Published in final edited form as:

J Am Stat Assoc. 2019 ; 114(525): 453–465. doi:10.1080/01621459.2017.1415908.

Adaptive Bayesian Time-Frequency Analysis of Multivariate Time Series

Zeda Li, PhD and

Paul H. Chook Department of Information Systems and Statistics, Baruch College, City University of New York

Robert Krafty*

Department of Biostatistics, University of Pittsburgh

Abstract

This article introduces a nonparametric approach to multivariate time-varying power spectrum analysis. The procedure adaptively partitions a time series into an unknown number of approximately stationary segments, where some spectral components may remain unchanged across segments, allowing components to evolve differently over time. Local spectra within segments are fit through Whittle likelihood based penalized spline models of modified Cholesky components, which provide flexible nonparametric estimates that preserve positive definite structures of spectral matrices. The approach is formulated in a Bayesian framework, in which the number and location of partitions are random, and relies on reversible jump Markov chain and Hamiltonian Monte Carlo methods that can adapt to the unknown number of segments and parameters. By averaging over the distribution of partitions, the approach can approximate both abrupt and slow-varying changes in spectral matrices. Empirical performance is evaluated in simulation studies and illustrated through analyses of electroencephalography during sleep and of the El Niño-Southern Oscillation.

Keywords

Locally Stationary Process; Modified Cholesky Decomposition; Nonstationary Multivariate Time Series; Reversible Jump Markov Chain Monte Carlo; Penalized Splines; Spectral Analysis

1 Introduction

Understanding second-order frequency domain properties of multivariate time series is essential to addressing scientific questions in a variety of fields. The time series of interest are often nonstationary and the manner in which frequency domain properties evolve over time can provide important scientific information. An example of such data comes from sleep medicine, where physicians commonly record patients' brain activity during sleep via

* (rkrafty@pitt.edu).

Supplementary Material

Supplementary materials are available online, including a pdf file that describes convergence assessment of the RJMCMC, further details of the analysis of the EEG data discussed in Section 5.1, additional simulation studies, and full details of the sampling scheme. Both R and Matlab code for implementing the proposed procedure are provided.

multi-channel electroencephalography (EEG) to inform the treatment of sleep disorders. Brain activity is dynamic, so that the data are inherently nonstationary, and the frequency domain properties of EEG provide interpretable biological information. An example of an epoch of two-channel EEG is displayed in Figure 1.

Methods for estimating the time-varying spectrum of a multivariate time series can be roughly grouped into three categories. The first category consists of estimators in which second-order frequency domain structures evolve continuously over time. This includes the slowly-evolving multivariate locally stationary process, which can be analyzed parametrically by fitting time series models with time-varying parameters (Dahlhaus, 2000), and nonparametrically via the bivariate smoothing of spectral components as functions of frequency and time (Guo and Dai, 2006). Additional nonparametric methods based on locally stationary wavelet models have been developed that, although also assume that second-order structures evolve continuously, are better localized in time and frequency to capture more rapid but continuous changes (Sanderson et al., 2010; Park et al., 2014). The second category of methods consist of estimators that are piecewise stationary. Approaches within this second category typically divide a time series into approximately stationary segments, then obtain estimates of local spectra within segments. These methods include both parametric approaches, such as fitting piecewise vector autoregressive models (Davis et al., 2006), and nonparametric approaches, such as using the multivariate smooth localized complex exponential (SLEX) library (Ombao et al., 2005). The final category, which is the focus of this article, are adaptive Bayesian methods that can approximate both abrupt and slowly-varying changes.

A number of Bayesian approaches to spectrum analysis have been explored in the literature. These include methods for analyzing stationary (Carter and Kohn, 1996; Cadonna et al., 2017; Choudhuri et al., 2004; Rosen and Stoffer, 2007; Macaro and Prado, 2014; Krafty et al., 2017) and nonstationary (Rosen et al., 2012; Zhang, 2016; Bruce et al., 2017) time series. Adaptive spectral analysis was introduced by Rosen et al. (2012) as a Bayesian approach to univariate nonstationary spectrum analysis, and was latter extended to the multivariate nonstationary setting by Zhang (2016). Under this approach, a time series is adaptively partitioned into a random number of approximately stationary segments, local spectra are estimated within in approximately stationary segments, and time-varying spectral estimates are obtained by averaging local estimates over the distribution of partitions.

The property that spectral matrices must be non-negative definite makes the time-varying spectral analysis of a multivariate time series a challenging problem. In particular, to assure that estimates are positive definite, most existing approaches require a common temporal structure across spectral components. For piecewise stationary based estimates, this equates to all spectral components having the same partitions, and inhibits the accurate modeling of processes where one component, say the cross-spectrum between two series, changes while other components, such as the individual spectra of the two series, remain stationary. An exception to this limitation is the method of Guo and Dai (2006), which models Cholesky components of spectral matrices to assure estimates are positive definite. However, their approach is highly dependent on slowly-varying temporal smoothness and cannot capture rapid or abrupt changes.

The goal of this article is to expand the scope of processes that can be accurately analyzed by introducing a novel multivariate time-frequency analysis procedure that can capture both abrupt- and smoothly-changing spectral dynamics while flexibly allowing components to evolve differently in both time and frequency. We introduce an adaptive spectral analysis procedure that begins by approximating a multivariate locally stationary time series through piecewise stationary processes with a random number of segments and partitions. Using reversible jump Markov chain Monte Carlo (RJMCMC) and Hamiltonian Monte Carlo (HMC) techniques, the approach estimates the unknown number and location of partitions and fits penalized splines to the modified Cholesky components of local spectral matrices. By modeling modified Cholesky components, the method is able to allow some components to remain stationary across a partition point while producing positive definite estimates. The sample generated from the sampling algorithm can be used to investigate the number and location of change points and, by averaging estimates over the distribution of partitions, not only produces estimates that can capture abrupt changes, but also effectively approximates slowly-varying dynamics. To the best of our knowledge, this is the first methodology for multivariate spectral analysis that flexibly allows spectral components to evolve differently over time while capturing both abrupt and slowly-varying dynamics. In addition to allowing components to evolve differently over time, empirical examples suggest that the proposed approach is computationally faster and can produce more accurate estimates compared to other approaches that can approximate both abrupt and slowly-varying dynamics (Zhang, 2016).

The rest of this article is as follows. A model for time-varying multivariate power spectra is introduced in Section 2. In Section 3, prior distributions for model parameters are formulated and a sampling scheme is outlined. Simulation studies are presented in Section 4 to illustrate the proposed procedure's ability to capture both slow-varying and abrupt changes, to estimate the number and location of partitions, and to compare its performance to that of other estimators. The use of the procedure in real-world applications is illustrated in Section 5, where it is used to analyze EEG from a patient during sleep and to analyze the indicators of El Nino-Southern Oscillation. Concluding remarks are given in Section 6. A detailed description of the sampling scheme is given in the supplementary materials.

2 The Model

2.1 Time-Varying Spectrum

This article considers the time-varying spectral analysis of a locally stationary N -dimensional time series defined through a Cramér representation with time-varying transfer function. The use of time-varying transfer functions to generalize power spectra for stationary time series to the nonstationary setting was introduced by Priestley (1965) for univariate time series, and extended to multivariate time series by Dahlhaus (2000) and Guo and Dai (2006). A time-varying transfer function $A(u, \omega)$ is a function of scaled time $u \in [0, 1]$ and frequency $\omega \in \mathbb{R}$ such that $A(u, \omega)$ is a nonsingular $N \times N$ complex valued matrix that is periodic and Hermitian as a function of frequency, or $A(u, \omega) = A(u, \omega + 1)$ and $A(u, \omega) = A(u, -\omega)^*$, where A^* is the complex conjugate of A . Formally, we consider $\mathbb{R}^N$ -valued time series of length T , $\{X_t : t = 1, \dots, T\}$, of the form

$$X_t = \int_{-1/2}^{1/2} A(t/T, \omega) \exp(2\pi i \omega t) dZ(\omega),$$

for a time-varying transfer function $A(u, \omega)$ and an N -dimensional mean-zero orthogonal process $Z(\omega)$ that is Hermitian and where $E\{dZ(\omega)dZ^*(\zeta)\}$ is the identity matrix if $\omega = \zeta$ and zero otherwise. Regularity conditions on the distribution of Z must be assumed to assure that estimators are well behaved; we assume that the cumulants of dZ exist and are bounded for all orders (Guo and Dai, 2006, Definition 2.1.1).

The focus of this article is on estimating the time-varying spectrum

$$f(u, \omega) = A(u, \omega)A(u, \omega)^*, \quad u \in [0, 1], \omega \in \mathbb{R}.$$

The time-varying spectrum $f(u, \omega)$ is a positive definite Hermitian $N \times N$ matrix. We assume it satisfies some regularity conditions both as a function of ω and as a function of u . For every u , we assume that each component of $f(u, \cdot)$ possesses a square-integrable first derivative as a function of frequency. For every ω , we assume that each component of $f(\cdot, \omega)$ is continuous as a function of scaled time at all but a possible finite number of points.

We refer to this class of time series as locally stationary and note that it differs somewhat from the locally stationary time series models considered by Dahlhaus (2000) and Guo and Dai (2006). First, Dahlhaus (2000) assumed a series of transfer functions $A_{t,T}^0(\omega)$ that converge to a large sample transfer function $A(u, \omega)$ in order to allow for the fitting of time-varying parametric models. Since we are considering nonparametric estimation, in a manner similar to Guo and Dai (2006), we define our model directly using $A(u, \omega)$. Second, the models of Dahlhaus (2000) and Guo and Dai (2006) require the time-varying spectrum to be continuous in both time and frequency. The proposed model is more flexible and allows for components of spectral matrices to evolve not only continuously, but also abruptly in time.

2.2 Piecewise Stationary Approximation

Our analysis of a locally stationary time series $\{X_t : t = 1, \dots, T\}$ begins by noting that it can be well approximated by a piecewise stationary process (Adak, 1998; Guo and Dai, 2006). Consider a partition of the time series into m segments defined by partition points $\delta = (\delta_0, \dots, \delta_m)'$ with $\delta_0 = 0$ and $\delta_m = T$ such that X_t is approximately stationary within the segments $\{t : \delta_{q-1} < t \leq \delta_q\}$ for $q = 1, \dots, m$. Then

$$X_t \approx \sum_{q=1}^m \int_{-1/2}^{1/2} A_q(\omega) \exp(2\pi i \omega t) dZ(\omega),$$

where $A_q(\omega) = A(u_q, \omega) \mathbb{I}(\delta_{q-1} < t \leq \delta_q)$, $\mathbb{I}(\cdot)$ is the indicator function, and $u_q = (\delta_q + \delta_{q-1})/2T$ is the scaled midpoint of the q th segment. Within the q th segment, the time series is approximately second-order stationary with local power spectrum $f(u_q, \omega) = A_q(\omega)A_q(\omega)^*$.

Conditional on an approximately stationary partition δ , known approaches and properties for stationary time series can be applied. In particular, the large sample distribution of the discrete Fourier transform of a stationary time series (Brillinger, 2002, Theorem 4.4.1) that provides the Whittle likelihood (Whittle, 1953) allows for the formulation of a product of Whittle likelihoods. Define the local discrete Fourier transform at frequency ℓ within segment q as

$$y_{q\ell} = n_q^{-1/2} \sum_{t=\delta_{q-1}+1}^{\delta_q} X_t \exp(-2\pi i \omega_{q\ell} t), \quad \ell = 1, \dots, L_q, \quad q = 1, \dots, m,$$

where $n_q = \delta_q - \delta_{q-1}$ is the number of observations in segment q , $\omega_{q\ell} = \ell n_q$ are the Fourier frequencies, and $L_q = \lfloor (n_q - 1)/2 \rfloor$. Given a partitioning into m segments δ , $y_{q\ell}$ are approximately independent zero-mean complex multivariate Gaussian random variables with covariance matrices $f(u_q, \omega_{q\ell})$, leading to a log-likelihood that can be approximated by a sum of local log-Whittle likelihoods

$$\mathcal{L}(Y | f, \delta, m) \approx - \sum_{q=1}^m \sum_{\ell=1}^{L_q} \left\{ \log |f(u_q, \omega_{q\ell})| + y_{q\ell}^* f^{-1}(u_q, \omega_{q\ell}) y_{q\ell} \right\}.$$

Throughout this article, we use the notation where Y represents the collection of all local discrete Fourier transforms.

Two issues concerning $\mathcal{L}$ should be noted. First, the local Whittle likelihood is a large sample approximation whose validity depends on the data being sufficiently close to stationary within $(\delta_{q-1}, \delta_q]$ and on n_q being sufficiently large when data are non-Gaussian. A discussion of this issue is presented in Section 6 and empirical investigations are presented in supplementary material. Second, as opposed to methods such as those considered by Ombao et al. (2005) and Davis et al. (2006), where the partitions are treated as fixed, here the number of segments m and partition δ are random variables whose prior distributions are given in Section 3.1. Estimators and inference will be obtained by averaging over the posterior distribution of partitions and number of segments.

2.3 The Modified Cholesky Decomposition

The spectral matrix is positive definite and, to nonparametrically allow for flexible smoothing among the different components while preserving positive-definiteness, we model modified Cholesky components of local spectra via linear penalized splines. The use of other possible nonparametric local spectrum estimators is discussed in Section 6. The modified Cholesky decomposition represents a time-varying spectral matrix as

$$f^{-1}(u, \omega) = \Theta(u, \omega) \Psi(u, \omega)^{-1} \Theta(u, \omega)^*$$

for a complex-valued $N \times N$ lower triangular matrix $\Theta(u, \omega)$ with ones on the diagonal and a positive diagonal matrix $\Psi(u, \omega)$. For a piecewise stationary approximation with partition δ

into m segments, we define the local modified Cholesky decomposition as $\hat{\Gamma}^{-1}(u_q, \omega) = \Theta(u_q, \omega) \Psi(u_q, \omega)^{-1} \Theta(u_q, \omega)^*$, for $q = 1, \dots, m$, and let $\theta_{jkq}(\omega)$ and $\psi_{jjq}(\omega)$ be the jk and jj elements of $\Theta(u_q, \omega)$ and $\Psi(u_q, \omega)$, respectively. Then, for each segment, there are N^2 -components to estimate: $\Re\{\theta_{jkq}(\omega)\}$ for $j > k = 1, \dots, N$, $\Im\{\theta_{jkq}(\omega)\}$ for $j > k = 1, \dots, N$, and $\psi_{jjq}(\omega)$ for $j = 1, \dots, N$.

In addition to the restriction that spectral matrices are positive definite, they also possess restrictions as functions of frequency in that they are periodic and are Hermitian. This means that all components are periodic, real components of Θ and the diagonal components of Ψ are even functions, or $\Re\{\theta_{jkq}(\omega)\} = \Re\{\theta_{jkq}(-\omega)\}$ and $\psi_{jjq}(\omega) = \psi_{jjq}(-\omega)$, and imaginary components of Θ are odd, or $\Im\{\theta_{jkq}(\omega)\} = -\Im\{\theta_{jkq}(-\omega)\}$. We model the components using periodic even and odd linear splines by considering

$$\Re\{\theta_{jkq}(\omega)\} = c_{jkq0} + \sum_{s=1}^{S-1} c_{jkqs} \cos(2\pi s\omega), \quad (1)$$

$$\Im\{\theta_{jkq}(\omega)\} = \sum_{s=1}^S b_{jkqs} \sin(2\pi s\omega), \quad (2)$$

$$\log\{\psi_{jjq}(\omega)\} = d_{jjq0} + \sum_{s=1}^{S-1} d_{jjqs} \cos(2\pi s\omega). \quad (3)$$

The Fourier frequencies for each segment form an equally spaced grid, so that the Demmler-Reinsch bases for periodic even and odd smoothing splines for local periodograms are given by $\{\cos(2\pi s\omega); s = 0, \dots, (L_q - 1)\}$ and $\{\sin(2\pi s\omega); s = 1, \dots, L_q\}$, respectively (Schwarz and Krivobokova, 2016, Section 3). Since $\psi_{jjq}(\omega)$ is positive, it is modeled on the log-scale.

Not all Cholesky components are required to change at each partition point, enabling components to evolve differently over time. In Section 3.1, we define prior distributions for the components that change at each partition point, in addition to priors for the number and location of partition points and for the coefficients b_{jkqs} , c_{jkqs} and d_{jjqs} . Prior distributions for the coefficients are selected to regularize integrated squared first derivatives and formulate Bayesian linear penalized splines. As noted by Krafty and Collinge (2013) and by Zhang (2016), this Bayesian linear penalized spline model for local spectra differs somewhat from that used by Rosen and Stoffer (2007) for stationary time series. Rosen and Stoffer (2007) uses a model that is periodic, but not restricted to be odd or even. Accounting for these geometric restrictions has been shown to improve performance, especially at the boundary (Krafty and Collinge, 2013). The proposed model for local spectra is equivalent to the smoothing spline model for a stationary spectrum of Krafty and Collinge (2013) when S

$= L_q$. Since coefficients decay rapidly, using $S < L_q$ basis functions presents considerable computational savings without sacrificing model fit. In simulation studies and data analysis, we select $S = 10$, which accounts for at least 99.975% of the total variance of the full smoothing spline when $n_q = 10^4$ (Krafty et al., 2017).

3 Priors and Sampling Scheme

3.1 Priors

We first define prior distributions for the parameters of the temporal partitioning, then define priors for the parameters of the penalized spline model conditional on the partitioning. The time series is partitioned into m approximately stationary segments and we choose a discrete uniform $\mathcal{U}(1, M)$ prior for m , where M is a fixed large integer that represents a maximum number of possible segments. To assure that the local Whittle likelihood approximation holds, we chose a minimum number of time points per segment $n_{\min}$ and select $M = \lfloor T/n_{\min} \rfloor$. Conditional on the number of segments m , the prior for the partition δ is

$$\Pr(\delta | m) = \prod_{q=1}^{m-1} \Pr(\delta_q | \delta_{q-1}, m),$$

where $\Pr(\delta_q | \delta_{q-1}, m) = 1/\alpha_q$ and $\alpha_q = T - \delta_{q-1} - (m-q+1)n_{\min} + 1$ is the number of possible locations for the q th partition point. This prior places equal weight on all possible locations of a partition point conditional on previous partition points. Some Cholesky components can remain unchanged across successive segments to enable flexible smoothness across time; we assume a uniform prior for the set of components that change at a given partition point. More specifically, at least one of the N^2 -Cholesky components is required to change across a partition point, so that there are a total of $2^{N^2} - 1$ possible sets of Cholesky components that can change. Letting ϕ_q be the set of N^2 -Cholesky components that change at δ_q , we assume a prior distribution on ϕ_q such that probability of any nonempty set of the N^2 -components is $1/(2^{N^2} - 1)$.

The Cholesky components of local spectra within segments are modeled using the Demmler-Reinsch basis for linear splines so that smoothing priors that regularize the integrated square first derivative can be formulated as Gaussian priors with diagonal covariance matrices (Rosen and Stoffer, 2007). Let $\mathbf{c}_{jkq} = [c_{jkq0}, \dots, c_{jkq(s-1)}]'$, $\mathbf{b}_{jkq} = [b_{jkq1}, \dots, b_{jkqs}]'$, and $\mathbf{d}_{jjq} = [d_{jjq0}, \dots, d_{jjq(s-1)}]'$ be vectors of coefficients of basis functions for the Cholesky components within the q th segment defined in (1), (2), and (3). The priors on $\mathbf{C}_{jkq}$, $\mathbf{b}_{jkq}$, and $\mathbf{d}_{jjq}$ are $N(0, D_{cjkq})$, $N(0, D_{bjkq})$ and $N(0, D_{djjq})$, respectively, where

$$D_{cjkq} = \text{diag}[\sigma_{\alpha}^2, \lambda_{cjkq}^2 (2\pi)^{-2}, \dots, \lambda_{cjkq}^2 \{2\pi(S-1)\}^{-2}],$$

$$D_{bjkq} = \text{diag}\{\lambda_{bjkq}^2 (2\pi)^{-2}, \dots, \lambda_{bjkq}^2 (2\pi S)^{-2}\},$$

$$D_{djjq} = \text{diag}[\sigma_{\alpha}^2, \lambda_{djjq}^2 (2\pi)^{-2}, \dots, \lambda_{djjq}^2 \{2\pi(S-1)\}^{-2}].$$

It should be noted that, in a manner similar to Krafty and Collinge (2013), the smoothing parameters λ_{cjkq}^2 , λ_{bjkq}^2 and λ_{djjq}^2 are scaled by $(2\pi s)^{-2}$; this scaling was not used by Rosen and Stoffer (2007) or Zhang (2016). This adjustment is needed to produce penalized splines that regularize the integrated square first derivative of Cholesky components. Conditional on a smoothing parameter λ , the integrated square first derivative of a Cholesky component has a prior distribution that is proportional to a chi-square. A smoothing parameter controls the roughness of a local Cholesky component such that, as its smoothing parameter is shrunk towards zero, the component tends towards a constant function of frequency with probability 1. Based on Gelman (2006), we assume priors for λ_{cjkq}^2 , λ_{bjkq}^2 and λ_{djjq}^2 to be Uniform(0, κ), where κ is a known large constant. The parameter σ_{α}^2 is a hyperparameter for the prior variance of the intercepts of the Cholesky components and is set to be a fixed large value. We found results to be insensitive to the selection of hyperparameters, with $\kappa \in (10^3, 10^6)$ and $\sigma_{\alpha}^2 \in (10^3, 10^7)$ providing indistinguishable results in all empirical examples considered.

3.2 Description of the Sampling Scheme

We present a scheme based on RJMCMC (Green, 1995) and HMC (Neal, 2011; Gelman et al., 2013) techniques to sample from the joint posterior distribution of parameters. Each iteration consists of two types of moves: within-model and between-model moves. We briefly describe the sampling scheme in this section and provide technical details in supplementary material.

Between-model moves involve randomly proposing either a death where the number of segments m decrease by 1, or a birth where m increases by 1. Parameters drawn within the move are jointly accepted or rejected. We denote m^p and m^c as the proposed and current number of segments respectively. If a birth move is proposed, $m^p = m^c + 1$, and a new partition point is drawn by first sampling a current segment to be split, then randomly selecting the new partition point within the segment. Conditional on m^p and the new partition point, the set of Cholesky components that change at the new partition point, ϕ_q , is sampled from the discrete uniform distribution. Then, in each of the two new segments, the smoothing parameters corresponding to the Cholesky components that were selected to change between the segments are proposed. Finally, coefficients of basis functions of the selected Cholesky components for both new segments are drawn from approximated conditional Gaussian distributions.

If a death is proposed, then $m^p = m^c - 1$. A partition is randomly selected to be deleted and a new segment is formed by combining the two neighboring segments separated by this point. Smoothing parameters for components that were previously different across these two segments are used to form the corresponding smoothing parameters in the newly combined segment. Then, the corresponding coefficients of basis functions of the combined segments are drawn.

Within-model moves involve no change in the number of partitions or segments, namely, $m^p = m^c$. First, a partition δ_q is randomly selected and relocated. Then, basis function coefficients of all Cholesky components within the segments impacted by the move are drawn. To efficiently explore local parameter spaces, Hamiltonian Monte Carlo (HMC) is used to draw coefficients of basis functions of the Cholesky components. Lastly, impacted smoothing parameters are drawn from conditional inverse gamma distributions.

It is worth noting that our proposed method has a number of differences compared to the Bayesian method proposed by Zhang (2016). First, the proposed model and sampling scheme allow some components to remain stationary across a partition point through the parameter ϕ . This allows estimated spectral components to have different temporal structures. In a manner similar to other existing methods (Ombao et al., 2001; Davis et al., 2006), the method of Zhang (2016) requires all components to change at each partition point. Second, the unknown number and location of partitions are sampled via two different methods, RJMCMC and SSAMC, respectively. While SSAMC has been shown to improve upon the convergence of RJMCMC in some circumstances, particularly in highly complex sample spaces, empirical results suggest that this is not necessarily the case in many examples of nonstationary multivariate spectrum analysis. Liang (2009) noted that the smoothing steps in SSAMC could potentially hinder estimation when there exist abrupt jumps between different models, which is the case when a time series is piecewise stationary with abrupt changes. In this setting, one model is of most interest. Thus, once the sampler arrives at the true number of segments and roughly correct partition locations, ideally the sampler will stay at this “correct model” with a high probability rather than frequently jumping to other models. On the other hand, when we have a time series with a slowly-varying spectrum, SSAMC could be superior to RJMCMC in terms of moving between different models. However, for an average of piecewise models to more accurately approximate a slow-varying spectrum, one desires each of the piecewise models to have different partition points. In other words, one desires partitions to have high a probability of relocating. Since within-model moves are responsible for relocating partitions, they are more crucial in obtaining estimates that, in a frequentist sense, effectively approximate slowly-varying dynamics. Lastly, in the proposed method, within-model moves are sampled via HMC. The HMC techniques enable desirable within-model convergence and acceptance rates, and consequently enable the location of partition points to vary with sufficient probability for final estimates averaged over all iterations to effectively approximate slowly-varying changes.

4 Simulation Studies

In this section, we perform simulation studies under three settings to investigate the frequentist characteristics of the proposed method and to compare it to existing methods. The goal of the first two settings is to investigate performance in estimating time-vary spectra, first for a trivariate piecewise stationary process with one partition and then for a bivariate slow-varying process. The goal of the third setting is to investigate properties in identifying partition points for a piecewise stationary process with multiple partitions of

varying magnitudes. For all models, results are reported where hyperparameters of the proposed method were chosen as $n_{\min} = 60$, $\sigma_{\alpha}^2 = 10^4$ and $\kappa = 10^5$.

4.1 Piecewise Stationary Process

We simulated 250 independent trivariate piecewise stationary time series of length $T = 600$ as

$$X_t = \begin{cases} \epsilon_t + \Phi_{11}\epsilon_{t-1} + \Phi_{12}\epsilon_{t-2} & \text{if } 1 \leq t \leq 300 \\ \epsilon_t + \Phi_{21}\epsilon_{t-1} + \Phi_{22}\epsilon_{t-2} & \text{if } 301 \leq t \leq 600, \end{cases}$$

where

$$\Phi_{11} = \begin{pmatrix} 0.6 & 0 & 0 \\ 0.2 & -0.5 & 0 \\ 0.1 & 0.3 & 0.4 \end{pmatrix}, \quad \Phi_{21} = \begin{pmatrix} 0.6 & 0 & 0 \\ 0.2 & 0.5 & 0 \\ -0.1 & -0.3 & 0.4 \end{pmatrix},$$

$\Phi_{12} = \Phi_{22} = \text{diag}(0.3, 0.3, 0)$, and ϵ_t are independent zero-mean trivariate Gaussian random variable whose components have unit variance and pairwise correlation 0.5. The true power spectrum is $f(u, \omega) = \Phi(u, \omega) \Sigma \Phi(u, \omega)^*$ where $\Phi(u, \omega) = I + \Phi_{11} \exp(-2\pi i \omega) + \Phi_{12} \exp(-4\pi i \omega)$ for $u \in [0, 1/2]$, $\Phi(u, \omega) = I + \Phi_{21} \exp(-2\pi i \omega) + \Phi_{22} \exp(-4\pi i \omega)$ for $u \in (1/2, 1]$, and $\Sigma = \text{Cov}(\epsilon_t)$ (Priestley, 1981, Chapter 9.4). A simulated realization of this process is shown in Figure 2.

We investigate the performance of an estimator of a spectral matrix $f(u, \omega)$ by calculating the average squared error (ASE) of the estimated individual spectra, $f_{11}(u, \omega)$, $f_{22}(\eta, \omega)$, and $f_{33}(u, \omega)$, and the estimated coherences

$$\rho_{jk}^2(u, \omega) = \left| f_{jk}(u, \omega) \right|^2 / \left\{ f_{jj}(u, \omega) f_{kk}(u, \omega) \right\}, \quad \text{for } j = 2, 3 \quad \text{and} \quad \forall k < j.$$

The ASE is time points obtained by averaging square errors across the equally spaced grid of $T=600$ time points and 51 frequencies. For example, the ASE of $\hat{f}_{11}(\mu, \omega)$ is computed as

$$\text{ASE} = [T(L+1)]^{-1} \sum_{t=1}^T \sum_{\ell=0}^L \left[\hat{f}_{11}(t/T, \omega_{\ell}) - f_{11}(t/T, \omega_{\ell}) \right]^2,$$

where $L = 50$ and $\hat{f}_{11}(u, \omega)$ is the estimated posterior mean of $f_{11}(u, \omega)$. The individual spectra, coherences, and estimates from the realization displayed in Figure 2 using the proposed method are given in Figure 3.

For each simulated time series, the time-varying power spectrum was estimated by three methods: the proposed method with 12000 iterations and a burn-in period of 4000 iterations, the piecewise vector autoregressive method of Davis et al. (2006) with default tuning parameters, and the nonparametric smoothing spline ANOVA approach of Guo and Dai

(2006) with 10 prespecified blocks. The means and standard deviations of the ASE are presented in Table 1. It should be noted that the SSAMC-based method of Zhang (2016), whose empirical performance is investigated in Section 4.3 for a smaller number of replicates, was not implemented in this setting due its high computational cost, which makes its implementation impractical for 250 replicates.

Unsurprisingly, the smoothing spline ANOVA estimator, which is based on a slowly-varying model, exhibited poor performance relative to the other estimators in this abruptly changing setting. This is particularly true when estimating f_{22} , ρ_{31}^2 and ρ_{32}^2 , which experience abrupt and drastic changes. The proposed method had smaller mean ASE compared to the piecewise vector autoregressive method for all components. This can be attributed to the fact that, while the piecewise vector autoregressive estimator requires all components to change at each partition point, the proposed approach enables some components to remain unchanged. This allows the proposed method to more accurately estimate components that remain constant across time.

In addition, percentiles of the empirical distribution of the simulated sample provide a natural means to conduct inference on any function of spectral matrices. To evaluate frequentist performance in conducting inference, we computed 95% pointwise credible intervals and evaluated average coverage across time-frequency points. The mean (standard deviation) coverage of the 95% credible intervals for f_{11} , f_{22} , f_{33} , ρ_{21}^2 , ρ_{31}^2 , and ρ_{32}^2 were 95.2% (2.8%), 94.1% (3.4%), 95.8% (2.7%), 93.5% (3.2%), 92.7% (2.4%), and 92.4% (2.3%), respectively.

4.2 Slow-varying Process

In this section, we considered a bivariate process that possesses a slow-varying power spectrum. We simulated 250 time series as

$$X_t = \epsilon_t + \Phi_{1t}\epsilon_{t-1} + \Phi_{2t}\epsilon_{t-2}, \quad t = 1, \dots, 1024,$$

with

$$\Phi_{1t} = \begin{pmatrix} \phi_1(t) & -1 \\ -1 & \phi_2(t) \end{pmatrix}, \quad \Phi_{2t} = \begin{pmatrix} 0.5 & 0 \\ 0 & -1.2 \end{pmatrix},$$

where $\phi_1(t) = 1.122 \{1 - 1.781\sin(\pi t/2048)\}$, $\phi_2(t) = 1.122 \{1 - 1.781\cos(0.8\pi t/2048)\}$, and ϵ_t are independent zero-mean bivariate Gaussian random variables whose components have unit variance and pairwise correlation 0.2. A realization of the time series is plotted in Figure 4; the true time-varying individual spectra, coherence and estimates under the proposed procedure from this realization are displayed in Figure 5. As in Section 4.1, three methods were compared: the proposed method with 10000 iterations and a burn-in period of 2000 iterations, the piecewise vector autoregressive method of Davis et al. (2006) with default tuning parameters, and the nonparametric smoothing spline ANOVA approach of Guo and Dai (2006) with 16 predefined blocks. Table 2 displays the means and standard

deviations of the ASE for each of the three estimators. Both the proposed method and the smoothing spline ANOVA method outperformed the piecewise vector autoregressive approach in all components. Although the proposed procedure and the piecewise autoregressive estimator are both based on piecewise stationary approximations, the proposed method adaptively averages over the distribution of partitions to better capture the slowly-varying structure. The proposed method, which uses the Whittle likelihood, performed slightly better than the smoothing spline ANOVA approach, which utilizes penalized sums of squares. This is consistent with findings for stationary univariate (Pawitan and O'Sullivan, 1994), nonstationary univariate (Qin and Wang, 2008) and stationary multivariate (Krafty and Collinge, 2013) time series, in which Whittle likelihood based methods were found to be more efficient compared to sums of squares based methods. The mean (standard deviation) of the average coverage of 95% pointwise credible intervals for f_{11} , f_{22} , and ρ_{21}^2 under the proposed method are 96.7% (2.3%), 96.5% (2.1%), 94.8% (4.6%), respectively.

4.3 Estimating Multiple Partitions

In this section, the main focus is to investigate the frequentist properties of the proposed method in estimating the number and location of partition points, and to compare them to the SSAMC-based method of Zhang (2016). We simulate 20 time series from the following bivariate piecewise autoregressive process

$$X_t = \Phi_1 X_{t-1} + \Phi_2 X_{t-2} + \epsilon_t, \quad (4)$$

where

$$\begin{cases} \Phi_1 = \text{diag}(0.5, -0.6), \Phi_2 = \text{diag}(0, -0.5), \epsilon_t = \epsilon_t^{(1)} & \text{if } 1 \leq t \leq 400, \\ \Phi_1 = \text{diag}(0.5, 0.6), \Phi_2 = \text{diag}(0, -0.5), \epsilon_t = \epsilon_t^{(1)} & \text{if } 401 \leq t \leq 5000, \\ \Phi_1 = \text{diag}(0.5, 0.6), \Phi_2 = \text{diag}(0, -0.5), \epsilon_t = \epsilon_t^{(2)} & \text{if } 5001 \leq t \leq 10000, \\ \Phi_1 = \text{diag}(1.32, 0.6), \Phi_2 = \text{diag}(-0.81, -0.5), \epsilon_t = \epsilon_t^{(2)} & \text{if } 10001 \leq t \leq 12000, \end{cases}$$

and $\epsilon_t^{(1)}, \epsilon_t^{(2)}$ are independent zero-mean bivariate Gaussian random variables whose components have unit variance and pairwise correlation 0.5 and 0.8, respectively. The partitions are located far apart from each other, and the change at $t = 5000$, where there is only a slight change in off-diagonal elements of the error term, is subtle compared to the other changes.

Both methods were run using $S = 10$ basis functions for 10,000 iterations with a burn-in of 2,000. SSAMC is able to draw multiple samples in a single iteration and we consider three cases for SSAMC: drawing 1 sample, drawing 5 samples, and drawing 20 samples. Other parameters were selected as the default suggested by Zhang (2016). Table 3 reports the mean and standard deviation of estimated posterior probabilities of numbers of segments. The proposed method had highest estimated posterior probability of the correct number of

segments with $\widehat{\Pr}(m = 4|X) = 0.794$. The SSAMC sampler regularly transitioned between $m = 4, 5, 6$ segments, with $\widehat{\Pr}(m = 4|X)$ increasing with the number of samples. This finding coincides with the findings of Liang (2009), in which SSAMC has a higher probability of sampling from different models compared to RJMCMC. The proposed estimator had smaller ASE compared to SSAMC, which could possibly be attributed to the proposed method sampling more often from models with the true number of $m = 4$ segments compared to SSAMC. As more samples are drawn in each iteration, the ASE of SSAMC decreases. Table 3 also reports the mean and standard deviation of run times (second) per iteration using R Version 3.4 and Windows 10 on a desktop computer with a 3.6 GHz Intel Core i7 processor and 8 GB RAM. The proposed method was observed to be considerably more efficient in terms of computation time.

5 Applications

5.1 Analysis of EEG During Sleep

Neurophysiological activity during sleep is hypothesized to be responsible for many of sleep's rejuvenating properties. Consequently, to inform individualized treatments of disturbed sleep, physicians often use EEG to record patient's brain activity during sleep. We consider data from one such patient at the Sleep Disorders Clinic at St Vincent's University Hospital, Dublin who was being treated for obstructive sleep apnea.

As one sleeps, the brain cycles between periods of rapid-eye movement sleep (REM), which is characterized by movement of the eyes and is when dreaming is more likely to occur, and non-rapid-eye movement sleep (NREM), which includes deep sleep. Brain activity is inherently different during the different periods of sleep, so that the data are nonstationary; activity within each period and when transiting between periods plays a fundamental role in how sleep affects health and functioning. In clinical practice, EEG are visually inspected by trained technicians to identify periods of REM and NREM sleep (Iber et al., 2007). To gain a deeper understanding of brain activity during the transition from visually scored NREM to REM, we apply the proposed method to a 20 second long epoch of 2-channel (bivariate) EEG, with 10 seconds before and 10 seconds after the first visually scored transition between NREM into REM sleep. The EEG signal was recorded from bilateral central EEG leads with C3 and C4 referenced to linked mastoids A1 and A2 at a sampling rate of 128 HZ. The data were linearly detrended, filtered using a 0.02 HZ high-pass filter and a median filter for artifact removal, and standardized to unit variance. The resulting data, which are displayed in Figure 1, are a $N = 2$ dimensional time series of length $T = 2560$.

The proposed method was fit with hyperparameters $\sigma_\alpha^2 = \kappa = 10^5$, and run for 10,000 iterations with burn-in of 2,000. The mean (standard deviation) for the run time per iteration under the computer specification reported in Section 4.3 was 0.463 (0.189) seconds. Figure 6 displays estimated log-spectra and coherence. Spectra are displayed on the log scale to aid visualization. There is an abrupt change both in the spectra and in the coherence. The change occurs at $t = 1250$ (9.75 second) with a probability of 0.94, which is close to the visually scored transition from NREM sleep to REM sleep at $t = 1280$ (10 seconds). All components are nearly stationary before and after this change. It should be noted that the proposed

method provides an automated analysis of the nature of the temporal dynamics, providing evidence that the transition between sleep stages is sudden as opposed to gradually evolving.

Broadly, our results indicate that power at low frequencies decreases from NREM to REM sleep. The estimated log-spectra during NREM shows elevated power at 10–14 HZ. Known as sleep spindles, activity within these frequencies are a common characteristic of deep sleep. During REM, we observe characteristic peaks at beta power (16–31 HZ). In addition to these characteristics of univariate spectra, the analysis also uncovered changes in coherence, where peak coherence shifts from the delta band of frequencies less than 4 Hz during NREM to the beta band of 16–31 HZ during REM. Further analyses of these data and convergence diagnostics are provided in supplementary material.

5.2 Analysis of the El Niño-Southern Oscillation

The El Niño-Southern Oscillation (ENSO) is an irregularly periodical disruption in the climate system over the Tropical Pacific, which has impacts on global weather and climate systems. The process has not only been extensively studied in the geological literature, but also in the statistics literature (Shumway and Stoffer, 2011; Rosen and Stoffer, 2007; Rosen et al., 2012). However, these analyses have either been based on univariate time series analyses of individual ENSO indicators, or restricted by stationarity assumptions. In this section, we jointly consider three important indicative time series of ENSO observed from 1951 to 2016 (see Figure 7). The first indicator is the monthly Southern Oscillation Index (SOI), measuring mean sea-level pressure differences between Tahiti and Darwin station. The second indicator is Nino3.4, the monthly average sea surface temperature in the region 5S-5N, 120W-170W. The third indicator is the monthly mean sea-level pressure anomalies at Darwin station (DSLPA).

We ran the sampling procedures for 12,000 iterations with burn-in of 5,000 iterations. The mean (standard deviation) for the run time per iteration under the computer specification reported in Section 4.3 was 0.980 (0.505) seconds. The estimated posterior probability of the number of segments is presented in Table 4. The mode is $m = 2$ segments, which has an estimated posterior probability of 0.652, and the estimated posterior probability of $m = 1$ segment, or stationarity, is 0.329. The estimated posterior distribution of the location of partitions conditional on $m = 2$, $\hat{Pr}(\delta_{1,2} = t \mid x)$, is shown in Figure 8. The estimated conditional (on $m = 2$) posterior probability of this point occurring between 1980–1989 is 0.940, with non-zero probability across the entire interval, but with higher probability closer to 1980. These results suggest that the spectrum could have experienced changes in the 1980's.

Plots of estimated time-varying log-spectra and pairwise coherences appear in Figure 9. They indicate fairly pronounced changes in spectra of SOI and DSLPA, while the spectrum of Nino 3.4 seems close to stationary. The spectra of SOI and DSLPA, and all coherences, show shifts in power from low frequencies to higher frequencies. Further, the integrated total power of SOI increased. These are indicative of a higher rate of ENSO occurrences in recent years.

Previous analyses of ENSO have produced contradictory results. In the geology literature, the analyses of Solow (2006) and Nicholls (2008) found ENSO to be stationary. These results were further supported by the analysis of Rosen et al. (2012), who estimated that the posterior probabilities of stationarity to be 0.95, 0.93, and 0.99 for the SOI, Nino3.4, and DSLPA indices, respectively. However, the analyses of Trenberth and Hoar (1996) and Timmermann et al. (2004) indicate a change point between 1980–1989, and the analyses of Vecchi and Wittenberg (2010) and L’Heureux et al. (2013) suggest that the variability and amplitude in SOI has increased since 1980. Our results are more consistent with these analyses that suggest changes in ENSO during the 1980’s. There are several possible reasons for the difference between our findings and those of Rosen et al. (2012). First, we used data from 1951 to 2016, which contains five additional years of recent observation, and thus increases the ability to detect changes. Second, our method provides a joint analysis of the three series, taking into account possible changes in the coherence, while Rosen et al. (2012) analyzed each series univariately. It is possible that the change in the SOI series is too subtle to be individually identified. However, by simultaneously utilizing information in the coherence and individual spectra, our method could increase the power to detect smaller second-order changes. Lastly, it should be noted that, although our analysis can be interpreted as suggesting a change in the ENSO spectrum during the 1980’s, the posterior probability of stationarity was estimated as a nontrivial 0.332. In this sense, our analysis is consistent with the geology literature as a whole, where several researchers have reported a change in ENSO during the 1980’s, while others have reported ENSO being stationary over this period of time.

6 Concluding Remarks

We conclude this article with a discussion of some limitations and possible extensions of the proposed methodology. First, theoretically, the proposed methodology provides a means of conducting time-varying spectrum analysis on an N -dimensional time series of general dimension N . However, the number of Cholesky components grows quadratically with N . As illustrated in Section 3.3 of the supplementary materials, the proposed sampling scheme could be impractical for high-dimensions. A possible extension of the proposed method for the analysis of higher-dimension time series could be through the incorporation of time-dependent frequency domain principal components (Ombao and Ho, 2006) within the sampling algorithm. Second, the proposed procedure is designed for the analysis of a single multivariate time series. Many applications are concerned not with the analysis of a single multivariate time series, but of the analysis of replicated multivariate time series and in how their time-varying spectra are associated with other variables, such as clinical outcomes. The method developed by Fiecas and Ombao (2016) enables such an analysis when dynamics are smooth. A possible extension of the proposed procedure to the replicated time series setting could involve simultaneously dividing the grid of time and clinical outcome values into approximately stationary blocks in a manner similar that which was recently proposed in the univariate setting by Bruce et al. (2017). Third, the proposed approach relies upon the validity of the large sample distribution of the local Fourier transform. This large sample approximation requires data within each segment to be approximately stationary and for there to be a sufficient number of time points within each block when data are non-Gaussian.

The approximation could be invalid if the spectrum changes smoothly but extremely rapidly relative to the sampling rate; care should be taken when interpreting the results in such scenarios. Empirical investigations of this issue are presented in Sections 3.1 and 3.2 of the supplemental materials. A correction to the Whittle likelihood for non-Gaussian data was developed for parametric models by Kirch et al. (2017) and a possible topic for future research is the development of a correction for our nonparametric model. Fourth is another issue that can arise when approximately stationary blocks are small. The Gaussian approximation (see Section 3.2) provides an efficient and flexible way to sample coefficients of basis functions when the number of observations in a stationary block is moderate or large (Le Cam, 1953; Le Cam and Yang, 1990). For a time series with small stationary blocks, more advanced distributional approximation approaches, such as variational Bayes and expectation propagation, could be adapted, but with an increase in computational complexity. Lastly, the procedure models the components of local spectra using penalized splines. This approach assumes smoothness in spectra as functions of frequency, but other bases and regularizing priors could be used to adapt to specific features of a given problem. For instance, Bayesian wavelet shrinkage (Chipman et al., 1997) could be used if near line spectra are expected or B-spline bases with knots selected to focus attention within certain frequency bands could be used if it is scientifically known a priori that power is contained within certain frequencies.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NIH grant R01GM113243. The EEG data considered in the article is available through PhysioNet, which was supported by NIH grant R01GM104987. The authors thank the Referees, Associate Editor and Editor for providing insightful comments that greatly improved the article.

References

- Adak S (1998), "Time-dependent spectral analysis of non-stationary time series," *Journal of the American Statistical Association*, 93, 1488–1501.
- Brillinger DR (2002), *Time Series: Data Analysis and Theory*, Philadelphia: SIAM.
- Bruce SA, Hall MH, Buysse DJ, and Krafty RT (2017), "Conditional adaptive Bayesian spectral analysis of nonstationary biomedical time series," *Biometrics*, In Press.
- Cadonna A, Kottas A, and Prado R (2017), "Bayesian mixture modeling for spectral density estimation," *Statistics & Probability Letters*, 125, 189195.
- Carter CK and Kohn R (1996), "Semiparametric Bayesian inference for time series with mixed spectra," *Journal of the Royal Statistical Society, Series B*, 59, 255–268.
- Chipman HA, Kolaczyk ED, and McCulloch RE (1997), "Adaptive Bayesian wavelet shrinkage," *Journal of the American Statistical Association*, 92, 1413–1421.
- Choudhuri N, Ghosal S, and Roy A (2004), "Bayesian estimation of the spectral density of a time series," *Journal of American Statistical Association*, 99, 1050–1059.
- Dahlhaus R (2000), "A likelihood approximation for locally stationary process," *Annals of Statistics*, 28, 1762–1794.
- Davis R, Lee T, and Rodriguez-Yam G (2006), "Structural break estimation for nonstationary time series models," *Journal of American Statistical Association*, 101, 223–239.

- Fiecas M and Ombao H (2016), "Modeling the evolution of dynamic brain processes during an associative learning experiment," *Journal of the American Statistical Association*, 111, 1440–1453.
- Gelman A (2006), "Prior distributions for variance parameters in hierarchical models," *Bayesian Analysis*, 1, 1–19.
- Gelman A, Carlin J, Stern H, Dunson D, Vehtari A, and Rubin D (2013), *Bayesian Data Analysis*, Boca Baton, London, New York: Chapman & Hall, 3rd ed.
- Green P (1995), "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, 4, 711–732.
- Guo W and Dai M (2006), "Multivariate time-dependent spectral analysis using Cholesky decomposition," *Statistica Sinica*, 16, 825–845.
- Iber C, Ancoli-Israel S, Chesson A, and Quan SF (2007), *The AASM Manual for the Scoring of Sleep and Associated Events*, Westchester, I. L.: American Academy of Sleep Medicine.
- Kirch C, Edwards MC, Meier A, and Meyer R (2017), "Beyond Whittle: Nonparametric correction of a parametric likelihood with a focus on Bayesian time series analysis," *arXiv Preprint*, arXiv: 1701.04846.
- Krafty RT and Collinge WO (2013), "Penalized multivariate Whittle likelihood for power spectrum estimation," *Biometrika*, 100, 447–458.
- Krafty RT, Rosen O, Stoffer DS, Buysse DJ, and Hall MH (2017), "Conditional spectral analysis of replicated multiple time series with application to nocturnal physiology," *Journal of the American Statistical Association*, In Press.
- Le Cam L (1953), "On some asymptotic properties of maximum likelihood estimates and related Bayes estimates," *University of California Publications in Statistics*, 103, 367–386.
- Le Cam L and Yang G (1990), *Asymptotics in Statistics: Some Basic Concepts*, New York: Springer.
- L'Heureux M, Lee S, and Lyon B (2013), "Recent multidecadal strengthening of the Walker circulation across the tropical Pacific," *Nature Climate Change*, 3, 571–576.
- Liang F (2009), "Improving SAMC using smoothing methods: theory and applications to Bayesian model selection problems," *Annals of Statistics*, 37, 2626–2654.
- Macaro C and Prado R (2014), "Spectral decompositions of multiple time series: a Bayesian non-parametric approach," *Psychometrika*, 79, 105–129. [PubMed: 24154824]
- Neal R (2011), "MCMC using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo*, eds. Brooks S, Gelman A, Jones G, and Meng X-L, Boca Raton, FL: Chapman & Hall/CRC Press, pp. 1385–1399.
- Nicholls N (2008), "Recent trends in the seasonal and temporal behaviour of the El Nino-Southern Oscillation," *Geophysical Research Letters*, 35, L19703.
- Ombao H and Ho M.-h. R. (2006), "Time-dependent frequency domain principal components analysis of multichannel non-stationary signals," *Computational Statistics & Data Analysis*, 50, 2339–2360.
- Ombao H, Raz J, von Sachs R, and Malow B (2001), "Automatic statistical analysis of bivariate nonstationary time series," *Journal of the American Statistical Association*, 96, 543–560.
- Ombao H, von Sachs R, and Guo W (2005), "SLEX analysis of multivariate nonstationary time series," *Journal of the American Statistical Association*, 100, 519–531.
- Park T, Eckley IA, and Ombao HC (2014), "Estimating time-evolving partial coherence between signals via multivariate locally stationary wavelet processes," *IEEE Transactions on Signal Processing*, 62, 5240–5240.
- Pawitan Y and O'Sullivan F (1994), "Nonparametric spectral density estimation using penalized Whittle likelihood," *Journal of the American Statistical Association*, 89, 600–610.
- Priestley MB (1965), "Evolutionary spectral and non-stationary process," *Journal of the Royal Statistical Society, Series B*, 28, 228–240.
- Priestley MB (1981), *Spectral Analysis and Time Series*, Vol. 1 and 2, London: Academic Press.
- Qin L and Wang Y (2008), "Nonparametric spectral analysis with applications to seizure characterization using EEG time series," *Annals of Applied Statistics*, 2, 1432–1451.
- Rosen O and Stoffer D (2007), "Automatic estimation of multivariate spectra via smoothing splines," *Biometrika*, 94, 335–345.

- Rosen O, Wood S, and Stoffer D (2012), “AdaptSPEC: Adaptive spectral estimation for nonstationary time series,” *Journal of the American Statistical Association*, 107, 1575–1589.
- Sanderson J, Fryzlewicz P, and Jones MW (2010), “Estimating linear dependence between nonstationary time series using the locally stationary wavelet model,” *Biometrika*, 97, 435–446.
- Schwarz K and Krivobokova T (2016), “A unified framework for spline estimators,” *Biometrika*, 103, 121–131.
- Shumway R and Stoffer D (2011), *Time Series Analysis and Its Applications*, Springer: New York.
- Solow A (2006), “An ENSO shift revisited,” *Geophysical Research Letters*, 33, L22602.
- Timmermann A, Jin F, and Collins M (2004), “Intensification of the annual cycle in the tropical Pacific due to greenhouse warming,” *Geophysical Research Letters*, 31, L12208.
- Trenberth K and Hoar T (1996), “The 1990–1995 El Niño southern oscillation event: longest on record,” *Geophysical Research Letters*, 23, 57–60.
- Vecchi G and Wittenberg A (2010), “El Niño and our future climate: where do we stand?” *Wiley Interdisciplinary Reviews: Climate Change*, 1, 260–270.
- Whittle P (1953), “Estimation and information in stationary time series,” *Arkiv for Matematik*, 2, 423–434.
- Zhang S (2016), “Adaptive spectral estimation for nonstationary multivariate time series,” *Computational Statistics & Data Analysis*, 103, 330–349.

Figure 1:
Two-channel EEG from a patient during sleep.

Figure 2:
A simulated piecewise stationary trivariate time series.

Figure 3: Time-varying individual spectra and their estimates (top two rows), and time-varying coherences and their estimates (bottom two rows) from the piecewise stationary process in Figure 2.

Figure 4:
A simulated slow-varying bivariate time series.

Figure 5:
Time-varying individual spectra and coherence (top row), and estimates (bottom row) from the slowly-varying process time series in Figure 4.

Figure 6:

Estimate of the time-varying log spectra and coherence of the EEG displayed in Figure 1 (f_{11} and f_{22} indicate C3-A2 and C4-A1 channels, respectively)

Figure 7:
Plots of montly SOI, Nino3.4 and DSLPA from 1951 to 2016.

Figure 8:
Plot of $\hat{p}_r(\delta_{1,2} = t | x)$ for the ENSO analysis.

Figure 9:

First row: the estimated individual time-varying log spectrum of SOI, Nino3.4 and DSLPA time series. Second row: the estimated pairwise coherences (From left to right: between SOI and Niño3.4, SOI and DSLPA, Niño3.4 and DSLPA).

Table 1:

Results of the simulation study of the piecewise stationary process. Based on 250 repetitions, means (standard deviation) of average square error $\times 10^2$ of spectral estimates obtained through the proposed method, piecewise vector autoregressive modeling (Auto-PARM), and smoothing spline ANOVA (SmoothANOVA).

	f_{11}	f_{22}	f_{33}	ρ_{21}^2	ρ_{31}^2	ρ_{32}^2
Proposed	25.1 (28.6)	22.8 (20.5)	13.8 (13.3)	1.4 (0.8)	2.1 (0.9)	2.2 (1.1)
Auto-PARM	37.2 (33.2)	25.4 (16.2)	14.3 (7.7)	6.2 (2.4)	3.8 (1.1)	9.0 (1.3)
SmoothANOVA	39.8 (10.6)	93.4 (16.2)	27.3 (6.4)	9.6 (0.5)	12.8 (0.9)	13.4 (1.2)

Table 2:

Results of the simulation study of the slowly-varying process. Based on 250 repetitions, means (standard deviation) of average square error $\times 10^2$ of spectral estimates obtained through the proposed method, piecewise vector autoregressive modeling (Auto-PARM), and smoothing spline ANOVA (SmoothANOVA).

	f_{11}	f_{22}	ρ_{21}^2
Proposed	49.2 (19.3)	51.4 (26.3)	1.1 (0.2)
Auto-PARM	90.6 (26.6)	85.5 (34.5)	7.0 (3.2)
SmoothANOVA	51.8 (20.3)	61.2 (22.5)	1.6 (1.3)

Table 3:

Results of the simulation study of process (4) based on 20 repetitions: means (standard deviations) of posterior probabilities of numbers of segments m , of average square error (ASE) across all spectral components and of computational time per iteration for the proposed method and the SSAMC based method.

m	The proposed	SSAMC (1 sample)	SSAMC (5 Samples)	SSAMC (20 samples)
1	0 (0)	0 (0)	0 (0)	0 (0)
2	0 (0)	0 (0)	0 (0)	0 (0)
3	0 (0)	0.001 (0.003)	0 (0)	0 (0)
4	0.794 (0.325)	0.251 (0.273)	0.306 (0.316)	0.385 (0.307)
5	0.197 (0.334)	0.348 (0.310)	0.342 (0.275)	0.375 (0.248)
6	0.009 (0.017)	0.275 (0.322)	0.288 (0.263)	0.230 (0.290)
7	0(0)	0.125 (0.172)	0.064 (0.108)	0.010 (0.05)
ASE	1.665 (0.206)	3.387 (0.461)	2.934 (0.371)	2.734 (0.329)
Run Time (s)	1.489 (0.217)	5.119 (0.975)	22.337 (7.141)	62.459 (22.066)

Table 4:

Posterior probability of the number of segments of the ENSO series.

Number of segments (m)	1	2	3	4
Posterior probability	0.3286	0.6521	0.0096	0.0097

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript