Efficient modeling of correlated noise

III. Scalable methods for jointly modeling several observables' time series with Gaussian processes

J.-B. Delisle[✉], N. Unger[✉], N. C. Hara^{★✉}, and D. Ségransan[✉]

Département d'Astronomie, Université de Genève, Chemin Pegasi 51, 1290 Versoix, Switzerland
 e-mail: jean-baptiste.delisle@unige.ch

Received 3 August 2021 / Accepted 15 December 2021

ABSTRACT

The radial velocity method is a very productive technique used to detect and confirm extrasolar planets. The most recent spectrographs, such as ESPRESSO or EXPRES, have the potential to detect Earth-like planets around Sun-like stars. However, stellar activity can induce radial velocity variations that dilute or even mimic the signature of a planet. A widely recognized method for disentangling these signals is to model the radial velocity time series, jointly with stellar activity indicators, using Gaussian processes and their derivatives. However, such modeling is prohibitive in terms of computational resources for large data sets, as the cost typically scales as the total number of measurements cubed. Here, we present S+LEAF 2, a Gaussian process framework that can be used to jointly model several time series, with a computational cost that scales linearly with the data set size. This framework thus provides a state-of-the-art Gaussian process model, with tractable computations even for large data sets. We illustrate the power of this framework by reanalyzing the 246 HARPS radial velocity measurements of the nearby K2 dwarf HD 13808, together with two activity indicators. We reproduce the results of a previous analysis of these data, but with a strongly decreased computational cost (more than two order of magnitude). The gain would be even greater for larger data sets.

Key words. methods: data analysis – methods: statistical – methods: analytical – planets and satellites: general

1. Introduction

It is common in astronomy to indirectly detect a physical event or the presence of a body by searching for its signature in a data set and, more specifically, in a time series. Astronomical time series are typically corrupted by photon noise, which is uncorrelated: the noise values at two distinct times are statistically independent. In that case, as more data are acquired, the searched-for signal should emerge more clearly. However, in many cases, the data are also corrupted by correlated noise emerging from other physical events, contamination from the Earth's atmosphere, instrumental noise, etc. In some cases, the structure of this correlated noise can mimic the signal of interest, leading to spurious detections or to a poor estimation of the model parameters.

This situation is encountered in particular when searching for exoplanets in radial velocity (RV) data. The RV of a star is the star velocity projected onto the line of sight, measured thanks to the Doppler effect. The presence of a planetary companion induces a reflex motion of the star and thus a periodic pattern in the RV time series. The latest generation of spectrographs, such as ESPRESSO (Pepe et al. 2021) or EXPRES (Blackman et al. 2020), is able to reach a RV precision of the order of 10 cm s^{-1} , and has the potential to discover Earth-like planets around Sun-like stars. However, correlated noise of stellar origin complicates this task. The p -modes and granulation processes introduce correlated noise at different timescales (Dumusque et al. 2011, 2012). Furthermore, the random appearance of spots and facu-

lae at the surface of the star, combined with the star's rotation, introduces complex structure in the data, which might be difficult to disentangle from low-mass planets. At longer timescales (from hundreds to thousands of days), the stellar magnetic cycle also induces RV variations, as well as variations in activity indicators such as the flux in the calcium II H & K emission lines ($\log R'_{\text{HK}}$; Noyes et al. 1984) as in, for instance, (Queloz et al. 2001). The RV signals induced by these multiple physical processes of stellar origin are globally referred to as stellar activity. A common way to account for stellar activity is to model it through a Gaussian process (GP) model (Rasmussen & Williams 2006). The GP regression method allows for the modeling of complex processes by parametrizing the covariance between the measurements instead of defining a deterministic model of the physical processes. For a GP $G(t)$ measured at times t_i and t_j , the values $G(t_i)$ and $G(t_j)$ are assumed to be randomly drawn from a normal distribution, with covariance $C_{i,j} = k(t_i, t_j)$, where k is the chosen parametrized kernel. The GP is often assumed to be stationary, such that the kernel k only depends on the lag $\Delta t = |t_i - t_j|$ between two measurements. A commonly used kernel to model stellar activity is (Aigrain et al. 2012; Haywood et al. 2014; Rajpaul et al. 2015):

$$k(\Delta t) = \sigma^2 \exp \left(-\frac{\Delta t^2}{2\rho^2} - \frac{\sin^2 \left(\frac{\pi \Delta t}{P} \right)}{2\eta^2} \right), \quad (1)$$

where σ , P , ρ , and η are the hyperparameters of the GP which need to be adjusted. In the following, we refer to this kernel as the squared-exponential periodic (SEP) kernel.

★ NCCR CHEOPS fellow.

Gaussian processes are known to be able to represent a wide range of signals. As such, when their hyperparameters are left free, they are prone to absorb the signal of interest (planetary signal) along with the correlated noise (stellar activity). To avoid this drawback, [Rajpaul et al. \(2015\)](#) proposed a framework in which the RV time series is modeled jointly with activity indicators. Building on [Aigrain et al. \(2012\)](#), the authors assume that the activity-induced variations of the measurements depend linearly on an underlying Gaussian process $G(t)$ and its derivative $G'(t)$. The evolution of the RV and indicators is modeled as:

$$\begin{aligned}\Delta RV &= V_c G(t) + V_r G'(t), \\ \Delta BIS &= B_c G(t) + B_r G'(t), \\ \Delta \log R'_{HK} &= L_c G(t),\end{aligned}\quad (2)$$

for some constants V_c, V_r, B_c, B_r, L_c . The GP's hyperparameters are thus constrained by the three time series, instead of the RV alone. This reduces the risk of the GP overfitting, that is, the absorption of planetary signals, since those signals are only present in the RV time series. This framework can be straightforwardly generalized to account for additional indicators, for the combination of several GP with different amplitudes, or even for the second order derivatives of the GP (e.g., [Jones et al. 2017](#)).

While this framework is very powerful in modeling stellar activity, it represents a challenge in terms of computational cost. Indeed, computing the likelihood (or χ^2) of the model for a given set of hyperparameters requires us to solve a linear system involving the full covariance matrix of the measurements. For a time series of size n , the full covariance matrix – including RV, BIS, and $\log R'_{HK}$ measurements – has a size of $3n \times 3n$, and the cost of the solving typically scales as $O((3n)^3)$. This becomes prohibitive in terms of computer time for large data sets, especially in the context of Bayesian methods (MCMC, nested sampling, etc.), which might require billions of evaluations of the likelihood.

In the context of a GP applied to a single time series, [Ambikasaran \(2015\)](#) and [Foreman-Mackey \(2018\)](#) (see also [Rybicki & Press 1995](#)) have shown that the so-called *celerite* kernel,

$$k(\Delta t) = \sum_{s < n_c} (a_s \cos(\nu_s \Delta t) + b_s \sin(\nu_s \Delta t)) e^{-\lambda_s \Delta t}, \quad (3)$$

where n_c is the number of components, and a_s, b_s, λ_s , and ν_s are the kernel hyperparameters, can be represented as a semiseparable matrix. As a consequence, the computational cost of evaluating the likelihood scales linearly with the number of points ($O(n)$), allowing to apply these methods to large data sets. [Delisle et al. \(2020b\)](#) defined a more general class of covariance matrices with a similar linear scaling of the cost: the S+LEAF matrix, which is the sum of a semiseparable matrix and a LEAF matrix. The LEAF component, which has non-zero elements close to the diagonal, is particularly adapted to represent calibration noise (see [Delisle et al. 2020b](#)). [Gordon et al. \(2020\)](#) extended the *celerite* model to the case of two-dimensional data sets. This applies, in particular, to the case of several parallel time series (e.g., RV, BIS, and $\log R'_{HK}$), with measurements taken at the same times. However, [Gordon et al. \(2020\)](#) do not discuss the treatment of the derivatives of the GP, and thus of models similar to the [Rajpaul et al. \(2015\)](#) model (Eq. (2)).

In this study, we extend the *celerite* and S+LEAF models to account for the case of several time series, with independent calendars, modeled as a linear combination of several GP and their derivatives. This allows us to apply models similar to

the model used by [Rajpaul et al. \(2015\)](#), but with a linear scaling of the evaluation cost of the likelihood. We call this new model S+LEAF 2, as it is a generalization of the S+LEAF model ([Delisle et al. 2020b](#)).

In Sect. 2, we recall the main properties of the *celerite* and S+LEAF models. We then show in Sect. 3 how to model the derivative of a *celerite* GP. In Sect. 4, we extend the model to the case of multiple time series. We illustrate the power of this framework by reanalyzing the HARPS RV of the nearby K2 dwarf HD 13808 in Sect. 5. Finally, we discuss our methods and results in Sect. 6. An open-source reference implementation of our algorithms as C library with python wrappers is publicly available¹.

2. The *celerite* and S+LEAF models for homogeneous time series

We consider a time series of measurements (t_i, y_i) ($i = 1, \dots, n$), which can be modeled by a deterministic component, a GP component, and measurement noise. In the case of radial velocities, the deterministic component encompasses the reflex motion due to companions, the systematic velocity of the system, instruments offsets, and so on. The GP might be used to model physical mechanisms that are too poorly understood or constrained to be included in the deterministic part. This is typically the case of stellar activity at different timescales (oscillations, granulation, rotation, magnetic cycles, etc.). Finally the noise component encompasses photon noise, calibration noise, and so on. The time series can thus be expressed as:

$$y_i = m(t_i) + G(t_i) + \epsilon_i, \quad (4)$$

where m is the deterministic part of the model, G is the GP, and ϵ the noise. These three components of the model might depend on a set of parameters θ . Assuming the noise to also be Gaussian (but not necessarily white), the log-likelihood function of a given set of parameters θ is

$$\begin{aligned}\ln \mathcal{L}(\theta) &= \ln p(y|\theta) \\ &= -\frac{1}{2} (y - m_\theta)^T C_\theta^{-1} (y - m_\theta) \\ &\quad - \frac{1}{2} \ln \det(2\pi C_\theta).\end{aligned}\quad (5)$$

C is the total covariance matrix of the time series which can be split as

$$C = K + \Sigma, \quad (6)$$

where K is the covariance of the GP G and Σ is the covariance of the noise

$$\begin{aligned}K_{i,j} &= \text{cov}(G(t_i), G(t_j)) = k(t_i, t_j), \\ \Sigma_{i,j} &= \text{cov}(\epsilon_i, \epsilon_j),\end{aligned}\quad (7)$$

with k the GP kernel function.

2.1. The *celerite* model

The *celerite* model proposed by [Foreman-Mackey et al. \(2017\)](#) allows for very efficient computations of this model and, in particular, the likelihood and its gradient with respect to θ ,

¹ <https://gitlab.unige.ch/jean-baptiste.delisle/spleaf>

in the case of white noise (diagonal covariance matrix $\Sigma = \text{diag}(\sigma^2)$) and assuming the kernel function k to follow Eq. (3). In this case, the covariance matrix C is semiseparable (see Foreman-Mackey et al. 2017):

$$C = \text{diag}(A + \sigma^2) + \text{tril}(UV^T) + \text{triu}(VU^T), \quad (8)$$

where $\text{diag}(A)$ is the diagonal matrix built from the vector A of size n , U and V are $n \times r$ matrices (with $r = 2n_c$ the rank of U and V) defined as:

$$\begin{aligned} A_i &= \sum_{s < n_c} a_s, \\ U_{i,s} &= e^{-\lambda_s t_i} (a_s \cos(\nu_s t_i) + b_s \sin(\nu_s t_i)), \\ U_{i,n_c+s} &= e^{-\lambda_s t_i} (a_s \sin(\nu_s t_i) - b_s \cos(\nu_s t_i)), \\ V_{i,s} &= e^{\lambda_s t_i} \cos(\nu_s t_i), \\ V_{i,n_c+s} &= e^{\lambda_s t_i} \sin(\nu_s t_i). \end{aligned} \quad (9)$$

This low-rank representation of the covariance matrix allows us to use very efficient dedicated algorithms (Cholesky decomposition, solving, dot product, determinant, see Foreman-Mackey et al. 2017; Foreman-Mackey 2018), which are particularly useful in computing the likelihood. The memory footprint of the *celerite* model scales as $O(nr)$, and the computational cost scales as $O(nr^2)$. For comparison, the memory footprint of a naive representation of the same covariance matrix scales as $O(n^2)$ and the computational costs typically scales as $O(n^3)$.

The *celerite* algorithm is actually not restricted to kernels of the form of Eq. (3), but it can be applied as long as the kernel admits a semiseparable representation following Eq. (8). In particular, the Matérn 3/2 and Matérn 5/2 kernels can be represented using a semiseparable representation of rank 2 and 3 respectively (see Appendix A.1).

2.2. The S+LEAF model

The S+LEAF model (Delisle et al. 2020b) extends the *celerite* model in the case of correlated measurement noise. While in the *celerite* model, the covariance matrix Σ is assumed to be diagonal, the S+LEAF model allows us to account for a more general class of noise models, where Σ can be represented by a LEAF matrix. Such LEAF matrices are sparse matrices where non-zero values are close to the diagonal (see Delisle et al. 2020b). This encompasses banded matrices, block-diagonal matrices, staircase matrices, and so on. In the context of radial velocities, LEAF matrices are especially useful to model calibration noise (see Delisle et al. 2020b). Indeed, to obtain precise RV measurements, the instrument must be calibrated periodically, typically once per night. Thus, several measurements taken with the same instrument during the same night share the same calibration noise, which introduces a block-diagonal contribution to the covariance matrix.

As in the case of the *celerite* model, the S+LEAF model allows for a sparse representation of the covariance matrix and for efficient dedicated algorithms to compute the likelihood and its gradient with respect to the parameters (Delisle et al. 2020b). The memory footprint of the S+LEAF model scales as $O((r + \bar{b})n)$ and its computational cost as $O((r^2 + r\bar{b} + \bar{b}^2)n)$, where $\bar{b}$ is the average band width of the LEAF component.

3. Derivative of a *celerite*/s+LEAF Gaussian process

We consider a GP $G(t)$ whose kernel function k is stationary and admits a semiseparable decomposition (Eq. (8)). The covariance matrix of G is thus given by:

$$\text{cov}(G(t_i), G(t_j)) = k(t_i - t_j) = U(t_i)V(t_j) = U_i V_j, \quad (10)$$

for $t_i \geq t_j$, namely, for the lower triangular part. By way of symmetry, we have

$$\text{cov}(G(t_i), G(t_j)) = k(t_j - t_i) = V_i U_j, \quad (11)$$

for $t_i < t_j$, namely, for the upper triangular part. Assuming G to be differentiable, the covariance between G and G' is given by the partial derivatives of the kernel function:

$$\begin{aligned} \text{cov}(G'(t_i), G(t_j)) &= \frac{\partial k(t_i, t_j)}{\partial t_i} = k'(t_i - t_j) = U'(t_i)V(t_j) = U'_i V_j, \\ \text{cov}(G(t_i), G'(t_j)) &= \frac{\partial k(t_i, t_j)}{\partial t_j} = -k'(t_i - t_j) = U_i V'_j, \end{aligned} \quad (12)$$

where we assume $t_i \geq t_j$, and the primes denote the differentiation with respect to time. For these two equations to be valid, the semiseparable representation (U, V) must verify:

$$U'V^T = -UV'^T. \quad (13)$$

This relation is actually linked with the stationarity of the kernel. Indeed, since

$$k(t, t + \Delta t) = k(\Delta t) = U(t)V(t + \Delta t) \quad (14)$$

does not depend on t , we have

$$\frac{\partial k(t, t + \Delta t)}{\partial t} = 0 = U'(t)V(t + \Delta t) + U(t)V'(t + \Delta t). \quad (15)$$

In addition to this stationarity condition, for the GP G to be differentiable the kernel must verify $k'(0) = 0$ (e.g., Rasmussen & Williams 2006), thus:

$$U'_i V_i = U_i V'_i = 0. \quad (16)$$

From Eq. (12), we deduce that the covariance matrix between G and G' admits several equivalent antisymmetric semiseparable representations

$$\begin{aligned} \text{cov}(G', G) &= \text{tril}(U'V^T) - \text{triu}(VU'^T) \\ &= -\text{tril}(UV'^T) + \text{triu}(V'U^T), \\ &= \text{tril}(U'V^T) + \text{triu}(V'U^T), \\ &\dots \end{aligned} \quad (17)$$

Similar representations can be deduced for $\text{cov}(G, G')$ by using the relation

$$\text{cov}(G, G') = \text{cov}(G', G)^T = -\text{cov}(G', G). \quad (18)$$

For the covariance matrix of G' itself, we have (for $t_i \geq t_j$)

$$\begin{aligned} \text{cov}(G'(t_i), G'(t_j)) &= \frac{\partial^2 k(t_i, t_j)}{\partial t_i \partial t_j} = -k''(t_i - t_j) \\ &= U'_i V'_j = -U''_i V_j = -U_i V''_j. \end{aligned} \quad (19)$$

Therefore, the covariance matrix of G' also admits several symmetric semiseparable representations:

$$\begin{aligned} \text{cov}(G', G') &= \text{diag}(B) + \text{tril}(U'V'^T) + \text{triu}(V'U'^T) \\ &= \text{diag}(B) - \text{tril}(U''V'^T) - \text{triu}(VU''^T) \\ &= \text{diag}(B) - \text{tril}(UV''^T) - \text{triu}(V''U^T) \\ &\dots \end{aligned} \quad (20)$$

with

$$B_i = U'_i V'_i = -U''_i V_i = -U_i V''_i. \quad (21)$$

Hereinafter, we use the following representations:

$$\begin{aligned} \text{cov}(G', G) &= \text{tril}(U'V^T) + \text{triu}(V'U^T), \\ \text{cov}(G, G') &= \text{tril}(UV'^T) + \text{triu}(VU'^T), \\ \text{cov}(G', G') &= \text{diag}(B) + \text{tril}(U'V'^T) + \text{triu}(V'U'^T). \end{aligned} \quad (22)$$

Applying this reasoning to the *celerite* kernel of Eq. (3), we find (as per Eq. (9)):

$$\begin{aligned} k'(\Delta t) &= \sum_{s < n_c} (a'_s \cos(\nu_s \Delta t) + b'_s \sin(\nu_s \Delta t)) e^{-\lambda_s \Delta t}, \\ -k''(\Delta t) &= \sum_{s < n_c} (a''_s \cos(\nu_s \Delta t) + b''_s \sin(\nu_s \Delta t)) e^{-\lambda_s \Delta t}, \\ U'_{i,s} &= e^{-\lambda_s t_i} (a'_s \cos(\nu_s t_i) + b'_s \sin(\nu_s t_i)), \\ U'_{i,n_c+s} &= e^{-\lambda_s t_i} (a'_s \sin(\nu_s t_i) - b'_s \cos(\nu_s t_i)), \\ V'_{i,s} &= e^{\lambda_s t_i} (\lambda_s \cos(\nu_s t_i) - \nu_s \sin(\nu_s t_i)), \\ V'_{i,n_c+s} &= e^{\lambda_s t_i} (\lambda_s \sin(\nu_s t_i) + \nu_s \cos(\nu_s t_i)), \end{aligned} \quad (23)$$

with

$$\begin{aligned} a'_s &= \nu_s b_s - \lambda_s a_s, \\ b'_s &= -\nu_s a_s - \lambda_s b_s, \\ a''_s &= (\nu_s^2 - \lambda_s^2) a_s + 2\lambda_s \nu_s b_s, \\ b''_s &= (\nu_s^2 - \lambda_s^2) b_s - 2\lambda_s \nu_s a_s. \end{aligned} \quad (24)$$

In this case, the differentiability condition of Eq. (16) can be rewritten as

$$k'(0) = 0 = \sum_{s < n_c} a'_s. \quad (25)$$

Thus the initial parameters must verify

$$\sum_{s < n_c} \nu_s b_s - \lambda_s a_s = 0, \quad (26)$$

for the GP to be differentiable. In particular, in the case of a kernel including a single *celerite* component ($n_c = 1$), one must verify $a'_s = \nu_s b_s - \lambda_s a_s = 0$ and Eq. (23) is simplified into

$$\begin{aligned} k'(\Delta t) &= b'_s \sin(\nu_s \Delta t) e^{-\lambda_s \Delta t}, \\ U'_{i,s} &= b'_s e^{-\lambda_s t_i} \sin(\nu_s t_i), \\ U'_{i,n_c+s} &= -b'_s e^{-\lambda_s t_i} \cos(\nu_s t_i). \end{aligned} \quad (27)$$

The SHO kernel, which is a particular type of *celerite* kernel that only depends on three free parameters, always verifies this differentiability condition. It thus provides a smoother GP than the general *celerite* kernel, which is why Foreman-Mackey et al. (2017) recommended its application.

In the case of the Matérn 3/2 and 5/2 kernels, similar semiseparable decompositions can be achieved for the derivatives (see Appendix A.2). More generally, Eq. (22) provides the semiseparable decomposition of the derivatives for any differentiable semiseparable kernel.

4. s+LEAF 2: Extending *celerite*/s+LEAF to heterogeneous time series

Following Rajpaul et al. (2015), we assume that the time series of the radial velocities and the different indicators follow

$$Y_{i,j} = f_i(T_{i,j}) + \sum_k (\alpha_{k,i} G_k(T_{i,j}) + \beta_{k,i} G'_k(T_{i,j})) + \epsilon_{i,j}, \quad (28)$$

where $(T_{1..}, Y_{1..})$ is the RV time series and $(T_{i..}, Y_{i..})$ are the indicators time series ($i > 1$), f_i is the determinist part of the model for the time series i , G_k are independent GP, and ϵ is the measurement noise (including photon noise, calibration noise, etc.).

The times and number of measurements need not be the same for all time series, which implies that T and Y are not necessarily matrices but collections of vectors of variable length. In the case of the model presented in Eq. (2), the activity indicators (BIS and $\log R'_{\text{HK}}$) are typically extracted from the same spectra as the RV time series and, thus, they share the same sampling. However, activity indicators can be extracted from other instruments or techniques. For instance, Haywood et al. (2014) trained a GP on the CoRoT light curve of CoRoT 7 to then use it to model the impact of stellar activity on the HARPS radial velocity time series of the same star. While both data sets were roughly contemporaneous, the time series did not share the same sampling. In such a case, we could define $(T_{1..}, Y_{1..})$ as the RV time series, and $(T_{2..}, Y_{2..})$ as the photometric time series (with $T_1 \neq T_2$), and model both time series jointly according to Eq. (28).

For the sake of readability, we consider in the following a single GP G , such that

$$Y_{i,j} = f_i(T_{i,j}) + \alpha_i G(T_{i,j}) + \beta_i G'(T_{i,j}) + \epsilon_{i,j}, \quad (29)$$

but the reasoning holds in the more general case of Eq. (28). The covariance matrix corresponding to Eq. (29) is

$$\begin{aligned} \text{cov}(Y_{i,j}, Y_{l,m}) &= \alpha_i \alpha_l \text{cov}(G(T_{i,j}), G(T_{l,m})) \\ &\quad + \beta_i \alpha_l \text{cov}(G'(T_{i,j}), G(T_{l,m})) \\ &\quad + \alpha_i \beta_l \text{cov}(G(T_{i,j}), G'(T_{l,m})) \\ &\quad + \beta_i \beta_l \text{cov}(G'(T_{i,j}), G'(T_{l,m})) \\ &\quad + \text{cov}(\epsilon_{i,j}, \epsilon_{l,k}). \end{aligned} \quad (30)$$

The full covariance matrix is a $n \times n$ matrix, where n is the total number of measurements (radial velocities and indicators). The cost of evaluating the corresponding likelihood – which requires to solve a linear system involving the covariance matrix and computing its determinant – typically scales as $\mathcal{O}(n^3)$. Therefore, a direct evaluation of the covariance matrix becomes rapidly prohibitive in terms of memory footprint and computing time.

4.1. Semiseparable representation of the model

In order to construct a semiseparable representation of the covariance matrix of Eq. (30) we merge all time vectors T_i into a single vector t , and all data vectors Y_i into a single vector y . We additionally sort the measurements by increasing time; thus, the measurements of the different time series are completely mixed in the merged time series (t, y) . For a measurement (t_k, y_k) we denote by $\mathcal{I}_k$ (index of the original time series the measurement belongs to) and $\mathcal{J}_k$ (index of the measurement in this original time series) the corresponding couple of indices of this measurement in T, Y . The model of Eq. (29) can be rewritten as

$$y_k = f_{\mathcal{I}_k}(t_k) + \alpha_k G(t_k) + \beta_k G'(t_k) + \epsilon_{\mathcal{I}_k, \mathcal{J}_k}, \quad (31)$$

with $\alpha_k = \alpha_{I_k}$ and $\beta_k = \beta_{I_k}$. The corresponding covariance matrix is (as per Eq. (30)):

$$\begin{aligned} C = \text{cov}(y, y) = & (\alpha\alpha^T) * \text{cov}(G(t), G(t)) \\ & + (\beta\alpha^T) * \text{cov}(G'(t), G(t)) \\ & + (\alpha\beta^T) * \text{cov}(G(t), G'(t)) \\ & + (\beta\beta^T) * \text{cov}(G'(t), G'(t)) \\ & + \Sigma, \end{aligned} \quad (32)$$

where Σ is the covariance matrix of the measurement noise

$$\Sigma_{k,l} = \text{cov}(\epsilon_{I_k \sim \mathcal{I}_k}, \epsilon_{I_l \sim \mathcal{I}_l}), \quad (33)$$

and $X * Y$ is the Hadamard (or element-wise) product

$$(X * Y)_{i,j} = X_{i,j} Y_{i,j}. \quad (34)$$

We now assume that the GP G can be modeled by a semiseparable kernel (see Eq. (8)). We can thus use the merged time vector t to compute the semiseparable representation of the covariance matrix of $G(t)$, $G'(t)$ according to Eqs. (8), (9), (22), and (23) (see also Appendix A), and we obtain

$$\begin{aligned} C = & (\alpha\alpha^T) * (\text{diag}(A) + \text{tril}(UV^T) + \text{triu}(VU^T)) \\ & + (\beta\alpha^T) * (\text{tril}(U'V^T) + \text{triu}(V'U^T)) \\ & + (\alpha\beta^T) * (\text{tril}(UV'^T) + \text{triu}(VU'^T)) \\ & + (\beta\beta^T) * (\text{diag}(B) + \text{tril}(U'V'^T) + \text{triu}(V'U'^T)) \\ & + \Sigma. \end{aligned} \quad (35)$$

The Hadamard product $(\alpha\beta^T) * M$ can also be rewritten as

$$(\alpha\beta^T) * M = \text{diag}(\alpha) M \text{diag}(\beta), \quad (36)$$

thus

$$\begin{aligned} (\alpha\alpha^T) * \text{diag}(A) &= \text{diag}(\alpha^2 * A), \\ (\alpha\beta^T) * \text{tril}(UV^T) &= \text{tril}((\alpha * U)(\beta * V)^T), \\ (\alpha\beta^T) * \text{triu}(UV^T) &= \text{triu}((\alpha * U)(\beta * V)^T), \end{aligned} \quad (37)$$

where α^2 is the element-wise square of α ($\alpha^2 = \alpha * \alpha$) and $\alpha * U$ is the element-wise product of each column of U by α

$$(\alpha * U)_{k,s} = \alpha_k U_{k,s}. \quad (38)$$

Using these relations, Eq. (35) can be rewritten as

$$\begin{aligned} C = & \text{diag}(\alpha^2 * A) \\ & + \text{tril}((\alpha * U)(\alpha * V)^T) + \text{triu}((\alpha * V)(\alpha * U)^T) \\ & + \text{tril}((\beta * U')(\alpha * V)^T) + \text{triu}((\beta * V')(\alpha * U)^T) \\ & + \text{tril}((\alpha * U)(\beta * V')^T) + \text{triu}((\alpha * V)(\beta * U')^T) \\ & + \text{diag}(\beta^2 * B) \\ & + \text{tril}((\beta * U')(\beta * V')^T) + \text{triu}((\beta * V')(\beta * U')^T) \\ & + \Sigma. \end{aligned} \quad (39)$$

Finally, the latter expression can be factorized to obtain the semiseparable representation

$$C = \text{diag}(\mathcal{A}) + \text{tril}(\mathcal{U}\mathcal{V}^T) + \text{triu}(\mathcal{V}\mathcal{U}^T) + \Sigma, \quad (40)$$

with

$$\begin{aligned} \mathcal{A} &= \alpha^2 * A + \beta^2 * B, \\ \mathcal{U} &= \alpha * U + \beta * U', \\ \mathcal{V} &= \alpha * V + \beta * V'. \end{aligned} \quad (41)$$

The rank of this semiseparable representation of C (number of columns in $\mathcal{U}$, $\mathcal{V}$) is thus the same as the rank of the underlying GP (number of columns in U , V). In the case of the more general model of Eq. (28), with several independent GP, it is straightforward to compute the semiseparable representation of the covariance matrix by vertical concatenation of the matrix $\mathcal{U}$ and $\mathcal{V}$ corresponding to each independent GP and the total rank is the sum of all the GP ranks.

Keeping the rank of the covariance matrix as low as possible allows us to significantly improve the performances of the method. Indeed, the cost of likelihood evaluations with a semiseparable covariance matrix of rank r scales as $O(nr^2)$. It is thus remarkable to note that introducing the derivative of the GP in the model, as well as different coefficients α and β for the different time series, does not increase the rank of the semiseparable representation of the covariance matrix. The factorization performed between Eqs. (39) and (40) is the key step that allows us to keep the same rank as the underlying GP. This factorization is achieved thanks to the specific choice of semiseparable representation introduced in Eq. (22). Indeed, instead of using (U, V, U', V') to represent the covariance matrix, one could use (U, V, U', U'') or (U, V, V', V'') (see Eqs. (17) and (20)). However, with these choices of representation, the covariance matrix would not factorize in the same way, the rank of its semiseparable representation would be twice the rank of the underlying GP, and the cost would thus quadruple.

4.2. Measurement noise and LEAF component

Using the covariance matrix representation of Eq. (40), the results described in Sect. 2 in the case of a homogeneous time series can be extended to the case of heterogeneous time series depending on several independent GP and their derivatives. The *celerite* algorithms can be applied in the case of purely white noise (Σ diagonal) and the S+LEAF algorithms can be applied to the more general case of close-to-diagonal correlated noise. The LEAF component of the S+LEAF model simply needs to be defined on the merged time series (t, y) .

4.3. Computational cost and memory footprint

The computational cost of our model and its memory footprint are the same as the underlying *celerite*/S+LEAF representation. Therefore, the memory footprint scales as $O((r + \bar{b})n)$ and its computational cost as $O((r^2 + r\bar{b} + \bar{b}^2)n)$, where n is the total number of measurements (including radial velocities and indicators), r is the rank of $\mathcal{U}$ and $\mathcal{V}$, $\bar{b}$ is the average band width of the LEAF component (considering the merged time series (t, y)). This is to be compared with a naive implementation of the same model which has a memory footprint in $O(n^2)$ and a computational cost in $O(n^3)$.

4.4. Overflows and preconditioning

As explained in Ambikasaran (2015), Foreman-Mackey et al. (2017), Delisle et al. (2020b), a naive computer implementation

of the semiseparable decomposition of Eq. (9) can lead to overflows and underflows due to the exponential terms in the definition of U and V . However, Foreman-Mackey et al. (2017) proposed a simple preconditioning method to circumvent this issue in the case of the *celerite* model, which is also valid in the case of the S+LEAF model (Delisle et al. 2020b). Instead of using directly the matrices U and V , one might use the matrices $\tilde{U}$, $\tilde{V}$, and ϕ defined as

$$\begin{aligned}\tilde{U}_{i,s} &= a_s \cos(\nu_s t_i) + b_s \sin(\nu_s t_i), \\ \tilde{U}_{i,n_c+s} &= a_s \sin(\nu_s t_i) - b_s \cos(\nu_s t_i), \\ \tilde{V}_{i,s} &= \cos(\nu_s t_i), \\ \tilde{V}_{i,n_c+s} &= \sin(\nu_s t_i), \\ \phi_{i,s} &= \phi_{i,n_c+s} = e^{-\lambda_s(t_{i+1}-t_i)},\end{aligned}\quad (42)$$

such that

$$U_{i,s} V_{j,s} = \tilde{U}_{i,s} \tilde{V}_{j,s} \prod_{k=j}^{i-1} \phi_{k,s}, \quad (43)$$

and all the *celerite*/S+LEAF algorithms can be adapted to use this representation (see Foreman-Mackey et al. 2017; Foreman-Mackey 2018; Delisle et al. 2020b). Similarly, we define the matrices $\tilde{U}'$, $\tilde{V}'$, $\tilde{\mathcal{U}}$, and $\tilde{\mathcal{V}}$ as

$$\begin{aligned}\tilde{U}'_{i,s} &= a'_s \cos(\nu_s t_i) + b'_s \sin(\nu_s t_i), \\ \tilde{U}'_{i,n_c+s} &= a'_s \sin(\nu_s t_i) - b'_s \cos(\nu_s t_i), \\ \tilde{V}'_{i,s} &= \lambda_s \cos(\nu_s t_i) - \nu_s \sin(\nu_s t_i), \\ \tilde{V}'_{i,n_c+s} &= \lambda_s \sin(\nu_s t_i) + \nu_s \cos(\nu_s t_i), \\ \tilde{\mathcal{U}} &= \alpha * \tilde{U} + \beta * \tilde{U}', \\ \tilde{\mathcal{V}} &= \alpha * \tilde{V} + \beta * \tilde{V}',\end{aligned}\quad (44)$$

which allow to apply the overflow-proof version of *celerite*/S+LEAF algorithms. The same preconditioning method can be applied in the case of the Matérn 3/2 and 5/2 kernels (see Appendix A.3).

4.5. Efficient computation of the gradient

In most applications, one needs to explore the parameter space either to maximize the likelihood using optimization algorithms or to obtain samples from the posterior distribution of parameters using Bayesian methods (MCMC, nested sampling, etc.). In both cases, many algorithms have been designed that make use of the gradient of the likelihood with respect to the parameters to improve the convergence efficiency. Following Foreman-Mackey (2018) and Delisle et al. (2020b), we deduce gradient backpropagation algorithms for all the operations used to compute the likelihood. While a detailed presentation of these backpropagation algorithms would be cumbersome, we refer the reader to Delisle et al. (2020b) Appendix B for the general idea of the method, and to the reference S+LEAF 2 implementation² for further details.

5. Application: Reanalysis of HD 13808

In this section, we apply our algorithms to reanalyze the RV time series of HD 13808. This K2V dwarf is known to harbor

² <https://gitlab.unige.ch/jean-baptiste.delisle/spleaf>

two planet candidates (see Mayor et al. 2011) recently published as confirmed planets by Ahner et al. (2021). In the latter study, the authors defined several alternative models of stellar activity and performed a Bayesian model comparison letting the number of planets vary. In addition to confirming the two candidates, Ahner et al. (2021) concluded that the best stellar activity model was a GP trained simultaneously on the RV, BIS, and $\log R'_{\text{HK}}$ time series, following Eq. (2), as proposed by Rajpaul et al. (2015). Modeling this GP required the authors to solve for a 738×738 linear system billions of times, which was very demanding in terms of computational resources. Here, we reanalyze the 246 HARPS RV measurements of HD 13808, together with the indicators time series, but modeling the GP using S+LEAF 2.

5.1. Choice of kernel

The SEP kernel (see Eq. (1)), which was used by Ahner et al. (2021) for their analysis of HD 13808, is not semiseparable and thus cannot be modeled with S+LEAF 2. However, other quasiperiodic kernels, such as the SHO kernel proposed by Foreman-Mackey et al. (2017), admit a semiseparable representation. Here, we aim to reproduce the main characteristics of the SEP kernel but using a semiseparable kernel. The SEP kernel is the product of a squared-exponential and the exponential of a sinusoidal (see Eq. (1)):

$$k(\Delta t) = \sigma^2 \exp\left(-\frac{\Delta t^2}{2\rho^2}\right) \exp\left(-\frac{\sin^2\left(\frac{\pi \Delta t}{P}\right)}{2\eta^2}\right). \quad (45)$$

The second part can be expanded as a power series, assuming $2\eta \gtrsim 1$

$$\begin{aligned}k(\Delta t) &= \sigma^2 \exp\left(-\frac{\Delta t^2}{2\rho^2}\right) \left(1 - \frac{\sin^2\left(\frac{\pi \Delta t}{P}\right)}{2\eta^2} + \frac{\sin^4\left(\frac{\pi \Delta t}{P}\right)}{8\eta^4} + O(\eta^{-6})\right) \\ &= \sigma^2 \exp\left(-\frac{\Delta t^2}{2\rho^2}\right) \frac{1 + f \cos(\nu \Delta t) + \frac{f^2}{4} \cos(2\nu \Delta t)}{1 + f + \frac{f^2}{4}} + O(f^3),\end{aligned}\quad (46)$$

with $f = (2\eta)^{-2}$ and $\nu = 2\pi/P$. This kernel thus introduces some correlation at the rotation period P , but also at the harmonics ($P/2$, $P/3$, etc.), with amplitudes decaying rapidly (scaling as η^{-2n} for the harmonics P/n). The squared-exponential part implies that the correlations vanish over long timescales ($\Delta t \gg \rho$). The squared-exponential kernel is not semiseparable, but the Matérn 1/2 (simple exponential decay), 3/2, and 5/2 kernels admit a semiseparable decomposition (see Appendix A), and offer a similar decay of the correlation over long timescales. The SEP kernel could thus be roughly approximated by:

$$k(\Delta t) = \sigma^2 \exp\left(-\frac{\Delta t}{\rho}\right) \frac{1 + f \cos(\nu \Delta t) + \frac{f^2}{4} \cos(2\nu \Delta t)}{1 + f + \frac{f^2}{4}}. \quad (47)$$

However, a GP following this kernel would not be differentiable ($k'(0) = -\sigma^2/\rho \neq 0$). In order to ensure differentiability, we introduce a modified kernel, which is the combination of a Matérn 3/2 kernel and two underdamped SHO terms

$$k(\Delta t) = \sigma^2 \frac{k_{3/2}(\Delta t) + f k_{\text{SHO, fund.}}(\Delta t) + \frac{f^2}{4} k_{\text{SHO, harm.}}(\Delta t)}{1 + f + \frac{f^2}{4}}, \quad (48)$$

where

$$\begin{aligned}
 k_{3/2}(\Delta t) &= \exp\left(-\frac{\sqrt{3}\Delta t}{\rho}\right)\left(1 + \frac{\sqrt{3}\Delta t}{\rho}\right), \\
 k_{\text{SHO, fund.}}(\Delta t) &= \exp\left(-\frac{\Delta t}{\rho}\right)\left(\cos(\nu\Delta t) + \frac{1}{\nu\rho}\sin(\nu\Delta t)\right), \\
 k_{\text{SHO, harm.}}(\Delta t) &= \exp\left(-\frac{\Delta t}{\rho}\right)\left(\cos(2\nu\Delta t) + \frac{1}{2\nu\rho}\sin(2\nu\Delta t)\right). \quad (49)
 \end{aligned}$$

The kernel of Eq. (48), which we refer to as the Matérn 3/2 exponential periodic (MEP) kernel in the following, is differentiable and presents the main characteristics of the SEP kernel, while being semiseparable. The semiseparable representation of the MEP kernel is of rank $r = 6$. We note that similar kernels have already been used in the literature to model stellar activity in photometric time series (e.g., David et al. 2019; Gillen et al. 2020).

It should be noted that the MEP kernel is once mean square differentiable but not twice, which means that G' is well defined but it is not itself differentiable. Since the time series (RV and activity indicators) are modeled as combinations of G and G' , we could require G' to be differentiable to obtain a smoother model. Such a twice mean square differentiable kernel should satisfy $k^{(3)}(0) = 0$, in addition to the mandatory differentiability condition ($k'(0) = 0$). For the MEP kernel, $k_{\text{MEP}}^{(3)}(0)$ is non-zero, however, in practice this kernel seems to produce smooth time series (see Sect. 5.3). Nevertheless, it is possible to design semiseparable kernels that are rigorously twice differentiable. For instance, the Matérn 5/2 kernel is twice differentiable and semiseparable with a rank of 3 (see Appendix A). In Appendix B we present the exponential-sine (ES) and the exponential-sine periodic (ESP) kernels. Both kernels are twice differentiable and semiseparable. The ES kernel is of rank 3 and closely resemble the SE kernel, while the ESP kernel is of rank 15 and approximates the SEP kernel very well. We find very similar results when using the ESP kernel instead of the MEP kernel, while the cost of likelihood evaluations is roughly doubled because of the higher rank of the ESP kernel (see Appendix B). In the following, we thus adopt the MEP kernel as a replacement for the SEP kernel, to reduce the computational cost and since it produces smooth time series, at least in our case. In the general case, the ESP kernel would typically generate smoother time series than the MEP, but with a doubled cost.

Following Rajpaul et al. (2015) and Ahrer et al. (2021) we use the GP to model simultaneously the RV, BIS, and $\log R'_{\text{HK}}$ time series of HD 13808 according to (see Eq. (2)):

$$\begin{aligned}
 \Delta \text{RV} &= \alpha_{\text{RV}} G(t) + \beta_{\text{RV}} G'(t), \\
 \Delta \text{BIS} &= \alpha_{\text{BIS}} G(t) + \beta_{\text{BIS}} G'(t), \\
 \Delta \log R'_{\text{HK}} &= \alpha_{\log R'_{\text{HK}}} G(t), \quad (50)
 \end{aligned}$$

where the coefficients α, β and the kernel's hyperparameters (P, ρ, η) need to be determined.

5.2. Performances

To evaluate the performances of S+LEAF 2 for the model of Eqs. (48) and (50), we generated random RV, BIS, and $\log R'_{\text{HK}}$ time series and record the cost of the likelihood evaluation as a function of the number of generated measurements. We compared S+LEAF 2 performances using the MEP kernel with the computational cost of evaluating the likelihood using the full covariance matrix with the SEP kernel. The two implementations

Fig. 1. Cost of a likelihood evaluation as a function of the total number of measurements using S+LEAF 2 or the full covariance matrix (see Sect. 5.2).

Fig. 2. Precision of the linear solving operation as a function of the total number of measurements using S+LEAF 2 or the full covariance matrix (see Sect. 5.2).

were run on the same computer, using a single core. The results are shown in Fig. 1 and confirm the $O(n)$ scaling of S+LEAF 2, while the implementation of the full covariance matrix indeed scales as $O(n^3)$ for a large value of n . In the case of HD 13808, the total number of measurements is 738 (3×246), and based on Fig. 1, we obtain a gain of a factor ~ 130 in computing time by using S+LEAF 2 instead of the naive implementation. For larger data sets, the gain would be even greater.

In addition to these performances tests, we also ran numerical precision tests. We used the same GP model as for the performances tests and assessed the stability of S+LEAF 2 by computing $CC^{-1}x$, that is, applying the solving and dot product algorithms on a random merged time series x . In theory, we should find $CC^{-1}x = x$, however, due to the limited machine precision, numerical errors accumulate and the results stray slightly from x . We computed the root mean square (rms) of these errors, for the S+LEAF 2 methods and for the full covariance matrix (see Fig. 2). In both cases, numerical errors grow with the total number of measurements. However, the level and growth rate are lower when using S+LEAF 2. For a given number of measurements, the number of arithmetic operations is lower when using S+LEAF 2, which could explain the improvement in precision.

5.3. Periodogram and false alarm probability

We analyzed HD 13808 data using the periodogram and false alarm probability (FAP) approach presented in

Fig. 3. Periodograms of the raw RV time series of HD 13808 (*top*) as well as of the residuals after subtracting the 14.19 d (*center*) and the 53.7 d (*bottom*) planets.

Delisle et al. (2020a). For each frequency, all the linear parameters (here the RV and indicators offsets) are re-adjusted together with the amplitudes of the cosine and sine at the considered frequency (only applied to the RV time series). The framework of Delisle et al. (2020a) only requires slight adaptations to account for the joint fit of RV and activity indicators, which are detailed in Appendix C.

We first performed a fit of a base model including the offsets γ_{RV} , γ_{BIS} , $\gamma_{\log R'_{HK}}$, jitter terms added in quadrature to the measurements errorbars, σ_{RV} , σ_{BIS} , $\sigma_{\log R'_{HK}}$, the kernels hyperparameters, P_{GP} , ρ_{GP} , η_{GP} , and the amplitudes, α , β , by maximizing the likelihood. Then using the fitted noise parameters, we computed a first periodogram, reajusting the offsets for each considered frequency. The resulting periodogram is plotted in Fig. 3 (top). We observe a very significant peak at 14.19 d (FAP = 6.7×10^{-22}). We then fitted a Keplerian orbit to this signal, and reajusted all parameters. A second periodogram (Fig. 3, middle) was then computed on the residuals, still reajusting the offsets, but keeping the planetary and noise parameters fixed. A second significant peak is visible at 53.7 d (FAP = 7.7×10^{-12}). As for the first planet, we performed a global fit of all the parameters after including this planet in the model. These fitted parameters are presented in Table 1, the kernel function corresponding to the fitted hyperparameters is illustrated in Fig. 5, and the residuals of each time series superimposed with the GP prediction are shown in Figs. 6 and 7. Finally, the periodogram of the residuals after fitting both planets (Fig. 3, bottom) does not show any significant peak (FAP above 1%). Our conclusions, based on this periodogram and FAP approach, thus agree with the findings of Ahner et al. (2021), with two confirmed planets at 14.19 d and 53.7 d.

We note that the first harmonics part ($k_{SHO, harm.}$) in the MEP kernel (Eq. (48)) plays a key role in the GP modeling, even if its amplitude is significantly smaller than the fundamental term ($k_{SHO, fund.}$). Indeed, when neglecting the harmonics term (see Fig. 4), we find a third highly significant signal around 19 d (FAP = 4.3×10^{-5}), which corresponds to half the period of

Table 1. Maximum likelihood solution and POLYCHORD posterior for the model with a GP and two planets (at 14.19 and 53.7 d).

Parameter	(units)	Maximum likelihood	POLYCHORD posterior
P_{GP}	(d)	38.171 ± 0.246	$38.162^{+0.159}_{-0.149}$
ρ_{GP}	(d)	279.6 ± 59.0	$308.64^{+54.5}_{-57.5}$
η_{GP}		0.978 ± 0.125	$1.053^{+0.146}_{-0.131}$
γ_{RV}	(ms^{-1})	41095.100 ± 0.419	$41095.132^{+0.507}_{-0.516}$
σ_{RV}	(ms^{-1})	1.700 ± 0.113	$1.815^{+0.123}_{-0.117}$
α_{RV}	(ms^{-1})	1.157 ± 0.198	$1.373^{+0.274}_{-0.221}$
β_{RV}		20.55 ± 4.64	$22.123^{+5.64}_{-5.00}$
γ_{BIS}	(ms^{-1})	5.605 ± 0.939	$5.583^{+1.12}_{-1.11}$
σ_{BIS}	(ms^{-1})	1.830 ± 0.140	$1.775^{+0.159}_{-0.149}$
α_{BIS}	(ms^{-1})	2.699 ± 0.408	$3.104^{+0.484}_{-0.428}$
β_{BIS}		-27.08 ± 5.31	$-31.597^{+5.80}_{-6.70}$
$\gamma_{\log R'_{HK}}$		-4.8774 ± 0.0181	$-4.8780^{+0.0215}_{-0.0218}$
$\sigma_{\log R'_{HK}}$		0.007374 ± 0.000759	$0.007751^{+0.000856}_{-0.000804}$
$\alpha_{\log R'_{HK}}$		0.05273 ± 0.00767	$0.06050^{+0.00907}_{-0.00803}$
P_b	(d)	14.18538 ± 0.00199	$14.18657^{+0.00218}_{-0.00224}$
$\lambda_{0,b}$	(deg)	-49.55 ± 3.30	$-47.59^{+3.65}_{-3.71}$
K_b	(ms^{-1})	3.669 ± 0.194	$3.677^{+0.207}_{-0.206}$
e_b		0.0759 ± 0.0544	$0.0441^{+0.0477}_{-0.0304}$
ω_b	(deg)	213.6 ± 34.4	$224.1^{+73.2}_{-107}$
P_c	(d)	53.6956 ± 0.0437	$53.7126^{+0.0533}_{-0.0514}$
$\lambda_{0,c}$	(deg)	-8.64 ± 6.04	$-6.561^{+6.54}_{-6.65}$
K_c	(ms^{-1})	2.027 ± 0.191	$1.948^{+0.199}_{-0.202}$
e_c		0.165 ± 0.117	$0.0929^{+0.101}_{-0.0648}$
ω_c	(deg)	280.7 ± 32.0	$251.6^{+60.2}_{-137}$

Notes. The reference epoch is 2455 000 BJD. The 1- σ uncertainties provided together with the maximum likelihood solution are derived from the inverse of the Hessian matrix of the log-likelihood. The POLYCHORD posterior values and uncertainties correspond to the median and the 15.865 and 84.135 percentiles.

the GP ($P_{GP} \approx 38$ d, see Table 1). On the contrary, when including the harmonics part, the peak around 19 d is no more significant (see Fig. 3). This is not surprising since stellar activity is expected to introduce signals in the RV and indicators time series at the rotation period, but also at its harmonics, and, in particular, the first harmonics.

By design, a kernel which has power at the rotation harmonics will have a tendency to absorb signals at this period (here 19 d), even if those are not due to stellar activity. To further test whether the 19 d signal could be due to a planet, Hara et al. (2022a) checked the consistency of this signal over the timespan of the data set. The authors showed that the 19 d signal, while not statistically significant, exhibits a stable presence across the data set, which might motivate further observations.

5.4. Bayesian framework and false inclusion probability

In order to more directly compare our results with the study of Ahner et al. (2021), we performed a full Bayesian evidence analysis using the nested sampling algorithm POLYCHORD (Handley et al. 2015). Recently, POLYCHORD was used for radial velocity exoplanet detection by Ahner et al. (2021), Rajpaul et al. (2021), and Unger et al. (2021). We aimed at reproducing a similar analysis as Ahner et al. (2021), but using S+LEAF 2 and the GP model detailed in Sect. 5.1.

Fig. 4. Same as Fig. 3 but neglecting the harmonics component of the MEP kernel (see Eq. (48)).

Fig. 5. Kernel function used to model HD 13808’s activity (MEP, see Eq. (48)). The GP hyperparameters are taken from the best fit of the two planets solution (Table 1). For comparison, the SEP kernel, which the MEP kernel is design to roughly mimic, is also plotted using the same set of hyperparameters.

The prior distributions we used for all parameters are detailed in Table 2. We kept mostly the same priors used by [Ahrer et al. \(2021\)](#), with only a few changes. We set the amplitude for the RV term of the GP to be strictly positive to avoid degeneracies in the sign of the amplitudes. We also shifted the prior for the mean longitudes from $[0, 2\pi]$ to $[-\pi, \pi]$. Indeed, the mean longitude of planet c is very close to 0, and this shift avoids splitting the peak in the posterior, thus improving the convergence efficiency. For all runs we used a number of live points equal to $50n_{\text{dim}}$, that is, 50 times the number of free parameters. For the precision criterion (stopping criterion), we used 10^{-9} .

We ran POLYCHORD for models with zero up to three planets and each model was run five times to obtain estimates of the value and uncertainty of the evidence. The evidence for each model is presented in Table 3. We see a steady and significant increase in evidence up to the model with two planets. The two and three planet models present similar evidence

Fig. 6. GP prediction (conditional distribution) from the best fit of the two planets solution (Table 1). The prediction is plotted for the GP and its derivative (*top*) and the full GP prediction for the RV, BIS, and $\log R'_{\text{HK}}$ time series superimposed with the corresponding residuals (*bottom three plots*).

($\Delta \ln Z = 3.5 \pm 2.5$, compatible at 1.4σ). Moreover, even if the three planet model is marginally favored, no clear period for a third planet emerges from the posteriors. To illustrate this, we computed the false inclusion probability (FIP) periodogram ([Hara et al. 2022b](#)) from the posteriors of all our runs with zero to three planets (see Fig. 8). The true inclusion probability (TIP) provides the probability for the system to host at least one planet in a given (small) range of period. On the contrary, the FIP ($=1-\text{TIP}$) is the probability that the system does not host any planet in the given range of period. The FIP periodogram of Fig. 8 is computed with a window size of $1/(t_{\text{max}} - t_{\text{min}})$

Fig. 7. Zoom of Fig. 6 around epoch 2 453 720 BJD.

in frequency. We observe in Fig. 8 two very significant peaks (FIP less than 10^{-8}) around 14.1 d and 53.7 d, which confirms that these planets should be included in the model. Then, two smaller peaks around 12 d and 19 d are also visible but none of them is significant (FIP higher than 70%). The posterior distribution of parameters of the two-planet runs is given in Table 1. Our results are in agreement with the periodogram and FAP approach detailed in Sect. 5.3, and with [Ahner et al. \(2021\)](#) findings.

6. Conclusion

In this article, we have presented S+LEAF 2, a GP framework that is able to efficiently model multiple time series simultaneously. Classical GP models have a computational cost that scales as

Table 2. Prior distributions used for each parameter in the nested sampling runs with POLYCHORD.

Parameter	(units)	Prior	Lower bound	Upper bound
P_{GP}	(d)	$\mathcal{U}$	10	100
ρ_{GP}	(d)	$\log \mathcal{U}$	10	400
η_{GP}		$\log \mathcal{U}$	0.01	10
γ_{RV}	(ms ⁻¹)	$\mathcal{U}$	$\text{RV}_{\min}$	$\text{RV}_{\max}$
σ_{RV}	(ms ⁻¹)	$\mathcal{U}$	0	20
α_{RV}	(ms ⁻¹)	$\mathcal{U}$	0	rms(RV)
β_{RV}		$\mathcal{U}$	-10 rms(RV)	10 rms(RV)
γ_{BIS}	(ms ⁻¹)	$\mathcal{U}$	$\text{BIS}_{\min}$	$\text{BIS}_{\max}$
σ_{BIS}	(ms ⁻¹)	$\mathcal{U}$	0	rms(BIS)
α_{BIS}	(ms ⁻¹)	$\mathcal{U}$	-rms(BIS)	rms(BIS)
β_{BIS}		$\mathcal{U}$	-10 rms(BIS)	10 rms(BIS)
$\gamma_{\log R'_{\text{HK}}}$		$\mathcal{U}$	$\log R'_{\text{HK min}}$	$\log R'_{\text{HK max}}$
$\sigma_{\log R'_{\text{HK}}}$		$\mathcal{U}$	0	rms(log R'_{HK})
$\alpha_{\log R'_{\text{HK}}}$		$\mathcal{U}$	-rms(log R'_{HK})	rms(log R'_{HK})
P	(d)	$\log \mathcal{U}$	5	100
λ_0	(deg)	$\mathcal{U}$	$-\pi$	π
K	(ms ⁻¹)	$\log \mathcal{U}$	0.1	10
e		$\mathcal{U}$	0	1
ω	(deg)	$\mathcal{U}$	0	2π

Notes. $\mathcal{U}$ stands for uniform distribution, while $\log \mathcal{U}$ stands for log-uniform distribution.

Table 3. Evidence of each considered model in our POLYCHORD runs.

Model	$\log Z$	$\log Z_i - \log Z_2$
0 planet	-557.23 ± 0.17	-88.73 ± 1.4
1 planet	-497.3 ± 1.1	-28.8 ± 1.8
2 planets	-468.5 ± 1.4	0.0
3 planets	-465.0 ± 2.1	$+3.5 \pm 2.5$

Notes. Uncertainties are calculated as the standard deviation of five identical nested sampling runs for each model.

Fig. 8. FIP periodogram for HD 13808. In blue we represent the FIP (false inclusion probability) and in yellow the TIP (true inclusion probability).

the cube of the number of measurements, which makes them prohibitive in terms of computational resources for large data sets. The computational cost of our GP framework scales linearly with the data set size, which allows for tractable computations even for large data sets.

This work builds on previous studies that provided efficient GP models for single time series (in particular the *celerite* and S+LEAF models, see [Rybizki & Press 1995](#); [Ambikasaran 2015](#); [Foreman-Mackey et al. 2017](#); [Delisle et al. 2020b](#)). It is also inspired by a recent generalization of *celerite* to the case

of two-dimensional data sets (Gordon et al. 2020), but extends it by accounting for the GP derivatives. These derivatives are especially important when modeling the effect of stellar activity on RV time series (see Aigrain et al. 2012; Rajpaul et al. 2015). Our framework additionally accounts for time series that do not share the same calendar, which is useful to train a GP simultaneously on RV and photometric measurements taken with two different instruments (e.g., Haywood et al. 2014).

We applied our methods to reanalyze the RV time series of the nearby K2 dwarf HD 13808. Our results are very similar to a recent state-of-the-art study of the same system (Ahrer et al. 2021) and we confirm the two planets announced in this article. However, we have shown that using our framework allowed us to dramatically decrease (by more than two orders of magnitude) the computational cost of the GP modeling. While the data set analyzed here consists of 738 measurements (RV, BIS, and $\log R'_{\text{HK}}$ at 246 epochs), the gain of using S+LEAF 2 would be even greater for larger data sets. Some data sets (e.g., HARPS or HARPS-N Sun-as-a-star RV time series) that could not have been analyzed with such a GP modeling are now achievable with our GP framework.

Finally, it is worth noting that the results from the periodogram and FAP approach are consistent with the much more computer intensive Bayesian evidence calculations using nested sampling. This illustrates the power of the periodogram and FAP computation including a correlated noise model, as proposed by Delisle et al. (2020a).

Acknowledgements. We thank the referee, D. Foreman-Mackey, for his very constructive feedback that helped to improve this manuscript. We acknowledge financial support from the Swiss National Science Foundation (SNSF). This work has, in part, been carried out within the framework of the National Centre for Competence in Research PlanetS supported by SNSF.

References

- Ahrer, E., Queloz, D., Rajpaul, V. M., et al. 2021, *MNRAS*, **503**, 1248
Aigrain, S., Pont, F., & Zucker, S. 2012, *MNRAS*, **419**, 3147
Ambikasaran, S. 2015, *Numer. Linear Algebra Appl.*, **22**, 1102
Baluev, R. V. 2008, *MNRAS*, **385**, 1279
Blackman, R. T., Fischer, D. A., Jurgenson, C. A., et al. 2020, *AJ*, **159**, 238
David, T. J., Petigura, E. A., Luger, R., et al. 2019, *ApJ*, **885**, L12
Delisle, J. B., Hara, N., & Ségransan, D. 2020a, *A&A*, **635**, A83
Delisle, J. B., Hara, N., & Ségransan, D. 2020b, *A&A*, **638**, A95
Dumusque, X., Udry, S., Lovis, C., Santos, N. C., & Monteiro, M. J. P. F. G. 2011, *A&A*, **525**, A140
Dumusque, X., Pepe, F., Lovis, C., et al. 2012, *Nature*, **491**, 207
Foreman-Mackey, D. 2018, *Res. Notes Am. Astron. Soc.*, **2**, 31
Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, *AJ*, **154**, 220
Gillen, E., Briegal, J. T., Hodgkin, S. T., et al. 2020, *MNRAS*, **492**, 1008
Gordon, T. A., Agol, E., & Foreman-Mackey, D. 2020, *AJ*, **160**, 240
Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015, *MNRAS*, **453**, 4384
Hara, N. C., Delisle, J.-B., Unger, N., & Dumusque, X. 2022a, *A&A*, **658**, A177
Hara, N. C., Unger, N., Delisle, J.-B., Díaz, R., & Ségransan, D. 2022b, *A&A*, accepted, <https://doi.org/10.1051/0004-6361/202140543>
Haywood, R. D., Collier Cameron, A., Queloz, D., et al. 2014, *MNRAS*, **443**, 2517
Jones, D. E., Stenning, D. C., Ford, E. B., et al. 2017, ArXiv e-prints [arXiv:1711.01318]
Jordán, A., Eyheramendy, S., & Buchner, J. 2021, *Res. Notes Am. Astron. Soc.*, **5**, 107
Mayor, M., Marmier, M., Lovis, C., et al. 2011, ArXiv e-prints [arXiv:1109.2497]
Noyes, R. W., Hartmann, L. W., Baliunas, S. L., Duncan, D. K., & Vaughan, A. H. 1984, *ApJ*, **279**, 763
Pepe, F., Cristiani, S., Rebolo, R., et al. 2021, *A&A*, **645**, A96
Queloz, D., Henry, G. W., Sivan, J. P., et al. 2001, *A&A*, **379**, 279
Rajpaul, V., Aigrain, S., Osborne, M. A., Reece, S., & Roberts, S. 2015, *MNRAS*, **452**, 2269
Rajpaul, V. M., Buchhave, L. A., Lacedelli, G., et al. 2021, *MNRAS*, **507**, 1847
Rasmussen, C. E., & Williams, C. K. I. 2006, *Gaussian Processes for Machine Learning* (MIT Press)
Rybicki, G. B., & Press, W. H. 1995, *Phys. Rev. Lett.*, **74**, 1060
Unger, N., Ségransan, D., Queloz, D., et al. 2021, *A&A*, **654**, A104

Appendix A: Semiseparable representation of Matérn 3/2 and Matérn 5/2 kernels

The Matérn 3/2 and 5/2 covariances are widely used in various fields of statistics. Their kernel functions are written as

$$\begin{aligned} k_{3/2}(\Delta t) &= \sigma^2 (1 + \Delta x) \mathbf{e}^{-\Delta x}, \\ k_{5/2}(\Delta t) &= \sigma^2 \left(1 + \Delta x + \frac{1}{3} \Delta x^2 \right) \mathbf{e}^{-\Delta x}, \end{aligned} \quad (\text{A.1})$$

where x is the rescaled time:

$$\begin{aligned} x &= \sqrt{3} \frac{t}{\rho} && \text{for the Matérn 3/2 kernel,} \\ x &= \sqrt{5} \frac{t}{\rho} && \text{for the Matérn 5/2 kernel,} \end{aligned} \quad (\text{A.2})$$

and σ and ρ are the GP hyperparameters.

A.1. Semiseparable representation

For two times $t_i > t_j$, we have

$$k_{3/2}(t_i - t_j) = \sigma^2 \left((x_i \mathbf{e}^{-x_i}) \mathbf{e}^{x_j} + \mathbf{e}^{-x_i} ((1 - x_j) \mathbf{e}^{x_j}) \right), \quad (\text{A.3})$$

which provides the semiseparable representation of rank 2

$$\begin{aligned} A_i &= \sigma^2, \\ U_{i,1} &= \sigma^2 x_i \mathbf{e}^{-x_i}, & V_{i,1} &= \mathbf{e}^{x_i}, \\ U_{i,2} &= \sigma^2 \mathbf{e}^{-x_i}, & V_{i,2} &= (1 - x_i) \mathbf{e}^{x_i} \end{aligned} \quad (\text{A.4})$$

for the Matérn 3/2 kernel. Similarly, the Matérn 5/2 kernel admits the semiseparable representation of rank 3

$$\begin{aligned} A_i &= \sigma^2, \\ U_{i,1} &= \sigma^2 \left(x_i + \frac{x_i^2}{3} \right) \mathbf{e}^{-x_i}, & V_{i,1} &= \mathbf{e}^{x_i}, \\ U_{i,2} &= \sigma^2 \mathbf{e}^{-x_i}, & V_{i,2} &= \left(1 - x_i + \frac{x_i^2}{3} \right) \mathbf{e}^{x_i}, \\ U_{i,3} &= \sigma^2 x_i \mathbf{e}^{-x_i}, & V_{i,3} &= -\frac{2}{3} x_i \mathbf{e}^{x_i}. \end{aligned} \quad (\text{A.5})$$

These representations are not unique and the choice of splitting into the 2 (Matérn 3/2) or 3 (Matérn 5/2) semiseparable terms is arbitrary.

A.2. Derivative of a Matérn Gaussian process

A GP following the Matérn 3/2 or 5/2 kernel is always differentiable, independently of the hyperparameters. Following the same reasoning as in Sect. 3, we compute the derivatives of U and V , as well as the matrix B , which appear in the covariance matrix between the GP $G(t)$ and its derivative $G'(t)$ and in the

covariance matrix of $G'(t)$ itself (see Eq. (22)). We find

$$\begin{aligned} U'_{i,1} &= \frac{\sqrt{3}\sigma^2}{\rho} (1 - x_i) \mathbf{e}^{-x_i}, \\ U'_{i,2} &= -\frac{\sqrt{3}\sigma^2}{\rho} \mathbf{e}^{-x_i}, \\ V'_{i,1} &= \frac{\sqrt{3}}{\rho} \mathbf{e}^{x_i}, \\ V'_{i,2} &= -\frac{\sqrt{3}}{\rho} x_i \mathbf{e}^{x_i}, \\ B_i &= \frac{3\sigma^2}{\rho^2}, \end{aligned} \quad (\text{A.6})$$

for the Matérn 3/2 kernel, and

$$\begin{aligned} U'_{i,1} &= \frac{\sqrt{5}\sigma^2}{3\rho} (3 - x_i - x_i^2) \mathbf{e}^{-x_i}, \\ U'_{i,2} &= -\frac{\sqrt{5}\sigma^2}{\rho} \mathbf{e}^{-x_i}, \\ U'_{i,3} &= \frac{\sqrt{5}\sigma^2}{\rho} (1 - x_i) \mathbf{e}^{-x_i}, \\ V'_{i,1} &= \frac{\sqrt{5}}{\rho} \mathbf{e}^{x_i}, \\ V'_{i,2} &= \frac{\sqrt{5}}{3\rho} (x_i^2 - x_i) \mathbf{e}^{x_i}, \\ V'_{i,3} &= -\frac{2\sqrt{5}}{3\rho} (1 + x_i) \mathbf{e}^{x_i}, \\ B_i &= \frac{5\sigma^2}{3\rho^2}, \end{aligned} \quad (\text{A.7})$$

for the Matérn 5/2 kernel.

A.3. Overflows and preconditioning

To avoid overflows (see Sect. 4.4), these representations can be adapted to use the same preconditioning method as in the case of the classical *celerite* quasiperiodic terms. The preconditioning matrix ϕ is defined as

$$\phi_{i,s} = \mathbf{e}^{-(x_{i+1} - x_i)}, \quad (\text{A.8})$$

and the preconditioned matrices $\tilde{U}$, $\tilde{V}$, $\tilde{U}'$, and $\tilde{V}'$ are obtained by dropping the exponential terms from the definitions of U , V , U' , and V' . For instance, we have

$$\begin{aligned} \tilde{U}_{i,1} &= \sigma^2 x_i, & \tilde{V}_{i,1} &= 1, \\ \tilde{U}_{i,2} &= \sigma^2, & \tilde{V}_{i,2} &= 1 - x_i, \\ \tilde{U}'_{i,1} &= \frac{\sqrt{3}\sigma^2}{\rho} (1 - x_i), & \tilde{V}'_{i,1} &= \frac{\sqrt{3}}{\rho}, \\ \tilde{U}'_{i,2} &= -\frac{\sqrt{3}\sigma^2}{\rho}, & \tilde{V}'_{i,2} &= -\frac{\sqrt{3}}{\rho} x_i, \\ \phi_{i,1} &= \phi_{i,2} = \mathbf{e}^{-(x_{i+1} - x_i)}, \end{aligned} \quad (\text{A.9})$$

for the Matérn 3/2 covariance matrix.

While this preconditioning allows us to prevent overflows and underflows due to the exponential terms, weaker numerical instabilities could arise due to the presence of x_i and x_i^2 in the

definitions of the preconditioned matrices. The absolute values of the rescaled times x_i should thus be kept as small as possible to improve numerical stability. Since the Matérn kernels are stationary (i.e., k only depends on Δt), a reference time t_0 can be chosen arbitrarily, and the definition of x (Eq. (A.2)) can be adapted accordingly

$$\begin{aligned} x &= \sqrt{3} \frac{t - t_0}{\rho} && \text{for the Matérn 3/2 kernel,} \\ x &= \sqrt{5} \frac{t - t_0}{\rho} && \text{for the Matérn 5/2 kernel.} \end{aligned} \quad (\text{A.10})$$

For instance, we could use

$$t_0 = \frac{\min(t) + \max(t)}{2} \quad (\text{A.11})$$

to avoid x_i values that are too large. This might not be sufficient for a very large time span compared to the decay timescale, ρ ; particularly in the case of the Matérn 5/2 kernel, which contains quadratic terms (x_i^2). Recently, [Jordán et al. \(2021\)](#) proposed a state-space representation for the Matérn 3/2 and 5/2 kernels, which allows a similar linear scaling of the likelihood evaluation with improved numerical stability.

Appendix B: Twice mean square differentiable semiseparable kernels

In this appendix we construct several twice mean square differentiable semiseparable kernels. When modeling time series as combinations of a GP and its derivative, the chosen GP kernel must at least be once mean square differentiable for the model to be valid. However, using a twice mean square differentiable kernel ensures that the GP's derivative is itself differentiable, which typically generates smoother models.

B.1. Alternatives to the SE kernel

The Matérn 5/2 kernel is a widely spread kernel which has the advantage of being both twice differentiable and semiseparable with rank $r = 3$ (see Appendix A). If twice differentiability is required, it is thus a natural alternative to the squared-exponential (SE) kernel:

$$k_{\text{SE}}(\Delta t) = \sigma^2 \exp\left(-\frac{\Delta t^2}{2\rho^2}\right) \quad (\text{B.1})$$

since the latter cannot be modeled with *celerite*/S+LEAF. Higher order Matérn kernels, such as the Matérn 7/2 kernel could also be used as they are more than twice differentiable and admit semiseparable representations. However, the rank of their semiseparable representations would be higher, which would increase the cost of likelihood evaluations.

We additionally propose here the exponential-sine (ES) kernel

$$k_{\text{ES}}(\Delta t) = \sigma^2 e^{-\lambda \Delta t} \left(1 + \frac{1 - 2\mu^{-2}}{3} (\cos(\mu \lambda \Delta t) - 1) + \mu^{-1} \sin(\mu \lambda \Delta t) \right), \quad (\text{B.2})$$

which is also twice differentiable and semiseparable with rank 3. Its cost is thus similar to the Matérn 5/2 kernel. The corresponding power spectral density (PSD)

$$S_{\text{ES}}(\omega) = \frac{2\sqrt{2}}{3\sqrt{\pi}} \frac{\sigma^2 (1 + \mu^2)(4 + \mu^2)}{\lambda (1 + (\omega/\lambda)^2) \left((1 + (\omega/\lambda)^2 - \mu^2)^2 + 4\mu^2 \right)} \quad (\text{B.3})$$

Fig. B.1. Comparison of the kernel functions (*top*) and power spectral density (*bottom*) of the SE, ES, and Matérn 5/2 kernels. The timescale of the Matérn 5/2 kernel is adjusted such as to minimize the maximum deviation from the SE kernel.

is always positive (for $\lambda > 0$), which ensures the positive definiteness of the kernel (e.g., [Foreman-Mackey et al. 2017](#)). The parameters λ and μ can be chosen arbitrarily, but using

$$\begin{aligned} \lambda &\approx \frac{1.091}{\rho}, \\ \mu &\approx 1.327, \end{aligned} \quad (\text{B.4})$$

makes the deviation between the SE and the ES kernels below $0.009\sigma^2$ for all lags Δt . Figure B.1 illustrates this by comparing the SE, ES, and Matérn 5/2 kernels and PSD.

B.2. Alternatives to the SEP kernel

As seen in Eq. (46) the SEP kernel can be approximated by

$$k_{\text{SEP}}(\Delta t) \approx k_{\text{SE}}(\Delta t) \frac{1 + f \cos(\nu \Delta t) + \frac{f^2}{4} \cos(2\nu \Delta t)}{1 + f + \frac{f^2}{4}}. \quad (\text{B.5})$$

In this expression, the periodic part

$$k_{\text{P}}(\Delta t) = \frac{1 + f \cos(\nu \Delta t) + \frac{f^2}{4} \cos(2\nu \Delta t)}{1 + f + \frac{f^2}{4}} \quad (\text{B.6})$$

is semiseparable and verifies $k'_{\text{P}}(0) = k_{\text{P}}^{(3)}(0) = 0$. Thus, in order to obtain a twice differentiable semiseparable kernel similar to the SEP kernel, we simply need to replace the SE part in Eq. (B.5) by a Matérn 5/2 or ES kernel and we define

$$k_{5/2\text{P}}(\Delta t) = k_{5/2}(\Delta t) k_{\text{P}}(\Delta t), \quad (\text{B.7})$$

$$k_{\text{ESP}}(\Delta t) = k_{\text{ES}}(\Delta t) k_{\text{P}}(\Delta t). \quad (\text{B.8})$$

Fig. B.2. Same as Fig. 1 but including the ESP kernel in the performance comparison.

Fig. B.3. Same as Fig. 3 but using the ESP kernel instead of the MEP kernel.

Indeed, the product of two semiseparable terms is semiseparable (see Foreman-Mackey et al. 2017) and since

$$\begin{aligned} (fg)' &= f'g + fg', \\ (fg)^{(3)} &= f^{(3)}g + 3f''g' + 3f'g'' + fg^{(3)} \end{aligned} \quad (\text{B.9})$$

the first and third derivatives of $k_{5/2P}$ and k_{ESP} also cancel out at $\Delta t = 0$. The PSD of these two kernels are given by

$$S_{KP}(\omega) = \frac{1}{1 + f + \frac{f^2}{4}} \left(S_k(\omega) + f \frac{S_k(\omega + \nu) + S_k(\omega - \nu)}{2} + \frac{f^2}{4} \frac{S_k(\omega + 2\nu) + S_k(\omega - 2\nu)}{2} \right), \quad (\text{B.10})$$

where $k = 5/2$ or ES. Since the PSD (S_k) of the Matérn 5/2 and ES kernels are strictly positive for all frequencies and the coefficient $f = (2\eta)^{-2}$ is strictly positive, we find that $S_{5/2P}$ and S_{ESP} are also strictly positive for all frequencies. The two corresponding kernels are thus positive definite.

The rank of the semiseparable representations of $k_{5/2P}$ and k_{ESP} is $r = 15$, since they are the product of a rank 3 kernel

Fig. B.4. Same as Fig. 6 but using the ESP kernel instead of the MEP kernel.

(Matérn 5/2 or ES) and a rank 5 kernel (periodic part). As a comparison, the MEP kernel (see Eq. 48), which is not twice differentiable, is of a rank of 6.

We reproduced the analyses of Sects. 5.2 and 5.3 using the ESP kernel instead of the MEP kernel. The results are presented in Figs. B.2–B.6. The cost of likelihood evaluations using the ESP kernel is about twice the cost of using the MEP kernel (see Fig. B.2), which is still much more efficient than modeling the full covariance matrix. The periodograms (Fig. B.3), as well as the GP prediction (Figs. B.4 and B.5) are very similar to the ones obtained with the MEP kernel (Figs. 3, 6, and 7). Finally, we see in Fig. B.6 that the ESP kernel reproduces the SEP kernel very closely while the MEP kernel mimics it more roughly (see Fig. 5). However, these differences between the MEP and the

Fig. B.5. Same as Fig. 7 but using the ESP kernel instead of the MEP kernel.

Fig. B.6. Same as Fig. 5 but using the ESP kernel instead of the MEP kernel.

ESP (or SEP) kernels seem to have a very weak impact on our analysis, since the periodograms and GP prediction are similar in both cases.

Appendix C: Periodogram and FAP for heterogeneous time series

We consider here the case of an heterogeneous time series following Eq. (28) and we are aimed at detecting a periodic signal affecting the first time series Y_1 only. The frameworks of Baluev (2008) and Delisle et al. (2020a) defining a general class of linear periodograms with their associated analytical FAP approximations can be applied to the merged time series y of Eq. (31) with a slight modification. We thus refer to Delisle et al. (2020a) for the details of the framework and we focus here on the required adaptations. Following Delisle et al. (2020a), we compare the χ^2 of a linear base model $\mathcal{H}$ of p parameters with enlarged models $\mathcal{K}$ of $p + d$ parameters, parameterized by the frequency ν . The base model is defined as

$$\mathcal{H} : m_{\mathcal{H}}(\theta_{\mathcal{H}}) = \varphi_{\mathcal{H}} \theta_{\mathcal{H}}, \quad (\text{C.1})$$

where $\theta_{\mathcal{H}}$ is the vector of size p of the model parameters, $\varphi_{\mathcal{H}}$ is a $n \times p$ matrix, and n the total number of points in the merged time series y . The columns of $\varphi_{\mathcal{H}}$ are explanatory time series that are scaled by the linear parameters $\theta_{\mathcal{H}}$. For instance, if we consider the merged time series of RV, BIS, and $\log R'_{\text{HK}}$, and we include in the model an offset for each of these time series, we would have to define:

$$m_{\mathcal{H}} = \gamma_{\text{RV}} \delta_{\text{RV}} + \gamma_{\text{BIS}} \delta_{\text{BIS}} + \gamma_{\log R'_{\text{HK}}} \delta_{\log R'_{\text{HK}}}, \quad (\text{C.2})$$

where γ_i is the offset of time series i , and δ_i is equal to one for measurements belonging to time series i and zero otherwise. The matrix $\varphi_{\mathcal{H}}$ would thus be a $n \times 3$ matrix, whose first column would be filled with 1 for RV measurements and 0 otherwise, the second column would be equal to 1 for BIS measurements, and the last column for $\log R'_{\text{HK}}$ measurements. The vector of parameters would then be $\theta_{\mathcal{H}} = (\gamma_{\text{RV}}, \gamma_{\text{BIS}}, \gamma_{\log R'_{\text{HK}}})$.

The enlarged model $\mathcal{K}(\nu)$ is written as

$$\mathcal{K}(\nu) : m_{\mathcal{K}}(\nu, \theta_{\mathcal{K}}) = \varphi_{\mathcal{K}}(\nu) \theta_{\mathcal{K}}, \quad (\text{C.3})$$

where $\theta_{\mathcal{K}} = (\theta_{\mathcal{H}}, \theta)$ is the vector of size $p + d$ of the parameters and $\varphi_{\mathcal{K}}(\nu) = (\varphi_{\mathcal{H}}, \varphi(\nu))$ is a $n \times (p + d)$ matrix whose p first columns are those of $\varphi_{\mathcal{H}}$, and whose d last columns are functions of the frequency, ν . In the case of an homogeneous time series, as in Delisle et al. (2020a), one typically uses $\varphi(\nu) = (\cos(\nu t), \sin(\nu t))$ (with $d = 2$). The main difference in the case of heterogeneous time series is that we only search for a periodicity in the first time series (Y_1 , typically the RV time series). Thus, we would define $\varphi(\nu) = (\cos(\nu t) \delta_1, \sin(\nu t) \delta_1)$. All the results presented in Delisle et al. (2020a) remain valid when applied to the merged time series. The only difference is that due to the presence of zeroes in φ for all measurements not belonging to the first time series, the averaging used in the definition of the effective time series length (Delisle et al. 2020a Eq. (8)) is restricted to the first time series measurements

$$\langle X \rangle = \sum_{\substack{i,j \\ \delta_1(i)=\delta_1(j)=1}} C_{i,j}^{-1} X_{i,j}, \quad (\text{C.4})$$

where C^{-1} is the inverse of the full covariance matrix of the merged time series.