

Published in final edited form as:

Artif Intell. 2014 November 1; 216: 55–75. doi:10.1016/j.artint.2014.05.006.

Modeling the Complex Dynamics and Changing Correlations of Epileptic Events

Drausin F. Wulsin^a, Emily B. Fox^c, and Brian Litt^{a,b}

^aDepartment of Bioengineering, University of Pennsylvania, Philadelphia, PA

^bDepartment of Neurology, University of Pennsylvania, Philadelphia, PA

^cDepartment of Statistics, University of Washington, Seattle, WA

Abstract

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.

Keywords

Bayesian nonparametric; EEG; factorial hidden Markov model; graphical model; time series

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,

© 2014 Elsevier B.V. All rights reserved.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

[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 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. Figure 1 shows an example of a 4x8 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. Figure 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 magnetoencephalogram (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 importance of our methodology is multifold: (i) the output is interpretable to a practitioner and (ii) the parsings can be used to relate seizure types both within and between patients even with different electrode setups. Enabling such broad-scale automatic analysis, and identifying dynamics unique to sub-clinical seizures, can lead to new insights in epilepsy treatments.

Although we are motivated by the study of seizures from iEEG data, our work is much more broadly applicable in time series analysis. For example, perhaps one has a collection of stocks and wants to model shared dynamics between them while capturing changing correlations. The BP-AR-HMM was applied to the analysis of a collection of motion capture data assuming independence between individuals; our modeling extension could account for coordinated motion with a sparse dependency structure between individuals. Regardless, we find the impact in the neuroscience domain to be quite significant.

2. A Structured Bayesian Nonparametric Factorial AR-HMM

2.1. Dynamic Model

Consider an event with N univariate time series of length T . This event could be a seizure, where each time series is one of the iEEG voltage-recording channels. For clarity of exposition, we refer to the individual univariate time series as *channels* and the resulting N -dimensional multivariate time series (stacking up the channel series) as the *event*. We denote the scalar value for each channel i at each (discrete) time point t as $y_t^{(i)}$ and model it using an r -order AR-HMM [5]. That is, each channel is modeled via Markov switches between a set of AR dynamics. Denoting the latent state at time t by $z_t^{(i)}$, we have:

$$\begin{aligned} z_t^{(i)} &\sim \pi_{z_{t-1}^{(i)}}^{(i)}, \\ y_t^{(i)} &= \sum_{j=1}^r a_{z_t^{(i)},j}^{(i)} y_{t-j}^{(i)} + \varepsilon_t^{(i)} = \mathbf{a}_{z_t^{(i)}}^T \tilde{\mathbf{y}}_t^{(i)} + \varepsilon_t^{(i)}. \end{aligned} \quad (1)$$

Here, $\mathbf{a}_k = (a_{k,1}, \dots, a_{k,r})^T$ are the AR parameters for state k and π_k is the transition distribution from state k to any other state. We also introduce the notation $\tilde{\mathbf{y}}_t^{(i)}$ as the vector of r previous observations $(y_{t-1}^{(i)}, \dots, y_{t-r}^{(i)})^T$.

In contrast to a vector AR (VAR) HMM specification of the event, our modeling of channel dynamics separately as in Eq. (1) allows for (i) asynchronous switches and (ii) sharing of dynamic parameters between recordings with a potentially different number of channels. However, a key aspect of our data is the fact that the channels are correlated. Likewise, these correlations change as the patient progresses through various seizure *event states* (e.g., “resting”, “onset”, “offset”, ...). That is, the channels may display one innovation covariance before a seizure (e.g., relatively independent and low-magnitude) but quite a different covariance during a seizure (e.g., correlated, higher magnitude). To capture this, we jointly model the innovations $\varepsilon_t = (\varepsilon_t^{(1)}, \dots, \varepsilon_t^{(N)})^T$ driving the AR-HMMs of Eq. (1) as

$$\begin{aligned} Z_t &\sim \varphi_{Z_{t-1}}, \\ \varepsilon_t &\sim N(\mathbf{0}, \Delta_{Z_t}), \end{aligned} \quad (2)$$

where Z_t denotes a Markov-evolving event state distinct from the individual channel states $\{z_t^{(i)}\}$, φ_l the transition distributions, and Δ_k the event-state-specific channel covariance. That is each l describes a different set of channel relationships.

For compactness, we sometimes alternately write

$$\mathbf{y}_t = \mathbf{A}_{z_t} \tilde{\mathbf{Y}}_t + \varepsilon_t(Z_t), \quad (3)$$

where \mathbf{y}_t is the concatenation of N channel observations at time t and \mathbf{z}_t is the vector of concatenated channel states. The overall dynamic model is represented graphically in Figure 2.

Scaling to large electrode grids—To scale our model to a large number of channels, we consider a Gaussian graphical model (GGM) for ε_t capturing a sparse dependency structure amongst the channels. Let $G = (V, E)$ be an undirected graph with V the set of channel *nodes* i and E the set of *edges* with $(i, j) \in E$ if i and j are connected by an edge in the graph. Then, $[\Delta_l^{-1}]_{ij} = 0$ for all $(i, j) \notin E$, implying $\varepsilon_t^{(i)}$ is conditionally independent of $\varepsilon_t^{(j)}$ given $\varepsilon_t^{(k)}$ for all channels $k \neq i, j$. In our dynamic model of Eq. (1), statements of conditional independence of ε_t translate directly to statements of the observations \mathbf{y}_t .

In our application, we choose G based on the spatial adjacencies of channels in the electrode grid, as depicted in Figure 1 (bottom left). In addition to encoding the spatial proximities of iEEG electrodes, the graphical model also yields a sparse precision matrix Δ_l^{-1} , allowing for more efficient scaling to the large number of channels commonly present in iEEG. These computational efficiencies are made clear in Section 3.

Interpretation as a sparse factorial HMM—Recall that our formulation involves $N + 1$ independently evolving Markov chains: N chains for the channel states $z_t^{(i)}$ plus one for the event state sequence Z_t . As indicated by the observation model of Eq. (3), the $N + 1$ Markov chains jointly generate our observation vector \mathbf{y}_t leading to an interpretation of our formulation as a factorial HMM [14]. However, here we have a *sparse* dependency structure in how the Markov chains influence a given observation \mathbf{y}_t , as induced by the conditional independencies in ε_t encoded in the graph G . That is, $y_t^{(i)}$ only depends on Z_t the set of $z_t^{(i)}$ for which j is a neighbor of i in G .

2.2. Prior Specification

Emission parameters—As in the AR-HMM, we place a multivariate normal prior on the AR coefficients,

$$\mathbf{a}_k \sim N(\mathbf{m}_0, \Sigma_0), \quad (4)$$

with mean \mathbf{m}_0 and covariance Σ_0 . Throughout this work, we let $\mathbf{m}_0 = \mathbf{0}$.

For the channel covariances Δ_l with sparse precisions Δ_l^{-1} determined by the graph G , we specify a hyper-inverse Wishart (HIW) prior,

$$\Delta_l \sim HIW_G(b_0, D_0), \quad (5)$$

where b_0 denotes the degrees of freedom and D_0 the scale. The HIW prior [15] enforces hyper-Markov conditions specified by G .

Feature constrained channel transition distributions—A natural question is how many AR states are the channels switching between? Likewise, which are shared between the channels and which are unique? We expect to see similar dynamics present in the channels (sharing of AR processes), but also some differences. For example, maybe only some of the channels ever get excited into a certain state. To capture this structure, we take a

Bayesian nonparametric approach building on the beta process (BP) AR-HMM of Fox et al. [16]. Through the beta process prior [17], the BP-AR-HMM defines a shared library of infinitely many AR coefficients $\{\mathbf{a}_k\}$, but encourages each channel to only use a sparse subset of them.

The BP-AR-HMM specifically defines a featural model. Let $\mathbf{f}^{(i)}$ be a binary feature vector associated with channel i with $f_k^{(i)}=1$ indicating that channel i uses the dynamic described by \mathbf{a}_k . Formally, the feature assignments $f_k^{(i)}$ and their corresponding parameters \mathbf{a}_k are generated by a beta process random measure and the conjugate Bernoulli process (BeP),

$$\begin{aligned} B &\sim BP(1, B_0), \\ X^{(i)} &\sim BeP(B), \end{aligned} \quad (6)$$

with base measure B_0 over the parameter space $\Theta = \mathbb{R}^r$ for our r -order autoregressive parameters \mathbf{a}_k . As specified in Eq. (4), we take the normalized measure $B_0/B_0(\Theta)$ to be $N(\mathbf{m}_0, \Sigma_0)$. The discrete measures B and $X^{(i)}$ can be represented as

$$B = \sum_{k=1}^{\infty} \omega_k \delta_{\mathbf{a}_k}, \quad X^{(i)} = \sum_{k=1}^{\infty} f_k^{(i)} \delta_{\mathbf{a}_k}, \quad (7)$$

with $f_k^{(i)} \sim Ber(\omega_k)$.

The resulting feature vectors $\mathbf{f}^{(i)}$ constrain the set of available states $z_t^{(i)}$ can take by constraining each transition distributions, $\pi_j^{(i)}$, to be 0 when $f_k^{(i)}=0$. Specifically, the BP-AR-HMM defines $\pi_j^{(i)}$ by introducing a set of gamma random variables, $\eta_{jk}^{(i)}$, and setting

$$\eta_{jk}^{(i)} \sim Gamma(\gamma_c + \kappa_c \delta(j, \kappa)) \quad (8)$$

$$\pi_j^{(i)} = \frac{\eta_j^{(i)} \circ \mathbf{f}^{(i)}}{\sum_{k|f_k^{(i)}=1} \eta_{jk}^{(i)}}. \quad (9)$$

The positive elements of $\pi_j^{(i)}$ can also be thought of as a sample from a finite Dirichlet distribution with only $K^{(i)}$ dimensions, where $K^{(i)} = \sum_k f_k^{(i)}$ represents the number of states channel i uses. For convenience, we sometimes denote the set of transition variables $\{\eta_{jk}^{(i)}\}_j$ as $\boldsymbol{\eta}^{(i)}$. As in the sticky HDP-HMM of Fox et al. [18], the parameter κ_c encourages self-transitions (i.e., state j at time $t-1$ to state j at time t).

Unconstrained event transition distributions—We again take a Bayesian non-parametric approach to define the event state HMM, building on the sticky HDP-HMM [18]. In particular, the transition distributions ϕ_l are hierarchically defined as

$$\begin{aligned}\beta &\sim \text{stick}(\alpha), \\ \varphi_l &\sim DP(\alpha_e \beta + \kappa_e \mathbf{e}_l),\end{aligned}\quad (10)$$

where $\text{stick}(\alpha)$ refers to a stick-breaking measure, also known as $\text{GEM}(\alpha)$, with β generated by

$$\begin{aligned}\beta'_k &\sim \text{Beta}(1, \alpha), \quad k=1, 2, \dots, \\ \beta_k &= \beta'_k \prod_{\ell=1}^{k-1} (1 - \beta'_\ell), \quad k=1, 2, \dots, \\ &= \beta'_k \left(1 - \sum_{\ell=1}^{k-1} \beta'_\ell\right) \quad k=1, 2, \dots.\end{aligned}\quad (11)$$

Again, the sticky parameter κ_e promotes self-transitions, reducing state redundancy.

We term this model the *sparse factorial BP-AR-HMM*. Although the graph G can be arbitrarily structured, because of our motivating seizure modeling application with a focus on a spatial-based graph structure, we often describe the sparse factorial BP-AR-HMM as capturing *spatial* correlations. We depict this model in the directed acyclic graphs shown in Figure 3. Note that while we formally consider a model of only a single event for notational simplicity, our formulation scales straightforwardly to multiple independent events. In this case, everything except the library of AR states $\{\mathbf{a}_k\}$ becomes event-specific. If all events share the same channel setup, we can assume the channel covariances $\{\Sigma\}$ are shared as well.

3. Posterior Computations

Although the components of our model related to the individual channel dynamics are similar to those in the BP-AR-HMM, our posterior computations are significantly different due to the coupling of the Markov chains via the correlated innovations $\boldsymbol{\varepsilon}$. In the BP-AR-HMM, conditioned on the feature assignments, each time series is independent. Here, however, we are faced with a factorial HMM structure and the associated challenges. Yet the underlying graph structure of the channel dependencies mitigates the scale of these challenges.

Algorithm 1 Sparse factorial BP-AR-HMM master MCMC sampler

- 1: **for** each MCMC iteration **do**
- 2: get a random permutation \mathbf{h} of the channel indices,
- 3: **for** each channel $i \in \mathbf{h}$ **do**
- 4: sample feature indicators $\mathbf{f}^{(i)}$ as in Eq. (12)
- 5: sample state sequence $z_{1:T}^{(i)}$ as in Eq. (13)
- 6: sample state transition parameters $\boldsymbol{\eta}^{(i)}$ as in Eq. (14)
- 7: **end for**
- 8: sample event states sequence $Z_{1:T}$
- 9: sample event state transition parameters $\boldsymbol{\phi}$ as in Eq. (17)
- 10: sample channel AR parameters $\{\mathbf{a}_k\}$ as in Eq. (19)

- 11: sample channel $\{i\}$ as in Eq. (18)
 - 12: (sample hyperparameters $\gamma_c, \kappa_c, \alpha_c, \kappa_e, \gamma_e,$ and $\alpha_e = B_0(\Theta)$)
 - 13: **end for**
-

Conditioned on channel sequences $\{\mathbf{z}_{1:T}^{(i')}\}$, we can marginalize $z_{1:T}^{(i)}$; because of the graph structure, we need only condition on a *sparse* set of other channels i' (i.e., neighbors of channel i in the graph). This step is important for efficiently sampling the feature assignments $\mathbf{f}^{(i)}$.

At a high level, each MCMC iteration proceeds through sampling channel states, events states, dynamic model parameters, and hyperparameters. Algorithm 1 summarizes these steps, which we briefly describe below and more fully in Appendices B-D.

Individual channel variables

We harness the fact that we can compute the marginal likelihood of $\mathbf{y}_{1:T}$ given $\mathbf{f}^{(i)}$ and the neighborhood set of other channels $\mathbf{z}_{1:T}^{(i')}$ in order to block sample $\{\mathbf{f}^{(i)}, z_{1:T}^{(i)}\}$. That is, we first sample $\mathbf{f}^{(i)}$ marginalizing $\mathbf{z}_{1:T}^{(i)}$ and then sample $z_{1:T}^{(i)}$ given the sampled $\mathbf{f}^{(i)}$. Sampling the active features $\mathbf{f}^{(i)}$ for channel i follows as in Fox et al. [5], using the Indian buffet process (IBP) [19] predictive representation associated with the beta process, but using a likelihood term that conditions on neighboring channel state sequences $\mathbf{z}_{1:T}^{(i')}$ and observations $\mathbf{y}_{1:T}^{(i')}$. We additionally condition on the event state sequence $Z_{1:T}$ to define the sequence of distributions on the innovations. Generically, this yields

$$p\left(f_k^{(i)} | y_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{F}^{-ik}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\}\right) \propto p\left(f_k^{(i)} | \mathbf{F}^{-ik}\right) p\left(y_{1:T}^{(i)} | \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{F}^{-ik}, f_k^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\}\right). \quad (12)$$

Here, \mathbf{F}^{-ik} denotes the set of feature assignments not including $f_k^{(i)}$. The first term is given by the IBP prior and the second term is the marginal conditional likelihood (marginalizing $z_{1:T}^{(i)}$). Based on the derived marginal conditional likelihood, feature sampling follows similarly to that of Fox et al. [5].

Conditioned on $\mathbf{f}^{(i)}$, we block sample the state sequence $z_{1:T}^{(i)}$ using a backward filtering forward sampling algorithm (see Appendix A) based on a decomposition of the full conditional as

$$p\left(z_{1:T}^{(i)} | y_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\}\right) = p\left(z_1^{(i)} | y_1^{(i)}, \mathbf{y}_1^{(i')}, \mathbf{z}_1^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\}\right) \prod_{t=2}^T p\left(z_t^{(i)} | y_{t:T}^{(i)}, \mathbf{y}_{t:T}^{(i')}, z_{t-1}^{(i)}, \mathbf{z}_{t:T}^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\}\right). \quad (13)$$

For sampling the transition parameters $\boldsymbol{\eta}^{(i)}$, we follow Hughes et al. [20, Supplement] and sample from the full conditional

$$p(\eta_{jk}^{(i)} | z_{1:T}^{(i)}, f_k^{(i)}) \propto \frac{(\eta_{jk}^{(i)})^{n_{jk}^{(i)} + \gamma_c + \kappa_c \delta(j,k) - 1} e^{\eta_{jk}^{(i)}}}{\sum_{k' | f_k^{(i)} = 1} \eta_{jk'}^{(i)}}, \quad (14)$$

where $n_{jk}^{(i)}$ denotes the number of times channel i transitions from state j to state k . We sample $\eta_j^{(i)} = C_j^{(i)} \bar{\eta}_j^{(i)}$ from its posterior via two auxiliary variables,

$$\begin{aligned} \bar{\eta}_j^{(i)} &\sim \text{Dir}(\gamma_c + \mathbf{e}_j \kappa_c + \mathbf{n}_j^{(i)}) \\ C_j^{(i)} &\sim \text{Gamma}(K\gamma_c + \kappa_c, 1), \end{aligned} \quad (15)$$

where $\mathbf{n}_j^{(i)}$ gives the transition counts from state j in channel i .

Event variables $\{\phi_l, \beta, \mathbf{Z}_{1:T}\}$

Conditioned on the channel state sequences $\mathbf{z}_{1:T}$ and AR coefficients $\{\mathbf{a}_k\}$, we can compute an innovations sequence as $\boldsymbol{\varepsilon}_t = \mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t$, where we recall the definition of \mathbf{A}_k and $\tilde{\mathbf{Y}}_t$ from Eq. (3). These innovations are the observations of the sticky HDP-HMM of Eq. (2). For simplicity and to allow block-sampling of $\mathbf{z}_{1:T}$, we consider a weak limit approximation of the sticky HDP-HMM as in [18]. The top-level Dirichlet process is approximated by an L -dimensional Dirichlet distribution [21], inducing a finite Dirichlet for ϕ_l :

$$\begin{aligned} \beta &\sim \text{Dir}(\gamma_e/L, \dots, \gamma_e/L), \\ \phi_l &\sim \text{Dir}(\alpha_e \beta + \kappa_e \mathbf{e}_l). \end{aligned} \quad (16)$$

Here, L provides an upper bound on the number of states in the HDP-HMM. The weak limit approximation still encourages using a subset of these L states.

Based on the weak limit approximation, we first sample the parent transition distribution β as in [22, 18], followed by sampling each ϕ_l from its Dirichlet posterior,

$$p(\phi_l | Z_{1:T}, \beta) \propto \text{Dir}(\alpha_e \beta + \mathbf{e}_l \kappa_e + \mathbf{n}_l), \quad (17)$$

where \mathbf{n}_l is a vector of transition counts of $Z_{1:T}$ from state l to the L different states.

Using standard conjugacy results, based on ‘‘observations’’ $\boldsymbol{\varepsilon}_t = \mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t$ for t such that $Z_t = l$, the full conditional for Δ_l is given by

$$p(\Delta_l | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_k\}) \propto \text{HIW}_G(b_l, D_l), \quad (18)$$

Where

$$\begin{aligned} b_l &= b_0 + |\{t | Z_t = l, t = 1, \dots, T\}|, \\ D_l &= D_0 + \sum_{t | Z_t = l} \boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t^T. \end{aligned}$$

Details on how to efficiently sample from a HIW distribution are provided in [23].

Conditioned on the truncated HDP-HMM event transition distributions $\{\phi_l\}$ and emission parameters $\{\lambda_l\}$, we use a standard backward filtering forward sampling scheme to block sample $Z_{1:T}$.

AR coefficients, $\{\mathbf{a}_k\}$

Each observation \mathbf{y}_t is generated based on a *matrix* of AR parameters

$\mathbf{A}_{z_t} = [\mathbf{a}_{z_t^{(1)}} | \cdots | \mathbf{a}_{z_t^{(N)}}]$. Thus, sampling \mathbf{a}_k involves conditioning on $\{\mathbf{a}_{k'}\}_{k' \neq k}$ and disentangling the contribution of \mathbf{a}_k on each y_t . As derived in Appendix E, the full conditional for \mathbf{a}_k is a multivariate normal

$$p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) \propto \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (19)$$

where

$$\begin{aligned} \boldsymbol{\Sigma}_k^{-1} &= \boldsymbol{\Sigma}_0^{-1} + \sum_{t=1}^T \bar{\mathbf{Y}}_t^{(\mathbf{k}^+)} \Delta_{Z_t}^{-1(\mathbf{k}^+, \mathbf{k}^+)} \left(\bar{\mathbf{Y}}_t^{(\mathbf{k}^+)} \right)^T, \\ \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\mu}_k &= \sum_{t=1}^T \bar{\mathbf{Y}}_t^{(\mathbf{k}^+)} \left(\Delta_{Z_t}^{-1(\mathbf{k}^+, \mathbf{k}^+)} \mathbf{y}_t^{(\mathbf{k}^+)} + \Delta_{Z_t}^{-1(\mathbf{k}^+, \mathbf{k}^-)} \boldsymbol{\varepsilon}_t^{(\mathbf{k}^-)} \right). \end{aligned}$$

The vectors \mathbf{k}^+ and \mathbf{k}^- denote the indices of channels assigned and not assigned to state k at time t , respectively. We use these to index into the rows and columns of the vectors $\boldsymbol{\varepsilon}_t$, \mathbf{y}_t , and matrix Δ_{Z_t} . Each column of matrix $\bar{\mathbf{Y}}_t^{(\mathbf{k}^+)}$ is the previous r observations for one of the channels assigned to state k at time t .

Hyperparameters

See Appendix D for the prior and full conditionals of the hyperparameters γ_c , κ_c , a_e , κ_e , γ_e , and $a_c = B_0(\Theta)$.

4. Experiments

4.1. Simulation Experiments

To initially explore some characteristics of our sparse factorial BP-AR-HMM, we examined a small simulated dataset of six time series in a 2×3 spatial arrangement, with vertices connecting all adjacent nodes (i.e., two cliques of 4 nodes each). We generated an event of length 2000 time points as follows. We defined five first-order AR channel states linearly spaced between -0.9 and 0.9 and three event states with covariances shown in the bottom left of Figure 4. Channel and event state transition distributions were set to 0.99 and 0.9, respectively, for a self-transition and uniform between the other states. Channel feature indicators $f_k^{(i)}$ were simulated from an IBP with $a_c = 10$ (no channel had indicators exceeding the five specified states). The sampled $\mathbf{f}^{(i)}$ were then used to modify the channel state transition distributions by setting to 0 transitions to states with $f_k^{(i)} = 0$ and then

renormalizing. Using these feature-constrained transition distributions, we simulated sequences $z_{1:T}^{(i)}$ for each channel $i = 1, \dots, 6$ and for $T = 2000$. The event sequence $Z_{1:T}$ was likewise simulated. Based on these sampled state sequences, and using the defined state-specific AR coefficients and channel covariances, we generated observations $y_{1:T}$ as in Eq. (3).

We ran our MCMC sampler for 6000 iterations, discarding the first 1000 as burn-in and thinning the chain by 10. Figure 4 shows the generated data and its true states along with the inferred states and learned channel covariances for a representative posterior sample. The event state matching is almost perfect, and the channel state matching is quite good, though we see that the sampler added an additional (yellow) state in the middle of the first time series when it should have assigned that section to the cyan state. The scale and structure of the estimated event state covariances match the true covariances quite well. Furthermore, Table 1 shows how the posterior estimates of the channel state AR coefficients also center well around the true values.

4.2. Parsing a Seizure

We tested the sparse factorial BP-AR-HMM on two similar seizures (events) from a patient of the Children's Hospital of Pennsylvania. These seizures were chosen because qualitatively they displayed a variety of dynamics throughout the beginning, middle, and end of the seizure and thus are ideal for exploring the extent to which our sparse factorial BP-AR-HMM can parse a set of rich neurophysiologic signals. We used the 90 seconds of data after the clinically-determined starts of each seizure from 16 channels, whose spatial layout in the electrode grid is shown in Figure 5 along with the graph encoding our conditional independence assumptions. The data were low-pass filtered and downsampled from 200 to 50 Hz, preserving the clinically important signals but reducing the computational burden of posterior inference. The data was also scaled to have 99% of values within $[-10, 10]$ for numerical reasons. We examined a 5th-order sparse factorial BP-AR-HMM and ran 10 MCMC chains for 6000 iterations, discarding 1000 samples as burn-in and using 10-sample thinning.

The sparse factorial BP-AR-HMM inferred state sequences for the sample corresponding to a minimum expected Hamming distance criterion ([18]) are shown in Figure 5. The results were analyzed by a board-certified epileptologist who agreed with the model's judgement in identifying the subtle changes from the background dynamic (cyan) initially present in all channels. The model's grouping of spatially-proximate channels into similar state transition patterns (e.g., channels 03, 07, 11, 15) was clinically intuitive and consistent with his own reading of the raw EEG. Using only the raw EEG, and prior to disclosing our results, he independently identified roughly six points in the duration of the seizure where the dynamics fundamentally change. The three main event state transitions shown in Figure 5 occurred almost exactly at the same time as three of his own marked transitions. The fourth coincides with a major shift in the channel dynamics with most channels transitioning to the green dynamic. The other two transitions he marked that are not displayed occurred after this onset period. From this analysis, we see that our event states provide an important global summary

of the dynamics of the seizure that augments the information conveyed from the channel state sequences.

Clinical relevance—While interpreting these state sequences and covariances from the model, it is important to keep in mind that they are ultimately estimates of a system whose parsing even highly-trained physicians disagree upon. Nevertheless, we believe that the event states directly describe the activity of particular clinical interest.

In modeling the correlations between channels, the event states give insight into how different physiologic areas of the brain interact over the course of a seizure. In the clinical workup for resective brain surgery, these event states could help define and specifically quantify the full range of ways in which neurophysiologic regions initiate seizures and how others are recruited over the numerous seizures of a patient. In addition, given fixed model parameters, our model can fit the channel and event state sequences of an hour's worth of 64-channel EEG data in a matter of minutes on a single 8-core machine, possibly facilitating epileptologist EEG annotation of long-term monitoring records.

The ultimate clinical aim of this work, however, involves understanding the relationship between epileptic bursts and seizures. Because the event state aspect of our model involves a Markov assumption, the intrinsic length of the event has little bearing on the states assigned to particular time points. Thus, these event states allow us to straightforwardly compare the neurophysiologic relationship dynamics in short bursts (often less than two seconds long) to those in much longer seizures (on the order of two minutes long), as explored in Section 4.4. Prior to this analysis, we first examine the importance of our various model components by comparing to baseline alternatives.

4.3. Model Comparison

The advantages of a spatial model—We explored the extent to which the spatial information and sparse dependencies encoded in the HIW prior improves our predictions of heldout data relative to a number of baseline models. To assess the impact of the sparse dependencies induced by the Gaussian graphical model for $\boldsymbol{\varepsilon}_t$, we compare to a full-covariance model with an IW prior on $\boldsymbol{\Sigma}$ (dense factorial). For assessing the importance of spatial correlations, we additionally compare to two alternatives where channels evolve independently: the BP-AR-HMM of Fox et al. [5] and an AR-HMM without the feature-based modeling provided by the beta process [24]. Both of these models use inverse gamma (IG) priors on the individual channel innovation variances. We learned a set of AR coefficients $\{\mathbf{a}_k\}$ and event covariances $\{\boldsymbol{\Sigma}_l\}$ on one seizure and then computed the heldout log-likelihood on a separate seizure, constraining it use the learned model parameters from the training seizure.

For the training seizure, MCMC samples were collected over 5000 samples across 10 chains, each with a 1000-sample burn in and 10-sample thinning. To compute the predictive log-likelihood of the heldout seizure, we analytically marginalized the heldout event state sequence $Z_{1:T}$ but perform a Monte Carlo integration over the feature vectors $\mathbf{f}^{(t)}$ and channel states $\mathbf{z}_{1:T}$ using our MCMC sampler. For each original MCMC sample generated

from the training seizure, a secondary chain is run fixing $\{\mathbf{a}_k\}$ and $\{\mathbf{l}\}$ and sampling $z_t^{(i)}, Z_t, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}$, and $\{\boldsymbol{\phi}\}$ for the heldout seizure. We approximate $p(\mathbf{y}_{1:T} | \{\boldsymbol{\phi}\}, \{\mathbf{a}_k\}, \{\mathbf{l}\})$ by averaging the secondary chain's closed-form $p(\mathbf{y}_{1:T} | \mathbf{z}_{1:T}, \{\boldsymbol{\phi}\}, \{\mathbf{a}_k\}, \{\mathbf{l}\})$, described in Appendix B.

Figure 6 (left) shows how conditioning on the innovations of neighboring channels in the sparse factorial model improves the prediction of an individual channel, as seen by its reduced innovation trace relative to the original BP-AR-HMM. The quantitative benefits of accounting for these correlations are seen in our predictions of heldout events, as depicted in Figure 6 (right), which compares the heldout log-likelihoods for the original and the factorial models listed above. As expected, the factorial models have significantly larger predictive power than the original models. Though hard to see due to the large factorial/original difference, the BP-based model also improves on the standard non-feature-based AR-HMM. Performance of the sparse factorial model is also at least as competitive as a full-covariance model (dense factorial). We would expect to see even larger gains for electrode grids with more channels due to the parsimonious representation presented by the graphical model. Regardless, these results demonstrate that the assumptions of sparsity in the channel dependencies do not adversely affect our performance.

The advantages of sparse factorial dependencies—In addition to providing a parsimonious modeling tool, the sparse dependencies among channels induced by the HIW prior allow our computations to scale linearly to the large number of channels present in iEEG. We compared a dense factorial BP-AR-HMM (entailing a fully-connected spatial graph) and a sparse factorial BP-AR-HMM on five datasets of 8, 16, 32, 64, and 96 channels (from three 32-channel electrodes) from the same seizure used previously. We ran the two models on each of the five datasets for at least 1000 MCMC iterations, using a profiler to tabulate the time spent in each step of the MCMC iteration.

Figure 7 shows the average time required to calculate the channel likelihoods at each time point under each AR channel state. This computation is used both for calculating the marginal likelihood (averaging over all the state sequences $z_{1:T}^{(i)}$) required in active feature sampling as well as in sampling the state sequences $z_{1:T}^{(i)}$. In our sparse factorial model, each channel has a constant set of M dependencies, assuming M neighboring channels. As such, the channel likelihood computation at each time point has an $O(M)$ complexity, implying an $O(MN)$ complexity for calculating the likelihoods of all N channels at each time point. In contrast, the likelihood computation at each time point under the full covariance model had complexity $O(N)$, implying $O(N^2)$ for calculating all the channel likelihoods. For $M \ll N$, as is typically the case, our sparse dependency model is significantly more computationally efficient.

Anecdotally, we also found that the IW prior experiments—especially those with larger number of channels—tended to occasionally have numerical underflow problems associated with the inverse term $\Delta_{Z_t}^{-1}(\mathbf{i}', \mathbf{i}')$ in the conditional channel likelihood calculation. This underflow in the IW prior model calculations is not surprising since the matrices inverted

are of dimension $N - 1$ (for N channels), whereas in the HIW prior, the sparse spatial dependencies of the electrode grids make these matrices no larger than eight-by-eight.

4.4. Comparing Epileptic Events of Different Scales

We applied our sparse factorial BP-AR-HMM to six channels of iEEG over 15 events from a human patient with hippocampal depth electrodes. These events comprise 14 short sub-clinical epileptic bursts of roughly five to eight seconds and a final, 2-3 minute clinical seizure. Our hypothesis was that the sub-clinical bursts display initiation dynamics similar to those of a full, clinical seizure and thus contain information about the seizure-generation process.

The events were automatically extracted from the patient's continuous iEEG record by taking sections of iEEG whose median line-length feature [25] crossed a preset threshold, also including 10 seconds before and after each event. The iEEG was preprocessed in the same way as in the previous section. The six channels studied came from a depth electrode implanted in the left temporal lobe of the patient's brain. We ran our MCMC sampler jointly on the 15 events. In particular, the AR channel state and event state parameters, $\{\mathbf{a}_k\}$ and $\{\beta\}$, were shared between the 15 events such that the parsings of each recording jointly informed the posteriors of these shared parameters. The hyperparameter settings, number of MCMC iterations, chains, and thinning was as in the experiment of Section 4.2.

Figure 8 compares two of the 14 sub-clinical bursts and the onset of the single seizure. We have aligned the three events relative to the beginnings of the red event state common to all three, which we treat as the start of the epileptic activity. The individual channel states of the four middle channels are also all green throughout most of the red event state. It is interesting to note that at this time the fifth channel's activity in all three events is much lower than those of the three channels above it, yet it is still assigned to the green state and continues in that state along with the other three channels as the event state switches from the red to the lime green state in all three events. While clinical opinions can vary widely in EEG reading, a physician would most likely not consider this segment of the fifth channel similar to the other three, as our model consistently does. But on a relative voltage axis, the segments actually look quite similar. In a sense, the fifth channel has the same dynamics as the other three but just with smaller magnitude. This kind of relationship is difficult for the human EEG readers to identify and shows how models such as ours are capable of providing new perspectives not readily apparent to a human reader. Additionally, we note the similarities in event state transitions.

The similarities mentioned above, among others, suggest some relationship between these two different classes of epileptic events. However, all bursts make a notable departure from the seizure: a large one-second depolarization in the middle three channels, highlighted at the end by the magenta event state and followed shortly thereafter by the end of the event. Neither the states assigned by our model nor the iEEG itself indicates that dynamic present in the clinical seizure. This difference leads us to posit that perhaps these sub-clinical bursts are a kind of false-start seizure, with similar onset patterns but a disrupting discharge that prevents the event from escalating to a full-blown seizure.

5. Conclusion

In this work, we develop a sparse factorial BP-AR-HMM model that allows for shared dynamic regimes between a variable number of time series, asynchronous regime-switching, and an unknown dictionary of dynamic regimes. Key to our model is capturing the between-series correlations and their evolution with a Markov-switching multivariate innovations process. For scalability, we assume a sparse dependency structure between the time series using a Gaussian graphical model.

This model is inspired by challenges in modeling high-dimensional intracranial EEG time series of seizures, which contain a large variety of small- and large-scale dynamics that are of interest to clinicians. We demonstrate the value of this unsupervised model by showing its ability to parse seizures in a clinically intuitive manner and to produce state of the art out-of-sample predictions on the iEEG data. Finally, we show how our model allows for flexible, large-scale analysis of different classes of epileptic events, opening new, valuable clinical research directions previously too challenging to pursue. Such analyses have direct relevance to clinical decision-making for patients with epilepsy.

Algorithm 2 HMM forward-filtering algorithm for calculating $p(y_{1:T})$

- 1: let π_0 and $\pi = (\pi_1 | \dots | \pi_k)^T$ be the initial and transition distributions
 - 2: let $\mathbf{u}_t \in \mathbb{R}_+^K$ be the likelihoods $p(y_t | z_t = k)$ for each k
 - 3: let $\xi_t \in \mathbb{R}_+^K$ be the forward messages $p(y_{1:t-1}, z_{t-1} = k)$ at time t for each k
 - 4: normalize each \mathbf{u}_t to sum to 1, preventing underflow during computations
 - 5: **for** $t = 1, \dots, T$ **do**
 - 6: store the marginal log probability over \mathbf{u}_t in v_t : $v_t \leftarrow \log(\mathbf{1}^T \mathbf{u}_t)$
 - 7: normalize \mathbf{u}_t : $\tilde{\mathbf{u}}_t \leftarrow \mathbf{u}_t / \exp(v_t)$
 - 8: **end for**
 - 9: calculate and normalize initial forward messages
 - 10: $\xi \leftarrow \tilde{\mathbf{u}}_1 \circ \pi_0$
 - 11: $w_1 \leftarrow \mathbf{1}^T \xi_1$, $\xi_1^{\sim} \leftarrow \xi_1 / w_1$
 - 12: $v_1 \leftarrow v_1 + \log(w_1)$
 - 13: propagate messages forward through each time point
 - 14: **for** $t = 2, \dots, T$ **do**
 - 15: transmit messages forward: $\xi_t \leftarrow \tilde{\mathbf{u}}_t \circ (\pi^T \xi_{t-1}^{\sim})$
 - 16: normalize: $w_t \leftarrow \mathbf{1}^T \xi_t$, $\xi_t^{\sim} \leftarrow \xi_t / w_t$
 - 17: $v_t \leftarrow v_t + \log(w_t)$
 - 18: **end for**
 - 19: calculate final marginal log likelihood: $\log p(y_{1:T}) = \sum_{t=1}^T v_t$
-

Acknowledgments

This work is supported in part by AFOSR Grant FA9550-12-1-0453 and DARPA Grant FA9550-12-1-0406 negotiated by AFOSR, ONR Grant N00014-10-1-0746, NINDS RO1-NS041811, RO1-NS48598, and U24-NS063930-03, and the Mirowski Discovery Fund for Epilepsy Research.

Appendix A. HMM sum-product algorithm

Consider a hidden Markov model of a sequence $y_{1:T}$ with corresponding discrete states $z_{1:T}$, each of which takes one of K values. The joint probability of $y_{1:t}$ and $z_t = \ell$ is

$$p(y_{1:t}, z_t = \ell) = p(y_t | z_t = \ell) \sum_{k=1}^K p(y_{1:t-1}, z_{t-1} = k) p(z_t = \ell | z_{t-1} = k), \quad (\text{A.1})$$

sometimes called a *forward message*, which depends on a recursive call for $y_{1:t-1}$ and $z_{t-1} = k$ with

$$p(y_1, z_1 = k) = p(y_1 | z_1 = k) p(z_1 = k). \quad (\text{A.2})$$

Calculating the marginal likelihood $p(y_{1:T})$ simply involves one last marginalization over z_T ,

$$p(y_{1:T}) = \sum_{k=1}^K p(y_{1:T}, z_T = k). \quad (\text{A.3})$$

Algorithm 3 HMM backward-filtering forward-sampling algorithm for block-sampling $z_{1:T}$

- 1: let π_0 and $\pi = (\pi_1, \dots, \pi_K)^T$ be the initial and transition distributions
 - 2: let $\mathbf{u}_t \in \mathbb{R}_+^K$ be the likelihoods $p(y_t | z_t = k)$ for each k
 - 3: let $\zeta_t \in \mathbb{R}_+^K$ be the backward messages $p(y_{t+1:T} | z_t = k)$ at t for each k
 - 4: normalize each \mathbf{u}_t to sum to 1, preventing underflow during computations
 - 5: **for** $t = 1, \dots, T$ **do**
 - 6: $\tilde{\mathbf{u}}_t \leftarrow \mathbf{u}_t / (\mathbf{1}^T \mathbf{u}_t)$
 - 7: **end for**
 - 8: calculate backward messages over all time points
 - 9: $\tilde{\zeta}_T = \mathbf{1}$
 - 10: **for** $t = T - 1, \dots, 1$ **do**
 - 11: transmit messages backward : $\tau_{t+1} \leftarrow \tilde{\mathbf{u}}_{t+1} \circ \tilde{\zeta}_{t+1}, \zeta_t \leftarrow \pi \tau_{t+1}$
 - 12: normalize : $\tilde{\zeta}_t \leftarrow \zeta_t / (\mathbf{1}^T \zeta_t)$
 - 13: **end for**
 - 14: $\tau_1 \leftarrow \tilde{\mathbf{u}}_1 \circ \tilde{\zeta}_1$
 - 15: sample first time point
 - 16: $\mathbf{q}_1 \leftarrow \pi_0 \circ \tau_1, \mathbf{q}_1^{\sim} \leftarrow \mathbf{q}_1 / (\mathbf{1}^T \mathbf{q}_1)$
 - 17: $z_1 \sim \mathbf{q}_1^{\sim}$
 - 18: sample other time points
 - 19: **for** $t = 2, \dots, T$ **do**
 - 20: $\mathbf{q}_t \leftarrow \pi_{z_{t-1}} \circ \tau_t, \mathbf{q}_t^{\sim} \leftarrow \mathbf{q}_t / (\mathbf{1}^T \mathbf{q}_t)$
 - 21: $z_t \sim \mathbf{q}_t^{\sim}$
 - 22: **end for**
-

Algorithm 2 provides a numerically stable recipe for calculating this marginal likelihood.

We can sample the states $z_{1:T}$ from their joint distribution, also known as block sampling, via a similar recursive formulation. The conditional likelihood of the last $T - t$ samples given the state at t is

$$p(y_{t+1:T} | z_t = \ell) = \sum_{k=1}^K p(y_{t+1} | z_{t+1} = k) p(y_{t+2:T} | z_{t+1} = k) p(z_{t+1} = k | z_t = \ell), \quad (\text{A.4})$$

which depends recursively on the *backward messages* $p(y_{t+2:T} | z_{t+1} = k)$ for each $k \in \{1, \dots, K\}$. To sample $z_{1:T}$ at once we use the joint posterior distribution of the entire state sequence $z_{1:T}$, which factors into

$$p(z_{1:T} | y_{1:T}) = p(z_T | z_{T-1}, y_{1:T}) p(z_{T-1} | z_{T-2}, y_{1:T}) \cdots p(z_2 | z_1, y_{1:T}) p(z_1 | y_{1:T}) \quad (\text{A.5})$$

If we first sample z_1 , we can condition on it to then sample z_2 and continue in this fashion until we finish with z_T . The posterior for z_t is the product of the backward message, the likelihood of y_t given z_t , and the probability of z_t given z_{t-1} ,

$$p(z_t | y_{1:T}) \propto p(y_{t+1:T} | z_t) p(y_t | z_t) p(z_t | z_{t-1}), \quad (\text{A.6})$$

where $p(z_t | z_{t-1}) \triangleq p(z_1)$ for $t = 1$. A numerical stable recipe for this backward-filtering forward-sampling is given in Algorithm 3.

Appendix B. Individual channel variables posterior

Sampling the variables associated with the individual channel i involves first sampling active features $\mathbf{f}^{(i)}$ (while marginalizing $z_{1:T}^{(i)}$), then conditioning on these feature assignments $\mathbf{f}^{(i)}$ to block sample the state sequence $z_{1:T}^{(i)}$, and finally sampling the transition distribution $\pi^{(i)}$ given the feature indicators $\mathbf{f}^{(i)}$ and state sequence $z_{1:T}^{(i)}$. Explicit algorithms for this sampling are given in Wulsin [26, Section 4.2.1].

Channel marginal likelihood

Let $\mathbf{i}' \subseteq \{1, \dots, N\}$ index the neighboring channels in the graph upon which channel i is conditioned. The conditional likelihood of observation $y_t^{(i)}$ under AR model k given the neighboring observations $y_t^{(\mathbf{i}')}$ at time t is

$$p\left(y_t^{(i)} | \tilde{\mathbf{y}}_t^{(\mathbf{i}')}, \mathbf{y}_t^{(\mathbf{i}')}, z_t^{(i)} = k, \mathbf{z}_t^{(\mathbf{i}')}, Z_t, \{\mathbf{a}_k\}, \{\Delta_t\}\right) \propto N(\tilde{\mu}_t, \tilde{\sigma}_t^2) \quad (\text{B.1})$$

for

$$\begin{aligned}\tilde{\mu}_t &= \mathbf{a}_k^T \tilde{\mathbf{y}}_t^{(i)} + \Delta_{z_t}^{(i,i')} \Delta_{z_t}^{-1(i',i')} \left(\mathbf{y}_t^{(i')} - \mathbf{A}_{z^{(i')}} \tilde{\mathbf{Y}}_t^{(i')} \right), \\ \tilde{\sigma}_t^2 &= \Delta_{z_t}^{(i,i)} - \Delta_{z_t}^{(i,i')} \Delta_{z_t}^{-1(i',i')} \Delta_{z_t}^{(i',i)},\end{aligned}\quad (\text{B.2})$$

which follows from the conditional distribution of the multivariate normal [27, pg. 579]. Using the forward-filtering scheme (see Algorithm 2) to marginalize over the exponentially many state sequences $z_{1:T}^{(i)}$, we can calculate the channel marginal likelihood,

$$p \left(\mathbf{y}_{1:T}^{(i)} | \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right), \quad (\text{B.3})$$

of channel i 's observations over all $t = 1, \dots, T$ given the observations $\mathbf{y}_{1:T}^{(i')}$ and the assigned states $z_{1:T}^{(i')}$ of neighboring channels i' and given the event state sequence $Z_{1:T}$. As previously discussed, taking the non-zero elements of the infinite-dimensional transition distributions $\boldsymbol{\pi}^{(i)}$, derived from $\mathbf{f}^{(i)}$ and $\boldsymbol{\eta}^{(i)}$ as in Eq. (9), yields a set of $K^{(i)}$ -dimensional active feature transition distributions $\boldsymbol{\pi} = \{\boldsymbol{\pi}_j\}$, reducing this marginalization to a series of matrix-vector products.

Sampling active features, $\mathbf{f}^{(i)}$

We briefly describe the active feature sampling scheme given in detail by Fox et al. [5]. Recall that for our HIW-spatial BP-AR-HMM, we need to condition on neighboring channel state sequences $z_{1:T}^{(i')}$ and event state sequences $Z_{1:T}$. Sampling the feature indicators $\mathbf{f}^{(i)}$ for channel i via the Indian buffet process (IBP) involves considering those features shared by other channels and those unique to channel i . Let $K_+ = \sum_{k=1}^K \mathbf{1} \left(f_k^{(1)} \vee \dots \vee f_k^{(N)} \right)$ denote the total number of active features used by at least one of the channels. We consider the set of shared features across channels not including those specific to channel i as $\mathcal{S}^{(-i)} \subseteq \{1, \dots, K_+\}$ and the set of unique features for channel i as $U^{(i)} \subseteq \{1, \dots, K_+\} / \mathcal{S}^{(-i)}$.

Shared features

The posterior for each shared feature $k \in \mathcal{S}^{(-i)}$ for channel i is given by

$$\begin{aligned}p \left(f_k^{(i)} | \mathbf{y}_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{F}^{-ik}, \mathbf{f}_{k' \neq k}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right) &\propto \\ p \left(f_k^{(i)} | \mathbf{F}^{-ik} \right) p \left(\mathbf{y}_{1:T}^{(i)} | \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{f}_{k' \neq k}^{(i)}, f_k^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right),\end{aligned}\quad (\text{B.4})$$

where the marginal likelihood of $\mathbf{y}_{1:T}^{(i)}$ term (see Eq. B.3) follows from the sum-product algorithm. Recalling the form of the IBP posterior predictive distribution, we have

$p \left(f_k^{(i)} = 1 | \mathbf{F}^{-ik} \right) = m_k^{(-i)} / N$, where $m_k^{(-i)}$ denotes the number of other channels that use feature k . We use this posterior to formulate a Metropolis-Hastings proposal that flips the current indicator value $f_k^{(i)}$ to its complement $\bar{f}_k^{(i)}$ with probability, $\rho(\bar{f}_k^{(i)} | f_k^{(i)})$,

$$f_k^{(i)} = \begin{cases} \bar{f}_k^{(i)}, & \text{w.p. } \rho(\bar{f}_k^{(i)} | f_k^{(i)}), \\ f_k^{(i)}, & \text{otherwise,} \end{cases} \quad (\text{B.5})$$

where

$$\rho(\bar{f}_k^{(i)} | f_k^{(i)}) = \min \left(\frac{p(\bar{f}_k^{(i)} | y_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{F}^{-ik}, \mathbf{f}_{k' \neq k}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\})}{p(f_k^{(i)} | y_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, Z_{1:T}, \mathbf{F}^{-ik}, \mathbf{f}_{k' \neq k}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\})}, 1 \right). \quad (\text{B.6})$$

Unique features

We either propose a new feature or remove a unique feature for channel i using a birth and death reversible jump MCMC sampler [28, 29, 30] (see Fox et al. [16] for details relevant to the BP-AR-HMM). We denote the number of unique features for channel i as $n^{(i)} = |\mathcal{U}^{(i)}|$.

We define the vector of shared feature indicators as $\mathbf{f}_-^{(i)} = \mathbf{f}_{k' | k' \in \mathcal{S}^{(-i)}}^{(i)}$ and that for unique feature indicators as $\mathbf{f}_+^{(i)} = \mathbf{f}_{k' | k' \in \mathcal{U}^{(i)}}^{(i)}$, which together $(\mathbf{f}_-^{(i)}, \mathbf{f}_+^{(i)})$ define the full feature indicator vector $\mathbf{f}^{(i)}$ for channel i . Similarly, \mathbf{a}_+ and $\boldsymbol{\eta}_+$ describe the model dynamics and transition parameters associated with these unique features. We propose a new unique feature vector \mathbf{f}'_+ and corresponding model dynamics \mathbf{a}'_+ transition parameters $\boldsymbol{\eta}'_+$ (sampled from their priors in the case of feature birth) with a proposal distribution of

$$q(\mathbf{f}'_+, \mathbf{a}'_+, \boldsymbol{\eta}'_+ | \mathbf{f}_+^{(i)}, \{\mathbf{a}_k\}_+^{(i)}, \boldsymbol{\eta}_+^{(i)}) = q(\mathbf{f}'_+ | \mathbf{f}_+^{(i)}) q(\mathbf{a}'_+ | \mathbf{f}'_+, \mathbf{f}_+^{(i)}, \{\mathbf{a}_k\}_+^{(i)}) q(\boldsymbol{\eta}'_+ | \mathbf{f}'_+, \mathbf{f}_+^{(i)}, \boldsymbol{\eta}_+^{(i)}). \quad (\text{B.7})$$

A new unique feature is proposed with probability 0.5 and each existing unique feature is removed with probability $0.5/n^{(i)}$. This proposal is accepted with probability

$$\rho(\mathbf{f}'_+, \mathbf{a}'_+, \boldsymbol{\eta}'_+ | \mathbf{f}_+^{(i)}, \{\mathbf{a}_k\}_+^{(i)}, \boldsymbol{\eta}_+^{(i)}) = \min \left(\frac{p(y_{1:T}^{(i)} | \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, [\mathbf{f}_-^{(i)} \mathbf{f}'_+], \boldsymbol{\eta}^{(i)}, \boldsymbol{\eta}'_+, \{\mathbf{a}_k\}, \{\Delta_l\})}{p(y_{1:T}^{(i)} | \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, [\mathbf{f}_-^{(i)} \mathbf{f}_+], \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\})}, \frac{\text{Poisson}(n'_i | \alpha_c / N) q(\mathbf{f}'_+ | \mathbf{f}'_+)}{\text{Poisson}(n_i | \alpha_c / N) q(\mathbf{f}'_+ | \mathbf{f}_+^{(i)})}, 1 \right). \quad (\text{B.8})$$

Channel state sequence, $z_{1:T}^{(i)}$

We block sample the state sequence for all the time points of channel i , given that channel's feature-constrained transition distributions $\pi^{(i)}$, the state parameters $\{\mathbf{a}_k\}$, the observations

$y_{1:T}^{(i)}$ and the neighboring observations $\mathbf{y}_{1:T}^{(i')}$ and current states $\mathbf{z}_{1:T}^{(i')}$. The joint probability of the state sequence $z_{1:T}^{(i)}$ is given by

$$p \left(z_{1:T}^{(i)} | y_{1:T}^{(i)}, \mathbf{y}_{1:T}^{(i')}, \mathbf{z}_{1:T}^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right) = p \left(z_1^{(i)} | y_1^{(i)}, \mathbf{y}_1^{(i')}, \mathbf{z}_1^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right) \cdot \prod_{t=2}^T p \left(z_t^{(i)} | y_{t:T}^{(i)}, \mathbf{y}_{t:T}^{(i')}, z_{t-1}^{(i)}, \mathbf{z}_{t:T}^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right). \quad (\text{B.9})$$

Again following the backward filtering forward sampling scheme (Algorithm 3), at each time point t we sample state $z_t^{(i)}$ conditioned on $z_{t-1}^{(i)}$ by marginalizing $z_{t+1:T}^{(i)}$. The conditional probability of $z_t^{(i)}$ is given by

$$p \left(z_t^{(i)} | y_{t:T}^{(i)}, \mathbf{y}_{t:T}^{(i')}, z_{t-1}^{(i)}, Z_{1:T}, \mathbf{z}_{t:T}^{(i')}, \mathbf{f}^{(i)}, \boldsymbol{\eta}^{(i)}, \{\mathbf{a}_k\}, \{\Delta_l\} \right) \propto \text{Multi} \left(\tilde{\boldsymbol{\pi}}_{z_{t-1}^{(i)}}^{(i)} \circ \mathbf{u}_t^{(i)} \circ \boldsymbol{\psi}_t \right), \quad (\text{B.10})$$

where $\tilde{\boldsymbol{\pi}}_{z_{t-1}^{(i)}}^{(i)}$ is the transition distribution given the assigned state at $t-1$, $\mathbf{u}_t^{(i)} \in \mathbb{R}^{K^{(i)}}$ is the vector of likelihoods under each possible state at time t (as in Eq. (B.1)), and $\boldsymbol{\psi}_t \in \mathbb{R}^{K^{(i)}}$ is the vector of backwards messages (see Eq. (A.4)) from time point $t+1$ to t .

Channel transition parameters, $\boldsymbol{\eta}^{(i)}$

Following [20, Supplement], the posterior for the transition variable $\eta_{jk}^{(i)}$ is given by

$$p(\eta_{jk}^{(i)} | z_{1:T}^{(i)}, f_k^{(i)}) \propto \frac{(\eta_{jk}^{(i)})^{n_{jk}^{(i)} + \gamma_c + \kappa_c \delta(j,k) - 1} e^{\eta_{jk}^{(i)}}}{\sum_{k' | f_k^{(i)} = 1} \eta_{jk'}^{(i)}}, \quad (\text{B.11})$$

where $\eta_{jk}^{(i)}$ denotes the number of times channel i transitions from state j to state k . We can sample from this posterior via two auxiliary variables,

$$\begin{aligned} \bar{\boldsymbol{\eta}}_j^{(i)} &\sim \text{Dir}(\gamma_c + \kappa_c \mathbf{e}_j + \mathbf{n}_j^{(i)}) \\ C_j^{(i)} &\sim \text{Gamma}(K\gamma_c + \kappa_c, 1) \\ \boldsymbol{\eta}_j^{(i)} &= C_j^{(i)} \bar{\boldsymbol{\eta}}_j^{(i)}. \end{aligned} \quad (\text{B.12})$$

Appendix C. Event state variables posterior

Since we model the event state process with a (truncated approximation to the) HDP-HMM, inference is more straightforward than with the channel states. We block sample the event state sequence $Z_{1:T}$ and then sample the event state transition distributions ϕ .

Event marginal likelihood

Let \mathbf{z}_t denote the vector of N channel states at time t . Since the space of \mathbf{z}_t is exponentially large, we cannot integrate it out to compute the marginal conditional likelihood of the data given the event state sequence $Z_{1:T}$ (and model parameters). Instead, we consider the conditional likelihood of an observation at time t given channel states \mathbf{z}_t and event state $Z_t = l$,

$$p(\mathbf{y}_t | \tilde{\mathbf{Y}}_t, \mathbf{z}_t, Z_t=l, \{\mathbf{a}_k\}, \Delta_l) \propto N(\mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t, \Delta_l). \quad (\text{C.1})$$

Recalling Eq. (3), we see that this conditional likelihood of \mathbf{y}_t is equivalent to a zero-mean multivariate normal model on the channel innovations $\boldsymbol{\varepsilon}_t$,

$$p(\boldsymbol{\varepsilon}_t | Z_t=l, \Delta_l) \propto N(\mathbf{0}, \Delta_l).$$

As with the channel marginal likelihoods, we use the forward-filtering algorithm (see Algorithm 2) to marginalize over the possible event state sequences $Z_{1:T}$, yielding a likelihood conditional on the channel states \mathbf{z}_t and autoregressive parameters $\{\mathbf{a}_k\}$, in addition to the event transition distribution ϕ and event state covariances $\{\Delta_l\}$,

$$p(\mathbf{y}_{1:T} | \mathbf{z}_{1:T}, \boldsymbol{\varphi}, \{\mathbf{a}_k\}, \{\Delta_l\}). \quad (\text{C.2})$$

Event state sequence, $Z_{1:T}$

The mechanics of sampling the event state sequence $Z_{1:T}$ directly parallel those of sampling the individual channel state sequences $z_{1:T}^{(i)}$. The joint probability of the event state sequence is given by

$$p(Z_{1:T} | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, \boldsymbol{\varphi}, \{\mathbf{a}_k\}, \{\Delta_l\}) \propto p(Z_1 | \mathbf{y}_1, \mathbf{z}_1, \boldsymbol{\varphi}, \{\mathbf{a}_k\}, \{\Delta_l\}) \prod_{t=2}^T p(Z_t | \mathbf{y}_{t:T}, \mathbf{z}_{t:T}, Z_{t-1}, \boldsymbol{\varphi}, \{\mathbf{a}_k\}, \{\Delta_l\}). \quad (\text{C.3})$$

We again use the backward filtering forward sampling scheme of the sumproduct algorithm to block sample each event state whose conditional probability distribution over the L states is given by

$$p(Z_t | \tilde{\mathbf{Y}}_t, \mathbf{z}_t, \boldsymbol{\varphi}, \{\mathbf{a}_k\}, \{\Delta_l\}) \propto \text{Multi}(\boldsymbol{\varphi}_{Z_{t-1}} \circ \mathbf{v}_t \circ \boldsymbol{\psi}_t), \quad (\text{C.4})$$

where $\phi_{Z_{t-1}}$ the transition distribution given the assigned state at $t - 1$, $\mathbf{v}_t \in \mathbb{R}^L$ is the vector of likelihoods under each of the L possible states at time t (as in Eq. (C.1)), and $\boldsymbol{\psi}_t \in \mathbb{R}^L$ is again the vector of backwards messages from time point $t + 1$ to t . \circ denotes element-wise product.

Event transition parameters, $\boldsymbol{\phi}$

The Dirichlet posterior for the event state l 's transition distribution $\boldsymbol{\phi}_l$ simply involves the transition counts \mathbf{n}_l from event state l to all L states,

$$p(\boldsymbol{\phi}_l | Z_{1:T}, \boldsymbol{\beta}) \propto \text{Dir}(\alpha_e \boldsymbol{\beta} + \mathbf{e}_l \kappa_e + \mathbf{n}_l), \quad (\text{C.5})$$

for global weights $\boldsymbol{\beta}$, concentration parameter α_e , and self-transition parameter κ_e .

Global transition parameters, $\boldsymbol{\beta}$

The Dirichlet posterior of the global transition distribution $\boldsymbol{\beta}$ involves the auxiliary variables $(m_{\cdot 1}, \dots, m_{\cdot L})$,

$$p(\boldsymbol{\beta} | Z_{1:T}) \propto \text{Dir}(\gamma_e / L + \overline{m}_{\cdot 1}, \dots, \gamma_e / L + \overline{m}_{\cdot L}), \quad (\text{C.6})$$

where these auxiliary variables are defined as

$$\begin{aligned} \overline{m}_{ll'} &= \begin{cases} m_{ll'}, & l \neq l' \\ m_{ll} - w_l, & l = l' \end{cases} \\ m_{ll'} &= \sum_{r=1}^{n_{ll'}} \theta_r \\ \theta_r &\sim \text{Ber} \left(\frac{\alpha_e \beta_l + \kappa_e \delta(l, l')}{\alpha_e \beta_l + \delta(l, l') + r} \right), \\ w_l &\sim \text{Bin} \left(m_{ll'}, \frac{\rho_e}{\rho_e + \beta_l (1 - \rho_e)} \right) \end{aligned} \quad (\text{C.7})$$

and $m_{\cdot l'} = \sum_l m_{ll'}$. See Fox [31, Appendix A] for full derivations.

Appendix D. Hyperparameters posterior

Below we give brief descriptions for the MCMC sampling of the hyperparameters in our model. Full derivations are given in Fox [31, Section 5.2.3, Appendix C].

Channel dynamics model hyperparameters, γ_c, κ_c

We use Metropolis-Hastings steps to propose a new value γ'_c from gamma distributions with fixed variance $\sigma_{\gamma_c}^2$ and accept with probability $\min(r(\gamma'_c | \gamma_c), 1)$,

$$\begin{aligned} r(\gamma'_c | \gamma_c) &= \frac{p(\{\boldsymbol{\pi}^{(i)}\} | \gamma'_c, \kappa, \{\mathbf{f}^{(i)}\}) p(\gamma'_c | \gamma_c^2 / \sigma_{\gamma_c}^2, \gamma_c / \sigma_{\gamma_c}^2) p(\gamma_c | \gamma'_c, \sigma_{\gamma_c}^2)}{p(\{\boldsymbol{\pi}^{(i)}\} | \gamma_c, \kappa, \{\mathbf{f}^{(i)}\}) q(\gamma_c | \gamma_c^2 / \sigma_{\gamma_c}^2, \gamma_c / \sigma_{\gamma_c}^2) q(\gamma'_c | \gamma_c, \sigma_{\gamma_c}^2)} \\ &= \frac{p(\{\boldsymbol{\pi}^{(i)}\} | \gamma'_c, \kappa, \{\mathbf{f}^{(i)}\})}{p(\{\boldsymbol{\pi}^{(i)}\} | \gamma_c, \kappa, \{\mathbf{f}^{(i)}\})} \frac{\Gamma(\nu) \gamma_c^{\nu - \nu - a}}{\Gamma(\nu') \gamma_c^{\nu - \nu' - a}} \exp \left(-b(\gamma'_c - \gamma_c) \sigma_{\gamma_c}^{2(\nu - \nu')} \right), \end{aligned} \quad (\text{D.1})$$

where $\nu = \gamma_c^2 / \sigma_{\gamma_c}^2$, $\nu' = \gamma_c'^2 / \sigma_{\gamma_c'}^2$, and we have a Gamma(a, b) prior on γ_c . The likelihood term $p(\{\boldsymbol{\pi}^{(i)}\} | \gamma_c', \kappa, \{\mathbf{f}^{(i)}\})$ follows from the Dirichlet distribution and is given by

$$p(\{\boldsymbol{\pi}^{(i)}\} | \gamma_c', \kappa, \{\mathbf{f}^{(i)}\}) = \prod_{i=1}^N \prod_{k=i}^{K^{(i)}} \left(\frac{\Gamma(\gamma_c' K^{(i)} + \kappa_c)}{\left(\prod_{k'}^{K^{(i)}-1} \Gamma(\gamma_c')\right) \Gamma(\gamma_c' + \kappa_c)} \prod_{j=1}^{K^{(i)}} \left(\tilde{\pi}_{kj}^{(i)}\right)^{\gamma_c' + \kappa_c \delta(k,j) - 1} \right), \quad (\text{D.2})$$

where that for $p(\{\boldsymbol{\pi}^{(i)}\} | \gamma_c, \kappa, \{\mathbf{f}^{(i)}\})$ is similar. Recall that the transition parameters $\{\boldsymbol{\pi}^{(i)}\}$ are independent over i , and thus their likelihoods multiply. The proposal and acceptance ratio for κ_c is similar.

Channel active features model hyperparameter α_c

We place a Gamma($a_{\alpha_c}, b_{\alpha_c}$) prior on α_c , which implies a gamma posterior of the form

$$p(\alpha_c | \{\mathbf{f}^{(i)}\}) \propto \text{Gamma}(a_{\alpha_c} + K_+, b_{\alpha_c} + \sum_{i=1}^N (1/i)), \quad (\text{D.3})$$

where $K_+ = \sum_{k=1}^K \mathbb{1}(f_k^{(1)} \vee \dots \vee f_k^{(N)})$ denotes the number of channel states activated in at least one of the channels.

Event dynamics model hyperparameters, $\boldsymbol{\gamma}_e, \boldsymbol{\alpha}_e, \boldsymbol{\kappa}_e, \boldsymbol{\rho}_e$

Instead of sampling α_e and κ_e independently, we add an additional parameter $\rho_e = \kappa_e / (\alpha_e + \kappa_e)$ and sample $(\alpha_e + \kappa_e)$ and ρ_e , which is simpler than sampling α_e and κ_e independently.

$(\boldsymbol{\alpha}_e + \boldsymbol{\kappa}_e)$ With a Gamma($a_{\alpha_e + \kappa_e}, b_{\alpha_e + \kappa_e}$) prior on $(\alpha_e + \kappa_e)$, we use auxiliary variables $\{r_l\}_{l=1}^L$ and $\{s_l\}_{l=1}^L$ to define the posterior,

$$p(\alpha_e + \kappa_e | Z_{1:T}) \propto \text{Gamma} \left(a_{\alpha_e + \kappa_e} + \bar{m}_{..} - \sum_{l=1}^L s_l, b_{\alpha_e + \kappa_e} - \sum_{l=1}^L \log r_l \right), \quad (\text{D.4})$$

where $\bar{m}_{..} = \sum_{l,l'=1}^L \bar{m}_{ll'}$ is the sum over auxiliary variables $m_{ll'}$ defined in Eq. (C.7), and the auxiliary variables r_l and s_l are sampled as

$$\begin{aligned} r_l &\sim \text{Beta}(\alpha + \kappa + 1, n_l), \\ s_l &\sim \text{Ber}(n_l / (n_l + \alpha + \kappa)). \end{aligned}$$

$\boldsymbol{\rho}_e$ With a Beta(c_{ρ_e}, d_{ρ_e}) prior on ρ_e , we use auxiliary variables $\{w_l\}_{l=1}^L$ to define the posterior,

$$p(\rho_e | \bar{\mathbf{m}}, \boldsymbol{\beta}) \propto \text{Beta} \left(c_{\rho_e} + \sum_{l=1}^L w_l, d_{\rho_e} + \bar{m}_{..} - \sum_{l=1}^L w_l \right). \quad (\text{D.5})$$

For $w_{ls} \sim \text{Ber}(\rho)$ over $s = 1, \dots, m_{ll}$, the posterior of the auxiliary variable w_l is

$$p(w_l | \bar{m}_{ll}, \beta_l) \propto \text{Bin}(\bar{m}_{ll}, \rho_e + \beta_l(1 - \rho_e)) \quad (\text{D.6})$$

γ_e With a $\text{Gamma}(a_{\gamma_e}, b_{\gamma_e})$ prior on γ_e , we use auxiliary variables v and q to define the posterior,

$$p(\gamma_e | m) \propto \text{Gamma} \left(a_{\gamma_e} - q + \sum_{l=1}^L \mathbf{1}(\bar{m}_{..l} > 0), b_{\gamma_e} - \log v \right). \quad (\text{D.7})$$

The auxiliary variables are sampled via

$$\begin{aligned} v &\sim \text{Beta}(\gamma_e + 1, \bar{m}_{..}), \\ q &\sim \text{Ber}(\bar{m}_{..} / (\gamma_e + \bar{m}_{..})). \end{aligned}$$

Appendix E. Autoregressive state coefficients posterior

Recall that our prior on the autoregressive coefficients \mathbf{a}_k is a multivariate normal with zero mean and covariance Σ_0 ,

$$\begin{aligned} p(\mathbf{a}_k | \Sigma_0) &\propto \mathbf{N}(\mathbf{0}, \Sigma_0) \\ \log p(\mathbf{a}_k | \Sigma_0) &\propto -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k. \end{aligned} \quad (\text{E.1})$$

The conditional event likelihood given the channel states $\mathbf{z}_{1:T}$ and the event states $Z_{1:T}$ is

$$\begin{aligned} p(\mathbf{y}_{1:T} | \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_k\}, \{\Delta_l\}) &\propto \prod_{t=1}^T \mathbf{N}(\mathbf{y}_t; \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t, \Delta_{z_t}) \\ \log p(\mathbf{y}_{1:T} | \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_k\}, \{\Delta_l\}) &\propto -\frac{1}{2} \sum_{t=1}^T (\mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t)^T \Delta_{z_t}^{-1} (\mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t). \end{aligned} \quad (\text{E.2})$$

The product of these prior and likelihood terms is the joint distribution over \mathbf{a}_k and $\mathbf{y}_{1:T}$,

$$\begin{aligned} \log p(\mathbf{a}_k, \mathbf{y}_{1:T} | \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) &\propto \\ -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k - \frac{1}{2} \sum_{t=1}^T (\mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t)^T \Delta_{z_t}^{-1} (\mathbf{y}_t - \mathbf{A}_{\mathbf{z}_t} \tilde{\mathbf{Y}}_t). \end{aligned} \quad (\text{E.3})$$

We take a brief tangent to prove a useful identity,

Lemma Appendix E.1

Let the column vector $\mathbf{x} \in \mathbb{R}^m$ and the symmetric matrix $A \in \mathbb{S}^{m \times m}$ be defined as

$$\mathbf{x} = \begin{bmatrix} \mathbf{y} \\ \mathbf{z} \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} B & C \\ C^T & D \end{bmatrix},$$

where $\mathbf{y} \in \mathbb{R}^p$, $\mathbf{z} \in \mathbb{R}^q$, $B \in \mathbf{S}^{p \times p}$, $D \in \mathbf{S}^{q \times q}$, $C \in \mathbb{R}^{p \times q}$, and $m = p + q$. Then

$$\mathbf{x}^T A \mathbf{x} = \mathbf{y}^T B \mathbf{y} + \mathbf{z}^T D \mathbf{z} + 2 \mathbf{y}^T C \mathbf{z}. \quad (\text{E.4})$$

Proof

$$\begin{aligned} \mathbf{x}^T A \mathbf{x} &= \begin{bmatrix} \mathbf{y}^T & \mathbf{z}^T \end{bmatrix} \begin{bmatrix} B & C \\ C^T & D \end{bmatrix} \begin{bmatrix} \mathbf{y} \\ \mathbf{z} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{y}^T & \mathbf{z}^T \end{bmatrix} \begin{bmatrix} B \mathbf{y} + C \mathbf{z} \\ C^T \mathbf{y} + D \mathbf{z} \end{bmatrix} \\ &= \mathbf{y}^T B \mathbf{y} + \mathbf{y}^T C \mathbf{z} + \mathbf{z}^T C^T \mathbf{y} + \mathbf{z}^T D \mathbf{z} \\ &= \mathbf{y}^T B \mathbf{y} + \mathbf{z}^T D \mathbf{z} + 2 \mathbf{y}^T C \mathbf{z} \end{aligned}$$

Note that this identity also holds for any permutation p applied to the rows of \mathbf{x} and the rows and columns of A . We now can manipulate the likelihood term of Eq. (E.3) into a form that separates \mathbf{a}_k from $\mathbf{a}_{k' \neq k}$. Suppose that \mathbf{k}^+ denotes the indices of the N channels where $z_t^{(i)} = k$ and $\mathbf{k}^- = \{1, \dots, N\} / \mathbf{k}^+$ denotes those for whom $z_t^{(i)} \neq k$. Furthermore, we use the superscript indexing on these two sets of indices to select the corresponding portions of the \mathbf{y}_t vector and the \mathbf{A}_{z_t} , $\tilde{\mathbf{Y}}_t$, and $\Delta_{z_t}^{-1}$ matrices. We start by decomposing the likelihood term into three parts,

$$\begin{aligned} &(\mathbf{y}_t - \mathbf{A}_{z_t} \tilde{\mathbf{Y}}_t)^T \Delta_{z_t}^{-1} (\mathbf{y}_t - \mathbf{A}_{z_t} \tilde{\mathbf{Y}}_t) \propto \\ &\left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^+, \mathbf{k}^+)} \left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right) + \\ &2 \left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^+, \mathbf{k}^-)} \left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right) + \\ &\left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^-, \mathbf{k}^-)} \left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right), \end{aligned} \quad (\text{E.5})$$

which we then insert into our previous expression for the joint distribution of \mathbf{a}_k and $\mathbf{y}_{1:T}$ (Eq. (E.3)),

$$\begin{aligned} \log p(\mathbf{a}_k, \mathbf{y}_{1:T} | \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) &\propto -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k - \\ &\frac{1}{2} \sum_{t=1}^T \left\{ \left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^+, \mathbf{k}^+)} \left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right) + \right. \\ &2 \left(\mathbf{y}_t^{(\mathbf{k}^+)} - \mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^+, \mathbf{k}^-)} \left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right) + \\ &\left. \left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right)^T \Delta_{z_t}^{-1(\mathbf{k}^-, \mathbf{k}^-)} \left(\mathbf{y}_t^{(\mathbf{k}^-)} - \mathbf{A}_{z_t}^{(\mathbf{k}^-, \mathbf{k}^-)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^-, \mathbf{k}^-)} \right) \right\}. \end{aligned} \quad (\text{E.6})$$

Conditioning on $\mathbf{y}_{1:T}$ allows us to absorb the third term of the sum into the proportionality, and after replacing $\mathbf{A}_{z_t}^{(\mathbf{k}^+, \mathbf{k}^+)} \tilde{\mathbf{Y}}_t^{(\mathbf{k}^+, \mathbf{k}^+)}$ with a more explicit expression, we have

$$\begin{aligned}
\log p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) &\propto -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k - \\
&\frac{1}{2} \sum_{t=1}^T \left\{ \left(\mathbf{y}_t^{(k^+)} - \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right)^T \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \right. \\
&\left. \left(\mathbf{y}_t^{(k^+)} - \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right) + \right. \\
&2 \left(\mathbf{y}_t^{(k^+)} - \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right)^T \left(\Delta_{Z_t}^{-1(k^+, k^-)} \right) \right. \\
&\left. \left(\mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)} \right) \right\}, \tag{E.7}
\end{aligned}$$

which we can further expand to yield

$$\begin{aligned}
\log p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) &\propto -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k - \\
&\frac{1}{2} \sum_{t=1}^T \left\{ \left(\mathbf{y}_t^{(k^+)} \right)^T \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\mathbf{y}_t^{(k^+)} \right) + \right. \\
&\left(\mathbf{a}_k^T \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right] \right) \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right) - \\
&2 \left(\mathbf{y}_t^{(k^+)} \right)^T \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right) \left. \right\} - \\
&\sum_{t=1}^T \left\{ \mathbf{y}_t^{(k^+)} \Delta_{Z_t}^{-1(k^+, k^-)} \left(\mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)} \right) - \right. \\
&\left. \left(\mathbf{a}_k^T \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right] \right) \left(\Delta_{Z_t}^{-1(k^+, k^-)} \right) \left(\mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)} \right) \right\}. \tag{E.8}
\end{aligned}$$

Absorbing more terms unrelated to \mathbf{a}_k into the proportionality, we have

$$\begin{aligned}
\log p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) &\propto -\frac{1}{2} \mathbf{a}_k^T \Sigma_0^{-1} \mathbf{a}_k - \\
&\frac{1}{2} \sum_{t=1}^T \left\{ \left(\mathbf{a}_k^T \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right] \right) \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right) - \right. \\
&2 \left(\mathbf{y}_t^{(k^+)} \right)^T \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right]^T \mathbf{a}_k \right) \left. \right\} - \\
&\sum_{t=1}^T \left\{ - \left(\mathbf{a}_k^T \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k^+|}^+)} \right] \right) \left(\Delta_{Z_t}^{-1(k^+, k^-)} \right) \left(\mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)} \right) \right\}, \tag{E.9}
\end{aligned}$$

which after some rearranging gives

$$\begin{aligned}
& \log p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) \propto \\
& -\frac{1}{2} \mathbf{a}_k^T \left\{ \sum_0^{-1} + \sum_{t=1}^T \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k|}^+)} \right] \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k|}^+)} \right]^T \right\} \mathbf{a}_k + \\
& \mathbf{a}_k^T \sum_{t=1}^T \left\{ \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k|}^+)} \right] \left(\Delta_{Z_t}^{-1(k^+, k^+)} \right) \left(\mathbf{y}_t^{(k^+)} \right) + \right. \\
& \left. \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k|}^+)} \right] \left(\Delta_{Z_t}^{-1(k^+, k^-)} \right) \left(\mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)} \right) \right\}. \tag{E.10}
\end{aligned}$$

Before completing the square, we will find it useful to introduce a bit more notation to simplify the expression,

$$\tilde{\mathbf{Y}}_t^{(k^+)} = \left[\tilde{\mathbf{y}}_t^{(k_1^+)} | \dots | \tilde{\mathbf{y}}_t^{(k_{|k|}^+)} \right], \quad \boldsymbol{\varepsilon}_t^{(k^-)} = \mathbf{y}_t^{(k^-)} - \mathbf{A}_{Z_t}^{(k^-, k^-)} \tilde{\mathbf{Y}}_t^{(k^-, k^-)},$$

yielding

$$\begin{aligned}
& \log p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_l\}) \propto \\
& -\frac{1}{2} \mathbf{a}_k^T \left\{ \sum_0^{-1} + \sum_{t=1}^T \tilde{\mathbf{Y}}_t^{(k^+)} \Delta_{Z_t}^{-1(k^+, k^+)} \tilde{\mathbf{Y}}_t^{T(k^+)} \right\} \mathbf{a}_k + \\
& \mathbf{a}_k^T \left\{ \sum_{t=1}^T \tilde{\mathbf{Y}}_t^{(k^+)} \left(\Delta_{Z_t}^{-1(k^+, k^+)} \mathbf{y}_t^{(k^+)} + \Delta_{Z_t}^{-1(k^+, k^-)} \boldsymbol{\varepsilon}_t^{(k^-)} \right) \right\}. \tag{E.11}
\end{aligned}$$

We desire an expression in the $-\frac{1}{2}(\mathbf{a}_k - \boldsymbol{\mu}_k)^T \sum_k^{-1} (\mathbf{a}_k - \boldsymbol{\mu}_k)$ for unknown $\boldsymbol{\mu}_k$ and \sum_k^{-1} so that it conforms to the multivariate normal density with mean $\boldsymbol{\mu}_k$ and precision \sum_k^{-1} . We already have our \sum_k^{-1} value from the quadratic term above,

$$\sum_k^{-1} = \sum_0^{-1} + \sum_{t=1}^T \tilde{\mathbf{Y}}_t^{(k^+)} \Delta_{Z_t}^{-1(k^+, k^+)} \tilde{\mathbf{Y}}_t^{T(k^+)}, \tag{E.12}$$

which allows us to solve the cross-term for $\boldsymbol{\mu}_k$,

$$\begin{aligned}
& -\frac{1}{2} (-2 \boldsymbol{\mu}_k^T \sum_k^{-1} \mathbf{a}_k) = \mathbf{a}_k^T \left(\sum_{t=1}^T \tilde{\mathbf{Y}}_t^{(k^+)} \left(\Delta_{Z_t}^{-1(k^+, k^+)} \mathbf{y}_t^{(k^+)} + \Delta_{Z_t}^{-1(k^+, k^-)} \boldsymbol{\varepsilon}_t^{(k^-)} \right) \right), \\
& \sum_k^{-1} \boldsymbol{\mu}_k = \sum_{t=1}^T \tilde{\mathbf{Y}}_t^{(k^+)} \left(\Delta_{Z_t}^{-1(k^+, k^+)} \mathbf{y}_t^{(k^+)} + \Delta_{Z_t}^{-1(k^+, k^-)} \boldsymbol{\varepsilon}_t^{(k^-)} \right). \tag{E.13}
\end{aligned}$$

We can pull the final required $-2 \boldsymbol{\mu}_k^T \sum_k^{-1} \mathbf{a}_k$ term from the proportionality and complete the square. Thus, we have the form of the posterior for \mathbf{a}_k ,

$$p(\mathbf{a}_k | \mathbf{y}_{1:T}, \mathbf{z}_{1:T}, Z_{1:T}, \{\mathbf{a}_{k'}\}_{k' \neq k}, \{\Delta_t\}) \propto \exp\left(-\frac{1}{2}(\mathbf{a}_k - \boldsymbol{\mu}_k)^T \sum_k^{-1} (\mathbf{a}_k - \boldsymbol{\mu}_k)\right) \propto N(\boldsymbol{\mu}_k, \sum_k), \quad (\text{E.14})$$

where

$$\sum_k^{-1} = \sum_0^{-1} + \sum_{t=1}^T \overline{\mathbf{Y}}_t^{(\mathbf{k}^+)} \Delta_{Z_t}^{-1}(\mathbf{k}^+, \mathbf{k}^+) \overline{\mathbf{Y}}_t^{T(\mathbf{k}^+)} \quad (\text{E.15})$$

$$\sum_k^{-1} \boldsymbol{\mu}_k = \sum_{t=1}^T \overline{\mathbf{Y}}_t^{(\mathbf{k}^+)} \left(\Delta_{Z_t}^{-1}(\mathbf{k}^+, \mathbf{k}^+) \mathbf{y}_t^{(\mathbf{k}^+)} + \Delta_{Z_t}^{-1}(\mathbf{k}^+, \mathbf{k}^-) \varepsilon_t^{(\mathbf{k}^-)} \right).$$

Appendix F. Experiment Parameters

Below, we give explicit values for the various priors and parameters used in our experiments.

Table F.2

Parameters used in sparse factorial BP-AR-HMM simulation experiment

parameter	description	value
N	number of time series per event	6
r	AR model order	1
\mathbf{m}_0	\mathbf{a}_k N prior mean	0
Σ_0	\mathbf{a}_k N prior covariance	$0.1 \cdot I_{1 \times 1}$
L	truncated number of event states	20
b_0	$_1$ HIW prior degrees of freedom	$N + 3$
D_0	$_1$ HIW prior scale	$0.05(b_0 - N - 1)(I_{N \times N} + 1)$
$(a_{a_e + \kappa_e}, b_{a_e + \kappa_e})$	$a_e + \kappa_e$ gamma prior	(1,1)
$(a_{\gamma_e}, b_{\gamma_e})$	γ_e gamma prior	(1,1)
(c_{ρ_e}, d_{ρ_e})	ρ_e beta prior	(1,1)
$(a_{\gamma_c}, b_{\gamma_c})$	γ_c gamma prior	(1,1)
$(a_{\kappa_c}, b_{\kappa_c})$	κ_c gamma prior	(1000,1)
$\sigma_{\gamma_c}^2$	γ_c Metropolis-Hastings proposal variance	1
$\sigma_{\kappa_c}^2$	κ_c Metropolis-Hastings proposal variance	100
(a_{a_c}, b_{a_c})	a_c gamma prior	(1,1)

Table F.3

Parameters used in epileptic seizures and bursts experiments. When applicable, the same parameters were used for the standard BP-AR-HMM as in the correlated BP-AR-HMMs. The analysis of two two seizures involved 16 iEEG channels, and the analysis of the 15 bursts and single seizure involved 6 iEEG channels.

parameter	description	value
N	number of time series per event	16 and 6
r	AR model order	5
\mathbf{m}_0	\mathbf{a}_k N prior mean	0
Σ_0	\mathbf{a}_k N prior covariance	$\text{Cov}(\{y_t^{(i)}\}_{\forall t,i})$
L	truncated number of event states	30
b_0	l (H)IW prior degrees of freedom	$N + 3$
D_0	l (H)IW prior scale	$(b_0 - N - 1) \cdot \text{Cov}(\{y_{t+1} - y_t\}_{\forall t})$
$(a_{\alpha_e + \kappa_e}, b_{\alpha_e + \kappa_e})$	$\alpha_e + \kappa_e$ gamma prior	(1,1)
$(a_{\gamma_e}, b_{\gamma_e})$	γ_e gamma prior	(1,1)
(c_{ρ_e}, d_{ρ_e})	ρ_e beta prior	(1,1)
$(a_{\gamma_c}, b_{\gamma_c})$	γ_c gamma prior	(1,1)
$(a_{\kappa_c}, b_{\kappa_c})$	κ_c gamma prior	(1000,1)
$\sigma_{\gamma_c}^2$	γ_c Metropolis-Hastings proposal variance	1
$\sigma_{\kappa_c}^2$	Metropolis-Hastings proposal variance	100
(a_{a_c}, b_{a_c})	gamma prior	(1,1)

References

1. D'Alessandro M, Vachtsevanos G, Esteller R, Echaz J, Cranstoun S, Worrell G, Parish L, Litt B. A multi-feature and multi-channel univariate selection process for seizure prediction. *Clinical Neurophysiology*. 2005; 116:506–516. [PubMed: 15721064]
2. Mirowski P, Madhavan D, Lecun Y, Kuzniecky R. Classification of patterns of EEG synchronization for seizure prediction. *Clinical Neurophysiology*. 2009; 120:1927–1940. [PubMed: 19837629]
3. Litt B, Esteller R, Echaz J, D'Alessandro M, Shor R, Henry T, Pennell P, Epstein C, Bakay R, Dichter M, Vachtsevanos G. Epileptic seizures may begin hours in advance of clinical onset: a report of five patients. *Neuron*. 2001; 30:51–64. [PubMed: 11343644]
4. Krystal AD, Prado R, West M. New methods of time series analysis of non-stationary EEG data: eigenstructure decompositions of time varying autoregressions. *Clinical Neurophysiology*. 1999; 110:2197–2206. [PubMed: 10616127]
5. Fox EB, Sudderth EB, Jordan MI, Willsky AS. Sharing features among dynamical systems with beta processes. *Advances in Neural Information Processing Systems*. 2009
6. Schindler K, Leung H, Elger CE, Lehnertz K. Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial EEG. *Brain*. 2007; 130:65–77. [PubMed: 17082199]
7. Schiff SJ, Sauer T, Kumar R, Weinstein SL. Neuronal spatiotemporal pattern discrimination: the dynamical evolution of seizures. *NeuroImage*. 2005; 28:1043–1055. [PubMed: 16198127]

8. Prado R, Molina F, Huerta G. Multivariate time series modeling and classification via hierarchical VAR mixtures. *Computational Statistics & Data Analysis*. 2006; 51:1445–1462.
9. Van Gael, J.; Saatici, Y.; Teh, YW.; Ghahramani, Z. Beam sampling for the infinite hidden Markov model. *Proceedings of the 25th International Conference on Machine Learning*; 2008.
10. Heller, K.; Teh, YW.; Gorur, D. The Infinite Hierarchical Hidden Markov Model. *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*; 2009.
11. Doshi-Velez, F.; Wingate, D.; Tenenbaum, J.; Roy, N. Infinite dynamic bayesian networks. *Proceedings of the 28th International Conference on Machine Learning*; 2011.
12. Dong, W.; Pentland, A.; Heller, KA. Graph-Coupled HMMs for Modeling the Spread of Infection. *Proceedings of the 28th Conference Conference on Uncertainty in Artificial Intelligence*; 2012.
13. Wulsin, DF.; Fox, EB.; Litt, B. Parsing Epileptic Events Using a Markov Switching Process for Correlated Time Series. *Proceedings of the 30th International Conference on Machine Learning*; 2013.
14. Ghahramani Z, Jordan MI. Factorial hidden Markov models. *Machine Learning*. 1997; 29:245–273.
15. Dawid AP, Lauritzen SL. Hyper Markov laws in the statistical analysis of decomposable graphical models. *The Annals of Statistics*. 1993; 21:1272–1317.
16. Fox EB, Sudderth EB, Jordan MI, Willsky AS. Joint Modeling of Multiple Related Time Series via the Beta Process. *Arxiv preprint arXiv:1111.4226v1*. 2011
17. Thibaux, R.; Jordan, MI. Hierarchical beta processes and the Indian buffet process. *Proceedings of the 10th Conference on Artificial Intelligence and Statistics*; 2007.
18. Fox EB, Sudderth EB, Jordan MI, Willsky AS. A sticky HDP-HMM with application to speaker diarization. *The Annals of Applied Statistics*. 2011; 5:1020–1056.
19. Griffiths TL, Ghahramani Z. Infinite latent feature models and the Indian buffet process. *Gatsby Computational Neuroscience Unit*. 2005 Technical Report #2005-001.
20. Hughes MC, Fox EB, Sudderth EB. Effective split-merge Monte Carlo methods for nonparametric models of sequential data. *Advances in Neural Information Processing Systems*. 2012
21. Ishwaran H, Zarepour M. Exact and approximate sum representations for the Dirichlet process. *Canadian Journal of Statistics*. 2002; 30:269–283.
22. Teh YW, Jordan MI, Beal MJ, Blei DM. Hierarchical Dirichlet processes. *Journal of the American Statistical Association*. 2006; 101:1566–1581.
23. Carvalho CM, Massam H, West M. Simulation of hyper-inverse Wishart distributions in graphical models. *Biometrika*. 2007; 94:647–659.
24. Fox EB, Sudderth EB, Jordan MI, Willsky AS. Bayesian nonparametric inference of switching dynamic linear models. *IEEE Transactions on Signal Processing*. 2011; 59:1569–1585.
25. Esteller, R.; Echauz, J.; Tchong, T.; Litt, B.; Pless, B. Line length: an efficient feature for seizure onset detection. *Proceedings of the 23rd Engineering in Medicine and Biology Society Conference*; 2001.
26. Wulsin, DF. PhD thesis. University of Pennsylvania; 2013. Bayesian Nonparametric Modeling of Epileptic Events.
27. Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DS. *Bayesian Data Analysis*. second ed. Chapman & Hall/CRC; Boca Raton, Florida: 2004.
28. Green PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*. 1995; 82:711–732.
29. Richardson S, Green PJ. On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society Series B*. 1997; 59:731–792.
30. Brooks, SP.; Gelman, A.; Jones, GL.; Meng, X., editors. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC; Boca Raton, Florida: 2011.
31. Fox, EB. PhD thesis. Massachusetts Institute of Technology; 2009. Bayesian Nonparametric Learning of Complex Dynamical Phenomena.




Figure 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.




Figure 2.

Graphical model of our factorial AR-HMM. The N channel states $z_t^{(i)}$ evolve according to independent Markov processes (transition distributions omitted for simplicity) and index the AR dynamic parameters a_k used in generating observation $y_t^{(i)}$. The Markov-evolving event state Z_t indexes the graph-structured covariance Δ_l of the correlated AR innovations resulting in multivariate observations $y_t = [y_t^{(1)}, \dots, y_t^{(N)}]^T$ sharing the same conditional independencies.




Figure 3. Referencing the channel state and event state sequences of Figure 2, here we depict the graphical model associated with our Bayesian nonparametric prior specification of Section 2.2. **(top)** The channel i feature indicators $\mathbf{f}^{(i)}$ are samples from a Bernoulli process with weights $\{\omega_k\}$ and constrain the channel transition distributions $\pi^{(i)}$. Channel states $z_t^{(i)}$ evolve independently for each channel according to these feature-constrained transition distributions $\pi^{(i)}$. **(bottom)** The event state Z_t evolves independently of each channel i 's state $z_t^{(i)}$ according to transition distributions $\{\phi_l\}$, which are coupled by global transition distribution β .

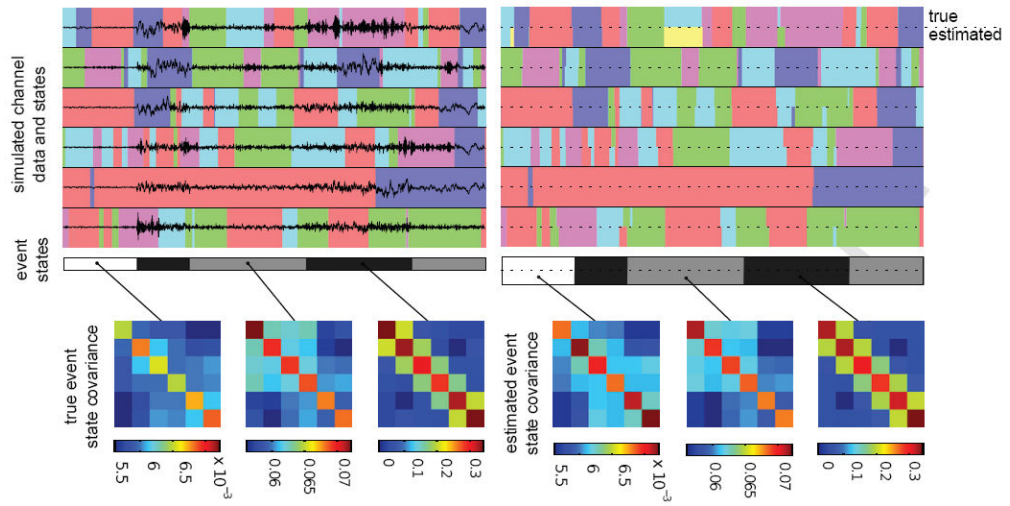


Figure 4. (top left) The six simulated channel time series overlaid on the five true channel states, denoted by different colors; the three true event states are shown in grayscale in the bar below. (top right) The true and estimated channel (color) and event (grayscale) states shown below for comparison after 6000 MCMC iterations. The true (bottom left) and estimated (bottom right) event state innovation covariances.

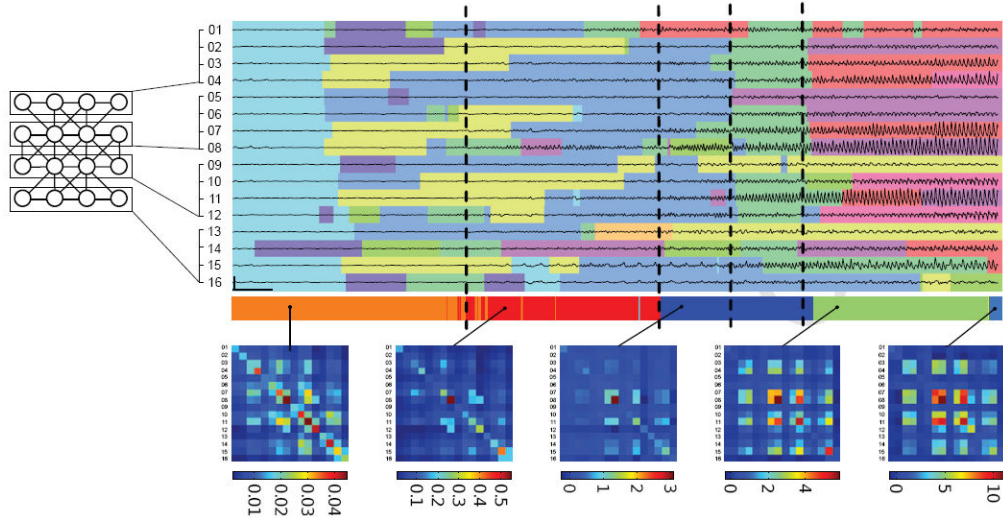


Figure 5. The graph used for a 16 channel iEEG electrode and corresponding traces over 25 seconds of a seizure onset with colors indicating the inferred channel states. The event states are shown below along with the associated innovation covariances. Vertical dashed lines indicate the EEG transition times marked independently by an epileptologist. Vertical and horizontal scale bars denote 1 mV and 1 second, respectively.




Figure 6.

(left) An example 16-channel clip of iEEG with the middle section of one channel zoomed in and innovations from the original BP-AR-HMM and our sparse factorial BP-AR-HMM shown below. (right) Boxplots of the heldout event log-likelihoods from the two original and factorial models with mean and median posterior likelihood given in green and red lines. Boxes denote the middle 50% prediction interval.




Figure 7. The average time per MCMC iteration required to calculate all of the channel likelihoods under each AR model at each time point.

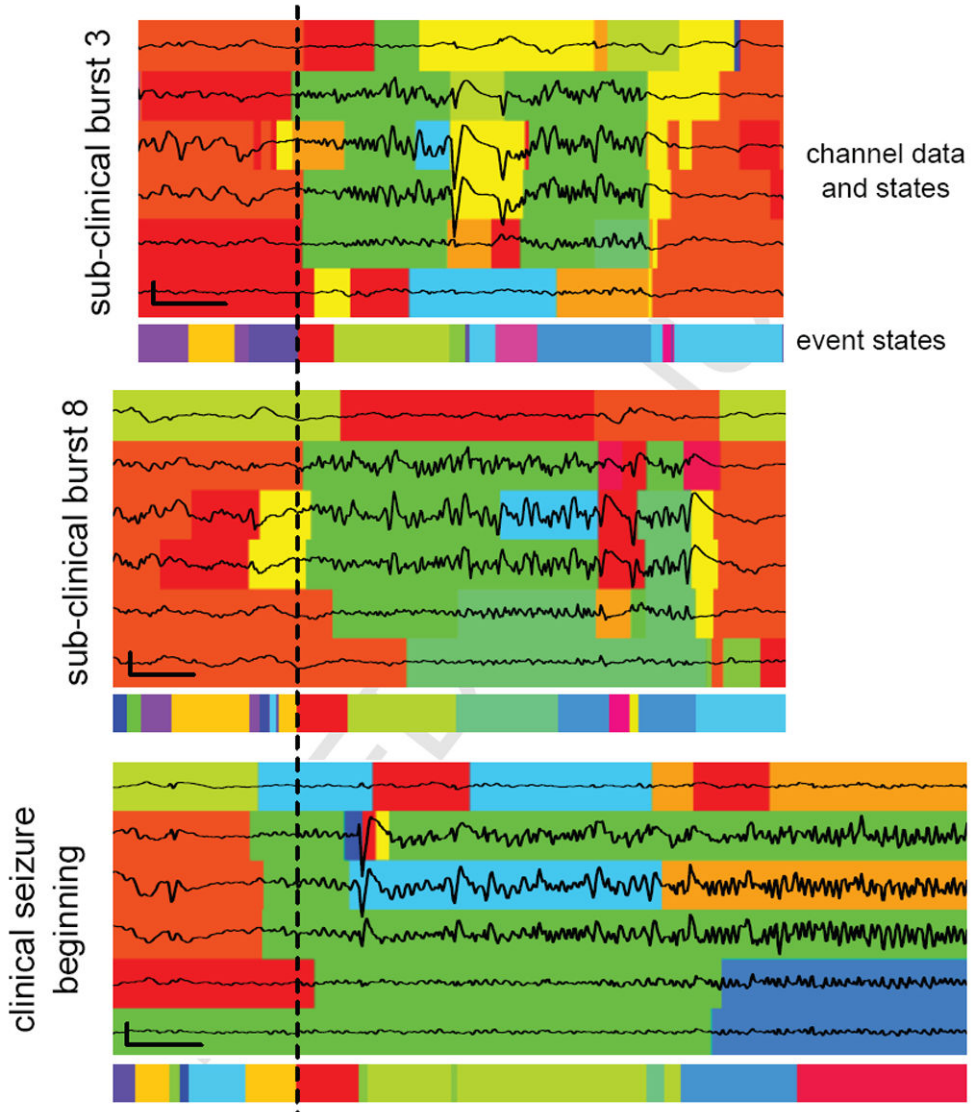


Figure 8. 6 iEEG traces from two sub-clinical bursts and onset of the single seizure with colors indicating inferred channel and event states. The dashed lines indicates the start of the red state in the three events. Vertical and horizontal scale bars denote 1mv and 1 second, respectively.

Table 1

The true and estimated values for the channel state coefficients in the simulated dataset. We include the posterior mean and 95% credible interval.

channel state	true a_k	post. a_k mean	post. a_k 95% interval
1	-0.900	-0.906	[-0.917, -0.896]
2	-0.450	-0.456	[-0.474, -0.436]
3	0	-0.009	[-0.038, 0.020]
4	0.450	0.445	[0.425, 0.466]
5	0.900	0.902	[0.890, 0.913]