

# Decoding Epileptogenesis in a Reduced State Space

François G. Meyer<sup>1</sup>, Alexander M. Benison<sup>2</sup>, Zachariah Smith<sup>2</sup>, and Daniel S. Barth<sup>2</sup>

<sup>1</sup>Electrical Engineering & Applied Mathematics, <sup>2</sup>Psychology & Neuroscience  
University of Colorado at Boulder, Boulder, CO 80309

**Abstract**—We describe here the recent results of a multidisciplinary effort to design a biomarker that can actively and continuously decode the progressive changes in neuronal organization leading to epilepsy, a process known as epileptogenesis. Using an animal model of acquired epilepsy, we chronically record hippocampal evoked potentials elicited by an auditory stimulus. Using a set of reduced coordinates, our algorithm can identify universal smooth low-dimensional configurations of the auditory evoked potentials that correspond to distinct stages of epileptogenesis. We use a hidden Markov model to learn the dynamics of the evoked potential, as it evolves along these smooth low-dimensional subsets. We provide experimental evidence that the biomarker is able to exploit subtle changes in the evoked potential to reliably decode the stage of epileptogenesis and predict whether an animal will eventually recover from the injury, or develop spontaneous seizures.

## I. INTRODUCTION

### A. The Challenge: Decoding Epileptogenesis

Epilepsy is a neurological disease that is characterized by the occurrence of several unprovoked seizures. Despite various causes of epilepsy and varying degrees of disease severity in the human population, hippocampal sclerosis is the most consistent neuropathological feature of temporal lobe epilepsy [1]. Animal models have been developed to study the neuronal changes underlying the clinical manifestations of epilepsy (chronic-spontaneous seizures). One popular model relies on controlled administration of a convulsant drug (e.g., pilocarpine) to induce status-epilepticus, a life-threatening condition in humans. This condition is followed by a latent seizure-free period of weeks to months, where progressive neuronal damage and network reorganization eventually lead to the development of spontaneous seizures. Most of our understanding of the progression of epilepsy, or *epileptogenesis*, is derived from such animal models. It is therefore critical to define a *biomarker* to monitor epileptogenesis. An accurate biomarker would be invaluable for the design of novel anti-epileptogenic drugs, and could eventually be translated into a diagnostic tool for humans.

### B. From Passive Recording to Active Sensing

Current efforts toward the development of a reliable and predictive biomarker of epileptogenesis fall in three main classes: molecular and cellular biomarkers [2], imaging biomarkers [3], and electrophysiological biomarkers [4]. The present work focuses on the last class of biomarkers that rely on recordings of the electrical activity associated with

neuronal firing. The development of electrophysiological biomarkers has focused on the analysis of both epileptiform spikes [5], and high frequency oscillations [6]. Interictal spikes are sharp electrical impulses. Their morphology has been shown to be correlated with the progression of epileptogenesis [5]. Recent experiments combined with computational models [5] suggest that epileptogenesis systematically modifies the morphology of the spikes in a predictable manner. High frequency oscillations detected in local field potentials are created by synchronously bursting pyramidal cells, and have been observed in healthy patients when the frequency of the ripples is lower than 250 Hz. Conversely, frequencies in the range of 250Hz to 800Hz are considered to be pathological ripples and are reliable biomarkers for the epileptogenic zone [7]. While recordings of spontaneous neuronal spiking can be indicative of neuronal excitability, and therefore correlate with the propensity for seizures, such methods cannot actively probe the hippocampal circuit in living animals during epileptogenesis. Consequently, one could argue that the *passive* electrophysiological recordings may not provide enough information to observe early changes in neuronal excitability associated with epileptogenesis.

The present work addresses this limitation and proposes for the first time a “computational biomarker” that relies on actively probing the excitability of the hippocampus using an auditory stimulus. We advocate an *active* approach whereby we probe the excitability of the hippocampus, a brain region known to be epileptogenic. Because the hippocampus also receives several sensory inputs (through the entorhinal cortex), we propose to record in the hippocampus the evoked potential elicited by an auditory stimulus. Since epilepsy does not modify the primary auditory cortex, any alterations in the evoked potential should be indicative of neuronal changes in the hippocampus underlying epileptogenesis. To quantify the property of this biomarker, we use a pilocarpine animal model of temporal lobe epilepsy, and chronically record hippocampal auditory evoked potentials during epileptogenesis.

We design a decoding algorithm to demonstrate that changes in the morphology of the hippocampal auditory evoked potential have universal predictive value and can be used to accurately quantify the progression of epilepsy. The authors are not aware of any work that uses machine learning methods to construct an *active* biomarker of epileptogenesis (but see [5] for approaches based on passive recordings).

condition	baseline	silent	latent	chronic
duration	1 week	72 hours	weeks – months	2–6 months
evoked potential	$h_r^{(t)}$	$h_r^{(t)}$	$h_r^{(t)}$	$h_r^{(t)}$
intervention	pilocarpine	paraldehyde		
event		status epilepticus (SE)		first spontaneous seizure

Figure 1. Timeline and nomenclature of the different conditions.

## II. THE ANIMAL MODEL

All procedures were performed in accordance with the University of Colorado Institutional Animal Care and Use Committee guidelines for the humane use of laboratory rats in biological research. Twenty-four male Sprague-Dawley rats (200–250 gm) were implanted with a hippocampal wire electrode, a ground screw, and a reference screw. The 24 rats were divided into two groups: a group of 17 rats that received lithium-pilocarpine, and a control group of 7 rats. The control group was composed of 2 rats that received all drug injections associated with the lithium-pilocarpine model except for pilocarpine, which was substituted with saline; and 5 rats that received no drugs. After full recovery from the electrode implantation (2 weeks) and at least one additional week of chronic recording of baseline video/EEG, 17 rats were given an injection of lithium chloride followed by an injection of pilocarpine hydrochloride 24 hours later (see Fig. 1). After one hour of status epilepticus, the animals were administered a dose of paraldehyde to terminate convulsions. Throughout the experiment, every  $\Delta t = 30$  min, an auditory stimulus, composed of a sequence of 120 square-wave clicks (0.1 ms duration, 2 sec ISI, 45dB SPL), was played in a top-mounted speaker. The 300 ms hippocampal responses to each click were filtered and sampled at 10 kHz, and the average of the 120 responses was computed. In the remainder of the paper, we denote by  $h^{(k)}$  the average evoked potential, measured at time  $k\Delta t$ . To further simplify the exposition,  $h^{(k)}$  is simply referred to as the *evoked potential measured at time  $k$* .

Figure 1 provides a detailed timeline, along with the nomenclature of the different periods associated with the progress and eventual onset of epilepsy. The period before the injection of pilocarpine is called *baseline*. Conversely, the period following the first spontaneous seizure is called *chronic*. We further define the *silent* period to be the 72 hour period of recovery immediately following the termination of status epilepticus, and the *latent* period to be the remaining period leading to the eventual onset of the first spontaneous seizure.

## III. OVERVIEW OF THE DECODING FRAMEWORK

We give here a brief overview of the decoding approach. Given an animal  $r = 1, \dots, 24$ , we consider the sequence of evoked potentials  $h_r^{(0)}, h_r^{(1)}, \dots$ . The first stage involves the construction of a denoised representation of  $h_r^{(k)}$  (see ① in Fig. 2). We decompose  $h_r^{(k)}(t)$  using a discrete stationary wavelet transform and retain a vector of  $s$  (carefully chosen)

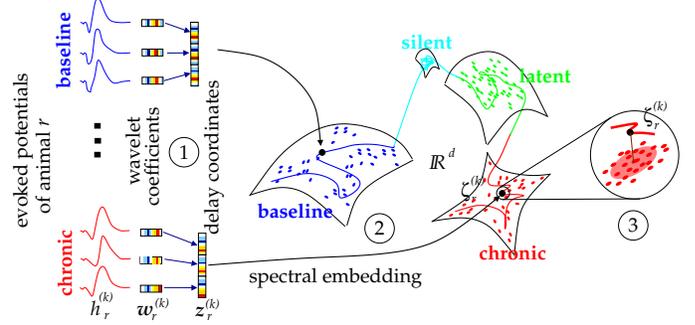


Figure 2. Overview of the decoding algorithm. ①: a vector of wavelet coefficients,  $w_r^{(k)}$ , is computed from the evoked potential  $h_r^{(k)}$ . A vector of time-delay wavelets coordinates,  $z_r^{(k)}$ , is formed by concatenating  $\tau$  consecutive  $w_r^{(k)}$ . ②: spectral embedding maps  $w_r^{(k)}$  to  $\zeta_r^{(k)}$ . ③: the distance between  $\zeta_r^{(k)}$  and the low-dimensional structure formed by each condition is computed.

wavelet and scaling coefficients,  $w_r^{(k)}$ , (see ① in Fig. 2). The second stage involves characterizing the association between the condition of the disease (baseline, silent, latent, or chronic) and the vector of wavelet coefficients  $w_r^{(k)}$ . We tackle this question by lifting  $w_r^{(k)}$  into  $\mathbb{R}^{\tau \times s}$  using time-delay embedding: we concatenate the consecutive vectors  $w_r^{(k)}, \dots, w_r^{(k+\tau-1)}$  to form a  $\tau \times s$  vector,  $z_r^{(k)}$ , of *time-delay wavelet coordinates* (see ① in Fig. 2). Low-dimensional structures, which uniquely characterize the stage of epileptogenesis, emerge in the high-dimensional space. We use spectral embedding to parameterize these low-dimensional structures, and map  $z_r^{(k)}$  to  $\zeta_r^{(k)}$  (see ② in Fig. 2). The first decoding stage involves geometrically computing the likelihood that a given vector  $\zeta_r^{(k)}$  corresponds to one of the four conditions. To this end, we quantify the distance of  $\zeta_r^{(k)}$  to the low-dimensional cluster formed by each condition (see ③ in Fig. 2). In the final decoding stage, we use a hidden Markov model to capture the intrinsic dynamics of epileptogenesis.

To alleviate the notation, and unless we explicitly compare or combine several animals, we dispense with the subscript  $r$ , denoting the dependency on rat  $r$ , when we discuss the analysis of the evoked potentials for a fixed animal.

## IV. DENOISING THE INPUT DATA

To compensate for the alteration of the tissue impedance around the electrodes, the evoked potentials  $h_r^{(k)}, k = 0, 1, \dots$  were normalized for each animal. For a given animal  $r$ , all evoked potentials for that animal were rescaled, in such a way that the average energy computed during the baseline condition,  $\langle h_r^2 \rangle$ , was one. In order to use the noisy evoked potentials to predict the state of epileptogenesis, we extract a denoised representation of  $h^{(k)}$ . We use a discrete stationary wavelet transform (CDF 9-7) to compute a redundant representation of  $h^{(k)}$ . We used a “leave-one-animal-out” cross-validation to determine the time intervals and the scales of the wavelet coefficients that resulted in the

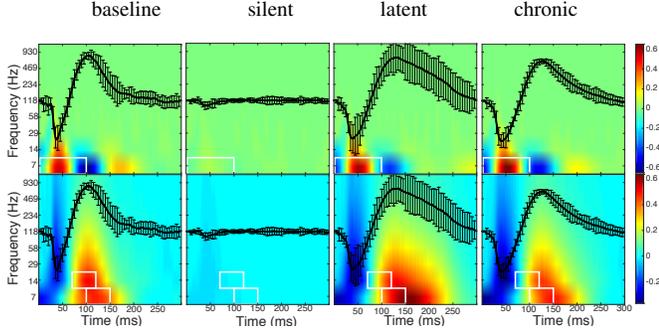


Figure 3. Top: average wavelet coefficients with the average evoked potential (across all animals). Bottom: average approximation coefficients. Only the scales  $j = 3$  (top) to 10 (bottom) are displayed. White rectangles delineate the time-frequency blocks used to construct  $\mathbf{w}^{(k)}$ .

the most accurate prediction of the condition of the animal that was not part of the training data. For most of the animals, the wavelet coefficients in the time-frequency region  $[0, 100] \text{ ms} \times [5, 10] \text{ Hz}$ , and the approximation coefficients from the time-frequency regions  $[70, 120] \text{ ms} \times [10, 20] \text{ Hz}$  and  $[100, 150] \text{ ms} \times [5, 10] \text{ Hz}$  (see Fig. 3) resulted in either the optimal, or very close to optimal, decoding performances. To simplify the decoding, we therefore kept the same time-frequency regions for all animals. In summary, each evoked potential was represented with a feature vector  $\mathbf{w}^{(k)}$  of 2,000 entries composed of 1,000 wavelet coefficients and 1,000 approximation coefficients. This representation is consistent with reports of disruption in the  $\theta$  rhythm (4-12 Hz) during the latent period preceding the onset of epilepsy [8].

## V. UNIVERSAL CONFIGURATIONS OF EPILEPTOGENESIS

We now describe the construction of universal (stable across all animals) low-dimensional smooth sets that are formed by all the hippocampal auditory evoked potentials collected during the same stage of epileptogenesis. These sets are created by lifting the wavelet coefficients  $\mathbf{w}^{(k)}$  of each  $h^{(k)}$  into high-dimension. This lifting effectively creates smooth low-dimensional coherent structures that can then be parameterized with a drastically smaller number of coordinates.

### A. Time-Delay Embedding of the Wavelet Coordinates

Given the time series  $\{\mathbf{w}^{(k)}, k = 0, 1, \dots\}$  for a given animal, we analyze the dynamics of this time series by considering the time-delay wavelet coordinates formed by concatenating  $\tau$  consecutive vectors of wavelet coefficients,

$$\mathbf{z}^{(k)} = [\mathbf{w}^{(k)} \quad \mathbf{w}^{(k+1)} \quad \dots \quad \mathbf{w}^{(k+\tau-1)}]. \quad (1)$$

We characterize the dynamical changes in the evoked potentials by studying the geometric structures formed by the trajectory of  $\mathbf{z}^{(k)}$  in  $\mathbb{R}^{\tau \times s}$ , as  $k$  evolves. In practice, the number of time-delay vectors,  $\tau$ , is determined using cross-validation. We confirmed that  $\tau = 1$ , to wit no time-delay

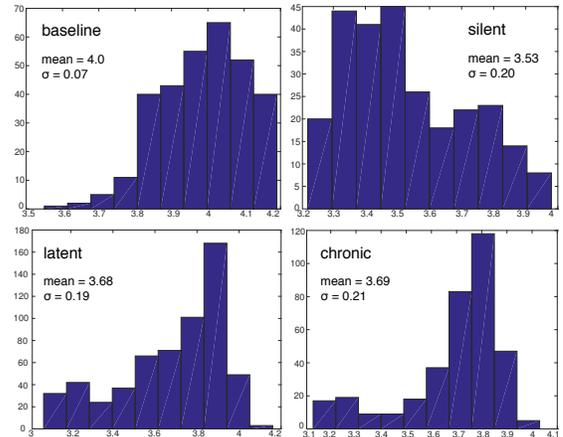


Figure 4. Distribution of the local dimension of the state space according to the condition. Top: baseline and silent; bottom: latent and chronic.

embedding, yielded poor prediction of the animal’s condition. While  $\tau$  was optimized for each animal, the average of the optimal values was  $\tau = 12$ , which corresponds to six hours.

To learn the geometry of the set formed by the different trajectories  $\mathbf{z}_r^{(k)}, k = 0, 1, \dots$ , we consider the union – over all the animals in the training set – of the vectors  $\mathbf{z}_r^{(k)}$ , and define the set

$$\mathbf{Z} = \bigcup_{r \in \text{training animals}} \{\mathbf{z}_r^{(k)}, k = 0, 1, \dots\}. \quad (2)$$

### B. What is the local dimension of the state-space?

Because we have only a limited number of training data, we need to drastically reduce the dimensionality of the state space (on average  $\tau \times s = 24,000$ ) in order to reliably decode the condition of each animal. We used a kernel-based version of the correlation dimension [9] to compute the local dimension of the point-cloud formed by the state space  $\mathbf{Z}$ . We discovered that the points  $\mathbf{z}_r^{(k)}$  were organized along low dimensional subsets that correspond to well defined stages of epileptogenesis and exhibit different local dimensions. Figure 4 displays histograms (as well as the mean, and standard deviation) of the estimates of the local dimensions for the four different conditions. These experiments suggest that we should be able to represent the state-space associated with the dynamic of the disease using  $d = 5$  dimensions. However, our results (not shown) indicate that the traditional linear approach, PCA, provides a very poor parameterization of the set  $\mathbf{Z}$ .

### C. Nonlinear Parameterization of the State Space

In order to identify the separate regions of  $\mathbb{R}^{\tau \times s}$  that correspond to the four conditions, we seek a smooth low-dimensional parameterization of  $\mathbf{Z}$ . A nonlinear approach spectral embedding [10], yields a low-dimensional parameterization of  $\mathbf{Z}$  that naturally clusters the different conditions

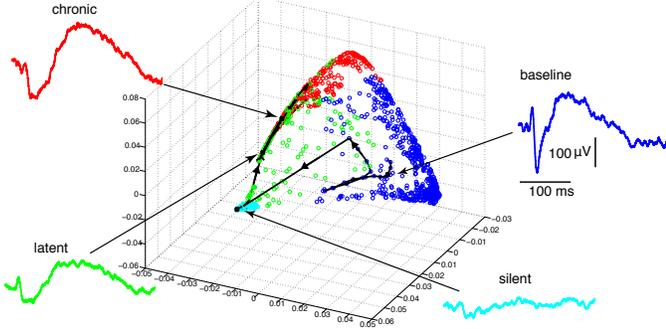


Figure 5. The training set of evoked potentials,  $\mathbf{Z}$ , displayed using the reduced coordinates  $\zeta^{(k)}$ . Each condition, baseline (blue), silent (cyan), latent (green), and chronic (red), forms a coherent sub-cloud.

into coherent disconnected smooth subsets. Briefly, we define a similarity matrix  $\mathcal{K}$  that quantifies how any two evoked potentials  $h^{(k)}$  and  $h^{(l)}$  extracted from the same or from different conditions, and from the same or from a different animal, at the respective times  $k$  and  $l$  co-vary,

$$\mathcal{K}(k, l) = \exp\left(-\|z^{(k)} - z^{(l)}\|^2 / \sigma^2\right). \quad (3)$$

The scaling constant  $\sigma$  was chosen to be a multiple of the median distance  $\|z^{(k)} - z^{(l)}\|$ ,  $k, l = 0, 1, \dots$ . We used the  $d$  eigenvectors of  $\mathcal{K}$  that optimally separated the dataset  $\mathbf{Z}$  into four clusters related to the four conditions. The  $i^{\text{th}}$  eigenvector of  $\mathcal{K}$  provided the  $i^{\text{th}}$  reduced coordinate,  $\zeta^{(k)}(i)$ , of  $z^{(k)}$ ,

$$\zeta^{(k)} = [\zeta^{(k)}(1) \quad \dots \quad \zeta^{(k)}(d)]^T. \quad (4)$$

As explained in section V-B, we use  $d = 5$ . Figure 5 displays the nonlinear parameterization of  $\mathbf{Z}$ . Each dot represents an evoked potential  $h_r^{(k)}$  measured at a time  $k$  from an animal  $r$ , and parameterized by the reduced coordinates  $\zeta_r^{(k)}$ . The color indicates the condition during which  $h_r^{(k)}$  was recorded. When collapsed across animals and time of recordings, four distinct clusters can be visually discerned, which can be interpreted in terms of the corresponding conditions. The evoked potentials in the silent condition are grouped together around a low-dimensional structure. The latent condition displays the largest spread: some evoked potentials are morphologically close to silent evoked potentials; while others are close to chronic evoked potentials. These visual observations are consistent with the computation of the local dimensionality described in section V-B. To further help with the interpretation of this nonlinear representation of  $\mathbf{Z}$ , we trace the trajectory of  $\zeta_r^{(k)}$ ,  $k = 0, 1, \dots$  for a single rat (H37) over several weeks of recordings (see thick black line in Fig. 5). The animal first spends several days in the baseline cluster (blue), then briskly traverses the entire space to reach the silent cluster (cyan) after status epilepticus has been induced. Eventually, as the animal recovers, it joins the latent condition (green), and advances toward the chronic cluster (red). Four representative evoked potentials

are shown along this trajectory, confirming changes in the morphology of  $h^{(k)}$  during epileptogenesis.

#### D. Learning Epileptogenesis from the Reduced Coordinates

The natural division of  $\mathbf{Z}$  into coherent subsets, which correspond to well defined stages of epileptogenesis, suggests that a purely geometric algorithm could be used to quantify the development of epilepsy. Indeed, given an unclassified evoked potential,  $\zeta^{(k)}$ , the distance from  $\zeta^{(k)}$  to each of the four clusters (baseline, silent, latent, and chronic) provides an estimate of the likelihood of being in the corresponding condition at time  $k$ .

We denote by  $\mathbf{Z}^c$ ,  $c = 0, 1, 2, 3$ , the four clusters formed by the evoked potentials in  $\mathbf{Z}$  that were collected during the baseline, silent, latent, and chronic conditions, respectively. We found that using a mixture of probabilistic principal components [11] to represent each  $\mathbf{Z}^c$  resulted in a remarkably low dimensional parameterization of that condition. Indeed, the mixture model is able to describe with a small number of components the geometric structure created by each  $\mathbf{Z}^c$ . Furthermore, unlike other models, the mixture model does not require the stringent assumption that  $\mathbf{Z}^c$  be a smooth sub-manifold.

Formally, for each  $c = 0, 1, 2, 3$ , we use a different mixture of probabilistic principal components to parameterize  $\mathbf{Z}^c$ . To reduce the number of indices, we do not make the dependency on  $c$  explicit in the following discussion. For a given cluster  $\mathbf{Z}^c$ , the mixture is composed of  $M$  local Gaussian densities, each describing the local principal directions with a  $d \times q$  matrix,  $U_i$ , around a collection of centers  $\bar{\zeta}_i$ ,  $i = 1, \dots, M$ . The local spread of the  $\zeta_r^{(k)}$  around  $\bar{\zeta}_i$  is given by the covariance matrix  $U_i U_i^T + \varepsilon_i^2 \mathbf{I}$ . The mixture model associated with the cluster  $\mathbf{Z}^c$  can be written as

$$\zeta_r^{(k)} \sim \sum_{i=1}^M \pi_i \mathcal{N}(U_i \mathbf{n}_i + \bar{\zeta}_i, \varepsilon_i^2 \mathbf{I}), \quad (5)$$

where the  $\mathbf{n}_i$ ,  $i = 1, \dots, M$  are  $q$ -dimensional latent Gaussian variables, and the mixture proportions  $\pi_i$  are non negative and sum to one. The parameters of the model are estimated using the EM algorithm [11]. For a given animal  $r$ , we compute the posterior probability that the model (5) associated with condition  $c$  would generate the reduced coordinates  $\zeta_r^{(k)}$ ,

$$\gamma_r^{(k)}(c) = \text{Prob}(c | \zeta_r^{(k)}). \quad (6)$$

We use the posterior probability of class membership  $\gamma_r^{(k)}(c)$  as a raw measurement of the progression of epileptogenesis. Because the instantaneous posterior probability  $\gamma_r^{(k)}(c)$  is oblivious to the past trajectory  $\{\zeta_r^{(k-1)}, \zeta_r^{(k-2)}, \dots\}$ , it can be noisy. To address this limitation, and enforce the temporal coherence that defines the progression of the disease, we introduce a Markovian description of the dynamics of the changes in the evoked potential  $h^{(k)}$  as a function of time.

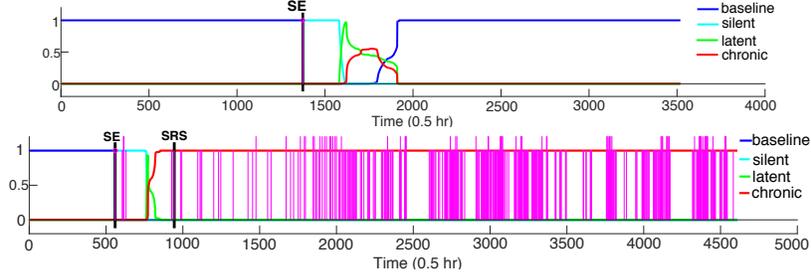


Figure 6. Biomarker probabilities  $\pi^{(k)}(c)$  of being in condition  $c$  for a *non-epileptic* (top) and an *epileptic* (bottom) animal. Seizure count (clipped at 1) is displayed in magenta. The black bar SE corresponds to *Status Epilepticus*; the second bar SRS corresponds to the first Spontaneous Recurring Seizure.

## VI. THE DYNAMICS OF EPILEPTOGENESIS

We use a hidden Markov Model to capture both the dynamics of epileptogenesis, and the resulting changes in the shape of  $h_r^{(k)}$  triggered by this hidden condition. We define  $x_r^{(k)}$  to be a discrete random variable that encodes the state of animal  $r$  at time  $k$ ,  $x_r^{(k)} \in \{0, 1, 2, 3\}$ , where the states 0, 1, 2, and 3 encode the baseline, silent, latent, and chronic conditions, respectively. We further assume that  $x_r^{(k)}$  is a Markov chain with probability transition matrix  $P_{i,j} = \text{Prob}(x_r^{(k+1)} = i | x_r^{(k)} = j)$ . We define  $y_r^{(k)}$  to be a discrete random variable, which takes its values in  $\{0, 1, 2, 3\}$ , and encodes the most likely condition of animal  $r$  at time  $k$ , given the measurement  $h_r^{(k)}$ ,

$$y_r^{(k)} = \underset{c=0,1,2,3}{\text{argmax}} \gamma_r^{(k)}(c), \quad (7)$$

where  $\gamma_r^{(k)}(c)$  is the posterior probability that condition  $c$  generates the reduced coordinates  $\zeta_r^{(k)}$ , and is given by (6). We denote by  $Q_{i,j} = \text{Prob}(y_r^{(k)} = i | x_r^{(k)} = j)$  the measurement probability distribution matrix. The probability transition matrix  $P$  and the measurement probability matrix  $Q$  were estimated from the training data. The trained hidden Markov model was then used in a “reverse mode” [12] to estimate the posterior probability of animal  $r$  to be in condition  $c$ ,

$$\pi_r^{(k)}(c) = \text{Prob}(x_r^{(k)} = c | y_r^{(k)}), \quad c = 0, 1, 2, 3. \quad (8)$$

Because the hidden Markov model was trained on the epileptic animals, we regularized the probability transition matrix  $P$  to allow a transition, which was never observed in the training data, from chronic to baseline with a very small probability.

Figure 6 displays the four probabilities  $\pi^{(k)}$  of being in baseline, silent, latent, and chronic state for a *non-epileptic* animal (top), and an *epileptic* animal (bottom). We plot the seizure count (clipped at 1) in magenta for the epileptic animal. The biomarker clearly identifies the baseline period (blue), and is able to detect changes in the evoked potential  $h^{(k)}$  that are indicative of neuronal alterations leading to epilepsy, before the first spontaneous recurring

seizure (SRS), and shifts from latent (green) to chronic (red) before SRS. The probability of being in the chronic state,  $\pi^{(k)}(3)$ , remains at one after SRS, indicating that the animal is irreversibly in the chronic state, as confirmed by the uninterrupted sequence of seizures.

## VII. RESULTS

We evaluated our approach using a “leave-one-animal-out” cross-validation: the algorithm was trained on 11 epileptic rats, and evaluated with one animal that was not part of the training data. For each test animal  $r$ , the decoding algorithm computed from  $h_r^{(k)}$  the posterior probability  $\pi_r^{(k)}(c)$ , given by (8). Figure 7 displays the temporal profile of the median probability  $\overline{\pi^{(k)}}(c)$  computed over all the epileptic animals (left,  $N=12$ ), and the non epileptic animals (right,  $N=5$ ). The computation of the median probability  $\overline{\pi^{(k)}}(c)$  required care: because each condition varied in length across the different animals, we could not use the beginning of each recording as the origin, and bluntly average across the animals.

As an alternative, we used two physiologically meaningful temporal landmarks to align the time series  $\{\pi_r^{(k)}(c), k = 0, 1, \dots\}$  of the different animals before averaging. We used the time of the insult (status epilepticus) to study the transition from baseline to silent, and we also used the onset of the first spontaneous seizure to study the transition from latent to chronic. As shown in Figs. 7-left and Fig. 8, this led to two different median probability profiles for each condition. In each figure, the origin corresponds to the physiological landmark used to realign the time series. To help visually compare the two figures, we added the latent period to Fig. 7-left, and the silent period to Fig. 8. We note that the median probability for the silent and latent periods are different in both figures: indeed, the large variation in the onset of the first spontaneous seizure leads to different realignment results.

**Epileptic animals.** The baseline and chronic periods were well defined for the 12 epileptic animals (see Fig. 7-left and 8), and corresponded to the period before status epilepticus, and the period after the first spontaneous seizure, respectively. The *sensitivity* was found to be 1 for both the baseline

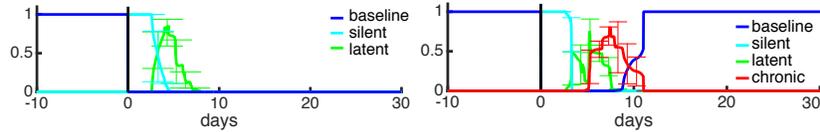


Figure 7. Median probability  $\overline{\pi^{(k)}}(c)$  of being in condition  $c$  for the epileptic animals (left), and the non epileptic animals (right). Day 0 (black bar) corresponds to status epilepticus across all animals.

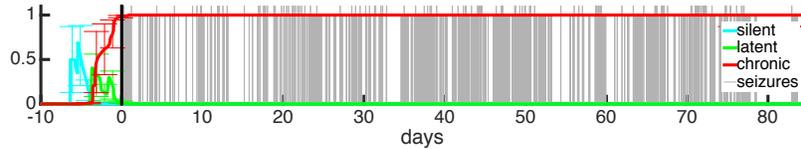


Figure 8. Median probability  $\overline{\pi^{(k)}}(c)$  of being in condition  $c$  for the epileptic animals. Day 0 (black bar) corresponds to the first spontaneous seizure across all animals. Seizure count (clipped at 1) is displayed in light grey.

and chronic periods. The *specificity* was 1 for the baseline period, and 0.59 for the chronic period, indicating that the probability  $\pi^{(k)}(3)$  rises to 1 before the first spontaneous seizure (see Fig. 8). Indeed, the biomarker is able to detect changes in the evoked potential  $h^{(k)}$  that are indicative of neuronal alterations leading to epilepsy, before the first spontaneous seizure.

**Non-epileptic animals.** Five of the 17 animals, which received the convulsant (pilocarpine), never developed spontaneous seizures. After the injections of paraldehyde, and the subsequent recovery from status epilepticus, these animals returned to a baseline condition. These animals were interesting, as they allowed us to address a common limitation of machine learning based biomarkers with regard to specificity of the biomarker to the general population. The biomarker was able to predict the return to baseline condition (see Fig. 7-right) in spite of the fact that the training dataset never included non-epileptic rats. This result is significant since it indicates that the low-dimensional clusters formed by embedding the time-delay wavelet coordinates have a universal aspect that can be reliably used to decode the status of the disease.

**Control animals.** The 5 rats that received no drugs were classified as being in the baseline condition with probability 1 at all time (results not shown). The 2 rats that received all drug injections except for pilocarpine, were also classified as being in the baseline conditions at all time, except for a brief suppression of amplitude, classified as silent (results not shown).

#### ACKNOWLEDGEMENT

FGM was supported in part by NSF DMS 1407340.

#### REFERENCES

- [1] A. Pitkänen and K. Lukasiuk, “Mechanisms of epileptogenesis and potential treatment targets,” *The Lancet Neurology*, vol. 10, no. 2, pp. 173–186, 2011.
- [2] K. Lukasiuk and A. J. Becker, “Molecular biomarkers of epileptogenesis,” *Neurotherapeutics*, vol. 11, no. 2, pp. 319–323, 2014.
- [3] S. R. Shultz, T. J. O’Brien, M. Stefanidou, and R. I. Kuzniecky, “Neuroimaging the epileptogenic process,” *Neurotherapeutics*, vol. 11, no. 2, pp. 347–357, 2014.
- [4] R. Staba, M. Stead, and G. Worrell, “Electrophysiological biomarkers of epilepsy,” *Neurotherapeutics*, vol. 11, no. 2, pp. 334–346, 2014.
- [5] C. Huneau, P. Benquet, G. Dieuset, A. Biraben, B. Martin, and F. Wendling, “Shape features of epileptic spikes are a marker of epileptogenesis in mice,” *Epilepsia*, vol. 54, no. 12, pp. 2219–2227, 2013.
- [6] J. Engel and F. da Silva, “High-frequency oscillations—where we are and where we need to go,” *Prog. Neurobiol.*, vol. 98, no. 3, pp. 316–318, 2012.
- [7] L. M. de la Prida, R. J. Staba, and J. A. Dian, “Conundrums of high-frequency oscillations (80–800 hz) in the epileptic brain,” *Journal of Clinical Neurophysiology*, vol. 32, no. 3, pp. 207–219, 2015.
- [8] L. Chauvière, N. Raftafi, C. Thinus-Blanc, F. Bartolomei, M. Esclapez, and C. Bernard, “Early deficits in spatial memory and theta rhythm in experimental temporal lobe epilepsy,” *J. Neurosci.*, vol. 29, no. 17, pp. 5402–5410, 2009.
- [9] M. Hein and J.-Y. Audibert, “Intrinsic dimensionality estimation of submanifolds in  $R^d$ ,” in *Proceedings of the 22nd international conference on Machine learning*. ACM, 2005, pp. 289–296.
- [10] M. Ghil, M. Allen, M. Dettinger, K. Ide, D. Kondrashov, M. Mann, A. W. Robertson, A. Saunders, Y. Tian, F. Varadi *et al.*, *Spectral dimensionality reduction*. Springer, 2006.
- [11] M. Tipping and C. Bishop, “Mixtures of probabilistic principal component analyzers,” *Neural computation*, vol. 11, no. 2, pp. 443–482, 1999.
- [12] O. Cappé, E. Moulines, and T. Rydén, *Inference in Hidden Markov Models*. Springer, 2009.