HHS Public Access

Author manuscript

J Neural Eng. Author manuscript; available in PMC 2017 June 01.

Published in final edited form as:

J Neural Eng. 2016 June ; 13(3): 036011. doi:10.1088/1741-2560/13/3/036011.

A novel seizure detection algorithm informed by hidden Markov model event states

Steven Baldassano[1,2], **Drausin Wulsin**[1], **Hoameng Ung**[1,2], **Tyler Blevins**[1,2], **Mesha-Gay Brown**[3], **Emily Fox**[4], and **Brian Litt**[1,2,5]

[1]Department of Bioengineering, University of Pennsylvania

[2]Center for Neuroengineering and Therapeutics, University of Pennsylvania

[3]Department of Neurology, University of Colorado Denver

[4]Department of Statistics, University of Washington

[5]Department of Neurology, University of Pennsylvania

Abstract

Objective—Recently the FDA approved the first responsive, closed-loop intracranial device to treat epilepsy. Because these devices must respond within seconds of seizure onset and not miss events, they are tuned to have high sensitivity, leading to frequent false positive stimulations and decreased battery life. In this work, we propose a more robust seizure detection model.

Approach—We use a Bayesian nonparametric Markov switching process to parse intracranial EEG (iEEG) data into distinct dynamic event states. Each event state is then modeled as a multidimensional Gaussian distribution to allow for predictive state assignment. By detecting event states highly specific for seizure onset zones, the method can identify precise regions of iEEG data associated with the transition to seizure activity, reducing false positive detections associated with interictal bursts. The seizure detection algorithm was translated to a real-time application and validated in a small pilot study using 391 days of continuous iEEG data from 2 dogs with naturally occurring, multifocal epilepsy. A feature-based seizure detector modeled after the NeuroPace RNS System was developed as a control.

Main results—Our novel seizure detection method demonstrated an improvement in false negative rate (0/55 seizures missed vs 2/55 seizures missed) as well as a significantly reduced false positive rate (0.0012/hour vs 0.058/hr). All seizures were detected an average of 12.1 ± 6.9 seconds before the onset of unequivocal epileptic activity (UEO).

Significance—This algorithm represents a computationally inexpensive, individualized, real-time detection method suitable for implantable antiepileptic devices that may considerably reduce false positive rate relative to current industry standards.

Corresponding Author Contact Information: Steven Baldassano, MD/PhD Candidate, Department of Bioengineering, University of Pennsylvania, 210 South 33rd Street, Suite 240 Skirkanich Hall, Philadelphia, PA 19104-6321, ; Email: stevennb@mail.med.upenn.edu

Conflicts of Interest: Dr. Brian Litt has licensed technology to NeuroPace, Inc. through the University of Pennsylvania

1. Introduction

Epilepsy affects over 60 million individuals worldwide, with one quarter of patients having disease refractory to standard therapies including medication and surgery[1]. Automated seizure detection algorithms have been studied for decades to improve the diagnosis and treatment of epilepsy[2]. More recently, these algorithms have been applied to closed-loop implantable devices designed to detect and electrically stimulate to abort epileptic activity[3]. These devices help satisfy a need for additional strategies for seizure control, with the potential to improve the quality of life of millions of patients suffering from epileptic seizures.

The first closed-loop implantable therapeutic device to be approved by the FDA for treatment of epilepsy is the NeuroPace Responsive Neurostimulation System, in 2013. This device, like other similar systems in development, uses real-time intracranial EEG (iEEG) data as input to an algorithm to detect onset of epileptic activity and trigger targeted electrical stimulation to arrest potential seizures. While such devices show promising clinical results, they are often limited by the efficacy of the detection algorithms. The algorithms used in these devices are typically dependent on extracting and analyzing specific "features" of the EEG signal, such as amplitude, line length[4], and area under the curve. In order to successfully and reliably avert epileptic activity, the detection algorithm must detect seizure onset with sufficient latency prior to clinical symptoms to provide an opportunity for intervention. A clinically useful system will detect all seizures and is thus required to be highly sensitive; however, these devices have been hampered by high false positive rates, causing unnecessary stimulation and increased frequency of repeat surgery to replace spent batteries.

One potential source of false positive detections by feature-based methods is the occurrence of sub-clinical epileptiform "bursts," also known in the literature by terms such as brief ictal rhythmic discharges (B(I)RDs), and others[5]. These events represent an abnormal EEG finding without obvious clinical manifestations, and often occur with greater frequency than seizures. Although the underlying pathology of these discharges remains uncertain, burst activity is associated with epilepsy[6], neonatal seizures[5], and brain trauma[7], and indicates poorer prognosis in long-term clinical outcomes[8]. Recent work demonstrates that the dynamics of burst events closely mimic those of the seizure onset zone[9], suggesting that bursts may represent the arrest of nascent seizures, making them prime candidates for false detection.

Previous research by Wulsin and Fox[10] demonstrated that epileptic events can be modeled using a Bayesian nonparametric hidden Markov switching process. This approach enables parsing of EEG data based on underlying brain dynamics. In this work, we present a novel seizure detection algorithm based on this method of EEG analysis. By isolating and modeling specific epochs of EEG associated with transition to seizure activity, we are able to detect seizure onset in real time in a personalized manner not reliant on feature extraction. Using data recorded from dogs with naturally occurring epilepsy we demonstrate that this seizure detection algorithm may represent a substantial improvement in detection specificity with minimal on-line computational requirements.

2. Methods

2.1 Animals

The animals implanted in this study have been described earlier[11,12]. Mixed hounds with spontaneous seizures were implanted with a continuous intracranial recording device designed and manufactured by NeuroVista Inc. (Seattle, Washington). Standard human-sized strip electrodes with a total of 16 contacts were implanted in the subdural space to cover both hemispheres of the canine neocortex (Figure 1). All dogs had normal neurological examinations and MRI. The dogs were housed in the University of Minnesota canine epilepsy monitoring unit and continuously monitored (24 hours/day) with video and iEEG. See Supplemental Materials for more information regarding surgical technique and implanted device design.

2.2 Overview of algorithm

This novel seizure detection algorithm is designed to respond to the overall behavior of the EEG data rather than to extracted features. The model, which is described below, uses an ergodic hidden Markov process to parse regions of the iEEG to different states, including the pre-seizure state. This method is applied to a training dataset in order to identify iEEG states characteristic of the immediate pre-seizure state and to optimize model parameters. These states are then approximated using Gaussian models to allow for real-time, unsupervised seizure detection in a testing dataset.

Parsing iEEG recording using an Autoregressive Hidden Markov Model—In order to parse complex epileptic behavior into distinct dynamical regimes, we relied on a Bayesian nonparametric autoregressive Markov switching process. Due to the non-stationary behavior of EEG, a time-varying autoregressive (AR) process is used to model each channel's activity. The model also mimics focal changes in EEG by allowing for shared dynamical states among a variable number of iEEG channels and asynchronous state switching among channels via the Beta process. Notably, our data is characterized by inter-channel correlations, which may change over time as the patient progresses through various seizure event states. This model is described fully in Wulsin, 2014[9].

Using AR-HMM to define event states and to identify "States of Interest"—While the AR-HMM determines a dynamical state at each time point for each individual channel, in this work we relied on overall event states in order to capture global brain dynamics. These states are determined from the channel state transitions by a Dirichlet process as described in Wulsin, 2014[9]. Each time point of iEEG data was parsed into one of 30 event states. The number of event states was empirically chosen to capture a sufficiently wide range of EEG behaviors. Preliminary studies using this model to parse seizure activity demonstrated successful characterization of seizure dynamics, with identification of dynamical transitions in agreement with those identified manually by a board-certified epileptologist (Figure 2). These event states are used to identify personalized "states of interest" (SOIs) disproportionately enriched in pre-seizure zones, defined as the 30-second window prior to the unequivocal epileptic onset (UEO) for each seizure. The UEO is defined as the earliest time that seizure activity is evident to an epileptologist viewing EEG data

without prior knowledge of seizure occurrence. In each dog, three SOIs were identified with consistent appearance in pre-seizure zones with specificity >99%; however, the final number of SOIs chosen may vary among subjects based on individual seizure onset iEEG profiles. Notably, these SOIs were not found to occur more frequently during burst activity than at baseline ($p = 0.21$).

Adaptation of States of Interest for Online Seizure Detection—Identification of pre-seizure SOIs is not directly useful for real-time seizure detection for several reasons. First of all, determination of event states for each time point by the AR-HMM requires the entire time series to be analyzed at once as the state determinations are not independent in time, thereby preventing predictive use. Secondly, determination of event states by this model is computationally intensive; even if an approximate predictive model were designed, the hardware demands for such computation would render it unsuitable for use in an implantable device.

Therefore, in order to translate this approach into a real-time detector, we developed a model informed by existing AR-HMM event states designed to predict future event states. The emission probabilities of each state are modeled as a Gaussian distribution, resulting in an N-dimensional vector mean ($\boldsymbol{\mu}$) and an N×N covariance matrix ($\boldsymbol{\Sigma}$) associated with each state, where N is the number of channels (Figure 3). These newly defined Gaussian distributions are then used to calculate to which state an incoming time point would most likely belong using maximum likelihood estimation. This calculation allows for the assignment of data to approximated event states in real time with minimal computational overhead. The model also included three event states designed to identify artifacts, generated by modeling artifact data with a three-component Gaussian mixture model.

2.3 Data segmentation and testing method

Model testing was carried out to simulate online detection. Each data set was sequentially segmented into a 30-day “burn-in period” to be discarded, a 60-day training data set, and the remainder to be used as a testing data set. The first 30 days of recording are removed due to significant short-term impedance changes post-implantation[13]. The training period in Dog 1 contained 14 seizures and 61 bursts, while the training period in Dog 2 contained 12 seizures and 524 bursts. This segmentation produced a testing dataset for Dog 1 of length 337 days with 20 seizures and 407 bursts, and a testing dataset for Dog 2 of 54 days with 35 seizures and 389 bursts.

The training data set was segmented into true event states using the AR-HMM. These states were then used as described previously in creating of Gaussian models to be used for online detection. Thus, the dataset used for algorithm testing was kept separate from all data used to inform the model.

Pre-seizure states were identified based on a sliding window of incoming data. If the percentage of points within the window identified as SOIs exceeds a specified threshold, the detector signals that a pre-seizure state is present. The specific window length and threshold value were optimized over the training data set by sampling the parameter space to provide the fewest false positives possible while ensuring that all seizures were detected prior to the

UEO (zero false negatives). In this testing, once a seizure is flagged the detector is deactivated for five minutes in order to prevent multiple detections of the same event. This method is demonstrated in Figure 4.

2.4 Development of control detectors

A feature-based detector modeled after the NeuroPace seizure detection system[14] was developed to serve as a control. This detector was responsive to signal line length[4], halfwave[15], and area under the curve[2]. Thresholds for each feature were determined graphically by plotting feature values over time in order to ensure that the threshold chosen was both sensitive and specific for seizure onset. The specific binary operations among the three features used to detect seizure onset were determined through optimization on the training data to limit false positives while preserving a false negative rate of zero.

A state-of-the-art seizure detection algorithm was also used for comparison. This algorithm, developed by Michael Hills, was the winner of the 2014 Kaggle competition for seizure detection[16]. This model extracts cross-channel correlations in the time and frequency domains as features, which are then used to train per-patient Random Forest classifiers. Test data is classified in one-second windows as ictal or interictal. The code for this algorithm is freely available, and was adapted to allow for seizure detection over continuous recordings. Note that this algorithm is not designed to satisfy the computational limitations of real-time detection in an implantable device, but rather represents the current upper limit for offline detection accuracy.

3. Results

The efficacy of each algorithm was assessed by the false negative and false positive rates, as well as the latency of each seizure call (Table 1). The latency was measured relative to the UEO of each seizure. Of the twenty seizures in the testing dataset for Dog 1, all twenty were detected by both the HMM-Gaussian model and the feature-based detector with average latencies of 12.1 ± 69 seconds and 18.5 ± 4.9 seconds before the marked UEO, respectively. Over the 337 days of recorded data, the HMM-Gaussian model returned 5 false positives (6.2×10^{-4}/hr; 0.25/seizure), while the feature-based detector returned 116 false positives (1.4×10^{-2}/hr; 5.8/seizure). For Dog 2, all thirty-five of the seizures were detected by the HMM-Gaussian model with an average latency of 10.7 ± 8.1 seconds before the UEO while the feature-based model detected thirty-three of thirty-five with average latency of 19.0 ±12.7 seconds before the UEO. Over 54 days of recorded data, the HMM-Gaussian model returned 6 false positives (4.6×10^{-3}/hr; 0.17/seizure) while the feature-based model returned 430 (0.33/hr; 12.3/seizure). Over both dogs, the seizure flags by the HMM-Gaussian method ranged in latency from 4 to 24 seconds before the UEO; the feature-based model ranged from 8–23 seconds before the UEO. In Dog 1, the rate of false positives in the feature-based detector (1.4×10^{-2}/hr) matches the published rate of the NeuroPace device (1.3×10^{-2}/hr)[14], suggesting that this model is an appropriate control for comparison to devices used in practice. It is important to note that each animal tested had multiple seizure onset types that, while falling into a range of similar morphologies, were variable enough in their temporal characteristics as to challenge standard seizure detection algorithms. Figure 5 depicts the

range of seizure onset patterns in one test animal. The performance metrics of the Kaggle-winning algorithm on this dataset are similar to those of the HMM-Gaussian model in both dogs, demonstrating that our method performs comparably to this state-of-the-art detector over these datasets.

In an effort to decrease false positives produced by the feature-based algorithm, the feature thresholds were revised to make the model more robust to bursts. Parameters were plotted as described above, and rather than selecting feature cutoffs based on baseline values, cutoffs were chosen to exceed the maximum activity seen during bursts in the training set. This revised model showed modestly improved false positive rates and slightly decreased latencies. In Dog 1, 71 false positives (8.8×10^{-3}/hr; 3.55/seizure) were detected with latency 15.7±3.8 seconds before the UEO; in Dog 2, 232 false positives (0.11/hr; 6.63/seizure) were detected with latency 15.6±5.1 seconds before the UEO.

Given the efficacy of the HMM-Gaussian model using the full electrode grid, we sought to examine the utility of this model using a subset of electrodes. Long-term implanted neural recording devices are hampered by inflammation and resulting gliosis at electrode sites[17–19]. Limitation of the number of recording electrodes offers the potential to decrease local tissue damage associated with long-term implantation while also decreasing surgical implant complexity and device cost. We investigated four configurations of electrode subsets in each animal: 8 leads confined to either the left or right hemisphere, 8 leads spread uniformly over the entire cortex, and 4 leads targeted most closely to the seizure onset zone. The presumed seizure onset zone was determined as the location with the earliest electrographic change prior to seizure propagation. The results of these trials are shown in Table 2. We discovered that selection of an appropriate subset of electrodes provided comparable results to those obtained using the full grid in Dog 1, with high accuracy achieved using only 4 optimally chosen electrodes. However, all electrode subsets selected in Dog 2 resulted in significant decreases in detection accuracy with multiple missed seizures. The decision to use a smaller number of implanted electrodes must be based on the individual subject's seizure profile, with greater efficacy expected in subjects with clear seizure onset zones, stereotyped pre-seizure electrographic changes, and few interictal bursts.

4. Discussion

The high rate of false positives inherent in current feature-based detection has hampered use of these systems in practice, as it is associated with unnecessary stimulation and decreased battery life[3]. While both the feature-based and HMM-Gaussian detectors correctly identified all seizures in Dog 1 and the vast majority of seizures in Dog 2, the HMM-Gaussian detector demonstrated a drastic reduction in false positive detections. Both detectors were able to consistently identify seizures in a clinically useful manner for closed-loop stimulation, well before the UEO; however, the feature-based detector was able to detect seizures 5–10 seconds earlier on average. Although proof of effective suppression of seizure symptoms must be definitely evaluated in *in vivo* studies, this finding provides a high level of confidence that the HMM-Gaussian detector consistently identifies seizure onset zones before seizure generalization occurs, at a time when clinical intervention is possible. This belief is supported by the fact that the NeuroPace seizure detector, which has shown a degree

of efficacy in symptom suppression in practice (seizure reduction of 40% relative to baseline)[20], has a published latency of 5.01 seconds after the UEO[14]. Clinical testing is needed to determine whether the rate of seizure aversion might be improved with earlier seizure detection and stimulation. Furthermore, the detection algorithm we have designed is ideal for incorporation into an implantable device. All computationally difficult calculations are performed externally during model setup for analysis of training data. Categorization of incoming data into estimated event states in real time requires only a single matrix multiplication per data point, allowing for high time-resolution sampling with minimal hardware requirements.

False positive detections flagged by the feature-based detector are not distributed at random throughout the recording. Rather, these calls are clustered in areas of high seizure activity. In particular, false positives tend to occur during or in close association with bursts (Figure 6), and retraining of the algorithm to limit detection of bursts effected a significant (but not complete) reduction in false positives. In contrast, the HMM-Gaussian seizure detector is based on identification of event states that are specifically chosen to be absent from bursts and surrounding background. This method greatly increases robustness to bursts, thereby eliminating a major source of false positive readings. In addition to accomplishing our primary goal of reliable seizure detection with greatly reduced false positive rate, we have shown that it is possible to identify specific epochs of iEEG behavior that are useful for distinguishing bursts from nascent seizures. A more thorough investigation of these periods may yield interesting information regarding the brain dynamics that modulate the variable arrest of propagation of epileptic activity.

The performance of the HMM-Gaussian algorithm was also compared to that of the winning seizure detection algorithm from the Kaggle competition, the Hills algorithm. Both methods detected all seizures with nearly identical numbers of false positive detections and no significant difference in detection latency. However, it is important to recognize that the Hills algorithm cannot be directly applied for use in an implantable device. Given the current hardware limitations of these devices, it is unfeasible to locally carry out bandpower computation, signal correlation, and classification using a Random Forest of 3000 trees in real time[21]. Nevertheless, the performance of the Hills algorithm over a range of human and canine datasets justifies its use as a benchmark for attainable detection accuracy[16]. Our work has demonstrated that the HMM-Gaussian algorithm can match the performance of this state-of-the-art offline detection method while dramatically reducing computational requirements during live detection.

Interaction between the computationally intensive training method we use in this paper and its low-computational overhead implementation suitable for an implantable devices is worthy of comment. With increasing availability of central, cloud-based computational and data integration, collection of data from individual devices, central training, and periodic updates of implantable devices is now a reality. Data for this study were all stored and processed on http://ieeg.org, an NIH-funded, cloud-based platform hosted on Amazon's Elastic Computing Cloud (EC2). Our laboratory has used this platform to collect dispersed data for training from remote sites, for central data review, annotation and algorithm training, and then individual device updates. One important implication of our study is that it

validates this type of paradigm for both research and its implementation in clinical care[22]. This work is currently formally under way as a testbed for a variety of device implementations.

5. Limitations

One limitation of this study is the small number of test animals, despite the prolonged recording periods and large number of seizure detections used in testing and validating algorithm performance. For this reason it is important that we not overemphasize the significance of our findings. The study presented is meant to elucidate the potential of a new computational method for responsive implanted antiepileptic devices, but is not meant to definitively validate its utility. It is important to note that seizures demonstrated considerable variability in the animals' studies, with a variety of onset patterns and locations (Figure 4) that challenged the NeuroPace-like detector. The animals studied also had a significant number of interictal epileptiform bursts that challenged standard detection algorithms. For these reasons we feel that the present study, despite being validated on only two prolonged recordings, demonstrates that the method is promising. In addition, the epileptic profiles of the dogs used in this study differ significantly. Dog 2 was characterized by considerably increased seizure and burst activity, likely responsible for the elevated rate of false positive detections produced by both algorithms and demonstrating that the accuracy of these algorithms is dataset-dependent. Certainly, validation on a much larger data set of continuous animal or human recordings, as may be available from the NeuroVista human data set[23], would be required to definitively prove potential benefit over current detection strategies.

It is also important to recognize that the direct clinical impact of a drastic reduction in false positive detections is unclear, and that devices that stimulate much more selectively could potentially provide less seizure control in their current embodiment. While it stands to reason that precise seizure detection may provide relief from symptoms without unnecessary stimulation, it is also possible that the efficacy of the NeuroPace device in practice stems in part from generation of a low-seizure state secondary to frequent stimulation. It had been demonstrated that frequent electrical stimulation may cause persistent alterations in brain function in epilepsy regardless of whether these stimulations are correlated with seizure onset[24–26]. *In vivo* testing is required to parse the relative importance of on-and off-target stimulation in seizure prevention, and to determine if highly targeted stimulation alone is clinically effective. It may also be that parameters not accounted for in the current first generation responsive devices, such as the relation between phase of target waveform and stimulation or synchronization with two-dimensional wavefronts, may be required to optimize selective responsive stimulation. In addition, characterization of changes in the distribution of iEEG states in the days to weeks preceding a seizure may provide avenues for identifying periods of high seizure susceptibility, rather than pinpointing precise seizure onset states. It is possible that identifying and targeting these epochs of epileptic vulnerability could prove to be more effective for modulatory stimulation than specific immediately pre-seizure states.

An important advantage of the HMM-Gaussian detection system is that it is tuned to an individual's particular seizure onset pattern. Event states are chosen as important for seizure detection based on their role in a particular patient, allowing this method to capture the individual's specific seizure morphology. However, this approach could, in theory, result in missed seizures if seizure morphology or location of onset were to change significantly over time. Long-term recording of electrographic signals may also be affected by chronic local inflammation or gradual changes in electrode impedance. While we did not observe any decrease in sensitivity, extended studies must be carried out to determine if model retraining is necessary. Implementation of this model in a portable device would provide extensive patient recordings such that it may be retrained on recent pre-seizure periods to ensure continued accuracy.

6. Conclusion

In this work, we present a novel algorithm for individualized seizure detection suitable for use in a closed loop, implantable system. We have demonstrated that modeling seizure activity using an autoregressive hidden Markov model may provide insights into novel methods of characterizing and analyzing EEG data. This algorithm represents a substantial improvement in accuracy of seizure detection over the industry standard, achieving a nearly 98% reduction in false positive rate while slightly improving detection sensitivity. Furthermore, the performance of this model closely matches that of the winning algorithm from the Kaggle seizure detection competition. Further *in vivo* study must be carried out in order to assess the potential clinical implications of this technology. This work also demonstrates potential for a new pipeline for individualized device data collection, training, and reprogramming utilizing a central cloud-based platform.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by National Institutes of Health (NIH) and the Mirowski Family Foundation grants through the University of Pennsylvania and Mayo Clinic. National Institutes of Health (NIH) (1-P20-NS-080181-01); NIH (U01-NS-073557-01A1). Canine iEEG data was collected by Ned Patterson (University of Minnesota), Charles Vite (University of Pennsylvania), Ben Brinkman (Mayo Clinic), and Daniel Crepeau (Mayo Clinic). The International Epilepsy Electrophysiology Portal is funded by the NIH (5-U24-NS-063930-05).

References

1. Kwan P, Brodie MJ. Early identification of refractory epilepsy. N. Engl. J. Med. 2000; 342:314–319. [PubMed: 10660394]
2. Gotman J. Automatic recognition of epileptic seizures in the EEG. Electroencephalogr. Clin. Neurophysiol. 1982; 54:530–540. [PubMed: 6181976]
3. Stacey WC, Litt B. Technology insight: neuroengineering and epilepsy-designing devices for seizure control. Nat. Clin. Pract. Neurol. 2008; 4:190–201. [PubMed: 18301414]
4. Esteller, R.; Echauz, J.; Tcheng, T.; Litt, B.; Pless, B. 2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Vol. 2. IEEE; 2001. Line length: an efficient feature for seizure onset detection; p. 1707-1710.

5. Nagarajan L, Palumbo L, Ghosh S. Brief electroencephalography rhythmic discharges (BERDs) in the neonate with seizures: their significance and prognostic implications. J. Child Neurol. 2011; 26:1529–1533. [PubMed: 21652591]
6. Gambardella A, et al. Usefulness of focal rhythmic discharges on scalp EEG of patients with focal cortical dysplasia and intractable epilepsy. Electroencephalogr. Clin. Neurophysiol. 1996; 98:243–249. [PubMed: 8641147]
7. Claassen J, et al. Electrographic seizures and periodic discharges after intracerebral hemorrhage. Neurology. 2007; 69:1356–1365. [PubMed: 17893296]
8. Orta DSJ, Chiappa KH, Quiroz AZ, Costello DJ, Cole AJ. Prognostic implications of periodic epileptiform discharges. Arch. Neurol. 2009; 66:985–991. [PubMed: 19667220]
9. Wulsin DF, Fox EB, Litt B. Modeling the Complex Dynamics and Changing Correlations of Epileptic Events. Artif. Intell. 2014; 216:55–75. [PubMed: 25284825]
10. Wulsin DF. Bayesian Nonparametric Modeling of Epileptic Events. Dissertations available from ProQuest. 2013:1–184. at <http://repository.upenn.edu/dissertations/AAI3566348>.
11. Davis KA, et al. A novel implanted device to wirelessly record and analyze continuous intracranial canine EEG. Epilepsy Res. 2011; 96:116–122. [PubMed: 21676591]
12. Howbert JJ, et al. Forecasting seizures in dogs with naturally occurring epilepsy. PLoS One. 2014; 9:e81920. [PubMed: 24416133]
13. Sillay KA, et al. Long-term measurement of impedance in chronically implanted depth and subdural electrodes during responsive neurostimulation in humans. Brain Stimul. 2013; 6:718–726. [PubMed: 23538208]
14. Echauz J, et al. Long Term Validation of Detection Algorithms Suitable for an Implantable Device. Epilepsia. 2001; 42(suppl. 7):35–36. [PubMed: 11887966]
15. Osorio I, Frei MG, Wilkinson SB. Real-Time Automated Detection and Quantitative Analysis of Seizures and Short-Term Prediction of Clinical Onset. Epilepsia. 1998; 39:615–627. [PubMed: 9637604]
16. Hills M. Seizure detection using FFT, temportal and spectral coefficients, eigenvalues, and Random Forest. 2014
17. Polikov VS, Tresco PA, Reichert WM. Response of brain tissue to chronically implanted neural electrodes. J. Neurosci. Methods. 2005; 148:1–18. [PubMed: 16198003]
18. Biran R, Martin DC, Tresco PA. Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. Exp. Neurol. 2005; 195:115–126. [PubMed: 16045910]
19. Green RA, Lovell NH, Wallace GG, Poole-Warren LA. Conducting polymers for neural interfaces: challenges in developing an effective long-term implant. Biomaterials. 2008; 29:3393–3399. [PubMed: 18501423]
20. Morrell MJ. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology. 2011; 77:1295–1304. [PubMed: 21917777]
21. Sun FT, Morrell MJ, Wharen RE. Responsive cortical stimulation for the treatment of epilepsy. Neurotherapeutics. 2008; 5:68–74. [PubMed: 18164485]
22. Wagenaar, JB.; Brinkmann, BH.; Ives, Z.; Worrell, GA.; Litt, B. 2013 6th International IEEE/EMBS Conference on Neural Engineering (NER). IEEE; 2013. A multimodal platform for cloud-based collaborative research; p. 1386-1389.
23. Cook MJ, et al. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. Lancet. Neurol. 2013; 12:563–571. [PubMed: 23642342]
24. Girgis M. Kindling as a model for limbic epilepsy. Neuroscience. 1981; 6:1695–1706. [PubMed: 7301119]
25. Goddard GV, McIntyre DC, Leech CK. A permanent change in brain function resulting from daily electrical stimulation. Exp. Neurol. 1969; 25:295–330. [PubMed: 4981856]
26. Theodore WH, Fisher RS. Brain stimulation for epilepsy. Lancet. Neurol. 2004; 3:111–118. [PubMed: 14747003]

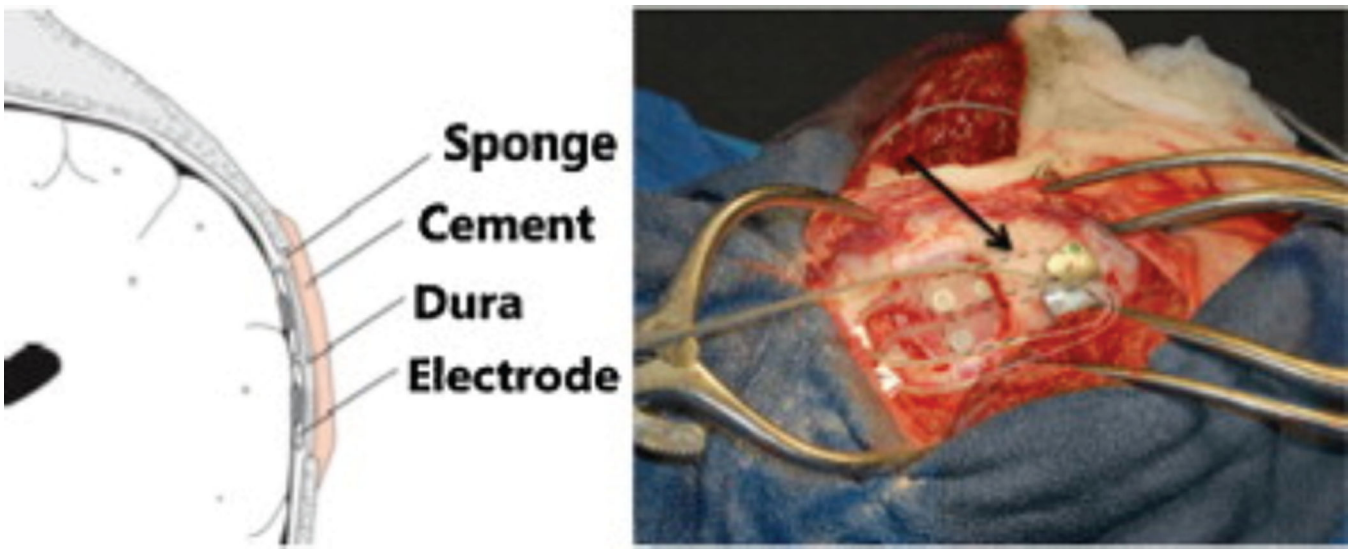

Figure 1.

Representation of electrode implantation location in canine neocortex. Figure reprinted with permission from Davis, Kathryn *et al.* A novel implanted device to wirelessly record and analyze continuous intracranial canine EEG. *Epilepsy Research* **96**, 116–122 (2014).

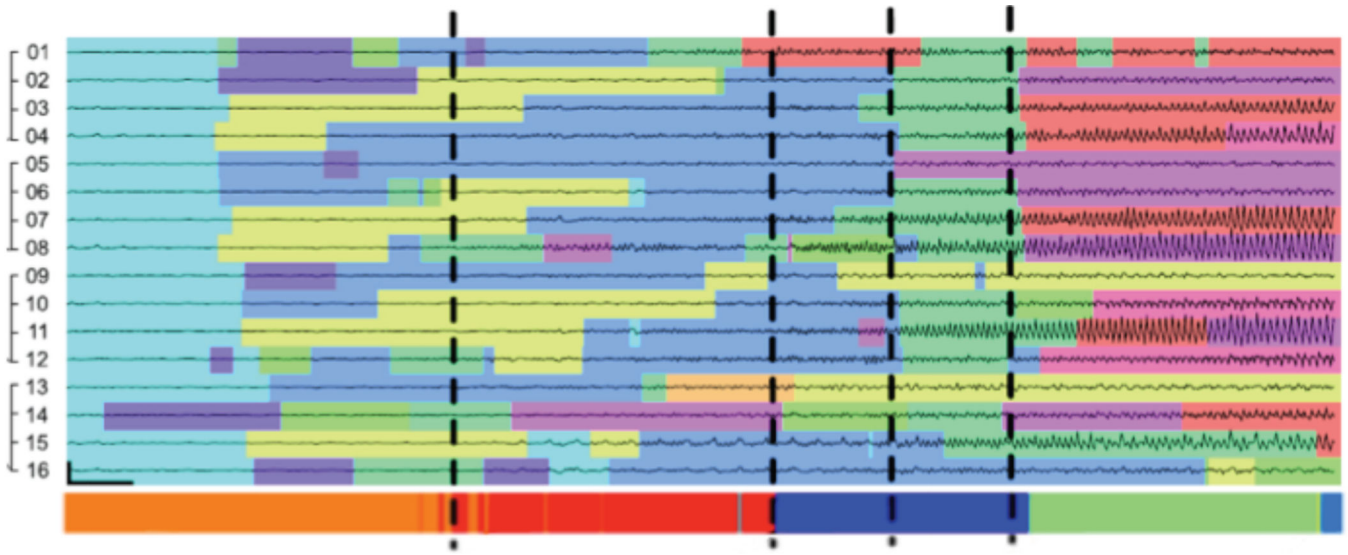

Figure 2.

Data traces from the 16 channel iEEG electrode over 25 seconds of a seizure onset with colors indicating the inferred channel states. Vertical dashed lines indicate the EEG transition times marked independently by an epileptologist. Figure adapted from Wulsin, et al., Modeling the Complex Dynamics and Changing Correlations of Epileptic Events, *Artif. Intell.* **216**, 55–75 (2014).

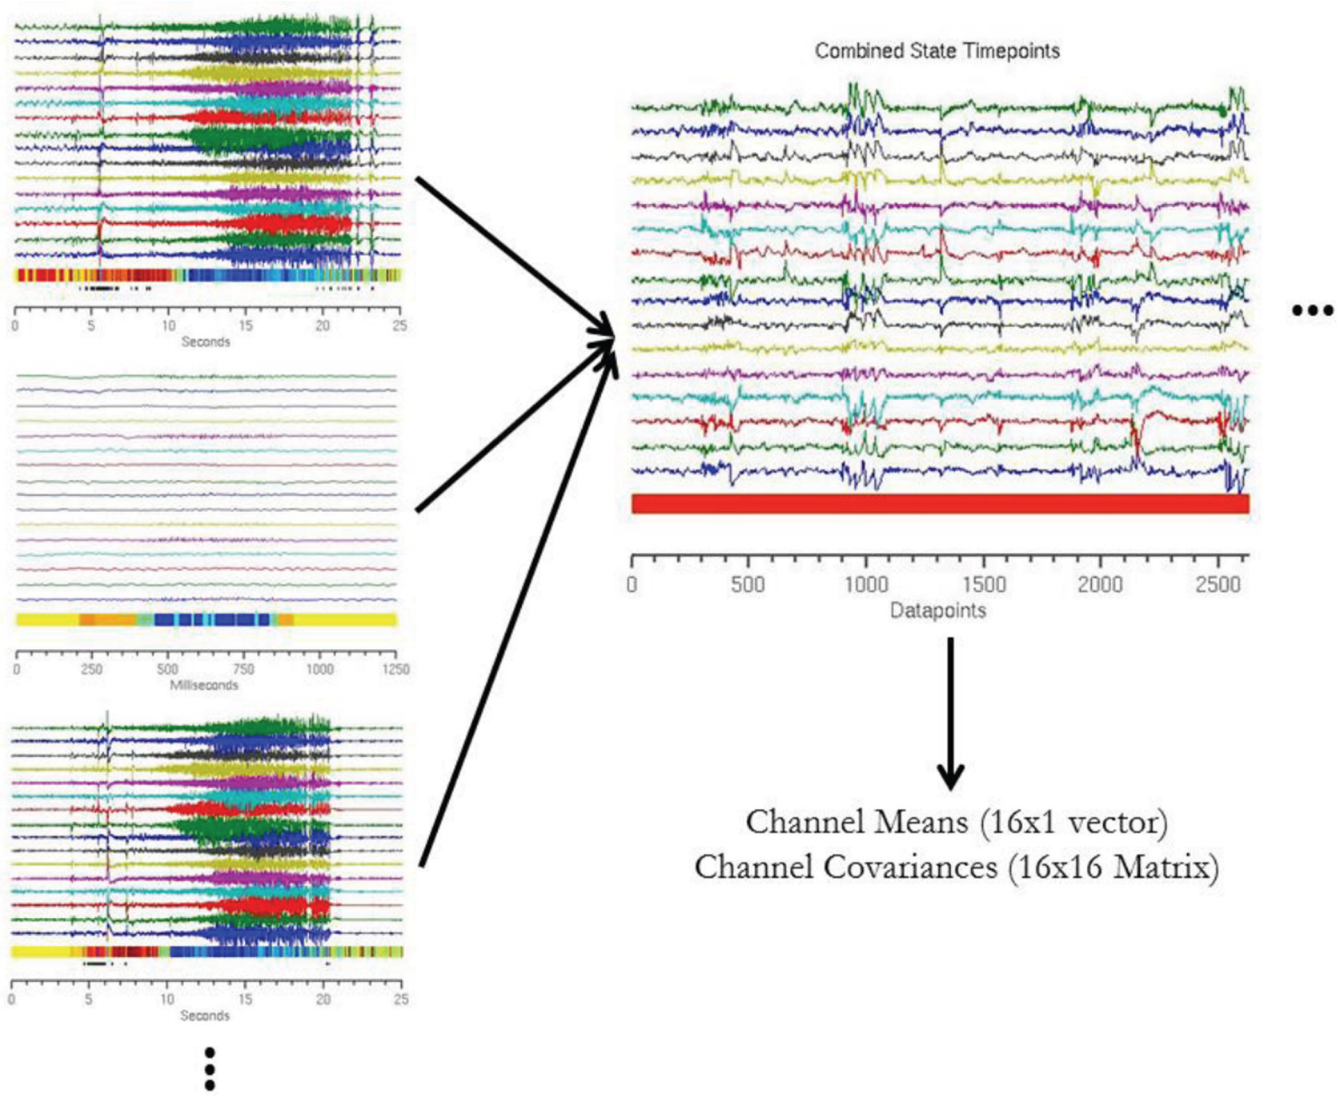

Figure 3.

Schematic of construction of Gaussian distributions. Each time point in the training dataset is assigned an event state based on the AR-HMM. Three sample recording clips are shown, with time points color coded by event state. The time points are then segregated by state to generate a unique dataset for each event state. These datasets are then modeled as a 16-dimensional Gaussian distribution. Figure above illustrates this operation for a single state (color coded as red).

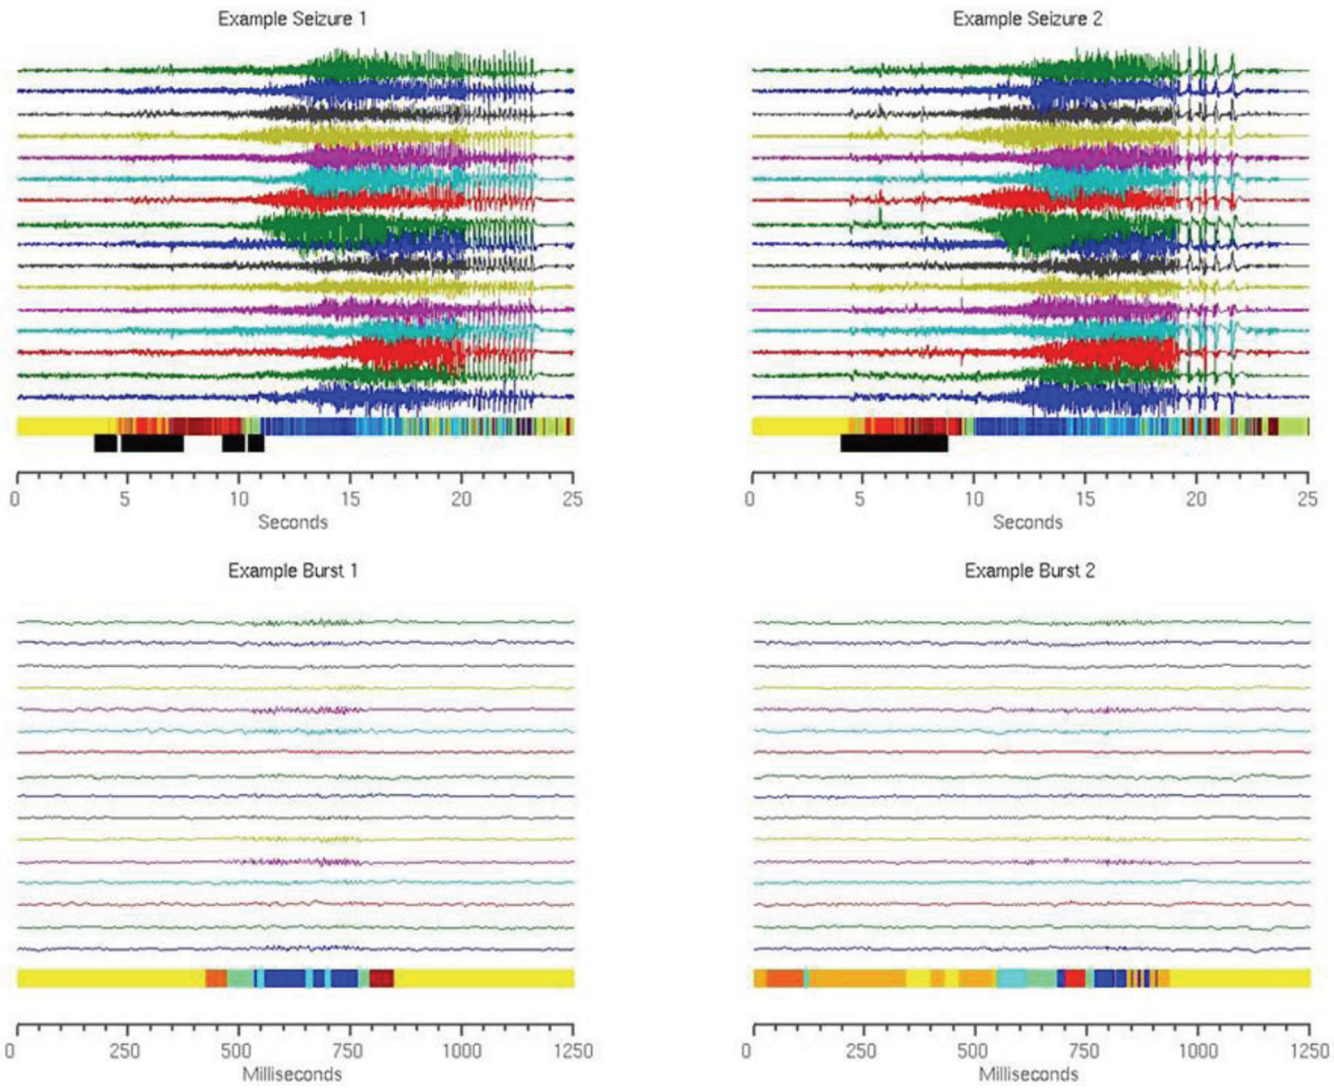

Figure 4.

Sample seizures and bursts from Dog 1. Each time point is color coded by real event state as determined by the AR-HMM. The same seizures and bursts were evaluated using the online Gaussian detection method to determine approximate event states and to identify time points at which the most likely state is an SOI. Time points that fall in sliding windows in which the percentage of predicted SOIs exceeds a set threshold are marked with black bars. These windows are sensitive and specific for seizure onset zones.

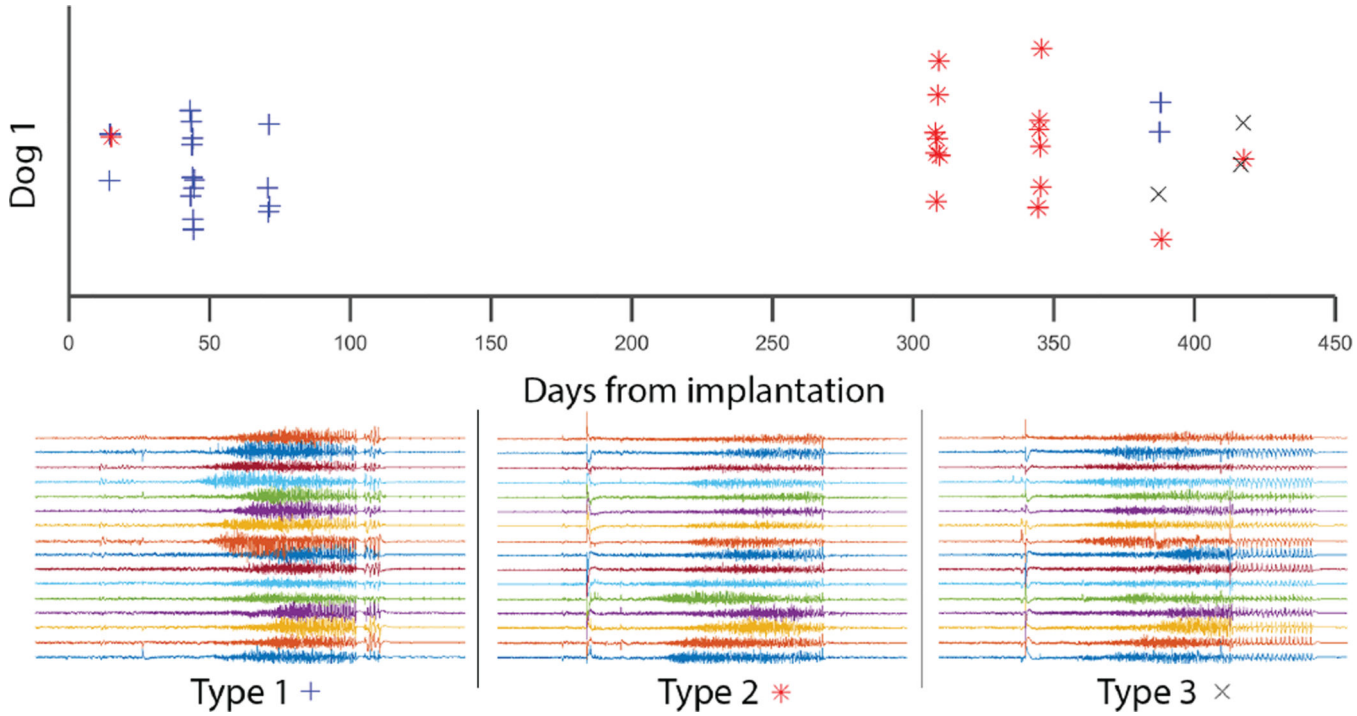

Figure 5.

Seizure onset patterns and locations are variable over the course of the recordings. Three major categories of seizure patterns are demonstrated in Dog 1 as determined by manual review by an epileptologist. Each seizure type had a different initial onset location.

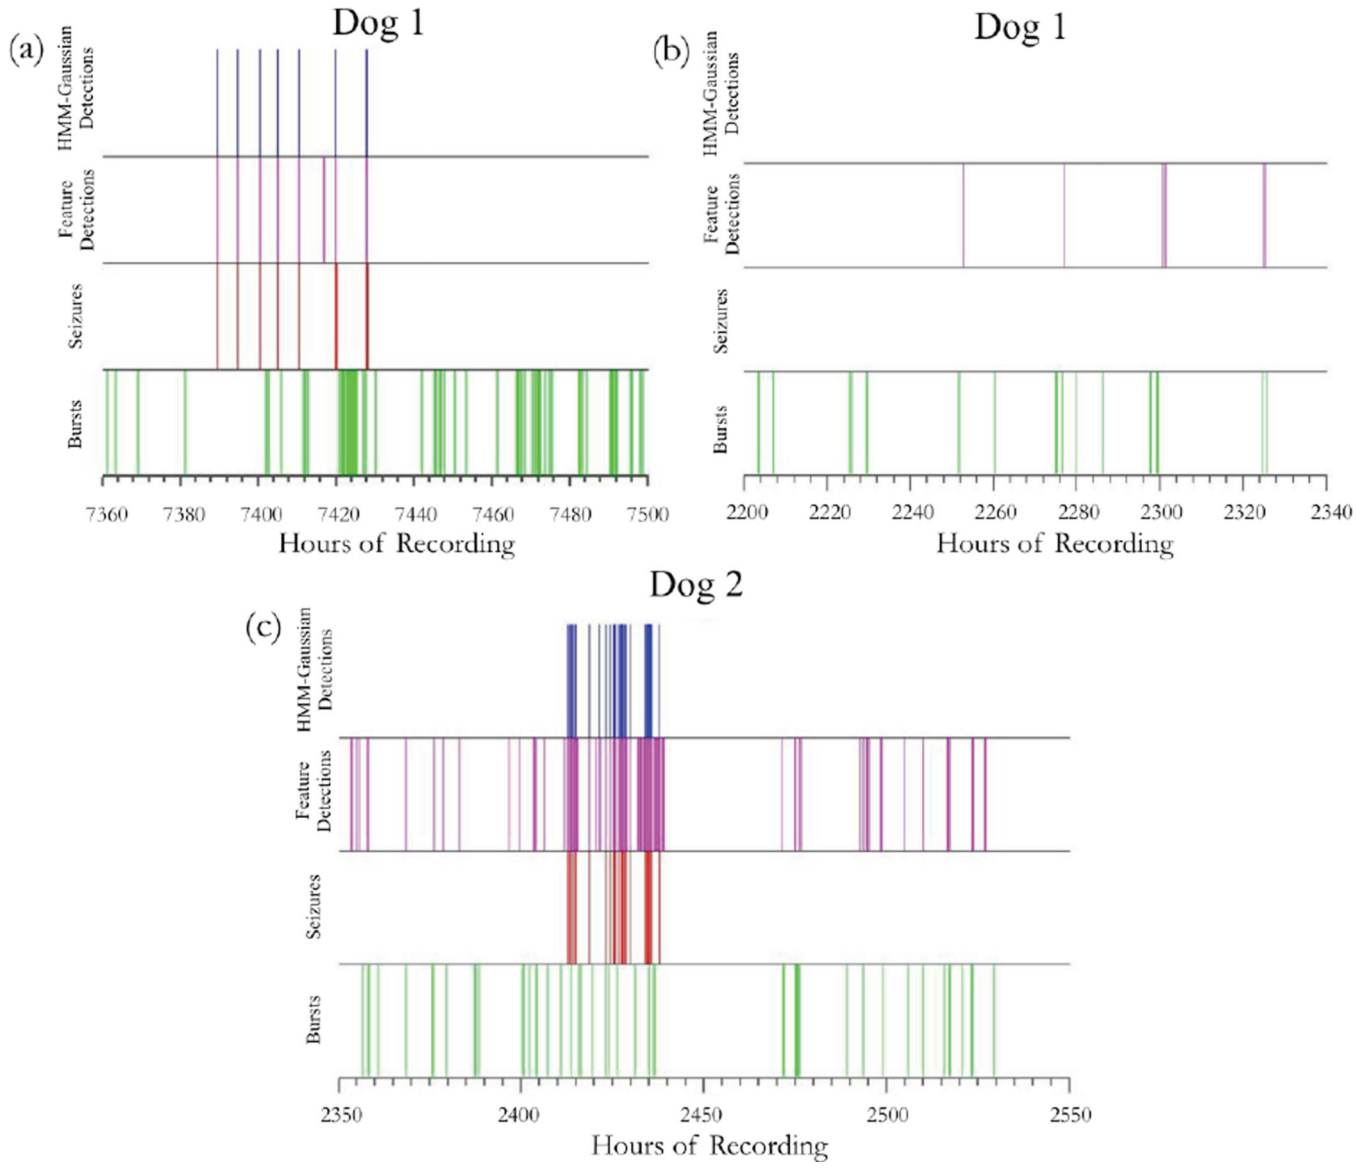

Figure 6.

Representative timeline excerpts from Dog 1 (a and b) and Dog 2 (c). The top row represents detected seizures by the HMM-Gaussian detector, the second row represents detected seizures by the feature-based, NeuroPace-like detector, the third row represents real seizures, and the bottom row represents real bursts. Note that false positives flagged by the NeuroPace-like detector tend to cluster with interictal bursts.

Table 1

Performance metrics of HMM-Gaussian, feature-based, and Kaggle-winning methods.

Dog	Method	FN	FP (rate/hr)	Latency (s)
1	HMM-Gaussian	0	5 (6.2e-4)	−12.1 ± 6.9
1	Feature	0	116 (1.4e-2)	−18.5 ± 4.9
1	Feature *	0	71 (8.8e-3)	−15.7 ± 3.8
1	Kaggle	0	3 (3.7e-4)	−10.1 ± 5.5
2	HMM-Gaussian	0	6 (4.6e-3)	−10.7 ± 8.1
2	Feature	2 (0.057)	430 (0.33)	−19.0 ± 12.7
2	Feature *	2 (0.057)	232 (0.18)	−15.6 ± 5.1
2	Kaggle	0	7 (5.4e-3)	−8.6 ± 4.2

Feature* indicates that the method was trained specifically to limit false positives during bursts.

FN = false negatives (missed seizures); FP = false positives. Latency is measured relative to UEO.

Table 2

Performance metrics of HMM-Gaussian method using electrode subsets.

Dog	Leads	FN	FP (rate/hr)	Latency (s)
1	Left	0	10 (1.2e-3)	−13.5 ± 6.0
1	Right	7	1 (1.2e-4)	−11.0 ± 8.9
1	Distributed	1	4 (4.9e-4)	−12.4 ± 6.7
1	Targeted	0	8 (9.9e-4)	−13.8 ± 5.8
2	Left	3	37 (1.8e-2)	−9.3 ± 7.6
2	Right	5	33 (1.6e-2)	−11.7 ± 8.8
2	Distributed	2	22 (1.0e-2)	−9.4 ± 8.3
2	Targeted	3	30 (1.4e-2)	−15.0 ± 10.3