Submitted to the Annals of Applied Statistics

FREQUENCY SELECTION IN BAYESIAN SPECTRAL MODELING OF TIME SERIES DATA WITH APPLICATIONS TO WEARABLE DEVICE MEASUREMENTS

BY BENIAMINO HADJ-AMAR[1,a], VAISHNAV KRISHNAN[2,b], AND MARINA VANNUCCI[3,c]

[1]*Department of Epidemiology and Biostatistics, University of South Carolina, Columbia, SC,*[a]*hadjamar@mailbox.sc.edu*

[2]*Departments of Neurology, Neuroscience, and Psychiatry & Behavioral Sciences, Baylor College of Medicine, Houston, TX,*[b]*vaishnav.krishnan@bcm.edu*

[3]*Department of Statistics, Rice University, Houston, TX,*[c]*marina@rice.edu*

This paper introduces a Bayesian spike-and-slab framework for spectral analysis of time series data. The proposed method combines frequency selection and dimensionality reduction with a refined grid of candidate frequencies, enabling high-resolution recovery of oscillatory components while promoting sparsity through a structured spike-and-slab prior. A stochastic search algorithm efficiently explores the posterior space, yielding posterior inclusion probabilities that quantify the relevance of each frequency. We extend the framework to multivariate signals via a hierarchical prior on frequency inclusion patterns, allowing the model to capture both shared and component-specific rhythms across multiple time series. Extensive simulation studies demonstrate the method's robustness and superior performance in frequency estimation and spectral power reconstruction compared to existing approaches. Applied to actigraphy data from individuals with partial-onset seizures, the univariate model identifies clinically relevant circadian and ultradian rhythms. In a second application, for the joint analysis of physical activity and skin temperature from a healthy individual, the multivariate model reveals partially overlapping rhythmic components consistent with known physiological coupling. This work establishes a powerful and interpretable approach to spectral analysis, with broad applicability to wearable data, chronobiology, and personalized health monitoring.

1. Introduction. Understanding the periodic patterns in cyclical phenomena provides critical insights into the rhythmic, and potentially predictable drivers of variability. Assessments of periodicity in a variety of measurable biological signals can enhance patient care by forecasting physiological events (Grant, Newman and Kriegsfeld, 2020) and/or adverse events (Friedrichs-Maeder et al., 2024; Vieluf et al., 2024). Among a plethora of such biological signals that can be derived from wearable devices, continuously collected accelerometry at the wrist ("wrist actigraphy") has received the greatest emphasis. Actigraphy offers a unique, non-invasive view into an individual's rest-activity cycles by capturing metrics related to sleep duration, timing, and daytime movement levels. In medical research, this data can reveal patterns that align with changes in seizure frequency and the prevalence of psychiatric symptoms, such as depression or anxiety (Lau et al., 2024) and also correlate with aspects of neurological disability, such as seizure frequency and side effect burden (Abboud et al., 2023). Increasingly, actigraphy is being complemented by additional wearable-derived signals—such as skin temperature, heart rate, or electrodermal activity—that provide orthogonal yet physiologically related information. When analyzed jointly, these modalities can reveal partially shared rhythms and enhance the detection of underlying chronobiological patterns, motivating the need for multivariate modeling frameworks.

Keywords and phrases: Frequency selection, Fourier, Multivariate Data, Rest-Activity Rhythms, Spike-and-Slab, Stochastic search MCMC.

Despite the promise of actigraphy, a challenge remains in accurately and automatically identifying the specific frequencies, amplitudes, and phases within activity data. These frequency components, acting as signatures within the data, may offer a novel biomarker for monitoring physiological conditions, potentially linking variations in rest-activity rhythms to neurological states. Current research thus demands a methodology capable of extracting these distinct periodicities directly from actigraphy time series, enabling more precise, data-driven interpretations of rest-activity dynamics. Moreover, multivariate signals—such as physical activity and skin temperature—can exhibit overlapping but distinct rhythms, motivating the need for joint modeling frameworks that capture shared and modality-specific dynamics. Robust, data-adaptive techniques for frequency selection are crucial to advance our ability to use actigraphy as a tool for continuous, real-world physiological monitoring in clinical settings.

Spectral analysis has long been a cornerstone for understanding stationary time series, with the periodogram being one of the earliest tools. Developed as the squared magnitude of the Fourier transform, the periodogram provides a raw estimate of the power spectral density but often requires smoothing as the estimate is not consistent (Brockwell and Davis, 1991; Shumway, Stoffer and Stoffer, 2000). Smoothing techniques, such as the Welch method and multi-taper spectral estimation, improve stability by averaging across overlapping segments or using orthogonal tapers (Percival and Walden, 1993). Fourier decomposition remains central, where time series are expressed as a sum of sinusoidal components. This framework underpins classical techniques like the discrete Fourier transform (DFT) and its variants, which form the basis for harmonic analysis. Representing a time series as the sum of sine and cosine functions at different frequencies provides a detailed and flexible framework for signal representation . However, including all Fourier frequencies can lead to several drawbacks (Tseng, Shrikumar and Kundaje, 2020; Lange, Brunton and Kutz, 2021): (i) overfitting, where the model captures noise rather than meaningful periodic signals; (ii) computational scalability issues, as the number of coefficients grows with the length of the time series; and (iii) poor interpretability, as the model does not automatically identify the relevant frequencies that drive overall variation in the data. Hence, selecting the appropriate subset of frequencies is essential to balance model accuracy and computational efficiency.

Bayesian methods have been used to address challenges in spectral analysis by incorporating prior information and enabling probabilistic approaches to spectral estimation. For example, Bretthorst (1988) introduced a Bayesian framework for spectral analysis, applying it to univariate stationary processes through prior distributions on spectral densities. Choudhuri, Ghosal and Roy (2004) utilized Dirichlet process priors and Bernstein polynomials for flexible spectral density estimation. In the multivariate setting, Bayesian spectral modeling has been extended to capture shared and component-specific structure across multiple time series. For instance, Cadonna, Kottas and Prado (2019) proposed a hierarchical framework that models the log-spectral density functions of multiple signals, allowing for borrowing of strength across series while retaining individual spectral characteristics. More recently, Hu and Prado (2023) developed computationally efficient methods for multivariate spectral estimation using smoothing priors, enabling the identification of common spectral features across high-dimensional signals with uncertainty quantification. However, these methods do not directly address the challenge of automatically selecting the significant peaks in the spectrum, and hence the frequencies that drive the variation in the data. More recently, Granados-Garcia et al. (2022) proposed a Bayesian nonparametric mixture of autoregressive kernels that can data-adaptively identify both spectral peak locations and their associated bandwidths. In the frequentist literature, penalized techniques like LASSO and Elastic Net (Tibshirani, 1996; Zou and Hastie, 2005) may be applied in the context of a saturated Fourier regression, where all possible Fourier frequencies are initially included, to encourage sparsity and improve interpretability (Incremona and De Nicolao, 2019; Ha and Zhang, 2019). However, while these

methods can help simplify the model, they can also lead to overfitting by selecting more frequencies than necessary, making it harder to identify the true underlying periodic signals.

In this article, we propose a Bayesian spectral modeling approach that uses discrete spike-and-slab priors for simultaneous frequency selection and dimensionality reduction. The model uses a bivariate spike-and-slab prior over sinusoidal coefficients, enabling joint frequency selection across sine and cosine terms. A minimum-separation constraint excludes near-redundant components, while frequency-specific inclusion probabilities allow adaptation to heterogeneous spectral features. Our prior model is complemented by incorporating a refined grid of candidate frequencies, to improve the resolution of the model. Traditional methods relying on Fourier frequencies often lack the precision needed to capture true periodicities, especially when fine temporal details are critical. By introducing a finer grid of candidate frequencies, the model achieves more accurate recovery of periodic signals and enhanced spectral estimates. While this refinement increases the number of frequency candidates, it provides significant gains in precision and accuracy, as demonstrated in our simulation studies. Together, the spike-and-slab prior and the refined frequency grid enable a powerful and flexible framework for spectral analysis. To model multivariate signals, we extend the univariate prior by introducing a hierarchical structure on frequency inclusion indicators that enables flexible learning of shared and component-specific periodicities across time series. By placing a Dirichlet prior on inclusion patterns, the model borrows strength across signals while allowing asynchronous rhythms. For posterior inference, we implement a stochastic search algorithm that efficiently explores the space of frequency models by iteratively adding, deleting, or swapping frequencies. We further improve computational efficiency by focusing on the most relevant regions of the posterior space, guided by informed initialization based on dominant frequencies in the data. Our proposed spike-and-slab prior modeling framework allows to quantify posterior probabilities of inclusion for each frequency and the posterior distribution over the number of frequencies, offering insight into the model's complexity.

We evaluate the performance of our proposed approach through extensive simulation studies, demonstrating accuracy in identifying frequencies that drive variation in the data, robustness across diverse scenarios, and clear advantages over existing alternative methods. This includes both univariate and multivariate settings, with simulations illustrating the method's ability to recover shared and component-specific rhythms across multiple time series. We then present two real-data applications. In the first, we apply the model to univariate actigraphy data collected from individuals with partial-onset seizures. We show that the model effectively identifies distinct frequency components by leveraging sparsity while preserving key oscillatory features. Specifically, the approach enables the robust extraction of circadian and ultradian rhythms, accurately capturing both dominant and subtle periodicities in the data. In the second application, we analyze bivariate data from a healthy individual, jointly modeling physical activity and skin temperature. This multivariate analysis reveals both shared and modality-specific components, capturing well-known physiological coupling between behavioral and thermoregulatory rhythms. By leveraging a structured spike-and-slab prior, the model promotes sparsity while flexibly capturing multiscale spectral structure.

The remainder of the paper is organized as follows: Section 2 introduces the model and prior specification, Section 3 details the proposed MCMC algorithm and discusses posterior inference. Section 4 presents the simulation studies, and Section 5 illustrates the applications to measurements from wearable devices. Section 6 concludes the paper with a discussion.

2. Methods.

2.1. *Spectral Modeling.* Consider a time series realization $\boldsymbol{y} = (y_1, \ldots, y_T)$, with T denoting the length of the time series. For simplicity, we assume T is even. In spectral analysis,

the observed time series can be exactly expressed as a weighted sum of sines and cosines oscillating at the *Fourier frequencies* $\tilde{\omega}_j = j/T$, for $j = 1, \ldots, M_{max}$, with $M_{max} = (T/2 - 1)$ being the maximum number of frequencies included in the model. Hence, we can formulate a saturated regression model involving all Fourier coefficients as

$$y_t = \sum_{j=1}^{M_{\max}} \left\{ \beta_{j1} \cos(2\pi\tilde{\omega}_j t) + \beta_{j2} \sin(2\pi\tilde{\omega}_j t) \right\} + \varepsilon_t, \tag{2.1}$$

for $t = 1, \ldots, T$. Here, the linear coefficients β_{j1} and β_{j2} determine amplitude and phase of sine and cosine at frequency ω_j; the power can be summarized as $P(\omega_j) = \beta_{j1}^2 + \beta_{j2}^2$, which denotes the relative contribution of frequency ω_j to the overall variation in the data. Further, we assume normally distributed errors with zero mean and variance σ^2, i.e. $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$.

2.1.1. *Dense Grid of Candidate Frequencies.* Frequency domain analysis is inherently constrained by the duration of the input time series. According to the Fourier uncertainty principle, the resolution of frequency domain output improves with longer time series segments, as the number of available frequency bins increases with the temporal integration window (Percival and Walden, 1993). Typically, one would consider the set of Fourier frequencies, $\tilde{\omega}_j = j/T$, as potential candidates for inclusion in a model. However, this approach limits the precision with which one can recover true periodicities. For instance, in cases where fine temporal details are needed, relying solely on Fourier frequencies may result in estimates that are close but not precise. To address this problem, we consider a dense grid of candidate frequencies, $\bar{\omega}_j = j \cdot \Delta\omega$ for $j = 0, 1, 2, \ldots, L$, where $\Delta\omega$ controls the grid resolution, and L is chosen such that $\bar{\omega}_j \in (0, 0.5)$, since the contribution of each frequency is symmetric and periodic, so frequencies above 0.5 are aliased with those below and do not provide additional information. By decreasing $\Delta\omega$, we enhance the frequency resolution, enabling more precise recovery of true periodicities. In section 4 we show on simulated data how a denser grid not only improves the estimate of the location of the frequencies but also their estimated power.

This choice provides a denser discretization of the frequency domain, which can improve the ability to localize true periodicities compared to approaches restricted to the Fourier frequencies. As shown in Keil et al. (2022), finer grids improve resolution, particularly in cases where high temporal resolution competes with the need for fine frequency distinctions. We note that Wu and Zhou (2024) also advocate the use of dense frequency grids with sufficiently fine mesh size to achieve improved frequency localization and near-parametric estimation accuracy in complex or non-stationary time series. By improving the grid resolution, we mitigate aliasing effects, as highlighted by the Nyquist sampling theorem, and achieve a more precise balance between temporal and frequency resolution (Keil et al., 2022). However, it is important to note that as the grid becomes finer, the number of potential frequencies increases, leading to higher computational costs. Therefore, the more precise localization of oscillatory components comes with a trade-off in terms of computational efficiency.

2.2. *Frequency Selection through Spike-and-Slab Priors.* We propose an approach based on discrete spike-and-slab priors that allows for simultaneous frequency selection and dimensionality reduction. Spike-and-slab priors are commonly used in regression models to achieve variable selection with the goal of identifing the most relevant predictors from a set of available covariates (Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Brown, Vannucci and Fearn, 1998). These priors combine two components: a *spike* that concentrates probability mass at zero (indicating that a variable is not included in the model) and a *slab* that has its mass distributed over a broad range of possible values. In the regression context, these sparsity-inducing mixture priors are applied to the regression coefficients, allowing irrelevant

covariates to be excluded by assigning their coefficients a value of zero. Spike-and-slab priors offer several benefits, including the ability to facilitate complex modeling through MCMC-based stochastic search, provide accurate predictions via model averaging, adapt seamlessly to multivariate and nonlinear frameworks, and address high-dimensional problems where the number of covariates surpasses the available data. Additionally, they enable the integration of prior knowledge into the modeling process. See Tadesse and Vannucci (2021) for a comprehensive review of theoretical, methodological and computational aspects of these priors. In our modeling framework, the spike part of the prior corresponds to assuming that some frequencies may not be relevant in the analysis; that is, some coefficients in the spectral decomposition (Eq. 2.1) should be set to zeros. On the other hand, the slab part allows the inclusion of frequencies with non-zero coefficients, recognizing that certain frequencies may have a meaningful impact on the signal and should be included in the model.

Let z_j be the latent inclusion indicator variable for frequency ω_j, such that $z_j = 1$ if ω_j significantly contributes to driving the overall oscillatory behavior of the time series, and $z_j = 0$ otherwise. We assume a discrete spike-and-slab prior on the coefficients $\boldsymbol{\beta}_j = (\beta_{j1}, \beta_{j2})$ as

$$[\boldsymbol{\beta}_j \mid z_j] \sim z_j \cdot \mathcal{N}_2(\mathbf{0}, \sigma_\beta^2 \boldsymbol{I}_2) + (1 - z_j) \cdot \boldsymbol{\delta}_0(\boldsymbol{\beta}_j), \tag{2.2}$$

where $\boldsymbol{\delta}_0$ denotes the bivariate Dirac mass at zero, that is, $\boldsymbol{\delta}_0(\boldsymbol{\beta}_j) = 1$ if $\boldsymbol{\beta}_j = \mathbf{0}$ and 0 otherwise. The slab variance σ_β^2 is a hyperparameter controlling the prior dispersion of the active coefficients. In the applications of this paper, we fix σ_β^2 to a moderately large value, to avoid over-shrinkage. We enforce sparsity and improve interpretability by assuming a constrained joint prior on the inclusion vector $\mathbf{z} = (z_1, \ldots, z_{M_{\max}})$ defined as

$$p(\mathbf{z} \mid \boldsymbol{\alpha}) = \left[\prod_{j=1}^{M_{\max}} \alpha_j^{z_j} (1 - \alpha_j)^{1 - z_j} \right] \cdot \mathbb{I}\left(\min_{j<k:\, z_j = z_k = 1} |\omega_j - \omega_k| \geq d \right), \tag{2.3}$$

where $\mathbb{I}(\cdot)$ is the indicator function and d is a user-defined minimum frequency separation threshold. This constraint ensures that any pair of active frequencies in the model are separated by at least d, thereby reducing collinearity among basis elements and improving identifiability. Configurations violating this constraint receive zero prior probability and are excluded from the posterior support. A similar principle is discussed by Hadj-Amar et al. (2020). This minimum-separation rule reflects both modeling and computational considerations: nearby frequencies are often difficult to distinguish given finite sample sizes, and including them may lead to overfitting or unstable inference. By integrating the separation condition directly into the prior, we guarantee that all valid configurations reside on a well-behaved, interpretable, and identifiable subspace of the model space. This formulation leads to a parsimonious and structured representation that facilitates automatic frequency selection while preserving model coherence. See Section 2.2.1 below for further discussion on how to set the parameter d. Furthermore, we assume $\alpha_j \sim \text{Beta}(a, b)$ and integrate the α_j's out analytically, leading to a collapsed marginal prior on z_j given by $p(z_j = 1) = \frac{a}{a+b}$ and $p(z_j = 0) = \frac{b}{a+b}$, noting that the minimum-separation constraint remains in effect under the marginal prior, as configurations violating the threshold continue to receive zero probability. By allowing frequency-specific inclusion probabilities α_j we achieve greater flexibility than using a single global parameter α, as this allows the model to accommodate heterogeneity in the likelihood of inclusion across different regions of the frequency spectrum. We complete the prior model with an inverse-gamma prior on the residual variance, $\sigma^2 \sim \text{IG}(\gamma_0/2, \nu_0/2)$.

2.2.1. *Minimum Separation Constraint: Specification and Rationale.* As introduced in the prior in Eq. (2.3), the inclusion vector $\boldsymbol{z} = (z_1, \ldots, z_{M_{\max}})$ is restricted to configurations

in which any pair of included frequencies are separated by at least d bins on the refined frequency grid. We now discuss the rationale for setting this parameter and its role in the model. In our implementation, we fix the minimum separation to $d = 5$, which corresponds to a frequency gap of approximately five bins on the refined grid. This choice reflects both biological plausibility and statistical stability. In the context of our applications to wearable data, we aim to capture dominant oscillations such as the circadian rhythm (approximately 24 hours), ultradian rhythms (e.g., 12 hours), and infradian cycles. On a fine frequency grid—such as one constructed with resolution $1/N$—adjacent bins may correspond to very similar periods. For instance, frequencies near the circadian range might correspond to periods of 24.64, 24.30, 23.96, 23.63, and 23.31 hours. Including both 23.96 and 23.31 hours in the same model would be physiologically implausible, as they reflect nearly indistinguishable periodicities from a biological standpoint. Formally, the separation constraint enforces the condition:

$$\min_{j<k:\, z_j=z_k=1} |\omega_j - \omega_k| \geq d \cdot \Delta\omega,$$

where $\Delta\omega$ is the frequency bin width. This ensures that the selected frequency components $\{\omega_j : z_j = 1\}$ are well separated, reducing redundancy and improving identifiability of latent periodicities. In practice, this has two major benefits: (i) it avoids collinearity among the basis vectors associated with nearby frequencies, thereby improving the stability of posterior inference; and (ii) it concentrates posterior inclusion probabilities (PIPs) around dominant components. Without the constraint, the model may spread posterior mass across a cluster of collinear frequencies, making interpretation difficult and diluting the PIP of each.

Importantly, while d is a user-specified parameter, its effects are interpretable and predictable. Smaller values of d increase flexibility but may lead to overfitting or redundancy; larger values promote parsimony but may miss closely spaced true signals. In our application, we found $d = 5$ to provide a balanced compromise, yielding stable posterior summaries and interpretable rhythm estimates across a range of subjects. Note that as the frequency grid becomes more refined (i.e., with smaller $\Delta\omega$), a proportionally larger value of d helps maintain identifiability by preventing the inclusion of highly correlated or physiologically indistinguishable components. The impact of d in settings with closely spaced frequencies is further investigated in a simulation study reported in the Supplementary Material.

2.3. *Extension to Multivariate Data with Shared Spectral Structure.* The univariate setting introduced earlier models an individual time series as a sum of sinusoidal components, with spike-and-slab priors guiding frequency selection. In multivariate settings, where multiple related time series are observed, it is often reasonable to expect that some oscillatory components recur across datasets. Applying the univariate model to each time series separately would ignore these potential commonalities. We propose an extension to a joint model that retains component-specific flexibility for each time series while introducing a hierarchical prior to couple frequency inclusion across datasets.

Let $\mathbf{y}_t = (y_{t1}, \ldots, y_{tD})^\top \in \mathbb{R}^D$ denote the observed values of a D-dimensional time series at time $t = 1, \ldots, T$. For each component $i \in \{1, \ldots, D\}$, we model the time series as a linear combination of sinusoidal functions with component-specific coefficients:

$$y_{ti} = \sum_{j=1}^{M_{\max}} \left\{ \beta_{ji}^{(1)} \cos(2\pi\omega_j t) + \beta_{ji}^{(2)} \sin(2\pi\omega_j t) \right\} + \varepsilon_{ti}, \quad \varepsilon_{ti} \sim \mathcal{N}(0, \sigma_i^2),$$

where the set of candidate frequencies ω_j has already been introduced in the univariate formulation, and $\boldsymbol{\beta}_{ji} = (\beta_{ji}^{(1)}, \beta_{ji}^{(2)})^\top$ are the component-specific coefficients determining the contribution of each frequency to component i. Here σ_i^2 is the residual variance for the i-th

component. In order to extend the spike-and-slab prior to the multivariate setting, we introduce binary indicators $z_{ji} \in \{0,1\}$ for each frequency ω_j and component i, such that $\boldsymbol{\beta}_{ji}$ is drawn from a slab prior if $z_{ji} = 1$ and set to zero otherwise. This allows each component of the multivariate signal to selectively include or exclude any frequency. Rather than treating each z_{ji} independently, we define a joint inclusion pattern $\mathbf{z}_j = (z_{j1}, \ldots, z_{jD}) \in \{0,1\}^D$ across all components for each frequency ω_j, as described below.

2.3.1. *Structured Frequency Inclusion via Hierarchical Priors.* To introduce structured information sharing across components, we place a hierarchical prior on the joint inclusion pattern $\mathbf{z}_j$ for each frequency ω_j. Specifically, we assume that $\mathbf{z}_j$ is drawn from a categorical distribution with probability vector $\boldsymbol{\pi}$, which itself is given a Dirichlet prior, e.g., $\mathbf{z}_j \sim \text{Categorical}(\boldsymbol{\pi})$ and $\boldsymbol{\pi} \sim \text{Dirichlet}(\boldsymbol{\alpha})$. For example, in the bivariate case ($D = 2$), the vector $\boldsymbol{\pi} = (\pi_{00}, \pi_{10}, \pi_{01}, \pi_{11})$ encodes the prior probabilities of all possible inclusion configurations across the two components. This hierarchical formulation couples frequency selection across components, allowing the model to borrow strength and favor synchronized inclusion patterns when supported by the data, while still permitting component-specific selections when necessary. For instance, a high value of π_{11} increases the probability that the same frequency is selected in both components, reflecting potential physiological synchronization or common periodic regulation. On the other hand, higher values of π_{10} or π_{01} allow for frequency inclusion in only one component, thus preserving the capacity for component-specific rhythms. By inferring $\boldsymbol{\pi}$ from the data, the model can adaptively learn whether spectral patterns are shared or distinct across components. Furthermore, the Dirichlet prior on $\boldsymbol{\pi}$ introduces further modeling flexibility by enabling us to encode prior beliefs about the prevalence of different patterns. Specifically, the hyperparameters $\boldsymbol{\alpha} = (\alpha_{00}, \alpha_{10}, \alpha_{01}, \alpha_{11})$ control the prior weight placed on each configuration: larger values of α_{00} promote sparsity by favoring exclusion from both components, while higher values of α_{11} encourage joint inclusion. This allows the prior to reflect different expectations about frequency sharing based on domain knowledge or modeling goals.

As in the univariate setting, we impose the same minimum-separation constraint on the selected frequencies within each component. Specifically, for any two frequencies ω_j and ω_k included in the same component i, we require that they be separated by at least d units on the frequency grid. This constraint is incorporated into the joint prior as:

$$p(\{\mathbf{z}_j\}_{j=1}^{M_{\max}} \mid \boldsymbol{\pi}) \propto \left[\prod_{j=1}^{M_{\max}} \pi(\mathbf{z}_j)\right] \cdot \mathbb{I}\left(\min_{j<k; z_{ji}=z_{ki}=1} |\omega_j - \omega_k| \geq d \text{ for all } i \in \{1, \ldots, D\}\right).$$

This enforces the separation constraint in a component-wise manner, ensuring that active frequencies remain well-separated in each dimension without redundancy.

While the inclusion indicators $\mathbf{z}_j$ govern whether a frequency is selected in each component, the model retains flexibility at the coefficient level. Specifically, when $z_{ji} = 1$, the corresponding component i maintains its own coefficient vector $\boldsymbol{\beta}_{ji}$, allowing for differences in amplitude and phase across components even when the same frequency is shared. This formulation captures both synchrony and heterogeneity: shared frequencies can exhibit aligned periodic behavior across components, yet magnitude and phase of oscillations may vary. This allows the model to support flexible and interpretable multivariate spectral decompositions.

3. Markov Chain Monte Carlo Sampler. We first describe the MCMC sampler for the univariate model of Section 2.1. Let $\boldsymbol{\beta} = (\boldsymbol{\beta}_1', \ldots, \boldsymbol{\beta}_{M_{max}}')'$ be the vector of coefficients corresponding to the entire set of frequencies $\boldsymbol{\omega} = (\tilde{\omega}_1, \ldots, \tilde{\omega}_{M_{max}})$, where $\boldsymbol{\beta}_j = (\beta_{j1}, \beta_{j2})'$, and

let $\boldsymbol{z} = (z_1, \ldots, z_{M_{max}})$ denote the vector of inclusion indicators for the frequencies. Then, Bayesian inference is based on the following factorization of the joint posterior distribution,

$$p(\boldsymbol{z}, \boldsymbol{\beta}, \sigma^2 | \boldsymbol{y}, \boldsymbol{\omega}) \propto p(\boldsymbol{y} | \boldsymbol{z}, \boldsymbol{\beta}, \sigma^2, \boldsymbol{\omega}) p(\boldsymbol{z}) p(\boldsymbol{\beta}_{\boldsymbol{z}} | \boldsymbol{z}) p(\sigma^2). \tag{3.1}$$

We devise a Metropolis-within-Gibbs sampler that iteratively alternates between sampling the indicators $\boldsymbol{z}$, the variance σ^2 and the active coefficients $\boldsymbol{\beta}_{\boldsymbol{z}}$ corresponding to the included frequencies, namely $\boldsymbol{\beta}_{\boldsymbol{z}} = \{\boldsymbol{\beta}_j : z_j = 1,\ j = 1, \ldots, M_{max}\}$ for $\boldsymbol{\omega}_{\boldsymbol{z}} = \{\tilde{\omega}_j : z_j = 1,\ j = 1, \ldots, M_{max}\}$. At a generic MCMC iteration we update the parameters as follows:

• **Jointly updating active frequency indicators $\boldsymbol{z}$ and linear basis coefficients $\boldsymbol{\beta}_{\boldsymbol{z}}$**: We perform a joint sampling of $\boldsymbol{z}$ and $\boldsymbol{\beta}_{\boldsymbol{z}}$ following the frequently employed stochastic search MCMC algorithm of Savitsky, Vannucci and Sha (2011), that cleverly avoids having to deal with the changing dimensions of the parameter space via a joint update of the inclusion indicators and the corresponding coefficients. This algorithm involves add-delete-swap moves, enabling the exploration of the posterior space by sequentially visiting models with different numbers of frequencies. Importantly, all proposed moves are embedded within a Metropolis–Hastings (MH) framework with appropriate acceptance probabilities, ensuring that the resulting Markov chain satisfies detailed balance and leaves the posterior distribution invariant. At each step, the new model differs from the previous one by including and/or excluding one frequency. The indices for modification are selected uniformly at random among valid candidates, conditional on the minimum separation constraint provided in Eq. (2.3). Thus, the move set is defined on the support of the prior. To elaborate, at each iteration the new model is created from the previous one by randomly selecting one of the following moves:

1. *Adding or deleting a frequency*: Select one of the M_{max} indices in $\boldsymbol{z}^{curr}$ at random and modify its value (add: $0 \to 1$; delete: $1 \to 0$). This either includes a new frequency in the model or deletes a frequency currently included. This move is carried out jointly for the pair of coefficients $\boldsymbol{\beta}_j = (\beta_{j1}, \beta_{j2})$, that is, they are simultaneously added or removed, hence ensuring that the corresponding frequency ω_j is added or removed, respectively, from the model. When adding a frequency, the new pair of coefficients $\boldsymbol{\beta}_j^{prop}$ is sampled from the conjugate Normal posterior distribution (see Eq. 3.3 below) of the entire vector $\boldsymbol{\beta}_{\boldsymbol{z}}$ of active linear coefficients, conditional on the proposed active frequencies determined by $\boldsymbol{z}^{prop}$. The j-th components are then selected as the proposed linear coefficients. The acceptance probability for adding a frequency is given by

$$\min\left[1, \frac{p(\boldsymbol{y} | \boldsymbol{z}^{prop}, \boldsymbol{\beta}^{prop}, \sigma^2, \boldsymbol{\omega}) p(z_j^{prop}) \prod_{i=1}^{2} p(\beta_{ji}^{prop} | z_j^{prop})}{p(\boldsymbol{y} | \boldsymbol{z}^{curr}, \boldsymbol{\beta}^{curr}, \sigma^2, \boldsymbol{\omega}) p(z_j^{curr})}\right], \tag{3.2}$$

   whereas the one for deleting a frequency is calculated similarly by inverting Eq. (3.2).
2. *Swapping frequencies*: Select a 0 and a 1 at random from $\boldsymbol{z}^{curr}$, then switch their values. This process adds a new frequency to the model while removing a currently included one. Naturally, this is reflected in the corresponding pair of coefficients (β_{j1}, β_{j2}).

• **Updating the active linear basis coefficients $\boldsymbol{\beta}_{\boldsymbol{z}}$**: We perform a standard Bayesian update of the active linear coefficients $\boldsymbol{\beta}_{\boldsymbol{z}}$. This update is not required for ergodicity, but it helps to speed up the convergence of the chain by supplying relatively more sampled values of the parameters at a given iterations (Savitsky, Vannucci and Sha, 2011). Specifically, the posterior distribution of $\boldsymbol{\beta}_{\boldsymbol{z}}$ conditional on σ^2, $\boldsymbol{y}$, and $\boldsymbol{\omega}_{\boldsymbol{z}}$ follows a multivariate normal distribution, $\boldsymbol{\beta}_{\boldsymbol{z}} | \sigma^2, \boldsymbol{y}, \boldsymbol{\omega}_{\boldsymbol{z}} \sim \mathcal{N}_{2M}(\boldsymbol{\lambda}_T, \boldsymbol{\Sigma}_T)$, where M is the number of active frequencies. The parameters of these distributions are defined as:

$$\boldsymbol{\Sigma}_T^{-1} = \frac{1}{\sigma_\beta^2} \boldsymbol{I}_{2M} + \frac{1}{\sigma^2} \boldsymbol{X}(\boldsymbol{\omega}_{\boldsymbol{z}})' \boldsymbol{X}(\boldsymbol{\omega}_{\boldsymbol{z}}), \quad \boldsymbol{\lambda}_T = \frac{1}{\sigma^2} \boldsymbol{\Sigma}_T^{-1} \boldsymbol{X}(\boldsymbol{\omega}_{\boldsymbol{z}})' \boldsymbol{y}. \tag{3.3}$$

Here, the spectral matrix $\boldsymbol{X}(\boldsymbol{\omega_z})$ is the matrix of basis functions with rows given by

$$\boldsymbol{x}_t(\boldsymbol{\omega_z}) = \big(\cos(2\pi\tilde{\omega}_1), \sin(2\pi\tilde{\omega}_1), \ldots, \cos(2\pi\tilde{\omega}_M), \sin(2\pi\tilde{\omega}_M)\big), \quad t = 1, \ldots, T.$$

• **Updating the residual variance** σ^2: This is standard, as the posterior distribution of σ^2 follows an inverse-gamma distribution, $\sigma^2|\boldsymbol{y}, \boldsymbol{\beta_z}, \boldsymbol{\omega_z} \sim IG(\gamma_T, \nu_T)$, with updated parameters

$$\gamma_T = \frac{T + \gamma_0}{2}, \quad \nu_T = \frac{1}{2}\left(\nu_0 + \sum_{t=1}^{T}\{y_t - \boldsymbol{x}_t(\boldsymbol{\omega_z})'\boldsymbol{\beta_z}\}^2\right). \tag{3.4}$$

As noted above, the minimum-separation constraint is not merely imposed on the proposal mechanism but is embedded in the prior (Eq. 2.3). Therefore, the posterior is zero outside this constrained support, and the sampler only explores valid configurations. Since the proposal mechanism ensures that each transition remains within this support and the move set connects all valid configurations (e.g., from $(1, 0, 0, \ldots)$ to $(1, 0, \ldots, 1)$ is reachable through a finite number of add/delete steps that maintain spacing), the Markov chain is irreducible over the support of the posterior. Hence, the sampler is valid and converges to the correct target distribution. In our implementation, the list of valid candidate indices for add/delete and swap moves is dynamically updated at each MCMC step to reflect the current configuration and separation constraint. This ensures the proposal mechanism remains efficient and valid without revisiting zero-probability states. Furthermore, to improve efficiency of the sampler and expedite convergence, we propose initializing the frequencies in the model using the locations of the top M_{start} peaks from the periodogram. This approach leverages the periodogram as a simple yet effective tool for identifying dominant frequencies in the data, providing informed starting values and reducing time spent by the sampler exploring low-probability areas.

3.1. *MCMC Sampler for Multivariate Data.* The sampler scheme described above can be extended to the the case of multiple time series. In the setting described in Section 2.3 each component retains its own set of Fourier coefficients and residual variance and inference is coupled across series through a shared prior over joint frequency inclusion patterns. Accordingly, the sampler enables (i) structured cross-component frequency selection, (ii) flexible reallocation across dimensions, and (iii) Bayesian learning of inclusion structure. We provide a short description of the algorithm here, with full details in the Supplementary Material.

At each iteration, the sampler proposes a local modification to the binary inclusion vector associated with a single frequency—either by toggling inclusion in one component (add/delete) or by swapping the inclusion status between two components, which may belong to the same or different dimensions. Crucially, these swap moves can reallocate spectral features across different components, allowing the sampler to explore cross-series spectral sharing in a flexible, data-driven way. Each proposal is accepted only if it respects a minimum-separation constraint within each component, and the MH ratio is computed using only the components affected by the change, which significantly improves computational efficiency. In line with the univariate sampler, each proposed inclusion move is paired with a joint update of the corresponding Fourier coefficients, sampled from their conditional Gaussian distribution under the proposed model. This pair $(\mathbf{z}_j^{\text{prop}}, \boldsymbol{\beta}_j^{\text{prop}})$ is then accepted or rejected using a MH step. After the inclusion indicators are updated, we perform a second stage of updates: we re-sample all active Fourier coefficients and component-specific residual variances from their full conjugate posteriors. This two-step strategy improves mixing by refining the coefficients under the current model configuration. Finally, the parameter of the Dirichlet prior over inclusion configurations is updated via its conjugate posterior, allowing the model to adaptively learn the relative prevalence of inclusion patterns across components. This hierarchical structure encourages shared frequency selection when supported by the data, while preserving the flexibility to model component-specific spectral features.

3.2. *Posterior Inference.* The procedures described in this section apply to both the univariate and multivariate models introduced earlier. From the MCMC output, we estimate Posterior Probabilities of Inclusion (PPIs) for each frequency ω_j, formally defined as the posterior probability that $z_j = 1$, and estimated from the MCMC samples as $P(z_j = 1 \mid \boldsymbol{y}, \cdot) = \frac{1}{S}\sum_{s=1}^{S} z_j^{(s)}$, with S the number of MCMC iterations and $z_j^{(s)}$ the value of the latent indicator z_j at the s-th iteration. This quantity provides a probabilistic assessment of the importance of each frequency, allowing us to prioritize those with higher PPIs as key contributors to the signal. We follow Barbieri and Berger (2004), and select key frequencies as those with PPI greater than 0.5. Alternatively, one could impose a more stricter rule based on high posterior inclusion probabilities (e.g., PPI > 0.9) to promote additional parsimony in the model. The model also provides a posterior distribution over the total number of selected frequencies, which serves as a natural measure of model complexity. This is computed as:

$$P(M = j \mid \boldsymbol{y}, \cdot) = \frac{1}{S}\sum_{s=1}^{S} \mathbb{I}\left(\sum_{j=1}^{M_{\max}} z_j^{(s)} = j\right),$$

where $\mathbb{I}(\cdot)$ is the indicator function, which equals 1 if the sum of the latent indicators at iteration s is equal to j, and 0 otherwise. In the multivariate setting, this distribution can be computed separately for each component, or jointly over inclusion patterns, depending on the structure of interest (see Section 2.3).

4. Simulation Studies. We conduct simulation studies to assess the performance of our proposed method, hereafter referred to as **spectralSS** (spectral spike-and-slab).

4.1. *Illustrative Example.* In this simulation example, we generate a time series with $T = 512$ time points, where the periodic behavior is driven by $M = 4$ relevant frequencies $\boldsymbol{\omega} = (1/67, 1/21, 1/13, 1/8)$. The corresponding coefficients are $\beta_{11} = 0.6$, $\beta_{12} = -0.6$, $\beta_{21} = 1.0$, $\beta_{22} = 0.3$, $\beta_{31} = 1.0$, $\beta_{32} = 0.2$, $\beta_{41} = 1.0$, and $\beta_{42} = -1.0$, while the residual variance is set to $\sigma^2 = 1.5$. This selection of frequencies captures a range of oscillatory behaviors, incorporating both high-power and low-power components, as well as very low and relatively high frequencies, providing a diverse representation of the underlying signal. Figure 1 (a) shows a realization of the process alongside the underlying signal.

We set the hyperparameters for the inclusion probability α_j, following a Beta distribution, as $a = 1$ and $b = 10$, reflecting mild prior knowledge about sparsity. We choose an uninformative prior for the residual variance, setting $\gamma_0 = \nu_0 = 0.001$, and specify the slab variance of the linear coefficients as weakly informative with $\sigma_\beta^2 = 10$. When running the sampler, we selected a finely spaced grid of candidate frequencies, defined as $\omega_j = j \cdot \Delta\omega$ for $j = 0, 1, 2, \ldots, L$, with a step size of $\Delta\omega = 0.0001$ and a total of $L = 5000$ frequency points. In Section 4.1.2 below we show how this choice improves performances with respect to standard Fourier frequencies, which produce worse results in terms of both the accuracy of frequency estimation and the corresponding power. We ran the MCMC sampler for 50,000 iterations, discarding the first 25,000 as burn-in. For the initial values of the periodogram, we selected the two most prominent peaks and imposed a minimum distance between selected frequencies of $d = 3$ to avoid closely spaced frequencies. Additionally, we experimented with an alternative initialization by including all frequencies in the model from the start, to assess whether the sampler could correctly identify the relevant frequencies and corresponding linear coefficients, as detailed in the Supplementary Material. Although this approach worked, initializing the sampler with only the two peaks significantly reduced the computational time required for convergence. With this parameterization, the sampler was run in approximately 50 seconds using software implemented in Julia. Convergence diagnostics for the MCMC

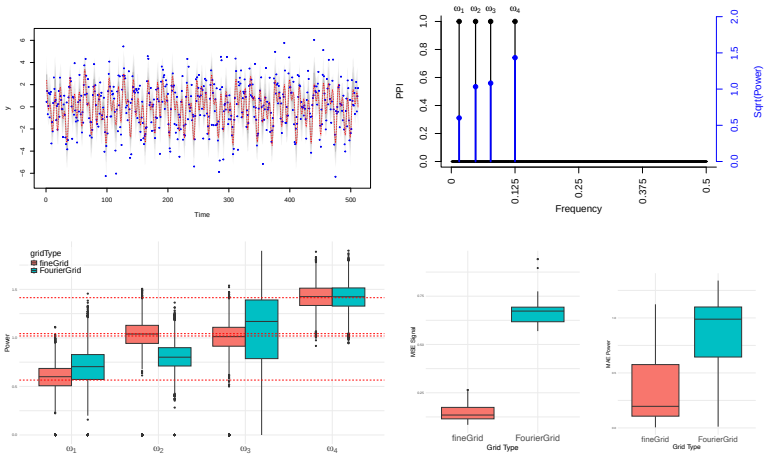

Fig 1: **Simulation Study**. (a) Simulated data (blue dots), true signal (dotted red line), and 100 draws from the estimated posterior predictive distribution (grey lines). (b) PPIs of the frequencies (black lines) and estimated square root of the power (blue lines) for frequencies with PPI> 0.5. (c) Box plots of the estimated square root of power $\hat{P}(\omega_j)$ at each frequency, conditioned on $M = 4$, with true power values shown as red dotted lines, for a finer grid and the classical Fourier grid. (d) Comparison of error metrics between a finer grid and the classical Fourier grid (MSE of the signal on the left and MAE of the power on the right).

sampler, including trace plots and the Heidelberger and Welch (Heidelberger and Welch, 1981) test to assess the model's reliability, are detailed in the Supplementary Material.

The mode of the posterior probability for the number of frequencies resulted in $p(M = 4|\boldsymbol{y}, \cdot) = 0.78$. Conditional on the modal number of frequencies, the estimated vector of significant frequencies were $\boldsymbol{\omega} = (0.0154, 0.0476, 0.0772, 0.125)$ and corresponding power $P(\boldsymbol{\omega}) = (0.744, 1.033, 1.082, 1.434)$. Figure 1 (a) displays a graphical posterior predictive check, consisting of the observations alongside 100 draws from the estimated posterior predictive (Gelman, 2007), showing that the model correctly captures the underlying signal. Figure 1 (b) illustrates the posterior inclusion probabilities (PPIs) of the frequencies together with the estimated square root of the power of the selected frequencies $\hat{P}(\omega_j) = \sqrt{\hat{\beta}_{j1}^2 + \hat{\beta}_{j2}^2})$ for those frequencies with PPI > 0.5. Using a more stringent threshold (PPI > 0.9) yields the same set of selected frequencies in this simulation study. Therefore, the model correctly identifies the number of frequencies and their location, as well as the corresponding magnitude.

4.1.1. *Comparison Study.* We evaluate the performance of our discrete spike-and-slab approach in identifying relevant frequencies against regularization techniques that operate on the same set of Fourier frequencies. Specifically, our methodology is benchmarked against four alternative approaches, each representing a different strategy for frequency selection in time series analysis: (i) **spectralLASSO**: This approach adopts the same saturated regression model framework as our proposed method. However, instead of using a spike-and-slab prior for frequency selection, variable selection is achieved via LASSO penalization (Tibshirani, 1996). This penalization method imposes an ℓ_1-norm constraint on the linear coefficients, which promotes sparsity by shrinking many of the Fourier coefficients to zero. (ii) **spectralEN** (spectral Elastic Net): Like LASSO, Elastic Net (Zou and Hastie, 2005) begins with a saturated regression model over all Fourier frequencies. However, it applies both ℓ_1- and ℓ_2-norm constraints. While this dual-penalty structure can be advantageous in

certain applications, the Elastic Net approach may include irrelevant frequencies due to its tendency to retain correlated predictors, potentially affecting model interpretability and precision in frequency selection. (iii) **spectralSCAD**: This method replaces the LASSO penalty with the Smoothly Clipped Absolute Deviation (SCAD, Fan and Li 2001) penalty, a non-convex regularization technique known to reduce shrinkage bias for large coefficients while maintaining sparsity. SCAD can offer improved frequency selection when the true signal is sparse and the Fourier basis is highly correlated. (iv) **spectralMCP**: This approach employs the Minimax Concave Penalty (MCP, Zhang 2010), another non-convex alternative designed to balance sparsity and estimation bias. MCP aggressively shrinks small coefficients while preserving larger ones, often leading to more precise frequency recovery under low signal-to-noise scenarios. We implemented spectralLASSO, spectralEN, spectralSCAD, and spectralMCP ourselves, using the `glmnet` and `ncvreg` packages in `R`. We utilized the following measures to assess accuracy in estimating both frequency locations and spectral power: (i) Mean Squared Error of the signal, calculated as $MSE_S = \frac{1}{T}\sum_{t=1}^{T}(f_t - \hat{f}_t)^2$, (ii) Absolute Error of the power, $AE_P = \left|\sum_{j=1}^{M} P(\omega_j) - \sum_{j=1}^{\hat{M}} \hat{P}(\omega_j)\right|$, and (iii) Absolute Error of the frequencies, $AE_F = \left|\sum_{j=1}^{M} \omega_j - \sum_{j=1}^{\hat{M}} \hat{\omega}_j\right|$, where the hat notation denotes estimated parameters. Here, f and $\hat{f}$ represent the true and estimated signals, respectively (e.g., $f_t = \sum_{j=1}^{M}\{\beta_{j1}\cos(2\pi\omega_j t) + \beta_{j2}\sin(2\pi\omega_j t)\}$).

The performance of our proposed approach across 50 simulations, compared with spectralLASSO, spectralEN, spectralSCAD, and spectralMCP, is illustrated in Figure 2. We present boxplots of (a) $\log(AE_F)$, (b) $\log(MSE_S)$, (c) $\log(AE_P)$, (d) the logarithm of the selected number of frequencies, and (e) log computational time. Our proposed method, **spectralSS**, demonstrates superior accuracy across all performance metrics. In particular, it achieves the lowest median error in frequency estimation, signal reconstruction, and power spectrum approximation, while also recovering the true number of frequencies more consistently than the other approaches. Although spectralSS is computationally more intensive than the regularization-based methods, the trade-off results in significantly improved performance. Among the penalization-based competitors, the non-convex methods—spectralSCAD and spectralMCP—generally outperform spectralLASSO and spectralEN in terms of accuracy, but tend to be less sparse or less stable in frequency selection than spectralSS. Notably, spectralEN exhibits a tendency to overestimate the number of active frequencies. This scenario considers a high-dimensional setting with a very small sample size relative to the number of candidate frequencies; specifically, the saturated spectral design matrix has dimension $512 \times 10{,}000$. Despite this challenging regime, our approach shows remarkable power and precision, consistently identifying the relevant components with minimal estimation error.

4.1.2. *Finer Grid of Frequencies.* Next, we investigate the advantages of selecting a finer grid of frequencies in comparison to the classical Fourier grid. To explore this, we simulated 50 time series based on the same simulation scenario described above, using both types of frequency grids. Figure 1 (c) presents box plots of the estimated square root of the power $\hat{P}(\omega_j)$ at each frequency, for each frequency grid, with the results conditioned on the estimated number of frequencies $M = 4$. The dotted horizontal red lines correspond to the true values of the power. For this example, selecting a finer frequency grid generally outperforms the Fourier grid, except for the frequency 0.125 (i.e., periodicity 8) where the results are almost identical. Figure1 (d) displays box plots of the MSE_S of the signal and the AE_P for both methods. Thus, the finer frequency grid consistently outperforms the classical Fourier grid in terms of both signal reconstruction and power estimation.

In the Supplementary Material we further investigate the behavior of the proposed method in scenarios where oscillatory components are closely spaced in the frequency domain and

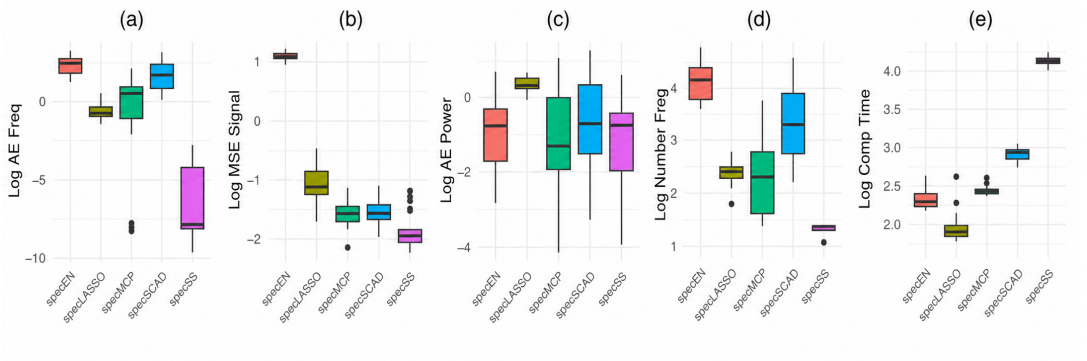

Fig 2: **Simulation Study.** Comparison of the performance of our proposed method (spectralSS) across 50 simulations against spectralLASSO, spectralEN, spectralSCAD and spectralMCP. Boxplots illustrate (a) $\log(AE_F)$, (b) $\log(MSE_S)$, (c) $\log(AE_P)$, (d) the logarithm of the selected number of frequencies, and (e) log computational time (in seconds).

compare its performance with alternative methods. We also examine the effectiveness of our proposed method in detecting spectral peaks when there is a mismatch between the model and the data-generating process. Finally, we present results from a sensitivity analysis.

4.2. *Bivariate Time Series.* Next, we consider a bivariate simulation with $T = 512$ time points, where each component exhibits distinct but partially overlapping periodic behavior. Specifically, the first series is generated using $M_1 = 2$ active frequencies, $\boldsymbol{\omega}_1 = (1/67,\ 1/13)$, while the second relies on $M_2 = 3$ active frequencies, $\boldsymbol{\omega}_2 = (1/21,\ 1/13,\ 1/6)$. The corresponding Fourier coefficients are set to $\boldsymbol{\beta}_{11} = (1.0,\ -0.4)$, $\boldsymbol{\beta}_{21} = (0.8,\ 1.0)$ for the first dimension, and $\boldsymbol{\beta}_{12} = (0.9,\ 0.5)$, $\boldsymbol{\beta}_{22} = (-1.0,\ 0.8)$, $\boldsymbol{\beta}_{32} = (0.7,\ -0.9)$ for the second. Independent Gaussian noise is added to each series, with component-specific variances set to $\sigma_1^2 = 1$ and $\sigma_2^2 = 1$. Notably, the two components share a common frequency at $1/13$, although their amplitudes and phases differ, yielding both shared and dimension-specific structure. We specified the Dirichlet hyperparameter for $\boldsymbol{\pi}$ as $\boldsymbol{\alpha} = (10.0,\ 3.0,\ 3.0,\ 3.0)$, leading to a prior that favors configurations with fewer jointly selected frequencies across components. Figure 3 (a) shows a realization of the process alongside the underlying signal.

The posterior mode for the number of active frequencies was $M_1 = 2$ and $M_2 = 3$ for the two components, with respective posterior probabilities $p(M_1 = 2 \mid \mathbf{y}, \cdot) = 0.87$ and $p(M_2 = 3 \mid \mathbf{y}, \cdot) = 0.82$. Conditional on these modal values, the estimated sets of dominant frequencies were $\boldsymbol{\omega}_1 = (0.016, 0.076)$ and $\boldsymbol{\omega}_2 = (0.047, 0.076, 0.166)$, with corresponding estimated signal power $P(\boldsymbol{\omega}_1) = (1.15, 0.76)$ and $P(\boldsymbol{\omega}_2) = (0.90, 0.61, 1.07)$. Figure 3(b) displays the PPIs across the candidate frequency grid for each component, along with the estimated square root of the signal power for frequencies with PPI > 0.5. The model successfully recovers the true number and location of the generating frequencies in both series and, crucially, identifies a shared frequency at $\omega = 0.076$ that is included in both components but with distinct amplitudes—highlighting its ability to capture partially overlapping yet component-specific spectral structure. In the Supplementary Materials we further show that similar results are obtained when the data are generated with correlated errors across components, despite the independence assumption in the fitted model.

5. Applications to Wearable Device Data. Wearable biosensors offer an unprecedented opportunity to non-invasively monitor physiological rhythms in real-world settings (De Zambotti et al., 2024). We apply our proposed framework on two distinct datasets: (i) actigraphy data from epilepsy patients, and (ii) temperature-activity data from a healthy subject.

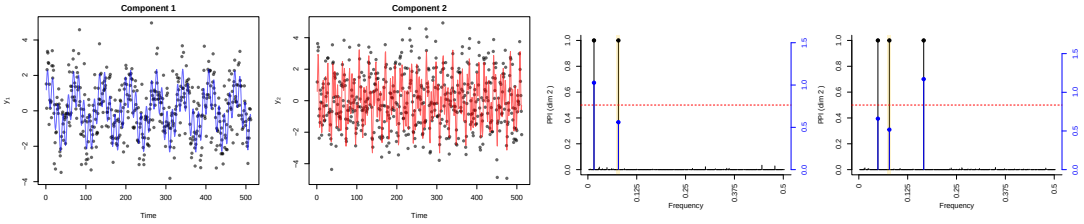

Fig 3: **Simulation Study.** (a) Simulated observations (dots) and true underlying signal (lines) for the two components. (b) Posterior inclusion probabilities (PPIs; black lines) and estimated square root of signal power (blue lines) for frequencies with PPI greater than 0.5. The shaded region highlights a frequency that is jointly selected across both components but with distinct amplitudes, illustrating the model's ability to recover partially shared spectral features.

5.1. *Wearable Activity Monitoring in Epilepsy Studies.* Actigraphy data provide a non-invasive means of tracking rest-activity cycles, capturing key metrics such as sleep duration and daytime movement levels. In epilepsy research, these data can reveal patterns associated with seizure frequency and psychiatric symptoms like depression and anxiety (Abboud et al., 2023). Identifying the specific frequencies, amplitudes, and phases within these cycles presents a challenge, yet these components may hold promise as a biomarkers that may encode states of neurological or psychiatric compromise (Smagula et al., 2024).

5.1.1. *Study Framework and Data Preprocessing.* Our analysis focused on data previously introduced by Abboud et al. (2023). Participants were enrolled through the Baylor Comprehensive Epilepsy Clinic in Houston, TX. The study targeted adults aged 18 to 75 with partial-onset seizures. Each participant wore an FDA-approved Actiwatch-2 device (Philips Respironics), highly durable and widely trusted in clinical research, continuously on their non-dominant wrist, including during sleep and showers. The Actiwatch-2 uses piezoelectric sensors calibrated to measure wrist movement via a single-axis accelerometer. For this analysis, we considered six individuals. We conducted a series of preprocessing steps on the raw activity data, adapting established methods (Huang et al., 2018; Abboud et al., 2023). First, we applied a logarithmic transformation (adding one to each value). Next, we used a moving average filter with a 15-point window, equivalent to a 15-minute span, followed by downsampling by averaging every five consecutive points, yielding a 5-minute sampling rate. This process resulted in time series with lengths T ranging from 1725 to 2301 data points.

5.1.2. *Parameter Settings.* For our proposed methodology, we selected a grid of Fourier frequencies, ensuring the grid was sufficiently dense given the length of the time series. We set $a = 1$ and $b = 10{,}000$ to promote a high degree of prior sparsity for the frequencies. This choice is guided by the sensitivity analysis presented in the Supplementary Material, where we evaluated the effect of varying b on frequency recovery across a wide range of sample sizes. Notably, for time series lengths comparable to those encountered in the real data applications, the setting $b = 10{,}000$ provided a favorable balance—yielding low error in frequency selection and signal reconstruction, while closely approximating the true number of active frequencies. While alternative strategies such as empirical Bayes via MMLE or tuning via cross-validation could be explored, our empirical results indicate that the chosen value of b offers robust and competitive performance in the considered regimes. Additionally, we imposed a minimum frequency distance of $d = 5$ to ensure sufficient separation between selected components. As discussed in Section 2.2.1, this constraint is particularly relevant in biological time series, such as actigraphy, where dominant rhythms like circadian ($\sim$ 24

hours), ultradian (e.g., 12 hours), and infradian cycles are of interest. On a fine frequency grid, nearby bins often correspond to periods that are nearly indistinguishable in physiological terms (e.g., 24.1 vs 23.7 hours). Setting $d = 5$ reduces the risk of selecting redundant or collinear frequencies in such dense regions, enhancing interpretability while stabilizing inference. The MCMC sampler was run for $50,000$ iterations, with the first $25,000$ discarded as burn-in. Posterior predictive graphical checks, presented in the Supplementary Material, demonstrate that the estimated model accurately captures the structure of the observed data.

5.1.3. *Results.* Figure 4 presents the results for Subject 1, while the corresponding results for the remaining five subjects are shown in the Supplementary Material. The first panel displays the observed standardized actigraphy time series. The second panel shows the estimated frequencies alongside their PPIs and corresponding square-root spectral power, with the raw periodogram in the third panel serving as a baseline for comparison. The fourth panel illustrates the posterior distribution over the number of selected frequencies. These results highlight the efficacy of the proposed spike-and-slab prior in identifying the most significant frequencies that contribute to the observed oscillatory patterns. The number of selected frequencies varies across individuals, reflecting differences in the periodic structure of their actigraphy data. Some subjects exhibit relatively simple oscillatory patterns characterized by a few dominant frequencies, while others display more intricate dynamics with a greater number of selected periodic components. As expected, all subjects exhibited a strong circadian (24-hour) periodicity, indicating the robust presence of diurnal rest-activity cycles. Subjects 3, 4, and 5, shown in the Supplementary Material, share similar oscillatory patterns, with ultradian (less than 24 hours) oscillations observed at approximately 12 and 6 hours. Subject 6 exhibits a comparatively simpler oscillatory pattern, characterized primarily by periodicities at approximately 24 and 12 hours. Notably, the periodogram for Subject 4 appears noisier compared to the others. In contrast, Subject 2 demonstrates an additional periodicity at approximately 13 hours, while Subject 1 presents a more complex oscillatory pattern. This is evidenced by the greater number of selected frequencies, which include ultradian oscillations at approximately 12, 8, 6, and 3.6 hours, in addition to the dominant circadian rhythm.

Our findings hold significant implications for monitoring circadian and ultradian rhythms in clinical populations. By automatically identifying key periodic components, the proposed methodology could serve as a powerful tool for tracking variations in rest-activity rhythms, potentially offering new insights into the relationship between these rhythms and neurological states, such as seizure frequency and psychiatric symptoms, including depression or anxiety. Indeed, the connection between mental health and actigraphy has primarily focused on circadian rhythms, while interindividual variability in the expression of ultradian rhythms has not been extensively explored. This is because traditional periodography often struggles to distinguish genuine ultradian rhythms from artifactual periodicities, as the dominance of the circadian oscillator can obscure ultradian rhythmicity. Furthermore, exogenous factors, such as a regular 8-hour eating schedule, can also produce pronounced 8-hour ultradian rhythms. Recent findings (Adhyapak et al., 2024) reveal unique signatures of ultradian rhythms expression between individuals, with these signatures also exhibiting temporal variability.

We further compared our proposed approach to spectralLASSO and spectralEN. SpectralLASSO identified an average of 322 frequencies across the six subjects, whereas spectralEN identified an average of 353 frequencies. This indicates that both methodologies struggle to automatically isolate the dominant frequencies from the broader set of candidates, with spectralEN tending to select a larger number of frequencies than spectralLASSO. These comparisons highlight the strengths of the spike-and-slab prior, which effectively balances sparsity and model fidelity. By doing so, it enables the identification of a parsimonious set of frequencies that closely align with prominent peaks in the periodogram, while also mitigating the

risk of overfitting to noise. Additionally, the posterior probability intervals provide a clear and interpretable measure of uncertainty in frequency estimation. These findings underscore the advantages of selection approaches over traditional spectral analysis methods, as they automate the selection of meaningful frequencies in a principled manner.

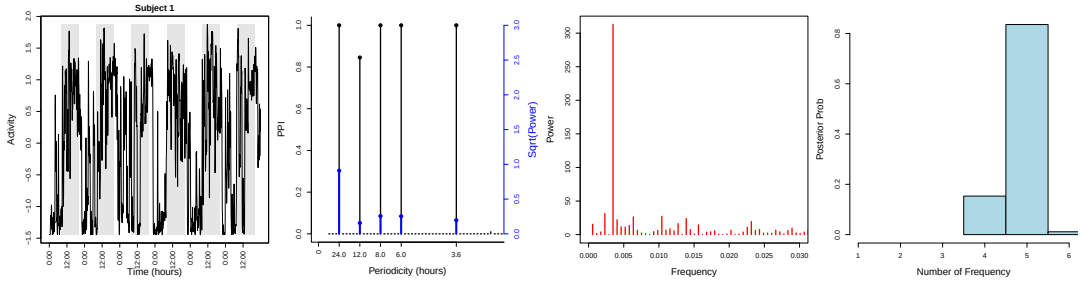

Fig 4: **Wearable Activity Monitoring.** Left panel: standardized activity time series, with gray bands indicating 8 AM–8 PM; second panel: estimated periodicities with posterior inclusion probabilities and square-root spectral power; third: raw periodogram; and fourth: posterior distribution of the number of selected frequencies.

5.2. *Bivariate Case Study.* Skin temperature and physical activity are two circadian-regulated signals with known physiological coupling. While core body temperature usually peaks in the late afternoon, skin temperature at the body's surface tends to rise during the early night, helping the body release heat and fall asleep more easily. This antiphase dynamic reflects the underlying circadian drive on thermoregulation and behavior, and has been leveraged in personalized medicine approaches such as chronotherapy. In recent years, mobile health platforms have enabled simultaneous and continuous telemetric monitoring of skin surface temperature and physical activity, offering a unique opportunity to study these rhythms. The data analyzed here were originally collected using the PiCADo mobile e-health platform and have been studied separately in previous works. Under a stationarity assumption, Komarzynski et al. (2018) analyzed the skin temperature time series and identified both strong 12-hour and 24-hour rhythms. Hadj-Amar et al. (2020) applied a nonstationary oscillatory model with an unknown change point to the same data, demonstrating that skin temperature alone can be used to detect sleep periods and reveal ultradian oscillations during the night—oscillations that are consistent with transitions between REM and non-REM sleep stages. Huang et al. (2018) modeled the activity data using a circadian hidden Markov model, while Hadj-Amar, Jewson and Fiecas (2023) employed a semi-Markov switching process for activity dynamics. To our knowledge, this is the first time that temperature and activity data from this dataset have been jointly analyzed within a unified statistical framework. Our preprocessing procedure follows Komarzynski et al. (2018), applying a square-root transformation to the raw activity counts and subsequently standardizing both time series. The resulting bivariate signals, spanning four consecutive days, are displayed in Figure 5(a).

5.2.1. *Results.* We adopt the same hyperparameter settings of the univariate application (Section 5.1) and specify the Dirichlet prior parameter $\boldsymbol{\alpha}$ for the inclusion patterns as in the bivariate simulation study (Section 4.2). The posterior distribution over the number of selected frequencies indicates that physical activity most likely contains six dominant periodic components, with posterior mode $P(M_1 = 6) = 0.35$, followed by seven components with

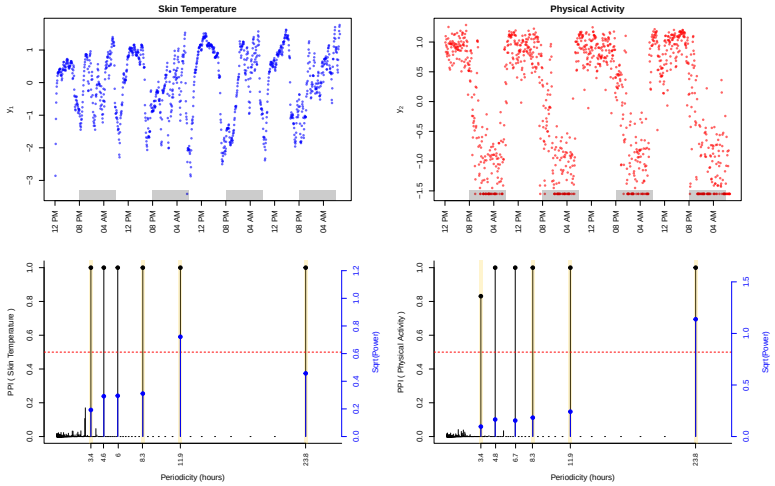

Fig 5: **Bivariate Case Study.** (a) Time series of skin temperature and physical activity with nighttime intervals shaded in gray (8 PM to 8 AM). (b) PPIs (black) and estimated power (blue) across periodic components, with shared periodicities highlighted in orange.

$P(M_1 = 7) = 0.29$. In contrast, skin temperature shows a posterior mode at seven components ($P(M_2 = 7) = 0.33$), with substantial probability mass at six and eight components ($P(M_2 = 6) = 0.23$, $P(M_2 = 8) = 0.23$). Posterior predictive checks indicate that the model faithfully reproduces the temporal dynamics of both signals, as evidenced by close alignment between observed data and posterior predictive draws (see Supplementary Material).

Figure 5(b) displays PPIs and estimated signal strength for the periodic components of physical activity (right panel) and skin temperature (left panel). In the activity signal, a dominant circadian rhythm emerges at approximately 24 hours. In contrast, while the temperature signal also shows evidence of a circadian component, its posterior amplitude is notably lower, suggesting reduced circadian influence relative to activity. Both modalities reveal rich ultradian dynamics. In particular, common components are identified near 3.4 and 8.3 hours, although their relative contribution differs between signals. For instance, the 8.3-hour oscillation exhibits high amplitude in temperature but is comparatively weaker in activity, whereas the 3.4-hour component, though less prominent, shows similar magnitude across both modalities. The presence of a shared ultradian rhythm near 3.4 hours is biologically meaningful, as it likely reflects coupled regulation between physical activity and thermoregulatory systems (Goh et al., 2019). In contrast, temperature shows additional specific components—such as at 6 hours and 11.9 hours—that are absent or significantly smaller in the activity signal, suggesting internal thermoregulatory cycles that operate independently of physical movement. The 11.9-hour component, consistent with a circahemidian rhythm, may emerge as a subharmonic of the underlying circadian oscillator (Grant et al., 2018). Moreover, some ultradian components appear modality-specific, including the 4.6-hour and 6-hour rhythms in temperature and the 4.8-hour and 6.7-hour components in activity, underscoring partially independent rhythmic processes. Overall, our analysis reveals that temperature and activity signals follow different, yet partly overlapping, rhythmic patterns. While physical activity primarily captures the circadian cycle, skin temperature exhibits a broader range of ultradian periodicities, emerging as a complementary biomarker for chronophysiological dynamics. The bivariate framework offers a flexible model of how physiological rhythms work together or differ.

6. Discussion. We have introduced a Bayesian spike-and-slab framework for spectral analysis, offering a robust and interpretable approach to identifying meaningful periodicities

in time series data. The proposed approach integrates frequency selection and dimensionality reduction using a refined grid of candidate frequencies, achieving high-resolution recovery of oscillatory components while enforcing sparsity through a spike-and-slab prior. We have designed an efficient stochastic search algorithm to explore the posterior space and estimate posterior inclusion probabilities to assess the significance of each frequency. Using simulated data, we have shown superior performance of the proposed method in identifying relevant frequencies and estimating spectral power compared to alternative approaches. Additionally, we have shown that using a finer frequency grid enhances accuracy in estimating power and reduces error metrics relative to the classical Fourier grid, highlighting the advantages of higher-resolution frequency selection. An application to actigraphy data from individuals with partial-onset seizures demonstrated the model's ability to identify clinically relevant rhythms, highlighting its potential as a tool for linking rest–activity patterns to neurological and psychiatric states. We have further extended the framework to the multivariate setting through a hierarchical prior on frequency inclusion, enabling the identification of both shared and component-specific rhythms across multiple signals. Applied to bivariate wearable measurements from a healthy individual to jointly model wrist activity and skin temperature, the model successfully captured partially overlapping ultradian components and known physiological coupling between thermoregulatory and behavioral rhythms. These results underscore the flexibility of the proposed method for multichannel physiological data and support its use in future applications of chronobiological modeling and digital health analytics.

Extensions in several directions may further broaden the utility and robustness of the proposed framework. First, the current model introduces a minimum separation parameter d that enforces a spacing constraint between selected frequencies through the indicator function $I(|\omega_j - \omega_k| \geq d)$. This parameter acts as a resolution control that prevents the model from selecting clusters of nearly indistinguishable frequencies on a refined grid. An interesting extension would be to treat d as a random parameter and allow the data to inform the degree of separation enforced by the model. However, because the constraint directly determines the set of admissible configurations of the inclusion indicators, different values of d correspond to different model spaces. In particular, the support of the prior on the inclusion indicators depends on d, and the constrained prior distribution involves a normalization constant equal to the number of admissible frequency configurations satisfying the separation constraint. This quantity depends on d and is combinatorial in nature, which would substantially complicate posterior computation and the design of the MCMC algorithm. For this reason, in the present work d is treated as a user-specified resolution parameter. Nevertheless, allowing the degree of frequency separation to be learned from the data is an interesting research direction.

A second limitation concerns scalability in the multivariate extension. In the proposed formulation, each candidate frequency ω_j is associated with a joint inclusion pattern $(z_j^{(1)}, \ldots, z_j^{(D)})$ across the D components, leading to 2^D possible configurations. As a result, the number of possible joint patterns grows exponentially with the dimension D, increasing both the number of parameters in the prior specification and the computational cost of exploring the corresponding model space through MCMC. While this formulation allows the model to capture shared and component-specific oscillatory behavior in a coherent probabilistic framework, the exponential growth in the number of configurations may limit scalability when the number of components becomes large. In practice, the approach remains computationally feasible for small to moderate values of D, such as the bivariate case considered in our experiments, where the joint inclusion structure provides a flexible mechanism for modeling partially shared rhythms across signals. More scalable formulations for higher-dimensional multivariate time series, potentially through more parsimonious joint priors or structured dependence across components, is another promising direction for future work.

Additional extensions could enhance the modeling framework. Alternative basis functions, such as wavelets or splines, could improve performance in the presence of nonstationary or localized temporal features. More flexible error structures, such as autoregressive or state-space models, would allow to better capture residual temporal dependence when needed. From a theoretical standpoint, posterior contraction rates and model selection consistency under spike-and-slab priors in the spectral domain could provide valuable insights into the estimator's properties. Additional gains could come from integrating time-varying covariates or subject-level information, for a deeper exploration of how rhythmic dynamics relate to external influences. Finally, improvements in computational efficiency via parallelization or variational inference would support applications to larger datasets. Although developed for wearable device measurements, this methodology could be adapted to detect periodic structure in other scientific domains, including neuroscience, finance, and meteorology.

Supplementary Materials. Supplementary materials available online include additional simulation details, convergence diagnostics, and posterior predictive checks. The *Julia* files used for the analysis can be accessed at https://github.com/Beniamino92/BayesFreqSelect. Real data used in the applications can be obtained from the corresponding author, upon reasonable request.

REFERENCES

ABBOUD, M. A., KAMEN, J. L., BASS, J. S., LIN, L., GAVVALA, J. R., RAO, S., SMAGULA, S. F. and KRISHNAN, V. (2023). Actigraphic correlates of neuropsychiatric symptoms in adults with focal epilepsy. *Epilepsia* **64** 1640–1652.

ADHYAPAK, N., ABBOUD, M. A., RAO, P. S., KAR, A., MIGNOT, E., DELUCCA, G., SMAGULA, S. F. and KRISHNAN, V. (2024). Stability and Volatility of Human Rest-Activity Rhythms: Insights from Very Long Actograms (VLAs). *medRxiv*.

BARBIERI, M. M. and BERGER, J. O. (2004). Optimal predictive model selection. *Annals of Statistics* **32** 870-897.

BRETTHORST, G. L. (1988). Excerpts from Bayesian spectrum analysis and parameter estimation. In *Maximum-Entropy and Bayesian Methods in Science and Engineering: Foundations* 75–145. Springer.

BROCKWELL, P. J. and DAVIS, R. A. (1991). *Time series: theory and methods*. Springer science & business media.

BROWN, P. J., VANNUCCI, M. and FEARN, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society, Series B* **60** 627–641.

CADONNA, A., KOTTAS, A. and PRADO, R. (2019). Bayesian spectral modeling for multiple time series. *Journal of the American Statistical Association* **114** 1838-1853.

CHOUDHURI, N., GHOSAL, S. and ROY, A. (2004). Bayesian estimation of the spectral density of a time series. *Journal of the American Statistical Association* **99** 1050–1059.

DE ZAMBOTTI, M., GOLDSTEIN, C., COOK, J., MENGHINI, L., ALTINI, M., CHENG, P. and ROBILLARD, R. (2024). State of the science and recommendations for using wearable technology in sleep and circadian research. *Sleep* **47** zsad325.

FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association* **96** 1348–1360.

FRIEDRICHS-MAEDER, C., PROIX, T., TCHENG, T. K., SKARPAAS, T., RAO, V. R. and BAUD, M. O. (2024). Seizure Cycles under Pharmacotherapy. *Annals of neurology* **95** 743–753.

GELMAN, A. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.

GEORGE, E. I. and MCCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association* **88** 881–889.

GOH, G. H., MALONEY, S. K., MARK, P. J. and BLACHE, D. (2019). Episodic ultradian events—ultradian rhythms. *Biology* **8** 15.

GRANADOS-GARCIA, G., FIECAS, M., BABAK, S., FORTIN, N. J. and OMBAO, H. (2022). Brain waves analysis via a non-parametric Bayesian mixture of autoregressive kernels. *Computational statistics & data analysis* **174** 107409.

GRANT, A. D., NEWMAN, M. and KRIEGSFELD, L. J. (2020). Ultradian rhythms in heart rate variability and distal body temperature anticipate onset of the luteinizing hormone surge. *Scientific reports* **10** 20378.

GRANT, A. D., WILSTERMAN, K., SMARR, B. L. and KRIEGSFELD, L. J. (2018). Evidence for a coupled oscillator model of endocrine ultradian rhythms. *Journal of biological rhythms* **33** 475–496.

HA, H. and ZHANG, H. (2019). Performance-influence model for highly configurable software with fourier learning and lasso regression. In *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)* 470–480. IEEE.

HADJ-AMAR, B., JEWSON, J. and FIECAS, M. (2023). Bayesian approximations to hidden semi-Markov models for telemetric monitoring of physical activity. *Bayesian Analysis* **18** 547–577.

HADJ-AMAR, B., RAND, B. F., FIECAS, M., LÉVI, F. and HUCKSTEPP, R. (2020). Bayesian model search for nonstationary periodic time series. *Journal of the American Statistical Association* **115** 1320–1335.

HEIDELBERGER, P. and WELCH, P. D. (1981). A spectral method for confidence interval generation and run length control in simulations. *Communications of the Association for Computing Machinery* **24** 233–245.

HU, Z. and PRADO, R. (2023). Fast Bayesian inference on spectral analysis of multivariate stationary time series. *Computational Statistics & Data Analysis* **178** 107596.

HUANG, Q., COHEN, D., KOMARZYNSKI, S., LI, X.-M., INNOMINATO, P., LÉVI, F. and FINKENSTÄDT, B. (2018). Hidden Markov models for monitoring circadian rhythmicity in telemetric activity data. *Journal of The Royal Society Interface* **15** 20170885.

INCREMONA, A. and DE NICOLAO, G. (2019). Spectral characterization of the multi-seasonal component of the Italian electric load: a LASSO-FFT approach. *IEEE Control Systems Letters* **4** 187–192.

KEIL, A., BERNAT, E. M., COHEN, M. X., DING, M., FABIANI, M., GRATTON, G., KAPPENMAN, E. S., MARIS, E., MATHEWSON, K. E., WARD, R. T. et al. (2022). Recommendations and publication guidelines for studies using frequency domain and time-frequency domain analyses of neural time series. *Psychophysiology* **59** e14052.

KOMARZYNSKI, S., HUANG, Q., INNOMINATO, P. F., BOUCHAHDA, M., BILSKI, K., ULUSAKARYA, A., BEAU, J., MOCQUERY, M., ADAMS, J., KARABOUE, A. et al. (2018). Picado: a mobile e-health platform for real-time monitoring of circadian dynamics of symptoms and physiological functions in cancer patients. *NPJ Digital Medicine* **1** 1–10.

LANGE, H., BRUNTON, S. L. and KUTZ, J. N. (2021). From fourier to koopman: Spectral methods for long-term time series prediction. *Journal of Machine Learning Research* **22** 1–38.

LAU, S. C., ZHANG, G., RUESCHMAN, M., LI, X., IRWIN, M. R., KRAFTY, R. T., MCCALL, W. V., SKIDMORE, E., PATEL, S. R., REDLINE, S. et al. (2024). Sleep–wake behavioral characteristics associated with depression symptoms: findings from the Multi-Ethnic Study of Atherosclerosis. *Sleep* **47** zsae045.

MITCHELL, T. J. and BEAUCHAMP, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the american statistical association* **83** 1023–1032.

PERCIVAL, D. B. and WALDEN, A. T. (1993). *Spectral analysis for physical applications*. cambridge university press.

SAVITSKY, T., VANNUCCI, M. and SHA, N. (2011). Variable selection for nonparametric Gaussian process priors: Models and computational strategies. *Statistical science: a review journal of the Institute of Mathematical Statistics* **26** 130.

SHUMWAY, R. H., STOFFER, D. S. and STOFFER, D. S. (2000). *Time series analysis and its applications* **3**. Springer.

SMAGULA, S. F., ZHANG, G., KRAFTY, R. T., RAMOS, A., SOTRES-ALVAREZ, D., RODAKOWSKI, J., GALLO, L. C., LAMAR, M., GUJRAL, S., FISCHER, D. et al. (2024). Sleep-wake behaviors associated with cognitive performance in middle-aged participants of the Hispanic Community Health Study/Study of Latinos. *Sleep Health* **10** 500-507.

TADESSE, M. G. and VANNUCCI, M. (2021). Handbook of Bayesian variable selection. *Chapman and Hall/CRC*.

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **58** 267–288.

TSENG, A., SHRIKUMAR, A. and KUNDAJE, A. (2020). Fourier-transform-based attribution priors improve the interpretability and stability of deep learning models for genomics. *Advances in Neural Information Processing Systems* **33** 1913–1923.

VIELUF, S., CANTLEY, S., KRISHNAN, V. and LODDENKEMPER, T. (2024). Ultradian rhythms in accelerometric and autonomic data vary based on seizure occurrence in paediatric epilepsy patients. *Brain Communications* **6** fcae034.

WU, H.-T. and ZHOU, Z. (2024). Frequency detection and change point estimation for time series of complex oscillation. *Journal of the American Statistical Association* **119** 1945–1956.

ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics* **38** 894—942.

ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **67** 301–320.