UNIVERSITY^{OF} BIRMINGHAM University of Birmingham Research at Birmingham

Saltire

Sebastian, Daniel; Triaud, Amaury; Brogi, Matteo

DOI: 10.1093/mnras/stad3765

License: Creative Commons: Attribution (CC BY)

Document Version Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Sebastian, D, Triaud, A & Brogi, M 2023, 'Saltire: A model to measure dynamical masses for high-contrast binaries and exoplanets with high-resolution spectroscopy', *Monthly Notices of the Royal Astronomical Society*, vol. 527, no. 4, stad3765, pp. 10921–10936. https://doi.org/10.1093/mnras/stad3765

Link to publication on Research at Birmingham portal

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

•Users may freely distribute the URL that is used to identify this publication.

Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)

•Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

D. Sebastian^(b),¹* A. H. M. J. Triaud^(b) and M. Brogi^(b),^{2,3,4}

¹School of Physics & Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK
 ²Dipartimento di Fisica, Universitá degli Studi di Torino, Via Pietro Giuria 1, I-10125 Torino, Italy
 ³Department of Physics, University of Warwick, Coventry CV4 7AL, UK
 ⁴INAF – Osservatorio Astrofisico di Torino, Via Osservatorio 20, I-10025 Pino Torinese, Italy

Accepted 2023 December 2. Received 2023 November 2; in original form 2023 August 7

ABSTRACT

High-resolution cross-correlation methods are widely used to discover and to characterise atomic and molecular species in exoplanet atmospheres. The characteristic cross-correlation signal is typically represented as a function of the velocity of the system, and the semi-amplitude of the planet's orbit. We present Saltire, a fast and simple model that accurately reproduces the shape of such cross-correlation signals, allowing a direct fit to the data by using a minimum set of parameters. We show how to use this model on the detection of atmospheric carbon monoxide in archival data of the hot Jupiter τ Boötis b, and how Saltire can be used to estimate the semi-amplitude and rest velocity of high brightness ratio binaries. By including the shape of the signal, we demonstrate that our model allows to robustly derive the signal position up to 10 times more accurate, compared to conventional methods. Furthermore, we discuss the impact of correlated noise and demonstrate that Saltire is a robust tool for estimating systematic uncertainties on the signal position. Saltire opens a new door to analyse high signal-to-noise data to accurately study atmospheric dynamics and to measure precise dynamical masses for exoplanets and faint stellar companions. We show that the phase-resolved shape of the atmospheric cross-correlation signal can accurately be reproduced, allowing studies of phase-dependent signal changes and to disentangle them from noise and data aliases.

Key words: techniques: spectroscopic – exoplanets – planets and satellites: atmospheres – binaries: spectroscopic.

1 INTRODUCTION

In the past decades, many tools have been developed to characterize planetary atmospheres using ground- and space-based spectrophotometry, as well as ground-based high-resolution spectroscopy. One important and widely used method to detect exoplanet atmospheres is high-resolution cross-correlation spectroscopy (HRCCS; Snellen et al. 2010), which has been used to make robust detections of atomic and molecular species within exoplanet atmospheres, and used to produce empirical information about exoplanet atmospheric dynamics. An advantage of this method is that it is not limited to transiting exoplanets.

Using the HRCCS method, planetary atmospheric signals are retrieved by obtaining consecutive spectroscopic measurements of high resolution, either in transmission during the planet's transit or close to its superior conjunction to measure the planet's dayside thermal emission (see Birkby 2018, for a review). The significance of a detection mainly depends on two aspects: (i) the signal's strength grows with the square root of the number of atmospheric lines that have been correlated, thus favouring spectral observations with wider wavelength range and/or higher spectral resolutions; (ii) the signal's amplitude is a function of the contrast ratio of the exoplanet's

atmosphere to the stellar photosphere, thus favouring close-in gas giants with high equilibrium temperatures ($T_{eq} \ge 2200$ K), so-called ultra-hot Jupiters (UHJs; Parmentier et al. 2018). A crucial aspect of the HRCCS method is that spectra are corrected from telluric and stellar contributions, which is generally possible because the planet signal is moving fast relative to them (by several km s⁻¹) within the span of a typical observation. The planetary signal is thus not correlated to the spectrograph or the stellar rest frame. There are several different methods for telluric and stellar line fitting (e.g. Brogi et al. 2012, 2016), or singular value decomposition (Kalman 1996), which have been applied successfully to many exoplanets and instrumental set-ups (e.g. Birkby et al. 2013; de Kok et al. 2013; Piskorz et al. 2016; Cheverall, Madhusudhan & Holmberg 2023).

https://doi.org/10.1093/mnras/stad3765

1.1 K_c-V_{rest} mapping

In the HRCCS framework, spectra corrected for telluric and stellar contributions are cross-correlated with a model spectrum of the planet's atmosphere or with a model of a specific atomic or molecular species, moved to the planet's rest frame. This method requires the semi-amplitude (K_c) of the planet's orbit to be known exactly, which is usually not the case. Typically, a map of cross-correlation functions (CCF) in the K_c - V_{rest} plane is generated ('CCF map' hereafter), which probes the cross-correlation signal for several possible semi-

Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

^{*} E-mail: D.Sebastian.1@bham.ac.uk

amplitudes and rest velocities (V_{rest} ; e.g. Brogi et al. 2012; de Kok et al. 2013).

Detecting the planet's atmospheric signal in the K_c-V_{rest} plane using HRCCS thus literately turns the planet and its host star into a double-lined spectroscopic binary as the planet's semi-amplitude can be measured. This allows to measure the mass ratio between host star and planet. For transiting planets, this allows even to model independent, dynamical planet masses (Birkby 2018).

Unfortunately, in these maps, K_c is often highly correlated to V_{rest} , which can cause large uncertainties of the planet's semi-amplitude. Typical uncertainties of K_c are of the order of tens of kilometres (e.g. Brogi et al. 2018; Bello-Arufe et al. 2022; Yan et al. 2023). Furthermore, it has been suggested that different species in the planet's atmosphere show different semi-amplitudes and rest velocities, different to the systemic velocity, both in infrared (IR; Brogi et al. 2023) and optical regimes (e.g. Nugroho et al. 2020). Atmospheric models reveal that the position and shape of the atmospheric signal in the CCF map can be used to disentangle the temperature structure, rotation, and dynamics of a planet's atmosphere, especially during transit observations (e.g. Wardenier et al. 2021; Kesseli et al. 2022). Thus, characterizing the accurate position of the planetary signal within the CCF map is essential to empirically understand these processes and essential to place constraints on atmospheric models for close-in giant planets.

Typically, the planet's semi-amplitude and rest velocity are determined either by fitting a Gaussian to the cross-correlation map at the maximum significance level (e.g. Prinoth et al. 2022), or by statistically analysing the maximum likelihood of the detection (e.g. Brogi & Line 2019; Brogi et al. 2023). Detections with state-ofthe-art instrumentation have a generally low detection significance (typically < 10 σ). Future instruments on large telescopes – like the Extremely Lage Telescope (ELT) – will result in a higher detection significance of atmospheric signals in the CCF maps and thus will allow to disentangle the planet's orbit from dynamics in the atmosphere.

1.2 Binaries as targets of opportunity

A typical hot Jupiter in a very close orbit would have a contrast ratio of the order of 10^{-4} to 10^{-5} between the planetary atmosphere and the stellar photosphere. In Sebastian et al. (inprepration), we showed that eclipsing binaries of F-, G-, and K-type stars, orbited by late-type M-dwarf companions, can be used to test the achievable accuracy of the HRCCS method that is required to derive dynamical masses. First, the contrast ratio between both stars is 10^{-3} to 10^{-4} , which is comparable to exoplanet atmospheres, making such binaries good proxies to exoplanets. Secondly, the spectrum of an M-star is well known, and we can use optimized model spectra, observed highresolution spectra, or even line masks – optimized for M-dwarfs – as templates to generate CCF maps.

A large sample of hundreds of eclipsing binaries with low mass (EBLM) companions, originally detected by the Wide Angle Search for Planets (WASP) survey (Pollacco et al. 2006), has been spectroscopically observed in the EBLM project (Triaud et al. 2013, 2017). The faint secondary implies that the binary is a single-lined binary. This allows to use stabilized spectrographs and line lists, just like for single stars, to derive very precise radial velocities of the primary. The goal of this project is to derive precise orbital parameters from radial velocity (RV) observations in order to measure fundamental parameters for the secondary stars such as precise mass, and radius of low-mass M-dwarfs (e.g. Gill et al. 2019; von Boetticher et al. 2019; Swayne et al. 2021; Sebastian et al. 2023). As single-lined

systems, only parameters relative to the primary can be obtained, but as double-lined systems, which the HRCCS method gives access to, absolute parameters can be measured. In addition, these binaries are monitored to search for circumbinary exoplanets (Triaud et al. 2022; Standing et al. 2023). We note that binaries of any type, with large brightness differences can benefit from our approach (e.g. Sun-like companions to red giants).

In this paper, we explore how to model the shape of a CCF map from a typical HRCCS observation. We introduce the Saltire model, a versatile model that allows (i) to fit an HRCCS CCF map and determine accurate parameters for K_c and V_{rest} and (ii) to predict CCF maps from an HRCCS observation of exoplanet atmospheres.

2 MODELLING HRCCS OBSERVATIONS

To model a CCF map, we first need to understand how it is created. It is a superposition of all spectroscopic observations that (i) have been post-processed by correcting for telluric and stellar lines from the host star and (ii) have been cross-correlated by a template spectrum. This is similarly true when analysing faint companions for EBLM binaries as it is for exo-atmospheric observations in transmission or dayside emission spectroscopy. The shape of this superposition depends on two main effects: (i) Changes of the relative contribution of the atmosphere for each measurement. This can be caused not only by intrinsic processes, e.g. due to secondary eclipses, phase curves, or temperature gradients during transit observations, but also due to changing observing conditions or artefacts related to post-processing of the data, like phase-dependent stretching of the atmospheric signal due to the telluric and host star removal processes. (ii) The orbital reflex motion of the companion, which is moving on a Keplerian orbit.

In the first step, we can assume negligible changes in the relative contribution. This is particularly true for high-contrast binaries where the companion is an M-dwarf, and where the phase-dependent contribution of its absorption spectrum can be assumed to be constant and known to a first order (neglecting stellar variability, reflected light, and ellipsoidal contributions for now). In this case, the superposition of all cross-correlated spectra will reach a maximum CCF signal in the companion's rest frame, hence when all spectra have been corrected according to the companion's reflex motion. Typically, in the HRCCS method, the orbit parameters are assumed to be fixed (neglecting uncertainties of the orbital parameter determination). In case of close-in giant planets, this is even simplified to a fixed sinusoidal motion.

The orbital parameters of binaries from the EBLM project are typically very well determined from high-resolution spectroscopy. Assuming two-body motion, the companion shares all Keplerian orbital parameters, except the companion's argument of periastron (ω_c) and semi-amplitude (K_c) . The parameter ω_c can be derived from the known primary orbit parameters: $\omega_c = \omega - \pi$. The barycentric radial velocity of the companion will then take the form

$$V_{r,c} = K_c[\cos(\nu + \omega_c) + e\cos(\omega_c)], \tag{1}$$

with ν the true anomaly being determined by the period (*P*), the time of periastron ($T_{0,peri}$), the eccentricity (*e*), and ω . K_c is not known at first, since it depends on the exact knowledge of the binaries' mass ratio. For this reason, CCF maps typically scan for a range of K_c , which includes the expected semi-amplitude of the companion. To demonstrate this, we simulate 123 observations with the Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations (ESPRESSO, Pepe et al. 2010) of the M-dwarf companion of the EBLM binary TOI-1338/BEBOP-1



Figure 1. Upper panel: CCFs for spectra, shifted to match different semiamplitudes of the secondary. Lower panel: Combined CCFs for each semiamplitude.

(Standing et al. 2023) covering its full orbital period (Obs1 hereafter). A simulation has the advantage of avoiding to include the spectrum of the primary and the need to include any post-processing in the analysis. The M-dwarf companion is simulated with PHOENIX model spectra (Husser et al. 2013) with $T_{\text{eff}} = 3300 \text{ K}$, log $g_{\star} = 5.0$, and [Fe/H] = 0.0. We first correct wavelengths from vacuum to air, following equation (10) from Husser et al. (2013), and then we match the spectral resolution to ESPRESSO ($R \sim 140\,000$) using the implementation in the ISPEC package (Blanco-Cuaresma et al. 2014), which convolves the spectra with a Gaussian kernel. Each spectrum is then shifted to match the companions' reflex motion using equation (1) as well as the primary's stellar and orbit parameters P = 14.608558 d, $T_{0,\text{peri}} = 2458206.16755$ BJD, e = 0.155522, and $\omega = 2.05572$ rad (Standing et al. 2023), and the 'true' semiamplitude $K_{c,true} = 77.84 \text{ km s}^{-1}$, estimated from the mass ratio (q = M_c/M_1) and the primary's semi-amplitude ($K_{c,true} = K_1 q$). Finally, we cross-correlate the model spectra with an M2-dwarf line list, a type typically used for high-precision radial velocity measurements implemented by the ESPRESSO data reduction pipeline.¹ For the cross-correlation, we convert the spectra to have a uniform sampling of $500 \,\mathrm{m \, s^{-1}}$, which is close to the pixel resolution of the onedimensional (1D) ESPRESSO spectra and we sample the CCF for a velocity range between -30 and $30 \,\mathrm{km \, s^{-1}}$ with a spacing of $1.5 \,\mathrm{km}\,\mathrm{s}^{-1}$

Fig. 1 shows the CCFs for each data point when shifting the spectra for $K_c = 76.7, 73.6$, and 64.3 km s^{-1} , respectively. The CCFs are less aligned to the companion's rest velocity, the further K_c differs from $K_{c,true}$. When we average each CCF to a combined CCF, a decrease of the combined CCF contrast as well as an asymmetry in the profile are clearly noticeable. By design, the width and shape of the combined CCF are defined by the shape of the single CCFs, as well as the average velocity shift of all spectra. This average velocity shift depends on the orbital parameters as well as on the phases the orbit is sampled. For a well-sampled circular orbit, the width is directly correlated to $K_c-K_{c,true}$. The asymmetry of the profile is thus primary caused by the small orbital eccentricity of the simulated binary.



Figure 2. CCF map shapes from simulated EBLM observations and their dependence from phase coverage. Upper panel: Simulated reflex motion of an EBLM binary. Black dots: Primary star. Blue crosses: Simulated secondary. Left: Full data set (Obs1). Right: Partial sample of the same data set (Obs2). Middle panels: CCF maps for both observations, each composed of combined CCFs for a range of K_c values, of the M-dwarf companion. Dashed lines and labels: Position of combined CCFs for K_c , used in Fig. 1. Lower panels: Best-fitting Saltire models. Left: Fit for Obs1. Right: Fit for Obs2.

We now produce a CCF map and show it in Fig. 2. This map represents the combined CCF functions for a range of K_c between 55 and 100 km s⁻¹ calculated in steps of 1.5 km s⁻¹. The positions of the three combined CCFs, displayed in Fig. 1, are now part of this map and their positions are highlighted for clarity. The resulting CCF map takes a cross-like shape, a saltire. This process resembles to what happens when optical telescopes are being focused. Far from the solution, the signal is defined and gets 'focused' at the correct semiamplitude and rest velocity. We thus refer to the process of averaging the CCF signal for a certain semi-amplitude as *K*-focusing.

Similar implementations have been described in literature for more than a decade now. For example, Lafarga et al. (2023) describe their 'fast' process, which involves the cross-correlation in the telluric rest frame first, and combining the resulting CCFs after moving them into the companion's rest frame, using a certain semi-amplitude. Independent from the implementation, *K*-focusing is always the underlying process to create CCF maps in the K_c - V_{rest} plane.

On the right panel of Fig. 2, we show how the shape of the CCF map changes when using a subset of 15 observations out of the 123 from the initial simulation that still cover the binary's orbit well (Obs2 hereafter).

As described above, for real observations our assumption of identical contribution does not hold. We thus need to take into account a weighting of the contribution of each observation to the CCF map.

¹The ESPRESSO pipeline has been publicly released on https://www.eso. org/sci/software/pipelines/espresso/espresso-pipe-recipes.html.

This is typically the root mean square (RMS) noise of each spectrum, but could also include for example phase-curve information of a specific molecule. Later in the paper (Section 5), we use actual observations where this effect is accounted for.

3 THE Saltire MODEL

We introduce Saltire,² a simple and scalable model to fit CCF maps, created from HRCCS observations. For this model, we assume that we can analytically express the CCF of the companion's contribution in velocity space at each time of the orbit. For precise radial velocity measurements, the CCF is most often fitted with a simple Gaussian function, which works well for solar-type stars (Baranne et al. 1996; Pepe et al. 2002). For M-dwarfs, side lobes are observed for reasons that are not completely understood and the simple Gaussian assumption does not hold. Bourrier et al. (2018) showed that a double Gaussian fit can improve the CCF fitting of M-dwarf spectra. Side lobes have also been observed for CCF signals of exoplanet atmospheres (see e.g. fig. 8 in Brogi et al. 2016). Furthermore, we show in Sebastian et al. (inprepration) that data post-processing methods, such as singular value decomposition, can introduce a stretching of the signal, which also produces artificial side lobes on the CCF function. We thus simulate the CCF signal as a function of the relative rest velocity using a double Gaussian defined as

$$f(v) = h + (A_1 \mathcal{N}_1(\mu, \sigma_1) + A_2 \mathcal{N}_2(\mu, \sigma_2)),$$
(2)

with the mean height of the CCF outside the signal h, N_1 and N_2 the two Gaussian functions with respective intensities A_1 and A_2 and standard deviations σ_1 and σ_2 at identical velocity (μ). To create the side lobes, A_1 and A_2 must have opposite signs. By fixing the velocity shift between the core (1) and the lobe component (2), we create an asymptotic problem. When both amplitudes are allowed to increase, they would maintain a similar differential amplitude. This causes an increase of the calculation time for a least-squares minimization algorithm. We avoid this by parametrising both amplitudes as

$$A_1 = \frac{\Sigma}{(\Delta+1)}; \ A_2 = \begin{cases} 0, & \text{if } \Delta = 0, \\ \frac{\Sigma}{\left(\frac{1}{\Delta}+1\right)}, & \text{otherwise.} \end{cases}$$
(3)

Hereby, we introduce the quotient $\Delta = A_2/A_1$, as well as the sum $\Sigma = A_1 + A_2$ of both intensities. The parameter Σ now represents the effective contrast of the CCF, which is independent from the amplitude of the side lobes. It is negative for expected absorption and positive for emission spectra. We define Δ within the range $-1 < \Delta \le 0$, which ensures that both amplitudes have an opposite sign. While $\Delta = 0$ will basically turn the function into a single Gaussian, $\Delta \sim -1$ would describe the asymptotic case with both amplitudes being very large. By setting priors for Δ , we can thus improve the convergence of the fit.

Similar to Fig. 1, we can now derive the CCF signal f(v) for each observation time. The relative rest velocity of the CCF signal can be derived from the known Keplerian orbital parameters and equation (1) as

$$v = V_{r,c}(K_c) - V_{r,c}(K_{c,true}) + V_{rest},$$
 (4)

for a given K_c . This adds two free parameters to our model, the true semi-amplitude ($K_{c,true}$) and the rest velocity (V_{rest}) of the detection. Combining the CCF functions f(v) for all observation times, we can derive a combined CCF function $f'(K_c)$ for each sampled value of K_c .

This process is identical to the *K*-focusing, described in Section 2. Therefore, we can combine all sampled $f'(K_c)$ into a model CCF map. This step is by design identical to what we described for simulated combined CCFs in Section 2. This method assumes that all RV trends, e.g. due to a third body in a wider orbit, have been removed prior to the analysis. A good understanding of the primary star's radial velocity trends is typically the case for high-precision spectroscopy of binaries or planet host stars.

Our definition of V_{rest} as the rest velocity of the detection means that the systemic velocity (V_{sys}) has not been removed. Since it is a free parameter in this model, it can simply be applied to the often-used velocity definition with V_{sys} being removed prior to the measurement.

3.1 Fitting CCF maps made easy

The Saltire model can be used to fit CCF maps. We implement the model as a residual function using the lmfit framework to apply least-squares fitting (Newville et al. 2016). In this section, we summarize the basic usage of the model and input parameters. Due to the model construction, the only free parameters are K_c , V_{rest} , and the parameters of the 1D function f(v), which are h, Σ , Δ , σ_1 , and σ_2 . Fixed parameters are simply identical to the parameters that have been used to create the CCF map of the simulated data described earlier (Obs1 and Obs2). Here, we divide three sets of input parameters: (i) the CCF map dimensions, which are the sampled velocity and semi-amplitude ranges $[V_{rest}]$ and $[K_c]$, respectively, (ii) the observation contribution parameters (obs), which are observation times and weights for each observation, and (iii) the fixed orbit parameters (fixpar), which include the period (P), time of periastron ($T_{0,peri}$), eccentricity (e), and argument of periastron (ω). The fixed orbit parameters are assumed to be arrays with length of the semi-amplitude range $[K_c]$, to allow fitting of possible different orbit parameters for each sampled semi-amplitude. Additionally, Saltire has a built-in ability to replace the double Gaussian function by a double Lorentzian function. Due to similarities between the parameters of both functions, the change can be done by simply setting the func keyword.

To apply priors, we make use of the lmfit param object. For Δ , the maximal value is zero (no side lobes) and the minimal value should be close, but > -0.5 (strong side lobes) to avoid the fit entering the asymptotic space. We also ensure that the expected width of the side lobes σ_2 is always larger than the line width σ_1 . Despite these recommendations, all free parameters and their limits can be freely selected prior to the fit.

The fitting is implemented by using the saltirise function. It returns the residuals normalized by the CCF errors (data model)/error as a flattened array for each combination of V_{rest} , and $K_{\rm c}$, or the initial model, if no data are provided. Despite, the leastsquares fit as implemented in lmfit allows to derive correlations and error propagation for each free parameter, Saltire also implements a run_MCMC function allowing a statisitical evaluation of the fit. This is to sample the posterior probability distribution of the Saltire free model parameters, using the Markov chain Monte Carlo (MCMC) code EMCEE (Foreman-Mackey et al. 2013). With the MCMC, an extra parameter is fitted, the CCF-jitter (σ_{iit}) of the data that optimally weights the fit in the log-likelihood function. The jitter is a way to avoid underestimated uncertainties, and accounts for RMS noise, as well as for correlated signals like spurious CCF correlations of stellar lines and/or post-processing artefacts. Uniform priors can be set using the same param object described above, but adding σ_{iit} as additional parameter.

²Saltire PYTHON code and documentation are available on Github.

4 APPLICATION TO SIMULATED HRCCS EBLM DATA

We estimate how well the model fits CCF maps by using two simulated data sets of a high-contrast binary, as described in Section 2, Obs1 and Obs2. We first do a least-squares minimization, followed by an MCMC sampling. As starting parameters, we use input values for K_c and V_{rest} and CCF function parameters listed in Table 1 with wide uniform priors. The least-squares minimization results are used as starting parameters for the MCMC sampling, also using wide uniform priors. We run 42 parallel chains with 4000 samples each. The first 1500 samples of each walker are rejected (the burn-in) and the resulting samples are thinned by a factor 5, resulting in a posterior distribution of 21 000 independent samples for each parameter. The reported parameters are derived as the 50th percentile, and errors by averaging the 15.8655/84.1345 percentiles of the posterior distribution.

Fig. 2 shows Saltire models generated for both data sets. By its very design, the model computes different CCF map shapes corresponding to different data sampling and phase coverage. Table 1 shows the results of the CCF map fit with Saltire for both observations. The resulting MCMC samples for both sets of simulated observations are displayed in Figs A1 and A2.

Fig. 3 reveals the shape of the function by performing cuts through the CCF maps as well as the best-fitting Saltire models. Cuts for constant semi-amplitudes are shown for (a) $K_c = 76.7 \text{ km s}^{-1}$, (b) 73.6 km s⁻¹, and (c) 64.3 km s⁻¹. Since we use synthetic spectra, no side lobes are generated. Due to the absence of side lobes, the parameters Δ and σ_2 are not well constrained. We find that repeating the fit with both parameters fixed does not change the result for the other parameters and uncertainties. A cut is shown for constant rest velocity close to the maximum CCF contrast (d) $V_{\text{rest}} = 0 \text{ km s}^{-1}$. This step is used in most studies to determine K_c . Fig. 4 shows the same for Obs2, highlighting the shape differences for both samples, which can best be seen for cuts b and c, hence for larger deviations from $K_{c,\text{true}}$.

The measurements for K_c returned by the model agree with the inserted value $K_{c,true}$ for the simulated observations, within 1σ . Despite being in statistical agreement, we note that there is a small difference of -12 m s^{-1} between Obs1 and Obs2. We also find that the values returned for V_{rest} by the model show a difference of 24 m s^{-1} (2.6σ). The origin of these discrepancies is discussed in Section 4.1. We note that the measured V_{rest} shows a similar offset of about 450 m s^{-1} for both Obs1 and Obs2. This offset is not surprising, since (i) the model spectra have been corrected from vacuum wavelengths, and (ii) the M-dwarf line list has been compiled using observed spectra. Both effects can be credited for this systematic offset.

To test the influence of the shape difference between Obs1 and Obs2, we used a simple 1D Gaussian fit to measure K_c for the rest velocity of maximum CCF contrast (cut d). This test also simulates what is traditionally employed within the literature to measure K_c from CCF maps. The best-fitting value for K_c we find this way is 77.786 \pm 0.088 and 77.694 \pm 0.069 km s⁻¹ for Obs1 and Obs2, respectively. These measurements are 54 m s⁻¹ (0.6 σ) and 146 m s⁻¹ (2.1 σ) from the inserted value $K_{c,true}$, respectively. Thus, we can confirm that the shape difference between both data sets results in a systematic deviation, even at maximum CCF contrast, which – if not modelled properly – will result in discrepant measurements that can be of the order of 10 times larger, compared to what we can achieve with the current implementation of Saltire.

4.1 K-focusing of CCF wiggles

Since the simulated observations are by definition noiseless, we ought to retrieve exact values, for both observations, but we do not. Therefore, the differences we retrieve must be created by the analysis. Spurious CCF correlations between the template and other parts of the spectrum are a prime suspect. They are a feature of CCFs (nicknamed *wiggles*; Lalitha et al. 2024).

In real observations, these wiggles are quasi-stationary, thus show only small changes due to changing observing conditions. In our simulated observations, these are exactly identical, for each observation. When creating the CCF map, these wiggles undergo the same *K*-focusing process (see Section 2), as the main CCF signal. We can now use these simulated observations to study the effects introduced by wiggles. The simulated stationary wiggles can be seen in the residuals panels of Fig. 3 for Obs1 and Fig. 4 for Obs2. They show a decreasing amplitude for larger deviations from $K_{c,true}$ (from a to c). This is expected from the *K*-focusing process, which basically causes a smearing by averaging them for increasing RV shifts. We also note that the overall shape of the wiggles for Obs1 and Obs2 appear very similar. This is caused by the fact that the orbit parameters and selected range for K_c are identical for both observations.

A crucial aspect is that the shape – originated by the *K*-focusing process – is also dependent on the phase coverage. Since both observations have a different phase coverage – similar to the companions' CCF shape – the wiggles must be different for both observations. Fig. 5 shows the difference between the residual maps of Obs1 and Obs2 with the inserted value of $K_{c,true}$ highlighted. Wiggles of the CCF map are identical at $K_{c,true}$, for both observations, thus are largely removed in this position, but differ increasingly – due to *K*-focusing – with larger semi-amplitude deviations. Fig. 5 thus shows that residual correlations caused by the wiggles must be the dominating structure.

This explains why we measure slightly different values for K_c and V_{rest} : The differences shown in Fig. 5 are currently not modelled by Saltire. Nevertheless, we show how 1D wiggles at $K_{c,true}$ propagate in the two-dimensional (2D) CCF map due to the K-focusing process. Modelling the 2D wiggles is not the goal of this paper, but future applications to the Saltire model could include a model of these quasi-stationary wiggles into the CCF function f(v) and thus use them as an analytical tool to disentangle between the CCF signal and other aliases from post-processing with the goal to measure yet more accurate parameters of the companion, using CCF maps.

5 APPLICATION TO EXOPLANET HRCCS OBSERVATIONS

Providing a framework to accurately measure atomic species and molecules in exoplanet atmospheres will be one of the main use cases for Saltire. We show its application to the close-in gas giant planet τ Boötis b. Brogi et al. (2012) (B12 hereafter) reported the detection of carbon monoxide (CO) in the atmosphere of this planet by observing the dayside emission of the planet. They used observation with the cryogenic high-resolution cross-dispersed infrared echelle spectrograph (CRIRES, Kaeufl et al. 2004) with a wavelength range of 2.287–2.345 µm and a resolution of $R \sim 100\,000$ to monitor the planet during three nights. These data cover the orbital phase before, close to, and after the phases of superior conjunction. The data were

Parameters	Ob	os1	Obs2		Input fit parameter
	Least squares	MCMC	Least squares	MCMC	
$K_{\rm c}({\rm kms^{-1}})$	77.841 ± 0.011	77.840 ± 0.012	77.828 ± 0.013	77.828 ± 0.013	77.840 (inserted)
$V_{\rm rest}({\rm kms^{-1}})$	0.4691 ± 0.0088	0.4690 ± 0.0083	0.4448 ± 0.0092	0.4446 ± 0.0090	0.0 (inserted)
Σ	$(-3.506 \pm 0.016)e - 04$	$(-3.510 \pm 0.015)e - 04$	$(-3.504 \pm 0.014)e - 04$	$(-3.507 \pm 0.013)e - 04$	-3e-04
Δ	-0.42 ± 2.31	-0.26 ± 0.14	-0.49 ± 1.04	-0.31 ± 0.13	-5e-3
$\sigma_1 ({\rm km}{\rm s}^{-1})$	2.32 ± 0.56	2.267 ± 0.055	2.38 ± 0.28	2.314 ± 0.049	2.4
$\sigma_2 (\mathrm{km}\mathrm{s}^{-1})$	2.6 ± 1.4	2.71 ± 0.18	2.66 ± 0.60	2.78 ± 0.15	4.58
h	$1 + (8.52 \pm 0.16)e - 6$	$1 + (8.53 \pm 0.15)e - 6$	$1 + (8.31 \pm 0.17)e - 6$	$1 + (8.23 \pm 0.17)e - 6$	1.0
$\sigma_{ m jit}$	-	$(3.729 \pm 0.079)e - 06$	-	$(4.155 \pm 0.083)e - 06$	-

Table 1. Saltire fit results for both simulated EBLM data sets.



Figure 3. Upper panel, left: CCF map for Obs1. Dashed lines and labels: Positions of cuts though the data, shown on the right. Right: Cuts through the data, for different semi-amplitudes (a, b, and c), and for the rest velocity of maximum CCF contrast (d). Black dots: Data for simulated data set Obs1. Red line: Best-fitting Saltire model. Blue dashed line: Mark of the input semi-amplitude of the EBLM companion. Lower panel, left: Residual CCF map after removing the Saltire model. Dashed lines and labels: Positions of cuts though the residuals, shown on the right. Right: Cuts through the residuals at similar position as above. Error bars represent uncertainties, derived from 2D Saltire fit to the CCF map.

post-processed identically as in B12, to clean telluric and stellar lines. The orbital radial velocity parameters have been presented in several studies (e.g. Butler et al. 2006; B12; Borsa et al. 2015; Justesen & Albrecht 2019). These also report a radial velocity trend of the F6V host star, due to a wide M-dwarf companion. We use the circular orbital solution presented by B12 for this demonstration as fixed parameters (P = 3.312433 d, $T_{0,peri} = 2455652.108$ HJD). Here, we neglect the expected trend, which is less than 1 m s⁻¹ and thus below the velocity resolution of CRIRES during the short observing period for the B12 data. Similarly to Section 4, we crosscorrelate the post-processed spectra with a line mask. The line mask is generated using line positions from the HITRAN database (Gordon et al. 2017) for ${}^{12}C^{16}O$, which were weighted by the CO model, described by B12. Using a CO line mask allows us to derive the mean line profile of the exoplanet atmosphere without prior assumption of its shape. We sample K_c from 50 to 170 km s⁻¹ in steps of 1.5 km s⁻¹, and V_{rest} from -70 to 50 km s⁻¹ in steps of 1.5 km s⁻¹. Our selected step width corresponds to the pixel resolution of the CRIRES data. Both ranges are selected to include the position of the CO detection, reported in B12 ($K_c = 110 \pm 3.2 \text{ km s}^{-1}$), as well as the systemic velocity of the host star reported in Donati et al. (2008) $(V_{\text{sys}} = -16.4 \pm 0.1 \text{ km s}^{-1})$.³ We co-add the CCF maps from different CRIRES detectors similarly to B12 by weighting them by the relative line intensities of the CO lines in each detector.

We obtain the best-fitting Saltire model to the combined CCF map by least-squares minimization using the input parameters given in Table 2. The resulting fit parameters are used as starting parameters to sample the posterior distribution using the MCMC with the same number of walkers and samples, as described in Section 4. The 50th percentile and average errors from the posterior distributions are presented in Table 2. For the parameters K_c and V_{rest} , we additionally report systematic uncertainties in excess of the MCMC errors (see Section 5.1 for a detailed discussion).

Fig. 6 shows the resulting CCF map for the combined observations as well as the best-fitting Saltire model. We can use the bestfitting parameters to estimate the significance of the detection as the CCF contrast (Σ) divided by the jitter term (σ_{jit}). We find that

³In B12, this systemic velocity was used as fixed input parameter for V_{rest} to optimize the orbital phase shift.



Figure 4. Upper panel, left: CCF map for Obs2. Dashed lines and labels: Positions of cuts though the data, shown on the right. Right: Cuts through the data, for different semi-amplitudes (a, b, and c), and for the rest velocity of maximum CCF contrast (d). Black dots: Data for simulated data set Obs2. Red line: Best-fitting Saltire model. Blue dashed line: Mark of the input semi-amplitude of the EBLM companion. Lower panel, left: Residual CCF map after removing the Saltire model. Dashed lines and labels: Positions of cuts though the residuals, shown on the right. Right: Cuts through the residuals at similar position as above. Error bars represent uncertainties, derived from 2D Saltire fit to the CCF map.



Figure 5. Difference between the residual maps from Obs1 and Obs2. Remaining structures are mainly caused by residual wiggles due to the phase differences between both observations. Dashed line: Position of inserted semi-amplitude ($K_{c,true}$).

the CO signal is resolved with an detection significance of 4.6σ at the expected position. With Saltire, we can precisely model the expected shape of this CO detection. It shows three cross-like structures, which are mostly buried in the noise for these data. The lower panel of Fig. 6 shows a cut of the combined Saltire model at maximum contrast (a). The double Gaussian structure of our best-fitting model function is well fitted, due to the side lobe-like structure

Table 2. Saltire fit results for atmospheric CO detection of τ Boötis b. The uncertainties for the parameters K_c and V_{rest} are a combination of MCMC fit errors and systematic uncertainties.

Parameters MCMC results Input fit pa	Input fit parameters	
$K_{\rm c} ({\rm kms^{-1}})$ 111.8 ± 2.4 (0.27, 2.39) ^b 110) ^a	
$V_{\text{rest}} (\text{km s}^{-1}) = -16.86 \pm 0.50 (0.12, 0.48)^{\text{b}} = -16$.4 ^a	
Σ (-4.70 ± 0.22)e-5 -4e-	-05	
$\Delta -0.468 \pm 0.024 -0$.1	
$\sigma_1 (\mathrm{km}\mathrm{s}^{-1})$ 2.85 ± 0.11 1.1	5	
$\sigma_2 (\mathrm{km}\mathrm{s}^{-1})$ 5.92 ± 0.25 3.4	5	
h $1 + (3.1 \pm 1.7)e^{-7}$ 1.0	0	
$\sigma_{\rm jit}$ (1.022 ± 0.0088)e-5 –		

^{*a*}B12.

^bError components: (fit error, systematic uncertainty).

of the data. The posteriors of the MCMC parameter for the width of the side lobes (σ_2) are well defined (visualized in Fig. A3). The parameter Δ shows a trend towards -0.5. Due to the asymptotic nature of the double Gaussian function, increasing the prior towards -1 will still result in the lowest value being favoured, as soon as side lobes are fitted. Since this will by design only allow marginal changes of the fit close to the numeric precision level, we keep the lower bound of this prior at -0.49.

The cuts demonstrate that the assumption of a Gaussian shape also holds for these exoplanet observations. We show the best fit for a cut through the semimajor axis close to the maximum CO detection (b) $V_{\text{rest}} = -16.5 \text{ km s}^{-1}$.

Our Saltire measurements for K_c and V_{rest} are statistically identical to the result reported by B12 (within 1σ). This is not surprising, since we used the same data, as well as the same orbital parameters, that have been derived by B12 to match the systemic velocity of the host star. Small differences below the uncertainties are not surprising as well, given our use of a different method for the



Figure 6. Upper panel, left: CO CCF map of τ Boötis b, reproduced from Brogi et al. (2012) data. Right: Saltire best-fitting model. Dashed lines indicate the position of maximum CCF contrast of the combined CCF map. The line orientation indicates the cuts – shown in the lower panels – for clarity. Lower panel: Cuts through the CCF map at maximum CCF contrast for both axis. Grey dots: CCF map data and residual data (lower panel). Error bars represent the jitter parameter, derived from MCMC sampling. Red line: Best Saltire model fit. Red shaded lines: Saltire models from posterior samples.

cross-correlation, which is based on a line mask and not on a model spectrum.

Similarly to Section 4, we fit a 1D Gaussian model to the rest velocity of maximum CCF contrast which results in $K_c = 111.29 \pm 0.58 \,\mathrm{km \, s^{-1}}$. The 1D Gaussian fit differs by about 540 m s⁻¹ from the Saltire fit (being statistically in agreement, when adding systematic uncertainties of 2–3 km s⁻¹, see Section 5.1).

5.1 Uncertainties - model robustness versus correlated noise

The V_{rest} and K_{c} sampling of the CCF map might be crucial when estimating fit errors. The uncertainty determination of the least-squares fit, as well as the jitter term derived by the MCMC, relies on the assumption that the uncertainties of adjacent data points are uncorrelated with each other. Violating the assumption of uncorrelated data will result in underestimated uncertainties from the fit.

For the V_{rest} range, this can easily be achieved by selecting a sampling that is wider than the average pixel resolution of the instrument. In this case, each data point originates from a superposition of different pixels of the detector fulfilling the requirement of uncorrelated data.

The K_c range is more difficult. As noted by Hoeijmakers et al. (2020), data in the K_c range are partly correlated, affecting the error determination using noise statistics. This correlation is a result of the *K*-focusing process. Each adjacent data point in K_c direction is a superposition of the same spectra. Those have been shifted following the reflex motion, at a different semi-amplitude, which is defined by the K_c spacing. For close-in exoplanets with circular or low-eccentricity orbits for instance, the maximum RV shift is

close to the K_c spacing itself, and minimum RV shifts are close to zero. For a closely sampled orbit, or observations during a transit/eclipse, some spectra, and sometimes all of them, are only shifted by velocities below the pixel resolution. Notably, even white noise of the spectra will, thus, become part of the correlated signal in the CCF map. To study the influence of correlation on uncertainties and measured parameters, we derive the geometrical correlation of each data point in the CCF map as the fraction of combined spectra that have been RV-shifted by less than the instrumental resolution.

For the simulated high-contrast binary (Obs1, see Section 2), a K_c sampling of $1.5 \,\mathrm{km \, s^{-1}}$ – as we use in this study – results in 21 per cent correlation between each adjacent data point in the K_c range. For the CO detection of τ Boötis b, all adjacent data points are shifted below the pixel resolution (100 per cent correlation) – due to the observation at orbital phases close to superior conjunction.

We explore in the following two sections the effects of correlated data on the uncertainties when fitting CCF maps with Saltire.

5.1.1 K_c sampling

First, we test for the impact of the K_c sampling on the K_c parameter, returned by Saltire. For this, we split the CCF map evenly in K_c space. For instance, taking every second row of the CCF map will result in two partial CCF maps, each with half the K_c sampling, hence reducing the correlation between adjacent data points. The more maps we create, the less correlated are each data from one another. By analysing each of these partial CCF maps independently, we can test for systematics not taken into account by the fit of the original CCF map.

For each set of partial maps, we then fit the Saltire model and extract K_c and its uncertainty using the least-squares fit. By doing so, we can compare the scatter of the returned K_c values, for similarly correlated data to the average fit error returned by Saltire. For τ Boötis b data, we start from a CCF map close to the pixel resolution (1.5 km s^{-1}) .

The upper panel of Fig. 7 shows the dependence of the returned K_c values as well as the average K_c value for differently correlated maps. The average value shows an increasing uncertainty the less the CCF map is sampled in K_c range. We derive the uncertainty by adding quadratically the average of the individual fit errors to the standard deviation of the individual measurements of K_c . The contributions of both components are shown in the lower panel of Fig. 7. Both uncertainties follow the same trend, which is governed by the data set sampling. They become comparable for a K_c sampling larger than 10 km s⁻¹ (<25 per cent correlation). We then repeat the same test for Obs1 of our simulated high-contrast binary, starting from a CCF map with a K_c sampling of 0.5 km s⁻¹. In this noiseless observation, the fit uncertainties returned by Saltire are always larger than the deviations (RMS) of the individual measurements. Also, no residual deviations of K_c are observed.

Using this test, we do not find systematic uncertainties in obvious excess of the reported fit uncertainties. We note that this test relies on the sampling of the signal itself, resulting in a lower accuracy the more sparsely the signal is sampled in the K_c range. It is, thus, sensitive to detect systematics from correlated data that have a correlation length scale in K_c smaller than the signal itself.

We further show in the upper panel of Fig. 7 that the average K_c parameter is largely invariant from the K_c sampling. Saltire can therefore be used to derive accurate parameters even for sparsely sampled data sets, which might be an opportunity to decrease



Figure 7. Uncertainty analysis from K_c spacing for τ Boötis b as well as for the simulated high-contrast binary (Obs1). Upper panel: Measurement of K_c from partial maps with different K_c spacing. Grey dots: Individual measurements for partial maps. Coloured dots: Average K_c of individual measurements. Colour coding: K_c sampling, of partial map. Lower panel: Components of uncertainties. Blue crosses: Average uncertainties from individual fits. Coloured triangles: RMS of individual K_c measurements. Colour coding is as above panel.

computation costs – which are usually involved – when deriving highly sampled CCF maps.

5.1.2 Global correlated noise

Secondly, we test the global impact of correlated noise on the measurement of K_c and V_{rest} . As mentioned above, by design of the *K*-focusing process, noise in the post-processed spectra can become correlated signal, for higher geometric correlations. Assuming that Poisson (white) noise is random for each spectrum, we can test this effect by splitting the data set into independent samples. Correlated white noise will thus be different in each sample, allowing to estimate global effective uncertainties by measuring the CCF position of the CO signal for each sample. Different from the first test, this will measure the global uncertainties for all data points in the CCF map (independent from their geometric correlation).

We split the CRIRES observations in three samples by selecting random data from each night (without repetition). In this way, each data sample consists of different measurements, but covers similar orbital phases. Then, we derive the CCF maps in the same way we did before for the whole data set. As we show in Fig. B1, the CO signal is detected for each sample at about 3σ significance. We measure the position of the CO signal position using a least-squares fit with Saltire. Similarly to the first test, we derive the RMS scatter of the position parameters.

We find that the resulting RMS scatter for K_c is 2.39 km s⁻¹ and for V_{rest} 0.48 km s⁻¹. These are about a factor of 10 larger, compared to the fit uncertainties we found earlier from the MCMC fit of the whole data set. We also repeat the 1D Gaussian fit for each of the samples and find for K_c similar uncertainties for the returned RMS. In Table 2, we add these systematic uncertainties for V_{rest} and K_c in quadrature to the MCMC uncertainties.

Summarizing both of the above tests for systematic uncertainties, we find clear systematics when using data with different noise budget, but negligible systematics when analysing the combined CCF map for correlation length scales that are smaller than or similar to the detectable CO signal itself.

We conclude that, first, systematics indeed originate from correlated noise, produced by the *K*-focusing process. This noise can be caused not only by white noise, noise from data post-processing, but also from underfitted signal (See Section 4.1). Secondly, the correlation length scale in the K_c range is similar to or larger than the CO signal itself, thus must lead to underestimated uncertainties when fitting the CCF map alone.

5.2 Phase-dependent modelling

We use the best-fitting parameters in Table 2 to evaluate the Saltire model separately for the phases of the three CRIRES observations at orbital phases before, during, and after superior conjunction. Fig. 8 shows the CCF maps for each observation, as well as the resulting models. The dependence of the CCF signal from the orbital phase is clearly visible. Despite not being a fit to the data, the model matches these observations rather well, with the absorption spectrum clearly detected in each observation.

Assuming a stationary CO spectrum of the planet atmosphere, we would expect aliases originating from wiggles (here, spurious correlations of the CO line mask with different parts of the CO spectrum, or with other atmospheric lines with similar patterns), which are quasi-stationary at each observed phase (see Section 4.1 for a discussion). This assumption can be used as a tool to disentangle wiggles from other aliases from correlated noise or imperfect postprocessing of the data by simply comparing residual aliases of observations at different orbital phases and including quasi-stationary components into the Saltire model.

Fig. 8, shows cuts (a) through the CCF maps for the semiamplitude of maximum CCF contrast (of the combined CCF map). The detection significance (3σ) for each CRIRES observation does not allow to securely disentangle the origin of the aliases, which, nevertheless, will be possible as soon as higher signal-to-noise data will become available.

Measuring the Keplerian semi-amplitude from atmospheric signals obtained near phases of transits or eclipses introduces large uncertainties (e.g. Brogi et al. 2018). This can be understood by the middle panel of Fig. 8, which shows the data during phases of superior conjunction. The CCF signal for both – data and model – shows a very small correlation with the rest velocity, but covers a wide range of semi-amplitudes (see cut b). Measuring the mean position of the CCF signal by only evaluating the maximum detection significance can lead to large deviations produced by small correlated noise components along the chosen V_{rest} axis. Using Saltire has clearly the advantage of being less biased to such small aliases, by taking the whole shape into account.



Figure 8. Left panels: CO CCF maps of τ Boötis b, reproduced from B12 data for different observation phases. From top to bottom: Before, during, and after superior conjunction. Dashed lines indicate the position of maximum CCF contrast of the combined CCF map. Second panel from left: Saltire models, for similar observation phases, evaluated from best-fitting parameters of the combined fit. Second panel from right: Cuts through the CCF maps for the semi-amplitude of maximum CCF contrast of the combined CCF map. Grey dots: CCF map data. Error bars are estimated as the MCMC jitter term of the combined sample times $\sqrt{3}$. Blue line: Saltire model evaluated from best-fitting parameters of the combined fit. Right panel: Similar cuts through the CCF maps, but for the rest velocity of maximum CCF contrast of the combined CCF map.

6 GOOD PRACTICES

Based on our modelling and the tests we described in Section 5.1, we recommend the following considerations when creating cross-correlation maps and fitting them with Saltire:

(i) Spectra need to be corrected from trends of the systemic velocity before creating the CCF map.

(ii) For the V_{rest} range of the CCF map, use the pixel resolution of the spectrograph as finest sampling to avoid additional correlations of the data. The sampling in the K_c range does not affect the results, as long as the CCF detection is well sampled.

(iii) As for any least-squares algorithm, choose your starting parameters and prior bounds close to the expected values. If you have clear side lobes, choose a value of Δ close to -0.5. Allowing -1 will increase the computational time significantly.

(iv) If fitting for h, allow for a proper range. Too large a range will result in no improvement of the least-squares fit, with no errors being calculated. This will not affect MCMC results.

(v) The returned uncertainties are precision errors only. Systematic changes of the peak position due to correlated noise of the data are likely. Thus, they also need to be taken into account (e.g. by measuring independent data samples) when reporting fit uncertainties. (vi) A word of caution: As we described in Section 5.1, noise is correlated by the *K*-focusing process. In other words, any noise structure in the CCF map can be mistaken as signal and fitted using the Saltire model. Especially for very low detection significance ($< 3\sigma$), fitting CCF maps could result in a best fit for a correlated noise structure. Again, measuring independent data samples can help to verify such detections.

7 DISCUSSION

HRCCS observations are widely used to detect and to map atomic and molecular species in exoplanet atmospheres both in the semiamplitude and rest velocity range. A precise measurement of the companion's signal in K_c – V_{rest} CCF maps allows (i) to constrain the orbit of the companion, and thus to turn a giant exoplanet or low-mass binary companion into a double-lined system, which allows to derive dynamical masses of both components, and (ii) to constrain atmospheric dynamics by measuring the signal position of different species. Measuring this position is difficult. It depends very much on the detection significance, and also on the orbital phase observed. Transit transmission observations especially suffer from large uncertainties when measuring K_c solely from the position of the maximum CCF signal (Brogi et al. 2018). Aiming to improve the measurement of the CCF signals' position, we present Saltire – an easy-to-use model to predict the phase-dependent shape of CCF maps, created for HRCCS observations of both exoplanet atmospheres, as well as of high-contrast binaries. The model allows to predict CCF maps only by assuming a 1D line shape. We show that the 2D shape for a CCF map results from a superpositioning, for this 1D function in the K_c-V_{rest} frame, which we call *K*-focusing process. This is exactly the same process that is in use for decades to generate CCF maps from HRCCS observations. Making use of the fact that the parameters that have been used to create the CCF map are fixed, we can use this model to efficiently fit complex 2D shapes with a minimal set of parameters. In this work, we realize this 1D shape by a double Gaussian function. The Saltire model is designed to facilitate any 1D line shape (e.g. Lorentzian or more complex shapes).

We demonstrate and evaluate the model for a simulated binary system, as well as for archival IR observations of the atmosphere of the gas giant exoplanet τ Boötis b. We use line masks to create the CCF maps, which allows us to derive average line profiles for the exoplanet. While the Gaussian approximation is expected to be a good fit to measure stellar spectra, we show that it can also be used to model the detection of atmospheric molecules in gas giants. By predicting the phase-dependent shape of an HRCCS observation, Saltire can be used to plan HRCCS observations by simulating the combined CCF map expected from a certain observation strategy, but it can also be used to precisely fit the signal positing.

With Saltire, we can recover the correct semi-amplitude with an accuracy better than 20 m s^{-1} , for a simulated noiseless observation. The companions' own signal ($\sim 100 \text{ km s}^{-1}$) is typically of the order of 5000 times greater than our method's accuracy, giving the potential to measure very accurate and precise dynamical masses. We also show that this accuracy can be achieved independently from the phase coverage of the observation, which allows to measure the semi-amplitude 10 times more accurate compared to a simple 1D Gaussian fit, which is often used in literature. This phase independence of Saltire is specially important for exoplanet observations, which usually cover only a part of the planet's orbit, resulting in asymmetric shapes of the CCF signal.

We show on simulated observations that in the absence of noise and post-processing aliases, the main source of uncertainties are aliases from spurious correlations with other lines/sets of lines within the spectrum (wiggles). With Saltire, we currently underfit the CCF signal, by not fitting for these aliases in the CCF.

Including them in Saltire would allow to better model the residual structure of high signal-to-noise CCF maps and improve the parameter determination, especially for data with a short wavelength coverage, which typically enhances such correlations. In-depth analysis of these wiggles, as well as possible methods to model them have recently been discussed for the detection of individual ionized, atomic species in the atmosphere of KELT-9b (Borsato et al. 2023), as well as for high-precision RV measurements of double-lined binaries (Lalitha et al. 2024). We show how such wiggles are represented in high signal-to-noise CCF maps due to the *K*-focusing process, and discuss the importance of including a phase-dependent modelling of these structures when deriving high-precision parameters. We propose that Saltire can therefore serve as a diagnostic tool to disentangle wiggles from other aliases, e.g. due to post-processing of data for high signal-to-noise data.

For the detection of CO in the atmosphere of τ Boötis b (5 σ detection), we recover the semi-amplitude better than 1σ , compared to the literature values. We achieve uncertainties of

2.4 km s⁻¹ for the planet's semi-amplitude and 0.5 km s⁻¹ for the rest velocity.

We find that the main source of uncertainty is correlated noise, which is a by-product of the *K*-focusing process and, thus, is typically present in HRCCS CCF maps. This noise can be caused by white noise, data post-processing, or by wiggles. The estimated uncertainties for τ Boötis b are of the order of 10 times larger than the fit errors. This is because, different from statistical noise in the CCF map, correlated noise becomes a CCF signal itself, thus fitting CCF maps will result in fit errors, which are underestimated.

We show that it is possible to estimate the impact of correlated white noise by analysing independent, partial data sets using Saltire. That this is even possible for the CO detection in τ Boötis b, validates this approach to be extended to other low signal-to-noise detections from exoplanet atmospheres. This shows that Saltire is an important tool to constrain the signal positions in CCF maps from HRCCS observations. We further show that Saltire is able to robustly return precise semi-amplitudes for different samplings in the K_c range of the CCF map offering possibilities to lower the computational costs when deriving CCF maps.

We compare our results with Saltire to a simple 1D Gaussian fit of the same CCF map. Since both methods are limited by the impact of correlated noise, no substantial improvement can be achieved. Nevertheless, as we can clearly show for our noiseless simulations, Saltire will allow to derive much more precise parameters for high signal-to-noise HRCCS detections.

Using observations covering a significant portion of the planet's orbit, phase-resolved atmospheric properties can be retrieved from UHJs, which offer a large signal-to-noise ratio. This has been successfully demonstrated for extremely irradiated planets like HD 209458 b (Beltz et al. 2021) and KELT-9b (Pino et al. 2022). In Saltire, we implement a weighting for the contribution of each observation. While this is used to account for uncertainties in the data, in the future Saltire could be used to simulate the phase-dependent contribution of single observations to combined CCF maps.

ACKNOWLEDGEMENTS

This research was funded from the European Research Council (ERC) through the European Union's Horizon 2020 research and innovation programme (grant agreement no. 803193/BEBOP). MB acknowledges partial support from the STFC research grant ST/T000406/1.

DATA AVAILABILITY

Most data underlying this article are available online as indicated in the specific section or reference. Data obtained with European Southern Observatory (ESO) telescopes are available in the ESO Science Archive Facility, at http://archive.eso.org/cms.html. Corrected data or CCF maps underlying this article will be shared on reasonable request to the corresponding author. Modelled data can be generated using the Saltire code available on Github.

REFERENCES

Baranne A. et al., 1996, A&AS, 119, 373

Bello-Arufe A., Buchhave L. A., Mendonça J. M., Tronsgaard R., Heng K., Jens Hoeijmakers H., Mayo A. W., 2022, A&A, 662, A51

Beltz H., Rauscher E., Brogi M., Kempton E. M. R., 2021, AJ, 161, 1

Birkby J. L., 2018, preprint (arXiv:1806.04617)

Birkby J. L., de Kok R. J., Brogi M., de Mooij E. J. W., Schwarz H., Albrecht S., Snellen I. A. G., 2013, MNRAS, 436, L35

Blanco-Cuaresma S., Soubiran C., Heiter U., Jofré P., 2014, A&A, 569, A111

Borsa F. et al., 2015, A&A, 578, A64

- Borsato N. W., Hoeijmakers H. J., Prinoth B., Thorsbro B., Forsberg R., Kitzmann D., Jones K., Heng K., 2023, A&A, 673, A158
- Bourrier V. et al., 2018, Nature, 553, 477
- Brogi M., Line M. R., 2019, AJ, 157, 114
- Brogi M., Snellen I. A. G., de Kok R. J., Albrecht S., Birkby J., de Mooij E. J. W., 2012, Nature, 486, 502
- Brogi M., de Kok R. J., Albrecht S., Snellen I. A. G., Birkby J. L., Schwarz H., 2016, ApJ, 817, 106
- Brogi M., Giacobbe P., Guilluy G., de Kok R. J., Sozzetti A., Mancini L., Bonomo A. S., 2018, A&A, 615, A16
- Brogi M. et al., 2023, AJ, 165, 91
- Butler R. P. et al., 2006, ApJ, 646, 505
- Cheverall C. J., Madhusudhan N., Holmberg M., 2023, MNRAS, 522, 661
- de Kok R. J., Brogi M., Snellen I. A. G., Birkby J., Albrecht S., de Mooij E. J. W., 2013, A&A, 554, A82
- Donati J. F. et al., 2008, MNRAS, 385, 1179
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
- Gill S. et al., 2019, A&A, 626, A119
- Gordon I. et al., 2017, J. Quant. Spectrosc. Radiat. Transfer, 203, 3
- Hoeijmakers H. J. et al., 2020, A&A, 641, A123
- Husser T. O., Wende-von Berg S., Dreizler S., Homeier D., Reiners A., Barman T., Hauschildt P. H., 2013, A&A, 553, A6
- Justesen A. B., Albrecht S., 2019, A&A, 625, A59
- Kaeufl H.-U. et al., 2004, in Moorwood A. F. M., Masanori I., eds, Proc. SPIE Conf. Ser. Vol. 5492, Ground-based Instrumentation for Astronomy. SPIE, Bellingham, p. 1218
- Kalman D., 1996, Coll. Math. J., 27, 2
- Kesseli A. Y., Snellen I. A. G., Casasayas-Barris N., Mollière P., Sánchez-López A., 2022, AJ, 163, 107

- Lafarga M. et al., 2023, MNRAS, 521, 1233
- Lalitha S. et al., 2024, MNRAS, 527, 2261
- Newville M., Stensitzki T., Allen D. B., Rawlik M., Ingargiola A., Nelson A., 2016, Astrophysics Source Code Library, record ascl:1606.014
- Nugroho S. K., Gibson N. P., de Mooij E. J. W., Watson C. A., Kawahara H., Merritt S., 2020, MNRAS, 496, 504
- Parmentier V. et al., 2018, A&A, 617, A110
- Pepe F., Mayor M., Galland F., Naef D., Queloz D., Santos N. C., Udry S., Burnet M., 2002, A&A, 388, 632
- Pepe F. A. et al., 2010, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III. SPIE, Bellingham, p. 77350F
- Pino L. et al., 2022, A&A, 668, A176
- Piskorz D. et al., 2016, ApJ, 832, 131
- Pollacco D. L. et al., 2006, PASP, 118, 1407
- Prinoth B. et al., 2022, Nat. Astron., 6, 449
- Sebastian D. et al., 2023, MNRAS, 519, 3546
- Snellen I. A. G., de Kok R. J., de Mooij E. J. W., Albrecht S., 2010, Nature, 465, 1049
- Standing M. R. et al., 2023, Nat. Astron., 7, 702
- Swayne M. I. et al., 2021, MNRAS, 506, 306
- Triaud A. H. M. J. et al., 2013, A&A, 549, A18
- Triaud A. H. M. J. et al., 2017, A&A, 608, A129
- Triaud A. H. M. J. et al., 2022, MNRAS, 511, 3561
- von Boetticher A. et al., 2019, A&A, 625, A150
- Wardenier J. P., Parmentier V., Lee E. K. H., Line M. R., Gharib-Nezhad E., 2021, MNRAS, 506, 1258

Yan F. et al., 2023, A&A, 672, A107

APPENDIX A: MCMC SAMPLES FOR SIMULATED HIGH-CONTRAST BINARY AND τ BOÖTIS B OBSERVATIONS

Saltire 10933



Figure A1. Corner plot of MCMC samples for Saltire fit to CCF map of simulated EBLM binary observation (Obs1).



Figure A2. Corner plot of MCMC samples for Saltire fit to CCF map of simulated EBLM binary observation (Obs2).

Saltire 10935



Figure A3. Corner plot of MCMC samples for Saltire fit to combined CCF map of τ Boötis b.

APPENDIX B: PARTIAL OBSERVATIONS OF τ BOÖTIS B



Figure B1. Combined CCF maps for τ Boötis b. Each sample consists of different spectra, covering a similar orbital phase of the planet. White dashed lines mark the position of the CO signal. Correlated white noise creates unique structures for each of these CCF maps, sometimes even stronger than the CO signal itself.

This paper has been typeset from a TFX/LATFX file prepared by the author.