Published in final edited form as:

Ann Appl Stat. 2021 September ; 15(3): 1171–1193. doi:10.1214/21-AOAS1455.

Identifying the Recurrence of Sleep Apnea Using A Harmonic Hidden Markov Model

Beniamino Hadj-Amar¹, Bärbel Finkenstädt¹, Mark Fiecas², Robert Huckstepp³

Beniamino Hadj-Amar: Beniamino.Hadj-Amar@rice.edu; Bärbel Finkenstädt: B.F.Finkenstadt@warwick.ac.uk; Mark Fiecas: mfeicas@umn.edu; Robert Huckstepp: R.Huckstepp@warwick.ac.uk

¹Department of Statistics, University of Warwick

²School of Public Health, Division of Biostatistics, University of Minnesota

³School of Life Sciences, University of Warwick

Abstract

We propose to model time-varying periodic and oscillatory processes by means of a hidden Markov model where the states are defined through the spectral properties of a periodic regime. The number of states is unknown along with the relevant periodicities, the role and number of which may vary across states. We address this inference problem by a Bayesian nonparametric hidden Markov model assuming a sticky hierarchical Dirichlet process for the switching dynamics between different states while the periodicities characterizing each state are explored by means of a trans-dimensional Markov chain Monte Carlo sampling step. We develop the full Bayesian inference algorithm and illustrate the use of our proposed methodology for different simulation studies as well as an application related to respiratory research which focuses on the detection of apnea instances in human breathing traces.

Keywords and phrases

Sleep Apnea; Time-Varying Frequencies; Reversible-Jump MCMC; Bayesian Non-parametrics; Hierarchical Dirichlet process

1 Introduction

Statistical methodology for identifying periodicities in time series can provide meaningful information about the underlying physical process. Non-stationary behavior seems to be the norm rather than the exception for physiological time series as timevarying periodicities and other forms of rich dynamical patterns are commonly observed in response to external perturbations and pathological states. For example, body temperature and rest activity might exhibit changes in their periodic patterns as an individual experiences a disruption in its circadian timing system (Krauchi and Wirz-Justice, 1994; Komarzynski et al., 2018). Heart rate variability and electroencephalography are other examples of data that are often characterized by time-changing spectral properties, the quantification of which can provide

¹GEM is an abbreviation for Griffiths, Engen and McCloskey, see Ignatov (1982); Perman, Pitman and Yor (1992); Pitman (1996) for background.

valuable information about the well-being of a subject (Malik, 1996; West, Prado and Krystal, 1999; Cohen, 2014; Bruce et al., 2018). This paper is motivated by modeling airflow trace data obtained from a sleep apnea study, where our objective is to identify and model the recurrence of different periodicities, which are indicative of the apneic and hyponeic events.

1.1 A Case Study on Sleep Apnea in Humans

Our study focuses on *sleep apnea* (Heinzer et al., 2015), a chronic respiratory disorder characterized by recurrent episodes of temporary (≥ 2 breaths) cessation of breathing during sleep (about 10 seconds in humans). Sleep apnea negatively affects several organ systems, such as the heart and kidneys in the long term. It is also associated with an increased likelihood of hypertension, stroke, several types of dementia, cardiovascular diseases, daytime sleepiness, depression and a diminished quality of life (Ancoli-Israel et al., 1991; Teran-Santos et al., 1999; Peker et al., 2002; Young, Peppard and Gottlieb, 2002; Yaggi et al., 2005; Cooke et al., 2009; Dewan, Nieto and Somers, 2015). Instances of sleep apnea can be subclassified based on the degree of reduction in airflow to the lungs whereby *apneas* are classified as a reduction of airflow by 90% and *hypopneas* require a reduction in airflow by at least 30% (with a reduction of blood oxygen levels by at least 3%). For example, the airflow trace shown in Figure 1 was collected from a human over a time span of 5.5 minutes of continuous breathing. During this time, apneic and hyponeic events were simulated; apneas appear in the first and second minute and around the start of the fifth minute, where there are two instances of hypopneas in the first half of the second minute and at the start of the fourth minute as marked in Figure 1. Note these events were classified by eye by an experienced experimental researcher. Detecting apneic and hyponeic events during sleep is one of the primary interests of researchers and clinicians working in the field of sleep medicine and relevant healthcare (Berry et al., 2017). Manual classification is a time-consuming process, and hence there is a need of a data-driven approach for the automated classification of these types of events.

1.2 Hidden Markov Models and Spectral Analysis

Approaches to spectral analysis of nonstationary processes were first developed by Priestley (1965) who introduced the concept of *evolutionary spectra*, namely spectral density functions that are time-dependent as well as localized in the frequency domain. This modeling framework was formalized as a class of nonstationary time series called *locally stationary* (Dahlhaus et al., 1997). Locally stationary processes can be well approximated by *piecewise stationary* processes and several authors proposed to model the time-varying spectra of locally stationary time series through the piecewise constant spectra of the corresponding stationary segments (Adak, 1998; Ombao et al., 2001; Davis, Lee and Rodriguez-Yam, 2006). This framework was extended to a Bayesian setting by Rosen, Stoffer and Wood (2009) and Rosen, Wood and Stoffer (2012) who estimated a time-varying spectral density using a fixed number of smoothing splines and approximated the likelihood function via a product of local Whittle likelihoods (Whittle, 1957). Their methodology is based on the assumption that the time series are piecewise stationary, and the underlying spectral density for each partition is smooth over frequencies. In order to deal with changes in spectral densities with sharp peaks which can be observed for some physiological data

sets such as respiratory data, Hadj-Amar et al. (2020) proposed a change-point analysis where they introduced a Bayesian methodology for inferring changepoints along with the number and values of the periodicities affecting each segment. While these approaches allow us to analyse the spectral changing properties of a process from a retrospective and exploratory point of view, in order to develop a more comprehensive understanding of the process driving the data, further modeling assumptions are needed that quantify the probabilistic rules governing the transitions as well as recurrence of different oscillatory dynamic patterns. For example, in the context of experimental sleep apnea research, both, correctly classifying the states of apnea as well as quantifying their risk of recurrence, possibly in the context of real-time monitoring of patients, is of major interest to the development of treatments for breathing disorders.

Here, we address the switching dynamics between different oscillatory states in the framework of a hidden Markov model (HMM) that assumes a discrete latent state sequence whose transition probabilities follow a Markovian structure (see e.g. Rabiner (1989); Ephraim and Merhav (2002); Cappé, Moulines and Rydén (2005)). Conditioned on the state sequence, the observations are assumed to be independent and generated from a family of probability distributions, which hereafter we refer to as the *emission distributions*. HMMs are arguably among the most popular statistical approaches used for modeling time series data when the observations exhibit nonstationary characteristics that can be represented by an underlying and unobserved hidden process. These modeling approaches, also known as hidden Markov processes and Markov-switching models, became notable by the work of Baum and Petrie (1966) and Baum and Eagon (1967), and HMMs have since been successfully used in many different applications (Krogh et al., 1994; Yau et al., 2011; Langrock et al., 2013; Yaghouby and Sunderam, 2015; Huang et al., 2018).

As we are interested in modeling the recurrence of periodicities in the airflow trace data we propose a harmonic HMM where the discrete latent state sequence reflects the time-varying changes as well as recurrence of periodic regimes as defined by their spectral properties. Furthermore, we pursue a flexible nonparametric specification within a Bayesian approach by assuming the infinite-state hierarchical Dirichlet process (HDP) as a building block (Teh et al., 2006). The HDP-HMM approach places a Dirichlet process (DP) prior on the Markovian transition probabilities of the system, while allowing the atoms associated with the statespecific conditional DPs to be shared between each other, yielding an HMM with a countably infinite number of states. The HDP-HMM therefore not only provides a nonparametric specification of the transition distributions but also removes the need for specifying a priori the number of states. In our case study, while it is true that by looking at Figure 1 there seems to be a substantial difference between the apnea-states and normal breathing (i.e. neither apnea or hypopnea), it is also conceivable that normal breathing may exhibit many distinct periodic patterns, both with respect to this subject and possibly across different individuals. Yet we are interested in classifying, in an unsupervised fashion, the occurrence of apnea/hypopnea states, it might also be necessary to model states corresponding to normal breathing and other aspects that characterize respiration, such as a sigh for example. Paz and West (2013) reported at least 13 forms of breathing patterns including forms of apnea, further justifying the need to not pre-specify the number of hidden states and hence using a more versatile HDP for this application. We focus on the *sticky*

HDP-HMM by Fox et al. (2011), where an additional parameter is introduced to promote self-transition with the effect that the sticky HDP-HMM more realistically explains the switching dynamics between states that exhibit some temporal mode persistence. We hence extend the Bayesian methodology for the sticky HDP-HMM to a spectral representation of the states where inference for the variable dimensionality regarding the number of periodicities that characterize the emission distributions of the states is achieved by developing an appropriate form of trans-dimensional MCMC sampling step (Green, 1995).

This article introduces a dynamic oscillatory model that is remarkably flexible while being developed in a framework that is still computationally accessible. To the best of our knowledge, it is the first statistical methodology that exploits an HMM for analyzing the spectral properties of a time series while quantifying the probabilistic mechanism governing the transitions and recurrence of distinct dynamic patterns. The rest of the paper is organized as follows. Section 2 presents the model and the general framework of our Bayesian approach. Section 3 and 4 provide the inference scheme and simulation studies to show the performance of the proposed method. In Section 5, we illustrate the use of our approach to detect instances of apnea in human breathing traces.

2 A Sticky HDP-HMM with Oscillatory Emissions

We propose a Bayesian approach relevant for analyzing observations of oscillatory dynamical systems based on an HMM. The observational model follows Andrieu and Doucet (1999) and Hadj-Amar et al. (2020), where the state-dependent data generating process is expressed via a Gaussian harmonic regression with an unknown number of periodicities. We integrate this methodology with the nonparametric HDP-HMM model introduced by Teh et al. (2006), where the rows of the infinite-dimensional transition matrix are framed to enable linkage between the probabilities associated with each hidden state in a hierarchical manner. Temporal mode persistence of the latent state sequence is achieved using the sticky HDP-HMM formulation of Fox et al. (2011).

Let $\mathbf{y} = (y_1, \dots, y_T)'$ be a realization of a time series whose oscillatory behavior may switch dynamically over time and let $\mathbf{z} = (z_1, \dots, z_T)'$ denote the hidden discrete-valued states of the Markov chain that characterize the different periodic regimes, where z_t denotes the state of the Markov chain at time t . Any observation y_t given the state z_t is assumed to be conditionally independent of the observations and states at other time steps (Rabiner, 1989). Here, a highly flexible nonparametric approach is postulated by assuming that the state space is unbounded, i.e. has infinitely many states as in Beal, Ghahramani and Rasmussen (2002) and Teh et al. (2006). Thus, the Markovian structure on the state sequence $\mathbf{z}$ is given by

$$z_t | z_{t-1}, (\pi_j)_{j=1}^{\infty} \sim \pi_{z_t-1}, \quad t = 1, \dots, T, \quad (1)$$

where $\boldsymbol{\pi}_j = (\pi_{j1}, \pi_{j2}, \dots)$ represents the (countably infinite) state-specific vector of transition probabilities, and in particular $\pi_{jk} = p(z_t = k | z_{t-1} = j)$, where $p(\cdot)$ is used as a generic notation for probability density or mass function, whichever appropriate. We assume that the initial state has distribution $\boldsymbol{\pi}_0 = (\pi_{01}, \pi_{02}, \dots)$, namely $z_0 \sim \boldsymbol{\pi}_0$.

Next, assume that each state j represents a periodic regime that is characterized by d_j relevant periodicities whose frequencies are denoted by $\omega_j = (\omega_{j1}, \dots, \omega_{jd_j})'$, recalling that periodicity is the inverse of frequency. Let $\beta_j = (\beta'_{j1}, \dots, \beta'_{jd_j})'$ be the vector of linear coefficients that can be associated with the amplitude and phase corresponding to each frequency ω_{jl} that is of relevance to state j , where $\beta_{jl} = (\beta_{jl}^{(1)}, \beta_{jl}^{(2)})'$ and $l = 1, \dots, d_j$. Furthermore, let us define $\theta_j = (d_j, \omega'_j, \beta'_j, \sigma_j^2)'$, where σ_j^2 accounts for a state-specific variance. Then, each observation is assumed to be generated from the following emission distribution

$$y_t | z_t = j, (\theta_j)_{j=1}^\infty \sim \mathcal{N}(f_{tj}, \sigma_j^2), \quad t = 1, \dots, T, \quad (2)$$

where the mean function f_{tj} for state j at time t is (Andrieu and Doucet, 1999; Hadj-Amar et al., 2020) specified to be oscillatory, i.e.,

$$f_{tj} = x_t(\omega_j)' \beta_j, \quad (3)$$

and the vector of basis functions $x_t(\omega_j)$ is defined as

$$x_t(\omega_j) = (\cos(2\pi\omega_{j1}t), \sin(2\pi\omega_{j1}t), \dots, \cos(2\pi\omega_{jd_j}t), \sin(2\pi\omega_{jd_j}t))'. \quad (4)$$

The dimension of each oscillatory function depends on the unknown number d_j of periodicities relevant to state j . Given a pre-fixed upper bound for the number of relevant periodicities per state, $d_{\max}$, the parameter space Θ_j for the vector of emission parameters θ_j can be written as $\Theta_j = \cup_{d_j=1}^{d_{\max}} \{d_j\} \times \{\mathbb{R}^{2d_j} \times \Omega_{d_j} \times \mathbb{R}^+\}$, where $\Omega_{d_j} = (0, 0.5)^{d_j}$ denotes the sample space for the frequencies of the j -th state. Hadj-Amar et al. (2020) introduced this modeling approach for oscillatory data that show regime shifts in periodicity, amplitude and phase. They assume that, conditional on an unknown number of change-points at unknown positions, the time series process can be approximated by a sequence of segments, each with mean functions specified by Gaussian harmonic models of the form given in Equation (3). Here this approach will be combined with a nonparametric sticky HDP-HMM model (Fox et al., 2011) which provides a structure for modeling switching dynamics and connectivity between different states.

2.1 A Bayesian Nonparametric Framework for Unbounded Markov States

Dirichlet processes provide a simple description of clustering processes where the number of clusters is not fixed a priori. Suitably extended to a hierarchical DP, this form of stochastic process provides a foundation for the design of state-space models in which the number of modes is random and inferred from the data. In contrast to classic methods that assume a parametric prior on the number of states, or use model selection techniques to determine the number of regimes in an HMM, where here we follow Beal, Ghahramani and Rasmussen (2002); Teh et al. (2006) and Fox et al. (2011), and assume the number of states to be unknown. We therefore do not need to pre-specify the number of hidden states, which provides a more flexible modeling framework. The DP may be used in frameworks where an element of the model is a discrete random variable of unknown cardinality (Hjort et al.,

2010). The unbounded HMM (i.e., where the number of possible states is unknown) can be seen as an infinite mixture model, where the mixing proportions are modelled as DPs (Beal, Ghahramani and Rasmussen, 2002; Rasmussen and Ghahramani, 2002; Teh et al., 2006).

The current state z_t indexes a specific transition distribution π_{z_t} over the positive integers, whose probabilities are the mixing proportions for the choice of the next state z_{t+1} . To allow the same set of next states to be reachable from each of the current states, we introduce a set of state-specific DPs, whose atoms are shared between each other (Teh et al., 2006). As in Fox et al. (2011) we implement the sticky version by increasing the expected probability of self-transitions. In particular, the state-specific transition distribution π_j follows the HDP

$$\pi_j \mid \eta, \kappa, \alpha \sim \text{DP}\left(\eta + \kappa, \frac{\eta\alpha + \kappa\delta_j}{\eta + \kappa}\right), \quad (5)$$

where

$$\alpha \mid \gamma \sim \text{GEM}(\gamma).$$

Here, the sequence $\alpha = (\alpha_k)_{k=1}^{\infty}$ can be seen as a *global* probability distribution over the positive integers that ties together the transition distributions π_j and guarantees that they have the same support. We denote by $\text{GEM}(\gamma)$ ¹ the *stick-breaking construction* (Sethuraman, 1994; Pitman, 2002) of α as

$$\alpha_k = \nu_k \prod_{l=1}^{k-1} (1 - \nu_l), \quad (6)$$

where

$$\nu_k \mid \gamma \sim \text{Beta}(1, \gamma), \quad (7)$$

for $k = 1, 2, \dots$, and γ is a positive real number that controls the expected value of the number of elements in α with significant probability mass. Equations (6) and (7) can be motivated by the equivalent process where a stick of length one is split into lengths specified by the weights α_k , where the k^{th} proportion is a random fraction ν_k of the remaining stick after the preceding $(k-1)$ proportions have been constructed. The stick-breaking construction ensures that the sequence α satisfies $\sum_{k=1}^{\infty} \alpha_k = 1$ with probability one.

Conditional on α , the hierarchical structure given in Equation (5) indicates that the statespecific transition distributions π_j are distributed according to a DP with *concentration parameter* $\eta + \kappa$ and *base distribution* $(\eta\alpha + \kappa\delta_j)/(\eta + \kappa)$, that is itself a DP. Here, η is a positive real number that controls the variability of the π_j 's around α , while κ is a positive real number that inflates the expected probability of a self-transition (Fox et al., 2011), and δ_j denotes a unit-mass measure concentrated at j . By setting $\kappa = 0$ in Equation (5), we obtain the non-sticky HDP-HMM framework proposed by Teh et al. (2006). It was noted that this specification could result in an unrealistically rapid alternation between different (and often redundant) states. The *sticky* formulation of Fox et al. (2011) allows for more

temporal state persistence by inflating the expected probabilities of self-transitions by an amount proportional to κ , i.e.

$$\mathbb{E}[\pi_{jk} | \eta, \kappa, \alpha] = \frac{\eta}{\eta + \kappa} \alpha_k + \frac{\kappa}{\eta + \kappa} \delta(j, k),$$

where $\delta(j, k) = 1$ if $k = j$ and zero otherwise. Though the sticky parameter parameter increases temporal persistence of the hidden state sequence, several simulations studies (see e.g. Fox et al. 2011) suggest that the sticky HDP-HMM is still capable to capture and identify states that are short in duration by inferring a small probability of self-transition.

3 Inference

Our inference scheme is formulated within a full Bayesian framework, where our proposed sampler alternates between updating the emission and the HMM parameters. Section 3.1 presents a reversible jump MCMC based algorithm to obtain posterior samples of the emission parameters θ_j , where a trans-dimensional MCMC sampler is developed to explore subspaces of variable dimensionality regarding the number of periodicities that characterize state j . This part of the inference scheme follows a similar structure to the one presented by Andrieu and Doucet (1999) and Hadj-Amar et al. (2020). In Section 3.2 we address model search on the number of states by exploiting the *Chinese restaurant franchise with loyal customers* (Fox et al., 2011), a metaphor that provides the building blocks to perform Bayesian nonparametric inference for updating the HMM parameters. The details related to this part of the sampler follow the scheme presented in Fox et al. (2011). The resulting Gibbs sampler for the full estimation algorithm is described in Section 3.2.1 while in Section 3.2.2 we address the label switching problem related to our proposed approach.

3.1 Emission Parameters

Conditional on the state sequence z , the observations y are implicitly partitioned into a finite number of states, where each state refers to at least one segment of the time series. When a type of periodic behavior recurs over time, the corresponding state is necessarily related to more than one segment. Let $y_j^* = (y'_{j1}, y'_{j2}, \dots, y'_{jR_j})'$ be the vector of (non-adjacent) segments that are assigned to state j , where y_{jr} denotes the r^{th} segment of the time series for which $z_t = j$ and R_j is the total number of segments assigned to that state. Then, the likelihood of the emission parameter θ_j given the observations in y_j^* is

$$\mathcal{L}(\theta_j | y_j^*) = (2\pi\sigma_j^2)^{-T_j^*/2} \exp \left[-\frac{1}{2\sigma_j^2} \sum_{t \in I_j^*} \{y_t - x_t(\omega_j)' \beta_j\}^2 \right], \quad (8)$$

where I_j^* and T_j^* denote the set of time points and number of observations, respectively, associated with y_j^* .

Following Hadj-Amar et al. (2020), we assume independent Poisson prior distributions for the number of frequencies d_j for each state j , constrained on $1 \leq d_j \leq d_{max}$. Conditional on d_j , we choose a uniform prior for the frequencies $\omega_{j,l} \sim \text{Uniform}(0, \phi_\omega)$, $l = 1, \dots, d_j$ where $0 <$

$\phi_\omega 0.5$. The value of ϕ_ω can be chosen to be informative in the sense that it may reflect prior information about the significant frequencies that drive the overall variation in the data, for example ϕ_ω may be assumed to be in the low frequencies range $0 < \phi_\omega < 0.1$. Analogous to a Bayesian regression (Bishop, 2006), a zero-mean isotropic Gaussian prior is assumed for the coefficients of the j^{th} regime, $\beta_j \sim \mathcal{N}_{2d_j}(\mathbf{0}, \sigma_\beta^2 I)$, where the prior variance σ_β^2 is fixed at a relatively large value (e.g., in our case 10^2). The prior on the residual variance σ_j^2 of state j is specified as Inverse-Gamma $\left(\frac{\xi_0}{2}, \frac{\tau_0}{2}\right)$, where ξ_0 and τ_0 are fixed at small values, noticing that when $\xi_0 = \tau_0 = 0$ we obtain Jeffreys' uninformative prior (Bernardo and Smith, 2009).

Bayesian inference on θ_j is built upon the following factorization of the joint posterior distribution

$$p(\theta_j | y_j^*) = p(d_j | y_j^*) p(\omega_j | d_j, y_j^*) p(\beta_j | \omega_j, d_j, y_j^*) p(\sigma_j^2 | \beta_j, \omega_j, d_j, y_j^*). \quad (9)$$

Sampling from (9) gives rise to a model selection problem regarding the number of periodicities, thus requiring an inference algorithm that is able to explore subspaces of variable dimensionality. This will be addressed by the reversible-jump sampling step introduced in the following section.

3.1.1 Reversible-Jump Sampler—Here we provide the details for drawing θ_j from the posterior distribution $p(\theta_j | y_j^*)$ given in Equation (9). Our methodology follows Andrieu and Doucet (1999) and Hadj-Amar et al. (2020) and is based on the principles of reversible-jump MCMC introduced in Green (1995). Notice that, conditional on the state sequence $\mathbf{z}$, the emission parameters θ_j can be updated independently and in parallel for each of the current states. Hence, for the rest of this subsection and for ease of notation, we drop the subscript corresponding to the j^{th} state.

At each iteration of the algorithm, a random choice with probabilities given in (10) based on the current number of frequencies d will dictate whether to add a frequency (*birth step*) with probability b_d , remove a frequency (*death step*) with probability r_d , or update the frequencies (*within step*) with probability $\mu_d = 1 - b_d - r_d$ where

$$b_d = c \min\left\{1, \frac{p(d+1)}{p(d)}\right\}, \quad r_{d+1} = c \min\left\{1, \frac{p(d)}{p(d+1)}\right\}, \quad (10)$$

for some constant $c \in \left[0, \frac{1}{2}\right]$ and $P(d)$ is the prior probability. Here, as in Hadj-Amar et al. (2020), we fixed $c = 0.4$ but other values are admissible as long as c is not larger than 0.5 to guarantee that the sum of the probabilities does not exceed 1 for some values of c . Naturally, $b_{d_{\max}} = r_1 = 0$. An outline of these moves is as follows (further details are provided in Supplementary Material A).

Within-Model Move: Conditional on the number of frequencies d , the vector of frequencies ω is sampled, where we update the frequencies one-at-a-time using Metropolis-Hastings (M-H) steps, with target distribution

$$p(\omega | \beta, \sigma^2, d, y^*) \propto \exp \left[-\frac{1}{2\sigma^2} \sum_{t \in I^*} \{y_t - x_t(\omega)' \beta\}^2 \right] \mathbb{1}_{[\omega \in \Omega_d]}. \quad (11)$$

Specifically, the proposal distribution is a combination of a Normal random walk centred around the current frequency and a draw from the periodogram of $\hat{y}$, where $\hat{y}$ denotes a segment of data randomly chosen from y^* with probability proportional to the number of observations belonging to that segment. Naturally, when a state does not recur over time, i.e. when a state refers to only one segment of the time series, that segment is chosen with probability one. Next, updating the vector of linear coefficients β and the residual variance σ^2 is carried out as in the fashion of the usual normal Bayesian regression setting (Gelman et al., 2014). Hence, β is updated in a Gibbs step from

$$\beta | \omega, \sigma^2, d, y^* \sim \mathcal{N}_{2d}(\hat{\beta}, V_\beta), \quad (12)$$

where

$$\begin{aligned} V_\beta &= (\sigma_\beta^{-2} I + \sigma^{-2} X^*(\omega)' X^*(\omega))^{-1}, \\ \hat{\beta} &= V_\beta (\sigma^{-2} X^*(\omega)' y^*), \end{aligned} \quad (13)$$

and we denote with $X^*(\omega)$ the design matrix with rows given by $x_t(\omega)$ (Equation 4), for $t \in I^*$. Finally, σ^2 is drawn in a Gibbs step directly from

$$\sigma^2 | \beta, \omega, d, y^* \sim \text{Inverse-Gamma} \left(\frac{T + \xi_0}{2}, \frac{\tau_0 + \sum_{t \in I^*} \{y_t - x_t(\omega)' \beta\}^2}{2} \right). \quad (14)$$

Trans-Dimensional Moves: For these types of move, the number of periodicities is either proposed to increase by one (birth) or decrease by one (death) (Green, 1995). If a birth move is attempted, we have that $d^p = d^c + 1$, where we denote with superscripts c and p , the current and proposed values, respectively. The proposed vector of frequencies is obtained by drawing an additional frequency to be included in the current vector. On the other hand if a death move is chosen, we have that $d^p = d^c - 1$ and one of the current periodicities is randomly selected to be deleted. Conditional on the proposed vector of frequencies, the vector of linear coefficients and the residual variance are sampled as in the within-model move described above. For both birth and death moves, the updates are jointly accepted or rejected in a M-H step.

3.2 HMM Parameters

We explain how to perform posterior inference about the probability distribution α , the transition probabilities π_j and the state sequence z . The *Chinese restaurant franchise with loyal customers* presented by Fox et al. (2011), which extends the *Chinese restaurant franchise* introduced by Teh et al. (2006), is a metaphor that can be used to express the generative process behind the sticky version of the HDP and provides a general framework

for performing inference. A high level summary of the metaphor is as follows: in a *Chinese restaurant franchise* the analogy of a *Chinese restaurant process* (Aldous, 1985) is extended to a set of restaurants, where an infinite global menu of dishes is shared across these restaurants. The process of seating customers at tables happens in a similar way as for the Chinese restaurant process, but is restaurant-specific. The process of choosing dishes at a specific table happens franchise-wide, namely the dishes are selected with probability proportional to the number of tables (in the entire franchise) that have previously served that dish. However, in the Chinese restaurant franchise with loyal customers, each restaurant in the franchise has a speciality dish which may keep many generations of customers eating in the same restaurant.

Let $y_{j1}, \dots, y_{jN_j}$ denote the set of customers in restaurant j , where N_j is the number of customers in restaurant j and each customer is pre-allocated to a specific restaurant designated by that customer's group j . Let us also define indicator random variables t_{ji} and k_{jt} , such that t_{ji} indicates the table assignment for customer i in restaurant j , and k_{jt} the dish assignment for table t in restaurant j . In the Chinese restaurant franchise with loyal customers, customer i in restaurant j chooses a table via $t_{ji} \sim \tilde{\pi}_j$, where $\tilde{\pi}_j \sim \text{GEM}(\eta + \kappa)$, and η and κ are as in Section 2.1. Each table is assigned a dish via $k_{jt} \sim (\eta\alpha + \kappa\delta_j)/(\eta + \kappa)$, so that there is more weight on the house speciality dish, namely the dish that has the same index as the restaurant. Here, α follows a DP with concentration parameters γ and can be seen as a collection of ratings for the dishes served in the global menu. Note that in the HMM formulated in Equation (1), the value of the hidden state z_t corresponds to the dish index, i.e. $k_{jtji} = z_{ji} = z_t$, where we suppose there exist a bijection $f: t \rightarrow ji$ of time indexes t to restaurant-customer indexes ji . Furthermore, as suggested in Fox et al. (2011), we augment the space and introduce *considered* dishes $\bar{k}_{jt}$ and *override* variables o_{jt} so that we have the following generative process

$$\begin{aligned} \bar{k}_{jt} | \alpha &\sim \alpha \\ o_{jt} | \eta, \kappa &\sim \text{Bernoulli}\left(\frac{\kappa}{\eta + \kappa}\right) \\ k_{jt} | \bar{k}_{jt}, o_{jt} &= \begin{cases} \bar{k}_{jt}, & o_{jt} = 0, \\ j, & o_{jt} = 1. \end{cases} \end{aligned}$$

Thus, a table first considers a dish $\bar{k}_{jt}$ without taking into account the dish of the house, i.e. $\bar{k}_{jt}$ is chosen from the infinite buffet line according to the ratings provided by α . Then, the dish k_{jt} that is actually being served can be the house-speciality dish j , with probability $\rho = \kappa/(\eta + \kappa)$, or the initially considered dish $\bar{k}_{jt}$, with probability $1 - \rho$. As shown in Fox et al. (2011), table counts $\bar{m}_{jk}$ of considered dishes are sufficient statistics for updating the collection of dish ratings α , where $\bar{m}_{jk}$ denotes how many of the tables in restaurant j considered dish k . The sampling of $\bar{m}_{jk}$ is additionally simplified by introducing the table counts m_{jk} of served dishes and override variables o_{jt} . In the next section we describe a Gibbs sampler which alternates between updating the hidden states z , dish ratings α , transition probabilities π_j , newly introduced random variables m_{jk} , o_{jt} and $\bar{m}_{jk}$, emission parameters θ_j , as well as the hyperparameters γ , η and κ .

3.2.1 Gibbs Sampler—We follow Kivinen, Sudderth and Jordan (2007) and Fox et al. (2011) and consider a Gibbs sampler which uses finite approximations to the DP to allow sampling in blocks of the state sequence z . In particular, conditioned on observations y , transition probabilities π_j and emission parameters θ_j , the hidden states z are sampled using a variant of the well-known HMM forward-backward procedure (see Supplementary Material B) presented in Rabiner (1989). In order to use this scheme, we must truncate the countably infinite transition distributions π_j (and global menu α), and this is achieved using the $K_{\max}$ -limit approximation to a DP (Ishwaran and Zarepour, 2002), i.e. $\text{GEM}_{K_{\max}}(\gamma) := \text{Dir}(\gamma/K_{\max}, \dots, \gamma/K_{\max})$, where the truncation level $K_{\max}$ is a number that exceeds the total number of expected HMM states, and $\text{Dir}(\cdot)$ denote the Dirichlet distribution. Following Fox et al. (2011), conditioned on the state sequence z and collection of dish ratings α , we sample the auxiliary variables m_{jk} , O_{jt} and $\bar{m}_{jk}$ as described in Supplementary Material C.2. Dish ratings α and transition distributions π_j are then updated from the following posterior distributions

$$\alpha | \bar{m}, \gamma \sim \text{Dir}(\gamma/K_{\max} + \bar{m} \cdot 1, \dots, \gamma/K_{\max} + \bar{m} \cdot K_{\max})$$

$$\pi_j | z, \alpha, \eta, \kappa \sim \text{Dir}(\eta\alpha_1 + n_{j1}, \dots, \eta\alpha_j + \kappa + n_{jj}, \dots, \eta\alpha_{K_{\max}} + n_{jK_{\max}}),$$

for each state $j = 1, \dots, K_{\max}$. Here, $\bar{m}$ is the vector of table counts of considered dishes for the whole franchise, and marginal counts are described with dots, so that $\bar{m} \cdot k = \sum_{j=1}^{K_{\max}} \bar{m}_{jk}$ is the number of tables in the whole franchise considering dish k . We denote with n_{jk} the number of Markov chain transitions from state j to state k in the hidden sequence z . Next, given the state sequence z and transition probabilities π_j , we draw the emission parameters θ_j for each of the currently instantiated state as described in Section 3.1, where each reversible-jump MCMC update is run for several iterations. We also need to update the emission parameters for states which are not instantiated (namely, those states among $\{1, \dots, K_{\max}\}$ that are not represented during a particular iteration of the sampler), and hence we draw the corresponding emission parameters from their priors. For computational or modeling reasons, the latter may be also performed for those instantiated states that do not contain a minimum number of observations. Finally, we sample the hyperparameters γ , η and κ in a Gibbs step (see Supplementary Material C.3).

For the HDP-HMM, different procedures have been applied for sampling the hidden state sequence z . Teh et al. (2006) originally introduced an approach based on a Gibbs sampler which has been shown to suffer from slow mixing behavior due to strong correlations that is frequently observed in the data at nearby time points. Van Gael et al. (2008) presented a *beam sampling* algorithm that combines a slice sampler (Neal et al., 2003) with dynamic programming. This allows to constrain the number of reachable states at each MCMC iteration to a finite number, where the entire hidden sequence z is drawn in block using a form of forward-backward filtering scheme. However, Fox et al. (2011) showed that applications of the beam sampler to the HDP-HMM resulted in slower mixing rates compared to the forwardbackward procedure that we use in our truncated model. Recently, Tripuraneni et al. (2015) developed a particle Gibbs MCMC algorithm (Andrieu, Doucet and

Holenstein, 2010) which uses an efficient proposal and makes use of ancestor sampling to enhance the mixing rate.

3.2.2 Label Switching—The proposed approach may suffer from *label switching* (see e.g. Redner and Walker (1984); Stephens (2000); Jasra, Holmes and Stephens (2005)) since the likelihood is invariant under permutations of labelling of the mixture components, for both hidden state labels $\{1, \dots, K_{max}\}$ and frequency labels $\{1, \dots, d_{max}\}$ in each state. The label switching problem occurs when using Bayesian mixture models and needs to be addressed in order to draw meaningful inference about the posterior model parameters. In our multiple model search, the frequencies (and their corresponding linear coefficients) are identified by keeping them in ascending order for every iteration of the sampler. Posterior samples of the model parameters corresponding to different hidden states are post-processed (after the full estimation run) using the relabelling algorithm developed by Stephens (2000). The basic idea behind this algorithm is to find permutations of the MCMC samples in such a way that the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951) between the ‘true’ distribution on clusterings, say $P(\theta)$, and a matrix of classification probabilities, say Q , is minimized. The KL distance is given by $d(Q, P(\theta))_{KL} = \sum_t \sum_j p_{tj}(\theta) \log \frac{p_{tj}(\theta)}{q_{tj}}$, where $p_{tj}(\theta) = p(z_t = j | z_{t-1}, y, \pi, \theta)$ is part of the MCMC output obtained as in Supplementary Material C.1, and q_{tj} is the probability that observation t is assigned to class j . The algorithm iterates between estimating Q and the most likely permutation of the hidden labels for each MCMC iteration. We chose the strategy of Stephens (2000) since it has been shown to perform very efficiently in terms of finding the correct relabelling (see e.g. Rodriguez and Walker (2014)). However, it may be computationally quite intensive in memory since it requires the storage of a matrix of probabilities of dimension $N \times T \times K_{max}$, where N is the number of MCMC samples. Furthermore, at each iterative step, the algorithm requires to go over $K_{max}!$ permutations of the labels for each MCMC iteration, which might significantly slow down the computation when using large values of K_{max} . Related approaches to the label switching issue include pivotal reordering algorithms (Marin, Mengersen and Robert, 2005), label invariant loss functions (Celeux, Hurn and Robert, 2000; Hurn, Justel and Robert, 2003) and equivalent classes representative methods (Papastamoulis and Iliopoulos, 2010), where an overview of these strategies can be found in Rodriguez and Walker (2014).

4 Simulation Studies

This section presents results of simulation studies to explore the performance of our proposed methodology in two different settings. In the first scenario the data are generated from the model described in Section 2 and thus this simulation study provides a “sanity” check that the algorithm is indeed retrieving the correct pre-fixed parameters. We also investigate signal extraction for the case that the innovations come from a heavy-tailed t-distribution instead of a Gaussian. Our second study deals with artificial data from an HMM whose emission distributions are characterized by oscillatory dynamics generated by state-specific autoregressive (AR) time series models. Julia code that implements our procedure is available at <https://github.com/Beniamino92/HHMM>.

4.1 Illustrative Example

We generated a time series consisting of $T = 1450$ data points from a three-state HMM with the following transition probability matrix showing high probabilities of self transition along the diagonal and Gaussian oscillatory emissions as specified in Equation (2), where the parameters of each of the three regimes and the transition probability matrix are given in Supplementary Material A. A realization from this model is displayed in Figure 2. The prior mean on the number of frequencies d_j is set equal to 1 and we place a Gamma (1,0.01) prior on the concentration parameters γ and $(\gamma + \kappa)$, and a Beta (100,1) prior on the self-transition proportion ρ as in Fox et al. (2011). The maximum number of periodicities per regime $d_{\max}$ is set to 5, while the truncation level $K_{\max}$ for the DP approximation is set equal to 7. Also, we set $\phi_\omega = 0.25$ as a threshold for the uniform prior. The proposed estimation algorithm is run for 15,000 iterations, 3,000 of which were discarded as burn-in. At each iteration, for each instantiated set of emission parameters, 2 reversible-jump MCMC updates were performed. The full estimation algorithm took 31 minutes with a program written in Julia 0.62 on an Intel® Core™ i7 2.2 GHz Processor 8 GB RAM. For our experiments, we used the R package *label.switching* of Papastamoulis (2016) to post-process the MCMC output with the relabelling algorithm of Stephens (2000).

Table 1 (left panel) shows that our estimation algorithm successfully detects the correct number of states in the sense that a model with $k = 3$ regimes has the highest posterior probability. In addition, our approach correctly identifies the right number of frequencies in each regime, as shown in Table 1 (right panel). Table 2 displays the estimated posterior mean and standard deviation of the frequencies along with the square root of the power of the corresponding frequencies, where the results are conditional on three estimated states and the modal number of frequencies within each state. Here, the power of each frequency ω_{jl} is summarized by the amplitude $A_{jl} = \sqrt{\beta_{jl}^{(1)^2} + \beta_{jl}^{(2)^2}}$, namely the square root of the sum of squares of the corresponding linear coefficients (see, e.g., Shumway and Stoffer (2017)). Our proposed method seems to provide a good match between true and estimated values for both frequencies and their power, for this example. We also show in Figure 2 the state-specific estimated signal (Equation (3)), and the estimated state sequence using the method of Stephens (2000) (as a piecewise horizontal line). The rows of the estimated transition probability matrix were $\hat{\pi}_1 = (0.9921, 0.0073, 0.0006)$, $\hat{\pi}_2 = (0.0005, 0.9956, 0.0040)$ and $\hat{\pi}_3 = (0.0051, 0.0006, 0.9942)$. The high probabilities along the diagonal reflect the estimated posterior mean of the self transition parameter $\hat{\rho} = 0.9860$, which is indeed centered around the true probability of self-transition.

Diagnostics for verifying convergence were performed in several ways. For example, we observed that the MCMC samples of the likelihood of the HMM reached a stable regime, while initializing the Markov chains from overdispersed starting values (see Figure 3 (b)). This diagnostic might be very useful, for example, in determining the burn-in period. However, we note that it does not guarantee convergence since steady values of the log likelihood might be the result of a Markov chain being stuck in some local mode of the target posterior distribution. The likelihood of an HMM with Gaussian emissions can be expressed as

$$\mathcal{L}(z, \pi, \theta | y) = p(z_1 | y, \pi, \theta) \mathcal{N}(y_1; f_{1z_1}, \sigma_{z_1}^2) \prod_{t=2}^T p(z_t | z_{t-1}, y, \pi, \theta) \mathcal{N}(y_t; f_{tz_t}, \sigma_{z_t}^2), \quad (15)$$

where $\mathcal{N}(y_t; f_{jt}, \sigma_j^2)$ denotes the density of a Gaussian distribution with mean $f_{jt} = x_t(\omega_j)' \beta_j$ (as in Equation 3) and variance σ_j^2 , evaluated at y_t . Conditioned on the modal number of states, we also validated convergence for the state-specific emission parameters by analyzing trace plots and running averages of the corresponding MCMC samples, with acceptable results as each trace reached a stable regime. As an example, we show in Figure 3(a) trace plots (after burn-in) for the posterior values of the frequencies.

Finally, while we notice that we have not set $K_{\max}$ to a very large value, this choice has been made *a posteriori* after we found that the estimation algorithm assigned negligible probabilities to a large number of components. For example, in this simulation study, we initially set $K_{\max} = 20$ and ran the full estimation algorithm; after observing that the posterior probabilities for the number of distinct states were equal to zero for all the models with more than four hidden states, we re-ran the estimation algorithm with a smaller value for the HDP truncation, i.e. $K_{\max} = 7$, obtaining the same correct results. This yielded some benefits from a computational perspective, in particular in terms of facilitating storage and memory access of the posterior sample, and also speeding up the relabelling algorithm developed by Stephens (2000). Furthermore, we notice that even in the case where the maximum complexity of the model is set to $K_{\max} = 7$, we are still dealing with a framework that assumes a relatively high number of regimes (within a fairly complicated setting of oscillations in each state). Conventional models might not even be able to achieve satisfactory estimation performances when specifying such a large number of Markov modes.

Signal Extraction with Non-Gaussian Innovations—In many scientific experiments it may be of interest to extract the underlying signal that generates the observed time series and HMMs can be used to this end. Here, we study the performance of our proposed approach in estimating the time-varying oscillatory signal f_{jt} (Equation 3) when the Gaussian assumption of ε_t in Equation (2) is violated. In particular, we generated 20 time series, each consisting of 1024 observations from the same simulation setting introduced above, where the innovations were simulated from heavy-tailed t -distributions with 2, 3, 2 degrees of freedom for state 1, 2, 3, respectively. The linear basis coefficients were chosen to be $\beta_{11} = (3, 2)'$, $\beta_{21} = (1.2, 4.0)'$, $\beta_{31} = (1.0, 5.0)'$, $\beta_{32} = (4.0, 3.0)'$. As a measure of performance we computed the (in-sample) mean squared error $MSE = \frac{1}{1024} \sum_{t=1}^{1024} (f_{tz_t} - \hat{f}_{tz_t})^2$ between true and estimated signal and compared the proposed approach with the method of Hadj-Amar et al. (2020) referred to as AutoNOM (Automatic Nonstationary Oscillatory Modelling), which we believe is the state-of-the-art in extracting the signal of nonstationary periodic processes. Our proposed estimation algorithm was run with the same parameterization as above while AutoNOM was performed for 15,000 updates, 3,000 of which were discarded as burn-in, where we fixed 15 maximum number of change points and 5 maximum number of frequencies per segment. The prior means for the number of change-points and frequencies

per segment are fixed at 2 and 1, respectively, and the minimum distance between change-points is set at 10. For both methodologies, the estimated signal was obtained by averaging across MCMC iterations. AutoNOM was run using the Julia software provided by the authors at <https://github.com/Beniamino92/AutoNOM>.

Figure 4 (a) presents boxplots of the MSE values for AutoNOM and our proposed approach which will be referred to as HHMM (Harmonic Hidden Markov Model). It becomes clear that the estimates of the signal obtained using HHMM are superior to those obtained using AutoNOM. However, this result is not surprising as the two approaches make different assumptions. In particular, AutoNOM does not assume recurrence of a periodic behavior and hence needs to estimate the regime-specific modeling parameters each time it detects a new segment, while our HHMM has the advantage of using the same set of parameters whenever a particular periodic pattern recurs in the time series. Hence, we also compared the performance of the two approaches in extracting the signal (under non-Gaussian innovations) in a scenario where the time series do not exhibit recurrence. Specifically, we generated 15 time series manifesting two change-points (where the oscillatory behavior corresponding to the three different partitions is parameterized as above) and computed the MSE between true and estimated signal as we did in the previous scenario. The corresponding boxplots displayed in Figure 4 (b) show that the two approaches seem to perform in similar way, with AutoNOM being slightly more accurate than our HHMM. Moreover, we have further examined a scenario where the noise is generated from exponentially distributed random variables in order to introduce a large skew. Here we simulated 15 time series from a three-state HMM as above where now the innovations corresponding to the three different states are generated from exponentially distributed random variables with rates 0.5, 1, and 0.2, respectively. The draws from the exponential distribution are centered in such a way that they have mean zero, to avoid noise that takes on strictly positive values. Figure 4 (c) presents boxplots of the MSE values for AutoNOM and our proposed HHMM, showing that our methodology seems to be superior to AutoNOM, in terms of extracting the signal when the innovations are skewed. We conclude that both approaches have their own strengths. Our proposed procedure is superior to AutoNOM in the sense that the additional HMM provides a framework for modeling and explicitly quantifying the switching dynamics and connectivity between different states. On the other hand, AutoNOM is better suited to scenarios where there are nonstationarities arising from singular change-points and the observed oscillatory processes evolve without recurrent patterns.

4.2 Markov Switching Autoregressive Process

We now investigate the performance of our approach in detecting time-changing periodicities in a scenario where the data generating process shows large departures from our modeling assumptions. The HMM hypothesis which assumes conditionally independent observations given the hidden state sequence, such as the one formulated in Equation (2), may sometimes be inadequate in expressing the temporal dependencies occurring in some phenomena. A different class of HMMs that relaxes this assumption is given by the *Markov switching autoregressive process*, also referred to as the AR-HMM (Juang and Rabiner, 1985; Albert and Chib, 1993; Frühwirth-Schnatter, 2006), where an AR process

is associated with each state. This model is used to design autoregressive dynamics for the emission distributions while allowing the state transition mechanisms to follow a discrete state Markov chain.

We generated $T = 900$ observations from an AR-HMM with two hidden states and autoregressive order fixed at $p = 2$, that is

$$\begin{aligned} z_t &\sim \pi_{z_t - 1} \\ y_t &= \sum_{l=1}^p \psi_l^{(z_t)} y_{t-l} + \varepsilon^{(z_t)}_t, \end{aligned} \quad (16)$$

where $\pi_1 = (0.99, 0.01)$ and $\pi_2 = (0.01, 0.99)$. The AR parameterization $\psi^{(1)} = (1.91, -0.991)$ and $\psi^{(2)} = (1.71, -0.995)$ is chosen in such a way that the state-specific spectral density functions display a pronounced peakedness. Furthermore, $\varepsilon_t^{(1)} \stackrel{iid}{\sim} \mathcal{N}(0, 0.1^2)$ and $\varepsilon_t^{(2)} \stackrel{iid}{\sim} \mathcal{N}(0, 0.05^2)$. A realization from this model is shown in Figure 5 (top) as a blue solid line. Our proposed estimation algorithm was run for 15,000 iterations 5,000 of which are used as burn-in. At each iteration, we performed 2 reversible-jump MCMC updates for each instantiated set of emission parameters. The rate of the Poisson prior for the number of periodicities is fixed at 10^{-1} and the corresponding truncation level $d_{\max}$ was fixed to 3. The maximum number of states $K_{\max}$ was set equal to 10 whereas the rest of the hyperparameters are specified as in Section 4.1. Our procedure seems to overestimate the number of states, as a model with 8 regimes had the highest posterior probability $\hat{\pi}(k = 8 | y) = 97\%$. However, this is not entirely unexpected as a visual inspection of the realization displayed in Figure 5 (top) suggest more than two distinct spectral patterns in the sense that the phases, amplitudes and frequencies appear to vary stochastically within a regime. Figure 5 (bottom) shows the estimated time-varying frequency peak along with a 95% credible interval obtained from the posterior sample. The estimate was determined by first selecting the dominant frequency (i.e. the frequency with the highest posterior power) corresponding to each observation and then averaging the frequency estimates over MCMC iterations. While our approach identifies a larger number of states when the data were generated from an AR-HMM we note that the data generating process are very different from the assumptions of our model and the proposed procedure still provides a reasonable summary of the underlying time changing spectral properties observed in the data. Furthermore, by setting the truncation level $K_{\max}$ equal to 2, we retrieve the true transition probability matrix that generates the switching dynamics between the two different autoregressive patterns, as the vectors of transition probabilities obtained using our estimation algorithm are $\hat{\pi}_1 = (0.99, 0.01)$ and $\hat{\pi}_2 = (0.98, 0.02)$.

In addition, we simulated 10 time series from model (16) and computed the mean squared error $\text{MSE} = \frac{1}{900} \sum_{t=1}^{900} (\omega_t - \hat{\omega}_t)^2$ between the true time-varying frequency peak ω_t and its estimate $\hat{\omega}_t$ for the proposed approach, AutoNOM and the procedure of Rosen, Wood and Stoffer (2012), referred to as AdaptSPEC (Adaptive Spectral Estimation). For both AutoNOM and AdaptSPEC, we ran the algorithm for 15,000 MCMC iterations (5,000 of which were used as burn-in), fixed the maximum number of change-points at 15 and set the

minimum distance between change-points to 30. The number of spline basis functions for AdaptSPEC is set to 10. AutoNOM is performed using a Poisson prior with rate 10^{-1} for both number of frequencies and number of change-points. AdaptSPEC was performed using the R package *BayesSpec* provided by the authors. Boxplots of the MSE values for the three different methodologies are displayed in Figure 5 (right), showing that our proposed HHMM seems to outperform the other two approaches in detecting the time-varying frequency peak, for this example. However, our procedure finds some very short sequences (such as in Figure 5 (bottom) for $t \approx 200, 500, 700$) demonstrating that the sticky parameter might not always be adequate enough in capturing the correct temporal mode persistence of the latent state sequence. AutoNOM and AdaptSPEC are less prone to this problem as both methodologies are able to specify a minimum time distance between change-points; though, we acknowledge that this constraint might not be optimal when the observed data exhibit relatively rapid changes. We also notice that, not surprisingly, the estimates of the time-varying frequency peak obtained using AutoNOM and our HHMM, which are based on a line-spectrum model, are both superior than the ones obtained via the smoothing spline basis of AdaptSPEC, which is built upon a continuous-spectrum setting; this is consistent with the findings in Hadj-Amar et al. (2020). However, it is important to keep in mind that, while AutoNOM and Adapt-SPEC allow to retrospectively analyse the spectral changing properties of a process from an exploratory angle, unlike our proposed HHMM, they do not quantify a probabilistic mechanism for the recurrence of periodic dynamic patterns.

5 Analysis of the Airflow Trace Data

The airflow trace shown in Figure 1 was collected from a human over a time span of 5.5 minutes of continuous breathing and measured via a facemask attached to a pressure transducer. Airflow pressure signals were amplified using the NeuroLog system connected to a 1401 interface and acquired on a computer using *Spike2* software (Cambridge Electronic Design). The data are sampled at rate of 4 Hertz, i.e., 4 observations per second, for a total of 1314 data points. The air flow data was captured and amplified only, and no pre-processing was carried out afterwards. We notice that the signal is clean since the pressure transducer is sensitive enough to pick up breathing even from a mouse in a plethysmograph, and it was attached to a medical face-mask directly in front of the mouth. Therefore, the signal to noise ratio is extraordinarily high.

We fitted our HHMM to the time series displayed in Figure 1 for 200,000 iterations, 125,000 of which were discarded as burn-in, where at each iteration, we carried out 10 reversible-jump MCMC updates for each instantiated set of emission parameters. The truncation level $K_{\max}$ was set to 10, whereas the maximum number of frequencies per state $d_{\max}$ was fixed to 3. With respect to the harmonic regression part of the model, we specified the prior for the innovations σ_j^2 as Inverse-Gamma (3.11, 3.17), so that it is centered at the empirical variance 1.5 and has a standard deviation of $\sqrt{2.5}$. The Poisson prior for the number of frequencies d_j is chosen equal to 10^{-2} , to favour models with a small number of components, and the prior on the frequencies $\omega_{j,l}$ is set informative as Uniform(0, ϕ_ω), where $\phi_\omega = 0.3$ was selected by looking at the raw periodogram of the data and noting that the power of the frequencies was approximately zero for any values larger than 0.3. Finally, we specified weakly informative

prior on the linear basis coefficients β_j as $\mathcal{N}_{2d_j}(\mathbf{0}, \sigma_\beta^2 I)$, where the prior variance $\sigma_\beta^2 = 4$ is chosen in now such a way that the mass of the prior is concentrated in reasonable regions based on the data. Regarding the part of the model relative to the HDP, we specified for both concentration parameters $\eta + \kappa$ and γ a weakly informative hyperprior Gamma (1,0.01), so that the corresponding priors for the base measures favour a DP model with a small number of components (see stick breaking construction, Equations (6) and (7)). The prior on the self transition proportions ρ is specified informative as Beta (10^3 , 1), so that we force high probability of self transitions.

The posterior distribution over the number of states had a mode at 6, with posterior probabilities $\hat{\pi}(k = 6 | y) = 94\%$, $\hat{\pi}(k = 7 | y) = 5\%$ and $\hat{\pi}(k = 8 | y) = 1\%$. Indeed, it is conceivable that the state corresponding to *normal* breathing (i.e. neither apnea or hypopnea) may exhibit more than one distinct periodic pattern, which further justifies the need to use a nonparametric HMM. Paz and West (2013) reported at least 13 forms of breathing patterns including forms of apnea. Figure 6 shows the estimated hidden state sequence (piecewise horizontal line), where we highlight our model estimate for apnea state (red) and hypopnea state (blue) while reporting the ground truth at the top of the plot. We have also included a posterior predictive graphical check consisting of the airflow trace alongside 20 draws from the estimated posterior predictive (Gelman et al., 2014), where the latter is obtained by first drawing a sample path, and then, conditioned on the hidden state sequence, the predicted values are simulated from the appropriate emission distributions. Our model seems to be able to capture the underlying signal that characterizes this time series. Conditional on the modal number of regimes, the number of periodicities belonging to apnea and hypopnea had a posterior mode at 3 and 2, respectively. Conditional on the modal number of frequencies, Table 3 displays the posterior mean and standard deviation of periodicities (in seconds) and powers that characterize the two states classified as apnea and hypopnea, showing that apnea instances seem to be characterized by larger periods and lower amplitude than hypopnea.

Our estimation algorithm detected all known apnea and hypopnea instances. In order for them to qualify as a clinically relevant obstructive event they must have a minimum length of 10 seconds (Berry et al., 2017). Thus, we only highlight the clinically relevant instances in Figure 6, discarding sequences of duration less than 10 seconds. We also detected a *post sigh* apnea (after the third minute) which is a normal phenomenon to observe in a breathing trace and hence should not count as a disordered breathing event. Again, such an event after a sigh can be identified as a sigh is characterized by an amplitude which is always higher than any other respiratory event and hence can be easily detected. Subtracting the number of sighs from the total number of apneas/hypopneas results in a measure of all apneas of interest without the confounding data from post sigh apneas. A common score to indicate the severity of sleep apnea is given by the Apnea-Hypopnea Index (AHI) which consists of the number of apneas and hypopneas per hour of sleep (Ruehland et al., 2009). Our proposed approach provides a realistic estimate of the total number of apnea and hypopnea instances recurring in this case study. While an essential aim of this paper is to detect apnea instances retrospectively, which is currently a time consuming and demanding task as it is performed by eye, we have also investigated the out-of-sample predictive performance of our proposed approach in Supplementary Material E.

6 Summary and Discussion

In this paper we developed a novel HMM approach that can address the challenges of modeling periodic phenomena whose behavior switches and recurs dynamically over time. The number of states is assumed unknown as well as their relevant periodicities which may differ over the different regimes since each regime is represented by a different periodic pattern. To address flexibility in the number of states, we assumed a sticky HDP-HMM that penalises rapid changing dynamics of the process and provides effective control over the switching rate. The variable dimensionality with respect to the number of frequencies that specifies the different states is tackled by a reversible-jump MCMC algorithm.

While being noticeably flexible, the model proposed in this article is developed in a framework that is still computationally accessible. Naturally, an alternative strategy would involve fitting several finite-state HMMs and then performing in a second stage model selection by means of the marginal likelihood (Kass and Raftery, 1995). Nevertheless, reliably approximating this quantity from the posterior sample is not straightforward, though several techniques have been proposed in the literature to overcome this burden, such as sequential Monte Carlo (SMC, Jasra et al. 2008), population MCMC (PMCMC, Liang and Wong 2001; Jasra, Stephens and Holmes 2007) or bridge sampling (Meng and Wong, 1996; Meng and Schilling, 2002), where we refer the reader to Zhou, Johansen and Aston (2016) for an excellent summary of these developments. However, these methods involve algorithms that are often computationally challenging and difficult to implement efficiently, especially within our modeling framework.

We illustrated the use of our approach in a case study relevant to respiratory research, where our methodology was able to identify recurring instances of sleep apnea in human breathing traces. Despite the fact that here we have focused on the detection of apnea instances, our proposed methodology provides a very flexible and general framework to analyze different breathing patterns. A question of interest is whether similar dynamical patterns can be identified across a heterogeneous patient cohort, and be used for the prognosis of patients' health and progress. The growth of information and communication technologies permits new advancements in the health care system to facilitate support in the homes of patients in order to proactively enhance their health and well-being. We believe that our proposed HMM approach has the potential to aid the iterative feedback between clinical investigations in sleep apnea research and practice with computational, statistical and mathematical analysis.

As pointed out by a referee, apnea states have certain features, such as low amplitude and low frequency behaviour, that may suggest that assuming symmetry among the parameters in their prior distribution might not be an ideal modeling approach. However, as we discussed in this article, it is plausible that normal breathing is exhibiting more than one distinct periodic pattern. While it would be of interest to integrate prior knowledge into the model to fully remove permutations of labelling of the HDP-HMM mixture components, we believe it would not be trivial to fully characterize in an identifiable way the different states corresponding to normal breathing. These are early days for such data to be analyzed in this

way and we are at the beginning of being able to construct a catalogue of more informative priors that might help this type of analysis in future.

Although both parametric and nonparametric HMMs have been shown to be good models in addressing learning challenges in time series data, they have the drawback of limiting the *state duration distribution*, i.e., the distribution for the number of consecutive time points that the Markov chain spends in a given state, to a geometric form (Ephraim and Merhav, 2002). In addition, the self-transition bias of the sticky HDP-HMM used to increase temporal state persistence is shared among all states and thus does not allow for inferring state-specific duration features. In our application, learning the duration structure of a specific state may be of interest to health care providers, for example, in assessing the severity of sleep apnea. Future work will address extending our approach to a hidden semi-Markov model (HSMM) setting (Guédon, 2003; Yu, 2010; Johnson and Willsky, 2013), where the generative process of an HMM is augmented by introducing a random state duration time which is drawn from some state-specific distribution when the state is entered. However, this increased flexibility in modeling the state duration has the cost of increasing substantially the computational effort to compute the likelihood: the message-passing procedure for HSMMs requires $\mathcal{O}(T^2 K + TK^2)$ basic computations for a time series of length T and number of states K , whereas the corresponding forward-backward algorithm for HMMs requires only $\mathcal{O}(TK^2)$.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We wish to thank Maxwell Renna, Paul Jenkins and Jim Griffin for their insightful and valuable comments. The work presented in this article was developed as part of the first author's Ph.D. thesis at the University of Warwick and he is currently affiliated with the Department of Statistics at Rice University. B. Hadj-Amar was supported by the Oxford-Warwick Statistics Programme (OxWaSP) and the Engineering and Physical Sciences Research Council (EPSRC), Grant Number EP/L016710/1. R. Huckstepp was supported by the Medical Research Council (MRC), Grant Number MC/PC/15070.

References

- Adak S. Time-dependent spectral analysis of nonstationary time series. *Journal of the American Statistical Association*. 1998; 93 :1488–1501.
- Albert JH, Chib S. Bayes inference via Gibbs sampling of autoregressive time series subject to Markov mean and variance shifts. *Journal of Business & Economic Statistics*. 1993; 11 :1–15.
- Aldous, DJ. École d'Été de Probabilités de Saint-Flour XII-I 1983. Springer; 1985. 1–198.
- Ancoli-Israel S, Klauber MR, Butters N, Parker L, Kripke DF. Dementia in institutionalized elderly: relation to sleep apnea. *Journal of the American Geriatrics Society*. 1991; 39 :258–263. [PubMed: 2005339]
- Andrieu C, Doucet A. Joint Bayesian model selection and estimation of noisy sinusoids via reversible jump MCMC. *IEEE Transactions on Signal Processing*. 1999; 47 :2667–2676.
- Andrieu C, Doucet A, Holenstein R. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2010; 72 :269–342.
- Baum LE, Eagon JA. An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology. *Bulletin of the American Mathematical Society*. 1967; 73 :360–363.

- Baum LE, Petrie T. Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*. 1966; 37 :1554–1563.
- Beal MJ, Ghahramani Z, Rasmussen CE. The infinite hidden Markov model. *Advances in Neural Information Processing Systems*. 2002 :577–584.
- Bernardo, JM, Smith, AF. Bayesian theory. Vol. 405. John Wiley & Sons; 2009.
- Berry RB, Brooks R, Gamaldo C, Harding SM, Lloyd RM, Quan SF, Troester MT, Vaughn BV. AASM scoring manual updates for 2017 (version 2.4). *Journal of Clinical Sleep Medicine*. 2017; 13 :665–666. [PubMed: 28416048]
- Bishop, CM. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag; New York: 2006.
- Bruce SA, Hall MH, Buysse DJ, Krafty RT. Conditional adaptive Bayesian spectral analysis of nonstationary biomedical time series. *Biometrics*. 2018; 74 :260–269. [PubMed: 28482111]
- Cappé, O, Moulines, E, Rydén, T. *Inference in hidden Markov models*. Springer; 2005.
- Celeux G, Hurn M, Robert CP. Computational and inferential difficulties with mixture posterior distributions. *Journal of the American Statistical Association*. 2000; 95 :957–970.
- Cohen, MX. *Analyzing neural time series data: theory and practice*. MIT press; 2014.
- Cooke JR, Ayalon L, Palmer BW, Loredó JS, Corey-Bloom J, Natarajan L, Liu L, Ancoli-Israel S. Sustained use of CPAP slows deterioration of cognition, sleep, and mood in patients with Alzheimer's disease and obstructive sleep apnea: a preliminary study. *Journal of clinical sleep medicine*. 2009; 5 :305–309. [PubMed: 19968005]
- Dahlhaus R, et al. Fitting time series models to nonstationary processes. *The Annals of Statistics*. 1997; 25 :1–37.
- Davis RA, Lee TCM, Rodriguez-Yam GA. Structural break estimation for nonstationary time series models. *Journal of the American Statistical Association*. 2006; 101 :223–239.
- Dewan NA, Nieto FJ, Somers VK. Intermittent hypoxemia and OSA: implications for comorbidities. *Chest*. 2015; 147 :266–274. [PubMed: 25560865]
- Ephraim Y, Merhav N. Hidden Markov processes. *IEEE Transactions on Information Theory*. 2002; 48 :1518–1569.
- Fox EB, Sudderth EB, Jordan MI, Willsky AS. A sticky HDP-HMM with application to speaker diarization. *The Annals of Applied Statistics*. 2011 :1020–1056.
- Frühwirth-Schnatter, S. *Finite mixture and Markov switching models*. Springer Science & Business Media; 2006.
- Gelman, A, Carlin, JB, Stern, HS, Dunson, DB, Vehtari, A, Rubin, DB. *Bayesian Data Analysis*. Vol. 2. CRC press; Boca Raton, FL: 2014.
- Green PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*. 1995; 82 :711–732.
- Guédon Y. Estimating hidden semi-Markov chains from discrete sequences. *Journal of Computational and Graphical Statistics*. 2003; 12 :604–639.
- Hadj-Amar B, Rand BF, Fiecas M, Lévi F, Huckstepp R. Bayesian Model Search for Nonstationary Periodic Time Series. *Journal of the American Statistical Association*. 2020; 115 :1320–1335.
- Heinzer R, Vat S, Marques-Vidal P, Marti-Soler H, Andries D, Tobback N, Mooser V, Preisig M, Malhotra A, Waeber G, et al. Prevalence of sleep-disordered breathing in the general population: the HypnoLaus study. *The Lancet Respiratory Medicine*. 2015; 3 :310–318. [PubMed: 25682233]
- Hjort, NL, Holmes, C, Müller, P, Walker, SG. *Bayesian nonparametrics*. Vol. 28. Cambridge university Press; 2010.
- Huang Q, Cohen D, Komarzynski S, Li X-M, Innominato P, Lévi F, Finkenstädt B. Hidden Markov models for monitoring circadian rhythmicity in telemetric activity data. *Journal of The Royal Society Interface*. 2018; 15
- Hurn M, Justel A, Robert CP. Estimating mixtures of regressions. *Journal of Computational and Graphical Statistics*. 2003; 12 :55–79.
- Ignatov T. On a constant arising in the asymptotic theory of symmetric groups, and on Poisson-Dirichlet measures. *Theory of Probability & Its Applications*. 1982; 27 :136–147.

- Ishwaran H, Zarepour M. Exact and approximate sum representations for the Dirichlet process. *Canadian Journal of Statistics*. 2002; 30 :269–283.
- Jasra A, Holmes CC, Stephens DA. Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling. *Statistical Science*. 2005 :50–67.
- Jasra A, Stephens DA, Holmes CC. On population-based simulation for static inference. *Statistics and Computing*. 2007; 17 :263–279.
- Jasra A, Doucet A, Stephens DA, Holmes CC. Interacting sequential Monte Carlo samplers for trans-dimensional simulation. *Computational Statistics & Data Analysis*. 2008; 52 :1765–1791.
- Johnson MJ, Willsky AS. Bayesian nonparametric hidden semi-Markov models. *Journal of Machine Learning Research*. 2013; 14 :673–701.
- Juang B-H, Rabiner L. Mixture autoregressive hidden Markov models for speech signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing*. 1985; 33 :1404–1413.
- Kass RE, Raftery AE. Bayes factors. *Journal of the American Statistical Association*. 1995; 90 :773–795.
- Kivinen, JJ; Sudderth, EB; Jordan, MI. Learning multiscale representations of natural scenes using Dirichlet processes. 2007 IEEE 11th International Conference on Computer Vision; IEEE; 2007. 1–8.
- Komarzynski S, Huang Q, Innominato PF, Maurice M, Arbaud A, Beau J, Bouchahda M, Ulusakarya A, Beaumatin N, Breda G, et al. Relevance of a Mobile internet Platform for Capturing Inter-and Intrasubject Variabilities in Circadian Coordination During Daily Routine: pilot Study. *Journal of Medical Internet Research*. 2018; 20 :e204. [PubMed: 29704408]
- Krauchi K, Wirz-Justice A. Circadian rhythm of heat production, heart rate, and skin and core temperature under unmasking conditions in men. *American Journal of Physiology-Regulatory Integrative and Comparative Physiology*. 1994; 267 :R819–R829.
- Krogh A, Brown M, Mian IS, Sjölander K, Haussler D. Hidden Markov models in computational biology: applications to protein modeling. *Journal of Molecular Biology*. 1994; 235 :1501–1531. [PubMed: 8107089]
- Kullback S, Leibler RA. On information and sufficiency. *The Annals of Mathematical Statistics*. 1951; 22 :79–86.
- Langrock R, Swihart BJ, Caffo BS, Punjabi NM, Crainiceanu CM. Combining hidden Markov models for comparing the dynamics of multiple sleep electroencephalograms. *Statistics in Medicine*. 2013; 32 :3342–3356. [PubMed: 23348835]
- Liang F, Wong WH. Real-parameter evolutionary Monte Carlo with applications to Bayesian mixture models. *Journal of the American Statistical Association*. 2001; 96 :653–666.
- Malik M. Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. *Annals of Noninvasive Electrocardiology*. 1996; 1 :151–181.
- Marin J-M, Mengersen K, Robert CP. Bayesian modelling and inference on mixtures of distributions. *Handbook of statistics*. 2005; 25 :459–507.
- Meng X-L, Schilling S. Warp bridge sampling. *Journal of Computational and Graphical Statistics*. 2002; 11 :552–586.
- Meng X-L, Wong WH. Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*. 1996 :831–860.
- Neal RM, et al. Slice sampling. *The Annals of Statistics*. 2003; 31 :705–767.
- Ombao HC, Raz JA, von Sachs R, Malow BA. Automatic statistical analysis of bivariate nonstationary time series. *Journal of the American Statistical Association*. 2001; 96 :543–560.
- Papastamoulis P. label.switching: An R Package for Dealing with the Label Switching Problem in MCMC outputs. *Journal of Statistical Software, Code Snippets*. 2016; 69 :1–24.
- Papastamoulis P, Iliopoulos G. An artificial allocations based solution to the label switching problem in Bayesian analysis of mixtures of distributions. *Journal of Computational and Graphical Statistics*. 2010; 19 :313–331.
- Paz, JC, West, MP. *Acute Care Handbook for Physical Therapists*. Elsevier Health Sciences; 2013.

- Peker Y, Hedner J, Norum J, Kraiczi H, Carlson J. Increased incidence of cardiovascular disease in middle-aged men with obstructive sleep apnea: a 7-year follow-up. *American Journal of Respiratory and Critical Care Medicine*. 2002; 166 :159–165. [PubMed: 12119227]
- Perman M, Pitman J, Yor M. Size-biased sampling of Poisson point processes and excursions. *Probability Theory and Related Fields*. 1992; 92 :21–39.
- Pitman J. Blackwell-Macqueen urn scheme. *Statistics, Probability, and Game Theory: Papers in Honor of David Blackwell*. 1996; 30 :245.
- Pitman J. Poisson–Dirichlet and GEM invariant distributions for split-and-merge transformations of an interval partition. *Combinatorics, Probability and Computing*. 2002; 11 :501–514.
- Priestley MB. Evolutionary spectra and non-stationary processes. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1965; 27 :204–229.
- Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*. 1989; 77 :257–286.
- Rasmussen CE, Ghahramani Z. Infinite mixtures of Gaussian process experts. *Advances in Neural Information Processing Systems*. 2002 :881–888.
- Redner RA, Walker HF. Mixture densities, maximum likelihood and the EM algorithm. *SIAM review*. 1984; 26 :195–239.
- Rodríguez CE, Walker SG. Label switching in Bayesian mixture models: Deterministic relabeling strategies. *Journal of Computational and Graphical Statistics*. 2014; 23 :25–45.
- Rosen O, Stoffer DS, Wood S. Local spectral analysis via a Bayesian mixture of smoothing splines. *Journal of the American Statistical Association*. 2009; 104 :249–262.
- Rosen O, Wood S, Stoffer DS. AdaptSPEC: Adaptive spectral estimation for nonstationary time series. *Journal of the American Statistical Association*. 2012; 107 :1575–1589.
- Ruehland WR, Rochford PD, OâAZdonoghue FJ, Pierce RJ, Singh P, Thornton AT. The new AASM criteria for scoring hypopneas: impact on the apnea hypopnea index. *Sleep*. 2009; 32 :150–157. [PubMed: 19238801]
- Sethuraman J. A constructive definition of Dirichlet priors. *Statistica Sinica*. 1994 :639–650.
- Shumway, RH, Stoffer, DS. *Time series analysis and its applications: with R examples*. Springer; 2017.
- Stephens M. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2000; 62 :795–809.
- Teh YW, Jordan MI, Beal MJ, Blei DM. Hierarchical Dirichlet Processes. *Journal of the American Statistical Association*. 2006; 101 :1566–1581.
- Teran-Santos J, Jimenez-Gomez A, Cordero-Guevara J, Burgos-Santander CG. The association between sleep apnea and the risk of traffic accidents. *New England Journal of Medicine*. 1999; 340 :847–851.
- Tripuraneni N, Gu SS, Ge H, Ghahramani Z. Particle Gibbs for infinite hidden Markov models. *Advances in Neural Information Processing Systems*. 2015 :2395–2403.
- Van Gael, J; Saatci, Y; Teh, YW; Ghahramani, Z. *Proceedings of the 25th international conference on Machine learning*; ACM; 2008. 1088–1095.
- West M, Prado R, Krystal AD. Evaluation and comparison of EEG traces: Latent structure in nonstationary time series. *Journal of the American Statistical Association*. 1999; 94 :375–387.
- Whittle P. Curve and periodogram smoothing. *Journal of the Royal Statistical Society Series B (Methodological)*. 1957; 19 :38–63.
- Yaggi HK, Concato J, Kernan WN, Lichtman JH, Brass LM, Mohsenin V. Obstructive sleep apnea as a risk factor for stroke and death. *New England Journal of Medicine*. 2005; 353 :2034–2041.
- Yaghoubi F, Sunderam S. Quasi-supervised scoring of human sleep in polysomnograms using augmented input variables. *Computers in Biology and Medicine*. 2015; 59 :54–63. [PubMed: 25679475]
- Yau C, Papaspiliopoulos O, Roberts GO, Holmes C. Bayesian non-parametric hidden Markov models with applications in genomics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2011; 73 :37–57. [PubMed: 21687778]

- Young T, Peppard PE, Gottlieb DJ. Epidemiology of obstructive sleep apnea: a population health perspective. *American Journal of Respiratory and Critical Care Medicine*. 2002; 165 :1217–1239. [PubMed: 11991871]
- Yu S-Z. Hidden semi-Markov models. *Artificial Intelligence*. 2010; 174 :215–243.
- Zhou Y, Johansen AM, Aston JA. Toward automatic model comparison: an adaptive sequential Monte Carlo approach. *Journal of Computational and Graphical Statistics*. 2016; 25 :701–726.

Fig 1. Airflow trace collected over a period of five and half minutes of continuous breathing where instances of simulated apnea and hypopnea (highlighted on the graph) were recurring over time.

Fig 2.

Illustrative Example. Dots represent the simulated time series, where the different colors corresponds to (true) different regimes. The state-specific estimated oscillatory mean function is displayed as a solid curve, and the estimated state sequence as a piecewise horizontal line at the top part of the graph.

Fig 3. Illustrative Example.

(a) Trace plots (after burn-in of 3000 updates) for posterior sample of frequencies, conditional on modal number of states and number of frequencies in each state; red lines correspond to true values of the frequencies. (b) Trace plots (including burn-in) of the likelihood for three Markov chains initialized at different starting values (where the initial 100 updates are omitted from the graph)

Fig 4. Signal extraction with non-Gaussian innovations.

Boxplots of the MSE values for AutoNOM and our proposed HHMM when (a) the data exhibit recurrent patterns; (b) the data do not exhibit recurrent patterns; (c) the innovations are skewed.

Fig 5.

(Top) A realization from model (16), where the piecewise horizontal line represents the true state sequence. (Bottom) True time varying frequency peak (dotted red line) and the estimate provided by our proposed approach (solid blue line) where we highlight a 95% credible interval obtained from the posterior sample. (Right) Boxplots of the MSE values for AutoNOM, our proposed HHMM and AdaptSPEC.

Fig 6.

Case Study. Dots represent the airflow trace collected over a period of five and half minutes of continuous breathing. The grey lines represent draws from the estimated posterior predictive. The piecewise horizontal line corresponds to the estimated state sequence where we highlight the states corresponding to estimated apnea (red) and hypopnea (blue), while reporting the ground truth at the top of the plot.

Table 1
Illustrative example. (left panel) posterior probabilities for number of distinct states k ;
(right panel) posterior probabilities for number of frequencies in each state, conditioned
on $k = 3$.

k	$\hat{\pi}(k y)$
1	0.00
2	0.00
3	0.99
4	0.01
5	0.00
6	0.00
7	0.00

m	$\hat{\pi}(d_1 k = 3, y)$	$\hat{\pi}(d_2 k = 3, y)$	$\hat{\pi}(d_3 k = 3, y)$
1	0.99	1.00	0.01
2	0.01	0.00	0.99
3	0.00	0.00	0.00
4	0.00	0.00	0.00
5	0.00	0.00	0.00

Table 2
Illustrative Example. Estimated posterior mean (and standard deviation) of frequencies and square root of the power of the corresponding frequencies.

	ω_{11}	ω_{21}	ω_{31}	ω_{32}
True	0.0400	0.0526	0.0833	0.1250
Estimated	0.0399 ($8.8 \cdot 10^{-6}$)	0.0526 ($6.3 \cdot 10^{-6}$)	0.0833 ($9.6 \cdot 10^{-6}$)	0.1249 ($9.4 \cdot 10^{-6}$)
	A_{11}	A_{21}	A_{31}	A_{32}
True	1.131	0.283	1.414	1.414
Estimated	1.069 (0.029)	0.281 (0.004)	1.380 (0.022)	1.367 (0.022)

Table 3
Case study. Posterior mean and standard deviation of frequencies and corresponding powers that characterize the two states classified as apnea and hypopnea.

Apnea		Hypopnea	
Freq	Power	Freq	Power
0.0159 ($6.66 \cdot 10^{-5}$)	0.1376 ($1.81 \cdot 10^{-2}$)	0.0455 ($1.29 \cdot 10^{-4}$)	0.261 ($3.37 \cdot 10^{-2}$)
0.0353 ($3.73 \cdot 10^{-5}$)	0.2153 ($1.87 \cdot 10^{-2}$)	0.0542 ($4.72 \cdot 10^{-5}$)	0.620 ($3.4 \cdot 10^{-2}$)
0.0379 ($3.89 \cdot 10^{-5}$)	0.2065 ($1.97 \cdot 10^{-2}$)	-	-