Artificial Intelligence 216 (2014) 55–75

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Modeling the complex dynamics and changing correlations of epileptic events ☆

Drausin F. Wulsin [a,*], Emily B. Fox [c], Brian Litt [a,b]

[a] *Department of Bioengineering, University of Pennsylvania, Philadelphia, PA, United States*
[b] *Department of Neurology, University of Pennsylvania, Philadelphia, PA, United States*
[c] *Department of Statistics, University of Washington, Seattle, WA, United States*

A R T I C L E I N F O

Article history:
Received 14 January 2014
Received in revised form 19 May 2014
Accepted 20 May 2014
Available online 9 July 2014

Keywords:
Bayesian nonparametric
EEG
Factorial hidden Markov model
Graphical model
Time series

A B S T R A C T

Patients with epilepsy can manifest short, sub-clinical epileptic "bursts" in addition to full-blown clinical seizures. We believe the relationship between these two classes of events—something not previously studied quantitatively—could yield important insights into the nature and intrinsic dynamics of seizures. A goal of our work is to parse these complex epileptic events into distinct dynamic regimes. A challenge posed by the intracranial EEG (iEEG) data we study is the fact that the number and placement of electrodes can vary between patients. We develop a Bayesian nonparametric Markov switching process that allows for (i) shared dynamic regimes between a variable number of channels, (ii) asynchronous regime-switching, and (iii) an unknown dictionary of dynamic regimes. We encode a sparse and changing set of dependencies between the channels using a Markov-switching Gaussian graphical model for the innovations process driving the channel dynamics and demonstrate the importance of this model in parsing and out-of-sample predictions of iEEG data. We show that our model produces intuitive state assignments that can help automate clinical analysis of seizures and enable the comparison of sub-clinical bursts and full clinical seizures.

1. Introduction

Despite over three decades of research, we still have very little idea of what defines a seizure. This ignorance stems both from the complexity of epilepsy as a disease and a paucity of quantitative tools that are flexible enough to describe epileptic events but restrictive enough to distill intelligible information from them. Much of the recent machine learning work in electroencephalogram (EEG) analysis has focused on seizure prediction, [cf., 1,2], an important area of study but one that generally has not focused on parsing the EEG directly, as a human EEG reader would. Such parsings are central for diagnosis and relating various types of abnormal activity. Recent evidence shows that the range of epileptic events extends beyond clinical seizures to include shorter, sub-clinical "bursts" lasting fewer than 10 seconds [3]. What is the relationship between these shorter bursts and the longer seizures? In this work, we demonstrate that machine learning techniques can have substantial impact in this domain by unpacking how seizures begin, progress, and end.

In particular, we build a Bayesian nonparametric time series model to analyze intracranial EEG (iEEG) data. We take a modeling approach similar to a physician's in analyzing EEG events: look directly at the evolution of the raw EEG voltage

☆ This paper is an invited revision of a paper first published at the 30th International Conference on Machine Learning (ICML 2013).
* Corresponding author.

http://dx.doi.org/10.1016/j.artint.2014.05.006
0004-3702/

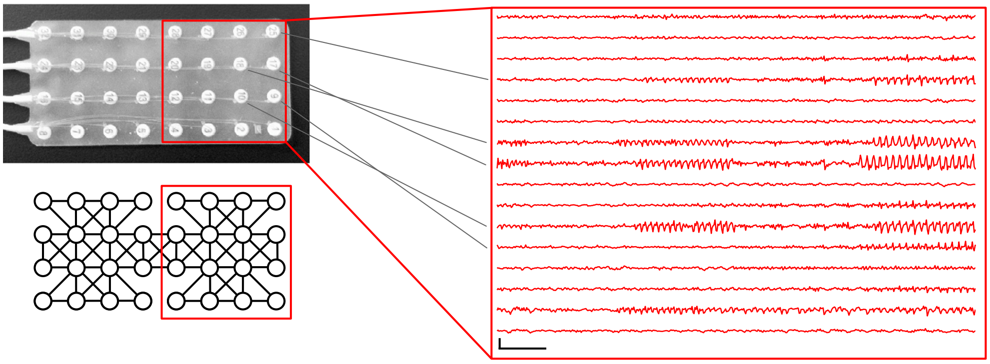

Fig. 1. An iEEG grid electrode and (**bottom left**) corresponding graphical model. (**middle**) Residual EEG values *after* subtracting predictions from a BP-AR-HMM assuming independent channels. All EEG scale bars indicate 1 mV vertically and 1 second horizontally.

traces. EEG signals exhibit nonstationary behavior during a variety of neurological events, and time-varying autoregressive (AR) processes have been proposed to model single channel data [4]. Here we aim to parse the recordings into interpretable regions of activity and thus propose to use autoregressive hidden Markov models (AR-HMMs) to define *locally* stationary processes. In the presence of multiple channels of simultaneous recordings, as is almost always the case in EEG, we wish to share AR states between the channels while allowing for asynchronous switches. The recent beta process (BP) AR-HMM of Fox et al. [5] provides a flexible model of such dynamics: a shared library of infinitely many possible AR states is defined and each time series uses a finite subset of the states. The process encourages sharing of AR states, while allowing for time-series-specific variability.

Conditioned on the selected AR dynamics, the BP-AR-HMM assumes independence between time series. In the case of iEEG, this assumption is almost assuredly false. Fig. 1 shows an example of a 4×8 intracranial electrode grid and the residual EEG traces of 16 channels *after* subtracting the predicted value in each channel using a conventional BP-AR-HMM. While the error term in some channels remains low throughout the recording, other channels—especially those spatially adjacent in the electrode grid—have very correlated error traces. We propose to capture correlations between channels by modeling a multivariate innovations process that drives independently evolving channel dynamics. We demonstrate the importance of accounting for this error trace in predicting heldout seizure recordings, making this a crucial modeling step before undertaking large-scale EEG analysis.

To aid in scaling to large electrode grids, we exploit a sparse dependency structure for the multivariate innovations process. In particular, we assume a graph with known vertex structure that encodes conditional independencies in the multivariate innovations process. The graph structure is based on the spatial adjacencies of the iEEG channels, with a few exceptions to make the graphical model fully decomposable. Fig. 1 (bottom left) shows an example of such a graphical model over the channels. Although the relative position of channels in the electrode grid is clear, determining the precise 3D location of each channel is extremely difficult. Furthermore, unlike in scalp EEG or magentoencephalogram (MEG), which have generally consistent channel positions from patient to patient, iEEG channels vary in number and position for each patient. These issues impede the use of alternative spatial and multivariate time series modeling techniques.

It is well-known that the correlations between EEG channels usually vary during the beginning, middle, and end of a seizure [6,7]. Prado et al. [8] employ a mixture-of-expert vector autoregressive (VAR) model to describe the different dynamics present in seven channels of scalp EEG. We take a similar approach by allowing for a Markov evolution to an underlying innovations covariance state.

An alternative modeling approach is to treat the channel recordings as a single multivariate time series, perhaps using a switching VAR process as in Prado et al. [8]. However, such an approach (i) assumes synchronous switches in dynamics between channels, (ii) scales poorly with the number of channels, and (iii) requires an identical number of channels between patients to share dynamics between event recordings.

Other work has explored nonparametric modeling of multiple time series. The infinite factorial HMM of Van Gael et al. [9] considers an infinite collection of chains each with a binary state space. The infinite hierarchical HMM [10] also involves infinitely many chains with finite state spaces, but with constrained transitions between the chains in a top down fashion. The infinite DBN of Doshi-Velez et al. [11] considers more general connection structures and arbitrary state spaces. Alternatively, the graph-coupled HMM of Dong et al. [12] allows graph-structured dependencies in the underlying states of some N Markov chains. Here, we consider a finite set of chains with infinite state spaces that evolve independently. The factorial structure combines the chain-specific AR dynamic states and the graph-structured innovations to generate the multivariate observations with sparse dependencies.

Expanding upon previous work [13], we show that our model for correlated time series has better out-of-sample predictions of iEEG data than standard AR- and BP-AR-HMMs and demonstrate the utility of our model in comparing short, sub-clinical epileptic bursts with longer, clinical seizures. Our inferred parsings of iEEG data concur with key features hand-annotated by clinicians but provide additional insight beyond what can be extracted from a visual read of the data. The