Hidden Markov models applied to LFPs from layer two and three of human cortex reveal highly stereotypical discrete states in epileptic seizures separated by more than an hour

Joaquín Rapela[1], Dmitrii Todorov[1]

***Abstract*— It has been previously demonstrated [5] that spiking activity, recorded from layers two and three of human cortex by micro-electrode arrays, is highly stereotypical across seizures separated by hours. Later it was shown [7] that local field potential (LFP) and multi-unit activity (MUA) extracted from the above micro-electrode array recordings during seizures are organized in slow discrete states. Here we re-examine this stereotypical discrete-state representation using hidden Markov models (HMMs) fitted to LFPs extracted from the same micro-electrode array recordings used in [5], [7]. We fitted different HMMs with autoregressive (AR) observation models to LFPs from three different 40-minutes recording blocks, each block containing a different seizure from one subject. We found some states of the HMMs that were active only during ictal periods, other states were active only during post-ictal periods, and still other states were active only during pre-ictal periods. The AR observation models corresponding to states active during a seizure stage (e.g., AR models corresponding to ictal states) were similar across different blocks (e.g., ictal AR models from block one were very similar to ictal AR models from blocks two and three) based on their power spectrums, autocorrelation functions, coefficients, and variances. This strong similarity was not only observed during seizures, but also during pre- and post-ictal periods. Thus, this report supports previous investigations asserting that electrophysiological activity is highly stereotypical across seizures and that it is organized in slow discrete states, and extends these assertions to pre- and post-ictal periods.**

I. Introduction

Using microelectrode-array recordings from cortical layers two and three of humans, Truccolo et al. [5] have previously shown that spiking activity of isolated single neurons is highly stereotypical across seizures separated by hours. Figure 3ab of that article showed that not only a neuron that fires strongly/weakly during a first seizure also fires strongly/weakly during a subsequent seizure, but also the precise structure of a neuron rasterplot is similar across seizure (up to a small time warping). Subsequently, Wagner, Truccolo, et al. [7] demonstrated that during epileptic seizures dynamics of LFP and MUA evolve in discrete states, which are consistent across seizures within each subject.

Here we revisit the stereotypicality and discrete state hypothesis of epileptic seizures using HMMs. Because these models represent the probability of a sample as a mixture of conditional probabilities given discrete states, they are well suited to study the discrete state hypothesis. Also, if the stereotypicality hypothesis holds, we expect to find similar features in HMMs fitted to recordings from different seizures.

Previous research has used HMMs for seizure detection/predictions in dogs [1] and rats [4]. The focus of this manuscript is on the characterization, and not the detection/prediction, of epileptic activity with unique microelectrode-array recordings from human cortex.

II. Results

We use recorded electrical activity from layers two and three from a patient suffering from epilepsy. We separately analyzed three 40-minute blocks of local field potentials (LFPs; Section IV), each block containing a different spike-and-wave seizure (Figure 1). In each block, we fitted a Hidden Markov model, with AR observation model (Section IV), to LFPs from channel number 64[1][2]. We selected the number of states of the HMM and the order of the AR observation model by fitting models with different number of states and different AR orders and selecting the model that maximized the cross-validated log-likelihood. This procedure lead to an HMM with 14 states and an AR order of 67.

To study the physiological relevance of these states, we first looked at the times at which each state was most probable during the recordings block using fractional occupancy (Section IV). A value of fractional occupancy close to one/zero for a given state at a given time indicates that the given state was most/least probable, given the observations, around the given time. We say that a state is active at a given time if its factional occupancy at the given time is different from zero.

We found that some hidden states were active only during the occurrence of seizures (i.e., ictal periods), others were active only before the occurrence of seizures (i.e., pre-ictal periods) and still others were active only after the occurrence of seizures (i.e., post-ictal periods). In all three blocks we found exactly three states that were active only during ictal periods. We call these states ictal states. Figure 2 plots the fractional occupancy versus time of the ictal states (each state is depicted by a different color; seizure start and end times are marked by dotted vertical lines). The factional occupancy

*We thank Prof. Wilson Truccolo and Prof. Sydney Cash for sharing with us their unique micro-electrode array recorings.

[1]Truccolo Lab & Department of Neuroscience & Carney Institute for Brain Science, Brown University, Providence RI, USA; Email: {joaquin_rapela,dmitrii_todorov}@brown.edu

[1]As a first approximation, to expedite computations, specially those related to model selection, we fitted our HMMs to LFPs from a single instead of multiple channels.

[2]Because LFPs extracted from our high-density microelectrode array recordings are highly correlated, similar results were obtained using other electrodes.

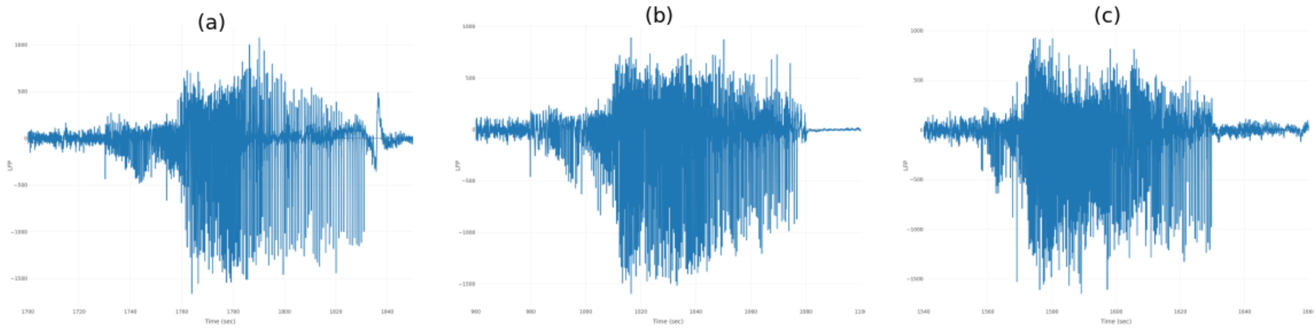

Fig. 1. LFPs of spike-and-wave seizures in blocks one (a), two (b) and three (c). Click on a panel to see the its interactive version.

curves of all ictal states were peaked. We call state A, B, C of a given block to the ictal state whose largest peak occurred first, second and third, respectively.

The factional occupancies of ictal states were remarkably similar across recording blocks. In each seizure, state A (red curve) was most probable (i.e., has largest fractional occupancy) around 20 seconds after seizure start, it was less probable 15 seconds later, and became more probable by the end of the seizure. States B and C were less probable than state A, their factional occupancy curves were unimodal, with peaks around the time when state A became less probable.

This shows that the timing of the three ictal states is similar across seizures separated by more than one hour. We next asked if statistical properties of the LFP corresponding to samples with different most probable states were different from each other and consistent across seizure blocks separated by more than one hour.

For each given state, s, we calculated the set $A(s)$ of samples whose most probable state was s. We then studied statistical properties of one-second-long epochs preceding samples in set $A(s)$. Below we considered two such properties, the power spectrum and the autocorrelation function (Section IV).

Figure 3 plots the power spectrum of the three ictal states for the three recordings blocks. In any block, the power spectrum of the different states were significantly different from each other. For instance, the power spectrum of state A was (frequency-wise) lower than that of state B, which in turn was lower than the power spectrum of state C. Also, the power spectrum of each state was similar across all recording blocks separated by more than one hour. For example, for all recording blocks (panels a-c) the power spectrum of state A (red curve) peaked at 5 Hz to a value around 30 dB and reached negative dB power at the highest frequencies. Further, as shown in Figure 4, within blocks autocorrelation functions were different for states with different labels and autocorrelation functions of states with the same label were remarkably similar between different blocks.

Another way of comparing states is by contrasting parameters of their AR models (e.g., coefficients, means and variances). Figure 5 plots the coefficients of the AR models associated with the ictal states in the three blocks, and gives the mean and variance of each state in the corresponding legend. Across blocks, coefficients tended to be similar, with maximum amplitude at lags around 150 ms and negligible magnitude for lags larger than 500 ms. Another regularity is that coefficients associated with state C had considerable larger magnitude than coefficients from states A and B, and that the AR model variance corresponding to state C was significantly larger than that of states B and C.

For brevity, the main body of this manuscript reports findings in ictal periods, and supplemental information [6] describes similar findings on pre- and post-ictal periods.

III. Discussion

Using HMMs we have found discrete states in epileptic seizures with statistical properties (e.g., autocorrelation function, power spectrum, coefficients and variance of HMMs) that are consistent across epileptic seizures separated by more than one hour (Figure 3, 4, 5).

Although appealing, our findings are very preliminary. Future research includes validating the generality of these first results in other subjects and seizure types, use an HMM observation model with combined LFP and MUA observations, model multiple channels, and handle longer duration recordings times (extending to several days or weeks).

These findings support previous observations in [5], [7] by showing that epileptic seizures are stereotypical and can be represented as sequences of discrete states. They extend these observations by revealing that discreetness and stereotypicality do not only hold in ictal periods, but also in pre- and post-ictal ones (supplementary information [6]).

Methodologically, the current investigation is substantially different from that in [7]. The semi-automatic segmentation algorithm used in this previous study operates on features extracted from LFPs and MUAs. Differently, our HMMs were fitted to raw LFPs, without feature engineering. Another difference is the the findings in [7] resulted from and ad-hoc non-probabilistic segmentation algorithm, while the current findings followed from a widely-used probabilistic model. That investigations using so different methodologies arrive to so similar findings suggests that the discreteness and stereotypicality of intra-cortical micro-array recordings of epileptic activity in humans are robust features and that they are not an artifact of the methodology used to observe them.

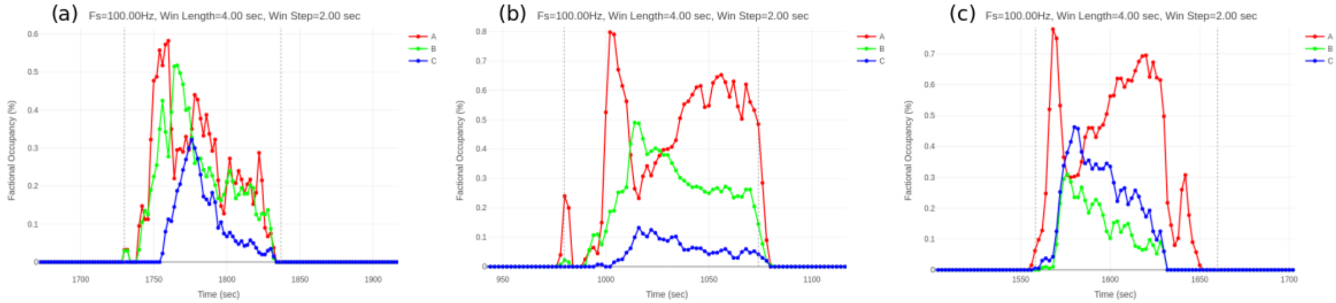

Fig. 2. Fractional occupancies of ictal states for blocks one (a), two (b) and three (c). Click on a panel to see its interactive version. For any ictal state, the shapes of its fractional occupancy curves were similar across the three recording blocks.

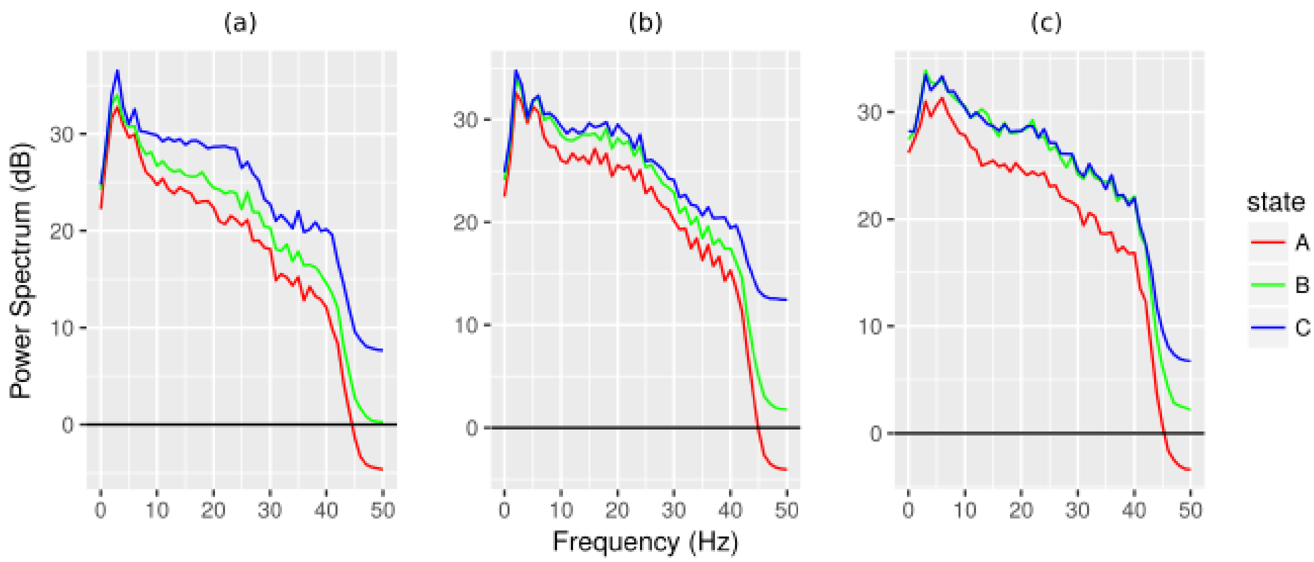

Fig. 3. Power spectrum of ictal states in blocks one (a), two (b) and three (c).

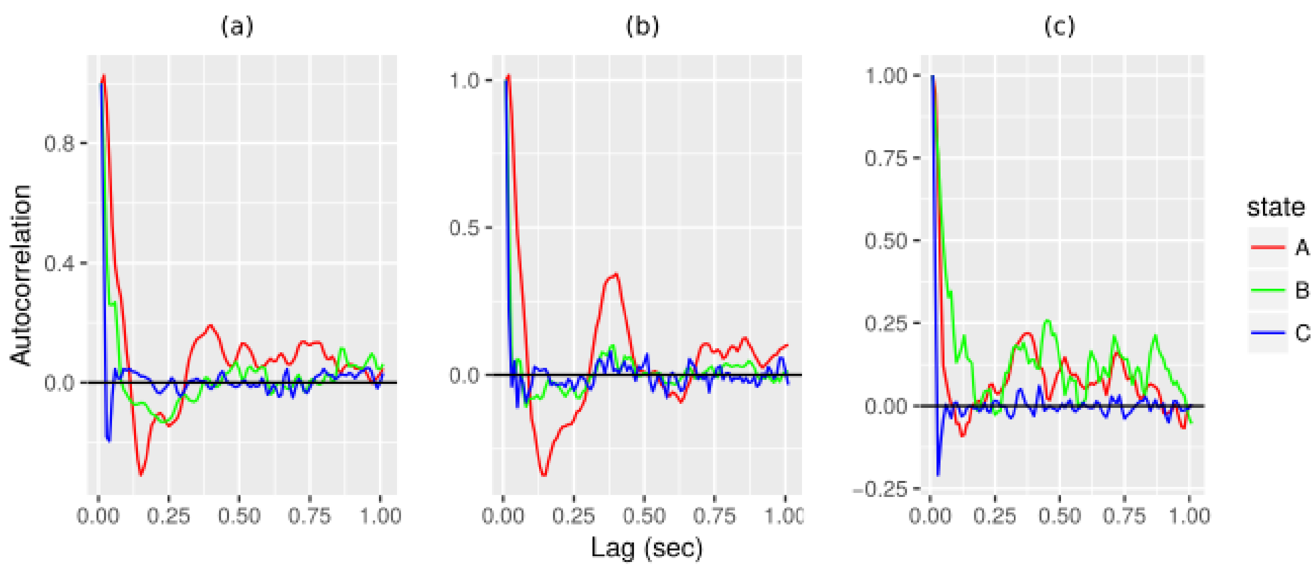

Fig. 4. Autocorrelation of ictal states in blocks one (a), two (b) and three (c).

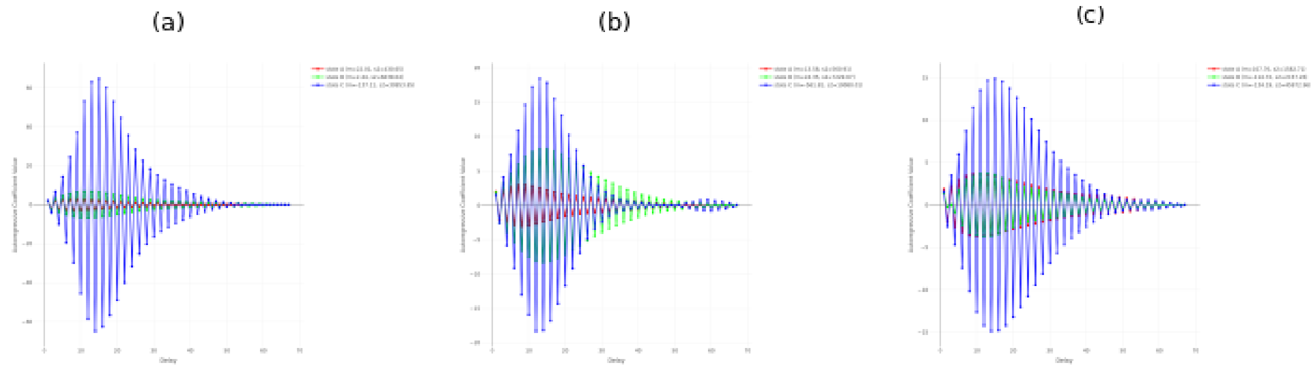

Fig. 5. Coefficients, means and variances of ictal states in blocks one (a), two (b) and three (c). Click on a panel to see the its interactive version.

IV. Methods

Recordings: We used broadband intracortically recorded field potentials (0.3 Hz-7.5 kHz; sampling at 30 kHz) from a 10×10 (4×4 mm^2) microelectrode array (96 recording electrodes plus 4 references) from cortical layers two and three of a person with focal epileptic seizures [5]. We characterized three 40-minute recordings blocks, each block contained one seizure. Seizure two occurred one hour and fifty minutes after seizure one and seizure three occurred two hours and fifty minutes after seizure two.

Pre-processing: We extracted LFPs by low-pass filtering the broadband field potentials (Butterwordth filter, order nine, cutoff frequency 500 Hz) and downsampling the result at a frequency of 2 kHz.

Hidden Markov model: To each block, we fitted a univariate HMM using expectation maximization [2] with an AR observation model [3][3].

Fractional occupancy: After fitting HMMs to LFPs we used the Viterbi algorithm [2] to find, for each sample time, the most probable state given the LFP observations. These most probable states can change rapidly among neighboring samples, making their visualization difficult. We overcame this difficulty using fractional occupancy, which for a given state, is the fraction of samples in a sliding window with most probable state equal to the given state. In this study we used a sliding window of length four seconds slid every two seconds.

State-dependent power spectrum and autocorrelation function: We used the Viterbi algorithm [2] to estimate the most probable state for each sample. For each state s, we built the set $A(s)$ of all samples at which s was the most probable state. Then, for each sample $q_s \in A(s)$ we build an epoch e_{q_s} of all samples preceding q_s by less than one second. To compute the power spectrum of a state s we averaged the square of the normalized Fourier transform of all epochs e_{q_s}, $P_s(\omega) = \text{median}_{q_s \in A(s)}(FT\{e_{q_s}\}(\omega))^2$. To compute the autocorrelation of a state s, we first removed from each sample s the mean of all samples with the same most probable state as s. Then, for each epoch e_{q_s}, with $q_s \in A(s)$, we computed the products between its value at lag zero (i.e., a sample from state s) and its value at all other lags, giving a new epoch of products $p_{q_s}(\tau) = e_{q_s}(0)e_{q_s}(\tau)$. We then summed all epochs of products across all samples of the state s, giving the (not normalized) autocovariance function for state s, $K_s(\tau) = \Sigma_{q_s \in A(s)} p_{q_s}(\tau)$. The autocorrelation function for state s is then $\rho_s(\tau) = K_s(\tau)/K_s(0)$.

[3]https://github.com/joacorapela/hiddenMarkovModels see example code under test/doTestAR.R

References

[1] Steven Baldassano, Drausin Wulsin, Hoameng Ung, Tyler Blevins, Mesha-Gay Brown, Emily Fox, and Brian Litt. A novel seizure detection algorithm informed by hidden markov model event states. *Journal of neural engineering*, 13(3):036011, 2016.

[2] Christopher M Bishop. *PATTERN RECOGNITION AND MACHINE LEARNING.* Springer-Verlag New York, 2016.

[3] Andrew M Fraser and Alexis Dimitriadis. Forecasting probability densities by using hidden markov models with mixed states. In Andreas S. Weigend and Neil A. Gershenfeld, editors, *Time series prediction: forecasting the future and understanding the past*, pages 265–282. Addison-Wesley, Reading, MA, 1993.

[4] Sridhar Sunderam, Ivan Osorio, Mark G Frei, and James F Watkins. Stochastic modeling and prediction of experimental seizures in sprague-dawley rats. *Journal of Clinical Neurophysiology*, 18(3):275–282, 2001.

[5] Wilson Truccolo, Jacob A Donoghue, Leigh R Hochberg, Emad N Eskandar, Joseph R Madsen, William S Anderson, Emery N Brown, Eric Halgren, and Sydney S Cash. Single-neuron dynamics in human focal epilepsy. *Nature neuroscience*, 14(5):635, 2011.

[6] `https://sccn.ucsd.edu/~rapela/embc19/hmm/paper/embc19_hmm_supplemental.pdf`.

[7] Fabien B Wagner, Emad N Eskandar, G Rees Cosgrove, Joseph R Madsen, Andrew S Blum, N Stevenson Potter, Leigh R Hochberg, Sydney S Cash, and Wilson Truccolo. Microscale spatiotemporal dynamics during neocortical propagation of human focal seizures. *Neuroimage*, 122:114–130, 2015.