



Optimally adaptive Bayesian spectral density estimation for stationary and nonstationary processes

Nick James¹ · Max Menzies²

Received: 24 February 2021 / Accepted: 4 May 2022 / Published online: 29 May 2022
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

This article improves on existing Bayesian methods to estimate the spectral density of stationary and nonstationary time series assuming a Gaussian process prior. By optimising an appropriate eigendecomposition using a smoothing spline covariance structure, our method more appropriately models data with both simple and complex periodic structure. We further justify the utility of this optimal eigendecomposition by investigating the performance of alternative covariance functions other than smoothing splines. We show that the optimal eigendecomposition provides a material improvement, while the other covariance functions under examination do not, all performing comparatively well as the smoothing spline. During our computational investigation, we introduce new validation metrics for the spectral density estimate, inspired from the physical sciences. We validate our models in an extensive simulation study and demonstrate superior performance with real data.

Keywords Spectral density estimation · Nonstationary · Reversible jump · Markov chain Monte Carlo · Gaussian process

1 Introduction

Spectral density estimation (SDE) is a common method to understand the autocovariance structure of a stationary time series and perhaps the key technique to detect periodicities in time series data. Of particular importance in the physical sciences, such as mass spectrometry (Todd 1991, 1995), are the peaks of a power spectrum, and the frequencies at which they occur. However, most real-world processes are nonstationary. This has necessitated the development of methods for SDE for nonstationary time series. In this article, we improve upon existing Bayesian techniques to estimate time-varying power spectra.

Our analysis takes place in a nonparametric Bayesian framework and implements a *reversible jump Markov chain Monte Carlo* (RJMCMC) algorithm, assuming a *Gaussian process* (GP) prior. Aligning with the existing literature on

locally stationary time series (Dahlhaus 1997; Adak 1998), we partition a given time series into a finite number of stationary segments. The true log power spectrum is estimated within the Whittle likelihood framework (Whittle 1957) and appropriate optimisation, using the local log periodograms. In doing so, no parametric assumption is made about the time series data itself (time domain), only about the log periodograms (frequency domain).

Our key improvement over existing techniques is the optimisation of an appropriate eigendecomposition, where a smoothing spline covariance structure is used. Therein, we maximise the Whittle likelihood with respect to the number of eigenvalues. This produces a suitable estimate of the power spectrum for both simple and complex spectra. To demonstrate the robustness of this smoothing spline optimisation, we also investigate alternative GP covariance functions. We specify to the stationary case and validate our estimates against analytically known autoregressive (AR) processes. We show that other covariance functions provide no benefit over the eigendecomposition of the smoothing spline. In the process thereof, we introduce new validation metrics that more appropriately reflect the needs of the physical sciences, where the amplitude and frequency of the peaks are paramount.

Our article builds off a rich history of nonparametric frequentist and Bayesian modelling to analyse and fit data. First,

✉ Max Menzies
max.menzies@alumni.harvard.edu

Nick James
nick.james@unimelb.edu.au

¹ School of Mathematics and Statistics The University of Melbourne, Parkville, VIC 3010, Australia

² Beijing Institute of Mathematical Sciences and Applications Tsinghua University, Beijing 101408, China

smoothing splines are the bedrock for such analysis. Coghburn and Davis (1974) and Wahba (1980) pioneered the use of smoothing splines for spectral density estimation in the stationary case, with improvements made by Gangopadhyay et al. (1999), Choudhuri et al. (2004) and others (Wood et al. 2002). Guo et al. (2003) used smoothing spline ANOVA in the locally stationary case. Variations in smoothing splines, such as B-splines (Eilers and Marx 1996), P-splines (Wood 2017) and others (Wahba 1990; Gu 2013), have been used across a range of implementations in the statistics literature.

Alternative covariance functions for GPs, other than smoothing splines, have been used in a wide range of applications to model and analyse data. Rasmussen and Williams (2005) provide an overview of modelling complex phenomena with combinations of covariance functions in GPs. Paciorek and Schervish (2003) perform GP regression and investigate smoothness over a class of nonstationary covariance functions, while Plagemann et al. (2008) utilise such functions within a Bayesian regression framework. More recently, nonparametric regression methods with automatic kernel learning have produced a GP regression framework for kernel learning (Duvenaud et al. 2013) and other frameworks (Lu et al. 2016; Wilson and Adams 2013). However, alternative covariance functions have been used infrequently for the problem of SDE.

More recently, smoothing splines and GPs have been used to study the spectra of nonstationary processes. Whereas the work of Wahba (1980) and Carter and Kohn (1997) places priors over spectra of stationary processes, more recent work has extended this to nonstationary time series, predicated on the work of Dahlhaus (1997) and Adak (1998). Reversible jump MCMC algorithms have been used to partition nonstationary time series into stationary segments. The number and location of these can be estimated flexibly from the data, and the uncertainty quantified, both in the time (Wood et al. 2011; James and Menzies 2021; Prakash et al. 2021) and frequency domain (Rosen et al. 2009, 2012).

In particular, Rosen, Wood and Stoffer have used a smoothing spline GP covariance structure and RJMCMC algorithms in conjunction effectively to perform SDE in the nonstationary setting. Their work partitions nonstationary time segments into locally stationary segments, first with predetermined partitions (Rosen et al. 2009) and then adaptively (Rosen et al. 2012), in each case using a product of local Whittle likelihoods to estimate the time series' time-varying power spectrum. The model includes a particular eigendecomposition of the covariance matrix associated with the smoothing spline. For computational savings, the first 10 eigenvectors are used—this number is fixed.

The main limitation of the aforementioned work is the global degree of smoothness induced by the fixed eigendecomposition. In our article's experiments, we show that this model generally underfits real-world processes. These

usually exhibit more complexity than the autoregressive processes that are canonically used in the validation of such SDE models. This leads to a systemic bias in the literature towards models that produce excessively smooth spectral estimates. Indeed, Hadj-Amar et al. (2019) examined the 10-eigenvector model of Rosen et al. (2012) and concluded that the fixed decomposition tends to produce estimates that are overly smooth. Our article ameliorates this problem by varying the number of eigenvalues used and optimising it by maximising the Whittle likelihood. Quantitatively, we produce better results than other models as measured by new and existing validation metrics, and qualitatively, our estimates provide a more appropriate level of smoothness upon visual inspection.

The article is structured as follows: in Sect. 2, we introduce the mathematical model, including our key improvement. In Sect. 3, we describe the validation metrics we use in the paper to evaluate our spectral density estimates, including existing metrics and a new framework for measuring distances between non-trivial sets of peaks. In Sect. 4, we describe our simulation study, incorporating both our chosen alternative covariance functions and introducing our new validation metrics. In Sect. 5, we apply our method to real data and observe more appropriate fitting of our model to real-world data with complex periodograms. In Sect. 6, we summarise our findings regarding the performance and appropriate smoothness produced by eigendecomposition optimisation versus variation in covariance functions in our GP model.

2 Mathematical model

2.1 Spectral density estimation of stationary time series

Let $(X_t)_{t=1,\dots,n}$ be a discrete real-valued time series.

Definition 1 (X_t) is stationary if

1. Each random variable X_t is integrable with a finite common mean μ for each t . By subtracting the mean, we may assume henceforth $\mu = 0$.
2. The autocovariance $\mathbb{E}[(X_t - \mu)(X_{t+k} - \mu)]$ is a function only of k , which we denote $\gamma(k)$.

We shall reserve the letter n to denote the length of a stationary time series and T to denote the length of a not-necessarily-stationary time series. Spectral analysis allows us to study the second-order properties, particularly periodicity, of a stationary time series expressed in the autocovariance structure. The power spectral density function (PSD) is defined as follows:

$$f(v) = \sum_{k=-\infty}^{\infty} \gamma(k) \exp(-2\pi i v k), \text{ for } -\frac{1}{2} \leq v \leq \frac{1}{2}. \quad (1)$$

Most important are the *Fourier frequencies* $v_j = \frac{j}{n}, j = 0, 1, \dots, n - 1$. Define the *discrete Fourier transform* (DFT) of the time series:

$$Z(v_j) = \frac{1}{\sqrt{n}} \sum_{t=1}^n X_t \exp(-2\pi i v_j t), \text{ for } j = 0, \dots, n - 1. \quad (2)$$

Whittle initially showed that under certain conditions, the DFTs have a *complex normal* (CN) distribution (Whittle 1957, 1954):

$$Z \sim \text{CN}(0, \mathbf{F}), \text{ where } \mathbf{F} = \text{diag}[f(v_0), \dots, f(v_{n-1})]. \quad (3)$$

For each Fourier frequency component, a noisy but unbiased estimate of the PSD $f(v_k)$ is the *periodogram* (Choudhuri et al. 2004), defined by $I(v_j) = |Z(v_j)|^2$. By symmetry, the periodogram contains $m = \lfloor \frac{n}{2} \rfloor + 1$ effective observations, corresponding to $j = 0, 1, \dots, m - 1$. Rosen et al. (2009, 2012) outline a signal plus noise representation of the periodogram:

$$\log I(v_j) = \log f(v_j) + \epsilon_j, \quad (4)$$

where ϵ_j is a $\log(\text{Exp}(1))$ random variable. Recall that a random variable Y is said to have distribution $\text{Exp}(1)$ if it has cumulative density function $\mathbb{P}(Y \leq y) = 1 - e^{-y}$. With this definition, the representation (4) assumes each noise variable ϵ_j (noise in the log periodogram) is an independent and identically distributed (i.i.d.) whose distribution is identical to $\log(Y)$. This model is justified by the following result: for a wide class of theoretical models, Theorem 10.3.2 of Brockwell and Davis (1991) proves that the vector of quotients

$$\left(\frac{I(v_0)}{f(v_0)}, \dots, \frac{I(v_{m-1})}{f(v_{m-1})} \right) \quad (5)$$

converges in distribution to a vector of i.i.d. $\text{Exp}(1)$ random variables. We reiterate that this theorem applies to the noise variables in the log periodogram, not the original time series. In particular, we make use of no parametric assumptions regarding the underlying process or noise of the original time series (time domain) and only of the log periodogram (frequency domain) as specified by (4).

For our applications, we assume the quotient $\frac{I}{f}$ is approximately exponentially distributed with mean 1. This is a simpler representation of the second moment of the distribution, reducing the problem of covariance estimation to a simpler problem of mean estimation. Note the periodogram

oscillates around the true spectral density, so there is a delicate balance between inferring the spectrum of a process and excessive smoothing, leading to negligible inference. We remark that unlike most statistical models, this does not assume Gaussian noise; in fact, by the aforementioned theoretical results, we assume exponentially distributed random noise instead.

2.2 Gaussian process regression

Definition 2 A *Gaussian process* is a collection of random variables, any finite collection of which have a multivariate Gaussian distribution (Rasmussen and Williams 2005). Such a process $f(\mathbf{x})$ is uniquely determined by its mean function $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ and covariance function $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$. We write $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$.

We assume the standard “noisy observation” Gaussian additive error model $y = f(\mathbf{x}) + \epsilon$, in which ϵ is i.i.d Gaussian noise with variance σ_n^2 (Rasmussen and Williams 2005). Then, the prior over the noisy observations is $\text{cov}(\mathbf{y}) = k(\mathbf{x}, \mathbf{x}') + \sigma_n^2 I$. The joint distribution of observed response values, \mathbf{y} , and function values evaluated at test points under the prior distribution, \mathbf{f}_* , are expressed as

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} k(\mathbf{x}, \mathbf{x}) + \sigma_n^2 I & k(\mathbf{x}, \mathbf{x}_*) \\ k(\mathbf{x}_*, \mathbf{x}) & k(\mathbf{x}_*, \mathbf{x}_*) \end{bmatrix} \right).$$

Finally, the GP regression predictive equations are as follows (Rasmussen and Williams 2005):

$$\mathbf{f}_* | \mathbf{x}, \mathbf{y}, \mathbf{x}' \sim \mathcal{N}(\bar{\mathbf{f}}_*, \text{cov}(\mathbf{f}_*)), \quad (6)$$

where the mean and covariance are defined as

$$\begin{aligned} \bar{\mathbf{f}}_* &:= \mathbb{E}[\mathbf{f}_* | \mathbf{x}, \mathbf{y}, \mathbf{x}_*] = k(\mathbf{x}_*, \mathbf{x}) [k(\mathbf{x}, \mathbf{x}) + \sigma_n^2 I]^{-1} \mathbf{y}, \\ \text{cov}(\mathbf{f}_*) &:= k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, \mathbf{x}) [k(\mathbf{x}, \mathbf{x}) + \sigma_n^2 I]^{-1} k(\mathbf{x}, \mathbf{x}_*). \end{aligned} \quad (7)$$

2.3 Model and priors: stationary case

We follow (Choudhuri et al. 2004; Rosen et al. 2012) and use the Whittle likelihood function to model the log periodogram within a Bayesian regression framework. Let $\mathbf{y} = (y(v_0), \dots, y(v_{m-1}))$ be the log of the periodogram, $\mathbf{f} = (f(v_0), \dots, f(v_{m-1}))$ be the PSD, and $\mathbf{g} = (g(v_0), \dots, g(v_{m-1}))$ be the log of the PSD. The likelihood of the log periodogram given the true power spectrum can be approximated by

$$p(\mathbf{y}|\mathbf{f}) = (2\pi)^{-m/2} \prod_{j=0}^{m-1} \exp\left(-\frac{1}{2} \left[\log f(v_j) + \frac{I(v_j)}{f(v_j)} \right]\right). \tag{9}$$

Following (Rosen et al. 2009, 2012), we rewrite Eq. (4) as

$$y(v_j) = g(v_j) + \epsilon_j, \tag{10}$$

To obtain a flexible estimate of the spectrum, we place a GP prior over the unknown function \mathbf{g} . That is, we assume $\mathbf{g} \sim N(\mathbf{0}, k(\mathbf{x}, \mathbf{x}_*))$. This is determined by its covariance function. Most of the SDE literature to date has used smoothing splines (Wahba 1980) or other spline varieties for this covariance structure. In subsequent sections, we investigate both smoothing splines and other covariance functions. We choose kernels that have been frequently examined elsewhere in the literature: the squared exponential, Matern_{3/2}, Matern_{5/2}, Sigmoid kernel and combinations of these such as the squared exponential + Sigmoid (Rasmussen and Williams 2005; Paciorek and Schervish 2003; Plagemann et al. 2008; Lu et al. 2016; Wilson and Adams 2013).

2.4 Nonstationary model and priors

In this section, we describe the model assumed throughout the paper and associated priors. For the majority of this paper, the precise form of nonstationarity we assume is a Dahlhaus piecewise stationary process (Dahlhaus 1997), which can be described as follows. Let a time series of length T , $(X_t)_{t=1, \dots, T}$, consist of an unknown number of segments m and change points $\tau_1, \dots, \tau_{m-1}$ between segments. For notational convenience, set $\tau_0 = 0, \tau_m = T$. Then, the entire time series X_t can be written

$$X_t = \sum_{i=1}^m \mathbf{1}_{[\tau_{i-1}+1, \tau_i]} X_t^i, \tag{11}$$

where each X_t^i is an independent and stationary time series over the interval $[\tau_{i-1} + 1, \tau_i]$. Following the notation of Sect. 2.3, let $n_i = \tau_i - \tau_{i-1}$ be the length of the i th segment, $m_i = \lfloor \frac{n_i}{2} \rfloor + 1$ be the effective length of each periodogram, and let $\mathbf{f}_i \in \mathbb{R}^{m_i}$ denote the PSD of the stationary time series X_t^i .

By independence of the processes, the Whittle likelihood approximation of the nonstationary model for a given partition $\boldsymbol{\tau} = (\tau_1, \dots, \tau_{m-1})$ is as follows:

$$L(\mathbf{f}_1, \dots, \mathbf{f}_m | \boldsymbol{\tau}, \mathbf{y}_1, \dots, \mathbf{y}_m) = \prod_{i=1}^m p(\mathbf{y}_i | \mathbf{f}_i), \tag{12}$$

where each local Whittle likelihood $p(\mathbf{y}_i | \mathbf{f}_i)$ is calculated according to (9).

Following Rosen et al. (2012), we place the following prior distributions on both the number of segments m and the segment partitions $\boldsymbol{\tau} \in \mathbb{R}^{m-1}$.

1. The prior on the number of segments m is an integer-valued uniform distribution with minimal number 1 and maximum number M .
2. The prior on each change point $\tau_i, i > 0$ is a uniform distribution conditional on the previous change point τ_{i-1} . Specifically, we allow τ_i to take any value with equal probability within a range $[\tau_{i-1} + t_{\min}, T]$. This requires the specification of a minimal possible segment length t_{\min} , which we set to be 40.
3. The priors on each log PSD $\mathbf{g}_i = \log \mathbf{f}_i, i = 1, \dots, m$ are independent and are specified as in Sect. 2.3.

2.5 MCMC implementation and choice of covariance functions

The primary contribution of this paper is an improved method for SDE of nonstationary time series. Our work uses the RJMCMC of Rosen et al. (2012), found in Appendix C. This scheme partitions the time series into locally stationary segments and models the log PSD $\log f(\cdot)$ of the segments with a Gaussian process with covariance matrix Ω . For computational savings, they employ an eigendecomposition $\Omega = QDQ^T$ and retain only the 10 largest eigenvectors in D . That is, they let D_{10} be the truncated diagonal matrix and set $X = QD_{10}^{1/2}$ as their design matrix. Let β be a vector of unknown regression coefficients with prior distribution $N(0, \tau^2 I)$ (where τ^2 is a smoothing parameter). Then, the successive estimate for the PSD is $\log f(\cdot) = X\beta$, iterated within the RJMCMC.

Given that the eigendecomposition is so critical, the exact number of eigenvalues should be estimated using the data. Using a fixed number of 10 eigenvalues (as do Rosen et al., detailed above) yields a smoothing spline that may be too smooth depending on the process being estimated. If the true spectrum is highly smooth, our estimator will exhibit too much variance; if the spectrum is particularly complex, our estimator will exhibit excessive bias. Instead, we select an optimal number of eigenvalues M by analysing posterior samples of the log spectrum for a range of eigenvectors, integrating over all possible values of regression coefficients, and maximising the marginal likelihood. We select our number of eigenvalues as follows:

$$\hat{M} = \operatorname{argmax}_{M \geq 1} \int p(\mathbf{y} | M, X_M, \beta) p(\beta | M, X_M) d\beta, \tag{13}$$

where \mathbf{y} is the log periodogram data, X_M is the design matrix under an eigendecomposition using M eigenvalues, and β is a vector of coefficients determining the weight for respective

eigenvectors. We refer to the corresponding eigendecomposition $\Omega = QD_{\hat{M}}Q^T$ as the *optimal smoothing spline*.

We remark that the determination of this \hat{M} requires post-processing after a sequence of independent MCMC runs, both in the stationary and nonstationary case, and in all implemented experiments throughout the paper. Specifically, \hat{M} is not computed within any individual MCMC iteration, but by analyzing separate chains for each possible number of eigenvalues M .

We also remark that the value of \hat{M} will generally be larger for more complex spectra, with a smaller number of eigenvectors being sufficient for relatively simple AR(1) and AR(2) processes. As for sample size, the relationship with \hat{M} may be a little more complex. A larger length of the time series n or T may increase the complexity of the spectrum and hence \hat{M} , but not necessarily. For example, sampling a very large number of data points from a simple AR(1) process would not yield a higher optimal value of \hat{M} . In fact, a small sample size could lead to additional difficulty. If an autoregressive process is of insufficient length, it may be difficult to extract information about the underlying signal in the periodogram. This is particularly important for time series that would generate complex spectra such as the AR(4) processes outlined in this paper. A larger length of the time series will also lead to greater computational cost when computing the periodogram. This could be costly in a reversible jump procedure (especially when there are relatively few segments), as this computation must occur many times until the algorithm’s termination.

In subsequent sections, we also explain the possibility of implementing a penalty on the number of eigenvalues. This would yield an alternative eigendecomposition we will refer to as the *penalised optimal smoothing spline*.

In subsequent sections, we examine the performance of the optimal smoothing spline against both simulated and real data. First, we specialise to the stationary case, where we may validate this optimal smoothing spline in the case of autoregressive processes with known analytic spectra. In this setting, there is no need for the partition of the time series, so the RJMCMC scheme reduces to a Metropolis–Hastings scheme described in Appendix 1. In this implementation, we are also free to substitute alternative GP covariance functions $k(\mathbf{x}, \mathbf{x}')$ (see Definition 2 and Appendix 1). The optimal smoothing spline compares favourably not only to the 10-eigenvector decomposition but also to the other covariance functions and provides further justification of its utility.

Having established that the optimal smoothing spline performs the best of our covariance functions in the stationary case, we proceed to a simulated nonstationary time series, once again using the RJMCMC. By examining three adjoined autoregressive processes with varying spectral complexity, we demonstrate the improvement in our method over the

10-eigenvector decomposition by validating each segment against its known analytic spectra. Finally, we proceed to test our optimal smoothing spline on real data, which is much less smooth than the autoregressive processes studied until this point. We show that the 10-eigenvector decomposition drastically underfits the many peaks of more complex real-world data. Whenever we use the reversible jump procedure to partition a nonstationary time series, we present our results conditional on the modal number of change points as a default.

2.6 Penalised likelihood

In this section, we discuss the possibility and potential advantages of imposing a penalty term in the selection of best number of eigenvalues M . Imposing a penalty term would inevitably select a smaller number of eigenvalues and carries several advantages:

1. If desired, it would keep the estimate of the PSD smoother.
2. This could reduce aliasing arising due to an insufficient number of data points.
3. When we are uncertain of ideal model complexity, we may prefer simpler models.
4. Different users may impose different costs of error. A user more sensitive to bias may favour the optimal, while a user more sensitive to variance may use a smaller number of eigenvalues.

Imposing penalty terms is an intricate and complex question depending on the exact problem at hand. In this paper, we take an approach informed by observations of the behaviour of the marginal log likelihood against the number of eigenvalues. Observe subsequent plots of the log likelihood, both in synthetic experiments (Figs. 1 and 2) and real data (Fig. 4). Initially, an increasing number of eigenvectors provide a substantial (approximately) linear increase in the marginal log likelihood, followed by a clear “elbow” in the plots, where the rate of increase drastically reduces or levels off altogether. We wish to isolate this elbow using a penalised approach. Specifically, we want to impose a linear penalty term on the log marginal likelihood—this will help isolate the value of M in which the growth of the log marginal likelihood changes from “fast” to “slow”.

For this purpose, let $L(M)$ denote the log marginal likelihood

$$\mathcal{L}(M) = \log \int p(\mathbf{y}|M, X_M, \beta)p(\beta|M, X_M)d\beta. \tag{14}$$

Consider an interval of consideration, $[M_{\min}, M_{\max}]$. In our implementation, $M_{\min} = 10$, while M_{\max} is usually 70. Let

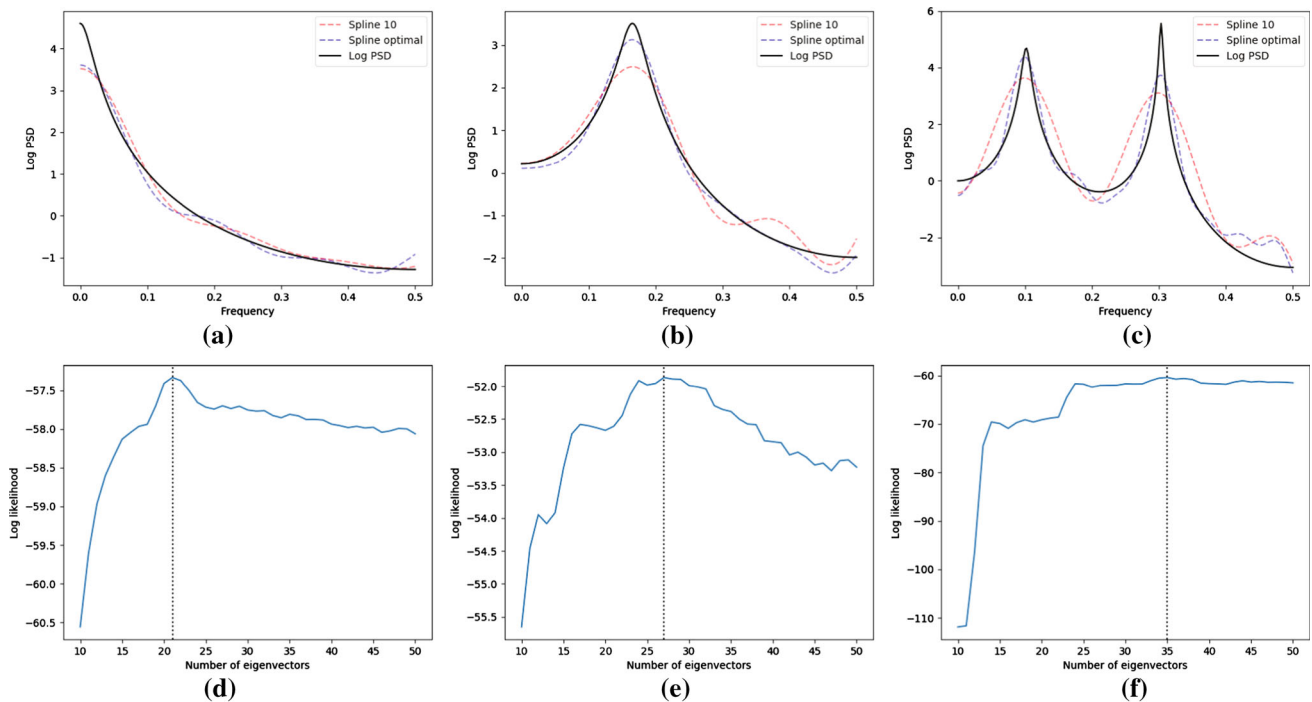


Fig. 1 Analytic and estimated log PSD for 10-eigenvector and optimal decomposition of the smoothing spline models for **a** AR(1), **b** AR(2), **c** AR(4). Log likelihood vs number of eigenvectors for **d** AR(1), **e** AR(2)

and **f** AR(4) process. The peak in the log likelihood determines the number of eigenvectors in the optimal model

$$\alpha = \frac{\mathcal{L}(M_{\max}) - \mathcal{L}(M_{\min})}{M_{\max} - M_{\min}} \tag{15}$$

be the “overall gradient” of the log likelihood curve. We impose the following penalty to determine the optimal penalised \hat{M}_p :

$$\hat{M}_p = \operatorname{argmax}_{M_{\min} \leq M \leq M_{\max}} \mathcal{L}(M) - \alpha M \tag{16}$$

$$= \operatorname{argmax}_{M_{\min} \leq M \leq M_{\max}} \log \left(\int p(\mathbf{y}|M, X_M, \beta) p(\beta|M, X_M) d\beta \right) - \alpha M. \tag{17}$$

Essentially, this penalised optimal \hat{M}_p aims to select the value of M at the elbow of the log marginal likelihood curve at “the point” (a heuristic) where (at least in the examples we observe) quick linear growth changes to much slower linear growth.

3 Validation metrics

In this section, we describe in detail both existing and new validation metrics—these will be used in the subsequent simulation study (Sect. 4) to compare our spectral density estimates $\hat{\mathbf{g}}$ of autoregressive processes with the known ana-

lytic log PSD **g**. First, we outline the existing metrics we use, and then we describe in detail a new framework for measuring distances between spectral peaks, including the possibility of multiple non-trivial peaks in more complex spectra.

3.1 Existing metrics

The three existing metrics we use are root mean squared error (RMSE), defined as $\sqrt{\frac{1}{m} \sum_{j=0}^{m-1} (\hat{g}(v_j) - g(v_j))^2}$, mean absolute error (MAE), defined as $\frac{1}{m} \sum_{j=0}^{m-1} |\hat{g}(v_j) - g(v_j)|$, and the Wasserstein distance between the finite sets $\{g(v_j)\}$ and $\{\hat{g}(v_j)\}$. Specifically, this converts a finite set to a uniform probability measure over its elements and computes the Wasserstein distance between these measures.

In further detail, let S be a finite subset of \mathbb{R} . To S , we can associate a uniform measure defined as a weighted sum of Dirac delta measures

$$\mu_S = \frac{1}{|S|} \sum_{s \in S} \delta_s. \tag{18}$$

Integrating μ_S yields a cumulative density function F and its associated quantile function F^{-1} . Concretely, if $S =$

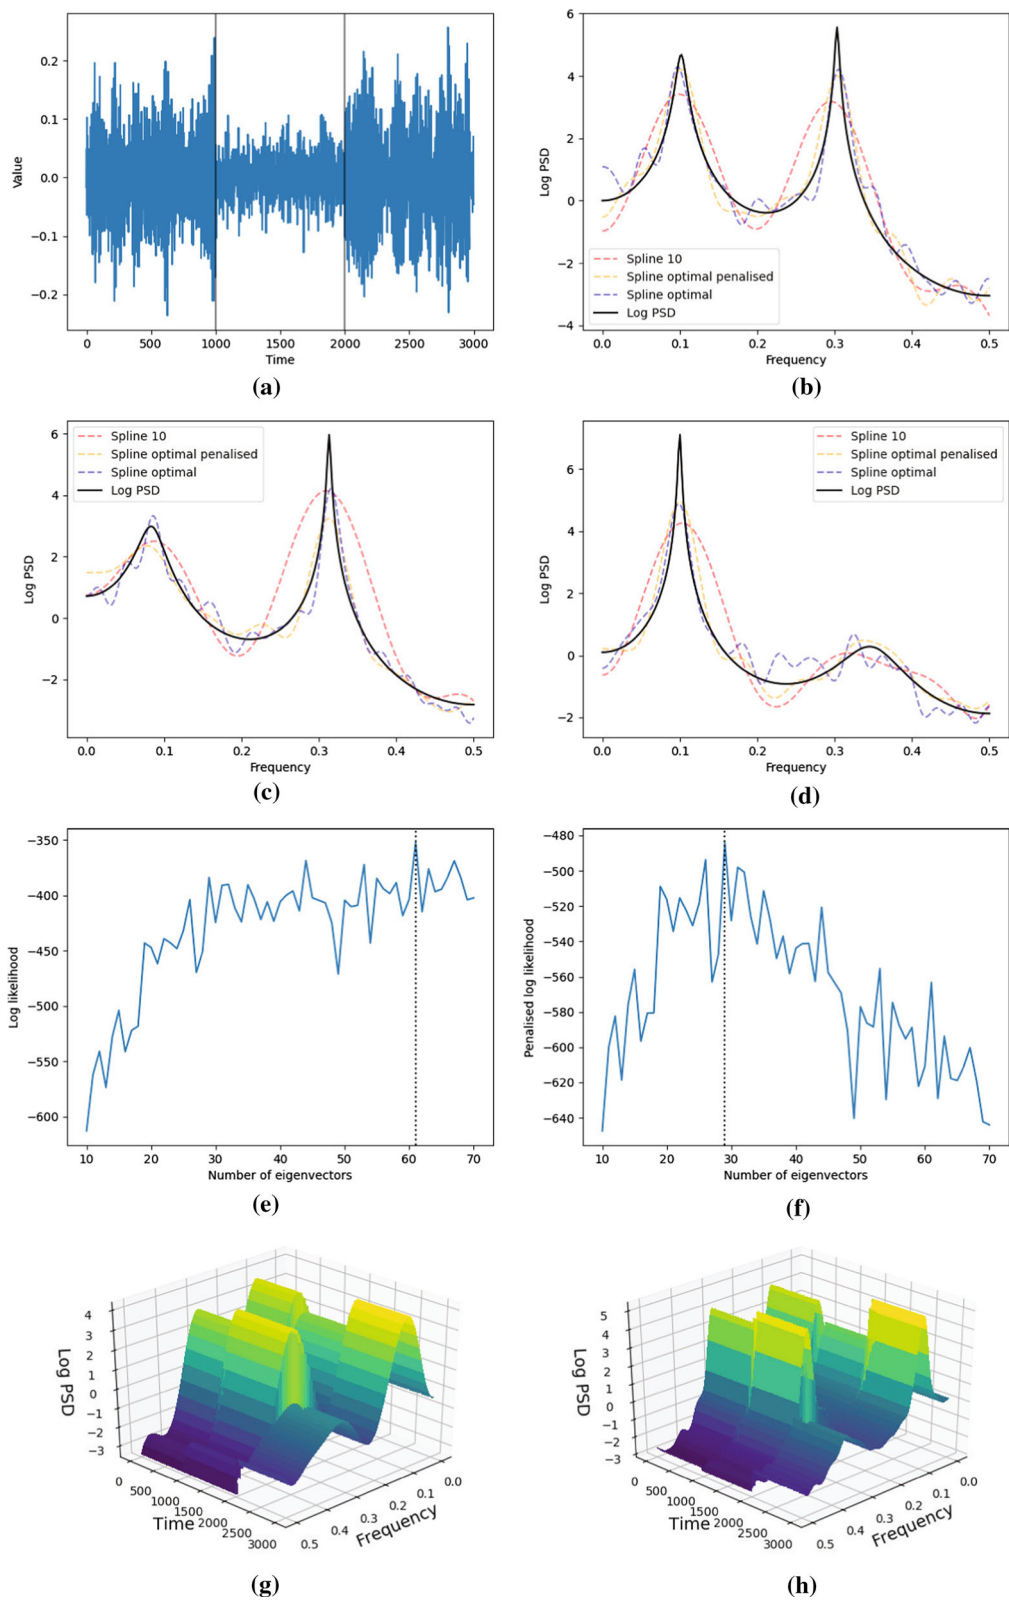


Fig. 2 Analytic and estimated log PSD for 10-eigenvector and optimal decomposition of smoothing spline models for nonstationary time series. **a** Realisation of the time series. Spectral estimates for **b** Segment 1. **c** Segment 2. **d** Segment 3, each using the optimal and penalised optimal spline. **e** Plot of log marginal likelihood vs number of eigenvectors,

using Eq. (14). The optimal number of eigenvectors is 61. **f** Plot of the penalised log likelihood function, using Eq. (17). The penalised optimal number of eigenvectors is 29. Spectral estimates for Spectral surface (time-varying PSD) for **g** 10-eigenvector and **h** optimal eigen-decomposition

$\{s_1, \dots, s_n\} \subset \mathbb{R}$, then

$$F = \sum_{j=1}^{n-1} \frac{j}{n} \mathbb{1}_{[s_j, s_{j+1})} + \mathbb{1}_{[s_n, \infty)}, \tag{19}$$

$$F^{-1} = \sum_{j=1}^n s_j \mathbb{1}_{(\frac{j-1}{n}, \frac{j}{n}]}. \tag{20}$$

Given two finite sets S, T , their L^1 -Wasserstein distance is computed as $d_W(S, T) = W_1(\mu_S, \mu_T)$ and be computed as

$$d_W(S, T) = W_1(\mu_S, \mu_T) = \int_0^1 |F^{-1} - G^{-1}| dx, \tag{21}$$

where F^{-1} and G^{-1} are defined as in (20) from S and T , respectively. For more detail, see (James and Menzies 2022), Appendix C.

3.2 Sets of spectral peaks and the proximity matching criterion

Our new validation metrics are inspired by applications in the physical sciences. There, a single dominant peak is often the greatest importance to observing scientists (Todd 1991). Such a peak has two attributes, its amplitude, $\max \mathbf{g}$ and the frequency at which it occurs, $\operatorname{argmax} \mathbf{g}$. With this in mind, we introduce the simplest form of two validation metrics, $|\max \hat{\mathbf{g}} - \max \mathbf{g}|$ and $|\operatorname{argmax} \hat{\mathbf{g}} - \operatorname{argmax} \mathbf{g}|$, measuring how well the spectral density estimation determines the amplitude and frequency of the greatest peak.

Next, we consider the possibility of multiple peaks in an analytic or estimated power spectrum, as is the case for an AR(4) process. The remainder of this section is dedicated to carefully exploring this possibility, including how to measure distance between (non-trivial) sets of estimated and analytical spectral peaks.

First, we consider the simplest possible case, assuming the following a priori:

1. the analytic and estimated power spectra have the same number of peaks, ρ_1, \dots, ρ_r for \mathbf{g} and $\hat{\rho}_1, \dots, \hat{\rho}_r$ for $\hat{\mathbf{g}}$;
2. there is a bijection between the two sets of peaks, for which each ρ_i is the closest peak of \mathbf{g} to $\hat{\rho}_i$, and vice versa.

Informally, this is the ideal situation, in which the peaks of \mathbf{g} can be matched one-to-one to their closest counterpart among the peaks of $\hat{\mathbf{g}}$, and vice versa. In this ideal case, we can define a validation metric between peaks' collective amplitudes and frequencies, respectively, for any $p \geq 1$ as follows:

$$\left(\frac{|g(\rho_1) - \hat{g}(\hat{\rho}_1)|^p + \dots + |g(\rho_r) - \hat{g}(\hat{\rho}_r)|^p}{r} \right)^{\frac{1}{p}}, \tag{22}$$

$$\left(\frac{|\rho_1 - \hat{\rho}_1|^p + \dots + |\rho_r - \hat{\rho}_r|^p}{r} \right)^{\frac{1}{p}}. \tag{23}$$

If $p = 2$, we return a (normalised) Euclidean distance between tuples (ρ_1, \dots, ρ_r) and $(\hat{\rho}_1, \dots, \hat{\rho}_r) \in \mathbb{R}^r$. If $p = 1$, we return an L^1 distance between the tuples. In subsequent paragraphs, we will see why the normalisation is appropriate for a more general framework.

More generally, we wish to outline a framework that specifically tests for whether this bijection exists, and define a validation measure that is well-defined regardless of whether the bijection exists. For this purpose, let S be the finite set of the (non-trivial) peaks of an analytic log power spectrum \mathbf{g} , and \hat{S} be the set of (non-trivial) peaks of an estimate $\hat{\mathbf{g}}$. The most naive definition of a peak as a local maximum of \mathbf{g} or $\hat{\mathbf{g}}$, respectively, may be unsuitable for this purpose—due to the wiggly nature of spectral density estimates, there may exist local maxima of $\hat{\mathbf{g}}$ that are insubstantial and should be excluded. An algorithmic framework for the refinement of peaks in the power spectra must be employed. A flexible algorithmic framework that does this for local maxima and minima is detailed in Appendix 2. For example, local maxima that are significantly less than the global maximum of the log PSD may be excluded as insubstantial under a variety of definitions and parameters. In our experiments with the AR(4) process, only a simple approach is needed: set $\delta = 2$; any local maximum $\hat{\rho}$ of $\hat{\mathbf{g}}$ with $\hat{g}(\hat{\rho}) < \max \hat{\mathbf{g}} - \delta$ is excluded. The same results are produced for any $\delta \in [2, 4]$, demonstrating the robustness of this refinement procedure. As this is applied to the log power spectrum, such an inequality is preserved when scaling the initial time series X_t by an affine transformation $X'_t = aX_t + b$. As an additional remark, we note that the precise number r of (non-trivial) peaks is inevitably determined by the precise refinement algorithm we use. We discuss further details in Appendix 2.

Having identified such sets S and \hat{S} , either through a choice of parameters within such a framework or by manual inspection, we define the *proximity matching criterion* as follows: say S and \hat{S} satisfy this criterion if:

1. there is a function $f : S \rightarrow \hat{S}$, where for each $s \in S$, $f(s)$ is the unique closest element of \hat{S} to s ;
2. there is a function $\hat{f} : \hat{S} \rightarrow S$, where conversely for each $\hat{s} \in \hat{S}$, $\hat{f}(\hat{s})$ is the unique closest element of S to \hat{s} ;
3. The functions f and \hat{f} are mutually inverse.

For example, suppose an analytic log PSD \mathbf{g} has a set of non-trivial peaks $S = \{0.1, 0.2\}$. If a certain estimate $\hat{\mathbf{g}}$ has a set of non-trivial peaks $\hat{S}_1 = \{0.11, 0.19\}$, then there exists a bijection between S and \hat{S}_1 that associates

with each element of S the closest element of \hat{S}_1 , and vice versa. The bijection can be written down explicitly as $f : S \rightarrow \hat{S}_1, f(0.1) = 0.11, f(0.2) = 0.19$ and $\hat{f} : \hat{S}_1 \rightarrow S, \hat{f}(0.11) = 0.1, \hat{f}(0.19) = 0.2$. Then, sets S and \hat{S}_1 satisfy the proximity matching criterion.

For a second example, suppose an alternative estimate \hat{g}_2 has a set of non-trivial peaks $\hat{S}_2 = \{0.11, 0.12\}$. Then, there is no bijection that matches each element of \hat{S}_2 to the closest element on S , as both 0.11 and 0.12 are closest to the same element $0.1 \in S$. Thus, no such proximity-matching bijection exists. For another example, if a third estimate \hat{g}_3 has a set of non-trivial peaks $\hat{S}_3 = \{0.11\}$, then no bijection exists with S , so the criterion is not satisfied.

This criterion formalises the ideal case described above, where the (non-trivial) peaks of \mathbf{g} and $\hat{\mathbf{g}}$ can be matched one-to-one and hence labelled ρ_1, \dots, ρ_r and $\hat{\rho}_1, \dots, \hat{\rho}_r$ via the mutually inverse bijective functions f and \hat{f} . Simply put, the criterion formalises the notion that one can naturally pair up the non-trivial peaks of \mathbf{g} and $\hat{\mathbf{g}}$ to facilitate a convenient comparison.

3.3 Measures between sets of non-trivial peaks

So far, we have defined a criterion that specifies the most convenient case, where one can conveniently measure distances between matched sets of peaks S and \hat{S} using (22) and (23). Finally, we define a measure between sets S and \hat{S} that makes sense in full generality, even when the proximity matching criterion fails. Again let $p \geq 1$; we adopt the *semi-metrics* defined in James et al. (2020) and let

$$d^p(S, \hat{S}) = \left(\frac{\sum_{\hat{s} \in \hat{S}} d(\hat{s}, S)^p}{2|\hat{S}|} + \frac{\sum_{s \in S} d(s, \hat{S})^p}{2|S|} \right)^{\frac{1}{p}}, \quad (24)$$

where $d(\hat{s}, S)$ is the minimal distance from \hat{s} to S , and vice versa. This validation semi-metric can be applied to measure the collective distance between the sets of peaks of an analytic and estimated spectra under any circumstance, conditional on an appropriate selection of non-trivial peaks. (Such selection may be performed algorithmically or by inspection). Moreover, in the case where the proximity matching criterion is satisfied and the sets of peaks can be matched up and labelled as ρ_1, \dots, ρ_r and $\hat{\rho}_1, \dots, \hat{\rho}_r$, it simply reduces to the metric (23), together with its normalising factor, between tuples (ρ_1, \dots, ρ_r) and $(\hat{\rho}_1, \dots, \hat{\rho}_r) \in \mathbb{R}^r$.

To summarise, we have defined a general validation (semi)-metric between the collective sets of peaks of an analytic and estimated power spectra \mathbf{g} and $\hat{\mathbf{g}}$, respectively. For full generality and precision, the sets of peaks must be selected from the data with a flexible algorithmic framework suitable for the application at hand and not merely chosen by observation from the shape of the analytic spectrum. In

the most suitable case, this semi-metric will reduce to what is still a new validation metric defined in Eq. (23).

4 Simulation study

In this section, we study the properties of our estimators with a detailed simulation study. We simulate three stationary autoregressive processes, AR(1), AR(2) and AR(4) and validate our spectral density estimates against their known analytic power spectra. For each experiment, we run 50 simulations with our MCMC scheme, either RJMCMC or Metropolis–Hastings, as described in Sect. 2. Each sampling scheme consists of 10,000 iterations, with 5000 used for burn-in, and each process contains $n = 500$ data points. Following the notation of Sect. 2, we set $m = 251$ and consider Fourier frequencies $\nu_j = \frac{j}{n}, j = 0, 1, \dots, m - 1$, the appropriate range of the periodogram and spectrum. As such, the range of the frequency axis in all plots is $0 \leq \nu \leq \frac{1}{2}$. Whenever we report results from our validation metrics (described in Sect. 3), we always average results across the 50 performed simulations. Henceforth, ϵ_t denotes a white noise random variable with distribution $N(0, 1)$.

First, we generate data from an AR(1) process $x_t = 0.9x_{t-1} + \epsilon_t$. We compute the log periodogram from the observed data and estimate the log PSD and GP hyperparameters with our MCMC scheme. Table 1 indicates that the best performing covariance functions are the squared exponential and Sigmoid. The log PSD has high power at low frequency and gradually declines in power when moving towards higher frequency components. The spectrum does not exhibit any sharp peaks and it is unsurprising that the highly smooth squared exponential performs well in estimating the log PSD. The worst performing covariance function is the 10-eigenvector decomposition of the smoothing spline of Rosen et al. (2012). There is little improvement with the optimal smoothing spline model in this case, as shown in Fig. 1a. This experiment suggests that for a simple spectrum, there is limited improvement in optimising the number of basis functions in the smoothing spline and alternative GP covariance functions provide superior performance.

Next, we generate data from an AR(2) process $x_t = 0.9x_{t-1} - 0.8x_{t-2} + \epsilon_t$, compute the log periodogram and estimate the log PSD and GP hyperparameters with the MCMC schemes. Table 1 indicates that several covariance functions perform similarly well: SQE+Sigmoid, squared exponential, both Matern functions and the optimal spline. The Sigmoid kernel and 10-eigenvector decomposition of the spline perform worse. The log PSD has a relatively flat gradient and a mild peak at $\nu \sim 0.15$. Once again, smoother covariance functions tend to perform better, while the Sigmoid covariance function, popularised for its ability to model abrupt changes in data modelling problems (Rasmussen and

Table 1 Results for synthetic data experiments

Covariance function	RMSE	MAE	Wasserstein	$ \operatorname{argmax} g - \operatorname{argmax} \hat{g} $	$ \max g - \max \hat{g} $
AR(1) model					
Spline ₁₀	0.26	0.20	0.18	0	0.90
Spline _{optimal}	0.25	0.19	0.16	0	0.93
Squared exponential	0.16	0.12	0.12	0	0.49
Matern _{3/2}	0.19	0.15	0.13	0	0.51
Matern _{5/2}	0.19	0.14	0.12	0	0.49
Sigmoid	0.15	0.11	0.11	0	0.34
SQE+Sigmoid	0.18	0.14	0.13	0	0.40
AR(2) model					
Spline ₁₀	0.33	0.26	0.24	0.0020	0.73
Spline _{optimal}	0.21	0.17	0.14	0.0020	0.36
Squared exponential	0.22	0.17	0.15	0.0020	0.56
Matern _{3/2}	0.21	0.17	0.14	0.0020	0.24
Matern _{5/2}	0.19	0.15	0.13	0.0020	0.37
Sigmoid	0.32	0.24	0.22	0.0040	0.85
SQE+Sigmoid	0.21	0.17	0.15	0.0020	0.53
AR(4) model					
Spline ₁₀	0.82	0.63	0.54	0.20	2.23
Spline _{optimal}	0.43	0.33	0.26	0.20	1.57
Squared exponential	0.45	0.31	0.27	0.20	1.98
Matern _{3/2}	0.37	0.27	0.21	0.20	1.75
Matern _{3/2}	0.40	0.29	0.24	0.20	1.90
Sigmoid	0.97	0.75	0.41	0.21	2.12
SQE+Sigmoid	0.45	0.32	0.26	0.20	2.06

Mean error over 50 simulations is recorded for various covariance functions and validation metrics. Each MCMC scheme consists of 10,000 iterations where 5000 are used for burn-in. For each AR process, $n = 1000$

Williams 2005), is a poor performer. Figure 1b demonstrates a substantial improvement in estimating the log PSD when an optimal number, larger than 10, of basis functions is used. In particular, the estimator does a better job of detecting the maximum amplitude of the power spectrum.

Third, we follow (Edwards et al. 2019), generate data from an AR(4) process $x_t = 0.9x_{t-1} - 0.9x_{t-2} + 0.9x_{t-3} - 0.9x_{t-4} + \epsilon_t$ and again estimate the log PSD and GP hyperparameters with the MCMC schemes. The log PSD of the AR(4) process is more complex than the prior two experiments. There are two peaks of different amplitudes in the spectrum, and the spectrum changes more abruptly than the AR(1) and AR(2) spectra. Table 1 indicates that the best performing covariance functions are the optimal smoothing spline and the two Matern functions. The worst performing covariance functions are the 10-eigenvector smoothing spline and the Sigmoid kernel. Once again, we see a substantial improvement in estimation when an optimal number of basis functions is used, displayed in Fig. 1c.

As the AR(4) process has two peaks, we may also test the proximity matching criterion of Sect. 3.2 and measure the distance between the collective sets of peaks of

the estimated and analytic power spectra. In Table 2, we note that in this relatively simple example with just two peaks, all covariance functions satisfy this criterion. As such, the more general semi-metric (24) reduces to the simpler metric between frequencies (23). We also include the collective distance between amplitudes in the table, as defined in (22). These distances show that the optimal smoothing spline does the best job at estimating both peaks at the same time.

Figure 1d-f plots the Whittle likelihood against the number of eigenvectors for the AR(1), AR(2) and AR(4) processes, respectively—this selects the optimal number of basis functions for each. We see that consistently more than 10 basis functions are required to optimally model the data; in fact, the choice of 10 basis functions in the eigendecomposition is only really appropriate for the AR(1) and AR(2) spectrum. Even the synthetic AR(4) spectrum, but more so real-world processes, benefits from an optimised decomposition that can better suit their complexity and lack of smoothness. Indeed, real-world time series may have multiple periodic components represented in complex power spectra with multiple peaks, of varying amplitudes. The

Table 2 Amplitude distance between collective sets of peaks of analytic and estimated AR(4) power spectra g and \hat{g} , respectively

AR(4) process: distance between sets of peaks	
Covariance function	Amplitude distance
Spline ₁₀	1.64
Spline _{optimal}	0.99
Squared exponential	1.24
Matern _{3/2}	1.18
Matern _{5/2}	1.31
Sigmoid kernel	1.57
SQE+Sigmoid	1.34

The proximity matching criterion is satisfied in every instance

fixed 10-eigenvector decomposition is excessively smooth for these more complex spectra. In Appendix A, we examine an alternative AR(4) process, known in the literature for its difficulties.

Next, we include a simulation study to demonstrate the efficacy of our improved method at both segmenting a non-stationary time series and then estimating the power spectrum of each locally stationary segment. We generate a piecewise autoregressive time series by concatenating three AR processes of segment length $n = 1000$:

$$x_t = \begin{cases} 0.9x_{t-1} - 0.9x_{t-2} + 0.9x_{t-3} - 0.9x_{t-4} + \epsilon_t, & 1 \leq t \leq 1000; \\ 0.8x_{t-1} - 0.6x_{t-2} + 0.9x_{t-3} - 0.8x_{t-4} + \epsilon_t, & 1001 \leq t \leq 2000; \\ 0.7x_{t-1} - 0.15x_{t-2} + 0.1x_{t-3} - 0.6x_{t-4} + \epsilon_t, & 2001 \leq t \leq 3000. \end{cases} \tag{25}$$

A realisation of this process is displayed in Fig. 2a. We remark that this is a more complex synthetic time series than that simulated by Rosen et al. (Figure 1, Rosen et al. 2012) in which they concatenate AR(1) and AR(2) processes. Indeed, an important part of the improvement in our methodology is the fact that their method suffers from limitations when the underlying spectra are more complex than that of AR(1)

or AR(2) processes, while our method has the flexibility to estimate the complex spectra of more involved processes and real-world data.

As before, we plot the Whittle likelihood against the number of eigenvectors in Fig. 2e—this selects the optimal number of basis functions for the optimal smoothing spline. We remark that the optimal number of basis functions is computed by optimising over the entire time-varying spectral surface. In this experiment, we also plot the penalised marginal likelihood function defined by (16) in Fig. 2f. These two figures clearly demonstrate the need for optimising the number of eigenvectors. Indeed, the 10-eigenvector decomposition provides the lowest likelihood of the data, while 61 eigenvectors are optimal and 29 are optimal in the penalised framework. Figure 2b–d shows that the optimal and penalised optimal smoothing spline models provide improved performance over each segment relative to the benchmark model of Rosen et al.

First, we report the posterior distributions of the number and locations of the change points. The three-segment model ($m = 3$) is selected with probability 1. The first change point (corresponding to the synthetic break at $t = 1000$) is estimated to have the following distribution: $t = 1032$ with probability 0.39, $t = 1035$ with probability 0.53, $t = 1070$ with probability 0.08. The second change point is estimated to be $t = 2002$ with probability 1.

In Fig. 2b, the optimal smoothing spline more effectively captures abrupt peaks in the PSD. We remark that the slight overfitting exhibited by the optimal spline would not lead to erroneous inference, as this overfitting only occurs at frequency components with limited power. In Fig. 2c, the optimal spline does a superior job at estimating both peaks in amplitude, while avoiding overfitting at frequency components that exhibit the most power. In Fig. 2d, we see two peaks in the PSD with significantly varying amplitude. There, the optimal model more appropriately determines the amplitude and frequency of the greatest peaks, just like our validation metrics in Table 1 recorded.

Figure 2g and h displays the two time-varying PSD estimated using the original 10-eigenvector decomposition and our optimal spline, respectively. Figure 2h clearly generates a more intricate surface than Fig. 2g, which is inappropriately

Table 3 Results for experiment on synthetic nonstationary time series described in Eq. (25)

Piecewise autoregressive model					
Method	RMSE	MAE	Wasserstein	$ \arg\max g - \arg\max \hat{g} $	$ \max g - \max \hat{g} $
Spline ₁₀	2.81	2.04	1.65	0.21	6.81
Spline _{optimal}	1.32	1.01	0.58	0.21	5.21
Spline _{penalised optimal}	1.33	0.96	0.70	0.20	6.22

Each validation metric is obtained by averaging the validation metrics on each segment, comparing the estimated log power spectra to the analytic known spectra. The optimal number of eigenvalues is 61, the penalised optimal is 29




Fig. 3 Estimate for the posterior distribution for a slowly varying autoregressive process defined in (26). The modal number of change points is 1, and conditional on this, the location of the change point is approximately uniform, subject to minimum bounds on each segment length

smooth. In particular, the scale of the z -axis demonstrates the optimised model’s improvement when estimating appropriate peaks at respective frequencies in the PSD.

Finally, we include an alternative synthetic experiment, where an autoregressive process slowly varies over time, rather than exhibiting clear synthetic breaks. We consider a $T = 500$ length time series examined by Rosen et al, defined by

$$x_t = (-0.5 + \frac{t}{T})x_{t-1} + \epsilon_t \tag{26}$$

In this instance, the optimal and penalised optimal number of eigenvalues is $M = 10$, as expected for a relatively simple (time-varying) AR(1) process. The modal number of segments is $m = 2$. Conditional on two segments (that is, one change point), the estimated posterior distribution of this change point is displayed in Fig. 3. As expected for a process with no clear change point, the posterior distribution is approximately uniform, subject to the constraints of a minimal possible segment length.

5 Real data

In this section, we analyse the spectra of two canonical datasets, sunspots (Choudhuri et al. 2004) and airline passenger data (Box et al. 2015). Both time series are well studied, with the former being particularly well studied in the spectral density estimation literature. We compare the performance of the optimal smoothing spline model and the 10-eigenvector smoothing spline. Unlike AR processes, real data do not possess analytic power spectra to validate estimates against. Both

time series require transformations to ensure they are stationary. For the sunspots data, following Choudhuri et al. (2004), we perform the following transformation:

$$S_t = y_t^{\frac{1}{2}} - \mu(\{y_t^{\frac{1}{2}} : t = 1, \dots, n\}).$$

That is, we take the square root of the observations and then mean-centre the resulting values. For the airline passenger data, we perform an original transformation based on first differences of fourth roots:

$$A_t = y_t^{\frac{1}{4}} - y_{t-1}^{\frac{1}{4}}.$$

The two transformed time series are displayed in Fig. 4a and b. As before, the number of basis functions is optimised to maximise the marginal Whittle likelihood. For the sunspots and airline passenger data, respectively, the optimal number of basis functions is 24 and 49, seen in Fig. 4c and d, respectively. The modal number of change points in both instances is 1, which is fitting as the transformed time series are approximately stationary. Thus, applying the stationary or nonstationary model from Sect. 2 yields identical results. The priors we use are detailed in that section.

Figure 4e and f demonstrates that our improved method models these complex processes substantially better than the existing 10-eigenvector model of Rosen et al. (2012), and that the latter does not have the complexity to model such processes. In Fig. 4e, the optimal smoothing spline captures the abruptness and amplitude of the one significant peak, as well as most of the movement of the log periodogram, while removing some of its characteristic noise. The 10-eigenvector spline fails to capture these features, with only slight recognition of that early peak. In Fig. 4f, the optimal smoothing spline captures all five peaks of this spectrum and most of its complex undulating behaviour; the 10-eigenvector spline fails to capture the prominent peaks in the data or provide any meaningful or accurate inference of the spectrum.

6 Conclusion

This paper has proposed and demonstrated the use of an improved algorithm for spectral density estimation of stationary and nonstationary time series. In addition to performing better quantitatively with respect to new and existing metrics, we have shown significantly better performance at appropriately matching the abruptness and complexity present in the PSD of real-world processes. The existing 10-eigenvector decomposition of Rosen et al. (2012) appropriately models smooth processes like the AR(1) and AR(2), but substantially underfits the real data we have presented; our optimal smoothing spline has a greater capability to fit both smooth and abruptly changing spectra.

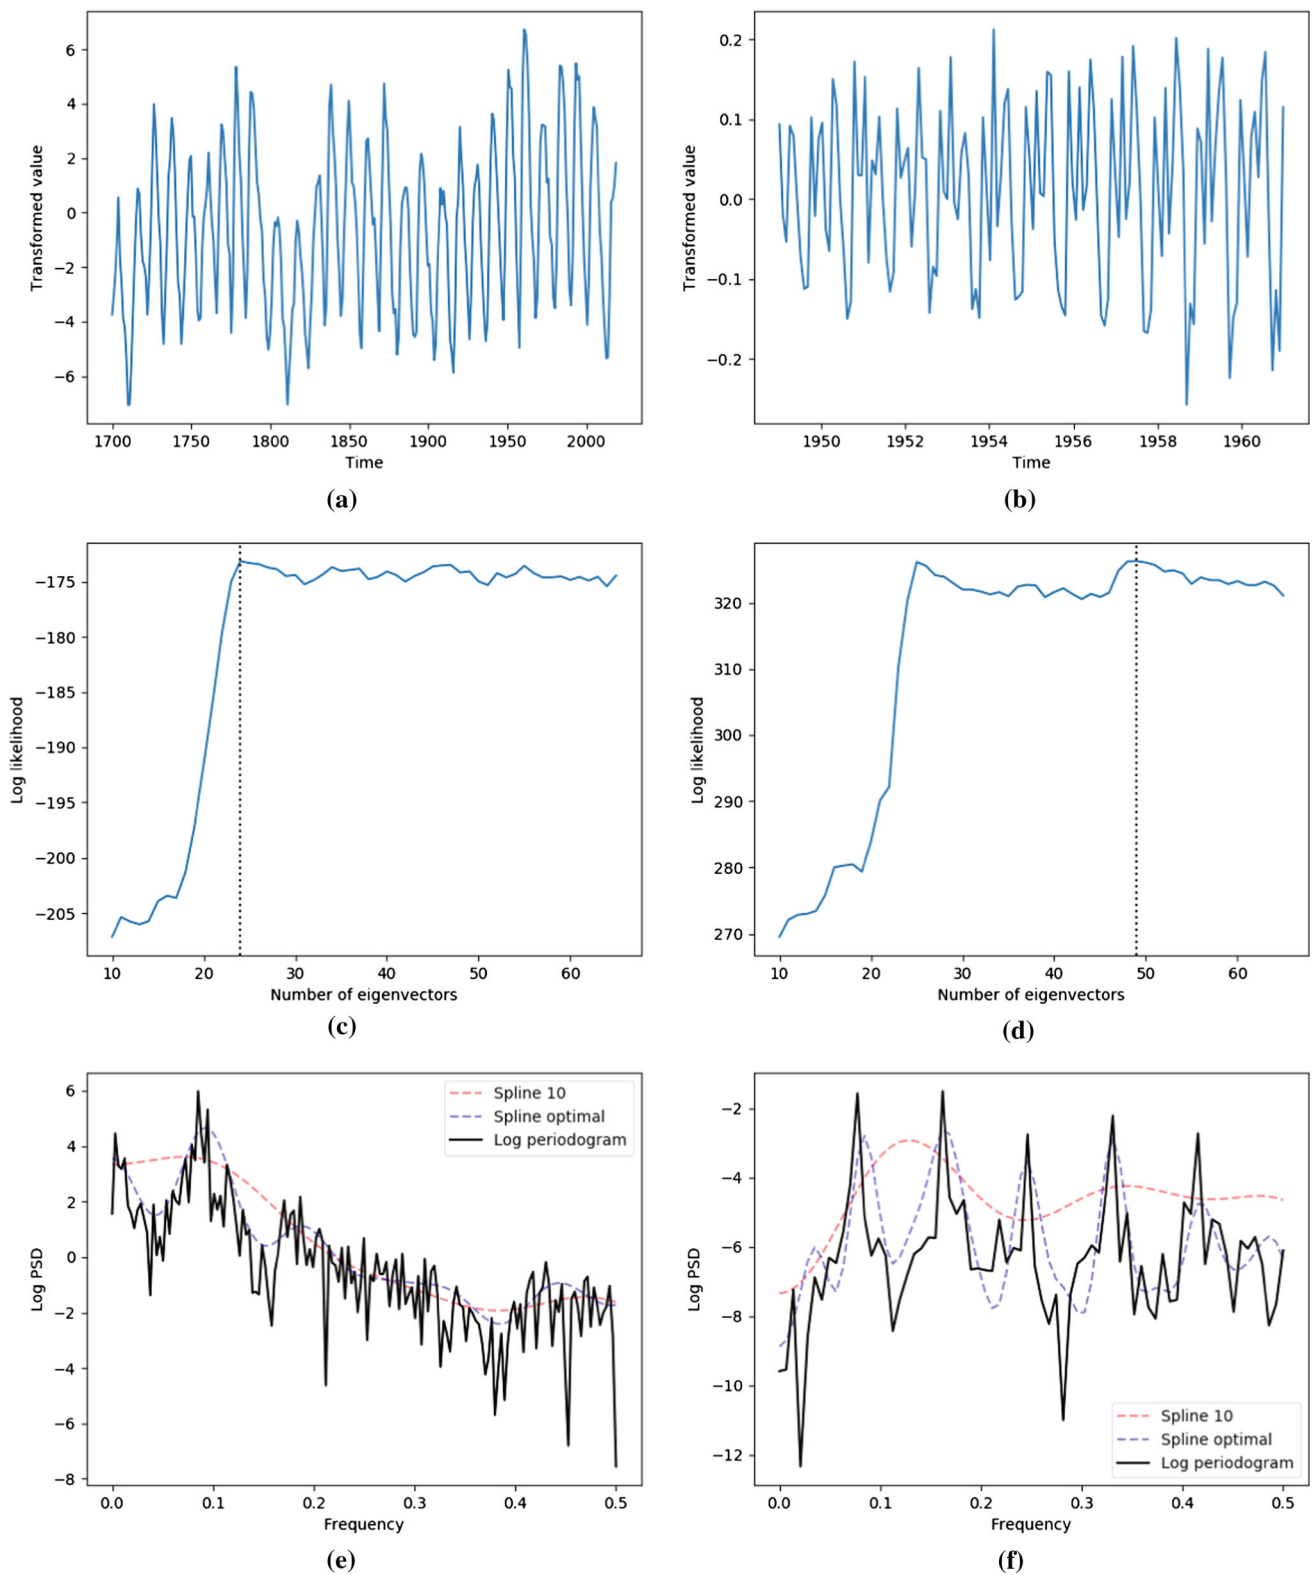


Fig. 4 Transformed time series for **a** sunspots, **b** airline passenger data. Log likelihood vs number of eigenvectors for **c** sunspots, **d** airline passenger data. Estimated log PSD for 10-eigenvector and optimal decomposition smoothing spline models for **e** sunspots, **f** airline passenger data

As further justification of the utility of the optimised eigendecomposition, we examine alternative covariance functions other than the smoothing spline and compare these against both the optimal and existing spline model in several settings. Our simulation study confirms that most well-known covariance functions perform acceptably well when analysing the power spectral density for stationary time series. Generally, smoother covariance functions such as the squared exponential and Matern perform better in the cases of smoother and simpler spectra, while models that combine stationary and nonstationary covariance functions such as the SQE+Sigmoid appear to generalise best across more complex spectral densities. The 10-eigenvector decomposition of the smoothing spline, as used within (Rosen et al. 2012), is the worst performing model. In particular, we demonstrate that eigenvector optimisation is what constitutes improvement in this case, not alternative covariance functions. Simply put, given that the eigendecomposition is so critical, it really should be learned from the data.

Our paper is not without limitations: first, while this paper improves upon the existing work of Rosen et al. (2012), there is no change to the underlying transdimensional sampling. Changing the transdimensional part of the model may lead to numerous further complications. If we were to make the algorithm locally adaptive in different segments of the time series, this could lead to an explosion in the number of model parameters and hence the computational cost of the procedure. In addition, allowing different numbers of eigenvalues in each segment could complicate the priors and likelihood computations that are involved in determining the segments themselves.

There are a variety of interesting directions for future research. First, future work could explore computationally feasible ways of creating a locally adaptive framework for PSD estimation in nonstationary time series, despite the challenges described above. That is, one could conceivably partition a locally stationary time series while at the same time optimising the number of basis functions used in the MCMC scheme within each segment, rather than over the entire time series period, as we have done. Future modelling frameworks could consider Bayesian and frequentist methods for identifying and learning the most appropriate covariance functions, or combinations thereof, to use within each segment of the time series. This could be performed within a Bayesian or non-Bayesian framework. We could also combine the eigendecomposition with the other covariance functions we have explored, including a learned choice of best covariance function through model averaging. Finally, our new validation metrics focused on collective sets of peaks, including the proximity matching criterion and algorithms for the appropriate refinement of peaks, could be employed in other contexts. These could be used together to measure a model's effectiveness at specifically identi-

fying a set of multiple peaks arising from more complex data.

Acknowledgements Many thanks to Lamiae Azizi, Sally Cripps and Alex Judge for helpful discussions.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Appendix A: Percival–Walden AR process

In this brief section, we apply our methodology to a rather challenging autoregressive process that has been highlighted several times in the literature (Box et al. 2015; Percival and Walden 1993) and is commonly known as the Percival–Walden AR(4). This process is defined as $x_t = 2.7607x_{t-1} - 3.8106x_{t-2} + 2.6535x_{t-3} - 0.9238x_{t-4} + \epsilon_t$, simulated over length $n = 1024$. As in Sect. 4, we simulate the process and validate our spectral density estimates against the known analytic power spectrum. In this experiment, the optimal and penalised optimal smoothing spline coincide, with 23 eigenvectors. The spectral estimates are plotted in Fig. 5 while validation metrics are provided in Table 4.

This experiment also provides an example where the proximity matching criterion of Sect. 3.2 fails. We can simply observe that the analytic log power spectrum \mathbf{g} has two peaks, while the spectral estimate $\hat{\mathbf{g}}$ for both Spline₁₀ and the (penalised) optimal smoothing spline have one peak each.

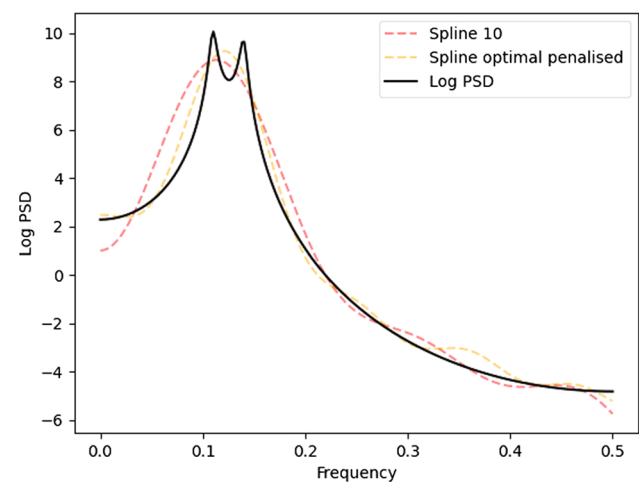


Fig. 5 Analytic and estimated log PSD for 10-eigenvector and optimal decomposition of the smoothing spline model for the Percival–Walden AR(4), defined in Appendix A. The optimal and penalised optimal smoothing splines coincide, with 23 eigenvectors. One can see that the spectral estimates each only have one spectral peak, while the analytic PSD has two, providing an example where the proximity matching criterion of Sect. 3.2 fails

Table 4 Results for synthetic experiments on the Percival–Walden AR(4), defined in Appendix A. The optimal and penalised optimal splines coincide, with 23 eigenvectors. In this instance, the proxim-

ity matching criterion from Sect. 3.2 fails, so the distance between sets of peaks in Eq. (24) must be used. We include this distance both between sets of frequencies and amplitudes

Percival–Walden process: validation metrics							
Method	RMSE	MAE	Wasserstein	$ \operatorname{argmax} g - \operatorname{argmax} \hat{g} $	$ \max g - \max \hat{g} $	Frequency distance (24)	Amplitude distance (24)
Spline ₁₀	0.91	0.61	0.51	0.002	1.17	0.0083	0.8575
Spline _{optimal}	0.50	0.36	0.31	0.01	0.81	0.013	0.136

As such, Table 4 includes the values of the semi-metric presented in (24) in Sect. 3.3. We observe that the (penalised) optimal smoothing spline provides a better approximation of the amplitudes of the two peaks than the existing method of Rosen et al.

Appendix B: Discussion of select existing methods

Alongside the statistics community, many signal processing practitioners and engineers have long been interested in the study of time series’ power spectra. Thus, it is worth noting the differences between a framework such as ours and frequentist or signal processing-based methods for power spectral density estimation.

We begin by describing Welch’s method, which is based upon Barlett’s method. Welch’s method aims to reduce noise in the resulting power spectral density estimate by sacrificing the degree of frequency resolution. The data are subdivided into overlapping segments, in each of which, a modified periodogram is computed. Each modified periodogram is averaged to produce a final estimate of the power spectral density. There are two model parameters in Welch’s method: the length of each segment and the degree of overlap in data points between adjacent segments.

Practically, Welch’s method has several limitations in comparison to Bayesian methods. First, the estimated power spectral density when using methods such as these may be less smooth (though not uniformly so). For scientists hoping to make observations with respect to the maximum amplitude and corresponding frequencies for any underlying time series, the often rough nature of Welch’s method may make inference more difficult. It is common in the Bayesian statistics literature to use a flexible prior function on the log of the power spectrum, such as a Gaussian process. The smoothness of the Gaussian process may be highly dependent on the covariance structure chosen by the modeller. Applying covariance functions such as the squared exponential and select Matern family variants allows for smooth interpolation in the resulting power spectral density estimate. Furthermore, recent research has shown that the variance is not a monotonic

decreasing function of the fraction of overlap within adjacent segments (Barbe et al. 2010). Second, Welch’s method is unable to algorithmically partition the time series based on changes in the power spectral density. Procedures such as the RJMCMC introduced in this paper identify points in time where the power spectral density has changed. Hence, one would be unable to determine locations in the time domain which correspond to changes in the underlying periodic nature of a process, if one were to use Welch’s method.

That said, many practitioners in the signal processing literature use techniques such as wavelets in the case of implementing spectral density estimation in a nonstationary setting. For instance, the continuous wavelet transform has been applied for spectral analysis in nonstationary signals. Wavelets overcome the obvious limitation with Fourier transformation-driven methods, where abrupt changes in time series’ behaviour is difficult to capture (due to its underlying construction as a sum of sinusoidal waves). Unlike sine waves, which smoothly oscillate, wavelets are derived from “step functions” that exist for a finite duration, allowing for the efficient capture of abrupt changes in modelling tasks.

Third, many would argue that a Bayesian framework such as ours provides a more principled approach to uncertainty quantification than frameworks such as Welch’s method. The methodology proposed in this paper consists of uncertainty surrounding the power spectral density estimate, in addition to uncertainty surrounding the change point location. One clear advantage of Welch’s method in comparison with the method we have proposed (and other MCMC-based methods), however, is a significant computational advantage. While there are certainly frequentist methods to estimate the uncertainty in traditional signal processing methods, there are always individuals who prefer the posterior distributions provided by Bayesian methods, including just from a psychological perspective, including the ability to make probabilistic statements about unknown parameters (Wasserman 2004).

Another commonly used framework for spectral density estimation is the multitaper method. Multitaper analysis is an extensions of traditional taper analysis, where time series are tapered before applying a Fourier transformation as a method of reducing potential bias resulting from spectral

leakage. The multitaper method averages over a variety of estimators with varying window functions (Thomson 1982; Mann and Lees 1996). This results in a power spectrum that exhibits reduced leakage and variance and retains important information from the initial and final sequences from the underlying time series. One major advantage of the multitaper method is that it can be applied in a fairly automatic manner, and is therefore appropriate in situations where many individual time series must be processed and a thorough analysis of each individual time series is not feasible. One possible limitation of the multitaper method is reduced spectral resolution. The multitaper method has proved to be an effective estimator in the presence of complex spectra. For example, Percival and Walden (1993) highlight the estimator’s effectiveness in detecting two peaks in the case of their AR(4) process described in Appendix A. As we saw, our methodology was unable to detect the two peaks. Of course, there are many techniques currently being used in addition to Welch’s method and the multitaper method described above. The choice between frequentist and Bayesian methods may depend on the precise problem and even the philosophical outlook of the practitioner. The literature is enriched by a robust continual development of both approaches.

Appendix C: Reversible jump sampling scheme

We follow Rosen et al. (2017, 2012) in our core implementation of the reversible jump sampling scheme. We remark that our method does not improve the trans-dimensional component of the model, described by the reversible jump scheme below. A time series partition is denoted $\xi_m = (\xi_{0,m}, \dots, \xi_{m,m})$ with m segments. We have a vector of *amplitude parameters* $\tau_m^2 = (\tau_{1,m}^2, \dots, \tau_{m,m}^2)'$ and *regression coefficients* $\beta_m = (\beta'_{1,m}, \dots, \beta'_{m,m})$ that we wish to estimate, for the j th component within a partition of m segments, $j = 1, \dots, m$. For notational simplicity, $\beta_{j,m}$, $j = 1, \dots, m$, is assumed to include the first entry, $\alpha_{0j,m}$. In the proceeding sections, superscripts c and p refer to current and proposed value in the sampling scheme.

First, we describe the **between-model moves**: let $\theta_m = (\xi'_m, \tau_m^2, \beta'_m)$ be the model parameters at some point in the sampling scheme and assume that the chain starts at $(m^c, \theta_{m^c}^c)$. The algorithm proposes the move to $(m^p, \theta_{m^p}^p)$, by drawing $(m^p, \theta_{m^p}^p)$ from the proposal distribution $q(m^p, \theta_{m^p}^p | m^c, \theta_{m^c}^c)$. That draw is accepted with probability

$$\alpha = \min \left\{ 1, \frac{p(m^p, \theta_{m^p}^p | \mathbf{x})q(m^c, \theta_{m^c}^c | m^p, \theta_{m^p}^p)}{p(m^c, \theta_{m^c}^c | \mathbf{x})q(m^p, \theta_{m^p}^p | m^c, \theta_{m^c}^c)} \right\},$$

with $p(\cdot)$ referring to a target distribution, the product of the likelihood and the prior. The target and proposal distributions will vary based on the type of move taken in the sampling scheme. First, $q(m^p, \theta_{m^p}^p | m^c, \theta_{m^c}^c)$ is described as follows:

$$\begin{aligned} q(m^p, \theta_{m^p}^p | m^c, \theta_{m^c}^c) &= q(m^p | m^c)q(\theta_{m^p}^p | m^p, m^c, \theta_{m^c}^c) \\ &= q(m^p | m^c)q(\xi_{m^p}^p, \tau_{m^p}^{2p}, \beta_{m^p}^p | m^p, m^c, \theta_{m^c}^c) \\ &= q(m^p | m^c)q(\xi_{m^p}^p | m^p, m^c, \theta_{m^c}^c)q(\tau_{m^p}^{2p} | \xi_{m^p}^p, m^p, m^c, \theta_{m^c}^c) \\ &\quad \times q(\beta_{m^p}^p | \tau_{m^p}^{2p}, \xi_{m^p}^p, m^p, m^c, \theta_{m^c}^c). \end{aligned}$$

To draw $(m^p, \theta_{m^p}^p)$ one must first draw m^p , followed by $\xi_{m^p}^p$, $\tau_{m^p}^{2p}$, and $\beta_{m^p}^p$. First, the number of segments m^p is drawn from the proposal distribution $q(m^p | m^c)$. Let M be the maximum number of segments and $m_{2,\min}^c$ be the number of current segments whose cardinality is at least $2t_{\min}$ data points. The proposal is as follows:

$$q(m^p = k | m^c) = \begin{cases} 1/2 & \text{if } k = m^c - 1, m^c + 1 \text{ and } m^c \neq 1, M, m_{2,\min}^c \neq 0 \\ 1 & \text{if } k = m^c - 1 \text{ and } m^c = M \text{ or } m_{2,\min}^c = 0 \\ 1 & \text{if } k = m^c + 1 \text{ and } m^c = 1 \end{cases}$$

Conditional on the proposed model m^p , a new partition $\xi_{m^p}^p$, a new vector of covariance amplitude parameters $\tau_{m^p}^{2p}$ and a new vector of regression coefficients, $\beta_{m^p}^p$ are proposed. In Rosen et al. (2012), τ^2 is referred to as a smoothing parameter. To impact the smoothness of the covariance function, the parameter would have to impact pairwise operations. Given that τ^2 sits outside the covariance matrix, we will refer to τ^2 as an amplitude parameter (akin to signal variance within the Gaussian process framework Rasmussen and Williams 2005).

Now, we describe the process of the **birth** of new segments. Suppose that $m^p = m^c + 1$. A time series partition,

$$\xi_{m^p}^p = (\xi_{0,m^c}^c, \dots, \xi_{k^*-1,m^c}^c, \xi_{k^*,m^p}^p, \xi_{k^*,m^c}^c, \dots, \xi_{m^c,m^c}^c)$$

is drawn from the proposal distribution $q(\xi_{m^p}^p | m^p, m^c, \theta_{m^c}^c)$. The algorithm proposes a partition by first selecting a random segment $j = k^*$ to split. Then, a point t^* within the segment $j = k^*$ is randomly selected to be the proposed partition point. This is subject to the constraint,

$$\xi_{k^*-1,m^c}^c + t_{\min} \leq t^* \leq \xi_{k^*,m^c}^c - t_{\min}. \text{ The proposal distribution is computed as follows:}$$

$$\begin{aligned} q(\xi_{j,m^p}^p = t^* | m^p, m^c, \xi_{m^c}^c) &= p(j = k^* | m^p, m^c, \xi_{m^c}^c) \\ p(\xi_{k^*,m^p}^p = t^* | j = k^*, m^p, m^c, \xi_{m^c}^c) &= \frac{1}{m_{2,\min}^c (n_{k^*,m^c} - 2t_{\min} + 1)}. \end{aligned}$$

The vector of amplitude parameters

$$\begin{aligned} \tau_{m^p}^{2p} &= (\tau_{1,m^c}^{2c}, \dots, \tau_{k^*-1,m^c}^{2c}, \\ \tau_{k^*,m^p}^{2p}, \tau_{k^*+1,m^p}^{2p}, \tau_{k^*+1,m^c}^{2c}, \dots, \tau_{m^c,m^c}^{2c}) \end{aligned}$$

is drawn from the proposal distribution

$q(\tau_{m^p}^{2p} | m^p, \xi_{m^p}^p, m^c, \theta_{m^c}^c) = q(\tau_{m^p}^{2p} | m^p, \tau_{m^c}^{2c})$. The algorithm is based on the reversible jump algorithm of Green (1995). It draws from a uniform distribution $u \sim U[0, 1]$ and defines τ_{k^*,m^p}^{2p} and τ_{k^*+1,m^p}^{2p} in terms of u and τ_{k^*,m^c}^{2c} as follows:

$$\tau_{k^*,m^p}^{2p} = \frac{u}{1-u} \tau_{k^*,m^c}^{2c}; \tag{27}$$

$$\tau_{k^*+1,m^p}^{2p} = \frac{1-u}{u} \tau_{k^*,m^c}^{2c}. \tag{28}$$

The vector of coefficients

$$\begin{aligned} \beta_{m^p}^p &= (\beta_{1,m^c}^c, \dots, \beta_{k^*-1,m^c}^c, \\ \beta_{k^*,m^p}^p, \beta_{k^*+1,m^p}^p, \beta_{k^*+1,m^c}^c, \dots, \beta_{m^c,m^c}^c) \end{aligned}$$

is drawn from the proposal distribution

$q(\beta_{m^p}^p | \tau_{m^p}^{2p}, \xi_{m^p}^p, m^p, m^c, \theta_{m^c}^c) = q(\beta_{m^p}^p | \tau_{m^p}^{2p}, \xi_{m^p}^p, m^p)$. The pair of vectors β_{k^*,m^p}^p and β_{k^*+1,m^p}^p are drawn from Gaussian approximations to the respective posterior conditional distributions $p(\beta_{k^*,m^p}^p | \mathbf{x}_{k^*}^p, \tau_{k^*,m^p}^{2p}, m^p)$ and $p(\beta_{k^*+1,m^p}^p | \mathbf{x}_{k^*+1}^p, \tau_{k^*+1,m^p}^{2p}, m^p)$, respectively. Here, $\mathbf{x}_{k^*}^p$ and $\mathbf{x}_{k^*+1}^p$ refer to the subsets of the time series with respective segments k^* and $k^* + 1$. $\xi_{m^p}^p$ will determine $\mathbf{x}_{*}^p = (\mathbf{x}_{k^*}^p, \mathbf{x}_{k^*+1}^p)'$. For the sake of exposition, we provide the following example: the coefficient β_{k^*,m^p}^p is drawn from the Gaussian distribution $N(\beta_{k^*}^{\max}, \Sigma_{k^*}^{\max})$, where $\beta_{k^*}^{\max}$ is defined as

$$\operatorname{argmax}_{\beta_{k^*,m^p}^p} p(\beta_{k^*,m^p}^p | \mathbf{x}_{k^*}^p, \tau_{k^*,m^p}^{2p}, m^p)$$

and

$$\Sigma_{k^*}^{\max} = - \left\{ \frac{\partial^2 \log p(\beta_{k^*,m^p}^p | \mathbf{x}_{k^*}^p, \tau_{k^*,m^p}^{2p}, m^p)}{\partial \beta_{k^*,m^p}^p \partial \beta_{k^*,m^p}^p} \Big|_{\beta_{k^*,m^p}^p = \beta_{k^*}^{\max}} \right\}^{-1}$$

For the birth move, the probability of acceptance is $\alpha = \min\{1, A\}$, where A is equal to

$$\begin{aligned} & \left| \frac{\partial(\tau_{k^*,m^p}^{2p}, \tau_{k^*+1,m^p}^{2p})(\tau_{k^*,m^c}^{2c}, u)}{\partial(\tau_{k^*,m^p}^{2p}, \tau_{k^*+1,m^p}^{2p})} \right| \frac{p(\theta_{m^p}^p | \mathbf{x}, m^p) p(\theta_{m^p}^p | m^p) p(m^p)}{p(\theta_{m^p}^p | \mathbf{x}, m^p) p(\theta_{m^c}^c | m^c) p(m^c)} \\ & \times \frac{p(m^c | m^p) p(\beta_{k^*,m^c}^c)}{p(m^p | m^c) p(\xi_{k^*,m^p}^p | m^p, m^c) p(u) p(\beta_{k^*,m^p}^p) p(\beta_{k^*+1,m^p}^p)}. \end{aligned}$$

Above, $p(u) = 1, 0 \leq u \leq 1$, while $p(\beta_{k^*,m^p}^p)$ and $p(\beta_{k^*+1,m^p}^p)$ are Gaussian proposal distributions $N(\beta_{k^*}^{\max}, \Sigma_{k^*}^{\max})$ and $N(\beta_{k^*+1}^{\max}, \Sigma_{k^*+1}^{\max})$, respectively. The Jacobian is computed as

$$\begin{aligned} & \left| \partial(\tau_{k^*,m^p}^{2p}, \tau_{k^*+1,m^p}^{2p})(\tau_{k^*,m^c}^{2c}, u) \right| \\ & = \frac{2\tau_{k^*,m^c}^{2c}}{u(1-u)} = 2(\tau_{k^*,m^p}^p + \tau_{k^*+1,m^p}^p)^2. \end{aligned}$$

Next, we describe the process of the **death** of new segments, that is, the reverse of a birth move, where $m^p = m^c - 1$. A time series partition

$$\xi_{m^p}^p = (\xi_{0,m^c}^c, \dots, \xi_{k^*-1,m^c}^c, \xi_{k^*+1,m^c}^c, \dots, \xi_{m^c,m^c}^c),$$

is proposed by randomly selecting a single partition from $m^c - 1$ candidates, and removing it. The partition point selected for removal is denoted $j = k^*$. There are $m^c - 1$ possible segments available for removal among the m^c segments currently in existence. The proposal may choose each partition point with equal probability, that is,

$$q(\xi_{j,m^p}^p | m^p, m^c, \xi_{m^c}^c) = \frac{1}{m^c - 1}.$$

The vector of amplitude parameters

$$\tau_{m^p}^{2p} = (\tau_{1,m^c}^{2c}, \dots, \tau_{k^*-1,m^c}^{2c}, \tau_{k^*,m^p}^{2c}, \tau_{k^*+2,m^c}^{2c}, \dots, \tau_{m^c,m^c}^{2c})$$

is drawn from the proposal distribution

$q(\tau_{m^p}^{2p} | m^p, \xi_{m^p}^p, m^c, \theta_{m^c}^c) = q(\tau_{m^p}^{2p} | m^p, \tau_{m^c}^{2c})$. One amplitude parameter τ_{k^*,m^p}^{2p} is formed from two candidate amplitude parameters, τ_{k^*,m^c}^{2c} and τ_{k^*+1,m^c}^{2c} . This is done by reversing the equations 27 and 28. That is,

$$\tau_{k^*,m^p}^{2p} = \sqrt{\tau_{k^*,m^c}^{2c} \tau_{k^*+1,m^c}^{2c}}.$$

Finally, the vector of regression coefficients,

$$\beta_{m^p}^p = (\beta_{1,m^c}^c, \dots, \beta_{k^*-1,m^c}^c, \beta_{k^*,m^p}^p, \beta_{k^*+2,m^c}^c, \dots, \beta_{m^c,m^c}^c)$$

is drawn from the proposal distribution

$q(\beta_{m^p}^p | \tau_{m^p}^{2p}, \xi_{m^p}^p, m^p, m^c, \theta_{m^c}^c) = q(\beta_{m^p}^p | \tau_{m^p}^{2p}, \xi_{m^p}^p, m^p)$. The vector of regression coefficients is drawn from a Gaussian approximation to the posterior distribution $p(\beta_{k^*,m^p}^p | \mathbf{x}, \tau_{k^*,m^p}^{2p}, \xi_{m^p}^p, m^p)$ following the same procedure for the vector of coefficients in the birth step. The probability of acceptance is the inverse of the analogous birth step. If the move is accepted then the following updates occur: $m^c = m^p$ and $\theta_{m^c}^c = \theta_{m^p}^p$.

Finally, we describe the **within-model moves**: henceforth, m is fixed; accordingly, notation describing the dependence on the number of segments is removed. There are two parts to a within-model move. First, a segment relocation is performed, and conditional on the relocation, the basis function coefficients are updated. The steps are jointly accepted or rejected with a Metropolis–Hastings step. The amplitude parameters are updated within a separate Gibbs sampling step.

The chain is assumed to be located at $\theta^c = (\xi^c, \beta^c)$. The proposed move $\theta^p = (\xi^p, \beta^p)$ is as follows: first, a partition point ξ_{k^*} is selected for relocation from $m - 1$ candidate partition points. Next, a position within the interval $[\xi_{k^*-1}, \xi_{k^*+1}]$ is selected, subject to the fact that the new location is at least t_{\min} data points away from ξ_{k^*-1} and ξ_{k^*+1} , so that

$$\Pr(\xi_{k^*}^p = t) = \Pr(j = k^*) \Pr(\xi_{k^*}^p = t | j = k^*),$$

where $\Pr(j = k^*) = (m - 1)^{-1}$. A mixture distribution for $\Pr(\xi_{k^*}^p = t | j = k^*)$ is constructed to explore the space most efficiently, so

$$\Pr(\xi_{k^*}^p = t | j = k^*) = \pi q_1(\xi_{k^*}^p = t | \xi_{k^*}^c) + (1 - \pi) q_2(\xi_{k^*}^p = t | \xi_{k^*}^c),$$

where $q_1(\xi_{k^*}^p = t | \xi_{k^*}^c) = (n_{k^*} + n_{k^*+1} - 2t_{\min} + 1)^{-1}$, $\xi_{k^*-1} + t_{\min} \leq t \leq \xi_{k^*+1} - t_{\min}$ and

$$q_2(\xi_{k^*}^p = t | \xi_{k^*}^c) = \begin{cases} 0 & \text{if } |t - \xi_{k^*}^c| > 1 \\ 1/3 & \text{if } |t - \xi_{k^*}^c| \leq 1, n_{k^*} \neq t_{\min} \text{ and } n_{k^*+1} \neq t_{\min} \\ 1/2 & \text{if } t - \xi_{k^*}^c \leq 1, n_{k^*} = t_{\min} \text{ and } n_{k^*+1} \neq t_{\min} \\ 1/2 & \text{if } \xi_{k^*}^c - t \leq 1, n_{k^*} \neq t_{\min} \text{ and } n_{k^*+1} = t_{\min} \\ 1 & \text{if } t = \xi_{k^*}^c, n_{k^*} = t_{\min} \text{ and } n_{k^*+1} = t_{\min} \end{cases}$$

The support of q_1 has $n_{k^*} + n_{k^*+1} - 2t_{\min} + 1$ data points while q_2 has at most three. The term q_2 alone would result in a high acceptance rate for the Metropolis–Hastings, but it would explore the parameter space slowly. The q_1 component allows for larger jumps, and produces a compromise between a high acceptance rate and thorough exploration of the parameter space.

Next, $\beta_j^p, j = k^*, k^* + 1$ is drawn from an approximation to $\prod_{j=k^*}^{k^*+1} p(\beta_j | \mathbf{x}_j^p, \tau_j^2)$, following the analogous step in the between-model move. The proposal distribution, which is evaluated at $\beta_j^p, j = k^*, k^* + 1$, is

$$q(\beta_*^p | \mathbf{x}_*^p, \tau_*^2) = \prod_{j=k^*}^{k^*+1} q(\beta_j^p | \mathbf{x}_j^p, \tau_j^2),$$

where $\beta_*^p = (\beta_{k^*}^p, \beta_{k^*+1}^p)'$ and $\tau_*^2 = (\tau_{k^*}^2, \tau_{k^*+1}^2)'$. The proposal distribution is evaluated at current values of β_*^c =

$(\beta_{k^*}^c, \beta_{k^*+1}^c)'$. β_*^p is accepted with probability

$$\alpha = \min \left\{ 1, \frac{p(\mathbf{x}_*^p | \beta_*^p) p(\beta_*^p | \tau_*^2) q(\beta_*^c | \mathbf{x}_*^c, \tau_*^2)}{p(\mathbf{x}_*^c | \beta_*^c) p(\beta_*^c | \tau_*^2) q(\beta_*^p | \mathbf{x}_*^p, \tau_*^2)} \right\},$$

where $\mathbf{x}_*^c = (\mathbf{x}_{k^*}^c, \mathbf{x}_{k^*+1}^c)$. When the draw is accepted, update the partition and regression coefficients $(\xi_{k^*}^c, \beta_*^c) = (\xi_{k^*}^p, \beta_*^p)$. Finally, draw τ^{2p} from

$$p(\tau_*^2 | \beta_*) = \prod_{j=k^*}^{k^*+1} p(\tau_j^2 | \beta_j).$$

This is a Gibbs sampling step, and accordingly the draw is accepted with probability 1.

Appendix D: Metropolis–Hastings algorithm

In this section, we describe the Metropolis–Hastings algorithm used in the stationary case for our simulation study (Sect. 4). As seen in Appendix C, the above RJMCMC reduces to a Metropolis–Hastings in the absence of the between-model moves.

We estimate the log of the spectral density by its posterior mean via a Bayesian approach and an adaptive MCMC algorithm:

$$\mathbb{E}(\mathbf{g} | \mathbf{y}) = \int \mathbb{E}(\mathbf{g} | \mathbf{y}, \theta) p(\theta | \mathbf{y}) d\theta \simeq \frac{1}{M} \sum_{j=1}^M \mathbb{E}(\hat{\mathbf{g}} | \mathbf{y}, \hat{\theta}^j). \tag{29}$$

Here, M is the number of post-burn-in iterations in the MCMC scheme; $\hat{\theta}^j$ are samples taken from the posterior distribution $p(\theta | \mathbf{y})$; $p(\cdot)$ is taken from a Gaussian distribution $N(\mu, \sigma^2)$ centred around $\mu = \theta^{[c]}$ that maximises the log marginal likelihood; σ is chosen arbitrarily; and $\hat{\mathbf{g}}$ is the forecast log spectrum.

Monte Carlo algorithms have been highly prominent for estimation of hyperparameters and spectral density in a non-parametric Bayesian framework. Metropolis et al. (1953) first proposed the Metropolis algorithm; this was generalised by Hastings in a more focused, statistical context (Hastings 1970). The random walk Metropolis–Hastings algorithm aims to sample from a target density π , given some candidate transition probability function $q(x, y)$. In our context, π represents the Whittle likelihood function multiplied by respective priors. The acceptance ratio is:

$$\alpha(x, y) = \begin{cases} \min \left(\frac{\pi(y)q(y,x)}{\pi(x)q(x,y)}, 1 \right) & \text{if } \pi(x)q(x, y) > 0 \\ 1 & \text{if } \pi(x)q(x, y) = 0. \end{cases} \tag{30}$$

Our MCMC scheme calculates the acceptance ratio every 50 iterations; based on an ideal acceptance ratio (Roberts and Rosenthal 2004), the step size is adjusted for each hyperparameter. The power spectrum and hyperparameters for the GP covariance function are calculated in each iteration of our sampling scheme and integrated over.

First, we initialise the values of our GP hyperparameters, $\theta^{[c]}$, our log PSD \hat{g}^c , our random perturbation size s^2 , and the adaptive adjustment for our step size ξ . Starting values for our GP hyperparameters are chosen based on maximising the marginal Whittle likelihood,

$$\theta^{[c]} = \operatorname{argmax}_{\theta} -(2\pi)^{-m/2} \prod_{j=0}^{m-1} \exp\left(-\frac{1}{2} \left[\log f(v_j) + \frac{I(v_j)}{f(v_j)} \right]\right). \tag{31}$$

The latent log PSD is modelled with a zero mean GP. That is, $\mathbf{g} \sim GP(0, k_{\theta}(x, x'))$, and we follow the notation of Rasmussen and Williams (2005) where $k(x, x')$ refers to any respective kernel. The adaptive MCMC algorithm samples from the posterior distribution of the log PSD and the posterior distribution of any candidate covariance function’s hyperparameters. First, the current and proposed values for the mean and covariance of the Gaussian process are computed. That is, the mean is computed:

$$\hat{g}^c = k_{\theta^c}(x', x)[k_{\theta^c}(x, x) + \sigma^2 I]^{-1} \log I^{(\circ)}, \tag{32}$$

and the covariance is computed:

$$\hat{V}^c = k_{\theta^c}(x', x') - k_{\theta^c}(x', x)[k_{\theta^c}(x, x) + \sigma^2 I]^{-1} k_{\theta^c}(x, x'). \tag{33}$$

New proposals for GP hyperparameters are determined via a random walk proposal distribution. A zero-mean Gaussian distribution is used to generate candidate perturbations, where s^2 is the variance of this Gaussian. That is

$$\theta^p \leftarrow q(\theta^p | \theta^c). \tag{34}$$

Having drawn the proposed GP hyperparameters, a proposed mean and covariance function of the log PSD are drawn from the posterior distribution of the GP. Both the proposed mean and covariance are computed similarly to the current values, simply replacing the values of the hyperparameters $\theta^c \leftarrow \theta^p$. So, the proposed mean of the log PSD is,

$$\hat{g}^p = k_{\theta^p}(x', x)[k_{\theta^p}(x, x) + \sigma^2 I]^{-1} \log I^{(\circ)} \tag{35}$$

and the proposed covariance is computed as follows

$$\hat{V}^p = k_{\theta^p}(x', x') - k_{\theta^p}(x', x)[k_{\theta^p}(x, x) + \sigma^2 I]^{-1} k_{\theta^p}(x, x'). \tag{36}$$

Having computed the proposed and current values of the log PSD, we update the current log PSD based on the Metropolis–Hastings transition kernel. First, we sample from a uniform distribution $u \sim U(0, 1)$ and compute our acceptance ratio,

$$\alpha = \min\left(1, \frac{p(\log I^{(\circ)} | \theta^p, \hat{g}^p) p(\hat{g}^p) q(\hat{g}^c | \log I^{(\circ)}, \theta^p)}{p(\log I^{(\circ)} | \theta^c, \hat{g}^c) p(\hat{g}^c) q(\hat{g}^p | \log I^{(\circ)}, \theta^c)}\right). \tag{37}$$

$p(\log I^{(\circ)} | \theta^p, \hat{g}^p)$ is our Whittle likelihood computation, the probability of the log PSD conditional on hyperparameters θ and our candidate estimate of the latent log PSD \hat{g} . $p(\hat{g}^p)$ represents the prior distribution on our latent log PSD and $q(\hat{g}^c | \log I^{(\circ)}, \theta^p)$ is our proposal distribution, representing the probability of the estimated log PSD conditional on the log periodogram and GP hyperparameters.

Should $u < \alpha$, we update the current values of the log PSD mean and spectrum to the proposed values. That is,

$$\hat{g}^{c+1} \leftarrow \hat{g}^p \tag{38}$$

$$\hat{V}^{c+1} \leftarrow \hat{V}^p. \tag{39}$$

If $u > \alpha$, both the mean and variance of the log PSD are kept at their current values,

$$\hat{g}^{c+1} \leftarrow \hat{g}^c \tag{40}$$

$$\hat{V}^{c+1} \leftarrow \hat{V}^c. \tag{41}$$

Importantly, modelling the log PSD with a GP prior does not mean that we are assuming a Gaussian error distribution around the spectrum. In actuality, proposed spectra are accepted and rejected through a Metropolis–Hastings procedure—resulting in log PSD samples being drawn from the true posterior distribution of the log PSD, $\log(\exp(1))$. Having sampled the log PSD, we then accept/reject candidate GP hyperparameters with another Metropolis–Hastings step. Our acceptance ratio is,

$$\alpha = \min\left(1, \frac{p(\log I^{(\circ)} | \theta^p) p(\theta^c | \theta^p)}{p(\log I^{(\circ)} | \theta^c) p(\theta^p | \theta^c)}\right), \tag{42}$$

where $p(\log I^{(\circ)} | \theta)$ represents the Whittle likelihood modelling the probability of the log PSD, $\log I^{(\circ)}$, conditional on hyperparameters θ . $p(\theta^p | \theta)$ is the prior distribution we place over GP hyperparameters, θ . Note that in this particular case, the symmetric proposal distributions cancel out and our

algorithm reduces simply to a Metropolis ratio. Again we follow the standard Metropolis–Hastings acceptance decision. If $u < \alpha$,

$$\theta^{c+1} \leftarrow \theta^p, \tag{43}$$

and the current hyperparameter values assume proposed values. Alternatively, if $u > \alpha$,

$$\theta^{c+1} \leftarrow \theta^c, \tag{44}$$

current value of the hyperparameters are not updated. Finally, following (Roberts and Rosenthal 2004) we implement an adaptive step-size within our random walk proposal. Every 50 iterations within our simulation, we compute the trailing acceptance ratio. An optimal acceptance ratio, Acc^{Opt} of 0.234 is targeted. If the acceptance ratio is too low, indicating that the step size may be too large, then the step size is systematically reduced. If $\text{Acceptance Ratio}_{(j-49):j} \forall j \in \{50, 100, 150, \dots, 10000\} < \text{Acc}^{Opt}$,

$$s^2 \leftarrow s^2 - \xi. \tag{45}$$

If the acceptance ratio is too high, indicating that the step size may be too small then step size is systematically increased. That is, when $\text{Acceptance Ratio}_{(j-49):j} \forall j \in \{50, 100, \dots, 10000\} > \text{Acc}^{Opt}$,

$$s^2 \leftarrow s^2 + \xi. \tag{46}$$

Finally, the log PSD and the respective analytic uncertainty bounds are determined by computing the median of the samples generated from the sampling procedure,

$$\mathcal{U}_{0.025}^{\text{final}} = \text{median}(\mathcal{U}_{0.025}^{5000:10000}) \tag{47}$$

$$\hat{g} = \text{median}(\hat{g}^{5000:10000}) \tag{48}$$

$$\mathcal{U}_{0.975}^{\text{final}} = \text{median}(\mathcal{U}_{0.975}^{5000:10000}). \tag{49}$$

Appendix E: Turning point algorithm

In this section, we provide more details for the identification of non-trivial peaks (local maxima). We aim to outline a broad and flexible framework for this purpose, in which the exact procedure may be altered according to the specific application. For example, one way to determine peaks of a given spectral estimate is simply by inspection. We aim to provide an algorithmic framework as an alternative to this.

Let g be an analytic or estimated log power spectral density function. We may begin, if necessary, by applying additional smoothing to this function (though this step is optional and can be omitted). Following (James et al. 2022), we apply a

two-step algorithm to the (possibly smoothed) function g , defined on $v_j = \frac{j}{n}$, $j = 0, 1, \dots, m - 1$. The first step produces an alternating sequence of local minima (troughs) and local maxima (peaks), which may include some immaterial turning points. The second step refines this sequence according to chosen conditions and parameters. The most important conditions to initially identify a peak or trough, respectively, are the following:

$$g(v_{j_0}) = \max\{g(v_j) : \max(0, j_0 - l) \leq j \leq \min(j_0 + l, m - 1)\}, \tag{50}$$

$$g(v_{j_0}) = \min\{g(v_j) : \max(0, j_0 - l) \leq j \leq \min(j_0 + l, m - 1)\}, \tag{51}$$

where l is a parameter to be chosen. Defining peaks and troughs according to this definition alone has some flaws, such as the potential for two consecutive peaks.

Instead, we implement an inductive procedure to choose an alternating sequence of peaks and troughs. Suppose j_0 is the last determined peak. We search in the period $j > j_0$ for the first of two cases: if we find a time $j_1 > j_0$ that satisfies (51) as well as a non-triviality condition $g(j_1) < g(j_0)$, we add j_1 to the set of troughs and proceed from there. If we find a time $j_1 > j_0$ that satisfies (50) and $g(j_0) \geq g(j_1)$, we ignore this lower peak as redundant; if we find a time $j_1 > j_0$ that satisfies (50) and $g(j_1) > g(j_0)$, we remove the peak j_0 , replace it with j_1 and continue from j_1 . A similar process applies from a trough at j_0 .

As a side remark, for an analytic log PSD g , we could simply use the analytical and differentiable form to find critical points as an alternative.

With either possibility, at this point, the function is assigned an alternating sequence of troughs and peaks. However, some turning points are immaterial and should be removed. Here, the framework can incorporate a flexible series of options to refine the set of peaks.

As mentioned in Sect. 3.2, one relatively simple option is simply to remove any local maximum (peak) $\hat{\rho}$ of \hat{g} with $\hat{g}(\hat{\rho}) < \max \hat{g} - \delta$, for some sensible constant δ . In our experiments, the same results are produced for any $\delta \in [2, 4]$, demonstrating the robustness of this relatively simple idea. Under an affine transformation of the original time series $X'_t = aX_t + b$, the log PSD \hat{g} changes by an additive constant, so this condition is unchanged when rescaling the original data.

This relatively simple condition is sufficient for our application. For the benefit of future work, we list some alternative options for an algorithmic refinement of the peaks (besides, of course, inspection as the simplest option). In previous work, we have analysed functions $v(t)$ that were necessarily valued only in non-negative reals. Thus, one may simply linearly the log PSD g so that its minimum value is zero. Then, numer-

ous options exist for refinement of non-trivial peaks (and troughs).

For example, let $t_1 < t_3$ be two peaks, necessarily separated by a trough. We select a parameter $\delta = 0.2$, and if the *peak ratio*, defined as $\frac{v(t_3)}{v(t_1)} < \delta$, we remove the peak t_3 . If two consecutive troughs t_2, t_4 remain, we remove t_2 if $v(t_2) > v(t_4)$, otherwise remove t_4 . That is, if the second peak has size less than δ of the first peak, we remove it. Alternatively, one may use this peak ratio on any peak, comparing it to the global max, rather than just comparing adjacent peaks. That is, let t_0 be the global maximum. Then, one could remove any peak t_1 with $\frac{v(t_1)}{v(t_0)} < \delta$.

Alternatively, we use appropriately defined gradient or log-gradient comparisons between points $t_1 < t_2$. For example, let

$$\text{log-grad}(t_1, t_2) = \frac{\log v(t_2) - \log v(t_1)}{t_2 - t_1}. \tag{52}$$

The numerator equals $\log(\frac{v(t_2)}{v(t_1)})$, a “logarithmic rate of change”. Unlike the standard rate of change given by $\frac{v(t_2)}{v(t_1)} - 1$, the logarithmic change is symmetrically between $(-\infty, \infty)$. Let t_1, t_2 be adjacent turning points (one a trough, one a peak). We choose a parameter ϵ ; if

$$|\text{log-grad}(t_1, t_2)| < \epsilon, \tag{53}$$

that is, the average logarithmic change is less than ϵ , we remove t_2 from our sets of peaks and troughs. If t_2 is not the final turning point, we also remove t_1 . After these refinement steps, we are left with an alternating sequence of non-trivial peaks and troughs. Finally, for this framework, we only need the peaks, so we simply discard the troughs.

As a final remark, only at the end is the final number r of non-trivial peaks determined. It is a function not only of the log PSD function \mathbf{g} , but also the precise conditions used to select and refine the (non-trivial) peaks.

References

Adak, S.: Time-dependent spectral analysis of nonstationary time series. *J. Am. Stat. Assoc.* **93**(444), 1488–1501 (1998). <https://doi.org/10.1080/01621459.1998.10473808>

Barbe, K., Pintelon, R., Schoukens, J.: Welch method revisited: non-parametric power spectrum estimation via circular overlap. *IEEE Trans. Signal Process.* **58**(2), 553–565 (2010). <https://doi.org/10.1109/tsp.2009.2031724>

Box, G.E.P., Jenkins, G.M., Reinsel, G.C., Ljung, G.M.: *Time Series Analysis: Forecasting and Control*. Wiley, Hoboken (2015)

Brockwell, P.J., Davis, R.A.: *Time Series: Theory and Methods*. Springer, New York (1991). <https://doi.org/10.1007/978-1-4419-0320-4>

Carter, C.K., Kohn, R.: Semiparametric Bayesian inference for time series with mixed spectra. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* **59**(1), 255–268 (1997). <https://doi.org/10.1111/1467-9868.00067>

Choudhuri, N., Ghosal, S., Roy, A.: Bayesian estimation of the spectral density of a time series. *J. Am. Stat. Assoc.* **99**(468), 1050–1059 (2004). <https://doi.org/10.1198/01621450400000557>

Cogburn, R., Davis, H.T.: Periodic splines and spectral estimation. *Ann. Stat.* **2**(6), 1108–1126 (1974). <https://doi.org/10.1214/aos/1176342868>

Dahlhaus, R.: Fitting time series models to nonstationary processes. *Ann. Stat.* **25**(1), 1–37 (1997). <https://doi.org/10.1214/aos/1034276620>

Duvenaud, D., Lloyd, J., Grosse, R., Tenenbaum, J., Ghahramani, Z.: Structure discovery in nonparametric regression through compositional kernel search. In: *Proceedings of the 30th International Conference on Machine Learning*, vol. 28, pp. 1166–1174 (2013)

Edwards, M.C., Meyer, R., Christensen, N.: Bayesian nonparametric spectral density estimation using B-spline priors. *Stat. Comput.* **29**(1), 67–78 (2019). <https://doi.org/10.1007/s11222-017-9796-9>

Eilers, P.H.C., Marx, B.D.: Flexible smoothing with B-splines and penalties. *Stat. Sci.* **11**(2), 89–121 (1996). <https://doi.org/10.1214/ss/1038425655>

Gangopadhyay, A., Mallick, B., Denison, D.: Estimation of spectral density of a stationary time series via an asymptotic representation of the periodogram. *J. Stat. Plan. Inference* **75**(2), 281–290 (1999). [https://doi.org/10.1016/s0378-3758\(98\)00148-7](https://doi.org/10.1016/s0378-3758(98)00148-7)

Green, P.J.: Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**(4), 711–732 (1995). <https://doi.org/10.1093/biomet/82.4.711>

Gu, C.: *Smoothing Spline ANOVA Models*. Springer, New York (2013). <https://doi.org/10.1007/978-1-4614-5369-7>

Guo, W., Dai, M., Ombao, H.C., von Sachs, R.: Smoothing spline ANOVA for time-dependent spectral analysis. *J. Am. Stat. Assoc.* **98**(463), 643–652 (2003). <https://doi.org/10.1198/01621450300000549>

Hadj-Amar, B., Rand, B.F., Fiecas, M., Lévi, F., Huckstepp, R.: Bayesian model search for nonstationary periodic time series. *J. Am. Stat. Assoc.* **115**(531), 1320–1335 (2019). <https://doi.org/10.1080/01621459.2019.1623043>

Hastings, W.K.: Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**(1), 97–109 (1970). <https://doi.org/10.1093/biomet/57.1.97>

James, N., Menzies, M.: A new measure between sets of probability distributions with applications to erratic financial behavior. *J. Stat. Mech: Theory Exp.* **2021**(12), 123404 (2021). <https://doi.org/10.1088/1742-5468/ac3d91>

James, N., Menzies, M.: Collective correlations, dynamics, and behavioural inconsistencies of the cryptocurrency market over time. *Nonlinear Dyn.* (2022). <https://doi.org/10.1007/s11071-021-07166-9>

James, N., Menzies, M., Azizi, L., Chan, J.: Novel semi-metrics for multivariate change point analysis and anomaly detection. *Physica D* **412**, 132636 (2020). <https://doi.org/10.1016/j.physd.2020.132636>

James, N., Menzies, M., Bondell, H.: Comparing the dynamics of COVID-19 infection and mortality in the United States, India, and Brazil. *Physica D* **432**, 133158 (2022). <https://doi.org/10.1016/j.physd.2022.133158>

Lu, J., Hoi, S.C., Wang, J., Zhao, P., Liu, Z.Y.: Large scale online kernel learning. *J. Mach. Learn. Res.* **17**(47), 1–43 (2016)

Mann, M.E., Lees, J.M.: Robust estimation of background noise and signal detection in climatic time series. *Clim. Change* **33**(3), 409–445 (1996). <https://doi.org/10.1007/bf00142586>

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**(6), 1087–1092 (1953). <https://doi.org/10.1063/1.1699114>

Paciorek, C.J., Schervish, M.J.: Nonstationary covariance functions for Gaussian process regression. In: *Proceedings of the 16th Interna-*

- tional Conference on Neural Information Processing Systems, pp. 273–280. MIT Press (2003)
- Percival, D.B., Walden, A.T.: Spectral Analysis for Physical Applications. Cambridge University Press, Cambridge (1993). <https://doi.org/10.1017/cbo9780511622762>
- Plagemann, C., Kersting, K., Burgard, W.: Nonstationary Gaussian Process regression using point estimates of local smoothness. In: Machine Learning and Knowledge Discovery in Databases, pp. 204–219. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87481-2_14
- Prakash, A., James, N., Menzies, M., Francis, G.: Structural clustering of volatility regimes for dynamic trading strategies. *Appl. Math. Finance* **28**(3), 236–274 (2021). <https://doi.org/10.1080/1350486x.2021.2007146>
- Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. MIT Press, Cambridge (2005)
- Roberts, G.O., Rosenthal, J.S.: General state space Markov chains and MCMC algorithms. *Probab. Surv.* **1**, 20–71 (2004). <https://doi.org/10.1214/154957804100000024>
- Rosen, O., Stoffer, D.S., Wood, S.: Local spectral analysis via a Bayesian mixture of smoothing splines. *J. Am. Stat. Assoc.* **104**(485), 249–262 (2009). <https://doi.org/10.1198/jasa.2009.0118>
- Rosen, O., Wood, S., Stoffer, D.: BayesSpec: Bayesian Spectral Analysis Techniques (2017). <https://CRAN.R-project.org/package=BayesSpec>. R package version 0.5.3
- Rosen, O., Wood, S., Stoffer, D.S.: AdaptSPEC: Adaptive spectral estimation for nonstationary time series. *J. Am. Stat. Assoc.* **107**(500), 1575–1589 (2012). <https://doi.org/10.1080/01621459.2012.716340>
- Thomson, D.: Spectrum estimation and harmonic analysis. *Proc. IEEE* **70**, 1055–1096 (1982)
- Todd, J.F.: Recommendations for nomenclature and symbolism for mass spectroscopy. *Int. J. Mass Spectrom. Ion Processes* **142**(3), 209–240 (1995). [https://doi.org/10.1016/0168-1176\(95\)93811-f](https://doi.org/10.1016/0168-1176(95)93811-f)
- Todd, J.F.J.: Recommendations for nomenclature and symbolism for mass spectroscopy (including an appendix of terms used in vacuum technology). (recommendations 1991). *Pure Appl. Chem.* **63**(10), 1541–1566 (1991). <https://doi.org/10.1351/pac199163101541>
- Wahba, G.: Automatic smoothing of the log periodogram. *J. Am. Stat. Assoc.* **75**(369), 122–132 (1980). <https://doi.org/10.1080/01621459.1980.10477441>
- Wahba, G.: Spline Models for Observational Data. Society for Industrial and Applied Mathematics (1990). <https://doi.org/10.1137/1.9781611970128>
- Wasserman, L.: All of Statistics: A Concise Course in Statistical Inference. Springer, New York (2004)
- Whittle, P.: On stationary processes in the plane. *Biometrika* **41**(3–4), 434–449 (1954). <https://doi.org/10.1093/biomet/41.3-4.434>
- Whittle, P.: Curve and periodogram smoothing. *J. R. Stat. Soc.: Ser. B (Methodol.)* **19**(1), 38–47 (1957). <https://doi.org/10.1111/j.2517-6161.1957.tb00242.x>
- Wilson, A.G., Adams, R.P.: Gaussian process kernels for pattern discovery and extrapolation. In: Proceedings of the 30th International Conference on Machine Learning, vol. 28, pp. 1067–1075 (2013)
- Wood, S., Rosen, O., Kohn, R.: Bayesian mixtures of autoregressive models. *J. Comput. Graph. Stat.* **20**(1), 174–195 (2011). <https://doi.org/10.1198/jcgs.2010.09174>
- Wood, S.A., Jian, W., Tanner, M.: Bayesian mixture of splines for spatially adaptive nonparametric regression. *Biometrika* **89**(3), 513–528 (2002). <https://doi.org/10.1093/biomet/89.3.513>
- Wood, S.N.: P-splines with derivative based penalties and tensor product smoothing of unevenly distributed data. *Stat. Comput.* **27**(4), 985–989 (2017). <https://doi.org/10.1007/s11222-016-9666-x>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.