33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Capturing the State Transitions of Seizure-Like Events using Hidden Markov Models

Mirna Guirgis, Demitre Serletis, Peter L. Carlen, and Berj L. Bardakjian, *Member, IEEE*

***Abstract*—The purpose of this study was to investigate the number of states present in the progression of a seizure-like event (SLE). Of particular interest is to determine if there are more than two clearly defined states, as this would suggest that there is a distinct state preceding an SLE. Whole-intact hippocampus from C57/BL mice was used to model epileptiform activity induced by the perfusion of a low Mg^{2+}/high K^{+} solution while extracellular field potentials were recorded from CA3 pyramidal neurons. Hidden Markov models (HMM) were used to model the state transitions of the recorded SLEs by incorporating various features of the Hilbert transform into the training algorithm; specifically, 2- and 3-state HMMs were explored. Although the 2-state model was able to distinguish between SLE and nonSLE behavior, it provided no improvements compared to visual inspection alone. However, the 3-state model was able to capture two distinct nonSLE states that visual inspection failed to discriminate. Moreover, by developing an HMM based system *a priori* knowledge of the state transitions was not required making this an ideal platform for seizure prediction algorithms.**

I. Introduction

THE degree of synchronization within neuronal networks of the brain can be a strong indicator of its overall health. Periods of high and/or extended synchronized neuronal discharges, designated as seizures or ictal events, cause transient interruptions of the brain's electrical activities [1]. These events may be spontaneous or triggered by any number of factors such as abnormal metabolic states (e.g., sleep deprivation, fever, etc.) or visual patterns [2]. Healthy brains respond by appropriately modifying the electrical responses whereas diseased brains (e.g., epileptic) may respond by transitioning into an ictal state. It is the nature of these transitions that is of particular interest when studying diseased brain dynamics.

Seizure onset prediction has been on the forefront of epilepsy research for several decades [3]. Although most patients are able to adequately control seizures through anticonvulsant medications, approximately 30% experience medically intractable epilepsy [4]. Thus, seizure prediction systems may provide more efficient and effective treatments. Prediction algorithms based on dynamic similarity [5], signal energy [6], largest Lyapunov exponent [7], and phase synchronization and cross correlation measures [8] have been developed. However, one of the difficulties with such algorithms is the dependence on comparing preictal data with baseline [9]. By definition this dependence incorporates bias into the system, as the time series being analyzed needs to be segmented prior to the analysis. All supervised learning algorithms are inherently biased in a similar manner. Thus, unsupervised learning would be the method of choice for an ideal seizure prediction system.

On the most primitive level, the notion of predicting the onset of a transition into an ictal state suggests that there is a preictal state wherein the brain is progressing towards ictal behavior. This would further imply that there are, at least, three distinct states – namely, preictal, ictal, and postictal (or baseline) – involved in the progression of a seizure-like event (SLE). It would then follow that an ideal seizure prediction system would be able to identify these states in an unsupervised manner.

The overarching aim of this work is to capture the dynamical states involved in epileptiform activity recorded from whole-intact hippocampal preparations. Extracellular field recordings were obtained from pyramidal neurons in the cornu ammonis region 3 (CA3) and analyzed offline by modeling them as Markov processes. Specifically, a hidden Markov model (HMM) based system is explored to determine if a two or three state model provides a more appropriate representation of SLE behavior. HMMs have been commonly used in speech recognition algorithms [10] but are also gaining prominence in brain state classification systems [11]. Such a system assumes that the underlying dynamical states are driving the observable changes recorded in the local field potentials and, more importantly, it does not require any *a priori* knowledge of the state progressions of the time series. Frequency ranges including delta (1–4 Hz), beta (15–40 Hz), and gamma (>40Hz) rhythms were incorporated into the system, which have been previously implicated in rodent epileptiform activity [12]. Features of the Hilbert transform – namely, amplitude, phase, and their respective derivatives – of the relevant frequency bands were used to construct the feature space training the HMMs.

Manuscript received April 15, 2011. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada under Grant 371588.

M. Guirgis is with the Institute of Biomaterials and Biomedical Engineering (IBBME), University of Toronto, Toronto, M5S 3G9, Canada (phone: +1-416-978-7855; e-mail: mirna.guirgis@utoronto.ca).

D. Serletis was with IBBME and the Department of Physiology (Dept. of Physiol.), University of Toronto, Toronto, M5S 1A8, Canada. He is now with the Division of Neurosurgery, Toronto Western Hospital, Toronto, M5T 2S8, Canada (e-mail: demitre.serletis@utoronto.ca).

P. L. Carlen is with IBBME, Dept. of Physiol., and the Toronto Western Research Institute, University Health Network, Toronto, M5T 2S8, Canada (e-mail: carlen@uhnresearch.ca).

B. L. Bardakjian is with IBBME and the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, M5S 3G4, Canada (e-mail: berj@cbl.utoronto.ca).

II. Materials and Methods

A. Hippocampal Preparation

Whole-intact hippocampal tissue was prepared from C57/BL (P10-14) mice in accordance with the animal care guidelines of the affiliated institutions. Animals were first anesthetized with halothane and decapitated. Once the brain was extracted the cerebellum was removed and the two hemispheres, separated along the sagittal plane, were submerged for 5 minutes in ice-cold (2-5^{o}C), oxygenated (95% O_2, 5% CO_2) artificial cerebrospinal fluid (ACSF), which contained (in mM): 123 NaCl, 2.5 KCl, 1.5 $CaCl_2 \bullet 2H_2O$, 2 $MgSO_4 \bullet 7H_2O$, 25 $NaHCO_3$, 1.2 $NaH_2PO_4 \bullet H_2O$, and 25 glucose (pH 7.4). The septal region and the ventral extension of each hemisphere were cut in order to detach the hippocampus while preserving the subiculum and entorhinal cortex [13]. The disconnected whole hippocampi were then submerged in room temperature, oxygenated ACSF for at least 1 hour prior to data collection.

B. Electrophysiological Recordings

Recordings were obtained from the prepared tissue using an RC-26 open bath recording chamber (Warner Instruments). Fine pins were used to secure the tissue in the chamber with its concave, medial surface facing downward. A steady flow of warmed (33.5 ± 0.5^{o}C), oxygenated (95% O_2, 5% CO_2) ACSF was perfused over the tissue at a rate of 2–3 ml/min. To minimize oxygen evaporation, the surface solution was also oxygenated (95% O_2, 5% CO_2). To induce seizure-like activity a low-Mg^{2+}/high-K^+ ACSF solution was used, which differed slightly from the standard ACSF previously described – namely, by using 0.25 mM $MgSO_4 \bullet 7H_2O$ and 5 mM KCl instead of 2 mM and 2.5 mM, respectively. This *in vitro* model is often used to induce epileptiform activity similar to that observed in *in vivo* electrographic seizures [14]. Application of low-Mg^{2+}/high-K^+ ACSF was implemented after at least 5 minutes of stable recording under the standard ACSF perfusion treatment. Continuous voltage recordings of the local network captured the transitions into SLEs for a maximum of 1 hour at which point standard ACSF was used as a washout.

Infrared differential interference contrast was used with an Olympus BX51WI upright microscope (Olympus Optical Co.) to guide electrodes to individual pyramidal neurons in the CA3. This region has been implicated as the driver of intra-hippocampal activity [15] and thus was the region of interest for this study. An Axopatch 200B amplifier (Axon Instruments) was used for extracellular field recordings. Glass electrodes with 3–5 MΩ resistance were made from borosilicate capillary tubing (World Precision Instruments) using a Narishige PP-830 vertical puller and were filled with standard ACSF. Data was collected with Clampex 10.2 software and analyzed off-line using MATLAB® and its publicly available HMM toolbox [16].

C. Hidden Markov Models

A hidden Markov model (HMM) is a nonparametric statistical approach of representing a Markov process wherein the observable output is dependent on the unobservable (hence, hidden) states. Thus, the underlying assumption governing HMMs is this output dependence on the hidden dynamical states of the system; the reader is referred to [10] for a detailed discussion. Briefly, a Q–state HMM, which is in state q_t at time t for $0 \le t \le T$, is characterized by $\boldsymbol{\lambda} = \{\boldsymbol{\pi}, \mathbf{A}, \mathbf{B}\}$. The initial state distribution is described by $\boldsymbol{\pi} = \{\pi_1, \ldots, \pi_Q\}$ where $\pi_j = \Pr(q_o = j)$ for $j = 1, \ldots, Q$ while the state transition matrix ($\mathbf{A} = \{a_{ij}\}$, where $a_{ij} = \Pr(q_t = j \mid q_{t-1} = i)$) describes the probability of transitioning from state q_i to state q_j for $1 \le i, j \le Q$. The observed emission matrix ($\mathbf{B} = \{b_j(\mathbf{O})\}$ where $\mathbf{O}$ is a K–dimensional feature vector for $j = 1, \ldots, Q$) describes the probability of observing the given output while being in the hidden state q_j. The entries of $\boldsymbol{\lambda}$ were initialized using the *k-means* algorithm [17] and were iteratively estimated with the well-established expectation-maximization algorithm [18]. Each feature vector ($\mathbf{O}$) was fit using an unsupervised mixture of Gaussians (MoG) where the multivariate Gaussian probability density function is defined as

$$\phi(\boldsymbol{O}, \mu_m, \Sigma_m) = \frac{1}{(2\pi)^{\frac{K}{2}} |\Sigma_m|^{\frac{1}{2}}} e^{-\left(\frac{1}{2}(\boldsymbol{O}-\mu_m)^T \Sigma_m^{-1}(\boldsymbol{O}-\mu_m)\right)} \quad (1)$$

This expression provides a measure of the distance between each feature vector $\mathbf{O}$ to the center of each Gaussian basis function m for $1 \le m \le M$ with mean μ_m and covariance Σ_m. As such, the emission probability is the sum of the weighted Gaussian probability densities over all the basis functions – namely,

$$b_j(\boldsymbol{O}) = \sum_{m=1}^{M} w_{jm} \phi(\boldsymbol{O}_m, \mu_{jm}, \Sigma_{jm}) \quad (2)$$

where w_{jk}, the mixture weights of the basis functions, sum to unity.

As the name implies, there are two stages involved in the expectation-maximization algorithm. The expectation stage evaluates the ability of the current model parameters to recreate the observed data. Of particular interest is the marginal posterior distribution for each state. This measure provides the time evolution of each state and is defined by

$$\gamma_t(i) = \frac{\alpha_t(i)\beta_t(i)}{\sum_i \alpha_t(i)\beta_t(i)} \quad (3)$$

where $\alpha_t(i)$ and $\beta_t(i)$ are the joint probabilities of observing all data up to time t at state i and the conditional probability of all future data from time $t+1$ onwards at state i, respectively. These probabilities are given by

$$\alpha_t(i) = \left[\sum_i \alpha_{t-1}(j) a_{ij}\right] b_i(\boldsymbol{O}_t) \quad (4a)$$

$$\beta_t(j) = \sum_i a_{ij} b_i(\boldsymbol{O}_{t+1}) \beta_{t+1}(i) \quad (4b)$$

During the maximization stage the model parameters are updated accordingly and the model is reevaluated. This process is terminated once the difference between log-likelihood (*LL*) values of successive iterations is below the predetermined threshold or when the number of iterations reaches the set maximum of 100. The *LL* provides a measure of the goodness-of-fit between the feature vector **O** and current model parameters and is defined by

$$LL = \log\left(P\left(\boldsymbol{O}\middle|a_{ij}, w_{jm}, \mu_{jm}, \Sigma_{jm}\right)\right) \quad (5)$$

Generally, *LL* convergence occurred within 30 iterations.

The training data utilized in this study consisted of a single recording of approximately 840 s containing two separate SLEs, each approximately 180 s. This provided the optimal training condition by exposing the system to the critical state transitions multiple times and thus better enabled the dynamics of the transitions to be captured. The main question of this study was to determine if there are more than two hidden states involved in the transitions between SLE and nonSLE behavior. Thus, the number of hidden states (Q) was varied; specifically, $Q = 2$ and $Q = 3$ were explored (henceforth referred to as the 2-state and 3-state model, respectively). Moreover, to avoid overfitting the number of basis functions was restricted to $M = 2$. Pilot studies also suggested that increasing this to $M = 3$ for either the 2-state or 3-state model yielded no significant difference.

III. Experimental Results

A. Feature Space Optimization

In constructing the feature space, an analytic signal was first computed by applying the Hilbert transform to the relevant frequency bands of the recorded field potentials; the reader is referred to [19] for a detailed discussion of the Hilbert transform. This allowed for the extraction of four distinct features – namely, amplitude, phase, and their respective first derivatives. The derivatives were obtained by applying a Gaussian filter to the desired signals. Specifically, this involved convolving each signal with the derivative of a Gaussian distribution. This is advantageous because it minimizes the underlying noise artifacts that are often amplified with traditional derivative methods. The specific frequency bands that were used were (in Hz): 1–4, 15–40, 40–58, 85–95, 95–130, and 130–155. Thus, the feature space incorporated both the delta (1–4 Hz), gamma (40–100 Hz), and super gamma (>100Hz) frequency bands which have been previously implicated in rodent epileptiform activity [12]. The division of the frequency bands was such that a single rhythm was captured in each band based on the fast Fourier transform.

B. Assessing Model Fit

In order to evaluate the ability of each model to capture the underlying dynamical states during the transition between SLE and nonSLE behavior, the marginal posterior distribution, as described by (3), of each state was computed for the respective models (Fig. 1). The 2-state model (Fig. 1(b)) was able to appropriately separate SLE and nonSLE behavior in a manner comparable to visual inspection of the local field potentials (Fig. 1(a)). These two regions in the local field potentials are noticeably different. Interestingly, the 3-state model (Fig. 1(c)) was able to distinguish between two separate nonSLE regions – namely, before and after the SLE itself. This is indicative of distinct dynamics in the region preceding the SLE compared to the one that follows, which visual inspection alone fails to recognize.

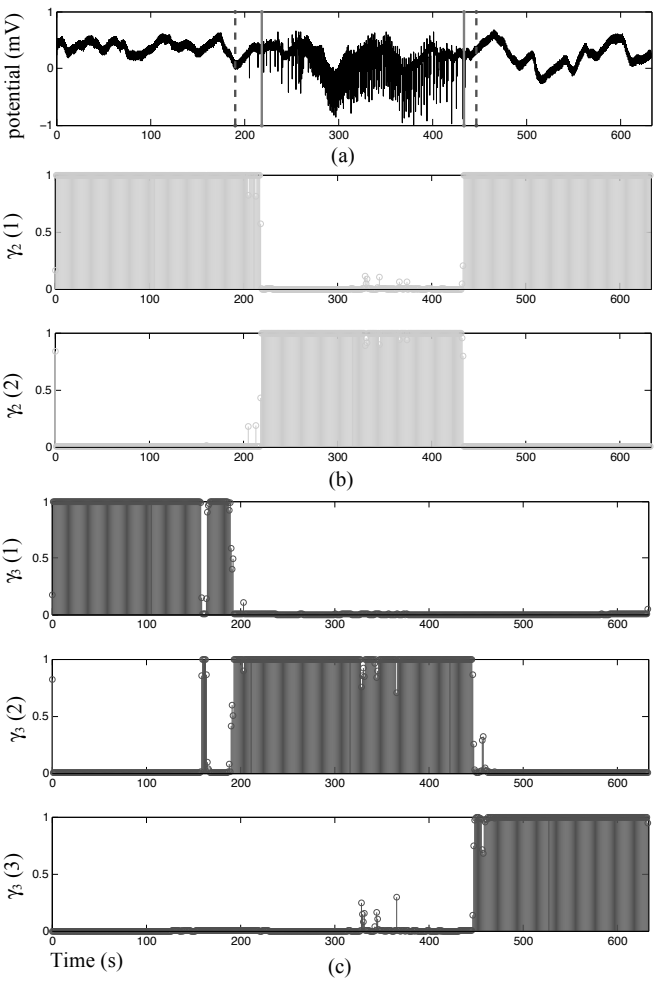

Fig. 1. SLE state classification via Markov modeling. (a) Extracellular field potentials were recorded from a pyramidal neuron in the CA3. A (b) 2-state and (c) 3-state HMM were explored. The marginal posterior distribution ($\gamma_Q(i)$) was computed for each state of the respective models where Q denotes the total number of model states and i denotes each individual state such that $1 \leq i \leq Q$. An expectation-maximization algorithm was used for training on a data set of ~840 s consisting of two separate SLEs, each ~180 s. The 2-state model was able to distinguish between SLE and nonSLE behavior. This, however, is not an improvement to visual inspection alone. Conversely, the 3-state model was able to capture two separate nonSLE states, one preceding the SLE (i.e., "preSLE") and the other immediately following (i.e., "postSLE"). Vertical solid and dashed lines in (a) indicate the SLE region identified by the 2- and 3-state model, respectively.

Moreover, the SLE state identified by both models (indicated by the vertical lines in Fig. 1(a)) differs in length. The 2-state model appears to transition into the SLE state later and transition out of this state earlier than its 3-state counterpart thus making the duration of the SLE state noticeably shorter. The 3-state model also appears to have an

initial spike into the SLE state around 160 s, which is not present in the 2-state model. Both models, however, do exhibit minor fluctuations in the center of the SLE state as well as around the transitions.

IV. Discussion

The advantage of using an HMM based system is in not requiring any *a priori* knowledge of when a state transition occurs. This eliminates any subjectivity in identifying particular states as well as provides a quantitative measure of how features differ between states.

Fig. 1(c) indicates that there is a significant difference between the dynamics of the region preceding an SLE compared to the one that follows afterward. Interestingly, the entrance into the SLE state occurs earlier in the 3-state model (Δt = ~30 s). There is also an earlier spike in the SLE state around 160 s. These observations suggest that there may be a region between the preSLE and SLE states that has not been fully captured in the model. Similarly, the transition out of the SLE state occurs earlier in the 2-state model (Δt = ~15 s). This again suggests that there may be an additional state between the SLE and postSLE states that was integrated into the existing model states. Thus, a natural extension of this work would be to incorporate more than three states into the model.

The minor fluctuations observed near the center of the SLE state in both models and around the state transitions suggest that feature space still has room for improvement. Perhaps by incorporating higher-order derivatives of the amplitude and phase these fluctuations would be eliminated.

One should note that if there were redundancy between states then there would be significant overlap between these states. In other words, the system would simultaneously be in two states because the feature space did not capture any significant differences to distinguish between them. Moreover, if the feature space were not able to capture the true dynamics of the time series then the confusion of the system would also be manifest in multiple transitions into and out of the states within a relatively short time period. Similarly, if too many states were incorporated into the system then the superfluous states would exhibit a near zero marginal posterior distribution for the entire time series. Thus the ability of the 3-state model to capture an additional distinct state compared to its 2-state counterpart is not merely the result of increasing the number of hidden states.

V. Conclusion

This study investigated the potential of 2- and 3-state Markov models, specifically first-order HMMs, to capture the dynamics involved in the transitions between SLE and nonSLE behavior. Although the 2-state model was able to distinguish between SLE and nonSLE behavior, it did not improve classification when compared to visual inspection alone. The 3-state model, however, was able to capture a distinct preSLE and postSLE state in addition to the SLE state common to both models. The variability in the length of the SLE state suggests that there may be additional states that were not distinctly captured but were rather integrated into the existing model states. This further suggests that incorporating more hidden states may yield a more precise classification of the states involved in the transitions between SLE and nonSLE behavior.

The practicality of an HMM based system lies in the unsupervised manner by which the states are identified. Eliminating the dependence on *a priori* knowledge of the states involved is a great benefit when it comes to prediction of seizure onset.

References

[1] D. Durand and M. Bikson, "Suppression and control of epileptiform activity by electrical stimulation: a review," *Proc. IEEE*, vol. 89, pp. 1065–1082, July 2001.

[2] P. Uhlhaas and W. Singer, "Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology" *Neuron*, vol. 52, pp.155-168, 2006.

[3] K. Lehnertz and C. Elger, "Can epileptic seizures be predicted? Evidence from nonlinear time series analysis of brain electrical activity," *Phys Rev Letters*, vol. 80, pp. 5019–5022, June 1998.

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

[5] V. Navarro *et al.*, "Seizure anticipation in human neocortical partial epilepsy," *Brain*, vol. 125, pp. 640–655, 2002.

[6] B. Litt *et al.*, "Epileptic seizures may begin hours in advance of clinical onset: a report of five patients," *Neuron*, vol. 30, pp. 51–64, 2001.

[7] L. Iasemidis, P. Pardalos, J. Sackellares, and D. Shiau, "Quadratic binary programming and dynamical sysem approach to determine the predictability of epileptic seizures," *J Comb Optimization*, vol. 5, pp. 9–26, 2001.

[8] F. Mormann, T. Kreuz, R. Andrzejak, P. David, K. Lehnertz, and C. Elger, "Epileptic seizures are preseded by a decrease in synchronization," *Epilepsy Res*, vol. 53, pp. 173–185, 2003.

[9] K. Lehnertz and B. Litt, "The first international collaborative workshop on seizure prediction: summary and data description," *Clin Neurophysiol*, vol. 116, pp. 493–505, 2005.

[10] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. of the IEEE*, vol. 77, pp. 257–286, 1989.

[11] S. Wong, A. Gardner, A. Krieger, and B. Litt, "A stochastic framework for evaluating seizure prediction algorithms using hidden Markov models," *J Neurophysiol*, vol. 97, pp. 2525–2532, 2007.

[12] A. Chiu, E. Kang, M. Derchansky, P. Carlen, and B. Bardakjian, "Online prediction of onsets of seizure-like events in hippocampal neural networks using wavelet artificial neural networks," *Ann Biomedical Eng*, vol. 34, pp. 282–294, 2006.

[13] C. Wu *et al.*, "Spontaneous rhythmic field potentials of isolated mouse hippocampal-subicular-entorhinal cortices in vitro," *J Physiol*, vol. 576, pp. 457–476, 2006.

[14] M. Derchansky *et al.*, "Model of frequent, recurrent, and spontaneous seizures in the intact mouse hippocampus," *Hippocampus*, vol. 14, pp. 935–947, 2004.

[15] M. Derchansky *et al.*, "Bidirectional multisite seizure propagation in the intact isolated hippocampus: the multifocality of the seizure 'focus'," *Neurobiol Dis*, vol. 23, pp. 312–328, 2006.

[16] K. Murphy. "Hidden Markov Model (HMM) Toolbox for Matlab," Internet: http://www.cs.ubc.ca/~murphyk/Software/HMM/hmm.html, June 8, 2000 [May 1, 2009].

[17] G. Schalk *et al.*, "Two-dimensional movement control using electrocorticographic signals in humans," *J of Neural Eng*, vol. 5, no. 1, pp. 75–84, 2008.

[18] L. Hochberg *et al.*, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," *Nature*, vol. 442, pp. 164–171, 2006.

[19] M. Johansson, "The Hilbert transform," M.A. thesis, Vaxjo University, Sweden, 1999.