

A Gaussian process model for stellar activity in 2D line profile time-series

Haochuan Yu ¹★, Suzanne Aigrain,¹ Baptiste Klein ¹, Michael Cretignier ¹, Florian Lienhard²
and Stephen J. Roberts³

¹Sub-department of Astrophysics, Department of Physics, University of Oxford, Oxford OX1 3RH, UK

²Department of Physics, ETH Zurich, Wolfgang-Pauli-Strasse 2, CH-8093 Zurich, Switzerland

³Machine Learning Research Group, Department of Engineering Science, University of Oxford, Oxford OX1 3PJ, UK

Accepted 2024 October 16. Received 2024 October 8; in original form 2024 July 12

ABSTRACT

Stellar active regions like spots and faculae can distort the shapes of spectral lines, inducing variations in the radial velocities that are often orders of magnitude larger than the signals from Earth-like planets. Efforts to mitigate these activity signals have hitherto focused on either the time or the velocity (wavelength) domains. We present a physics-driven Gaussian process (GP) framework to model activity signals directly in time series of line profiles or cross-correlation functions (CCFs). Unlike existing methods that correct activity signals in line profile time series, our approach exploits the time correlation between velocity (wavelength) bins in the line profile variations, and is based on a simplified but physically motivated model for the origin of these variations. When tested on both synthetic and real data sets with signal-to-noise ratios down to ~ 100 , our method was able to separate the planetary signal from the activity signal, even when their periods were identical. We also conducted injection/recovery tests using two years of realistically sampled HARPS-N solar data, demonstrating the ability of the method to accurately recover a signal induced by a 1.5-Earth mass planet with a semi-amplitude of 0.3 m s^{-1} and a period of 33 d during high solar activity.

Key words: line: profiles – methods: statistical – techniques: spectroscopic – Sun: activity – planets and satellites: detection.

1 INTRODUCTION

Stellar variability is currently the limiting factor on the detection and characterization of many exoplanets via the radial velocity (RV) method (e.g. Fischer et al. 2016; Crass et al. 2021). This variability arises from the interplay between magnetic fields, convection, and rotation. The most obvious component of this variability is associated with evolving magnetically active regions (spots and plage) on the rotating stellar surface, though (super-)granulation in the quiet photosphere also plays a significant role (e.g. Meunier et al. 2015; Meunier & Lagrange 2019; Al Moulla et al. 2023; O’Sullivan & Aigrain 2024). Over the past decade, considerable effort has been devoted to developing methods to model and mitigate these stellar RV signals in order to disentangle them from planetary signals (see Zhao et al. 2022 for a review).

Most of the methods proposed so far can be grouped into three broad categories: (a) time-domain methods, which model activity signals in RV time-series and, in some cases, a small number of additional activity indicators (e.g. Boisse et al. 2011; Haywood et al. 2014; Rajpaul et al. 2015; Barragán et al. 2022; Delisle et al. 2022), (b) velocity(wavelength)-domain methods, which typically work by separating pure Doppler shifts from line-shape variations in the cross-correlation function (CCF) or full spectrum (e.g. Collier Cameron et al. 2021; de Beurs et al. 2022; Zhao et al. 2022; Klein

et al. 2024; Liang, Winn & Melchior 2024; Zhao et al. 2024), (c) line-by-line methods, which exploit the fact that different spectral lines respond to activity in a different manner, while Doppler shifts affect all lines in the same way (e.g. Dumusque 2018; Cretignier et al. 2020; Cretignier, Dumusque & Pepe 2022; Lienhard et al. 2023; Shahaf & Zackay 2023). Each of these approaches has shown promising results, but also suffers from limitations. Time-domain methods require good sampling of the stellar rotation, active region evolution, and planetary orbit time-scales, and do not fully utilize the extensive information contained in the line-shape variations (and a fortiori how these variations vary from line to line). Velocity-domain and line-by-line (LBL) methods, on the other hand, typically do not account for the time-dependence of the activity signals in general. In this paper, we introduce a hybrid method, which models both the time- and the velocity-dependence of activity signals explicitly, in time-series of line-profiles (or CCFs).

To place both existing methods and the new approach that we present here in context, it is helpful to recall the standard workflow for stellar RV measurements for ultra-stable spectrographs, which is illustrated in Fig. 1. The workflow begins with a series of reduced, wavelength-calibrated, continuum-normalized high-dispersion spectra. In the most commonly used pipelines, such as the Data Reduction Software (DRS) used for the HARPS, HARPS-N, and ESPRESSO instruments (Pepe et al. 2021), each spectrum is cross-correlated with a template (usually a digital binary mask constructed from a line list, Baranne et al. 1996; Pepe et al. 2002), resulting in a CCF for each epoch. A simple function (usually a Gaussian) is then fit to

* E-mail: haochuan.yu@physics.ox.ac.uk

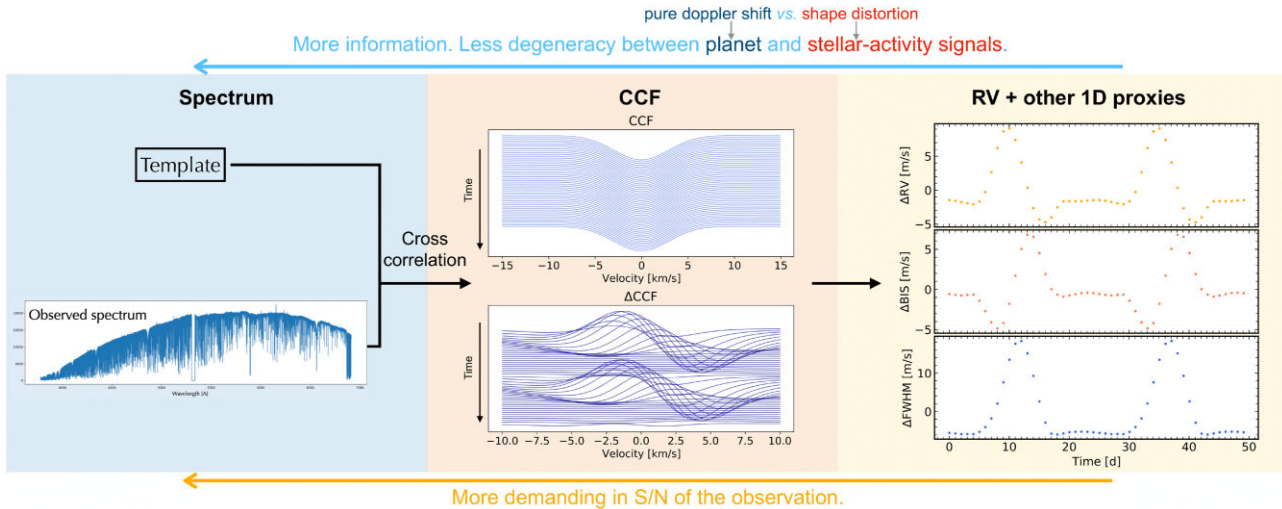


Figure 1. A typical workflow for RV extraction.

the CCF to derive the RV, together with the amplitude (contrast) and full width at half-maximum (FWHM) of the CCF. Additionally, a measurement of the CCF asymmetry, such as the bisector inverse slope (BIS) is usually extracted at the same time. Furthermore, activity indicators such as the Mount Wilson S -index (or $\log R'_{\text{HK}}$ computed from S -index) (Noyes et al. 1984), which quantify the amount of chromospheric emission in the cores of the Ca II H & K lines, are frequently extracted from the spectra alongside the RVs and line-shape parameters.

Alternatives to this workflow exist, particularly the LBL approach mentioned above, which bypasses the cross-correlation step and instead involves measuring separate RVs for thousands of individual spectral lines (the line lists used in LBL RVs are closely related to those used to construct the cross-correlation templates), or the template-free RV extraction method of Rajpaul, Aigrain & Buchhave (2020), which works by pairwise comparison of the spectra from different epochs, within individual orders or segments thereof. The RVs for individual lines or segments are of course much less precise than the global RVs, but they can then be combined carefully to minimize sensitivity to instrumental effects and stellar signals.

This workflow, in all its various incarnations, can be seen as a dimensionality reduction procedure: each spectrum (consisting of $\sim 10^5$ pixels) is reduced to a set of a few 10^3 LBL or segment RVs, or to a CCF (consisting of a few tens to hundreds of velocity bins), and then finally to a handful of numbers (one RV, plus a few line-shape and activity indicators). Stellar activity can be modelled and mitigated at any point in this process. Doing so at an earlier stage, when the dimensionality of the signal is larger, can be advantageous, as more information is retained, and the degeneracy between instrumental and stellar signals and pure Doppler shifts induced by companions is lower. However, this also requires that every spectrum has an excellent signal-to-noise ratio (SNR), which is very expensive in terms of telescope time. Furthermore, these methods tend not to take the time-dependence of activity signals into account. On the other hand, the time-domain methods, which start from the lowest-dimension end-result of the workflow, benefit from much more precise inputs (and hence have lower SNR requirements on the input spectra) and do use the time-domain information, but at the expense of discarding information about how activity signals affect the line shape and vary from line to line.

In this work, we therefore chose to start from an intermediate stage in the workflow, the line-profile or CCF stage. As the CCF combines information from thousands of spectral lines, its SNR is much higher than that of the spectra, so such a method should in principle be applicable to moderate SNR data. Most existing methods to model activity at the CCF level (Jones et al. 2017, Collier Cameron et al. 2021, John, Collier Cameron & Wilson 2022, Zhao et al. 2022, Klein et al. 2024) work by separating the CCF variations that can be explained by pure Doppler shifts from those that correspond to changes in the profile shape. However, despite their diminished amplitude, notable quasi-periodic variations tied to stellar activity still persist in the shift-driven RVs.

These methods in general do not explicitly model the time-dependence of the line-profile variations. Doing this seems like a natural progression, as the geometry of the problem (evolving active regions on a rotating sphere) imparts fairly strong constraints on the joint time- and velocity-dependence of the activity signal. These constraints form the basis for the FF' method of Aigrain, Pont & Zucker (2012) and its extension to modelling RV and activity indicator time-series using GPs (Rajpaul et al. 2015), and here we extend that model to a 2D time-series of line-profiles or CCFs. The hope is that the resulting model will provide a more complete description of the activity signal, while exploiting more of the information contained in the spectra, and being less sensitive to the time-sampling, than pure time-domain methods.

One subset of the existing approaches to model activity signals in time-series of line profiles does explicitly model their time-dependence: Doppler imaging (e.g. Donati et al. 2014; Petit et al. 2015; Yu et al. 2017; Klein et al. 2021, 2022) seeks to reconstruct a map of the inhomogeneities on the stellar surface that reproduces the line-profile variations. Limitations of the approach include that it does not explicitly model the effects of convective blueshift suppression in active regions and does not account for the potential evolution of active regions over time. A recently developed SpotCCF framework (Di Maio et al. 2024) use similar principles as Doppler imaging but restricts the number of spots in the reconstructed map of the stellar surface. This approach does allow flexibility for the spots to evolve, though do not specifically model the effects of convective blueshift suppression as well. The new method presented here addresses these limitations.

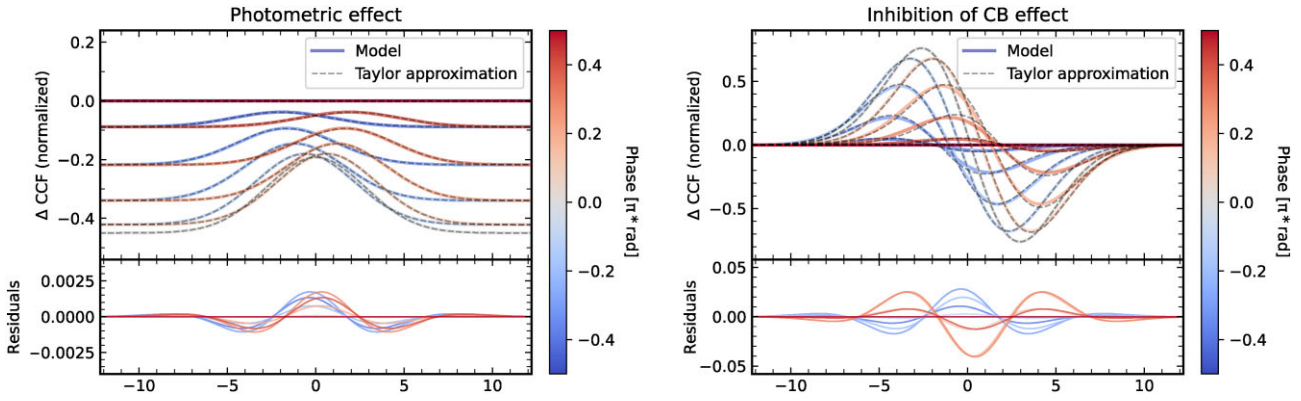


Figure 2. Perturbation of the line profile due to the photometric effect of a dark spot (left, ΔC_{phot}) and the ICB effect of a facula (right, ΔC_{icb}) for a single active region traversing the stellar disc along its equator. The coloured curves were evaluated according to equations (1) and (5), respectively, at different rotational phases. The dashed grey lines show the corresponding Taylor approximations in equations (2) and (6), truncated to 3rd order. The lower panels show the residuals between the line profiles evaluated from the original model and the Taylor approximations. In the left-hand panel, the vertical offsets between ΔCCFs from different rotation phases arise from the fact that we are estimating the line-profile using CCFs that are not continuum-normalized, and the continuum level therefore reflects the total flux variations induced by the active region.

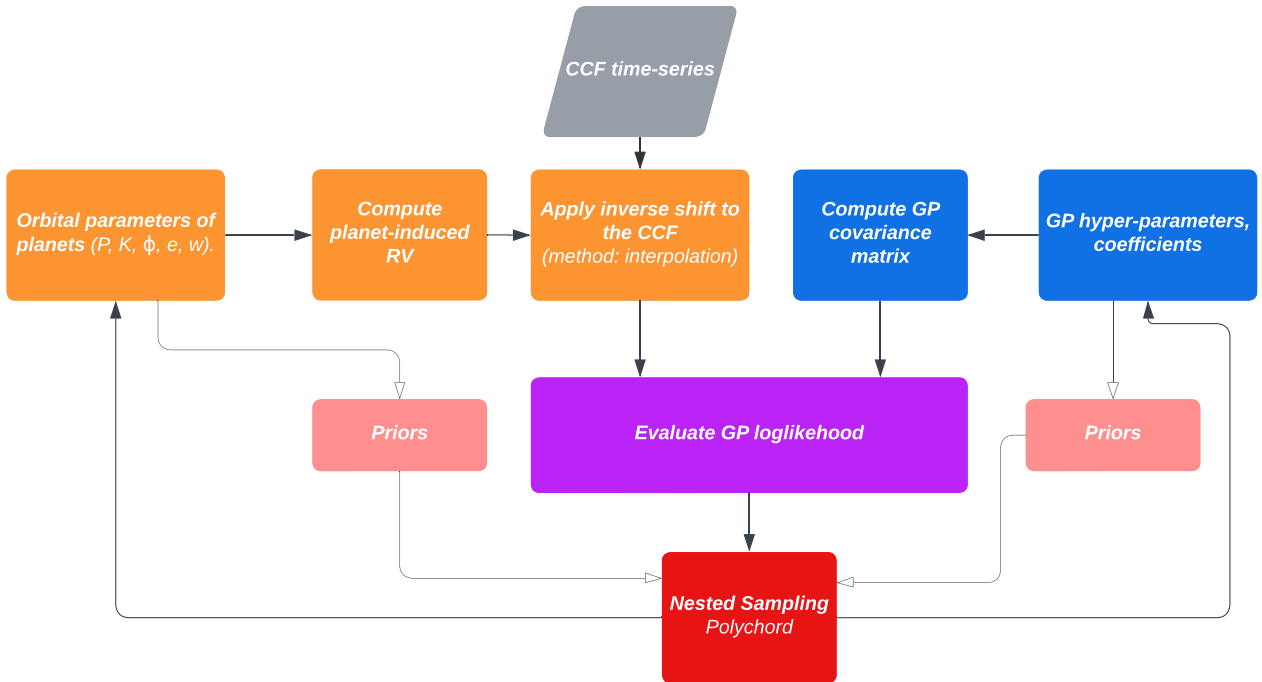


Figure 3. Flowchart illustrating the modelling workflow used in this work. The inputs are the observed time series of CCFs as well as the priors on all free parameters in the model. The output is the estimated posterior distribution of the free parameters.

Table 1. Key input parameters for simulating time-series of CCFs with SOAP2.

Case	P_{rot} (d)	N_{spots} –	Long ($^{\circ}$)	Lat ($^{\circ}$)	N_{faculae} –	Long ($^{\circ}$)	Lat ($^{\circ}$)
(i)	25.0	1	180	30	1	60	60
(ii)	25.0	2	180, –150	30, –15	2	60, –40	60, 70

In this paper, we introduce the CCF-GP framework, a novel approach for modelling stellar activity signals within time-series of line-profiles or CCFs in Section 2. We then present tests and applications of the framework on both synthetic and real data in

Section 3. We present our conclusions and suggest avenues for future work in Section 4.

2 A PHYSICS-BASED MODEL FOR LINE PROFILE VARIATIONS

We seek to model the perturbation to the mean profile of the absorption lines in a star’s spectrum induced by evolving active regions on the rotating stellar surface. Unlike previous works, which modelled either the time- or the velocity(wavelength)-dependence of activity signals, but not both, we construct a 2D model and apply it directly to the time-series of relative line profiles.




Figure 4. Left-hand panel: simulated reference-subtracted CCF (ΔCCF) with SOAP2 (case (i)). Middle panel: the best-fit CCF-GP model. Right-hand panel: the residuals after subtracting the best-fit model.

2.1 Preliminaries

Our model can be seen as an extension of the FF' framework introduced by Aigrain et al. (2012), which uses simple geometric arguments to elucidate the relationship between activity signals in photometry and in RV. We therefore start by reviewing the key elements of this framework.

2.1.1 The FF' framework

In the FF' framework, the quantity F represents the relative variation in flux due to the presence of one or more active regions on the visible hemisphere of the star. F is negative when there is a decrease in flux. Active regions can consist of dark spots, bright faculae, or a combination of the two. The rotation of the star causes the projected area of the active regions to change, and hence $F(t)$ to vary, where t denotes time. Other factors also contribute to the time-dependence of $F(t)$, namely the intrinsic evolution (growth and decay) of the active regions, and the difference in limb-darkening properties between the active regions and the rest of the photosphere, but these are treated as second-order effects in the FF' framework.

The photometric contribution of any given active region to the total stellar RV variations scales as the product of $F(t)$, and $V_{\text{loc}}(t)$, the line-of-sight velocity of the stellar surface at the location of the active region.

In addition to this photometric effect, magnetically active regions also affect the total RV because a locally enhanced magnetic field density leads to a local reduction in the velocity of convective upflows. In the absence of active regions, these upflows result in a net convective blueshift, which is strongest near the disc centre and weakest near the limb. This blueshift is locally partially suppressed in the presence of active regions (e.g. Gray 2008). This ‘Inhibition of the Convective Blueshift’ (ICB) effect is expected to scale as the product of the active region projected area, which varies like $F(t)$, and of the line-of-sight component of the up-flow velocity.

The key insight of Aigrain et al. (2012) is that, to first order, $V_{\text{loc}}(t)$ behaves as the first derivative of $F(t)$, denoted by $F'(t)$, while the line-of-sight component of the up-flow velocity behaves as $F(t)$. The total RV variation can therefore be approximated as the sum of two terms, one scaling as $F(t)F'(t)$ and one scaling as $F^2(t)$.

2.1.2 The multi-GP framework

Aigrain et al. (2012) also noted that the time-dependence of $F(t)$ can be described by a GP with a quasi-periodic kernel. Rajpaul et al. (2015) then adapted the FF' formalism to model simultaneously the

signatures of active regions in RV time-series and in individual time-series of activity indicators extracted from the same spectra. These include both chromospheric indices, such as the Mt. Wilson S -index (or $\log R'_{\text{HK}}$ computed from S -index) or line-shape indicators, which trace changes in the width or asymmetry of the mean spectral line profile. Each time series is modelled as the linear combination of a term in $G(t) \equiv F^2(t)$ and a term in $G'(t) \propto F(t)F'(t)$, where $G(t)$ is a quasi-periodic GP, and the coefficients of proportionality are free parameters of the model. As this formalism involves a multi-output GP (where a single, latent GP variable is used to model multiple output time-series), we refer to it as the ‘multi-GP’ approach. A number of extensions and public implementations of this framework have since been published (e.g. Jones et al. 2017; Barragán et al. 2022; Delisle et al. 2022; Nardiello et al. 2022; Hara & Delisle 2023).

2.2 Extension to line-profile variations

We now consider, instead of a discrete set of time-series (RV and one or more activity indicators), a 2D time-series of line profiles. More precisely, our starting point is a time-series of relative line-profiles, $\Delta C(v, t)$, where v denotes velocity and t denotes time, obtained by subtracting a reference line profile (representing the line profile in the absence of any active regions) $C_{\text{ref}}(v, t)$ from the original time-series. In other words, we seek to model the perturbation to the line profile induced by (one or more) active regions. We do not specify at this stage how the line-profiles are extracted from the spectra, or how the reference line profile is constructed; this will be discussed in detail in Section 2.4.

As in the FF' framework, we treat the effect on the stellar line profile of an active region moving across the stellar disc as the sum of two distinct terms: the photometric effect and the ICB effect.

2.2.1 The photometric effect

To first order, an active region (e.g. a dark spot) alters the mean stellar line profile by suppressing the flux emerging from a small region of the stellar surface that rotates with the star. The resulting distortion of the line-profile, $\Delta C_{\text{phot}}(v, t)$, can be modelled as:

$$\Delta C_{\text{phot}}(v, t) \propto F(t) C_{\text{loc}}(v - V_{\text{loc}}(t)), \quad (1)$$

where $C_{\text{loc}}(v)$ is the intrinsic line profile of the quiet photosphere (before rotational broadening). In equation (1), it is evaluated at the location of the spot, i.e. Doppler-shifted by $V_{\text{loc}}(t)$, the line-of-sight component of the stellar surface velocity at the location of the line

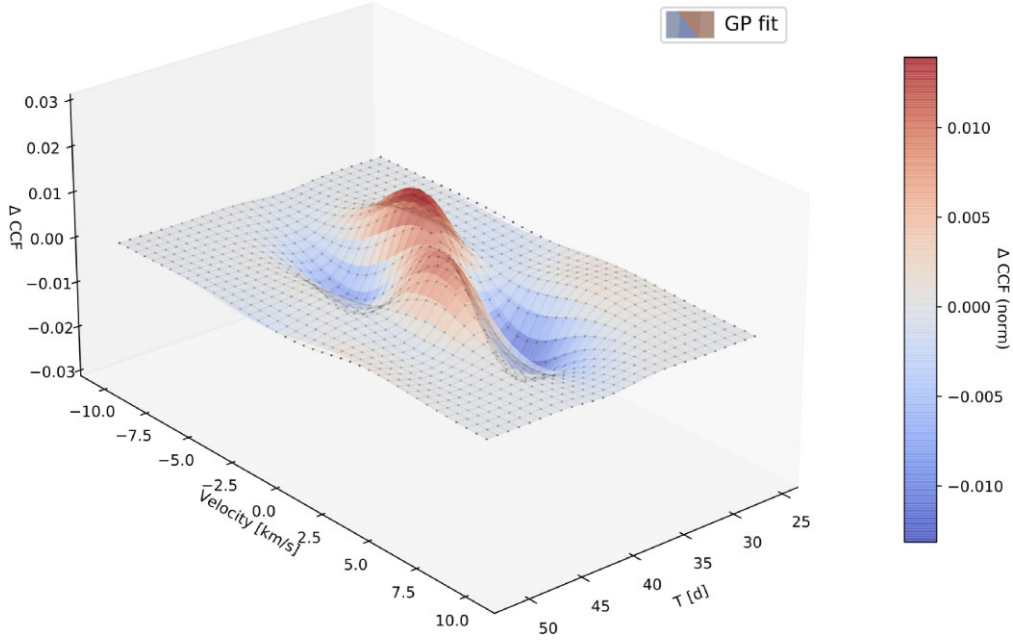


Figure 5. 3D visualization of the CCF-GP model fit to the referenced-subtracted CCF, ΔCCF simulated with SOAP2 (case (i)), zooming into one rotation period. The black markers represent the simulated data, and the coloured surface illustrates the model fit, with the colour gradient indicating ΔCCF values (positive in red and negative in blue).

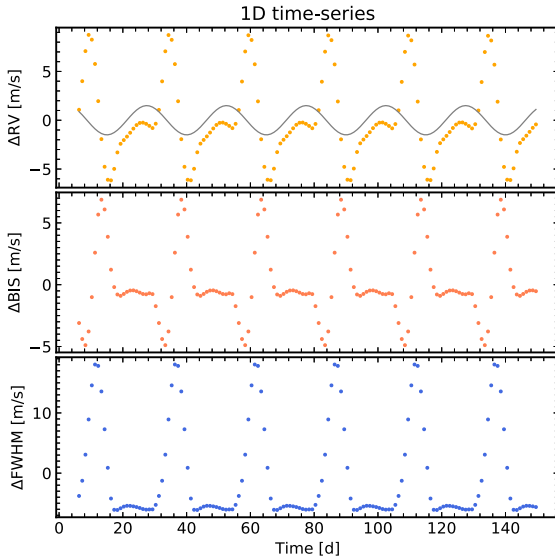


Figure 6. Mean-subtracted 1D time-series, including RV, BIS, and FWHM, which are derived from the planet-injected simulated CCFs generated with SOAP2 case (ii). The first panel also shows the planet-induced RV in grey.

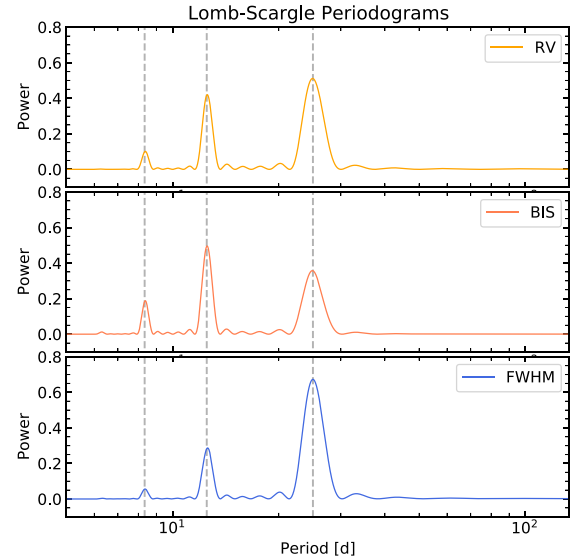


Figure 7. Lomb–Scargle periodogram of each 1D time series (including RV, BIS, and FWHM) shown in Fig. 6. The grey vertical markers indicate the positions of the 25 d rotation signal and its harmonics.

profile. As in the FF' framework, $F(t)$ is a time-dependent term that accounts for the spot's photometric contrast, fore-shortening, and limb-darkening. It can be thought of as the relative photometric perturbation caused by the spot in the bandpass probed by the spectrograph.

Following equation (1), we introduce a Taylor-series expansion of C_{loc} about $V_{\text{loc}} = 0$:

$$\Delta C_{\text{phot}}(v, t) \propto F(t) \sum_{k=0}^{\infty} \frac{V_{\text{loc}}^k(t)}{k!} C_{\text{loc}}^{(k)}(v), \quad (2)$$

where $C_{\text{loc}}^{(k)}(v)$ is the k^{th} derivative of C_{loc} with respect to v . The coloured curves in the left-hand panel of Fig. 2 show ΔC_{phot} computed according to equation (1) for a dark spot traversing the stellar disc along its equator. In this case, $F(t)$ is approximated as a pure sine function. The dashed grey lines show the Taylor approximation in equation (2), truncated to 3rd order.

As we have already seen, we can approximate $V_{\text{loc}}(t)$ as $\gamma F'(t)$, where γ is a constant coefficient. Therefore, the line-profile pertur-




Figure 8. Posterior distributions of the recovered orbital parameters for the simulated data set (SOAP2 case (ii)) from both CCF-GP (in blue) and multi-GP (in orange) models. Vertical red lines indicate the true values from the injections.

bation due to the photometric effect can be represented as

$$\Delta C_{\text{phot}}(v, t) = \sum_{k=0}^{\infty} u_k(t) p_k(v), \quad (3)$$

where the vectors $p_k(v) \equiv C_{\text{loc}}^{(k)}(v)$ are given by the local, quiet line-profile and its derivatives, while the vectors u are proportional to a combination of F and its derivatives, similar to the original FF' framework:

$$\begin{aligned} u_0(t) &\propto F(t), \\ u_1(t) &\propto \gamma F(t)F'(t), \\ u_2(t) &\propto \frac{1}{2}\gamma^2 F(t)F'^2(t), \\ u_3(t) &\propto \frac{1}{6}\gamma^3 F(t)F'^3(t), \dots \end{aligned} \quad (4)$$

2.2.2 Inhibition of convective blueshift effect

As previously mentioned, the magnetic field associated with active regions suppresses convection locally, leading to a local redshift in the line profile. Note that this is a very simplified representation: in reality, the magnetic field also alters the line profiles. For now, we model the ICB effect as a straightforward shift in the local line profile at the location of the active region.

Given this premise, the perturbation of the line-profile due to the ICB effect can be expressed as

$$\Delta C_{\text{cb}}(v, t) \propto F(t) [C_{\text{loc}}(v - V_{\text{loc}}(t) - \delta V_c) - C_{\text{loc}}(v - V_{\text{loc}}(t))], \quad (5)$$

where δV_c denotes the local redshift induced by the magnetic field. As in the case of the photometric effect, we then employ a Taylor-series

expansion:

$$\Delta C_{\text{cb}}(v, t) \propto F(t) \sum_{k=0}^{\infty} \frac{(V_{\text{loc}}(t) - \delta V_c)^k - V_{\text{loc}}^k(t)}{k!} C_{\text{loc}}^{(k)}(v). \quad (6)$$

The right-hand panel of Fig. 2 shows the ICB perturbation to the line profile, ΔC_{cb} , caused by a facula crossing the stellar disc along its equator, as given by equation (5). Again, here $F(t)$ is approximated as a pure sine function. The dashed grey lines show Taylor approximation of equation (6), once again truncated to 3rd order.

We can thus represent the line-profile perturbation due to the ICB effect in the same form as in equation (3):

$$\Delta C_{\text{cb}}(v, t) = \sum_{k=0}^{\infty} u_k(t) p_k(v), \quad (7)$$

but the vectors $u_k(t)$ are now given by

$$\begin{aligned} u_0(t) &= 0, \\ u_1(t) &\propto -\delta V_c F(t), \\ u_2(t) &\propto -\delta V_c \gamma F(t)F'(t) + \frac{1}{2}\delta V_c^2 F(t), \\ u_3(t) &\propto -\frac{1}{2}\delta V_c \gamma^2 F(t)F'^2(t) + \frac{1}{2}\delta V_c^2 \gamma F(t)F'(t) - \frac{1}{6}\delta V_c^3 F(t), \\ &\dots \end{aligned} \quad (8)$$

2.2.3 Full model

For a single active region, we can combine the photometric and ICB components of the perturbation by adding them together. The full model takes the same form as equations (3) and (7):

$$\Delta C(v, t) = \sum_{k=0}^{\infty} u_k(t) p_k(v), \quad (9)$$

and, as before, the vectors $p_k(v) \equiv C_{\text{loc}}^{(k)}(v)$ are given by the local, quiet line-profile and its derivatives, but the vectors $u_k(t)$ are given by:

$$\begin{aligned} u_0(t) &= \alpha F(t), \\ u_1(t) &= \alpha \gamma F(t)F'(t) - \beta \delta V_c F(t), \\ u_2(t) &= \frac{\alpha}{2} \gamma^2 F(t)F'^2(t) \\ &\quad + \beta [-\delta V_c \gamma F(t)F'(t) + \frac{1}{2}\delta V_c^2 F(t)], \\ u_3(t) &= \frac{\alpha}{6} \gamma^3 F(t)F'^3(t) \\ &\quad + \beta [-\frac{1}{2}\delta V_c \gamma^2 F(t)F'^2(t) + \frac{1}{2}\delta V_c^2 \gamma F(t)F'(t) - \frac{1}{6}\delta V_c^3 F(t)], \\ &\dots \end{aligned} \quad (10)$$

where α and β are constant coefficients.

For the case where the line profiles are continuum-normalized, we expect $\Delta C_{\text{norm}}(v, t)$ described by:

$$\Delta C_{\text{norm}}(v, t) = \frac{\Delta C(v, t) + C_{\text{ref}}(v, t)}{1 + \alpha F(t)} - C_{\text{ref}}(v, t), \quad (11)$$

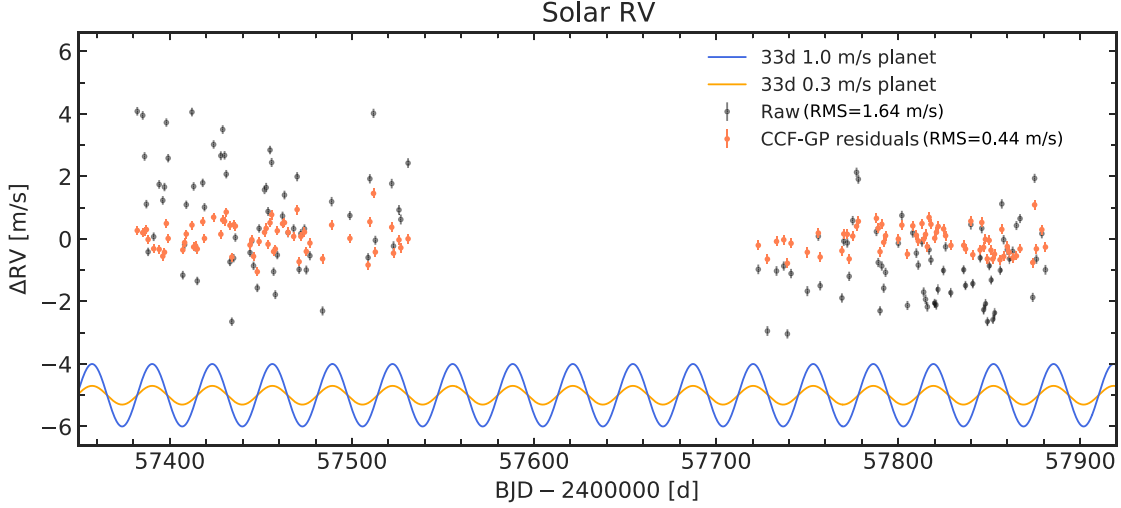
assuming the $C_{\text{ref}}(v, t)$ is continuum-normalized already.

2.2.4 Latent GP model for the $u_k(t)$

In the original multi-GP framework, which is tailored for 1D time-series data, each time-series is modelled as a sum of terms in F^2 and

Table 2. Evidence for combinations of model and simulated data set with 0 or 1 planet, generated from SOAP2 case (ii).

Case	CCF-GP log Z	CCF-GP $\Delta \log Z$	Multi-GP log Z	Multi-GP $\Delta \log Z$
0-planet data and 0-planet model	44232.20 \pm 0.33	-45.44	2348.28 \pm 0.30	-2.31
0-planet data and 1-planet model	44277.64 \pm 0.32	-	2350.59 \pm 0.29	-
1-planet data and 0-planet model	43941.19 \pm 0.24	-	2292.24 \pm 0.30	-
1-planet data and 1-planet model	44242.19 \pm 0.34	301.00	2349.71 \pm 0.29	57.47

**Figure 9.** The HARPS-N solar RVs used for injection/recovery tests (black markers) and the residual RVs (coral markers), which were measured from the CCFs after subtracting the best-fit CCF-GP model. Also shown are the injected 33 d planet signals for both semi-amplitude settings in blue (1.0 m s^{-1}) and orange (0.3 m s^{-1}), offset by -5 m s^{-1} in the figure.

FF' . These are modelled by defining a latent variable $G \equiv F^2$, over a GP prior is placed. Because $dG/dt = 2FF'$, the observed time series are linear combinations of the latent GP and its first derivative and are thus GPs themselves.

In our 2D model for the line profile variations, the expressions for the vectors $u_k(t)$ in equation (10) involve combinations of F and F' , F^2 , F^3 , \dots . These can no longer be expressed directly as linear combinations of a GP and its derivative(s). As a result, even if we assume that F , or some function of F (such as F^2 , F' , or F'^2) is a GP, the $u_k(t)$ are not GPs, and neither is the time-series of line-profile variations.

We explored two possible solutions to this problem. The first consists in placing a GP prior on F and jointly sampling F and F' from this prior, before computing F^2 , F^3 , \dots . These are then used to compute the $u_k(t)$ before the resulting full model is compared to the observed line-profile time-series. This approach allows us to reproduce the precise form of equation (10), but involves drawing a GP sample for every likelihood evaluation. The computational cost of such an approach proved to be prohibitive.

We therefore adopted a less accurate but cheaper solution, which consists in treating each $u_k(t)$ as a linear combination of some GP $G(t)$ and its time-derivative $G'(t)$, just as in the multi-GP framework described in Section 2.1.2. While this is clearly an approximation, our experience in modelling RV and activity indicator time-series with the multi-GP framework has shown that the combination of a GP and its first derivative is typically flexible enough to capture higher-order variations.

Therefore, in our final model, each $u_k(t)$ is approximated as $a_k G(t) + b_k G'(t)$ and $p_k(v) \equiv C_{\text{loc}}^{(k)}(v)$. We describe the procedure

for deriving $p_k(t)$ from the observations later in Section 2.4. Consequently, the line profile residuals at a given velocity are also a linear combination of G and G' :

$$\Delta C(v, t) = \sum_{k=0}^r p_k(v) a_k G(t) + \sum_{k=0}^r p_k(v) b_k G'(t), \quad (12)$$

where r is the number of terms used. We found that setting $r = 3$ provides a sufficiently accurate approximation for most cases while also efficiently reducing computational costs.

Thus, the matrix $\Delta \mathbf{C}$ is modelled as a latent GP model with n outputs, where n is the number of velocity bins, and the relationship between the output time-series is controlled by the $p_k(t)$, which corresponds to the emergent line profile of the quiet photosphere. Importantly, the total number of free parameters in the model is only, $2r + m$, where m is the number of parameters of the GP covariance function used.

2.2.5 Relationship to principal component analysis

The key insight in our model is that the line-profile perturbations can be modelled as a sum of separable functions of time and velocity. This is in some sense analogous to performing principal component analysis (PCA) on the matrix of line-profile residuals. Specifically, the singular value decomposition (SVD) of this matrix is given by

$$\Delta \mathbf{C} = \mathbf{U} \mathbf{S} \mathbf{P}, \quad (13)$$

where \mathbf{U} and \mathbf{P} are rectangular matrices of left- and right-singular vectors (eigenvectors) which depend only on time and on velocity, respectively, while \mathbf{S} is a diagonal matrix of singular values (square-root of eigenvalues). However, in SVD each set of singular vectors




Figure 10. Lomb–Scargle periodogram of the raw solar RV and the RV residuals shown in Fig. 9. The grey vertical markers indicate the positions of the 27 d solar-rotation signal and its harmonics.

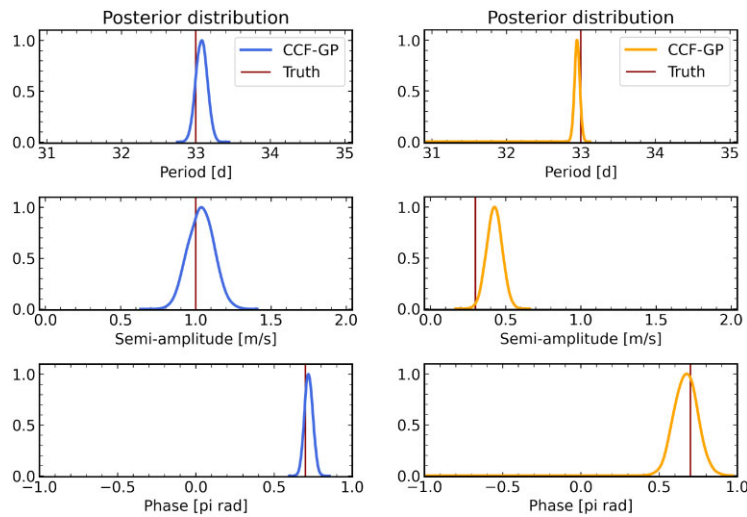


Figure 11. Posterior distribution of the extracted planetary parameters for the HARPS-N injection/recovery tests for planets with 33 d period and semi-amplitudes of 1.0 (left, blue) or 0.3 m s⁻¹ (right, orange). Vertical red lines indicate the true values from the injections.

forms an orthonormal basis, whereas this isn’t necessarily the case for the $u_k(t)$ and $p_k(v)$ vectors in our model.

2.2.6 Limitations

The physical model we have outlined in Section 2.2.3 describes the perturbation to the profile of a single absorption line caused by a single active region crossing the stellar disc and is in itself an approximation. Furthermore, there may be multiple active regions present at the same time at different locations on the stellar disc. The ratio of the a and b coefficients in equation (10) depends on the relative area and contrast spot and plage, which can differ between active regions, while the local redshift caused by plage (δV_c in equation (5)) is likely to depend on both the strength of the magnetic field in the plage and the viewing angle. Thus, treating the α and β coefficients in equation (10) as independent of time is not strictly correct. Another aspect that we have overlooked in our physically motivated model is that the observed line profile is not equal to the true line profile, but is convolved with the instrumental line profile.

For instance, when we incorporate the distortion of the line profile ΔC into the reference profile c_{ref} , both profiles have already been convolved with the instrumental profile. This can be slightly biased from the reality when the two unconvolved components are first summed and the convolution is performed after.

To some extent, we expect these imperfections in our physical model to be absorbed by the flexible latent GP formulation we have adopted in Section 2.2.4. Conversely, this means that the physical interpretation of the model becomes less tractable.

2.3 Disentangling activity signals from planets

A key purpose of the model introduced above is to disentangle activity signals from the Doppler shifts caused by planetary companions, and thereby to improve the detectability and characterization of the planets. To do this, we must explicitly account for the planet(s) when modelling the line-profile variations, as our model would otherwise have enough flexibility to absorb or modify some of the planet signals. The parameters of the planet(s) and the activity model must therefore be sampled jointly. In order to do this without escalating the

computational cost, it is vital to keep the GP as the bottom layer of the workflow to ensure the likelihood can still be computed analytically.

2.3.1 Workflow

Fig. 3 illustrates our overall workflow. At each step, we sample the parameters for the planet(s) and the activity model simultaneously. The $5p + 1$ planet parameters (where p is the number of planets, and the remaining parameter defines the systemic RV) are used to evaluate the planet-induced RV signal, and the line profiles are then shifted inversely by the corresponding amount.

To apply a shift to the line profiles, we initially fit a non-parametric function to each profile using least-squares cubic interpolation, as implemented in `Scipy` (Virtanen et al. 2020).¹ Subsequently, we apply an inverse shift to the velocity grid of each line profile to account for the RV signal induced by the planet. This shifted velocity grid is then input into the fitted non-parametric function to compute the shifted line profile for each time stamp.

In parallel, given the $2r + m$ activity parameters, a GP solver is used to evaluate the covariance matrix \mathbf{K} of the model and compute a marginal likelihood according to:

$$\log \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}) = -\frac{N_{\text{obs}}}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{K}| - \frac{1}{2} \mathbf{r}^T \mathbf{K}^{-1} \mathbf{r}, \quad (14)$$

where $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are the parameters of the planet and activity models, $\mathbf{r} = \mathbf{y} - \boldsymbol{\mu}$ contains the residuals of the observed line profiles \mathbf{y} subtracted by the reference (mean) line profile $\boldsymbol{\mu}$, ‘flattened’ into a one-dimensional vector. N_{obs} is the number of observations. As we are modelling relative line profiles, which have been shifted to account for the planet signal, the mean function of the GP is zero everywhere.

2.3.2 Practical implementation of GP

We use the fast GP solver `S + LEAF 2`² (Delisle et al. 2022) with a Matérn 3/2 exponential periodic (MEP) kernel. This serves as a fast approximation of the quasi-periodic kernel frequently used to describe rotational modulation active regions in stellar light curves and RV time-series, allowing the likelihood evaluation to scale linearly with the number of observations and velocity bins. We also include a constant term on the diagonal of the covariance matrix to absorb any imperfections in our model, which results in $m = 4$ parameters for the GP covariance function.

As we are sampling over a relatively large number of parameters and expect the likelihood hyper-surface to be multimodal in some cases, we sample the joint posterior over the parameters using the `PolyChord`³ nested sampler (Handley, Hobson & Lasenby 2015a, b). Nested sampling is generally considered more adept at handling multimodal problems, making it a more appropriate choice for this application than more standard Markov Chain Monte Carlo (MCMC) samplers, while the `polychord` implementation is known to perform especially well for large numbers of parameters.

¹<https://scipy.org/>

²<https://gitlab.unige.ch/Jean-Baptiste.Delisle/spleaf>

³<https://github.com/PolyChord/PolyChordLite>

2.4 Model inputs

The starting point for our method is a matrix $\Delta \mathbf{C}$ containing a time-series of relative line-profiles:

$$\Delta C_{ij} = C_{ij} - c_{\text{ref},i} = C(v_i, t_j) - c_{\text{ref}}(v_i), \quad (15)$$

obtained by subtracting a reference line profile \mathbf{c}_{ref} (representing the line profile in the absence of any active regions) from each row of the matrix of line profiles \mathbf{C} .

2.4.1 Extracting the line profiles

Except where stated otherwise, we equate each row of \mathbf{C} to the CCF produced by the standard ‘DRS’ (Data Reduction Software, used for HARPS, HARPS-N, and ESPRESSO data (e.g. Lovis 2007) for that spectrum. These CCFs are obtained by cross-correlating the spectra with a binary mask, which is zero everywhere except at the centre of individual spectral lines (Baranne et al. 1996, resulting in ‘inverted’ CCFs, which resemble an absorption line profile, e.g. the top-middle panel in Fig. 1). Note that if the CCFs are un-normalized, the continuum level (away from the centre of the line profile) varies with time.

In principle, it may be preferable to extract line profiles from the spectra using least squares deconvolution (LSD, e.g. Donati et al. 1997; Kochukhov, Makaganiuk & Piskunov 2010). However, we chose to use CCFs in this work because they are computed automatically by the DRS and thus readily available. In practice, we have found that, provided the line lists used in the creation of the CCF template and in the LSD process are similar, the resulting line profiles are also very similar.

2.4.2 Reference profile

Ideally, the reference profile \mathbf{c}_{ref} should be equal to the line profile in the absence of any active regions, so that the residuals $\Delta \mathbf{C}$ represent the perturbations induced by activity. In practice, we do not normally have spectra that are guaranteed to be activity-free. Therefore, the reference line profile is obtained by averaging the line profiles \mathbf{C} over the time axis, optionally selecting a subset with a lower activity level (as measured, for example, by the calcium S or $\log R'_{\text{HK}}$ indices).

2.4.3 Estimating the $p_k(v)$

We approximate the local, quiet line profile $C_{\text{loc}}^{(k)}(v)$ as the deconvolution of the reference profile \mathbf{c}_{ref} with a rotational broadening kernel estimated from the $v \sin i$ of the star. The effect of the rotational broadening becomes negligible for slow rotators when the $v \sin i$ is significantly smaller than the width of the instrumental profile.

3 TESTS AND APPLICATIONS

In this section, we analysed both synthetic and real data sets. Our objectives are fourfold: (a) to validate the framework by assessing its ability to model signatures from stellar spots and faculae; (b) to determine if the framework can more effectively disentangle planetary and stellar signals than existing 1D time series methods; (c) to evaluate its efficiency in detecting planets with low-amplitude signals; and (d) to demonstrate its capability to operate effectively at a SNR typical of most real-world observations, approximately 100.




Figure 12. As Figs 9 and 11, but the figure shows the injection/recovery test of the long-period planet at 101.2 d orbital period, and of a semi-amplitude of 0.5 m s^{-1} .

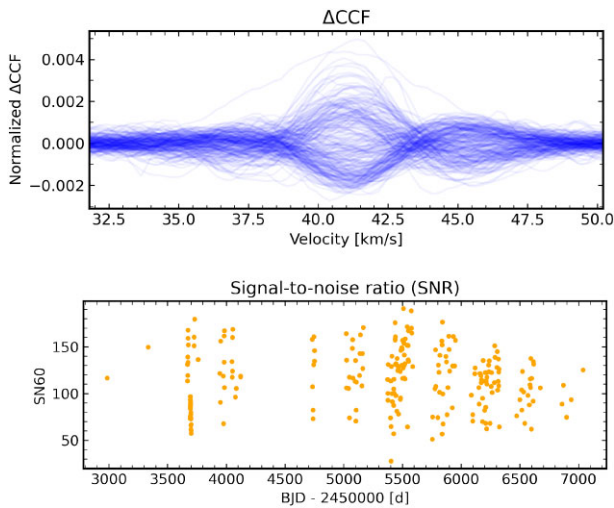


Figure 13. Reference-subtracted CCFs and the associated SNR per pixel (in échelle order 60) plotted over time. The SNR for these observations varies between approximately 50 and 180, with a median value of 100.

3.1 Simulated data from SOAP 2.0

Before implementing the suggested framework on real data, we tested it using simulated data from version 2 of the SOAP (Spot Oscillation and Planet) code (Boisse, Bonfils & Santos 2012; Dumusque, Boisse & Santos 2014). SOAP2⁴ enables users to generate time-series of CCFs distorted by active regions, namely spots and faculae, on the stellar surface. This is achieved by specifying the latitude, longitude, and size of each region, alongside the star’s basic properties such as its rotation period. The tests involved two configurations – cases (i) and (ii), detailed in Table 1, including the simulated rotation period of the star, the numbers and locations of the spots and faculae. For both cases, the sampling rate is one point per day, and the 150 d duration corresponds to six times the rotation period of the star.

For the first case, we simulated a stellar surface featuring one spot and one facula, primarily to validate the framework’s concept, specifically, its ability to track the signatures in the CCFs influenced

by both the spot and facula. The simulated reference-subtracted CCFs (ΔCCF) are shown in Fig. 4 in the left-hand panel, with the best-fit CCF-GP model in the middle panel and residuals in the right-hand panel. The typical amplitude of the variations in the ΔCCF is around 1 per cent. Fig. 5 shows a 3D representation of ΔCCF versus velocity and time, enabling a clearer visualization of the model fit. This figure highlights that the variations are correlated in both the velocity and time domains, and demonstrates that their correlations are effectively captured by our framework.

For the second case, our objective was to demonstrate the model’s ability to limit the degeneracy between the signals induced by stellar activity and planets. For this test, we simulated a more intricate and realistic stellar surface with two spots and two faculae, with the properties listed in Table 1. To challenge the framework, we then injected a planetary signal into the simulated CCFs, with exactly the same orbital period as the star’s rotation period (in this instance, 25 d). This planet signal had a semi-amplitude of 1.0 m s^{-1} , and its phase was arbitrarily set to 0.3π rad. The phase is defined with reference to $T = 0$, and phase = 0 corresponding to the position of inferior conjunction for the injected planet. Fig. 6 displays the mean-subtracted 1D proxies, including RV, BIS, and FWHM, derived from the planet-injected simulated CCFs. The first panel also shows the planet-induced RV in grey to enable its comparison to the overall RV signal.

We first computed the Lomb–Scargle periodogram (using *Astropy* (Astropy Collaboration et al. 2013, 2018, 2022) of each 1D time series, as shown in Fig. 7. The grey vertical markers indicate the positions of 25 d signals and their harmonics. Using the traditional approach of comparing the periodogram of the RV and 1D activity indicators, any peak in the RV periodogram that also appears in the periodograms of the activity indicators would be dismissed as likely arising from activity. Therefore, discerning signals induced by planets from those induced by stellar activity is impossible using periodograms alone when both share the same period.

We subsequently applied the CCF-GP framework to the CCF time series. For comparison, we also applied the multi-GP framework to the RV, BIS, and FWHM extracted from the same CCF time series. Both models incorporated the same planet model for RVs, featuring a single Keplerian. Crucially, the parameter priors in both the multi-GP and CCF-GP were identical and remained uninformative. Fig. 8 presents the posterior distribution of the recovered orbital parameters

⁴<https://www.astro.up.pt/resources/soap2/>

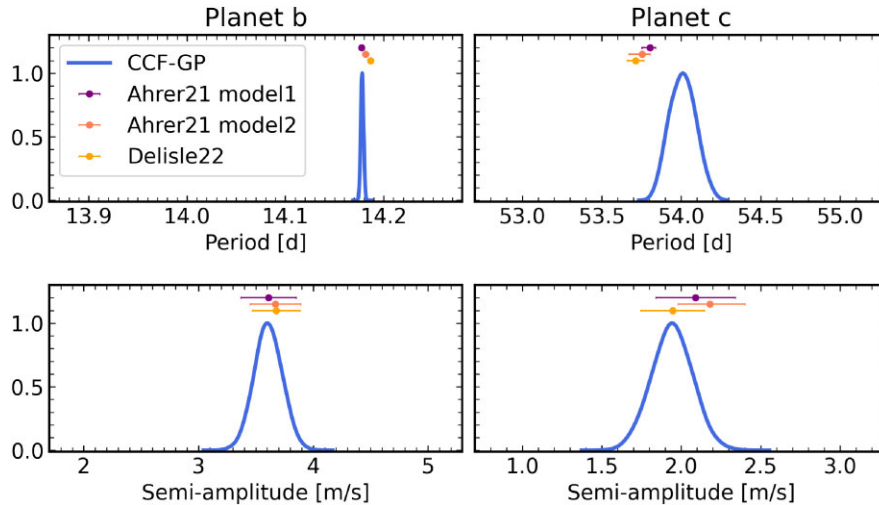


Figure 14. Posterior distributions of the derived planetary parameters from the CCF-GP analysis on the HD 13808 data. The derived values and associated uncertainties of the corresponding parameters from Ahrer et al. (2021) and Delisle et al. (2022) are indicated by the markers with errorbars.

from both CCF-GP (in blue) and multi-GP (in orange) models. Vertical red lines indicate the true values from the injections.

It is evident that both models adeptly recover the 25 d period. However, as the stellar rotation period is also at 25 d, to be certain that the detected signal is indeed that of the planet, we also required that the phase of the planet signal be retrieved. Observing the posterior distribution over the phase reveals a notable advantage for the CCF-GP, which captures the truth within a 1σ uncertainty range, while the phase determined via the multi-GP deviates by more than 0.1π rad from the true value. Regarding the semi-amplitude, both models exhibit a tendency to provide inflated estimates of $K = 1.86^{+0.05}_{-0.05} \text{ m s}^{-1}$ (CCF-GP) and $K = 2.57^{+0.26}_{-0.24} \text{ m s}^{-1}$ (multi-GP), though the CCF-GP derived value is closer to the truth of $K = 1.5 \text{ m s}^{-1}$. This inflation implies that the models are conservative in removing the activity signals, leaving a fragment of these signals intertwined with the planet signals. In the case of the CCF-GP model, this slightly conservative behaviour might come from the simplification of modelling the CCFs as vanilla Gaussian processes (as discussed earlier in Section 2), potentially neglecting minor yet pertinent non-Gaussian aspects of the signals.

With the same simulated data set, we conduct evidence comparisons on combinations of model and simulated data set with 0 or 1 planet with CCF-GP and multi-GP frameworks. The four cases are listed in Table 2, with their corresponding evidence $\log Z$ estimated from the nested sampling. We also provide $\Delta \log Z$, which is calculated as the evidence of the ‘correct’ model minus the ‘incorrect’ model. We find that for the data containing 1 injected planet, the model containing 1 planet (1-planet model) is always strongly preferred compared to the 0-planet model for both CCF-GP and multi-GP. However, for the data without planet injection, the 1-planet model is preferred to the 0-planet model for both CCF-GP and multi-GP, though such preference for the latter one is less presented. This suggests further avenues for improvement to prevent false-positive detections.

Overall, we find that in this specific test, the CCF-GP approach generally outperforms the multi-GP approach in disentangling activity and planetary signals, especially when their periods closely match the star’s rotation period or its harmonics.

3.2 Injection and recovery test on HARPS-N solar data

To assess the behaviour of the CCF-GP framework under more realistic conditions and to determine its ability to detect low-amplitude planet signals (while having a benchmark truth), we conducted planet injection and recovery tests using HARPS-N solar data.

The spectra were collected using the HARPS-N solar telescope (Cosentino et al. 2012; Dumusque et al. 2015; Phillips et al. 2016) with observations taken for several hours each day at 5 min cadence since 2015. Initial processing of the raw data was performed by Dumusque et al. (2021) employing the ESPRESSO DRS (Pepe et al. 2021). Following Collier Cameron et al. (2019), cloud-affected spectra were excluded, and the remainder were shifted to the solar rest frame through interpolation using the Solar barycentric RV solution derived with the JPL Horizons software (Giorgini et al. 1996). All sets of three consecutive 5 min exposures are combined to make one 15 min exposure for mitigating the effect of p -mode oscillations. We then simulated the temporal sampling one might expect for hypothetical stars observed from La Palma by adding 0.5 d to the observation times (to shift them to night times), then selected at most one observation each night when a hypothetical star located at $\text{RA} = 90^\circ$, $\text{DEC} = +30^\circ$ is observable at an airmass below 1.8. Our test utilized data from mid-2015 to mid-2017, corresponding to an active part of the solar magnetic activity cycle. There are 130 observations in total, with an average sampling rate of one point per 3.8 d, and a typical SNR of ~ 370 per 0.82 km s^{-1} velocity bin in échelle order 60 of the spectra. Fig. 9 shows the mean-subtracted RV measured from the selected data, which exhibits a root mean square (RMS) value of 1.64 m s^{-1} .

We first applied the CCF-GP framework to analyse the HARPS-N solar CCFs without injecting any planetary signals. This involved fitting the CCF-GP model to the data and subsequently subtracting it to get the CCF residuals. We then computed the RV residuals by fitting a Gaussian function to each of these residual CCFs, which are represented as coral markers in Fig. 9. The RV residuals have a RMS value of 0.44 m s^{-1} , with an RMS reduction factor of around 4 when compared to the raw solar RVs.

We also calculated the Lomb–Scargle periodograms of the raw and the residual RVs, shown in Fig. 10. The periodograms are normalized as half of the square root of the power spectrum density (PSD) values, and are thus called the semi-amplitude spectrum. Grey

vertical markers indicate the positions of the 27 d solar-rotation signal and their harmonics. Most of the frequency components, especially the ones related to the solar rotation period, are significantly reduced in the residual RVs.

We tested two scenarios by injecting a planet into the HARPS-N solar CCFs with semi-amplitudes of 1.0 and 0.3 m s^{-1} . In both cases, the orbital period was set to 33 d, and the phases were arbitrarily set to 0.7π rad. The phase is defined with reference to $\text{BJD} = 2450000$. These correspond to planets of masses around 5 and 1.5 Earth masses, respectively. Fig. 9 shows the RVs of the planet for both semi-amplitude settings in blue and orange. We then applied the CCF-GP framework on the planet-injected CCF. All priors on the free parameters are uninformative, which means this is a blind search (though we only tested a one-planet model).

Fig. 11 shows the posterior distribution of the extracted planetary parameters, with vertical red lines indicating the true values from the injections. For the 1.0 m s^{-1} case, all the planet parameters are accurately recovered. For the 0.3 m s^{-1} case, while the period and phase are well-retrieved, the semi-amplitude starts to be overestimated, with a derived value of $K = 0.42^{+0.04}_{-0.04} \text{m s}^{-1}$. This discrepancy can be attributed to multiple factors. As already discussed in Section 3.1, one plausible explanation is the overlooked non-Gaussianity of the CCF time series. Another factor is the presence of physical effects in the solar data that are not accounted for in our model, such as the bisector shape alterations at active regions, and the effects of supergranulation, etc.

We also tested a case specifically targeting long-period planets. For this case, we utilized the same HARPS solar RV data set as in the previous cases, but without accounting for the hypothetical star's observability constraints. The data contains 290 observations, with an average sampling rate of one data point per 1.7 d. This number of observations would take around four or five years to gather on the hypothetical star considered in the previous test. Once a full decade of HARPS-N solar data becomes publicly available, we will be able to simulate observations from a survey such as the Terra Hunting Experiment (Hall et al. 2018) more realistically.

For this test, we injected a planet with an orbital period of 101.2 d and a semi-amplitude of 0.5 m s^{-1} , corresponding to a planet of 3.6 Earth masses. The corresponding raw and planet RVs are shown in the left-hand panel of Fig. 12. The results of applying the CCF-GP model, as shown in the right-hand panel of the same figure, reveal that the recovered planet parameters align well with the true values.

3.3 Application to real observations

We applied the CCF-GP framework to HARPS observations of HD 13808 to demonstrate its efficacy on real-world observations with an average SNR of ~ 100 . HD 13808 is a system consisting of a K2V dwarf star orbited by two confirmed non-transiting Neptune-mass planets (Ahrer et al. 2021). The data, which are queried from the ESO archive, include 246 spectroscopic measurements taken with HARPS, mostly between 2005 and 2016. The data underwent pre-processing via the automated DRS pipeline. Fig. 13 presents the series of reference-subtracted CCFs and the associated SNR per pixel (in echelle order 60). The SNR for these observations varies between approximately 50 and 180, with a median value of around 100. We applied the CCF-GP framework, incorporating a two-planet RV model. As in previous tests, all priors on free parameters remained uninformative, implying that this was a blind detection test.

In Fig. 14, the posterior distributions of the extracted planetary parameters are displayed. For real observations, a definitive ‘true’ value remains elusive. Therefore, we compared the extracted values

with those reported in two recent publications, Ahrer et al. (2021) and Delisle et al. (2022). Delisle et al. (2022) implemented a modified version of the multi-GP with a fast MEP kernel. Their best-fit values for the planet parameters are indicated by the orange markers. Conversely, Ahrer et al. (2021) evaluated various models for activity, endorsing that the harmonic activity model (model 1 in Fig. 14), which simultaneously fits the 1st, 2nd, and 3rd harmonic of rotation period to the BIS and RVs, and the multi-GP model (model 2) as two of the most robust models; their results are indicated by purple and coral markers. The derived semi-amplitudes for planets b and c are $K_b = 3.60^{+0.12}_{-0.11} \text{m s}^{-1}$ and $K_c = 1.94^{+0.12}_{-0.12} \text{m s}^{-1}$, which are within one sigma deviation from the values reported in Ahrer et al. (2021) and Delisle et al. (2022). This example demonstrates that the CCF-GP framework can effectively work with data of averaged SNR of around 100, accommodating the majority of current observations and producing results consistent with existing approaches.

4 CONCLUSIONS

In this paper, we introduced the CCF-GP framework, a novel approach developed to model stellar activity signals within time-series of CCFs using GPs. This model addresses the distortions in CCFs due to activity regions on the stellar surface, i.e. spots and faculae, stemming from both the photometric effect and the inhibition of convective blueshift effect.

Upon testing on both synthetic and real data sets, with SNRs down to around 100, the framework showcased its robustness in distinguishing between activity and planetary signals, even when they share identical periods. Planet injection/recovery tests on the realistically sampled HARPS-N solar data highlight the framework's potential to push the detection limit of planetary signals down to $K = 0.3 \text{m s}^{-1}$ at $P = 33$ d (corresponding to a planet of 1.5-Earth mass) during the Sun's high activity phase in its magnetic cycle. The framework surpasses the performance of multi-GP models on RV and activity indicators in the tests conducted. We note that similar to the multi-GP approach, CCF-GP also requires adequate temporal sampling relative to the rotation period of the star for optimal performance.

In the future development of this model, we plan to explicitly address the issue of non-Gaussianity mentioned in Section 2.2.4. One approach as a substitute for using the current approximate GP model is to place a GP prior on F and use, e.g. Markov chain Monte Carlo (MCMC) methods to sample from it and then propagate the samples to approximate the exact stochastic process of the CCF model. However, this comes at the cost of increased computational demands.

Additionally, there is significant potential in enriching the model with more detailed physics. One would be the replacement of the current QP or MEP kernels with kernels that are obtained from physical parameters, such as the ‘impulse response’ GP kernel developed in Hara & Delisle (2023). Furthermore, the model could benefit from incorporating the differential responses of spectral lines to activity as well as the effects of supergranulation. The inherent flexibility of the model's mathematical framework ensures that these integrations can be implemented seamlessly.

ACKNOWLEDGEMENTS

HY, SA, and BK acknowledge funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No 865624, GPRV). MC acknowledges the SNSF support under grant P500PT_211024.

DATA AVAILABILITY

This work was based entirely on publicly available data downloaded from the Data Analysis Center for Exoplanets (DACE)⁵ and the ESO Science Archive.⁶

REFERENCES

- Ahrer E. et al., 2021, *MNRAS*, 503, 1248
 Aigrain S., Pont F., Zucker S., 2012, *MNRAS*, 419, 3147
 Al Moulla K., Dumusque X., Figueira P., Lo Curto G., Santos N. C., Wildi F., 2023, *A&A*, 669, A39
 Astropy Collaboration et al., 2013, *A&A*, 558, A33
 Astropy Collaboration et al., 2018, *AJ*, 156, 123
 Astropy Collaboration et al., 2022, *ApJ*, 935, 167
 Baranne A. et al., 1996, *A&AS*, 119, 373
 Barragán O., Aigrain S., Rajpaul V. M., Zicher N., 2022, *MNRAS*, 509, 866
 Boisse I., Bouchy F., Hébrard G., Bonfils X., Santos N., Vauclair S., 2011, *A&A*, 528, A4
 Boisse I., Bonfils X., Santos N. C., 2012, *A&A*, 545, A109
 Collier Cameron A. et al., 2019, *MNRAS*, 487, 1082
 Collier Cameron A. et al., 2021, *MNRAS*, 505, 1699
 Cosentino R. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, SPIE Conf. Ser. Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. SPIE, Bellingham. p. 84461V
 Crass J. et al., 2021, preprint (arXiv:2107.14291)
 Cretignier M., Dumusque X., Allart R., Pepe F., Lovis C., 2020, *A&A*, 633, A76
 Cretignier M., Dumusque X., Pepe F., 2022, *A&A*, 659, A68
 de Beurs Z. L. et al., 2022, *AJ*, 164, 49
 Delisle J. B., Unger N., Hara N. C., Ségransan D., 2022, *A&A*, 659, A182
 Di Maio C. et al., 2024, *A&A*, 683, A239
 Donati J. F., Semel M., Carter B. D., Rees D. E., Collier Cameron A., 1997, *MNRAS*, 291, 658
 Donati J. F. et al., 2014, *MNRAS*, 444, 3220
 Dumusque X., 2018, *A&A*, 620, A47
 Dumusque X., Boisse I., Santos N. C., 2014, *ApJ*, 796, 132
 Dumusque X. et al., 2015, *ApJ*, 814, L21
 Dumusque X. et al., 2021, *A&A*, 648, A103
 Fischer D. A. et al., 2016, *PASP*, 128, 066001
 Giorgini J. D. et al., 1996, *BAAS*, 28, 1158
 Gray D. F., 2008, *The Observation and Analysis of Stellar Photospheres*. Cambridge University Press, Cambridge, UK
 Hall R. D., Thompson S. J., Handley W., Queloz D., 2018, *MNRAS*, 479, 2968
 Handley W. J., Hobson M. P., Lasenby A. N., 2015a, *MNRAS*, 450, L61
 Handley W. J., Hobson M. P., Lasenby A. N., 2015b, *MNRAS*, 453, 4384
 Hara N. C., Delisle J.-B., 2023, preprint (arXiv:2304.08489)
 Haywood R. D. et al., 2014, *MNRAS*, 443, 2517
 John A. A., Collier Cameron A., Wilson T. G., 2022, *MNRAS*, 515, 3975
 Jones D. E., Stenning D. C., Ford E. B., Wolpert R. L., Loredó T. J., Gilbertson C., Dumusque X., 2017, preprint (arXiv:1711.01318)
 Klein B., Donati J.-F., Hébrard É. M., Zaire B., Folsom C. P., Morin J., Delfosse X., Bonfils X., 2021, *MNRAS*, 500, 1844
 Klein B. et al., 2022, *MNRAS*, 512, 5067
 Klein B. et al., 2024, *MNRAS*, 531, 4238
 Kochukhov O., Makaganiuk V., Piskunov N., 2010, *A&A*, 524, A5
 Liang Y., Winn J. N., Melchior P., 2024, *AJ*, 167, 23
 Lienhard F., Mortier A., Cegla H. M., Cameron A. C., Klein B., Watson C. A., 2023, *MNRAS*, 522, 5862
 Lovis C., 2007, PhD thesis, University of Geneva, Switzerland
 Meunier N., Lagrange A. M., 2019, *A&A*, 625, L6
 Meunier N., Lagrange A. M., Borgniet S., Rieutord M., 2015, *A&A*, 583, A118
 Nardiello D. et al., 2022, *A&A*, 664, A163
 Noyes R. W., Hartmann L. W., Baliunas S. L., Duncan D. K., Vaughan A. H., 1984, *ApJ*, 279, 763
 O’Sullivan N. K., Aigrain S., 2024, *MNRAS*, 531, 4181
 Pepe F., Mayor M., Galland F., Naef D., Queloz D., Santos N. C., Udry S., Burnet M., 2002, *A&A*, 388, 632
 Pepe F. et al., 2021, *A&A*, 645, A96
 Petit P. et al., 2015, *A&A*, 584, A84
 Phillips D. F. et al., 2016, in Navarro R., Burge J. H., eds, SPIE Conf. Ser. Vol. 9912, Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II. SPIE, Bellingham, 99126Z
 Rajpaul V., Aigrain S., Osborne M. A., Reece S., Roberts S., 2015, *MNRAS*, 452, 2269
 Rajpaul V. M., Aigrain S., Buchhave L. A., 2020, *MNRAS*, 492, 3960
 Shahaf S., Zackay B., 2023, *MNRAS*, 525, 6223
 Virtanen P. et al., 2020, *Nature Methods*, 17, 261
 Yu L. et al., 2017, *MNRAS*, 467, 1342
 Zhao L. L. et al., 2022, *AJ*, 163, 171
 Zhao Y. et al., 2024, *A&A*, 687, A281

⁵<https://dace.unige.ch/dashboard/>

⁶<https://archive.eso.org/cms.html>

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.