

LOW-PASS FILTERING AS BAYESIAN INFERENCE

Cristbal Valenzuela* Felipe Tobar*†

*Departament of Mathematical Engineering, Universidad de Chile.

†Center for Mathematical Modeling, Universidad de Chile.

ABSTRACT

We propose a Bayesian nonparametric method for low-pass filtering that can naturally handle unevenly-sampled and noise-corrupted observations. The proposed model is constructed as a latent-factor model for time series, where the latent factors are Gaussian processes with non-overlapping spectra. With this construction, the low-pass version of the time series can be identified as the low-frequency latent component, and therefore it can be found by means of Bayesian inference. We show that the model admits exact training and can be implemented with minimal numerical approximations. Finally, the proposed model is validated against standard linear filters on synthetic and real-world time series.

Index Terms— Spectral estimation, nonuniformly-sampled data, Gaussian process, low-pass filters, Bayesian inference.

1. INTRODUCTION

Monitoring the spectral content of a time series is of critical importance in real-world applications across a wide range of scientific disciplines. This is because the concentration of energy at a specific range of frequencies might be indicative of mechanical faults [1], cardiac anomalies [2], astronomical discoveries [3, 4], and whale calls from submarine audio recordings [5] to name a few.

The standard practice to isolate components within a specific frequency range from a time-series observation, referred to as *filtering*, is to convolve the observations with an object called *linear filter*. This convolution removes all frequencies that do not correspond to the desired frequency range, thus, *filtering out* unimportant frequencies. The theoretical rationale behind this approach is supported by the application of the Convolution Theorem [6] to power spectral densities (PSD): the PSD of a filtered time series corresponds to the PSD of the linear filter (user-designed) multiplied by the PSD of the observed time series (not controllable). This result allows for designing the linear filter so as to remove unwanted frequency components to then perform the numerical convolution. We refer to low-pass filtering when the range of frequencies to be removed are centred (symmetrically) around zero.

We identify two drawbacks of this standard approach to filtering. First, to perform the numerical convolution, the time series has to be uniformly sampled, that is, no missing observations or random acquisition times can be allowed. This is a rather stringent assumption, since in real-world applications missing data is commonplace due to mechanical or electrical failures and the sampling rate is given by the hardware. For instance, the observations of light curves in Astronomy are only available at some time instants due to climate conditions, orientation of the telescope and even the priority of the experiment within the agenda of the observatory. The second drawback of the convolution method is its implicit deterministic assumption: by computing a low-pass version of the time series as

a moving average, we are accumulating observation noise without properly accounting for the dispersion that this might cause.

We aim to address these two drawbacks by formulating the low-pass filtering problem as a Bayesian inference one. We model the observed time series as a mixture of three latent components: one of low frequencies, one of high frequencies, and an observation noise component. Then, we find the low-frequency component through probabilistic inference: we place a prior distribution on each component and then, using observations of the time series, we find the posterior distribution over the low-frequency component. In particular, we choose Gaussian processes [7] priors over the components to leverage the expressiveness of the GP formulation while introducing minimal numerical approximations, therefore, the proposed method will be referred to as Gaussian process low-pass filter (GPLP).

2. BACKGROUND: GAUSSIAN PROCESSES

2.1. Spectral representation of Gaussian processes

A Gaussian process (GP) [7] over the input set \mathbf{X} is a real-valued stochastic process $(f(x))_{x \in \mathbf{X}}$, such that for any finite subset of inputs $\{x_i\}_{i=1}^n \subset \mathbf{X}$, the random variables $\{f(x_i)\}_{i=1}^n$ are jointly Gaussian. Without loss of generality we choose $\mathbf{X} = \mathbb{R}^N$. In this sense, a GP defines a distribution over real-valued functions $f : \mathbf{X} \mapsto \mathbb{R}, x \mapsto f(x)$, that is uniquely determined by its mean function $m(x) = \mathbb{E}(f(x))$, typically assumed to be zero, and its covariance kernel $K(x, x') = \text{cov}(f(x), f(x'))$, $x, x' \in \mathbf{X}$.

The covariance kernel summarises the dynamic behaviour of the GP and thus it is key when designing GP models. In practice, we can rely upon the Wiener-Khinchin theorem [8], which states that an integrable function $K : \mathbb{R}^N \mapsto \mathbb{C}$ is the covariance function of a weakly-stationary mean-square-continuous stochastic process $f : \mathbb{R}^N \mapsto \mathbb{R}$ if and only if it admits the representation

$$K(\tau) = \int_{\mathbb{R}^n} e^{2\pi i \omega^T \tau} S(\omega) d\omega, \quad (1)$$

where $S(\omega)$ is a non-negative bounded function on \mathbb{R}^n and i denotes the imaginary unit. Henceforth, given a kernel K , we will refer to S as their power spectral density given by the above theorem; where the PSD $S(\omega)$ is the Fourier transform of the covariance kernel $K(\tau)$. This result allows us to encode spectral properties directly in the covariance function by first designing the PSD to then calculate the kernel as the inverse Fourier transform of the so designed PSD function.

The relationship between GPs and spectral representations has acquired attention recently in the machine learning community. For instance, covariance functions can be constructed in the spectral domain in parametric [9, 10] and nonparametric [11, 12] ways. Additionally, the harmonic structure of GPs has been exploited to develop

computationally-efficient sparse GP models by using *inducing variables* in the spectral domain [13, 14, 15]. More recently, GPs have also been considered to address the spectral estimation problem, in particular, for nonuniformly-sampled data and detection of periodicities [16, 17, 18]. An open challenge in the spectral treatment of GPs, is that learning frequency representations is prone to local optima, since one aims to approximate a periodogram; this has been partially addressed using Bayesian optimisation with derivatives [19] and derivative-free Monte Carlo methods [20]. In our case, however, training simply involves a standard square exponential kernel (presented next) and therefore optimisation is straightforward.

2.2. The square exponential case

The *de facto* covariance kernel for GP models is the square-exponential covariance denoted by

$$K_{SE}(x, x') = \sigma^2 \exp\left(-\frac{1}{2l^2} \|x - x'\|^2\right), \quad (2)$$

where the parameter σ^2 denotes the marginal variance of the process (i.e., the magnitude) and l denotes the lengthscale, that is, the range of correlation between values of the process: the larger the lengthscale the longer the range of temporal correlations.

The popularity of the SE kernel stems from its properties [7]. The paths generated by a GP with an SE kernel are (i) dense in the space of continuous functions, (ii) infinite-times differentiable a.e. and (iii) smooth, meaning that their power spectral density is concentrated around zero. In fact, due to the exponential form of the Fourier operator, the PSD of the SE kernel is also SE and given by

$$S_{SE}(\xi) = \mathcal{F}\{K_{SE}\}(\xi) = \sigma^2 \sqrt{2\pi} l^2 \exp(-2\pi^2 l^2 \xi^2), \quad (3)$$

where the lengthscale of this PSD (spectral domain) is now inversely proportional to the lengthscale of the covariance (temporal domain); this has key advantages when using GPs for spectral estimation [18, 14]. This can be understood intuitively: a process with long-range correlations has low frequency energy (smooth), whereas a kernel with short-length correlations necessarily has high frequency components (rough).

The generative model proposed in the next section will represent observed signals as a GP with SE covariance function composed of a mixture of (non overlapping) low- and high-frequency components.

3. A LATENT-COMPONENT GENERATIVE MODEL FOR BAYESIAN FILTERING

We propose the following generative model for a continuous-time (latent) signal $(f(t))_{t \in \mathbb{R}}$ as a mixture of two components of the form

$$f(t) = f_l(t) + f_h(t), \quad (4)$$

where f_l is a signal of low-frequency content and f_h one of high-frequency content.

3.1. Assumptions over the spectral components f_l and f_h

We model f_l and f_h as independent GPs with covariance kernels denoted respectively by K_l and K_h , and, accordingly, power spectral densities S_l and S_h . To discriminate between higher and lower frequencies, we impose the following restrictions over S_l and S_h :

1. The support of S_l , denoted by $\text{supp}(S_l)$, is compact and centred around the origin, meaning that f_l is a process of low-frequency content.

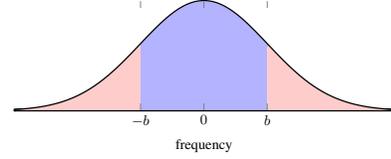


Fig. 1: Power spectral densities of the proposed model: The region contained inside the black line is the PSD of the process f , S_{SE} , whereas the regions in blue and red denote the PSDs of the low-frequency (S_l) and high-frequency (S_h) content respectively. Frequency b is the maximum frequency in the support of S_l and the minimum in the support of S_h (positive part).

2. The supports of the PSDs of f_l and f_h are non overlapping, that is, $\text{supp}(S_l) \cap \text{supp}(S_h) = \emptyset$. This implies that each frequency present in the signal f came, exclusively, from either f_l or f_h .
3. The sum of the component PSDs is a square-exponential kernel, that is, $S_l(\xi) + S_h(\xi) = S_{SE}(\xi)$ for some hyperparameters σ^2 and l as in eq. (2).

Notice that, as a consequence of the third restriction, the marginal distribution over the process f is a Gaussian process with an SE kernel $K_{SE} = K_l + K_h$ due to the linearity of the Fourier transform and the independence of f_l and f_h .

Figure 1 illustrates the PSDs of the components of high and low frequency. We have denoted by b the interface between the zones of low and high frequency, meaning that b is the highest frequency of the low-frequency signal and well as the lowest frequency of the high-frequency signal. Consequently, the bandwidth of the low-frequency part is $2b$.

3.2. Likelihood and model fitting

Assuming an independent sequence of Gaussian observation noise, the observations $(y(t))_{t \in \mathbb{R}}$ are then defined as

$$y(t) = f(t) + \eta(t), \quad \eta(t) \sim \mathcal{N}(0, \sigma_\eta^2), \quad (5)$$

Combining the observation model defined in eq. (5) with the GP-prior assumed for the spectral components, the marginal likelihood of the proposed model is Gaussian and therefore its hyperparameters can be obtained through minimisation of the negative log-likelihood (NLL). Notice that despite the elaborate frequency-wise construction of the latent process f through the non-overlapping spectra of the components f_l and f_h , the covariance kernel of f is square-exponential, thus allowing for straightforward model learning. Specifically, the NLL of the model is given by

$$\text{NLL}(\mathbf{y}|\mathbf{t}) = \log(2\pi|\Sigma_{\mathbf{y}}|) + \frac{1}{2}\mathbf{y}^\top \Sigma_{\mathbf{y}}^{-1} \mathbf{y}, \quad (6)$$

where $\mathbf{y} = [y_1, \dots, y_N]$ are the (noise corrupted and possibly missing) observations acquired at time instants $\mathbf{t} = [t_1, t_2, \dots, t_N]$, and $\Sigma_{\mathbf{y}}$ is the covariance matrix of \mathbf{y} defined by

$$\Sigma_{\mathbf{y}} = K_{SE}(\mathbf{t}, \mathbf{t}) + \sigma_\eta^2 \mathbf{I}, \quad (7)$$

therefore, the hyperparameters are those of the K_{SE} kernel and the noise variance σ_η^2 .

Finally, observe that the strict non-overlapping property of the components f_l and f_h is not problematic for training, in fact, the cutoff frequency b does not even appear for model training.

4. FILTERING AS POSTERIOR INFERENCE

Denote by b the required cut-off frequency of the low-pass filtering problem. Using the proposed model, we can assume that this cutoff frequency b is equal to the limit between the low- and high-frequency components. In this context, low-pass filtering problem is equivalent to performing inference over the low-frequency component f_l conditional to observations of the time series. Due to the assumptions made on the signal we refer to this approach as GP low-pass filter (GPLP).

Denoting the observations by $\mathbf{y} \in \mathbb{R}^n$, GPLP addresses low-pass filtering by computing the posterior distribution $p(f_l|\mathbf{y})$. Due to the self-conjugacy of the Gaussian distribution and its closure under additivity, this posterior is also a GP, with mean and covariance given by

$$m_{f_l|\mathbf{y}} = \Sigma_{f_l, \mathbf{y}} \Sigma_{\mathbf{y}}^{-1} \mathbf{y} \quad (8)$$

$$K_{f_l|\mathbf{y}} = K_{f_l} - \Sigma_{f_l, \mathbf{y}} \Sigma_{\mathbf{y}}^{-1} \Sigma_{f_l, \mathbf{y}}^\top, \quad (9)$$

where we have assumed zero mean for f_l and f_h (and therefore of y), $\Sigma_{\mathbf{y}}$ is the covariance of the observations defined in eq. (7), $\Sigma_{f_l, \mathbf{y}}$ denotes the covariance between f_l and \mathbf{y} , and K_l is the kernel of f_l .

Let us also note that the cross covariance $\Sigma_{f_l, \mathbf{y}}$ and the kernel K_l share the same expression. Denoting the covariance between the low-frequency process f_l at time t and the observation y at time t' by $\Sigma_{f_l, \mathbf{y}}(t, t')$, we obtain

$$\begin{aligned} \Sigma_{f_l, \mathbf{y}}(t, t') &= \mathbb{E}[f_l(t)(f_l(t') + f_h(t') + \eta(t'))] \\ &= \mathbb{E}[f_l(t)f_l(t')] \\ &= K_l(t, t'), \end{aligned}$$

since the processes f_l , f_h and η are independent Gaussian processes.

Therefore, the only critical quantity required to compute eqs. (8)-(9) is kernel K_l . Following the model proposed in eq. (4) and its assumptions, the PSD of f_l , denoted by S_l , can be obtained by multiplying the PSD of f with a rectangular function of width $2b$, that is,

$$S_l(\xi) = S_{SE}(\xi) \text{rect}\left(\frac{\xi}{2b}\right), \quad (10)$$

where we used the convention that $\text{rect}(\xi)$ is equal to one for $|\xi| < 1/2$ and 0 elsewhere. As a consequence, the kernel K_l can be calculated using the convolution theorem: (\star is the convolution operator)

$$\begin{aligned} K_l(t) &= \mathcal{F}^{-1}(S_l(\xi)) \\ &= \mathcal{F}^{-1}\left(S_{SE}(\xi) \text{rect}\left(\frac{\xi}{2b}\right)\right) \\ &= \mathcal{F}^{-1}(S_{SE}(\xi)) \star \mathcal{F}^{-1}\left(\text{rect}\left(\frac{\xi}{2b}\right)\right) \\ &= K_{SE}(t) \star \text{sinc}(2bt) \cdot 2b \\ &= 2b \cdot \int \sigma^2 \exp\left(-\frac{1}{2l^2}(t-\tau)^2\right) \frac{\sin(2\pi b\tau)}{2\pi b\tau} d\tau \\ &= \sigma^2 e^{-\frac{1}{2l^2}t^2} \Re\left(\text{erf}\left(\sqrt{2}bl\pi - i\frac{t}{\sqrt{2}l}\right)\right), \end{aligned} \quad (11)$$

where $\text{erf}(t)$ denotes the error function given by

$$\text{erf}(t) = \frac{1}{\sqrt{\pi}} \int_{-t}^t e^{-x^2} dx, \quad (12)$$

Using Taylor expansions, the error function can be calculated up to an arbitrary degree of accuracy [21].

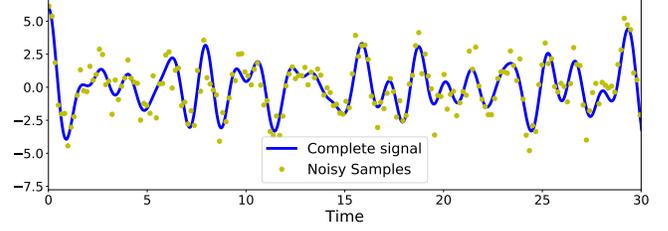


Fig. 2: The latent signal (blue) consists of 6 cosines as shown in eq. (13) and the observations are shown in yellow.

5. SIMULATIONS

The proposed model for Bayesian low-pass filtering using GPs, termed GPLP, is next validated using synthetic and real-world data. Our experimental validation aims to show that GPLP (i) successfully recovers low-frequency data from missing and noisy observation, (ii) provides accurate point-estimates with respect to the benchmarks, and (iii) places meaningful error bars. Our benchmarks include ground-truth signals and the Butterworth filter.

5.1. A synthetic time series with line spectra

We considered the line-spectra time series given by

$$f(t) = \sum_{\omega_i \in F_{\text{low}}} \cos(2\pi\omega_i t) + \sum_{\omega_j \in F_{\text{high}}} \cos(2\pi\omega_j t) \quad (13)$$

where the sets F_{low} and F_{high} are such that $\forall \omega_i \in F_{\text{low}}, \forall \omega_j \in F_{\text{high}} : \omega_i < \omega_j$. Simply put, F_{low} is a set of low frequencies and F_{high} a set of high frequencies—all these frequencies are in Hertz (Hz). Signals constructed in this way have sparse PSDs meaning that only a finite number of frequencies convey all the signal energy or information.

We chose $F_{\text{low}} = \{0.31, 0.38, 0.48\}$ and $F_{\text{high}} = \{0.51, 0.64, 0.75\}$ and simulated a path of $f(t)$, as defined in eq. (13) for 5000 evenly-spaced time indices in $t \in [-100, 100]$. The observation time-series \mathbf{y} consisted only in a 25% of the signal (again, evenly spaced) all of which were corrupted by Gaussian noise of std. dev. $\sigma_\eta = 1.0$. Fig. 2 shows the latent signal and the observation considered for this experiment.

We implemented the proposed GPLP to recover the low-frequency content of the original (latent) signal f only using the observations \mathbf{y} . We first trained the generative model as explained in Sec. 3.2 to find the hyper parameters l , σ^2 and σ_η^2 . We then chose the cutoff frequency to be $b = 0.495\text{Hz}$. We then computed the low-frequency covariance function to calculate the moments of the posterior distribution $p(f_h|\mathbf{y})$. Fig. 3 shows the learnt kernels and their corresponding PSDs. Notice how, just as illustrated in Fig. 1, the spectral densities of the latent low-frequency component is band-limited, supported only on $[-b, b]$ and tightly bounded by the (unfiltered) time series.

Fig. 4 shows the GPLP estimate compared against the ground truth and a low-pass version of the data using a Butterworth low-pass filter of order 10, with the same cutoff frequency; this filter is a standard in linear filtering. GPLP obtained a mean-squared error of 0.16 while the Butterworth low-pass filter gave a mean-squared error of 0.26, in addition to this marginal difference in performance, notice that GPLP provided accurate 95% error bars.

To further validate the ability of the proposed GPLP to filter out low-frequency spectral content, Fig. 5 shows the Fast Fourier Transform (FFT) of the low-pass versions of GPLP, Butterworth and the

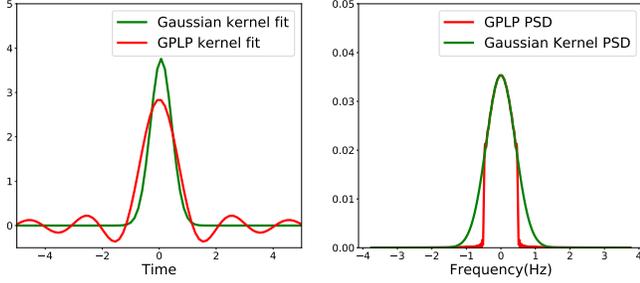


Fig. 3: Left: Learnt Gaussian kernel shown in green and the low-frequency kernel in red. Right: Learnt Gaussian PSD shown in green and low-frequency PSD in red.

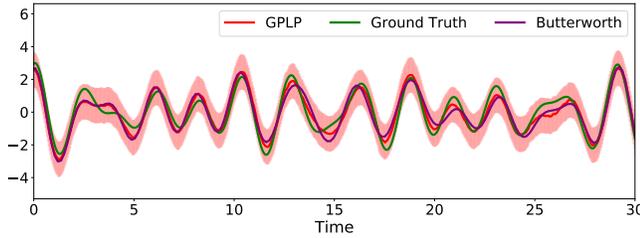


Fig. 4: Inference over the low-frequency component: posterior mean of the proposed GPLP (red), Butterworth low-pass filter (purple) and ground truth signal (green).

original signal. Notice from that GPLP successfully recovered the first three spectral components and rejected the higher ones.

5.2. Low-pass filtering of unevenly-sampled observations

A critical downside of the standard filtering techniques is that most of them require the data to be evenly spaced. Here is where the proposed model excels: as our method is based on an infinite-dimensional prior over continuous-time signals, missing observations are naturally handled by integrating out missing values.

We replicated the exact same setting as in Sec. 5.1 but considered **randomly-chosen** observations (Again, just 25% of the total number of points). Fig. 6 shows the posterior mean over the low-frequency component together with the 95% confidence interval and the ground truth, as well as the result in the frequency domain. The MSE of the GPLP estimate was 0.23, thus improving over Butterworth using evenly spaced data.

5.3. Filtering a heart-rate time series

We considered two 1800-sample heart-rate signals¹. One corresponding to a healthy subject and an unhealthy subject. The anomaly can be detected from the heart-rate signal by looking at the energy contained below 0.05 (Hz): if most of the energy is below this threshold, the subject is likely to suffer from congestive heart failure [22].

The aim of this experiment was to use GPLP to quantify the portion of energy below 0.05, as this reveals whether the signal corresponds to a healthy or unhealthy subject. We implemented GPLP

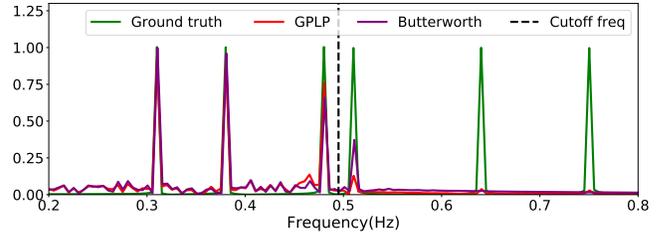


Fig. 5: FFT for full signal (green) and low-pass estimates: proposed GPLP (red) and Butterworth (purple). The cutoff frequency is shown in by a vertical dashed line.

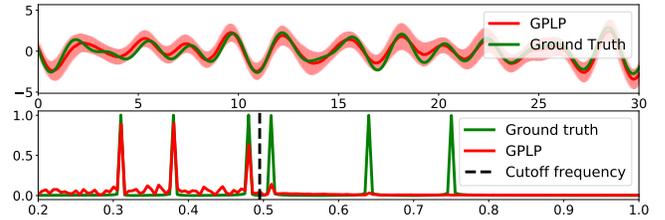


Fig. 6: Above: Low-pass filtering of unevenly-sampled observations. Proposed GPLP shown in red (mean and 95% error bars). Below: Same result, but shown in the frequency domain.

on both signals with a cutoff frequency of 0.05(Hz) and we found that the healthy signal has a 77% of its energy below 0.05(Hz) and that the unhealthy signal has a 97% of its energy below 0.05(Hz). Therefore, the GPLP method can discriminate between healthy and unhealthy subjects from the heart-rate signals.

6. DISCUSSION

We have proposed a Bayesian approach to low-pass filtering. The method is based on a latent-component generative model for time series, where the components are Gaussian processes with non-overlapping spectra. With this model, finding the low-pass version of a signal can be addressed from a Bayesian inference point of view. The main contribution over existing low-pass filters in the linear filter literature is that the proposed model offers an account of its own uncertainty and can naturally handle missing or noisy observations.

The proposed method has been validated empirically using synthetic and real-world data, where we have shown its ability to recover unbiased estimates of the true low-frequency signals (both in the evenly- and unevenly-sampled cases) and performed accurately with respect to its classical counterpart: the Butterworth filter.

7. ACKNOWLEDGEMENTS

This work was funded by projects Fondecyt-Iniciacin #11171165 and Conicyt-PIA #AFB170001 Center for Mathematical Modeling.

¹<http://ecg.mit.edu/time-series/>

8. REFERENCES

- [1] D.-H. Kwak, D.-H. Lee, J.-H. Ahn, and B.-H. Koh, "Fault detection of roller-bearings using signal processing and optimization algorithms," *Sensors (Basel, Switzerland)*, vol. 14, pp. 283–298, 12 2013.
- [2] F. Tobar, G. Rios, T. Valdivia, and P. Guerrero, "Recovering latent signals from a mixture of measurements using a Gaussian process prior," *IEEE Signal Processing Letters*, vol. 24, pp. 231–235, 2017.
- [3] S. E. Whitcomb and J. Keene, "Low-pass interference filters for submillimeter astronomy," *Appl. Opt.*, vol. 19, no. 2, pp. 197–198, Jan 1980.
- [4] L. Araya-Hernandez, J. F. Silva, A. Osses, and F. Tobar, "A Bayesian mixture-of-Gaussians model for astronomical observations in interferometry," in *Proc. of Chilecon*, 2017, pp. 1–5.
- [5] A. Cuevas, A. Veragua, S. EspaoJimenez, G. Chiang, and F. Tobar, "Unsupervised blue whale call detection using multiple time-frequency features," in *Proc. of Chilecon*, 2017, pp. 1–6.
- [6] Y. Katznelson, *An Introduction to Harmonic Analysis*, Cambridge Mathematical Library. Cambridge University Press, 3 edition, 2004.
- [7] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.
- [8] A. M. Yaglom, *Correlation Theory of Stationary and Related Random Functions*, Cambridge Mathematical Library. Springer-Verlag New York, 1 edition, 1987.
- [9] A. G. Wilson and R. P. Adams, "Gaussian process kernels for pattern discovery and extrapolation," in *Proc. International Conference on Machine Learning*, 2013.
- [10] G. Parra and F. Tobar, "Spectral Mixture Kernels for Multi-output Gaussian Processes," in *Neural Information Processing Systems*, 2017.
- [11] F. Tobar, T. Bui, and R. Turner, "Learning stationary time series using Gaussian processes with nonparametric kernels," in *Neural Information Processing Systems*, 2015, pp. 3483–3491.
- [12] F. Tobar, T. Bui, and R. Turner, "Design of covariance functions using inter-domain inducing variables," in *NIPS 2015 - Time Series Workshop*, December 2015.
- [13] M. Lázaro-Gredilla, J. Quiñero Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, "Sparse spectrum Gaussian process regression," *Journal of Machine Learning Research*, vol. 11, no. Jun, pp. 1865–1881, 2010.
- [14] M. Lázaro-Gredilla and A. Figueiras-Vidal, "Inter-domain Gaussian processes for sparse inference using inducing features," in *Advances in Neural Information Processing Systems* 22, pp. 1087–1095. 2009.
- [15] J. Hensman, N. Durrande, and A. Solin, "Variational fourier features for Gaussian processes," *Journal of Machine Learning Research*, vol. 18, no. 151, pp. 1–52, 2018.
- [16] N. Durrande, J. Hensman, M. Rattray, and N. D. Lawrence, "Detecting periodicities with Gaussian processes," *PeerJ Computer Science*, vol. 2, pp. e50, 2016.
- [17] Y. Wang, R. Khardon, and P. Protopapas, "Nonparametric Bayesian estimation of periodic light curves," *The Astrophysical Journal*, vol. 756, no. 1, pp. 67, 2012.
- [18] F. Tobar, "Bayesian nonparametric spectral estimation," in *Advances in Neural Information Processing Systems 31*, pp. 10148–10158. 2018.
- [19] J. Wu, M. Poloczek, A. G. Wilson, and P. Frazier, "Bayesian optimization with gradients," in *Advances in Neural Information Processing Systems 30*, pp. 5267–5278. 2017.
- [20] G. Rios and F. Tobar, "Learning non-Gaussian time series using the Box-Cox Gaussian process," in *Proc. of the International Joint Conference on Neural Networks*, 2018, pp. 1–8.
- [21] S. Chevillard, "The functions erf and erfc computed with arbitrary precision and explicit error bounds," *Information and Computation*, vol. 216, pp. 72 – 95, 2012, Special Issue: 8th Conference on Real Numbers and Computers.
- [22] L. Glass, P. Hunter, and A. McCulloch, *Theory of Heart: Biomechanics, Biophysics, and Nonlinear Dynamics of Cardiac Function*, 01 1991.