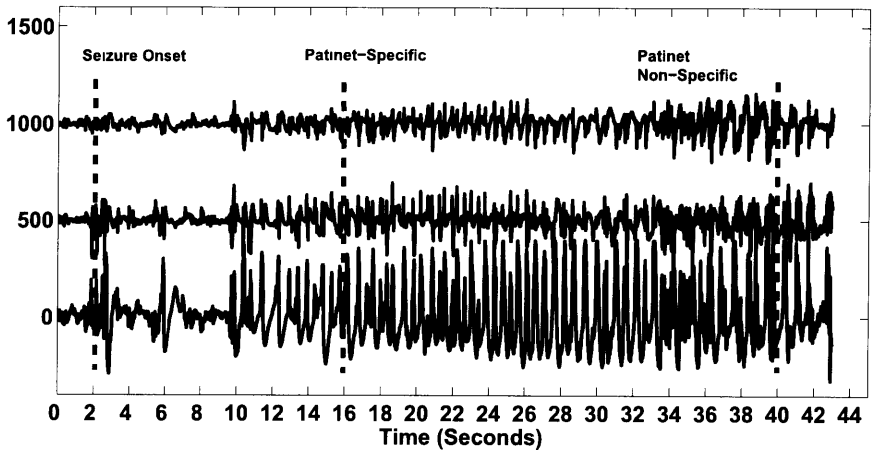

Figure 9-10: Example of a seizure within the iEEG of Patient 6. The seizure, which begins at 2 seconds, consists of a few spikes that evolve into a high-amplitude spike train. The patient-specific detector recognized the seizure at 16 seconds, while the patient non-specific detector did so at 40 seconds.

Figure 9-10 shows an iEEG tracing of a typical seizure from Patient 6. The seizure, which begins at 2 seconds, consist of a few spikes that evolve into a high-amplitude spike train. The patient-specific detector recognized the seizure at 16 seconds, while the patient non-specific detector did so at 40 seconds. Presumably, the patient non-specific detector had a long latency because the dominant frequency component present at seizure onset (1 Hz) is outside the frequency band monitored by the patient non-specific detector (15-37 Hz). The patient-specific detector was able to detect the seizure rapidly because it learned from training seizures the importance of low frequency information.

9.7.2 False Detections

Figure 9-11 shows an iEEG tracing of a seizure from Patient 5. The seizure, which begins at 6 seconds, consist of a high-frequency rhythm that increases in amplitude and decreases in frequency as the seizure progresses. Both the patient-specific and patient non-specific detectors could rapidly detect the onset of this patient's seizure.

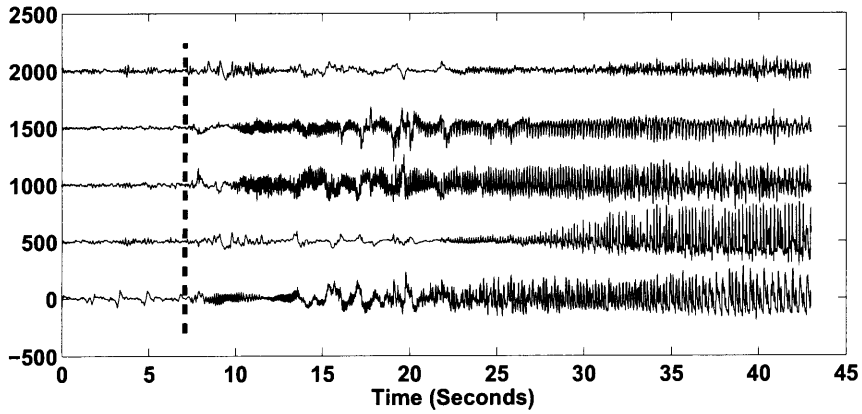

Figure 9-11: Example of a seizure within the iEEG of Patient 5. The seizure, which begins at 6 seconds, consists of a high-frequency rhythm that increases in amplitude and decreases in frequency as the seizure progresses.

During the non-seizure state, the iEEG of Patient 5 contained bursts of rhythmic activity such as that seen in Figure 9-12. The patient-specific detector did not declare these bursts as false detections because it had learned, through training, that they are a feature of the patient's baseline activity. On the other hand, the patient non-specific detector declared many of these bursts as seizure events even though they have spatial and spectral characteristics that differ significantly from those of the actual seizure.

9.8 Predicting Clinical Seizure Onset

Detecting the electrical onset of a seizure using iEEG could be used to alert a patient to a seizure prior to the development of debilitating clinical symptoms. The success of this application depends on the separation between the the electrical onset of a seizure within iEEG and the onset of clinical symptoms. This separation varies across

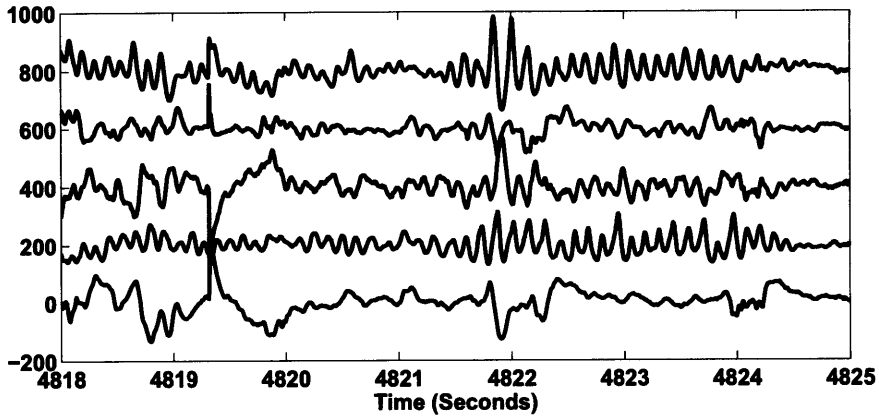

Figure 9-12: Example of a burst of rhythmic non-seizure activity within the iEEG of Patient 5. The patient non-specific detector declared this burst as a seizure event even though its spatial and spectral character differs significantly from that of the seizure shown in Figure 9-11.

patients. Figure 9-13 shows when, relative to the clinical onset, the patient-specific and patient non-specific detectors declare the onset of a seizure for each of the 17 test subjects. Negative latencies correspond to detections that preceded the clinical onset of a seizure and positive latencies correspond to detections that followed the clinical onset. An alarm raised by our patient-specific detector could provide patients 3 and 6 more than thirty seconds to prepare for the clinical onset of a seizure.

9.9 Implementation

Both the patient-specific and non-specific algorithms were implemented on the neurostimulator described in [2]. As discussed in section 9.2, the analog front-end of the stimulator is capable of processing two iEEG channels, and for each channel, extracting the spectral power within two bands with configurable center frequency and bandwidth. The digital back-end samples the analog power profile, assembles these samples into a feature vector, and classifies the feature vector using a support-vector machine. This signal chain is illustrated the Figure 9-14.

By delegating the feature extraction to the analog front-end, the digital back-end remains idle until it is necessary to sample the power profile and test for the presence

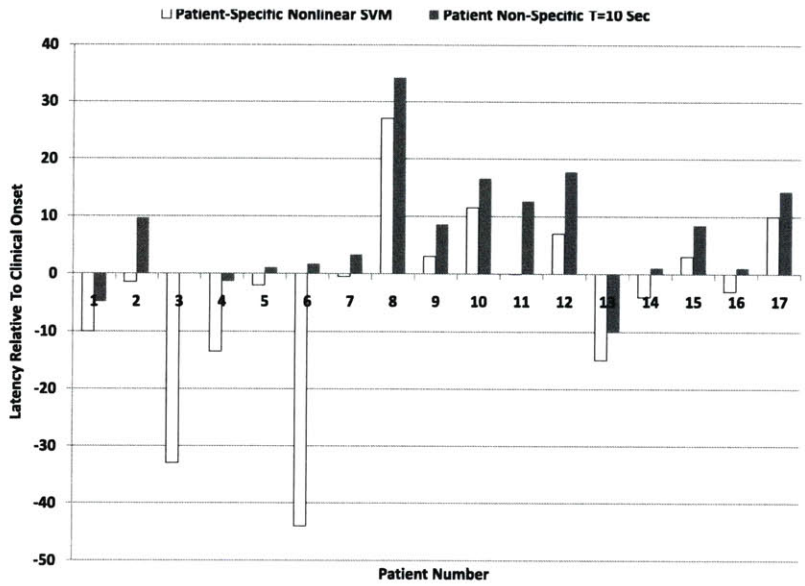

Figure 9-13: Detection delay relative to the clinical onset of a seizure. An alarm based on detecting the electrical onset of a seizure using our patient-specific method could provide some patients, such as patients 3 and 6, enough time to prepare for the clinical onset of a seizure.

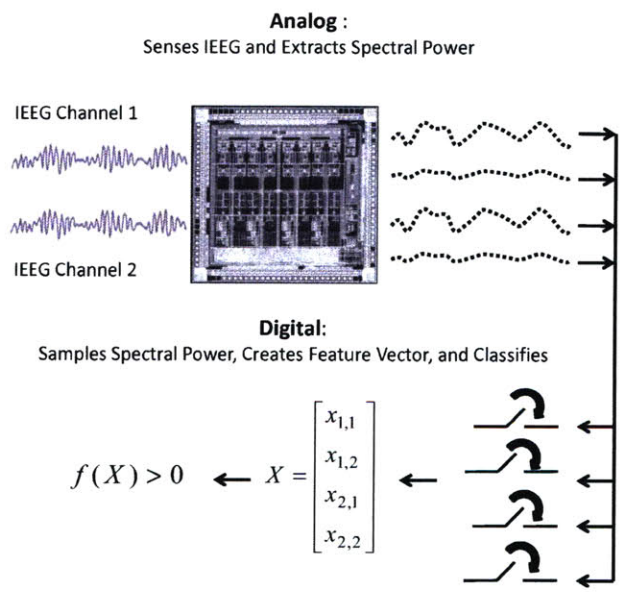

Figure 9-14: Implementation of our machine-learning based, patient-specific detector on the hardware described in [2]. An analog front-end processes two iEEG channels, and for each channel, extracts the spectral power within two configurable frequency bands. A digital back-end samples the analog power profile, assembles these samples into a feature vector, and classifies the feature vector using a support-vector machine.

of seizure activity. This sampling can occur at a rate as low as 1 Hz. Had feature extraction been delegated to the digital back-end, then sampling would have had to occur at a rate suitable for iEEG signals, which is typically 200-250 Hz.

Our measurements indicate that the total current consumed by the patient non-specific detector was $32\mu A$. The patient-specific detector consumed $12\mu A$ when using a linear SVM and $56\mu A$ when using a nonlinear SVM. These measurements suggest that the patient-specific detector employing a linear SVM exhibits the most favorable combination of accuracy, latency, and power consumption.

Chapter 10

Conclusion and Future Work

In this chapter, we conclude the thesis with a summary of its goals and contributions followed by proposed improvements and directions for future work.

10.1 Goals and Contributions

The goal of this thesis was to design, evaluate, and clinically test a seizure onset detection algorithm. The algorithm we developed contributes technically to the field of seizure detection in the following ways:

- **Enhanced Performance:** In Chapter 3 we presented a machine-learning based, patient-specific seizure onset detector. The detector uses a support-vector machine to classify a feature vector that automatically encodes the time evolution of spectral and spatial features within the scalp EEG. When trained on 2 or more seizures and tested on 844 hours of continuous scalp EEG from 23 pediatric subjects, our algorithm detected 96% of 163 test seizures with a median detection delay of 3 seconds (average 4.6 seconds) and a median false detection rate of 0.07 false detections per hour (average 0.13 false detection per hour).

- **Uses Multiple Physiologic Signals:** In Chapter 6 we showed how the detector presented in Chapter 3 can be extended with information from another

physiologic source whenever the scalp EEG alone is an unreliable indicator of seizure onset. This capability is important for the detection of seizures whose onsets lack the development of rhythmic activity and instead reflect physical sequelae of the seizure.

- **Minimal User Intervention:** Our algorithm does not require a user to define the values of parameters such as which EEG channels to monitor for seizure activity. In our approach, a user only needs to define the onset of activity associated with a seizure in a set of physiologic signals. The relationships between these signals that distinguish the seizure and non-seizure periods are automatically learned.

- **Suitable for Implantable Medical Devices:** In Chapter 9 we showed that the algorithm presented in Chapter 3 can be adapted for the detection of seizure onsets within intracranial EEG, and can be implemented on the low-power hardware of an implantable neurostimulator. When evaluated on 81 hours of intracranial EEG containing 61 seizures and gathered from 17 adult subjects, our algorithm detected 60/61 seizures within 8.5 seconds and declared a total of 19 false detections.

  In this thesis we also used our algorithm to enable delay-sensitive therapeutic and diagnostic applications:

- **Non-invasive Closed-Loop Control of the Vagus Nerve Stimulator:** In Chapter 7 we showed how our algorithm can be used to initiate vagus nerve stimulation in response to detecting the onset of a seizure using multiple physiologic signals. As an example of the system's capabilities, during an 81 hour clinical test of the system on a patient, the computerized system detected 5/5 seizures and initiated VNS within 5 seconds of the appearance of ictal discharges in the EEG.

- **Computerized Initiation of Ictal SPECT Studies:** In Chapter 8 we showed how our algorithm can be used to initiate a functional neuroimaging

study following seizure onset. The neuroimaging modality, ictal SPECT, is used to radiographically localize the cerebral origin of a seizure. Our system could initiate injection of the radiotracer used for ictal SPECT within 19.3 ± 2.3 seconds in 8/8 prospective trials, while the clinical team required 27.7 ± 8.5 seconds, and failed to initiate ictal SPECT in one of the trials.

10.2 Future Work

10.2.1 Feature Vector Enhancement

The feature vector developed in this thesis captures relationships between channels by concatenating univariate spectral features derived from each EEG channel. Recently, bivariate features that directly measure relationships between EEG channels, such as phase synchrony, have been shown to be effective in seizure prediction [36]. Moreover, graph-theoretic features of a network representation of EEG, such as clustering coefficient C (measure of local connectedness) and shortest path length L (measure of overall network integration), have been shown to distinguish between the seizure and non-seizure state [42]. Future work will investigate whether the addition of these features to our feature vector improves the latency, sensitivity, or specificity of the detector presented in Chapter 3.

10.2.2 Detecting Seizure Cessation

The detector presented in Chapter 3 signals the onset of a seizure, but is not designed to signal the end of a seizure. Detecting the cessation of a seizure allows one to compute seizure duration, and the length of the time interval between the end of one seizure and the beginning of the next seizure. These quantities have clinical significance. Seizures that persist for more than 10 minutes, or a cluster of seizures that are closely spaced in time signal a possible transition into status epilepticus [13]. Status epilepticus refers to a life-threatening condition in which the brain enters a state of persistent seizure activity. The detector in Chapter 3 learned the transition

between non-seizure and seizure activity using a feature vector that encodes the time evolution of spectral and spatial features within the scalp EEG. Future work will investigate whether the transition from seizure back to non-seizure activity can also be learned using a similar feature vector.

10.2.3 Closed-Loop, Non-invasive Brain Stimulation

In this thesis we illustrated the feasibility of initiating vagus nerve stimulation following the detection of seizure onset using non-invasive physiologic signals. Future work will investigate the feasibility and therapeutic benefit of using our detection algorithms to initiate other neurostimulation modalities such as as repetitive transcranial magnetic stimulation (rTMS), transcranial direct current stimulation (tDCS), and trigeminal nerve stimulation (TNS) [12]. This research will complement the development of detect-and-treat systems that deliver cortical or deep-brain stimulation based on the analysis of invasive physiologic signals (iEEG) [31, 41, 58].

Bibliography

[1] Rehan Akbani, Stephen Kwek, and Nathalie Japkowicz. *Applying Support-Vector Machines to Imbalanced Datasets.* Springer, 2004.

[2] Al-Thaddeus Avestruz, Wesley Santa, Dave Carlson, Randy Jensen, Scott Stanslaski, Alan Helfenstine, and Tim Denison. A 5μw/channel spectral analysis ic for chronic bidirectional brain-machine interfaces. *IEEE Journal Of Solid-State Circuits*, 43(12):3006–3024, 2008.

[3] E. Ben-Menachem. Vagus nerve stimulation for the treatment of epilepsy. *Lancet Neurology*, 1(8):477–482, 2002.

[4] P. Boon, K. Vonck, and P. Van Walleghem et al. Programmed and magnet-induced vagus nerve stimulation for refractory epilepsy. *Journal Of Clinical Neurophysiology*, 18(5):402–407, 2001.

[5] Martin J Brodie, Steven C Schachter, and Patrick Kwan. *Fast Facts: Epilepsy.* Health Press, third edition, 2005.

[6] Begley CE, Famulari M, Annegers JF, Lairson DR, Reynolds TF, and Coan S et al. The cost of epilepsy in the United States: An estimate from population-based and clinical survey. *Epilepsia*, 41(3):342–351, 2000.

[7] Alexander M. Chan, Felice T. Sun, Erem H. Boto, and Brett M. Wingeier. Automated seizure onset detection for accurate onset time determination in intracranial EEG. *Clinical Neurophysiology*, 119:2687–2696, 2008.

[8] Bernard S. Chang and Daniel H. Lowenstein. Epilepsy. *The New England Journal of Medicine*, 349(13):1257–1266.

[9] Nello Cristianini and John Shawe-Taylor. *An Introduction To Support Vector Machines And Other Kernel Based Learning Methods.* Cambridge University Press, 2000.

[10] Barbara Cysyk, Jehuda Sepkuty, Ronald Lesser, and Ali Civelek. Truly ictal spect is of major importance for reliable localization of seizure focus. *Epilepsia*, 38:Supplement 8:146, 1997.

[11] C.M. DeGiorgio, S.C. Shcahcter, A. Handforth, M. Salinsky, J. Thompson, and B. Uthman et al. Prospective long-term study of vagus nerve stimulation for the treatment of refractory seizures. *Epilepsia*, 41(9):1195–1200, 2000.

[12] C.M. DeGiorgio, A. Shewmon, D. Murray, and T. Whitehurst. Pilot study of trigeminal nerve stimulation (TNS) for epilepsy: a proof-of-concept trial. *Epilepsia*, 47(7):1213–1215, 2006.

[13] Frank W. Drislane. *Status Epilepticus.* Humana Press, 2005.

[14] Javier Echauz, Stephen Wong, Otis Smart, Andrew Gardner, Gregory Worrell, and Brian Litt. Computation applied to clinical epilepsy and antiepileptic devices. In Ivan Soltesz and Kevin Staley, editors, *Computational Neuroscience in Epilepsy*, part 32, pages 530–558. Academic Press, California, first edition, 2008.

[15] Michael Feichtinger, Hans Eder, Alexander Holl, Eva Korner, and Gerda Zmugg et al. Automatic and remote controlled ictal SPECT injection for seizure focus localization by use of a commercial contrast agent application pump. *Epilepsia*, 48(7):1409–1413.

[16] David M. Ficker. Sudden unexplained death and injury in epilepsy. *Epilepsia*, 41(Suppl. 2):S7–S12.

[17] David E. Friedman and Frank G. Gilliam. Seizure-related injuries are underrported in pharmacoresistant localization-related epilepsy. *Epilepsia*, 2009.

[18] J. Gotman. Automatic recognition of epileptic seizure in the EEG. *Electroencephalography and Clinical Neurophysiology*, 54:530–540, 1982.

[19] J. Gotman, J.R. Ives, and P. Gloor. Frequency content of EEG and EMG at seizure onset: Possibility of removal of EMG artifact by digital filtering. *Electroencephalography and Clinical Neurophysiology*, 52:626–639, 1981.

[20] Sukhi Grewal and Jean Gotman. An automatic warning system for epileptic seizures recorded on intracerebral EEGs. *Clinical Neurophysiology*, 116:2460–2472, 2005.

[21] The Vagus Nerve Stimulation Study Group. A randomized controlled trial of chronic vagus nerve stimulation for treatment of medically intractable seizures. *Neurology*, 45(2):224–230, 1995.

[22] E.J. Hammond, B.M. Uthman, S.A. Reid, and B.J. Wilder. Electrophysiological studies of cervical vagus nerve stimulation in humans: EEG effects. *Epilepsia*, 33(6):1013–1020, 1992.

[23] A. Handforth, C.M. DeGiorgio, S.C. Schachter, B.M. Uthman, D.K. Naritoku, and E.S. Tecoma et al. Vagus nerve stimulation therapy for partial onset seizures: a randomized active control trial. *Neurology*, 51(1):48–55, 1998.

[24] C. Hopee, A. Poepel, and C.E. Elger. Epilepsy: Accuracy of patient seizure counts. *Arc Neurol*, 64(11):1595–1599, 2007.

[25] G.L. Morris III. A retrospective analysis of the effects of magnet-activated stimulation in conjuction with vagus nerve stimulation therapy. *Epilepsy and Behavior*, 4(6):740–745, 2003.

[26] J.D. Jirsch, E. Urrestarazu, P. LeVan, A. Olivier, F. Dubeau, and J. Gotman. High-frequency oscillation during human focal seizures. *Brain*, 129:1593–1608, 2006.

[27] T. Joachims. *Text Categorization with Support-Vector Machines: Learning with Many Relevant Features.* Springer, 1998.

[28] T. Joachims. *Making large-Scale SVM Learning Practical. Advances in Kernel Methods - Support Vector Learning.* MIT-Press, 1999.

[29] Christophe Jouny and Gregory Bergey. Dynamics of epileptic seizures during evolution and propagation. In Ivan Soltesz and Kevin Staley, editors, *Computational Neuroscience in Epilepsy*, part 28, pages 457–469. Academic Press, California, first edition, 2008.

[30] Eric Kandel, J.H. Schwartz, and Thomas M. Jessell. *Principles Of Neural Science.* McGraw-Hill Medical, 2001.

[31] E.H. Kossoff, E.K. Ritzl, J.M. Politsky, A.M. Murro, J.R. Smith, and R.B. Duckrow et al. Effect of an external responsive neurostimulator on seizures and electrographic discharges during subdural electrode monitoring. *Epilepsia*, 45(12):1560–1567, 2004.

[32] D. Labar, J. Murphy, and E. Tecoma. Vagus nerve stimulation for medication-resistant generalized epilepsy. *Neurology*, 52(7):1510–1512, 1999.

[33] Fritz Leutmezer, Christiana Schernthaner, Stefanie Lurger, Klaus Pötzelberger, and Christoph Baumgartner. Electrocardiographic changes at the onset of epileptic seizures. *Epilepsia*, 44(3):348–354, 2003.

[34] P. LeVan, E. Urrestarzu, and J. Gotman. A system for automatic artifact removal in ictal scalp EEG based on independent component analysis and bayesian classification. *Clinical Neurophysiology*, 117:912–927, 2006.

[35] Ralph Meier, Heike Dittrich, Andreas Schulze-Bonhage, and Ad Aertsen. Detecting epileptic seizures in long-term human EEG: A new approach to automatic online and reat-time detection and classification of polymorphic seizure patterns. *Journal of Clinical Neurophysiology*, 25:119–131, 2008.

[36] Florian Mormann, Thomas Kreuz, Christoph Rieke, Ralph Andrzejak, Alexander Kraskov, Peter David, Christian Elger, and Klaus Lehnertz. On the predictability of epileptic seizures. *Clinical Neurophysiology*, 116:569–587, 2005.

[37] Anthony M. Murro, Don W. King, Joseph R. Smith, Brian B. Gallagher, Herman F. Flanigin, and Kimford Meador. Computerized seizure detection of complex partial seizures. *Clinical Neurophysiology*, 79(4):330–333, 1991.

[38] Ivan Osorio, Mark Frei, and Steven Wilkinson. Real-time automated detection and quantitative analysis of seizures and short term prediction of clinical onset. *Epilepsia*, 39(6):615–627, 1998.

[39] Steven Pacia and John Ebersole. Intracranial EEG substrates of scalp ictal patterns from temporal lobe foci. *Epilepsia*, 38(6):642–654, 1997.

[40] Shyamal Patel, Chiara Mancinelli, Anthony Dalton, Ben Patritti, Trudy Pang, Steven Schachter, and Paolo Bonato. Detecting epileptic seizures using wearable sensors. *35th Annual Northeast Bioengineering Conference*, pages 1–2, 2009.

[41] T.E. Peters, N.C. Bhavaraju, M.G. Frei, and I. Osorio. Network system for automated seizure detection and contingent delivery of therapy. *Journal of Clinical Neurophysiology*, 18(6):545–549, 2001.

[42] S.C. Ponten, F. Bartolomei, and C.J. Stam. Small-world networks and epilepsy: Graph theoretical analysis of intracerebrally recorded mesial temporal lobe seizures. *Clinical Neurophysiology*, 118:918–927, 2007.

[43] Hao Qu. *Self-adapting Algorithms for Seizure Detection during EEG Monitering.* PhD dissertation, McGill University, 1995.

[44] Hao Qu and Jean Gotman. Improvement in seizure detection performance by automatic adaptation to the EEG of each patient. *Electroencephalography and Clinical Neurophysiology*, 86:79–87, 1993.

[45] Hao Qu and Jean Gotman. A seizure warning system for long-term epilepsy monitoring. *Neurology*, 45:2250–2254, 1995.

[46] Hao Qu and Jean Gotman. A patient-specific algorithm for the detection of seizure onset in long-term EEG monitoring: Possible use as a warning device. *IEEE Transactions On Biomedical Engineering*, 44:115–122, 1997.

[47] Steven Rothman and Xiao-Feng Yang. Local cooling: A therapy for intractable neocortical epilepsy. *Epilepsy Currents*, 3(5):153–156, 2003.

[48] M.E. Saab and J. Gotman. A system to detect the onset of epileptic seizures in scalp EEG. *Clinical Neurophysiology*, 116:427–442, 2005.

[49] Steven Schachter. *Brainstorms: Epilepsy in Our Words. Personal Accounts of Living with Seizures.* Raven Press, 1993.

[50] Steven Schachter and Clifford Saper. Vagus nerve stimulation. *Epilepsia*, 39:677–686, 1998.

[51] Steven C. Schachter. *Vagus Nerve Stimulation.* Informa HealthCare, 2003.

[52] B. Schlkopf, P. Knirsch, C. Smola, and A. Burges. Fast approximation of support vector kernel expansions and an interpretation of clustering as approximation in feature spaces. In P.Levi, M. Schanz, R.J. Ahler, and F.May, editors, *Mustererkennung 1998-20 DAGM Symposium*, pages 124–132. Springer-Verlag, Berline, Germany, first edition, 1998.

[53] Jehuda P. Sepkuty, Ronal P. Lesser, Cahid, Barbara Cysyk, Robert Webber, and Roy Shipley. An automated injection system (with patient selection) for SPECT imaging in seizure localization. *Epilepsia*, 39(12):1350–1356.

[54] Eugene I. Shih, Ali H. Shoeb, and John V. Guttag. Sensor Selection for Energy-Efficient Ambulatory Medical Monitoring. In *Mobisys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 347–358, New York, NY, USA, 2009. ACM.

[55] Ali Shoeb. *Patient-Specific Seizure Onset Detection.* MEng dissertation, Massachusetts Institute of Technology, 2003.

[56] Ali Shoeb, Herman Edwards, Jack Connolly, Blaise Bourgeois, S. Ted Treves, and John Guttag. Patient-specific seizure onset detection. *Epilepsy and Behavior*, 5:483–498, 2004.

[57] Ali Shoeb, Trudy Pang, John Guttag, and Steven Schachter. Non-invasive computerized system for automatically initiating vagus nerve stimulation following patient-specific detection of seizures or epileptiform discharges. *International Journal of Neural Systems*, 19(3):157–172, 2009.

[58] Felice Sun, Martha Morrell, and Robert Wharen. Responsive cortical stimulation for the treatment of epilepsy. *Neurotherapeutics*, 5(1):68–74, 2008.

[59] James Tao, Armit Ray, Susan Ebersle, and John Ebersole. Intracranial EEG substrates of scalp eeg interictal spikes. *Epilepsia*, 46(5):669–676, 2005.

[60] William Theodore and Robert Fisher. Brain stimulation for epilepsy. *Lancet Neurology*, 3(2):111–118, 2004.

[61] Sergios Theodoridis and Konstantinos Koutroumbas. *Pattern Recognition*, chapter 3. Elsevier, fourth edition, 2009.

[62] Elena Urrestarazu, Rahul Chander, Francois Dubeau, and Jean Gotman. Interictal high-frequency oscillations (100-500hz) in the intracerebral EEG of epileptic patients. *Brain*, 130:2354–2366, 2007.

[63] G.I. Varghese, M.J. Purcaro, and J.E. Motelaw et al. Clinical use of ictal SPECT in secondarily generalized tonic-clonic seizures. *Brain*, 132:2102–2113.

[64] Scott Wilson, Mark Scheuer, Ronald Emerson, and Andrew Gabor. Seizure detection: Evaluation of the Reveal algorithm. *Clinical Neurophysiology*, 10:2280–2291, 2004.

[65] Scott B. Wilson. A neural network method for automatic and incremental learning applied to patient-dependant seizure detection. *Clinical Neurophysiology*, 116:1785–1795, 2005.

[66] Scott B. Wilson. Alogirthm architectures for patient dependant seizure detection. *Clinical Neurophysiology*, 117:1204–1216, 2006.

[67] Dixon Woodbury and J. Walter Woodbury. Effects of vagal stimulation on experimentally induced seizures in rats. *Epilepsia*, 31(Suppl. 2):S7–S19, 1990.

[68] Jacob Zabara. Inhibition of experimental seizures in canines by repetitive vagal stimulation. *Epilepsia*, 33(6):1005–1012, 1992.

[69] Feng Zhang, Li-Ping Wang, and Martin Brauner et al. Multimodal fast optical interrogation of neural circuitry. *Nature*, 446:633–639, 2007.