

Local Spectral Analysis via a Bayesian Mixture of Smoothing Splines

By Ori Rosen

Department of Mathematical Sciences, University of Texas at El Paso, El Paso, Texas 79968,
U.S.A.

David S. Stoffer

Department of Statistics, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, U.S.A.

and

Sally Wood

Australian Graduate School of Management, University of New South Wales, Sydney, New South
Wales, 2052, Australia

In many practical problems, time series are realizations of nonstationary random processes. These processes can often be modeled as processes with slowly changing dynamics or as piecewise stationary processes. In these cases, various approaches to estimating the time-varying spectral density have been proposed. Our approach in this paper is to estimate the log of the Dahlhaus local spectrum using a Bayesian mixture of splines. The basic idea of our approach is to first partition the data into small sections. We then assume that the log spectral density of the evolutionary process in any given partition is a mixture of individual log spectra. We use a mixture of smoothing splines model with time varying mixing weights to estimate the evolutionary log spectrum. The mixture model is fit using Markov chain Monte Carlo techniques that yield estimates of the log spectra of the individual subsections. In addition to an estimate of the local log spectral density, the method yields pointwise credible intervals. We use a reversible jump step to automatically determine the number of different spectral components.

Key Words: Evolutionary spectra; Locally stationary time series; Mixture of splines; Reversible jump MCMC.

1 Introduction

In many practical problems, the spectral analysis of time series that are realizations of nonstationary random processes is of interest. For example, analyzing the Southern Oscillation Index (SOI) time series is important in monitoring global climate change and predicting rainfall. Priestley (1965) was the first to introduce the concept of evolutionary spectra and a Cramér representation with a time-varying transfer function. This idea was later refined by Dahlhaus (1997) where he established an asymptotic framework for locally stationary processes. The definition of a Dahlhaus locally

stationary process will be given in the next section, but throughout the remainder of this section, we use the term locally stationary in a generic sense.

Several authors have considered the estimation of locally stationary processes under a variety of assumptions. The estimators that were developed by Dahlhaus for his evolutionary spectra are consistent, but the method is not computationally efficient and can be problematic when the time series is long. Chiann and Morettin (1999) proposed a wavelet based version of the estimator proposed by Dahlhaus. Various other approaches have been suggested to overcome the computational difficulty. Ombao et al. (2001) proposed nonparametric estimators based on smooth local exponential functions. Guo et al. (2003) extended the work of Ombao et al. (2001) to allow for simultaneous smoothing in both the time and frequency domains. Our primary goal is to establish a Bayesian approach for estimating the time-varying spectrum of a Dahlhaus locally stationary process in a semi-automatic fashion. We will make our goal more precise in the next section.

Another approach to analyzing locally stationary time series is to consider fitting piecewise autoregressive (AR) models; this approach was suggested by Kitagawa and Akaike (1978). Recently Davis, Lee and Rodriguez-Yam (2006) suggested fitting piecewise AR models using minimum description length and a genetic algorithm for solving the difficult optimization problem. Although Davis et al. (2006) showed that their simulation results for a few locally (and piecewise) stationary AR models perform better than those of Ombao et al. (2001), it is clear that, generally, a parametric technique will outperform a nonparametric technique when the parametric model is correct. If an investigator is certain the dynamics generating the data are locally stationary AR models, then we would suggest using a parametric approach. If, however, an investigator is uncertain about the dynamics of the process, a nonparametric procedure might be preferred. Indeed, some authors, for example Thomson (1990) in analyzing locally stationary geophysical series, warn strongly that AR spectra can be misleading.

One of the difficulties of the non-Bayesian approaches is that they do not jointly model the uncertainty surrounding the point estimates of the spectra and the point estimates of the number and location of the locally stationary segments. Thus, to make inference regarding the individual locally stationary spectra, plug-in point estimates of the number and location of the locally stationary segments are used. Such an approach underestimates the uncertainty surrounding estimates of the locally stationary spectra.

Bayesian approaches to modeling nonstationary time series include Lavielle (1998) and Punsakaya et al. (2002). Punsakaya et al. (2002) proposed a Bayesian method for fitting piecewise linear regression models like AR models. Their method is based on placing prior distributions on the number of changepoints, their locations and the order of the linear regression in each segment. The method is implemented via reversible jump MCMC methods. Unlike the work of Punsakaya et al. (2002), our approach is nonparametric.

There are four main contributions of our paper. First we present a novel method of modeling the time varying log spectral density using a mixture of a finite but unknown number of individual log spectra, where the mixture weights of the individual log spectra vary with time. The basic idea of our approach is to first partition the data into small sections. We then assume that in any given partition, the log spectral density of the process is a mixture of individual log spectra. We introduce nonstationarity into our model by allowing the mixing weights of the individual log spectra to change smoothly across partitions.

The second contribution is that we perform model averaging rather than model selection by allowing the number of underlying individual spectra to be unknown and to vary from $j = 1, \dots, J$. To this end, we construct a Markov chain Monte Carlo (MCMC) scheme using the reversible jump method of Green (1995) to move between models of varying numbers of components. Thus, our resulting estimate of the time varying log spectra is averaged across models of varying numbers of components.

The third contribution of our paper is that we use the output from the MCMC scheme not only to estimate but also to make inference regarding the time varying log spectra, the number of underlying locally stationary spectra and the location and rate at which the time series changes from one locally stationary process to another.

Finally, segmentation methods such as those mentioned in our introduction are entirely local in that estimates are based only on the data in the segment (with perhaps some overlap from contiguous segments). Our method uses all of the data in the estimation and inference procedures. For example, in Section 7.2, we will analyze seismic data. In those examples, the data are amplitude modulated signal wherein the strength of the signal varies over time but the frequency of the signal remains constant. Any existing segmentation method would ignore this fact and simply base its estimates on the data in each segment. Our method will recognize the fact that there are similarities in the signal over the entire sample and will use the information in all of the data to estimate the evolutionary spectrum.

2 Locally Stationary Time Series

The Dahlhaus (1997) approach to modeling evolutionary spectra is based on a time dependent analog to the classical Cramér spectral representation of stationary processes. The following definition is for the case of a zero-mean process.

Definition 2.1 (Dahlhaus) *A sequence of zero-mean stochastic processes, $\{X_{t,N}; t = 1, \dots, N\}$, for $N \geq 1$, is called **locally stationary** with transfer function A^0 if there exists a representation*

$$X_{t,N} = \int_{-1/2}^{1/2} A_{t,N}^0(\nu) \exp(i2\pi\nu t) dZ(\nu), \quad (1)$$

where

(i) $Z(\nu)$ is a zero mean orthogonal increment process on $[-1/2, 1/2]$;

(ii) There exists a positive constant K and a smooth function $A : [0, 1] \times [-1/2, 1/2] \rightarrow \mathbb{C}$ with $A(u, \nu) = \overline{A(u, -\nu)}$ such that for all N ,

$$\sup_{t, \nu} |A_{t,N}^0(\nu) - A(t/N, \nu)| \leq K N^{-1}; \quad (2)$$

(iii) For all ν , $A(u, \nu)$ is continuous in u .

The Dahlhaus evolutionary spectrum can then be defined as follows.

Definition 2.2 (Dahlhaus) *The **evolutionary spectrum** of a Dahlhaus locally stationary process at time $u \in [0, 1]$ and frequency $\nu \in [-1/2, 1/2]$ is given by $f(u, \nu) = |A(u, \nu)|^2$.*

We refer the reader to Dahlhaus (1997) for examples of processes satisfying Definition 2.1. Note that the first argument of $A(u, \nu)$ is rescaled to live on the unit interval and increasing the number of observations, N , means we allow for more data to be observed locally in the sense of infill asymptotics, rather than in the sense of increasing domain asymptotics, i.e., observing more data in the future. Equation (2) deals with the smoothness of A in u such that it is allowed to change only slowly over time. The idea behind this is essentially that for each fixed N , we implicitly assume some local interval of stationarity about each time point and a smooth change from one interval to the next.

Because a Dahlhaus locally stationary process changes smoothly in time, u , we may approximate the evolutionary spectrum by a piecewise stationary process as follows. For a fixed N , consider a partition of the unit interval into S segments of equal length chosen in such a way that $n = N/S$ is small relative to N . We denote the corresponding intervals by $I_s = (\frac{s-1}{S}, \frac{s}{S}]$, for $s = 1, \dots, S$. Next, we define the piecewise stationary process

$$\tilde{X}_{t,N} = \sum_{s=1}^S X_{s,t'} \delta(t/N, I_s), \quad (3)$$

where $t' = 1, \dots, n$ [i.e., $t' = (t-1) \pmod{n} + 1$], and where $\delta(t/N, I_s) = 1$ if $t/N \in I_s$ and $\delta(t/N, I_s) = 0$ if $t/N \notin I_s$. The stationary process $X_{s,t'}$ is defined to have a spectral density $\tilde{f}_s(\nu)$, which is given by averaging the evolutionary spectrum in that interval; i.e.,

$$\tilde{f}_s(\nu) = \frac{1}{S} \int_{u \in I_s} f(u, \nu) du \quad (4)$$

for $s = 1, \dots, S$.

In the asymptotics, we may let $S \rightarrow \infty$, but $S/N \rightarrow 0$ as $N \rightarrow \infty$. Hence, because of the smoothness of $A(u, \nu)$ in u , the local spectral density of $\tilde{X}_{t,N}$, namely $\tilde{f}_s(\nu)$ for $t/N \in I_s$, and $s = 1, \dots, S$, well approximates the evolutionary spectrum, $f(u, \nu)$, of $X_{t,N}$, provided N and S are sufficiently large. In addition, Ombao et al. (2001, Theorem 1) argue that under mild conditions, locally stationary processes can be well approximated by piecewise stationary processes in the sense that the average mean square error between $X_{t,N}$ and $\tilde{X}_{t,N}$ is $O(n^2/N^2)$.

Our general technique relies on the approximation given by (3)–(4). To make our technique easier to understand, however, we first describe the case for a stationary process in the next section. After the stationary case has been discussed, we will then describe our technique for the Dahlhaus locally stationary case in the subsequent sections.

3 Spectral Estimation for Stationary Time Series

As previously indicated, our general approach to the problem of estimating local spectra is best understood by first explaining the technique for estimating the spectral density of a stationary process.

3.1 Model

Suppose that a stationary time series, $\{X_t\}$, has a spectral density given by $f(\nu)$, for $-1/2 < \nu \leq 1/2$. We assume that $f(\nu)$ is bounded and positive. Given a realization, x_1, \dots, x_n , the periodogram of the data at frequency ν (measured in cycles per unit time) is

$$I_n(\nu) = \frac{1}{n} \left| \sum_{t=1}^n x_t \exp(-2\pi i \nu t) \right|^2 .$$

Let $\nu_k = k/n$, for $k = 0, \dots, n-1$, be the Fourier frequencies. Whittle (1957) showed that, under appropriate conditions, for large n , the likelihood of $\mathbf{x} = (x_1, \dots, x_n)$ can be approximated as

$$p(\mathbf{x} \mid f) \propto \prod_{k=0}^{n-1} \exp \left\{ -\frac{1}{2} \left[\log f(\nu_k) + I_n(\nu_k) / f(\nu_k) \right] \right\} . \quad (5)$$

Let $y_n(\nu_k)$ be the log of the periodogram evaluated at the Fourier frequencies, then the representation in (5) suggests the log-linear model

$$y_n(\nu_k) = \log f(\nu_k) + \epsilon_k , \quad (6)$$

where $y_n(\nu_k) = \log I_n(\nu_k)$ for $k = 0, \dots, [n/2]$, where $[n/2]$ is the largest integer less than or equal $n/2$, the ϵ_k 's are independent, $\epsilon_k \sim \log(\chi_2^2/2)$ for $k = 1, \dots, [n/2] - 1$, and $\epsilon_k \sim \log(\chi_1^2)$ for

$k = 0, [n/2]$. Note that in (5), there are only $[n/2] + 1$ distinct observations since the spectral density and the periodogram are both even functions of ν . For ease of notation, in what follows, we assume that n is even.

It is seen from (6) that the log spectral density can be estimated nonparametrically with the log periodogram as the dependent variable. Wahba (1980) used a frequentist approach for estimating the log spectral density via cubic smoothing splines. Carter and Kohn (1997) used a Bayesian approach to modeling the log spectral density as cubic smoothing splines as well by expressing equation (6) in a state-space form. They approximated the error distribution in (6) by a mixture of five normal distributions and introduced latent distribution component indicators to facilitate the estimation.

3.2 Priors

Let $g(\nu_k) = \log f(\nu_k)$. To place a prior on $g(\nu_k)$ we follow Wahba (1990, p. 16) and express $g(\nu_k)$ as the sum of its linear and nonlinear components, so that

$$g(\nu_k) = \alpha_0 + \alpha_1 \nu_k + h(\nu_k) ,$$

where $h(\nu_k)$ is the nonlinear component. We place a smoothing spline prior on $h(\nu)$, which means that $h(\nu)$ is a zero mean Gaussian process with variance-covariance matrix $\tau^2 \Omega$, where the ij -th element of Ω is given by

$$\text{cov}(h(\nu_i), h(\nu_j)) = \tau^2 \omega_{ij} ,$$

where τ^2 is the smoothing parameter and

$$\omega_{ij} = \nu_i^2 (\nu_j - \nu_i/3)/2, \quad \nu_i \leq \nu_j .$$

The parameters α_0 and α_1 are the value of the log of spectral density and its first derivative at $\nu = 0$ respectively. The symmetry and periodicity of the spectral density mean that $(\partial \log f(\nu)/\partial \nu)|_{\nu=0} = 0$. Accordingly α_1 is set to be identically zero and the prior on α_0 is $N(0, c_\alpha)$, for some large c_α . To complete the prior specification on $g(\nu)$, we assume $p(\tau^2) \propto 1/\tau^2$.

To facilitate the computations we follow Wood, Jiang and Tanner (2002) and write $\mathbf{h} = (h(\nu_0), \dots, h(\nu_{n/2}))'$ as a linear combination of basis functions so that $\mathbf{h} = X\boldsymbol{\beta}$, where the columns of the design matrix X are the Demmler-Reinsch basis functions evaluated at the Fourier frequencies and $\boldsymbol{\beta}$ is an $(n/2 + 1) \times 1$ vector of regression coefficients. These basis functions are constructed by decomposing Ω as $\Omega = QDQ'$, where Q is the matrix of eigenvectors of Ω , and D is a diagonal matrix containing the eigenvalues of Ω . Letting $X = QD^{1/2}$ and setting the prior on $\boldsymbol{\beta}$ to be $p(\boldsymbol{\beta}) \sim N(0, \tau^2 I_{(n/2+1)})$, we have $X\boldsymbol{\beta} \sim N(0, \tau^2 \Omega)$ as required. In practice, one can retain only the columns of X corresponding to the largest eigenvalues of Ω resulting in computational saving without affecting the fit. The model in (6) can now be expressed as

$$\mathbf{y} = \alpha_0 \mathbf{1} + X\boldsymbol{\beta} + \boldsymbol{\epsilon} ,$$

where $\mathbf{1}$ is an $(n/2 + 1) \times 1$ vector of ones and $\boldsymbol{\epsilon} = (\epsilon_0, \dots, \epsilon_{n/2})'$ with the ϵ_k 's distributed as in (6). Note that this is the same model as in Carter and Kohn (1997). Both this paper and Carter and Kohn (1997) assume a smoothing spline prior on the log spectral density. We chose a different representation of this prior for ease of computation and to avoid the introduction of the latent component indicators used by Carter and Kohn (1997).

4 Spectral Estimation for Locally Stationary Time Series

In this section, we generalize the spectral analysis of stationary time series described in Section 3 to time series that are Dahlhaus locally stationary as defined in Section 2. Given observations $\{X_{t,N}; t = 1, \dots, N\}$ from a Dahlhaus locally stationary process with evolutionary spectrum $f(u, \nu)$, we consider the approximation process $\{\tilde{X}_{t,N}; t = 1, \dots, N\}$ defined in (3). In particular, given a segmentation S , we form the partitioned data $\{X_{s,t'}; t' = 1, \dots, n; s = 1, \dots, S\}$ as described in (3). Our goal is to use the partitioned data to estimate the local spectra, $\tilde{f}_s(\nu)$, as specified in (4), for $s = 1, \dots, S$ and to use the estimated local spectra as an estimate of the evolutionary spectrum $f(u, \nu)$. The choice of the number of segments, S , will be discussed after we present the model, but the basic idea is that the finer the partition of the unit time interval, the better the estimate of the evolutionary spectrum.

One obvious but primitive approach to this analysis is to estimate the local spectra, $\tilde{f}_s(\nu)$, individually in each segment $s = 1, \dots, S$ using the techniques presented in Section 3. Although this approach may be viable if, in fact, the data are Dahlhaus locally stationary and changing slowly over time, we would like to have an approach that is more robust to departures from this assumption. For example, if the process of interest is piecewise stationary, its time-varying spectra is not smooth and may make jumps at certain time points. If we segment such a process in such a way that a particular partition contains a jump, assuming the data are stationary in that segment will lead to a bad estimate. Another problem with the primitive method, as with all of the segmentation methods previously mentioned, is that it does not use the information from all of the data (or even data in contiguous segments), but rather only from the data in that segment. To overcome these drawbacks, our strategy is to model evolutionary spectra as a mixture of an unknown but finite number of spectra. We make these ideas precise in the following subsection.

4.1 Model

Given a segmentation S , our goal is the estimation of $g_s(\nu) = \log \tilde{f}_s(\nu)$. To this end, we first calculate the periodogram corresponding to each segment. Let $\mathbf{y}_s = (y_{s,0}, \dots, y_{s,n/2})'$ be the log-periodogram for segment s , $s = 1, \dots, S$, evaluated at the Fourier frequencies. We model these

observations as

$$y_{s,k} = g_s(\nu_k) + \epsilon_k$$

where $g_s(\nu)$ is the log of the spectral density for segment s , for $s = 1, \dots, S$, and the ϵ_k 's for $k = 0, \dots, [n/2]$ are distributed as in (6). We model $g_s(\nu)$ as a mixture of an unknown but finite number of spectra so that

$$g_s(\nu) = \sum_{j=1}^J g_{js}(\nu) \Pr(j),$$

where J is the maximum number of components, $\Pr(j)$ is the prior probability that the mixture contains j components, and $g_{js}(\nu)$ is the log of the spectral density of a mixture of j components in segment s . For a given number of mixture components j and segment s we model $g_{js}(\nu)$ as

$$g_{js}(\nu) = \sum_{r=1}^j \pi_{rjs} \log f_{rj}(\nu),$$

where $f_{rj}(\nu)$ is the spectral density of the r^{th} component and π_{rjs} is the unknown weight assigned to the r^{th} component in segment s , with $\sum_{r=1}^j \pi_{rjs} = 1$. Note that the spectral density $f_{rj}(\nu)$ is common to all segments. The value of π_{rjs} represents the probability that in a mixture of j components, the data in segment s have spectral density $f_{rj}(\nu)$. A key point to note is that these probabilities are parameterized to depend upon the segment s and are modeled using a multinomial logistic regression to be specified in Section 4.5. This means that, although the component spectra are common to all segments, a time varying estimate of the spectral density is obtained by allowing the weights of the common spectra to change across segments.

4.2 Segmentation

It is important to note that the choice of the number of segments, S , is not crucial to our estimation process subject to certain constraints. In theory there are potentially as many segments as there are data points. However, practically, we need a minimum number of observations in each segment to estimate the spectral density and for the Whittle approximation to the likelihood to hold. In simulations, we found that using a minimum of 64 observations in each segment gave reasonable results, i.e., the estimates are qualitatively close to the actual values, if the true local spectra have well separated peaks. If an investigator is interested in a finer grid of frequencies, however, then he/she may choose a larger (but relatively small) number of observations in a segment.

Many nonparametric local techniques, including the adaptive techniques that rely on orthogonal libraries and use entropy-based basis algorithms such as the Best Basis Algorithm of Coifman and Wickerhauser (1992), use an arbitrary maximum level of segmentation of the unit interval to initialize the local analysis. In addition, most of these techniques use dyadic segmentation for ease; by dyadic segmentation we mean intervals of $[0, 1]$ of the form $(k2^{-j}, (k+1)2^{-j})$, $k = 0, \dots, 2^j - 1$,




Figure 1: True (solid lines) and estimates (dashed lines) of the log spectral density. The left panel shows $\log(f_{12})$ which is the log of the spectral density of the time series $x_t = 0.9x_{t-1} + \epsilon_t$. The right panel shows $\log(f_{22})$, the log of the spectral density for $x_t = -0.9x_{t-1} + \epsilon_t$.

for levels $j = 1, 2, \dots, J$, where J is the maximum level. Although our method also requires picking a maximal segmentation through the choice of S , it is not crucial that the segmentation be dyadic.

Finally, but perhaps most importantly, is that the parameters of the mixing function π_{rjs} in our model are of more importance than the number of segments because these parameters control the location and rate at which the time series moves from one stationary process to another. The prior on these parameters is discussed in Section 4.5, and the sampling scheme used to explore the posterior distributions of these parameters is described in Section 5.

4.3 An Example

To illustrate our proposed method, consider a time series of length 1024 generated from the following piecewise stationary model

$$x_t = \begin{cases} 0.9x_{t-1} + \epsilon_t & \text{if } 1 \leq t \leq 450 \\ -0.9x_{t-1} + \epsilon_t & \text{if } 451 \leq t \leq 1024, \end{cases}$$

where $\epsilon_t \sim N(0, 1)$. For illustrative purposes, suppose we know that there are two components (i.e., $\Pr(j = 2) = 1$). We first divide the time series into non-overlapping segments each containing $n = 64$ data points. This gives a total of $S = 16$ segments. Our estimate of the log spectra in segment s for $s = 1, \dots, 16$ is

$$g_{2s}(\nu) = \pi_{12s} \log f_{12}(\nu) + (1 - \pi_{12s}) \log f_{22}(\nu). \tag{7}$$




Figure 2: Estimate of the mixing function π_{12s} in (7) based on 16 segments.

Figure 1 shows the true (solid line) and estimated (dashed line) log spectral density for $x_t = 0.9x_{t-1} + \epsilon_t$ (left panel) and for $x_t = -0.9x_{t-1} + \epsilon_t$ (right panel), $\epsilon_t \sim N(0, 1)$. Figure 2 plots the estimated mixing function π_{12s} as a function of the segment. This figure shows that the probability that the data have spectral density $f_{12}(\nu)$ is close to 0.91 at the beginning of the time series. This probability decreases to 0.5 by the 7th segment and is approximately 0.03 by the end of the time series.

4.4 Prior on the Component Spectra

As in Section 3.2, we express $\log f_{rj}(\nu)$ as

$$\log f_{rj}(\nu_k) = \alpha_{0rj} + h_{rj}(\nu_k)$$

and write $\mathbf{h}_{rj} = X\boldsymbol{\beta}_{rj}$. The priors on $\boldsymbol{\beta}_{rj}$, and on α_{0rj} for $r = 1, \dots, j$ and $j = 1, \dots, J$ are as given in Section 3.2, and are assumed to be independent across all r and j . The prior on τ_{rj}^2 is as in Section 3.2, and in addition we impose an ordering on τ_{rj}^2 for $r = 1, \dots, j$ and $j = 1, \dots, J$, so that for a given j , $\tau_{1j}^2 > \dots > \tau_{jj}^2$. This ensures that the likelihood is identified.

4.5 Prior on the Mixing Probabilities

The mixing probabilities are expressed using the multinomial linear logit model so that

$$\pi_{rjs} = \frac{\exp(\boldsymbol{\delta}'_{rj} \mathbf{u}_s)}{\sum_{h=1}^j \exp(\boldsymbol{\delta}'_{hj} \mathbf{u}_s)} \quad (8)$$

with parameters $\boldsymbol{\delta}_{rj}$, $r = 1, \dots, j$ and $j = 1, \dots, J$. In (8), $\mathbf{u}_s = (1, u_s)'$, where the covariate u_s is taken as $u_s = s/S$, and $\boldsymbol{\delta}_{rj} = (\delta_{0rj}, \delta_{1rj})'$. For identifiability, $\boldsymbol{\delta}_{1j}$ is set to zero. Such logistic weights are also used in the mixtures-of-experts model (Jacobs et al., 1991). The priors on $\boldsymbol{\delta}_{rj}$ for $r = 1, \dots, j$ and $j = 1, \dots, J$ are bivariate normal with zero mean and variance $\sigma_\delta^2 I_2$ and are assumed independent across all r and j . In all of our analyses, σ_δ^2 was equal to 4.

4.6 Prior on the number of components, j

Typically, we assume *a priori* that the maximum number of components is J and that $\Pr(j = k) = 1/J$ for $k = 1, \dots, J$. In the simulated examples of Section 6, J was set to 10, whereas in the real examples of Section 7, J was set to 20. We do, however, consider the use of Poisson priors in the example of Section 7.1.

5 Bayesian Inference

We estimate the log of the spectral density in segment s , for $s = 1, \dots, S$ by its posterior mean $E(g_s|\mathbf{y})$, with all unknown parameters integrated out and we use MCMC to perform the required multidimensional integration.

The expectation $E\{g_s(\nu)|\mathbf{y}\}$ is defined as

$$E\{g_s(\nu)|\mathbf{y}\} = \sum_{j=1}^J \int E\{g_s(\nu)|\mathbf{y}, \boldsymbol{\theta}_j, j\} p(\boldsymbol{\theta}_j|\mathbf{y}, j) d\boldsymbol{\theta}_j \Pr(j|\mathbf{y})$$

where $\boldsymbol{\theta}_j = (\boldsymbol{\alpha}'_j, \boldsymbol{\beta}'_j, \boldsymbol{\tau}'_j, \boldsymbol{\delta}'_j)'$, $\boldsymbol{\alpha}_j = (\alpha_{01}, \dots, \alpha_{0j})'$, $\boldsymbol{\beta}_j = (\boldsymbol{\beta}'_{1j}, \dots, \boldsymbol{\beta}'_{jj})'$, $\boldsymbol{\tau}_j = (\tau_{1j}^2, \dots, \tau_{jj}^2)'$ and $\boldsymbol{\delta}_j = (\boldsymbol{\delta}'_{2j}, \dots, \boldsymbol{\delta}'_{jj})'$. This integral cannot be evaluated explicitly and we use MCMC simulation to estimate it. In addition to the point estimates, we construct $(1 - \alpha)$ -level pointwise credible intervals for the log spectra by obtaining the $\alpha/2$ and $1 - \alpha/2$ percentiles of the MCMC fitted log spectra based on all the iterates after the burn-in period. Note that these credible intervals reflect the uncertainty surrounding not only our estimate of $\log f_{rj}$ but also our uncertainty surrounding the number of components J and the mixing probabilities π_{rsj} .

To simplify the simulation from the posterior distribution $p(\boldsymbol{\theta}_j, j|\mathbf{y})$, we introduce latent variables that are generated during the simulation. The first of these is the number of components j that are active at any point in the simulation. Then, given j , define the vector of indicator variables γ_{srj} for $s = 1, \dots, S$ and $r = 1, \dots, j$, where $\gamma_{srj} = 1$ if \mathbf{y}_s originated from the r th component, and $\gamma_{srj} = 0$, otherwise.

Sampling Scheme

We outline here the MCMC scheme for our model. More details are given in the Appendix. The

sampling scheme consists of two parts; a between-model move followed by a within-model move. The number of components j is first initialized, then conditional on this value, the other model parameters α_j , β_j , τ_j and δ_j are initialized.

1. Between Model Move

A new value of j is proposed, and conditional on this value, parameter values for α_j , β_j , τ_j and δ_j are proposed. These proposed values are then accepted or rejected using a Metropolis-Hastings step.

2. Within Model Move

Given the value of j , the parameters specific to a model of j components are then updated as follows.

- (a) Let $\beta_{rj}^* = (\alpha_{rj}, \beta'_{rj})'$, $r = 1, \dots, j$, $\beta_j^* = (\beta_{1j}^*, \dots, \beta_{jj}^*)'$ and $X^* = (\mathbf{1}, X)$. Generate β_j^* from $p(\beta_j^* \mid \tau_j, \gamma_j, X^*, \tilde{\mathbf{y}})$ via Metropolis-Hastings steps, where $\gamma_j = \{\gamma_{srj}\}$, for $r = 1, \dots, j$ and $s = 1, \dots, S$, are the component indicators, and $\tilde{\mathbf{y}} = (\mathbf{y}'_1, \dots, \mathbf{y}'_S)'$.
- (b) Generate τ_j from $p(\tau_j \mid \beta_j)$.
- (c) Generate δ_j from $p(\delta_j \mid \gamma_j, U)$ via a Metropolis-Hastings step, where U is the matrix whose s th row is \mathbf{u}'_s , $s = 1, \dots, S$.
- (d) Let $\gamma_{sj} = r$ if $\gamma_{srj} = 1$. Generate the component indicators from $p(\gamma_{sj} = r \mid \beta_j^*, \delta_j, X^*, \mathbf{y}_s)$.

The proposal densities for generating β_j^* and δ_j are multivariate normal. The resulting acceptance rates are around 30% and 80%, respectively.

6 Simulations

We conducted a simulation study using three settings, the first two of which are taken from Ombao et al. (2001). The first is a slowly-varying AR(2) process, the second is a piecewise stationary model and the third is a time varying AR(6) process. We generated 200 time series from the slowly-varying AR(2) process and 50 time series from each of the other two processes. For all three settings, each time series is of length $N = 1024$, divided into $S = 16$ segments of size $n = 64$. The Gibbs sampler for our model was then run for each time series with 100,000 iterations, including a burn-in period of 20,000 iterations. A run of 100,000 iterations takes about an hour on a 3.4 GHz Linux PC, using Fortran 90 code.

6.1 Slowly Varying AR(2) Process

In this simulation setting, 200 realizations are generated from the model

$$x_t = a_t x_{t-1} - 0.81 x_{t-2} + \epsilon_t, \quad t = 1, \dots, 1024, \quad (9)$$




Figure 3: Top: A realization of the slowly varying AR(2) process given in (9). Bottom: The coefficient a_t as a function of t for process (9). In both figures time is rescaled to the unit interval, $n = 1024$.




Figure 4: Left: True time-varying log-spectrum of process (9). Right: Estimated log-spectrum based on averaging the 200 fitted log-spectra

where $a_t = 0.8(1 - 0.5 \cos(\pi t/1024))$ and $\epsilon_t \stackrel{iid}{\sim} N(0, 1)$. In this case, the process is neither piecewise stationary, nor does it have a dyadic structure. Figure 3 shows a realization of the process (9) and a plot of a_t . Note that a_t changes slowly as a function of t , which in turn results in a slowly-varying spectrum as a function of t .

Figure 4 presents the true time-varying log spectrum of process (9) on the left-hand side, as well as the average of the 200 estimated log spectra, on the right. In this figure, darker shades correspond to higher power, and time has been rescaled to the unit interval. Each of these estimated

log spectra is obtained as described in Section 5. We see that our method does a good job, on average, in estimating the true time-varying spectrum of the process. Following Ombao et al. (2001) we use the averaged squared error (ASE) to assess the distance between the true log spectrum and its estimate. In our notation, ASE is given by

$$\text{ASE} = \{N(n/2 + 1)\}^{-1} \sum_{t=1}^N \sum_{k=0}^{n/2} \left\{ \log \widehat{f}(t/N, \nu_k) - \log f(t/N, \nu_k) \right\}^2, \quad (10)$$

where $f(\cdot, \cdot)$ is the true time-varying spectrum, and

$$\log \widehat{f}(t/N, \nu_k) = \sum_{s=1}^S \widehat{g}_s(\nu_k) \delta(t/N, I_s)$$

where $\widehat{g}_s(\nu_k)$ the estimate of the log of time-varying spectrum $g_s(\nu_k)$ as discussed in Section 5 and $\delta(t/N, I_s)$ is defined below (3). The average of the 200 ASE values corresponding to the 200 realizations was 0.261 with a standard deviation of 0.093. The 10th, 25th, 50th, 75th and 90th percentiles of the 200 ASE values were 0.17, 0.20, 0.23, 0.31 and 0.39, respectively.

To compare our method with the AutoSLEX method presented in Ombao et al. (2001), we repeated the study using AutoSLEX. One of the inputs to the AutoSLEX method is the maximum level of the dyadic segmentation, which determines the maximum number of segments. In this simulation, we set the maximum level to 4, which corresponds to segments of size $n = 64$. The average of the 200 ASE values in this case was 0.759 with a standard deviation of 0.070. The 10th, 25th, 50th, 75th and 90th percentiles of the 200 ASE values were 0.68, 0.72, 0.75, 0.80 and 0.84, respectively. (We note that these results are in conflict with those reported in Ombao et al., 2001, however, our results are based on the AutoSLEX Matlab programs obtained from one of the authors. The calculations of ASE for our method and for AutoSLEX were performed in an identical manner.) It is evident that our method shows a substantial improvement over AutoSLEX.

6.2 Piecewise Stationary Process

Although we assume local stationarity, our technique will also work in the case that the process is piecewise stationary. The following is an example of the viability of our method in this case. In this simulation setting, the time series are generated from the following model

$$x_t = \begin{cases} 0.9x_{t-1} + \epsilon_t & \text{if } 1 \leq t \leq 512 \\ 1.69x_{t-1} - 0.81x_{t-2} + \epsilon_t & \text{if } 513 \leq t \leq 768 \\ 1.32x_{t-1} - 0.81x_{t-2} + \epsilon_t & \text{if } 769 \leq t \leq 1024, \end{cases} \quad (11)$$

where $\epsilon_t \stackrel{iid}{\sim} N(0, 1)$. In model (11), the lengths of the stationary time series are powers of 2. Figure 6 shows the true time-varying log spectrum of process (11) on the left-hand side, as well as the average of the 50 estimated log spectra, on the right. It is clear from Figure 6 that, on average, our technique provides a good estimate of the true piecewise spectra in this example.




Figure 5: From left to right, individual log-spectrum estimates corresponding to the 10th, 50th and 90th percentiles of the 200 ASE values for process (9)




Figure 6: Left: True time-varying log-spectrum of process (11). Right: Estimated log-spectrum based on averaging the 50 fitted log-spectra

6.3 Time Varying AR(6) Process

In this example, we generate 50 realizations, each of length $N = 1024$, from a time varying autoregressive process of lag 6. In particular, this process can be expressed as $\phi_t(B)x_t = \epsilon_t$, $\epsilon_t \stackrel{iid}{\sim} N(0, 1)$, where $B^m x_t = x_{t-m}$, and

$$\phi_t(B) = 1 - \phi_{t1}B - \dots - \phi_{t6}B^6 . \quad (12)$$




Figure 7: Left: True time-varying log-spectrum of process based on (12). Middle: Estimated log-spectrum based on averaging the 50 fitted log-spectra. Right: Estimated log-spectrum based on a single realization

Assume the characteristic polynomial in (12) has three pairs of conjugate complex roots $(a_{tj}^{-1}, a_{tj}^{*-1})$, $j = 1, 2, 3$, where the superscript $*$ denotes the complex conjugate, i.e.,

$$\phi_t(B) = (1 - a_{t1}B)(1 - a_{t1}^*B)(1 - a_{t2}B)(1 - a_{t2}^*B)(1 - a_{t3}B)(1 - a_{t3}^*B).$$

To make these roots vary with time, let $a_{tj}^{-1} = A_j \exp(2\pi i \theta_{tj})$, $j = 1, 2, 3$, where $i = \sqrt{-1}$, and let the θ_{tj} s satisfy

$$\begin{aligned} \theta_{t1} &= .05 + (.1/(N-1))t \\ \theta_{t2} &= .25 \\ \theta_{t3} &= .45 - (.1/(N-1))t, \end{aligned}$$

for $t = 1, \dots, N$. The values of A_1 , A_2 and A_3 are 1.1, 1.12 and 1.1, respectively. The corresponding coefficients, ϕ_{tj} , $t = 1, \dots, N$, $j = 1, \dots, 6$, are obtained as the negative coefficients of the polynomial in (12).

We fitted our method to each of the 50 realizations. Figure 7 presents the true time varying log-spectrum on the left-hand side, the average of the 50 estimated log-spectra, in the middle, and an estimated log-spectrum based on a single realization. From these plots, it is evident that our method provides good estimates of the true time varying spectrum.

7 Applications

In this section, we give two examples to illustrate our methodology. The total number of iterations is 300,000 in the first example and 400,000 in the second.

7.1 Southern Oscillation Index (SOI)

In recent years there has been much research and debate on climate change. One area of research is the El Niño/Southern Oscillation (ENSO) phenomenon. ENSO is an irregular low frequency oscillation between a warm El Niño state and a cold La Niña state. The Southern Oscillation Index (SOI) is an indicator of the ENSO phenomenon and is calculated to be the standardized anomaly of the mean sea-level pressure difference between Tahiti and Darwin. The strong El Niños of 1982/83 together with the more frequent occurrences of El Niño in recent decades have raised the question of whether human-induced global warming has changed the structure of the ENSO time series. (Timmermann et al., 1999). Trenberth and Hoar (1996) used quarterly standardized seasonal Sea Level Pressure Anomalies at Darwin (DSLPA) as an indicator of ENSO and assumed an ARMA(3,1) model as the data-generating process. They tested explicitly for a change in the time series from 1981 onwards and found that given the values of the index from 1882-1981, the likelihood of observing the 1990-1995 ENSO is about 1:3000 years. The SOI is considered to be a better indicator of ENSO than the DSLPA (Chen, 1982) and therefore we use this index as our data.

The data, shown in Figure 8, are monthly values of the SOI from January 1876 to April 2008 and are available at <http://www.bom.gov.au/climate/current/soihtm1.shtml>. We divided the data




Figure 8: Monthly values of the Southern Oscillation Index (SOI) from January 1876 to April 2008

into 24 segments, each containing 64 observations (leaving out the first 52 observations), and fitted our model. We used a uniform prior for the number of components in the mixture, with an upper limit of 20 components, i.e., $Pr(j = k) = 0.05$ for $k = 1, \dots, 20$. Using this prior, our method estimates $Pr(j = 1|\mathbf{y}) = 0.75$, that is, the posterior probability of the time series being stationary is 0.75. To study the effect of the prior on the posterior distribution of the number of components,

we redid the analysis using a Poisson distribution, with mean λ , i.e., $Pr(j = k) = \frac{\lambda^k e^{-\lambda}}{k!}$ for three different values of λ . The results of this analysis appear in Table 1. This table shows that a single component has the highest posterior probability for all priors on the number of components. Even when $\lambda = 5$ and hence $Pr(j = 1) = 0.03$, the posterior probability is $Pr(j = 1|\mathbf{y}) = 0.66$. This finding is in contrast to Trenberth and Hoar (1996) who found that the time series is nonstationary. There are a number of explanations for this disparity.

Number of Components	Prior			
	Uniform	Poisson		
	$P(j = k) = 0.05$	$\lambda = 1$	$\lambda = 2$	$\lambda = 5$
1	0.755	0.94	0.75	0.66
2	0.241	0.06	0.24	0.33
3	0.004	0.00	0.01	0.01
4	0.000	0.00	0.00	0.00

Table 1: Posterior probabilities of the number of components in the model fitted to the SOI data as a function of the prior distribution.

The first is that we use the SOI rather than the DSLPA as an indicator of ENSO. To test if the cause of the difference in findings was due to using different indicators of ENSO, we repeated the analysis of Trenberth and Hoar using the SOI data. As in Trenberth and Hoar (1996), we grouped the monthly data for the period 1876 to 1979 into seasons and fitted a parametric model for the time series using AIC as the basis for model selection. Similar to Trenberth and Hoar, an ARMA(3,1) of the form

$$y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \phi_3 y_{t-3} + e_t + \theta_1 e_{t-1}, \quad e_t \sim N(0, \sigma^2) \quad (13)$$

was the model selected, with parameter estimates $\hat{\phi}_1 = 1.303$, $\hat{\phi}_2 = -0.2707$, $\hat{\phi}_3 = -0.1876$, $\theta_1 = -0.7524$ and $\hat{\sigma}^2 = 49$. In (13), y_t is the average of the SOI for the three months in season t . This model was then used to generate 1,000,000 years of data. Trenberth and Hoar focused on the period December 1989–May 1995, where 22 successive positive anomalies with a median of 0.94 mb were recorded, and asked the question “If the ARMA(3,1) model is the true model, how likely is it that we observe 22 successive negative DSLPA, and how likely is it that we observe a median value of 0.94 mb for any 22 season period?” The answers to these questions were about 1 in 8000 years and 1 in 3000 years, respectively. From this they concluded that the time series structure of the DSLPA had changed. Repeating the same analysis for the SOI data, we obtained the distribution of run lengths of either sign, see Figure 9. Our results are of a similar order of magnitude to those computed by Trenberth and Hoar, but comparison is difficult because the maximum run length for the SOI data was only 11, not 22. However, except for one slightly positive SOI (0.20 for December, January and February in 92/93), the maximum run length would have been 17. The frequencies of observing run lengths of 11, 17, 22 or greater appear in Table 2. This table shows that the estimated frequencies of events are very similar for both the SOI and the DSLPA, which is not surprising given

Run Length	Data	
	SOI	DSLPA
> 11	1 in 47 years	1 in 55 years
> 17	1 in 800 years	1 in 910 years
> 22	1 in 9200 years	1 in 8600 years

Table 2: Frequency of occurrence of length of a run in SOI.




Figure 9: Distribution of the number of runs for simulated values of the SOI assuming an ARMA(3,1) process is the correct model.

that the ARMA(3,1) estimated for the SOI is very similar to the ARMA(3,1) process estimated for the DSLPA (Trenberth and Hoar, 1996). However, different conclusions would be reached because the behavior of the SOI for the period in question, i.e., December 1989–May 1995, is different from the behavior of the DSLPA in the same period. A similar picture emerges for the test of the 22 season median value. The median value of the SOI from December 1989–May 1995 was -7.5 . Using the data-generating model given in (13), we computed that a 22 season median of -7.5 could be expected to happen about once in 130 years, a very different value from once in 3000 years.

There are a number of other differences in our analysis of the ENSO effect. Trenberth and Hoar (1996) assumed an ARMA(3,1) model for the data-generating process, whereas the method presented in this paper models changes in the frequency domain to avoid assuming a particular parametric model.

Another explanation is that Trenberth and Hoar (1996) tested explicitly if there had been a change from 1981 onwards. In our model, explicitly testing for a difference in the SOI from 1981 onwards is equivalent to assuming *a priori* that in a mixture of two components, $\pi_{12s} = 1$ for segments s corresponding to time periods prior to 1981 and $\pi_{12s} = 0$, otherwise. In effect, the uncertainty surrounding the number and location of change points has been ignored by Trenberth




Figure 10: Estimated log spectrum (solid line) and 95% posterior intervals (dashed line) against frequency for monthly values of the Southern Oscillation Index (SOI) from June 1878 to May 2006.

and Hoar (1996). In contrast, our methodology makes no assumptions regarding the location and number of any possible nonstationarity, and our estimate of $Pr(j = 1|\mathbf{y})$ integrates out the uncertainty regarding both the mixing function π and the number of possible change points.

A plot of the estimate of the log power together with 95% posterior intervals appears in Figure 10. Figure 10 shows a peak at about $\nu = 0.02$, which corresponds to a 48 month (4 year) cycle. We note that in our analysis, by modeling the evolution of the log spectra, we are modeling the evolution of the second moment. It may be that the SOI is a nonstationary process but the nonstationarity is apparent only in higher-order moments.

7.2 Earthquakes and Explosions

In this section we illustrate our methodology with seismic traces of an earthquake and a mining explosion taken from Shumway and Stoffer (2006). These data were measured at a recording station in Scandinavia, and are each of length 2048. Figures 11 and 12 show the data, along with their estimated log spectrum. These particular time series both consist of two waves, the compression wave, also known as primary or P wave, which is the start of the series and the shear, or S wave, which arrives at the midpoint of the series. The delay between the arrival of the P wave and the arrival of the S wave is used to get a quick and reasonably accurate estimate of the location of the event. The analysis of such seismic data is one of critical importance for monitoring a comprehensive test-ban treaty. As argued by many authors, e.g., Shumway and Stoffer (2006, Ch 7), distinguishing between the seismic traces of earthquakes and explosions is best accomplished




Figure 11: The earthquake time series (a), and its estimated time-varying log spectrum (b).

in the frequency domain. Hence, two problems are of interest in this analysis. The first one is to identify the arrival of the S wave and the second problem is to estimate the time-varying spectrum of the process, in order to distinguish between the seismic traces of explosions and earthquakes.

To illustrate our method, we fitted the mixture model to both time series after dividing them into 32 segments of 64 observations each and computing the corresponding log periodograms. The Gibbs sampler was run for 400000 iterations with a burn-in period of 200000 iterations. Table 3 presents the posterior probabilities of the number of components for the two time series. The S component for the earthquake shows power at the low frequencies only, and the power remains strong for a long time. In contrast, the explosion shows power at higher frequencies than the earthquake, and the power of the signals (P and S waves) does not last as long as in the case of the earthquake.

No. of Components	Earthquake	Explosion
1	0.000	0.000
2	0.960	0.098
3	0.040	0.700
4	0.000	0.193
> 5	0.000	0.009

Table 3: Posterior probabilities of the number of components




Figure 12: The explosion time series (a), and its estimated time-varying log spectrum (b).

Acknowledgements

The work of Ori Rosen and David Stoffer was supported, in part, by NSF grants DMS-0706723, DMS-0804140 and DMS-0805050.

Appendix: Details of the Sampling Scheme

First initialize the number of components, j . Conditional on this value, initialize the indicators, γ_{srj} , for $s = 1, \dots, S$ and $r = 1, \dots, j$ and $\theta_j = (\beta_j^*, \tau_j', \delta_j')'$, where $\beta_j^* = (\alpha_{01}, \beta'_{1j}, \dots, \alpha_{0j}, \beta'_{jj})'$, $\tau_j = (\tau_{1j}^2, \dots, \tau_{jj}^2)'$, and $\delta_j = (\delta'_{2j}, \dots, \delta'_{jj})'$. Note that $\theta_1 = (\beta_{11}^*, \tau_{11}^2)'$.

1. MOVING BETWEEN MODELS

Let j^c and j^p be the current and proposed number of components in the model, respectively. Let θ^c and θ^p be the corresponding parameter values, i.e. $\theta^c = \theta_{j^c}$ and $\theta^p = \theta_{j^p}$. Let $V^c = (j^c, \theta^c)$ and $V^p = (j^p, \theta^p)$. Using a proposal density $q(V^c, V^p)$, we draw V^p from $P(V^p \mid \tilde{y})$ as follows.

- (a) If $1 < j^c < J$, propose a value for j^p with transition probability $q(j^c \rightarrow j^p = j^c \pm 1) = 0.5$. If $j^c = 1$, then $q(j^c \rightarrow j^p = 2) = 1$, and if $j^c = J$, then $q(j^c \rightarrow j^p = J - 1) = 1$.

(b) If $j^p = j^c + 1$, then $\boldsymbol{\theta}^p = (\boldsymbol{\beta}'_{j^c}, \boldsymbol{\beta}'_{j^p j^p}, \tau'_{j^c}, \tau^2_{j^p j^p}, \boldsymbol{\delta}'_{j^c}, \boldsymbol{\delta}'_{j^p j^p})'$, where $\boldsymbol{\delta}_{j^p j^p}$, $\tau^2_{j^p j^p}$ and $\boldsymbol{\beta}^*_{j^p j^p}$ are generated as follows.

- i. Draw $\boldsymbol{\delta}_{j^p j^p}$ from $N(0, \sigma_\delta^2 I_2)$.
- ii. Draw $\tau^2_{j^p j^p}$ from $U(0, \tau^2_{j^c j^c})$
- iii. Draw $\alpha_{j^p j^p}$ from $N(0, \sigma_\alpha^2)$
- iv. Draw $\boldsymbol{\beta}_{j^p j^p}$ from $N(0, \tau^2_{j^p j^p} \cdot I_k)$, where k is the number of basis functions.
- v. Accept the proposed parameter values with probability

$$\epsilon = \min \left\{ 1, \frac{P(V^p | \tilde{\mathbf{y}})q(V^p, V^c)}{P(V^c | \tilde{\mathbf{y}})q(V^c, V^p)} \right\} \quad (\text{A.1})$$

(c) If $j^p = j^c - 1$, then

- i. $\boldsymbol{\theta}^p = \boldsymbol{\theta}_{j^c-1}$
- ii. Accept the proposed parameter values with probability (A.1).

2. UPDATING WITHIN A MODEL

Given the new values of j and $\boldsymbol{\theta}_j$, the model parameters are updated as follows.

(a) If $j > 1$ then

- i. Let $\gamma_{sj} = r$ if $\gamma_{srj} = 1$. Generate component indicators γ_{sj} , $s = 1, \dots, S$, from

$$p(\gamma_{sj} = r | \boldsymbol{\beta}_j^*, \boldsymbol{\delta}_j, X^*, \mathbf{y}_s) = \pi_{rj_s} p(\mathbf{y}_s | X^*, \boldsymbol{\beta}_{rj}^*) / \sum_{h=1}^j \pi_{hj_s} p(\mathbf{y}_s | X^*, \boldsymbol{\beta}_{hj}^*), \quad (\text{A.2})$$

where

$$\begin{aligned} P(\mathbf{y}_s | X^*, \boldsymbol{\beta}^*) &\propto \exp\left(\frac{y_{s,0} - \mu_0}{2} - \frac{1}{2} \exp(y_{s,0} - \mu_0)\right) \\ &\times \exp\left(\frac{y_{s,n/2} - \mu_{n/2}}{2} - \frac{1}{2} \exp(y_{s,n/2} - \mu_{n/2})\right) \\ &\times \prod_{k=1}^{n/2-1} \exp\left(y_{s,k} - \mu_k - \exp(y_{s,k} - \mu_k)\right) \end{aligned} \quad (\text{A.3})$$

and μ_k is the k th element of $X^* \boldsymbol{\beta}^*$.

- ii. Generate $\boldsymbol{\delta}_j$ from $p(\boldsymbol{\delta}_j | \boldsymbol{\gamma}_j, U)$ via a Metropolis-Hastings step. The conditional posterior distribution of $\boldsymbol{\delta}_j$ is

$$\begin{aligned} p(\boldsymbol{\delta}_j | \boldsymbol{\gamma}_j, U) &\propto p(\boldsymbol{\delta}_j) \prod_{s=1}^S \prod_{r=1}^j \pi_{rj_s}^{\gamma_{srj}} \\ &= p(\boldsymbol{\delta}_j) \prod_{s=1}^S \prod_{r=1}^j \left(\frac{\exp(\boldsymbol{\delta}'_{rj} \mathbf{u}_s)}{\sum_{h=1}^j \exp(\boldsymbol{\delta}'_{hj} \mathbf{u}_s)} \right)^{\gamma_{srj}}. \end{aligned} \quad (\text{A.4})$$

The proposal density is $N(\boldsymbol{\delta}_j^{\max}, \Sigma_\delta)$, where $\boldsymbol{\delta}_j^{\max}$ is the value of $\boldsymbol{\delta}_j$ maximizing $p(\boldsymbol{\delta}_j | \boldsymbol{\gamma}_j, U)$ in (A.4), and $\Sigma_\delta = \left\{ -\frac{\partial^2 \log p(\boldsymbol{\delta}_j | \boldsymbol{\gamma}_j, U)}{\partial \boldsymbol{\delta}_j \partial \boldsymbol{\delta}'_j} \Big|_{\boldsymbol{\delta}_j = \boldsymbol{\delta}_j^{\max}} \right\}^{-1}$.

(b) For any j :

i. Generate β_j^* from its conditional posterior distribution

$$\begin{aligned} p(\beta_j^* \mid \tau_j, \gamma_j, X^*, \tilde{\mathbf{y}}) &= p(\beta_j^*) \prod_{s=1}^S \prod_{r=1}^j p(\mathbf{y}_s \mid X^*, \beta_{rj}^*)^{\gamma_{srj}} \\ &= \prod_{s=1}^S \prod_{r=1}^j p(\beta_{rj}^*) p(\mathbf{y}_s \mid X^*, \beta_{rj}^*)^{\gamma_{srj}} \end{aligned} \quad (\text{A.5})$$

as follows. For $r = 1, \dots, j$:

Generate the vector β_{rj}^* via a Metropolis-Hastings step. The conditional posterior distribution of β_{rj}^* is

$$p(\beta_{rj}^* \mid \tau_{rj}^2, \gamma_j, X^*, \tilde{\mathbf{y}}) = p(\beta_{rj}^*) \prod_{s=1}^S p(\mathbf{y}_s \mid X^*, \beta_{rj}^*)^{\gamma_{srj}}.$$

The proposal density is $N(\beta_{rj}^{*\max}, \Sigma_{\beta_{rj}^*})$, where $\beta_{rj}^{*\max} = \arg \max_{\beta_{rj}^*} p(\beta_{rj}^* \mid \tau_{rj}^2, \gamma_j, X^*, \tilde{\mathbf{y}})$

$$\text{and } \Sigma_{\beta_{rj}^*} = \left\{ - \frac{\partial^2 \log p(\beta_{rj}^* \mid \tau_{rj}^2, \gamma_j, X^*, \tilde{\mathbf{y}})}{\partial \beta_{rj}^* \partial \beta_{rj}^{*\prime}} \Big|_{\beta_{rj}^* = \beta_{rj}^{*\max}} \right\}^{-1}.$$

ii. Generate τ_j from $p(\tau_j \mid \beta_j, \tilde{\mathbf{y}}) = p(\tau_j \mid \beta_j)$, such that $\tau_{1j}^2 \geq \tau_{2j}^2 \geq \dots \geq \tau_{jj}^2$. Note that $p(\tau_{rj}^2 \mid \beta_{rj}) = IG(\frac{k}{2}, \frac{1}{2} \beta_{rj}' \beta_{rj})$, where k is the number of basis functions.

References

- Adak, S. (1998). Time-dependent spectral analysis of nonstationary time series. *J. Am. Statist. Assoc.*, **93**, 1488–1501.
- Carter, C.K. and Kohn, R. (1997). Semiparametric Bayesian inference for time series with mixed spectra. *J.R. Statist. Soc. B*, **59**, 255–268.
- Chen, W.Y. (1982). Assessment of southern oscillation sea-level pressure indices. *Monthly Weather Review*, **110**, 800–807.
- Chiann, C. and Morettin, P.A. (1999). Estimation of time-varying linear systems. *Statist. Inference Stochastic Process.*, **2**, 253–285.
- Coifman, R. and Wickerhauser, M. (1992). Entropy based algorithms for best basis selection. *IEEE Transactions on Information Theory*, **38**, 712–718.
- Dahlhaus, R. (1997). Fitting time series models to nonstationary processes. *Annals of Statistics*, **25**, 1–37.
- Davis, R.A., Lee, T.C.M. and Rodriguez-Yam, G.A. (2006). Structural breaks estimation for nonstationary time series models. *J. Amer. Statist. Assoc.*, **101**, 223–239.
- Green, P.J. (1995). Reversible jump MCMC computation and Bayesian model determination. *Biometrika*, **82**, 711–732.
- Guo, W., Dai, M., Ombao, H., von Sachs, R. (2003). Smoothing spline ANOVA for time-dependent spectral analysis. *J. Amer. Statist. Assoc.*, **98**, 643–652.
- Jacobs, R.A., Jordan, M.I., Nowlan, S.J. and Hinton, G.E. (1991). Adaptive mixtures of local experts. *Neural Comp.*, **3**, 79–87.
- Kitagawa, G. and Akaike, H. (1978). A procedure for the modeling of nonstationary time series. *Ann. Inst. Statist. Math.*, **30**, 351–363.

- Kitagawa, G., and Gersch, W. (1996), Smoothness Priors Analysis of Time Series, Lecture Notes in Statistics (Vol. 116), New York: Springer-Verlag.
- Lavielle, M. (1998). Optimal segmentation of random processes. *IEEE Transactions on Signal Processing*, **46**, 1365-1373.
- Ombao, H.C., Raz, J.A., Von Sachs, R. and Malow, B.A. (2001). Automatic statistical analysis of bivariate nonstationary time series. *J. Am. Statist. Assoc.*, **96**, 543-560.
- Priestley, M. (1965). Evolutionary spectra and nonstationary processes. *J.R. Statist. Soc. B*, **27**, 204-237.
- Punskaya, E., Andrieu, C., Doucet, A. and Fitzgerald, W.J. (2002). Bayesian curve fitting using MCMC with applications to signal segmentation. *IEEE Transactions on Signal Processing*, **50**, 747-758.
- Shumway, R.H. and Stoffer, D.S. (2006). *Time Series Analysis and Its Applications: With R Examples*, Second Edition, Springer, New York.
- Thomson, D.J. (1990) Time series analysis of Holocene climate data. *Phil. Trans. Roy. Soc. A*, **330**, 601-616.
- Timmermann, A., Oberhuber, J., Bacher, A., Esch, M., Latif, M. and Roeckner, E. (1999). Increased El Niño frequency in a climate model forced by future greenhouse warming. *Nature*, **398**, 694-697.
- Trenberth, K.E. and Hoar, T.J. (1996). The 1990-1995 El Niño-southern oscillation event: longest on record. *Geophysical Research Letters*, **23**, 57-60.
- Wahba, G. (1980). Automatic smoothing of the log periodogram. *J. Am. Statist. Assoc.*, **75**, 122-132.
- Wahba, G. (1990). *Spline Models for Observational Data*. SIAM, Philadelphia. CBMS-NSF Regional Conference Series in Applied Mathematics, Vol. 59.
- Whittle, P (1957). Curve and periodogram smoothing. *J.R. Statist. Soc. B*, **19**, 38-47.
- Wood, S., Jiang, W. and Tanner, M.A. (2002). Bayesian mixture of splines for spatially adaptive nonparametric regression. *Biometrika*, **89**, 513-528.